Forward Selection Fama-MacBeth Regression with Higher-Order Asset Pricing Factors*

Nicola Borri, Denis Chetverikov, Yukun Liu, Aleh Tsyvinski

March 16, 2025

Abstract

We show that the higher-orders and their interactions of the common sparse linear factors can effectively subsume the factor zoo. To this extend, we propose a forward selection method to estimate a high-dimensional stochastic discount factor model, isolating the most relevant higher-order factors. Applying this approach to terms derived from six widely used factors (the Fama-French five-factor model and the momentum factor), we show that the resulting higher-order model significantly outperforms traditional benchmarks both in-sample and out-of-sample. Moreover, it effectively subsumes a majority of the factors from the extensive factor zoo, suggesting that the pricing power of most zoo factors is attributable to their exposure to higher-order terms of common linear factors.

JEL Codes: G10, G12, C10

Keywords: asset returns, higher-order interactions, forward selection Fama-MacBeth regression method, factor zoo

*Nicola Borri is with Luiss University, Rome. Denis Chetverikov is with the University of California, Los Angeles. Yukun Liu is with the University of Rochester, Simon Business School. Aleh Tsyvinski is with Yale University. We thank Stefano Giglio, Toby Moskowitz, and Whitney Newey for comments.

1 Introduction

The asset pricing literature has proposed hundreds of candidate factors that could potentially explain the cross-section of asset returns. Cochrane (2011), in his presidential address, refers to this extensive set of candidate factors as the "factor zoo." Sparse linear factor models, such as those introduced by Fama and French (2015) and Carhart (1997), struggle to adequately price this expansive factor zoo, implying they may not fully span the stochastic discount factor (SDF). In this paper, we show that the interactions of the higher-order terms of the common sparse linear factors can effectively subsume the factor zoo. We also propose a simple methodology that we called *forward selection Fama-MacBeth regression method* to determine important higher-order factors.

The prevalent focus in the literature on sparse linear factor models is primarily driven by considerations of parsimony and is motivated by the Taylor expansion of the SDF, where a first-order approximation might explain a substantial portion of its variation. However, there is no theoretical guarantee that a linear approximation sufficiently captures the variation of the SDF. We show instead that higher-order factors and their interactions can significantly better capture the SDF.

The main difficulty of studying the higher-order terms is the dimensionality issue. Consider, for example, the standard 6-factor model that includes the Fama-French 5 factors (Fama and French (2015)) and the momentum factor (Carhart (1997)). The higher-orders up to degree of 3 (4) already have 63 (114) factors. Therefore, we develop a forward selection method to choose factors that explain the stochastic discount factor that is suitable for the case when the number of potentially important factors is large. The method adds factors one by one into the set of selected factors, each time maximizing the R^2 of the cross-sectional regression (second step) of the Fama-MacBeth procedure. After the set of factors is selected, we estimate the stochastic discount factor loadings by running the cross-sectional regression of sample average returns on sample covariances between returns and selected factors. We derive the rate of convergence of the proposed estimator and also propose a method to debias it, leading to simple inference procedures. We refer to our method as the forward selection Fama-MacBeth regression method (FS-FMB).

Our method complements the Lasso-based method proposed in Feng et al. (2020). Both forward selection and Lasso are among the most important machine learning techniques suitable for estimation of high-dimensional models. Advantages of the forward selection method are the following. First, it is invariant to scaling of factors, which is important because Lasso-based methods yield results that depend non-trivially on how factors are scaled (see Nagel (2021)). Second, forward selection is highly interpretable: not only it gives the set of factors explaining the stochastic discount factor, but it also gives the order of importance of factors. Third, forward selection is very easy to implement from the computational point of view.

Specifically, we show that the performance of the common factors is dramatically improved with their higher-order factors. We start with 6 of the most common factors in the asset pricing literature. These factors include the Fama-French 5 factors of Fama and French (2015) and the momentum factor of Jegadeesh and Titman (1993), and we refer to this model as the FF5M. We consider higher-orders as well as interactions of these 6 factors up to degree of 3, leading to 57 candidate higher-order factors. For the test assets, we use a large cross-section of equities including 484 characteristic-managed portfolios from Kozak et al. (2020). We apply our FS-FMB to select the higher-order factors that are most important in pricing the cross-section of these 484 asset returns. The procedure selects 7 higher-order factors, including 2 second degree higher-orders and 5 interactions. We show that the inclusion of the higher-order factors leads to dramatic improvement in pricing the cross-section of asset returns – the adjusted cross-sectional R^2 increases from 31.2% for the baseline FF5M model to 59% for the selected higher-order factor model. Moreover, the intercept alpha of the higher-order factor model is not statistically significant at the 5% level.

We then examine the SDF loadings of the higher-order factors selected by the forward selection FMB regression with both the non-debiased estimates based on the standard FMB two-pass regression and the debiased estimates. Estimates from the standard FMB two-pass method are likely biased because of the omitted variable bias (see, e.g., Chernozhukov et al. (2018) and Giglio and Xiu (2021)). The directions of the point estimates are the same and the magnitudes are similar with both methods. The estimates based on the non-debiased method are all significant at the 5%

level, while the estimates based on the debiased method are all significant at the 10% level with 5 of the 7 SDF loading estimates being significant at the 5% level. Overall, the results show that the higher-order factors are important to empirically capture the SDF and significantly price assets in the cross-section.

In our baseline analysis, we demonstrate that higher-order factors are important components of the SDF and price cross-sectional asset returns in-sample. We further provide evidence of the out-of-sample performance of the selected higher-order factor model, where we conduct two types of out-of-sample tests. Firstly, we present an out-of-sample test in the asset space. We select higher-order factors and estimate SDF loadings from the training assets and then evaluate the out-of-sample performance using the remaining assets. Our results show that the selected higher-order factors and the performance of the higher-order factor model are stable in this out-of-sample exercise.

Secondly, we conduct an out-of-sample test in the time-series, where we split the sample into two equal subsamples based on the time series. We use the first half of the sample as the training sample and the second half of the sample as the test sample. Then, we estimate the model in the training sample and evaluate the out-of-sample performance in the test sample. Compared with the benchmark factor models including the CAPM, the FF3 factor model, the FF5 factor model, and the FF5M factor model, the out-of-sample R^2 of the higher-order factor model is significantly larger.

Next, we examine whether the selected higher-order factor model can subsume the pricing power of the factors in the factor zoo. We employ two empirical methods. First, we perform standard FMB regressions to test the significance of the SDF loading of each factor in the factor zoo, controlling for the higher-order factor model. Second, we construct factor mimicking portfolios for the seven higher-order factors and conduct time-series asset pricing tests on the zoo factors.

Under the first approach, we estimate each zoo factor's SDF loading while controlling for our higher-order factor model. We find that only 7 our of the 148 zoo factors from Jensen et al. (2023) remain significant at the 5% level. The second approach is based on the factor-mimicking

approach. The benefit is that we are able to convert the non-tradable higher-order factors to tradable factors but the projection inevitably loses information from the original factors. The results using the second approach confirms that the higher-order factor model significantly reduce the pricing power of the zoo factors. Overall, these results demonstrate that the higher-order factors selected by the forward Fama-MacBeth regression method effectively subsume the zoo factors, although the linear benchmark factors cannot.

We then examine the relationship between the selected higher-order factors and macroeconomic risks. First, we examine the correlations between the selected higher-order factors and two primary macroeconomic indicators – U.S. equity market returns and the CBOE VIX index – in the full sample and during the recession periods. For the first two higher-order factors (i.e., SMB2 and SMB2*Mom), the magnitude of the correlations with U.S. equity market returns are markedly higher during the NBER recession periods relative to the full sample. For the last four higher-order factors (i.e., Mkt-RF2, Mkt-RF2*RMW, Mkt-RF*SMB, and HML2*Mkt-RF), they tend to have high correlation in magnitude with the VIX index, especially during the NBER recession periods. These results show that the selected higher-order factors capture different macroeconomic risks, especially during economic downturns. Additionally, we examine the exposures of the 7 selected higher-order factors to a large set of common macroeconomic factors. We find that the these higher-order factors explain a sizable fraction of variations for 5 of the 7 higher-order factors, but not for SMB2*Mom or Mom2*RMW. Of the macroeconomic factors, factors related to uncertainties appear to be more important in accounting for the higher-order factors.

Additionally, we provide robustness results for our empirical analyses. First, we show that the performance of the higher-order factor model is extremely unlikely to be driven by randomness by simulating white-noise factors. Second, we use the selected higher-order factor model to subsume an alternative factor zoo of Chen and Zimmermann (2022) and reach similar conclusions. Third, we experiment alternative sets of candidate higher-order factors: (1) only include higher-orders, (2) only interactions, (3) up to degree 2, and (4) up to degree 4. We find that both higher-orders and interactions are important components of the SDF, especially interactions. Moreover, we show

that the selected higher-order factor model up to degree 2 significantly underperform, while the selected higher-order factor model up to degree 4 has similar performance as our baseline model up to degree 3.

Our paper relates to a vast literature on factor models in asset pricing. A major part of the literature focuses on linear factor models, such as the characteristic-based models (e.g., Fama and French (1993; 1996; 2015); Carhart; Koijen et al.) and statistical-based models (e.g., Connor and Korajczyk). A small subset of papers emphasizes non-linearity in explaining asset returns, such as Bansal and Viswanathan (1993), Harvey and Siddique (2000), and Dittmar (2002). We contribute to this literature by showing that higher-orders and higher-order interactions of common linear factors are important components of the SDF and price the cross-sectional asset returns.

The paper also contributes to the factor zoo literature. Cochrane (2011) identifies the factor zoo issue and calls for research to identify factors that are actually important for the SDF. Harvey et al. (2016) argue that, based on multiple hypothesis framework, many zoo factors are not significant. Harvey and Liu (2021) provides a bootstrap model selection framework that accounts for multiple hypothesis testing. McLean and Pontiff (2016) show that most of the zoo factors suffer from post publication bias. Green et al. (2017) uses FMB regression to evaluate 94 firm characteristics and argue that less than ten provide significant independent information to expected returns. Feng et al. (2020) propose a rigorous econometric framework to evaluate whether a new factor contributes to the SDF in light of the existing factor zoo. Our paper demonstrates that the higher-orders of the common linear factors are able to subsume the majority of the factor zoo.

Finally, our paper relates to the machine learning literature. Gu et al. (2020) show that machine learning methods are better in predicting future stock returns out-of-sample than OLS methods. Freyberger et al. (2020) propose a group LASSO method to choose firm characteristics that are most related to future returns. Kozak et al. (2020) show that machine learning methods applied to the factor zoo can better approximate the SDF. Lettau and Pelger (2020) propose a method for estimating latent asset pricing factors that fit both the time-series and the cross-section of asset returns. Giglio and Xiu (2021) propose a method to estimate risk premia of factors even if there is

missing factors. Borri et al. (2024) show that a a new non-linear single-factor asset pricing model motivated by the Kolmogorov-Arnold representation theorem is able to price assets from different asset classes. In this paper, we propose a forward selection method to estimate a high-dimensional stochastic discount factor (SDF) model, isolating the most relevant higher-order factors.

The rest of the paper is organized as follows. Section 2 discusses the forward selection Fama-MacBeth method. Section 3 presents the data, estimate, and evaluate our baseline model. Section 4 shows that the selected higher-order factor model effectively subsumes the majority of the factor zoo. Section 5 studies the relationship between the selected higher-order factors and macroeconomic risks. Section 6 presents additional robustness analyses and Section 7 concludes. The Online Appendix contains technical derivations and further empirical results.

2 Methodology

In this section, we set up the model. We introduce its main components—returns, factors, and the stochastic discount factor—and explain the relationship between them. We then present a forward selection Fama-MacBeth regression procedure as a method for selecting the factors that drive the stochastic discount factor. Using this selection procedure, we propose an estimator of the stochastic discount factor loadings. As we will see, the proposed estimator is biased. We therefore also explain how to debias it and how to perform inference on the stochastic discount factor loadings using the debiased estimator.

Notation. For any integer $p \geq 1$, we denote $[p] = \{1, \ldots, p\}$. For any vector $x = (x_1, \ldots, x_p)^\top \in \mathbb{R}^p$ and any set $\S \subset [p]$, we denote $x_\S = (x_i)_{i \in \S}$, the sub-vector of x consisting of all its components with indices in $\S$. For any matrix $A = (A_{i,j})_{i \in [n], j \in [p]} \in \mathbb{R}^{n \times p}$ and any set $\S \subset [p]$, we denote $A_\S = (A_{i,j})_{i \in [n], j \in \S}$, the sub-matrix of A consisting of all its columns with indices in $\S$. For any vector $x = (x_1, \ldots, x_p)^\top \in \mathbb{R}^p$, we let $\|x\|_1 = \sum_{j=1}^p |x_j|$, $\|x\|_2 = (\sum_{j=1}^p |x_j|^2)^{1/2}$, and $\|x\|_\infty = \max_{1 \leq j \leq p} |x_j|$ be its ℓ_1-, ℓ_2-, and ℓ_∞-norms. For any matrix $A = (A_{i,j})_{i \in [n], j \in [p]} \in \mathbb{R}^{n \times p}$, we let $\|A\|_2 = \sup_{v \in \mathbb{R}^p : \|v\|_2 = 1} \|Av\|$ and $\|A\|_{\infty,1} = \max_{1 \leq j \leq p} \sum_{i=1}^n |A_{i,j}|$ be its spectral and $\ell_{\infty,1}-$norms.

For any symmetric matrix $A \in \mathbb{R}^{p\times p}$, we let $\lambda_{\min}(A)$ and $\lambda_{\max}(A)$ be its smallest and largest eigenvalues.

2.1 Model

Let $r_{i,t} \in \mathbb{R}$ be the excess return of asset $i \in [N] = \{1, \ldots, N\}$ and time period $t \in [T] = \{1, \ldots, T\}$ and let $f_t = (f_{t,1}, \ldots, f_{t,p})^\top \in \mathbb{R}^p$ be the vector of factors. Throughout the paper, we assume that the returns and the factors are stationary over time. In addition, we assume that the vector of factors are able to span the stochastic discount factor (SDF):

$$m_t = 1 - \psi^\top (f_t - \mathbb{E}[f_t]), \quad \text{for all } t \in [T], \tag{1}$$

where m_t is the SDF and $\psi = (\psi_1, \ldots, \psi_p)^\top \in \mathbb{R}^p$ is the vector of SDF loadings.

By definition, since we work with the *excess* returns, m_t is orthogonal to the returns in the sense that $\mathbb{E}[m_t r_{i,t}] = 0$ for all $i \in [N]$. Using these equations in combination with 1, it then follows from Section 6.3 in Cochrane (2005) that for each asset $i \in [N]$, we have

$$r_{i,t} = \beta_i^\top (\gamma - \mathbb{E}[f_t]) + \beta_i^\top f_t + \varepsilon_{i,t}, \quad \text{for all } t \in [T], \tag{2}$$

where $\beta_i = (\beta_{i,1}, \ldots, \beta_{i,p})^\top \in \mathbb{R}^p$ is the vector of factor exposures corresponding to the factors f_t, $\gamma = (\gamma_1, \ldots, \gamma_p)^\top \in \mathbb{R}^p$ is the vector of corresponding factor risk premia, and $\varepsilon_{i,t} \in \mathbb{R}$ is the residual satisfying $\mathbb{E}[\varepsilon_{i,t}] = 0$ and $\mathbb{E}[\varepsilon_{i,t} f_t] = \mathbf{0}_p$. The vector of SDF loadings ψ is related to the vector of factor risk premia γ via the following equation:

$$\psi = (\mathbb{E}[(f_t - \mathbb{E}[f_t])(f_t - \mathbb{E}[f_t])^\top])^{-1}\gamma, \tag{3}$$

where we assume that the matrix $\mathbb{E}[(f_t - \mathbb{E}[f_t])(f_t - \mathbb{E}[f_t])^\top]$ is non-singular, so that its inverse exists. In addition, by using standard least-squares projection formulas together with 2 and 3, we have

$$\mathbb{E}[r_{i,t}] = \sum_{i=1}^{N} C_{i,j}\psi_j, \tag{4}$$

where $C_{i,j} = \mathbb{E}[(f_{t,j} - \mathbb{E}[f_{t,j}])(r_{i,t} - \mathbb{E}[r_{i,t}])] \in \mathbb{R}^p$ is the covariance between $f_{t,j}$ and $r_{i,t}$.

In this paper, a primary task is to propose a novel estimator of the vector of SDF loadings ψ that can be used in the case when the number of factors p is large, potentially much larger than the number of assets N and/or the number of time periods T. As we will see, the classic estimation techniques may fail in this case. In fact, one can show that the model cannot be consistently estimated without further assumptions in this case. Therefore, to make progress, and following the literature on estimation of high-dimensional models, we assume that the model is sparse, in the sense that most of the variation in the SDF in 1 can be explained by a relatively small number of factors. In other words, only a relatively small number of coefficients in the vector ψ are sufficiently different from zero. In the Appendix, we also discuss estimation of the vector of factor risk premia γ.

2.2 Selection & Estimation

To develop intuition for the proposed method, let us first recall the classic Fama-MacBeth procedure that can be used to estimate the vector of risk premia γ in the low-dimensional case, i.e. when p is small. The procedure consists of two steps. First, for each asset $i \in [N]$, we run a time series OLS regression of $r_{i,t}$ on f_t, with intercept included, to obtain estimates of factor exposures $\widehat{\beta}_i$. Second, we run a cross-sectional OLS regression of average return $\bar{r}_i = T^{-1}\sum_{t=1}^{T} r_{i,t}$ on $\widehat{\beta}_i$ (without intercept). The vector of slope estimates in this regression is then an estimator of the vector of risk premia γ. In turn, once the vector of risk premia is estimated, the vector of stochastic discount factor loadings ψ can be estimated using an empirical version of equation 3.

The procedure described above is consistent when the number of factors p is small. When the number of factors p is large, however, all three steps may fail as they require inverting singular/poorly-conditioned matrices. To deal with this problem, we propose a forward selection algorithm that

starts with no factors and then adds factors iteratively one by one with the goal of achieving the highest possible R^2 in the cross-sectional regression of the Fama-MacBeth procedure. We stop the algorithm once a desired number of factors is achieved or the growth in the R^2 becomes sufficiently small. Once a set of factors is selected, we perform the Fama-MacBeth procedure and use an empirical version of 3 to estimate ψ as described above on selected factors. We set an estimate of the SDF loadings for non-selected factors to be zero.

To define our method more precisely, for each subset $\mathcal{S}$ of $[p] = \{1, \ldots, p\}$, consider the Fama-MacBeth procedure using only factors in the set $\mathcal{S}$. In particular, define

$$\widehat{\beta}_{i,\S} = \left(\sum_{t=1}^{T} (f_{t,\S} - \bar{f}_{\S})(f_{t,\S} - \bar{f}_{\S})^{\top} \right)^{-1} \left(\sum_{t=1}^{T} (f_{t,\S} - \bar{f}_{\S}) r_{i,t} \right), \tag{5}$$

where $f_{t,\S} = (f_{t,j})_{j \in \S}$, $\bar{f}_{\S} = T^{-1} \sum_{t=1}^{T} f_{t,\S}$, and the inverse of the matrix is understood as the Moore-Penrose generalized inverse if the matrix is singular. Also, let $R^2_{FM}(\S)$ be the R^2 in the cross-sectional OLS regression of $\bar{r}_i$ on $\widehat{\beta}_{i,\S}$. We apply the forward selection algorithm with the aim to maximize $R^2_{FM}(\S)$. For a desired number of factors $\widehat{s}$, the algorithm takes the following form:

Algorithm 2.1 (Forward Selection Fama-MacBeth Regression)**.**

1. *Set* $\S = \emptyset$ *and* $Q = 0$.
2. *Set* $\widehat{j} = \arg\max_{j \in [p] \setminus \S} R^2_{FM}(\S \cup j)$.
3. *Set* $\S = \S \cup \widehat{j}$ *and* $Q = Q + 1$.
4. *If* $Q = \widehat{s}$*, then stop. Otherwise, proceed to Step 2.*

Alternatively, the algorithm can be stopped when adding an extra factor does not significantly increase the R^2, i.e. $R^2_{FM}(\S \cup \widehat{j}) - R^2_{FM}(\S) \leq \epsilon$ for some small value of the tolerance parameter ϵ and $\widehat{j}$ appearing on Step 2 of the algorithm. In practice, we use $\epsilon = 1\%$. Also, some factors that are deemed to be important, e.g. Fama-French factors, can be included in the algorithm on the

first step, i.e. we can set $\S = \S_0$ for some pre-specified set of factors $\S_0$ on Step 1 of the algorithm instead of setting $\S = \emptyset$.

The result of Algorithm 2.1 is the set of factors $\S$, which we denote as $\widehat{\S}$. Using this set, we estimate the vector of SDF loadings by $\widehat{\psi} = (\widehat{\psi}_1, \ldots, \widehat{\psi}_p)^\top$, where we set

$$\widehat{\psi}_{\widehat{\S}} = \left(\frac{1}{T} \sum_{t=1}^{T} (f_{t,\widehat{\S}} - \bar{f}_{\widehat{\S}})(f_{t,\widehat{\S}} - \bar{f}_{\widehat{\S}})^\top \right)^{-1} \left(\sum_{i=1}^{N} \widehat{\beta}_{i,\widehat{\S}} \widehat{\beta}_{i,\widehat{\S}}^\top \right)^{-1} \left(\sum_{i=1}^{N} \widehat{\beta}_{i,\widehat{\S}} \bar{r}_i \right) \tag{6}$$

and $\widehat{\psi}_j = 0$ for all $j \notin \widehat{\S}$. Also, substituting (5) into (6), it follows that $\widehat{\psi}_{\widehat{\S}}$ can be rewritten as

$$\widehat{\psi}_{\widehat{\S}} = \left(\sum_{i=1}^{N} \widehat{C}_{i,\widehat{\S}} \widehat{C}_{i,\widehat{\S}}^\top \right)^{-1} \left(\sum_{i=1}^{N} \widehat{C}_{i,\widehat{\S}} \bar{r}_i \right), \tag{7}$$

which means that $\widehat{\psi}_{\widehat{\S}}$ can be computed by cross-sectional OLS regression of $\bar{r}_i$ on $\widehat{C}_{i,\widehat{\S}}$. We will prove consistency and derive the rate of convergence of this estimator in Section **??** below.

Remark 2.1 (Equivalent Forms of Forward Selection Fama-MacBeth Regression)**.** *For each $i \in [N]$ and $j \in [p]$, let $\widehat{C}_{i,j} = T^{-1} \sum_{t=1}^{T} (f_{t,j} - \bar{f}_j) r_{i,t}$ be the time series sample covariance between the factor $f_{t,j}$ and the return $r_{i,t}$. Observe that as long as the matrix $\mathbb{E}[(f_t - \mathbb{E}[f_t])(f_t - \mathbb{E}[f_t])^\top]$ is non-singular, the number of factors $\widehat{s}$ used in Algorithm 2.1 is relatively small, and some other mild regularity conditions are satisfied, the matrices*

$$T^{-1} \sum_{t=1}^{T} (f_{t,\S} - \bar{f}_{\S})(f_{t,\S} - \bar{f}_{\S})^\top$$

*will be non-singular for all $\S$ used in Algorithm 2.1 with high probability by a suitable version of the law of large numbers. However, whenever this happens, $R^2_{FM}(\S)$ will be numerically equal to the R^2 in the cross-sectional OLS regression of $\bar{r}_i$ on $\widehat{C}_{i,\S}$, where $\widehat{C}_{i,\S} = (\widehat{C}_{i,j})_{j \in \S}$. Thus, we can equivalently maximize the R^2 in the latter regression using Algorithm 2.1 to obtain the same set of factors $\widehat{\S}$. This equivalent form of Algorithm 2.1, together with equation (**??**), explains that the algorithm is expected to select factors whose SDF loadings are sufficiently different from zero. The*

algorithm, however, may fail to select factors whose SDF loadings are close to zero, at least in finite samples.

Another equivalent form of Algorithm 2.1 can be obtained by replacing the multivariate betas (5) in the cross-sectional OLS regression by univariate betas. Indeed, for each $i \in [N]$ and $j \in [p]$, let $\widetilde{\beta}_{i,j}$ be the slope coefficient in the univariate time series OLS regression of $r_{i,t}$ on $f_{t,j}$, with intercept included:

$$\widetilde{\beta}_{i,j} = \left(\sum_{t=1}^{T}(f_{t,j} - \bar{f}_j)^2\right)^{-1}\left(\sum_{t=1}^{T}(f_{t,j} - \bar{f}_j)r_{i,t}\right).$$

Also, for each subset § of $[p]$, set $\widetilde{\beta}_{i,\S} = (\widetilde{\beta}_{i,j})_{j\in\S}$. Then observe that under the same regularity conditions as above, we will have $T^{-1}\sum_{t=1}^{T}(f_{t,j} - \bar{f}_j)^2 > 0$ for all $j \in [p]$ with high probability, and whenever this happens, the R^2 in the cross-sectional OLS regression of $\bar{r}_i$ on $\widetilde{\beta}_{i,\S}$ will be equal to the R^2 in the cross-sectional OLS regression of $\bar{r}_i$ on $\widehat{C}_{i,\S}$, yielding another equivalent form of Algorithm 2.1.

□

2.3 Debiased Estimation

Here, we explain how to debias the estimator $\widehat{\psi} = (\widehat{\psi}_1, \ldots, \widehat{\psi}_p)^\top$. To develop intuition for the debiasing procedure, fix any component of the vector $\widehat{\psi}$, say $\widehat{\psi}_j$, and observe that, as discussed in the previous section, Algorithm 2.1 may not select factors whose SDF loadings are close to zero. The estimator $\widehat{\psi}_j$ may thus be biased for two reasons. First, ψ_j may be close to zero itself, in which case Algorithm 2.1 may not be able to select it, leading to the bias in $\widehat{\psi}_j$ as long as $\psi_j \neq 0$. Second, even if ψ_j is not close to zero and the factor j is selected, Algorithm 2.1 may fail to select other factors k whose covariance with returns $C_{i,k}$ is cross-sectionally correlated with $C_{i,j}$, leading to the omitted variable bias in $\widehat{\psi}_j$ as long as $\psi_k \neq 0$, which follows from the OLS characterization of the estimator $\widehat{\psi}_{\widehat{\S}}$ in (7).

Two reasons for the bias of $\widehat{\psi}_j$ suggest that we can debias it by appropriately expanding the set $\widehat{\S}$ used to construct $\widehat{\psi}$ in (6). Specifically, we need to add to the set $\widehat{\S}$ the factor j itself and all factors k whose covariance with returns $C_{i,k}$ is cross-sectionally correlated with $C_{i,k}$. In practice,

since neither $C_{i,j}$ nor $C_{i,k}$ is observed, we work with sample covariances $\widehat{C}_{i,j}$ and $\widehat{C}_{i,k}$.

More formally, we proceed as follows. First, let $\widehat{\S}_1$ be the set § produced by Algorithm 2.1 and let $\widehat{\S}_2$ be the set produced by Algorithm 2.1 with $R^2_{FM}(\S)$ replaced by the R^2 in the cross-sectional OLS regression of $\widehat{C}_{i,j}$ on $\widehat{C}_{i,\S}$ (without intercept) and with $[p]$ replaced by $[p] \setminus \{j\}$. Second, set $\widehat{\S}_D = \widehat{\S}_1 \cup \widehat{\S}_2 \cup \{j\}$. Third, let $\widehat{\psi}$ be the estimator defined in Section 2.2 but with the set $\widehat{\S}$ there replaced by the $\widehat{\S}_D$ here. The corresponding component $\widehat{\psi}_j$ of the vector $\widehat{\psi}$ will be a debiased estimator of ψ_j, which we denote as $\widehat{\psi}_{D,j}$. By repeating this procedure for all $j \in [p]$, we obtain the full vector $\widehat{\psi}_D = (\widehat{\psi}_{D,1}, \ldots, \widehat{\psi}_{D,p})^\top$. In the next section, we will prove that for each $j \in [p]$, the estimator $\widehat{\psi}_{D,j}$ is $\sqrt{T}$-consistent and $\sqrt{T}(\widehat{\psi}_{D,j} - \psi_j)$ converges to a centered normal distribution. We will also provide a consistent estimator of the asymptotic variance.

2.4 Asymptotic Theory

To formulate asymptotic results for the estimator $\widehat{\psi}_D$, for all $j \in [p]$, define

$$\eta_j = \left(\mathbb{E}[(f_{t,-j} - \mathbb{E}[f_{t,-j}])(f_{t,-j} - \mathbb{E}[f_{t,-j}])^\top)]\right)^{-1} \mathbb{E}[(f_{t,-j} - \mathbb{E}[f_{t,-j}])f_{t,j}],$$

where $f_{t,-j} = (f_{t,1}, \ldots, f_{t,j-1}, f_{t,j+1}, \ldots, f_{t,p})^\top \in \mathbb{R}^{p-1}$. Also, for all $t \in [T]$, let

$$z_{t,j} = f_{t,j} - \mathbb{E}[f_{t,j}] - \eta^\top (f_{t,-j} - \mathbb{E}[f_{t,-j}])$$

be the residual from the least-squares projection of $f_{t,j} - \mathbb{E}[f_{t,j}]$ on $f_{t,-j} - \mathbb{E}[f_{t,-j}]$. In addition, let $\sigma^2_{z,j} = \mathbb{E}[z^2_{t,j}]$ be the variance of $z_{t,j}$.

Theorem 2.1. *Under appropriate regularity conditions, for each $j \in [p]$,*

$$\sqrt{T}(\widehat{\psi}_{D,j} - \psi_j) \to_d N(0, \sigma^2_{\psi,j}),$$

where

$$\sigma^2_{\psi,j} = \lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \sum_{s=1}^{T} \mathbb{E}\left[\frac{m_t z_{t,j} m_s z_{s,j}}{\sigma^4_{z,j}}\right]$$

Remark 2.2 (Relation to Feng et al. (2020))**.** *The asymptotic variance* $\sigma^2_{\psi,j}$ *appearing in this theorem is the same as the asymptotic variance in Theorem 1 of Feng et al. (2020). In fact, one can show that our debiased estimator* $\widehat{\psi}_{D,j}$ *is asymptotically equivalent to the corresponding estimator in that paper. The estimators, however, are expected to be different in finite samples.* □

3 Empirical Analysis

In this section, we present the data used in the empirical analysis and the main results of the selection of the high-order factors by the forward Fama-MacBeth (FMB) method.

3.1 Data

We present a description of the data used for the estimation of the forward selection FMB method. In order to estimate the model and perform cross-sectional asset pricing tests, we use data from various sources. Test assets are 484 characteristic-managed portfolios from Kozak et al. (2020).[1] These portfolios are constructed by weighting stocks based on their characteristic signals, which are based on cross-sectional ranks. Small firms (below 0.01% of market cap) are excluded. Baseline factors are the Fama and French (2015)'s five factors and the momentum factor of Jegadeesh and Titman (1993), which we obtain from Kenneth French's data library. We refer to these six baseline factors as the FF5M model. Furthermore, we consider factors for 153 characteristics in 13 themes, using data from the US equity market, constructed by Jensen, Kelly and Pedersen (2023).[2] We refer to these 153 factors as the factor zoo (see, e.g., Feng, Giglio and Xiu (2020)). The sample starts in October 1973 and ends in December 2019.

We construct factors based on the FF5M factors. We denote both the higher-orders and the interactions of FF5M as higher-order factors. Specifically, for any factor f_i in the FF5M, we first form all the f_i^2 and f_i^3 as candidate factors. Moreover, we form all the pairwise interactions of

[1] We obtain the portfolio data from Serhiy Kozak's web page at: https://serhiykozak.com/#/data-code.

[2] The data for the factor zoo is available through the Global Factor Data website at https://jkpfactors.com/. We use capped value weighted data for all factors and all countries av.

degree 2 and 3. That is, for any two factors f_i and f_j in the FF5M model, we form the pairwise interactions of the type $f_i \times f_j$ and $f_i \times f_j^2$ as candidate factors. Therefore, the total number of candidate higher-order factors is 57, where there are 6 f_i^2, 6 f_i^3, 15 $f_i \times f_j$, and 30 $f_i \times f_j^2$.

Table 1 presents summary statistics of the test assets and factors. For the mean, the standard deviation, and the quantiles, we report averages across assets. All returns are monthly, in excess of the US risk-free rate, and reported in percentage. The average mean return for the test assets is 0.65% monthly, with a standard deviation of 4.96%. The average return for the FF5M factors is lower and equal to 0.39% monthly, with a standard deviation of 3.16%. The average return for the factor zoo is even lower and equal to 0.23% monthly, with a standard deviation of 2.71%. The higher-order factors have a distribution characterized by more extreme realizations, as exemplified by the 95% quantile of the powers of degree 2 and 3 equal, respectively, to 42.75% and 167.67%. We note that the higher order factors are generally non-tradable. Figure A.1 in the Appendix plots the distributions of the returns of the high-order factors.

Table 1: **Summary Statistics**

This table reports summary statistics of the tests assets and factors. For the mean, the standard deviation and the quantiles, we report averages across assets. Returns are monthly and in excess of the US risk-free rate. Test assets are from Kozak et al. (2020). The Fama-French five factors and the momentum factor (FF5M) are from Ken French's data library. "Interactions" correspond to pairwise interactions of the FF5M factors of degree 2 (i.e., $f_i \times f_j$) and degree 3 (i.e., $f_i^2 \times f_j$). "Powers" correspond to the power of the individual FF5M factors of degree 2 (i.e., f_i^2) and degree 3 (i.e., f_i^3). The factor zoo includes 153 factors based on US equity data constructed by Jensen et al. (2023). The sample period is October 1973 to December 2019.

$R_{i,t}$ (%)	Mean	Std	$q_{5\%}$	$q_{50\%}$	$q_{95\%}$	N	T
Test Assets	0.65	4.97	-7.58	0.9	8.1	484	555
FF5M	0.39	3.16	-4.63	0.4	5.13	6	555
Interactions degree 2	-0.48	16.47	-17.82	-0.02	14.12	15	555
Interactions degree 3	0.53	20.14	-6.6	0.02	8.44	30	555
Powers degree 2	11.1	28.62	0.03	3.34	42.75	6	555
Powers degree 3	-9.83	581.32	-147.67	0.23	167.67	6	555
Factor zoo	0.23	2.71	-3.5	0.19	4.14	153	555

3.2 Selected Factors by Forward Selection Fama-MacBeth Regressions

Our goal is to study whether and which of the higher-order factors constructed from the six factors in the baseline FF5M are important in explaining the SDF. To select the most relevant higher-order factors to the SDF, we use the forward selection FMB method that we have discussed in detail in Section 2 above. We start with the six baseline FF5M factors and add the higher-order interactions one at a time.

We present the main estimation results from the second step of the forward selection FMB method in Table 2. Panel A shows the performance of the baseline models using the 484 characteristic-managed portfolios from Kozak et al. (2020). The baseline models include the standard CAPM, the Fama and French (1993)'s 3-factor model (FF3), the Fama and French (2015)'s 5-factor model (FF5) and the 5-factor model augmented with the momentum factor of Jegadeesh and Titman (1993) (FF5M). The adjusted cross-sectional R-squared for the CAPM is just 0.032, and the estimated intercept (0.009) is significantly different from zero with a t-statistic of 3.6. The adjusted cross-sectional R-squared reaches 0.312 for the FF5M model, while the estimate for the intercept (0.004) remains significantly different from zero with a t-statistic of 2.2.

Panel B refers to the forward selection FMB method. At each step, we augment the FF5M by one of the additional higher-order factors, which we select in order to maximize the adjusted cross-sectional R-squared. The selection procedure stops when the improvement in the R-squared in a given step is smaller than 1 pp. The forward FMB method selects 7 higher-order factors: the SMB factor squared (SMB2), the interaction between SMB2 and the momentum factor (SMB2*Mom), the interaction between the momentum factor squared and the profitability factor (Mom2*RMW), the market factor squared (Mkt-RF2), the interaction between Mkt-RF2 and the profitability factor (Mkt-RF2*RMW), the interaction between the market and size factors (Mkt-RF*SMB) and the interaction between the value factor squared and the market factor (HML2*Mkt-RF). Of these 7 higher-order factors that are selected, 2 are higher-orders (SMB2 and Mkt-RF2) and 5 are interactions (SMB2*Mom, Mom2*RMW, Mkt-RF2*RMW, Mkt-RF*SMB, and HML2*Mkt-RF).

Interestingly, none of the higher-orders of degree 3 are selected.

The adjusted cross-sectional R-squared increases from 31.2% for the baseline FF5M model to 41% in the first step when SMB2 is included, and gradually increases to 59%, in the last step when all 7 factors are included and the procedure stops. Across the seven steps, the estimate for the intercept is around 0.003, with a t-statistic in the final step of 1.819, which is below the threshold for significance at the 5% level.

Table 2: **Cross-Sectional Performance**

This table reports the results from the second step of the FMB method. Panel A refers to baseline models: the standard CAPM, the Fama-French 3-factor, the Fama-French 5-factor and the Fama-French 5-factor and momentum models. The table reports the adjusted cross-sectional R-squared, the estimate for the intercept (α) and associated t-statistic. The t-statistics is based on Newey-West corrected standard errors. Panel B refers to the forward selection Fama-MacBeth regression method. At each step, we augment the FF5M by one additional higher-order factor (h_j) which maximizes the adjusted cross-sectional R-squared. The method stops when the improvement in the R-squared in a given step is smaller than 1 pp.

Panel A: Baseline Models				
#	Model	Adj. R-squared	α	t-stat (α)
1	CAPM	0.032	0.009	3.635
2	FF3	0.126	0.01	4.701
3	FF5	0.275	0.006	2.899
4	FF5M	0.312	0.004	2.201

Panel B: Forward Selection Fama-MacBeth Regression Method				
Step	h_j	Adj. R-squared	α	t-stat (α)
1	SMB2	0.41	0.003	1.821
2	SMB2*Mom	0.467	0.004	2.02
3	Mom2*RMW	0.495	0.003	1.324
4	Mkt-RF2	0.529	0.003	1.769
5	Mkt-RF2*RMW	0.552	0.003	1.636
6	Mkt-Rf*SMB	0.572	0.004	2.161
7	HML2*Mkt-RF	0.587	0.003	1.819

We then examine the SDF loadings of the higher-order factors selected by the forward selection FMB method. The results are summarized in Table 3. Panel A shows the non-debiased estimates

based on the standard FM two-pass regression, while Panel B presents the debiased estimates as outlined in Section 2.3. The directions of the point estimates based on the non-debiased method and the debiased method are the same. Moreover, the estimates from the two methods are similar in magnitudes. The estimates based on the non-debiased method are all significant at the 5% level, while the estimates based on the debiased method are all significant at the 10% level with 5 of the 7 SDF loading estimates being significant at the 5% level. Overall, the results highlight that the higher-order factors are important components of the SDF and significantly price assets in the cross-section.

Table 3: **SDF loadings**

This table reports the estimates for the SDF loadings (b_j) associated with the selected higher-order factors. Panel A reports non-debiased estimates while Panel B reports debiased estimates. The t-statistics is based on Newey-West corrected standard errors. High-order factors are selected by the forward selection Fama-MacBeth regression method (see Table 2).

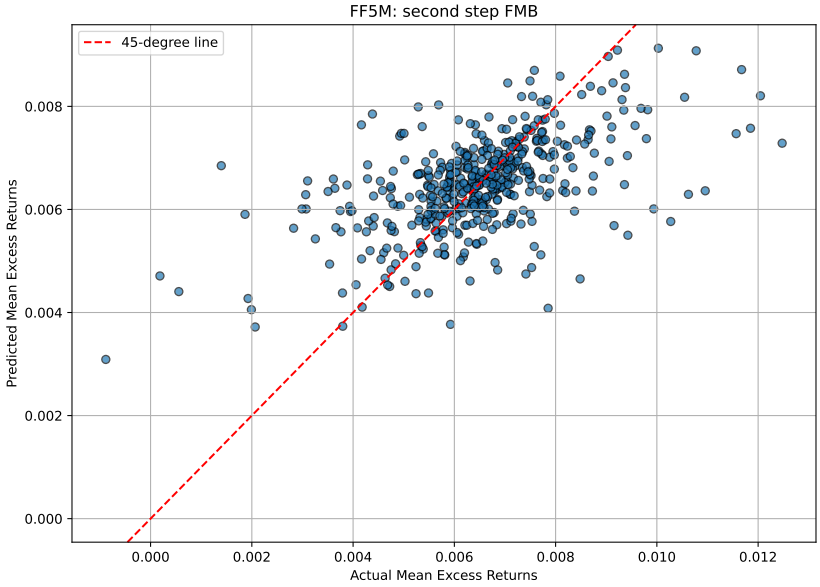
	Panel A: Non-Debiased Estimates		Panel B: Debiased Estimates	
h_j	b_j	$t(b_j)$	b_j	$t(b_j)$
SMB2	1.158	2.43	1.403	1.837
SMB2*Mom	-1.193	-4.074	-1.004	-1.812
Mom2*RMW	-1.38	-4.105	-1.178	-2.882
Mkt-RF2	1.637	5.619	1.878	4.606
Mkt-RF2*RMW	-3.53	-4.169	-3.87	-4.301
Mkt-RF*SMB	-1.404	-3.032	-2.012	-3.636
HML2*Mkt-RF	-1.385	-2.823	-1.606	-2.039

Figure 1 plots predicted against realized returns for all the test assets. Panel A shows the performance of the FF5M model and panel B presents the higher-order model. The FF5M has a limited ability in explaining the average returns (Panel A), in line with the relatively low R-squareds of the model. On the other hand, there is a strong positive relationship between the average realized returns and the predicted returns for the higher-order model selected by the forward selection FMB method.

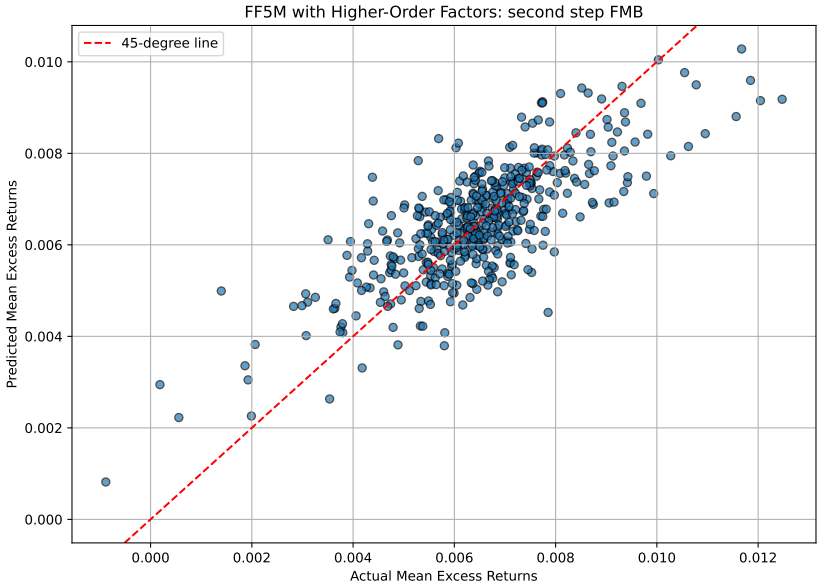
Figure 1: **Actual vs. Predicted Returns**

This figure plots actual mean returns of the test assets against the mean predicted returns from the FF5M model (subplot a) and the FF5M model augmented with the higher-order factors selected by the forward selection Fama-MacBeth regression method (subplot b, see 2). The predicted return is the OLS estimate of the betas times the estimated prices of risk.

(a) **FF5M model**



(b) **FF5M model with higher-order factors**



We note that the predicted return used in the graph is the OLS estimate of the betas from the first step of the two-pass FMB method times the estimated prices of risk. That is, we do not impose the no-arbitrage condition for the tradable factors that pins down the price of risk by the sample mean of the factor. Doing so would, by construction, result in an even worse fit by construction for the baseline FF5M models where all the factors are tradable. We present the results in Table A.1 in the Online Appendix where we impose the restriction for the tradable factors. The results confirm the importance of higher-order factors in pricing assets in the cross-section.

3.3 Out-of-Sample

The main results presented in Table 2 are based on in-sample estimation. We further provide evidence of the out-of-sample performance of the higher-order model selected by the forward selection FMB method. We consider two types of out-of-sample tests: (1) out-of-sample in terms of asset space and (2) out-of-sample in terms of time period estimations.

3.3.1 Asset space

In Table 4, we present an additional set of results based on cross-validation applied to the forward selection FMB method. The cross-validation is applied to the second step of the forward selection FMB procedure. This is an out-of-sample exercise in terms of asset space, where we select higher-order factors and estimate SDF loadings from the training assets, and then evaluate the out-of-sample performance in the out-of-sample assets.

Specifically, at each step of the forward selection FMB method, we divide the sample in 5 equal-sized, non-overlapping folds. For each fold k, we use the remaining 4 folds as the training set and the left-out fold as the test set. In each training step, we estimate the cross-sectional regression of average returns on the betas and obtain estimates for the risk prices. Next, we use the estimated risk prices to compute the adjusted R-squared on the test set. This process is repeated 5 times and the adjusted cross-sectional R-squared associated with a given higher-order factor is the average, at each step, of the adjusted cross-sectional R-squareds obtained in the 5-folds cross-

validation. Table 4 reveals that the forward selection FMB method with cross-validation (CV) selects a total of 8 higher-order factors, with a cross-sectional R-squared in the final step of 0.596. Moreover, the first six selected factors are the same of the baseline method, while the last selected factor corresponds to the last from the baseline method. Furthermore, the cross-validated adjusted cross-sectional R-squareds, which represent out-of-sample measures of fit of the cross-sectional model, is close to the in-sample R-squareds at 0.535. Overall, these results show that the selected higher-order factor model has stable out-of-sample performance.

Table 4: **Forward Selection FMB Regression Method with Cross-Validation**

This table reports the results from the second step of the FMB regression method for the forward selection Fama-MacBeth regression method with 5-fold cross-validation. At each step, we include one additional higher-order factor which maximizes the adjusted cross-sectional R-squared computed with cross-validation using 5-folds. The method stops when the improvement in the R-squared in a given step is smaller than 1 pp. For each step, we report the selected factor (h_j), the adjusted cross-sectional R-squared, the adjusted cross-validated (CV) cross-sectional R-squared, the estimate for the intercept (α) and associated t-statistic. The t-statistic is based on Newey-West corrected standard errors. The adjusted cross-validated cross-sectional R-squared is the average at each step of the adjusted R-squareds obtained in the 5-folds cross-validation.

Step	h_j	Adj. R-squared	Adj. CV R-squared	α	t-stat (α)
1	SMB2	0.41	0.366	0.003	1.821
2	SMB2*Mom	0.467	0.422	0.004	2.02
3	Mom2*RMW	0.495	0.443	0.003	1.324
4	Mkt-RF2	0.529	0.472	0.003	1.769
5	Mkt-RF2*RMW	0.552	0.495	0.003	1.636
6	Mkt-Rf*SMB	0.572	0.511	0.004	2.161
7	Mom2*HML	0.582	0.527	0.004	2.145
8	HML2*Mkt-RF	0.596	0.535	0.003	1.833

3.3.2 Time periods

We next evaluate the out-of-sample performance in the time-series of the higher-order factor model against the benchmark factor models (i.e., CAPM, FF3, FF5, and FF5M). Specifically, we first split the sample into two 50-50 subsamples based on the time series, where we denote the

first half of the sample as the training sample and the second half as the out-of-sample. Then, we estimate the covariances and the SDF loadings in the training sample based on the different factor models. We denote the R^2 from these estimations as R^2_{train}. Lastly, we compare the average asset returns in the out-of-sample to the predicted asset returns using covaraince and the SDF loadings estimated from the training sample to obtain the out-of-sample R^2, or R^2_{oos}. We note that both the R^2_{train} and the R^2_{oos} are likely to be lower than the full sample R^2 for any model due to estimation errors, especially related to average returns. It is well-known that estimations of the average equity returns are noisy and require a long time series. Therefore, our focus is comparing R^2_{oos} for the higher-order factor model against those of the benchmark factor models.

Table 5 presents the results. The R^2_{oos}s of CAPM and the FF3 factor model are close to zero. The R^2_{oos}s of the FF5 factor model and the FF5M factor model are 7% and 8.9%, representing an improvement from the CAPM and the FF3 factor model. The R^2_{oos} of higher-order factor model selected from the forward selection FMB regression is significantly larger than the benchmark factor models at 15.8%, which is about twice as high as the R^2_{oos}s of FF5 and FF5M factor models. Additionally, in untabuldated result, we repeat the exercise but select the higher-order factors in the training sample. The magnitude of the R^2_{oos} is close to Column (5) at 14.8%.

One concern is that the 50-50 split where the first half of the sample is the training sample and the second half of the sample is the out-of-sample may be arbitrary. As a robustness check, we repeat the analyses by randomly selecting, for 1,000 times, 50% of the sample as the training sample and the remaining as the out-of-sample. We report the results in Table A.2 in the Online Appendix. Table A.2 delivers the same message – theR^2_{oos} of the higher-order factor model is significantly higher than those of the benchmark factor models.

Table 5: **Comparison of R-squareds**

This table presents the in-sample and out-of-sample R^2 under different factor models. The factor models include the CAPM, the Fama-French 3-factor model (FF3), the Fama-French 5-factor model (FF5), the Fama-French 5-factor plus momentum factor model (FF5M), and the higher-order factor model selected by the forward selection FMB method. We first split the sample into two 50-50 subsamples based on the time series, where we denote the first half of the sample as the training sample and the second half as the out-of-sample. Then, we estimate the covariances and the SDF loadings in the training sample based on the different factor models. We denote the R^2 from these estimations as R^2_{train}. Lastly, we compare the average asset returns in the out-of-sample to the predicted asset returns using covaraiances and the SDF loadings estimated from the training sample to obtain the out-of-sample R^2, or R^2_{oos}. The table reports the R^2_{train} and R^2_{oos} for each model.

	(1) CAPM	(2) FF3	(3) FF5	(4) FF5M	(5) Higher-Order
R^2_{train}	0.003	0.204	0.273	0.312	0.449
R^2_{oos}	0.001	0.012	0.07	0.089	0.158

4 Factor Zoo

In this section, we show that the baseline FF5M model, augmented with the higher-order factors selected by the forward selection FMB method, effectively subsumes the majority of the factors in the factor zoo, or zoo factors.

We employ two empirical methods to support this argument. First, we perform Fama-MacBeth regressions to test the significance of the SDF loading of each factor in the factor zoo, controlling for our higher-order factors. Second, we construct factor mimicking portfolios for the seven higher-order factors to examine conduct time-series asset pricing tests on the zoo factors.

4.1 Cross-sectional method

The first empirical strategy to evaluate whether the high-order factors can subsume the zoo factors is based on the two-pass Fama-MacBeth procedure, motivated by the analysis in Feng et

al. (2020). Specifically, for each zoo factor, we estimate its SDF loading while controlling for our higher-order factor model, which includes the FF5M and the selected higher-order factors.

Figure 2 summarizes the results. Panel A plots the distribution of the absolute values of the estimates for the intercept from the cross-sectional regressions. Panel B displays the distribution of the absolute t-statistics associated with the SDF loading of the factor zoo, using Newey and West (1987) corrected standard errors. In both panels, the blue histogram represents the estimates obtained by controlling for the FF5M factors, while the green histogram represents the estimates obtained by additionally controlling for the higher-order factors selected by the forward regression method. Vertical dashed lines indicate the median values of each distribution – a vertical blue dashed-line for estimates obtained controlling for the FF5M factors and a vertical green dashed-line for estimates obtained controlling for the higher-order factor model. Additionally, we also include a vertical red dashed-line corresponds to the median value with no controls, or the factor model includes only the factor in the factor zoo.

Panel A shows that the intercept estimates tend to be highly statistically significant, with a median value of 4.00, when the factor model includes only the factor in the factor zoo. When the factor model includes both FF5M and the factor in the factor zoo, the median t-statistic decreases substantially but remains significant at the 5% level at 2.23. When the factor model further includes the higher-order factors, the median t-statistic is no longer statistically significant at the 5% level at 1.84.

Panel B shows that, when the factor model includes only the factor in the factor zoo, the median SDF loading estimate for the factor zoo is not statistically different from zero, with a t-statistic of 0.86. Note that these estimates are likely biased because of omitted factors (see, e.g., Feng et al. (2020)). Controlling for the FF5M, the median t-statistic increases substantially to 1.38. 34% of zoo factors become statistically significant at the 5% level with a t-statistic greater than 2 suggesting that FF5M helps to remove some bias in the factor estimates. Importantly, when we additionally control for the high-order factors, we find that for 95% of the zoo factors the risk price estimate is not statistically different from zero, with the median t-statistic dropping from

1.38 to only 0.71. Specifically, we find that only 7 out of 148 zoo factors remain significant at the 5% level. Of these seven factors, three belong to the "low leverage" category of Jensen et al. (2023) (*cash_at*, *rd5_at*, and *rd_sale*). The remaining significant factors are in the categories "size" (*rd_me*), "seasonality" (*seas_6_10an*), "accruals" (*cowc_grla*) and "quality" (*qmj*). Overall, these results reveal that the high-order factors selected by the forward regression method effectively subsume the majority of the zoo factors.

4.2 Time-series method

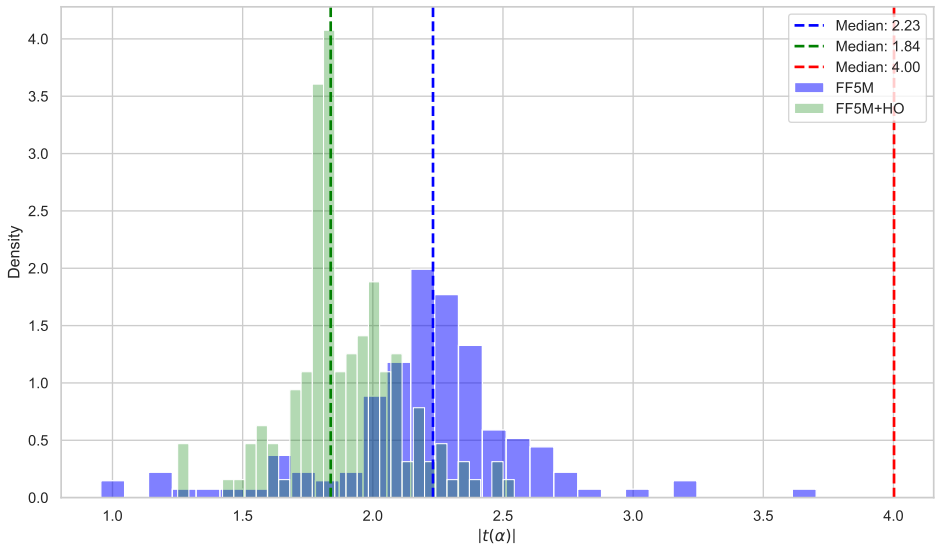
The second empirical strategy we employ to evaluate whether the high-order factors can subsume the factor zoo is based on time-series regressions, where we study whether the zoo factors can be spanned by the FF5M and the selected high-order factors.

Note that although the FF5M factors are tradable, the high-order factors are generally not directly tradable. Therefore, in order to estimate time-series spanning regressions, we must first construct factor mimicking portfolios of the higher-order factors based on tradable assets. In order to do so, we project each high-order factors onto the entire factor zoo and obtain its factor-mimicking portfolio as the regression's predicted values. This factor-mimicking portfolios only capture the information in the factor zoo that are related to the higher-order factors. The benefit of the factor-mimicking portfolio approach is that we are able to convert the non-tradable higher-order factors to tradable factors, but we note that the factor-mimicking portfolios of the higher-order factors inevitably loses information from the original factors. Figure A.2 shows the adjusted R^2 of projecting each selected higher-order factor onto the factor zoo. Mkt-RF*SMB has the lowest adjusted R^2 which is less than 40% and SMB2*Mom has the highest adjusted R^2 at more than 80%, suggesting the projections capture importance fractions, though to various degrees, of variations in the higher-order factors. Nevertheless, we use them to check if they capture pricing information of each factor.

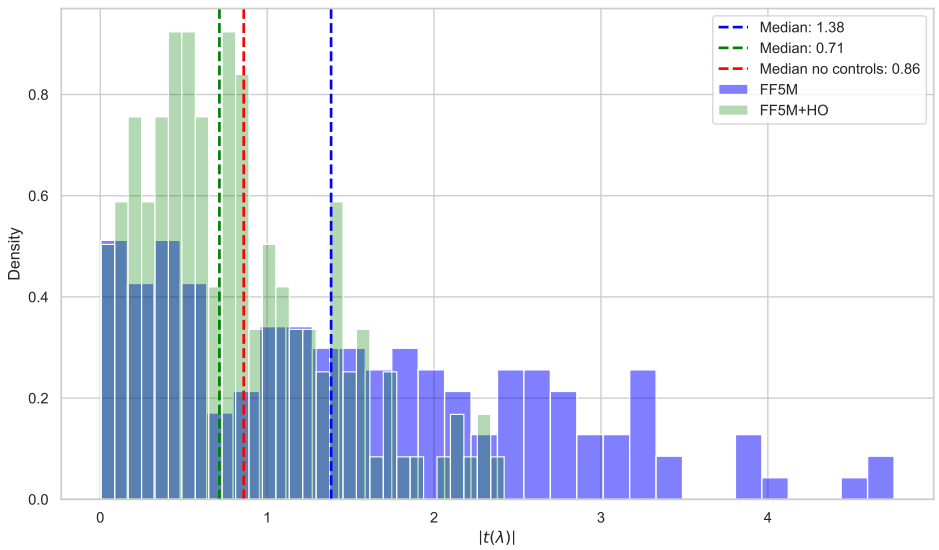
Figure 2: **Reducing the Factor Zoo in the Cross-Section**

This figure plots the distribution of absolute t-statistics associated with the risk prices ($t(\lambda)$) and the intercepts ($t(\alpha)$) of the factor zoo from Jensen et al. (2023) based on US equity data. For each zoo factor, we use the FMB method to estimate its risk price controlling for the FF5M (blue histogram) and the FF5M with the high-order factors selected by the forward selection FMB method (green histogram). The dashed-lines in blue and green correspond to the median absolute t-statistic for the FF5M and for the FF5M with high-order factors models, respectively. The dashed-red-line corresponds to the median absolute t-statistic with no controls. The t-statistics are based on Newey-West corrected standard errors.

(a) **Distribution of** $|t(\alpha)|$



(b) **Distribution of** $|t(\lambda)|$



To do so, we estimate a time-series regression for each factor in the factor zoo on the FF5M and the factor-mimicking portfolios of the high-order factors. The object of interest is the intercept. If the higher-order factors indeed span a given factor in the factor zoo, we expect to find a small and statistically insignificant intercept.

Figure 3 summarizes the results of the time-series spanning regressions. Panel A plots the distribution of the absolute t-statistics associated with the intercepts, and Panel B shows the distribution of the absolute magnitudes for the intercepts. The latter are annualized and expressed in percentage points.

Panel A shows that the median of the absolute t-statistic of the average returns for the zoo factors is statistically significant at the 5% level at 2.08. Controlling for the FF5M brings the median absolute t-statistic from 2.06 to 1.87, a decline of about 9%. Controlling for the high-order factors brings the median absolute t-statistic further down to 1.49, a decline of about 21% from 1.87. Panel B reveals a similar pattern. The median of the absolute intercept estimate of the zoo factors drops from 2.92 pp to 1.69 pp controlling for the FF5M model. It further decreases to 1.34 pp when the higher-order factors are included.

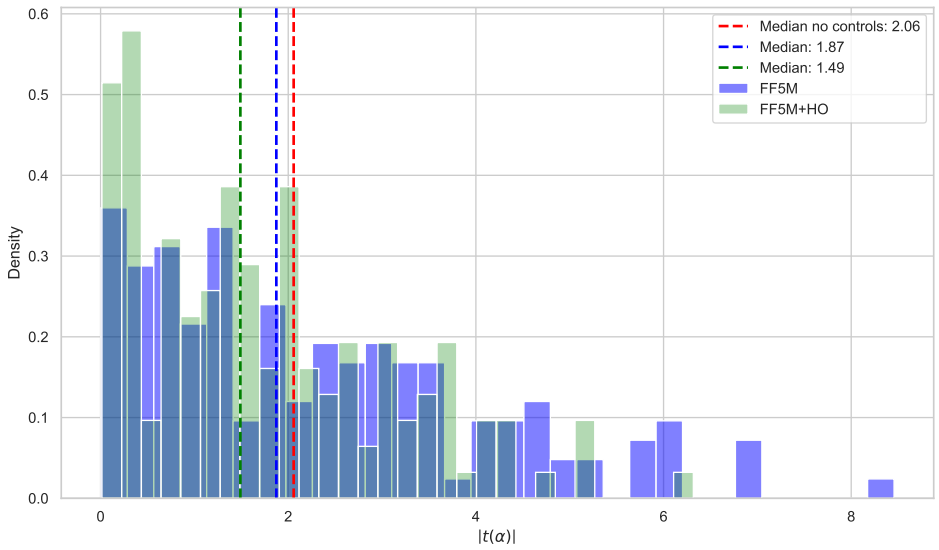
4.3 Relationship between factor zoo and higher-order factors

In the previous subsections, we show that the higher-order factor model selected by the forward selection FMB method is successful in subsuming the majority of the zoo factors. Mechanically, the ability of the selected higher-order factor model to subsume the factor zoo has to come from correlation between the zoo factors and some or all of the factors in the higher-order factor model. In this subsection, we provide information on the loadings of the zoo factors on the factors of the selected higher-order factor model, which sheds light on the relative importance of the higher-order factors in subsuming the zoo factors.

Figure 3: **Reducing the Factor Zoo in the Time-Series**

This figure presents the results of factor spanning time-series regressions for each zoo factor controlling for the FF5M factors (blue histogram) and the FF5M factors with the mimicking portfolios corresponding to the higher-order factors selected by the forward regression method (green histogram). Subplot (a) plots the distribution of the absolute t-statistics for the intercept (α). Subplots (b) plots the distribution of the absolute αs. We denote with dashed vertical lines the median value when controlling for the FF5M factors (blue-dashed-line), for the FF5M with the mimicking portfolios corresponding to the higher-order factors selected by the forward selection Fama-MacBeth regression method (green-dashed-line), and when we include no controls (red-dashed-line). The factors are 153 factors from the factor zoo constructed in Jensen et al. (2023). The t-statistic is based on Newey-West corrected standard errors.

(a) **Distribution of** $|t(\alpha)|$



(b) **Distribution of** $|\alpha|$

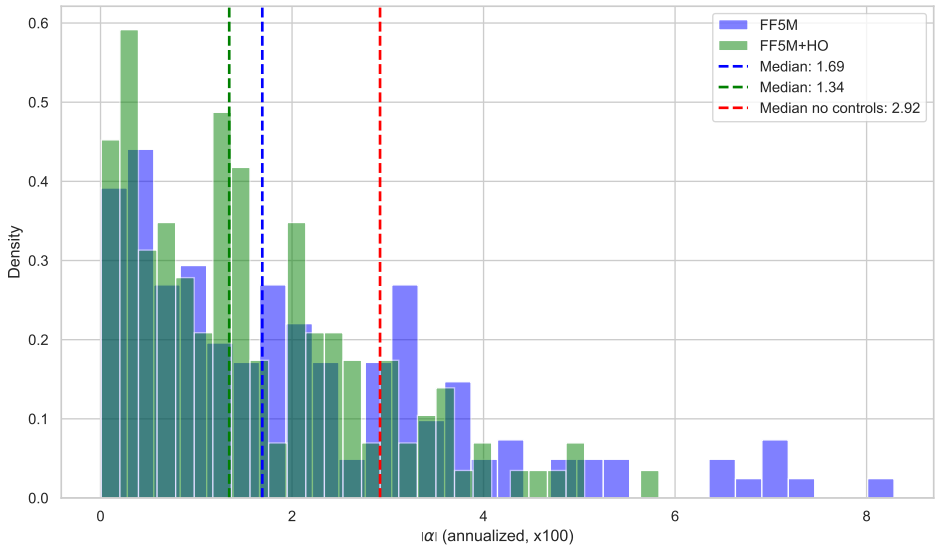


Table 6 documents the fraction of zoo factors with loading significantly different from zero at the 5% confidence level on each of the FF5M factors and the selected higher-order factors selected from the forward selection FM regression. Column (1) shows the fraction of significance for the factors in FF5M model only and Column (2) reports the fraction of significance for the factors in the full higher-order factor model. Column (1) shows that the fractions of significance range from 45.9% for Mkt-RF to 75% for RMW for the 6 factors in FF6M. Column (2) shows that the fractions of significance remain stable for the 6 factors in FF6M, when we further include the selected higher-order factors. The fractions of significance for the higher-order factors range from 27% for Mkt-RF2 to 55.4% for HML2*Mkt-RF, highlighting that it is not the case a specific higher-order factor fully subsumes the zoo factors.

Table 6: **Loadings of the Zoo Factors on Higher-order Factor Model**

This table reports the fraction of zoo factors with loading significantly different from zero at the 5% confidence level on each of the FF5M factors and higher-order factors selected by the forward selection FMB regression method. Significance is evaluated by t-statistics based on Newey-West corrected standard errors. The zoo factors are 148 factors constructed inJensen et al. (2023).

	(1)	(2)
Factor	Frac Sig 5%	Frac Sig 5%
Mkt-RF	0.459	0.459
SMB	0.568	0.628
HML	0.642	0.709
RMW	0.75	0.669
CMA	0.601	0.595
Mom	0.486	0.405
SMB2		0.446
SMB2*Mom		0.331
Mom2*RMW		0.338
Mkt-RF2		0.27
Mkt-RF2*RMW		0.358
Mkt-RF*SMB		0.304
HML2*Mkt-RF		0.554
# zoo factors	148	148

5 Higher-Order Factors and Macro Risk

In this section, we conduct analyses to better understand the economic forces the higher-order factors capture. In the first exercise, we examine the correlations between the selected higher-order factors and two important macroeconomic risks as proxied by the U.S. equity market returns and the CBOE VIX index.

In Table 7, we present the results. Panel A documents the correlations with U.S. equity market returns and Panel B shows the correlations with CBOE VIX index. Beyond correlations over the full sample, we also calculate the correlations during the NBER recession periods and the correlations corresponding to the bottom 10%, the middle 10%-90%, and the top 10% of higher-order factors' return distribution.

For the first two higher-order factors (i.e., SMB2 and SMB2*Mom), the magnitude of the correlations with U.S. equity market returns are markedly higher during the NBER recession periods relative to the full sample. For example, the correlation between SMB2 and U.S. equity market returns is only 0.016 for the full sample, but it jumps to 0.205 during the NBER recession periods. For the last four higher-order factors (i.e., Mkt-RF2, Mkt-RF2*RMW, Mkt-RF*SMB, and HML2*Mkt-RF), they tend to have high correlation in magnitude with the VIX index, especially during the NBER recession periods. For example, the correlation between Mkt-RF*RMW and VIX is 0.32 for the full sample and increases to 0.515 during the NBER recession periods. These results show that the selected higher-order factors capture different macroeconomic risks, especially during economic downturns.

Table 7: **Correlations with Macroeconomic Risk**

This table reports the correlation coefficients for each high-order factor selected by the forward regression method with the U.S. equity market return (Panel A) and with the CBOE VIX index. We present both unconditional (Full) correlation coefficients and conditional correlation coefficients for specific subsamples: one containing U.S. NBER recession periods, and others corresponding to the bottom and top 10% of each factor's return distribution, as well as the remaining 10%–90% quantile. Correlation coefficients involving the VIX index are estimated starting in January 1990, reflecting the index's availability from that date onward.

Factor (h_j)/Sample	Full	NBER Recessions	$q_{\to 0.1}(h_j)$	$q_{0.1\to 0.9}(h_j)$	$q_{0.9\to}(h_j)$
Panel A: Correlation with US Equity Market					
SMB2	0.016	0.205	-0.062	-0.01	0.1
SMB2*Mom	-0.032	-0.345	-0.119	-0.018	0.085
Mom2*RMW	-0.113	-0.109	-0.085	-0.077	-0.001
Mkt-RF2	-0.164	-0.108	0.388	0.199	-0.267
Mkt-RF2*RMW	-0.372	-0.287	0.118	0.027	-0.479
Mkt-RF*SMB	-0.206	-0.136	-0.293	-0.009	-0.299
HML2*Mkt-RF	0.524	0.636	0.357	0.64	0.607
Panel B: Correlation with VIX Index					
SMB2	0.134	0.03	0.076	0.189	0.005
SMB2*Mom	0.001	-0.07	-0.233	0.043	-0.053
Mom2*RMW	0.015	0.116	-0.07	0.195	0.236
Mkt-RF2	0.603	0.697	0.141	0.284	0.515
Mkt-RF2*RMW	0.32	0.515	-0.255	0.138	0.376
Mkt-RF*SMB	0.257	0.486	-0.298	0.197	0.438
HML2*Mkt-RF	-0.239	-0.437	-0.266	-0.334	0.104

Additionally, in Table 8, we examine the exposures of the selected higher-order factors to the common macroeconomic factors. The adjusted R^2s are relatively higher and above 20% for 5 of the 7 higher-order factors, but the adjusted R^2s for the SMB2*Mom and the Mom2*RMW factors are low at 9% and 1%, respectively. In particular, Mom2*RMW is not significantly exposed to any of the macroeconomic factors considered. Of the macroeconomic factors, factors related to uncertainties appear to be more important in accounting for the higher-order factors, especially financial uncertainty.

Table 8: **Exposure to Macroeconomic Factors**

This table reports the results of regressing each selected higher-order factor on a list of common macroeconomic factors. Significance of the coefficients is based on Newey-West corrected standard errors. *** denotes significance at 1%. ** denotes significance at 5%. * denotes significance at 10%. To each factor, the table associate a reference, a category, and the number of high-order factors with a significant exposure at the 10% confidence level.

			(1)	(2)	(3)	(4)	(5)	(6)	(7)	
Factor	Reference	Category	SMB2	SMB2*Mom	Mom2*RMW	Mkt-RF2	Mkt-RF2*RMW	Mkt-RF*SMB	HML2*Mkt-RF	# sig
Common idio-vol shock	Herskovic et al. (2016)	Volatility	0.31	0.77***	0.31	0.44	-0.16	0.2	0.42**	2
Consum growth	Chapman (1997)	Consum	1.06	0.25	-0.2	-0.25	0.05	-0.04	0.01	0
Consum growth squared	Chapman (1997)	Consum	3.56*	-3.08	1.46	-1.44	-1.46***	-2.99**	1.04	3
Intermediary capital return factor	He et al. (2017)	Inter	1.15	-5.29**	1.99	2.02*	0.27	1.26	1.92**	3
Leverage factor	Adrian et al. (2014)	Inter	-0.03	-0.23	0.11	0.73**	-0.11	0.3	0.23	1
Intermediary capital risk factor	He et al. (2017)	Inter	-0.75	0.56	1.97	0.09	0.01	-0.32	-1.50**	1
Equity market return	Fama and French (1993)	Fin Markets	-0.35	4.39	-5.34	-0.77	-1.33**	0.08	1.78***	2
Liquidty	Pástor and Stambaugh (2003)	Liquidity	-0.02	-0.37	0.12	-0.43	0.04	-0.78*	0.09	1
Polynomial of labor income growth	Dittmar (2002)	Macro	-69.94	353.57	-167.24	145.62	-88.83	-73.58	-121	0
Macro uncertainty	Bali et al. (2017)	Macro	-0.92*	0.02	0.55	-0.33	-0.29*	-0.25	-0.24	2
Financial uncertainty	Bali et al. (2017)	Macro	1.36**	0.38	-0.92	0.78***	0.19**	0.30**	0.13	4
Real uncertainty	Bali et al. (2017)	Macro	0.23	-1.69	0.59	1.64***	0.19	1	0.07	1
Labor income growth	Campbell (1996)	Macro	144.2	-703.78	329.56	-290.04	177.76	146.78	244.27	0
Change long-term govt yield	Sweeney and Warga (1986)	Macro	-5.21	-24.69**	13.85	9.34*	-3.21	10.21**	4.18	3
News implied volatility	Manela and Moreira (2017)	Macro	-0.02**	-0.01	0.01	0.01**	0	0	-0.00**	3
Investor sentiment (orth macro)	Baker and Wurgler (2006)	Sentiment	0.02	-0.03	0	0.12***	0.02*	0.03	-0.03	2
Investor sentiment (orth macro)	Huang et al. (2015)	Sentiment	-0.01	0.05	-0.05	-0.09***	0	-0.05**	-0.01	2
		Adj. R^2	0.21	0.09	0.01	0.42	0.21	0.27	0.46	

6 Additional Results

In this section, we provide additional results for our empirical analyses, and we discuss extensions of our baseline setup.

6.1 Simulation Result

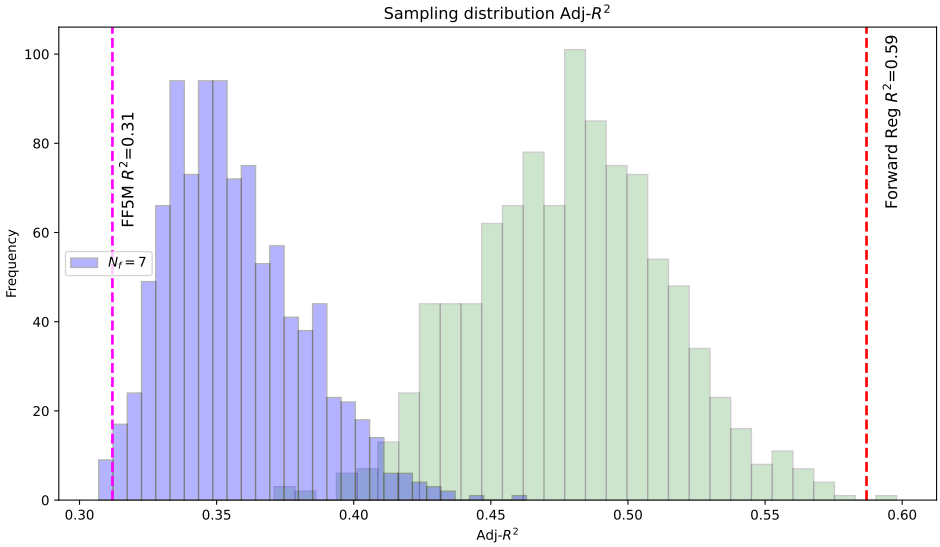
The selected higher-order model incorporates 7 higher-order factors beyond the traditional FF5M model. Our results indicate that the model performance significantly improves with the 7 higher-order factors where the R-squared of the second-pass regression increases from 31.2% to 58.7%. However, as noted by Lewellen et al. (2010), spurious factors can artificially inflate model fit in the second-pass Fama-MacBeth regression, particularly when the test assets exhibit a strong factor structure. Lewellen et al. (2010) illustrate this concern using the Fama-French 25 Size-BM portfolios which have a particularly strong factor structure. While our analysis mitigates this issue by employing a broad set of 484 test assets and demonstrating the superior out-of-sample performance, we further conduct simulation exercises to confirm that our results are not driven by spurious factors.

Therefore, we estimate the sampling distribution of cross-sectional R-squareds obtained using randomly generated factors, following the approach advocated by Lewellen et al. (2010). Specifically, in each simulation, we generate a set of 57 random factors, matching the number of candidate higher-order factors. These random factors are assumed to be $i.i.d. \sim N(0, \sigma^2)$, and we set σ^2 equal to the sample variance of the market factor (Mkt-RF). We then apply a forward regression procedure, sequentially selecting factors that maximize the adjusted cross-sectional R-squared when added to the FF5M factors. We repeat the procedure 1,000 times by generating 1,000 independent sets of factors.

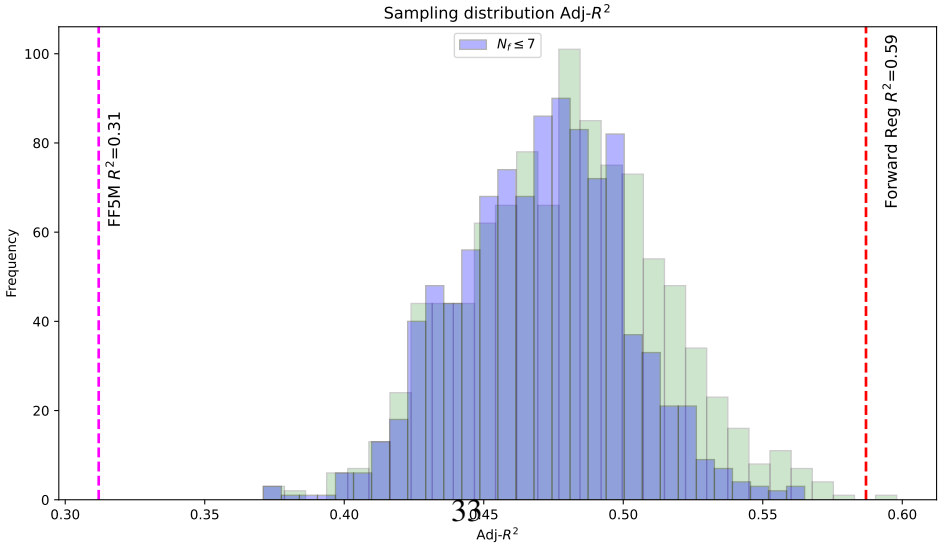
Figure 4: **Sampling Distribution of R-squareds**

This figure plots the sampling distribution (green histogram) of adjusted cross-sectional R-squareds obtained applying the forward selection FMB method to the FF5M factors and randomly generated higher-order factors. First, we randomly generate 57 factors $i.i.d. \sim N(0, \sigma(\text{Mkt-RF}))$. Next, at each step of the forward selection FMB method, we select one factor to maximizes the adjusted cross-sectional R-squared until $\Delta\left(\text{adj } R^2\right) \leq 1\%$. The procedure is repeated 1,000 times to obtain the sampling distribution of R-squareds. In subplot (a), we include the sampling distribution obtained in the case in which the number of additionally selected factors is constained to be smaller or equal to 7 (blue histogram). In subplot (b), we include the sampling distribution obtained in the case in which we augment the FF5M model with 7 randomly generated factors, repeating the procedure 1,000 times. The pink-dashed line corresponds to the adjusted cross-sectional R-squared of the FF5M model. The red-dashed line corresponds to the adjusted cross-sectional R-squared of the FF5M model with higher-order factors selected by the forward regression method (see Table 2).

(a)



(b)



The green histogram in figure 4 summarizes the sampling distribution of the resulting R-squareds from the 1,000 simulations. In the simulation exercise, we obtain a cross-sectional R-squared larger than 0.59 only in 0.1% of the simulations, meaning that the adjusted R-squared from the higher-order model is significant at the 0.1% level. This simulation exercise highlights the fact that it is extremely unlikely that our results are obtained by chance.

Note that the method can select a number of factors different than 7, which is the number of higher-order factors presented in Table 2. In fact, in the simulation exercise we conduct, the maximum number of random factors selected is 12. In Panel A, we compare the sampling distribution of R-squareds with that obtained with the additional constrained that the number of selected factors cannot be larger than 7 (blue histogram). In this case, we never observe R-squareds larger than 0.59 (maximum adjusted cross-sectional R-squared of 0.565).

Finally, in Panel B, we further compare the sampling distribution of R-squareds with that obtained by directly adding 7 randomly generated factors, $i.i.d. \sim N(0, \sigma^2)$, to the FF5M factors, repeating the procedure 1,000 times (blue histogram). In this case, the maximum adjusted cross-sectional R-squared is only 0.463.

6.2 An alternative factor zoo

In order to investigate the robustness of our results in terms of subsuming the factor zoo, we replicate the analysis based on the two-pass Fama-MacBeth procedure using an alternative set of zoo factors, based on the factors constructed by Chen and Zimmermann (2022)**.** Specifically, for each factor in Chen and Zimmermann (2022), we estimate its SDF loading while controlling for our higher-order factor model, which includes the FF5M and the selected higher-order factors. Figure 5 summarizes the results. Panel A shows that the intercept estimates tend to be highly statistically significant, with a median value of 3.97, when the factor model includes only the factor in the factor zoo. When the factor model includes both FF5M and the factor in the factor zoo, the median t-statistic decreases substantially but remains significant at the 5% level at 2.23. When the factor model further includes the higher-order factors, the median t-statistic is no longer statistically

significant at the 5% level at 1.87. Panel B reveals that, when the factor model includes only the factor in the factor zoo, the median SDF loading estimate for the factor zoo is not statistically different from zero, with a t-statistic of 0.80. Note that these estimates are likely biased because of omitted factors (see, e.g., Feng et al. (2020)). Controlling for the FF5M, the median t-statistic increases substantially to 1.91. 46% of zoo factors become statistically significant at the 5% level with a t-statistic greater than 2 suggesting that FF5M helps to remove some bias in the factor estimates. Importantly, when we additionally control for the high-order factors, we find that for 93% of the zoo factors the risk price estimate is not statistically different from zero, with the median t-statistic dropping from 1.79 to only 0.72. Specifically, we find that only 12 out 159 zoo factors remain significant at the 5% level. These results confirm our findings based on the factors in Jensen et al. (2023).[3]

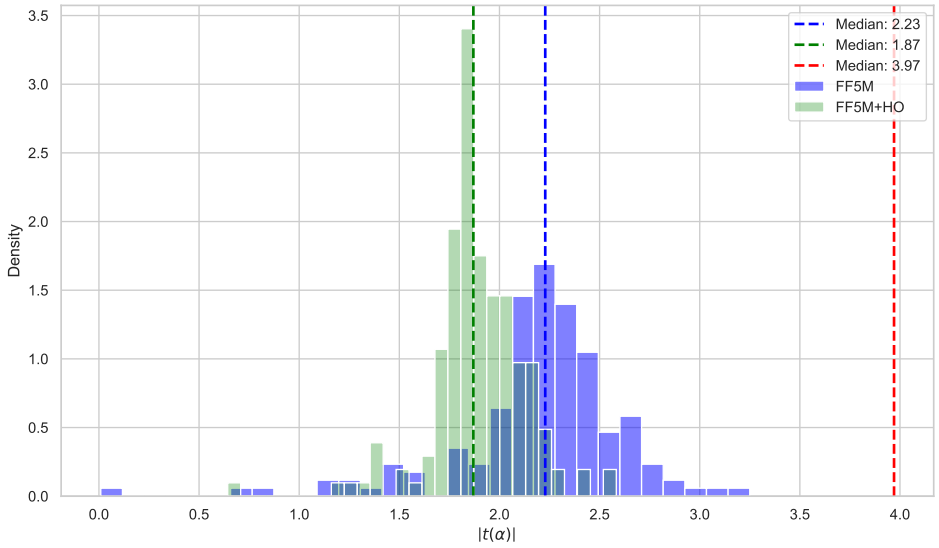
Table A.3 in the Online Appendix presents the fraction of zoo factors in Chen and Zimmermann (2022) with loading significantly different from zero at the 5% confidence level on each of the FF5M factors and the selected higher-order factors selected from the forward selection FM regression. Column (1) shows the fraction of significance for the factors in FF5M model only and Column (2) reports the fraction of significance for the factors in the full higher-order factor model. Column (1) shows that the fractions of significance range from 42.8% for CMA to 57.2% for SMB and RMW for the 6 factors in FF6M. Column (2) shows that the fractions of significance remain stable for the 6 factors in FF6M, when we further include the selected higher-order factors. The fractions of significance for the higher-order factors range from 20.8% for Mkt-RF2*RMW to 47.8% for SMB2*Mom. These results are consistent with the findings we obtain using the zoo factors in Jensen et al. (2023), further highlighting that it is not the case a specific higher-order factor fully subsumes the zoo factors.

[3]The factors in Chen and Zimmermann (2022) that remain significant at the 5% confidence level are: MomSeason06YrPlus, CBOperProf, ChInv, STreversal, DelFINL, DivYieldST, InvGrowth, GrSaleToGrInv, MomOffSeason16YrPlus, DebtIssuance, CompositeDebtIssuance, OrderBacklog.

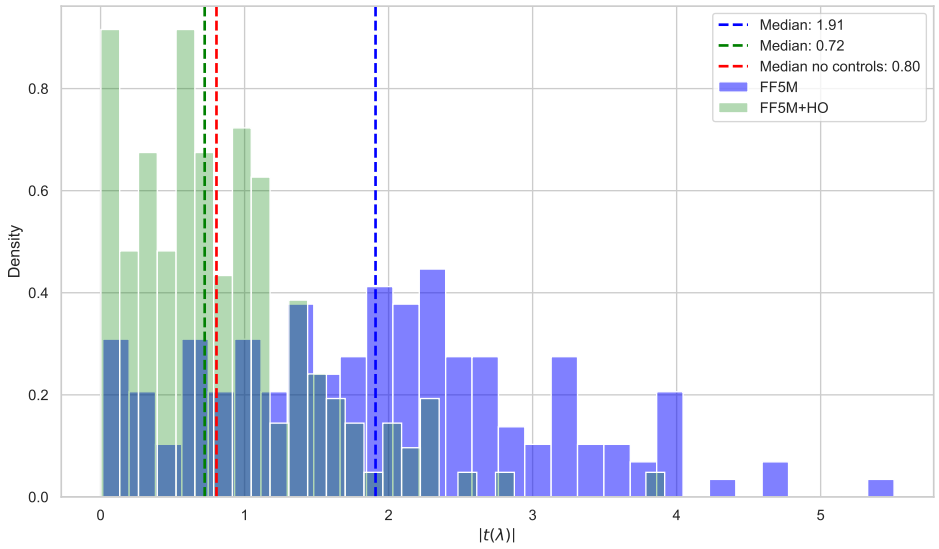
Figure 5: **Reducing Alternative Zoo Factors in the Cross-Section**

This figure plots the distribution of absolute t-statistics associated with the risk prices ($t(\lambda)$) and the intercepts ($t(\alpha)$) of the alternative zoo factors from Chen and Zimmermann (2022) based on US equity data. For each zoo factor, we use the FMB method to estimate its risk price controlling for the FF5M (blue histogram) and the FF5M with the high-order factors selected by the forward selection FMB method (green histogram). The dashed-lines in blue and gree correspond to the median absolute t-statistic for the FF5M and for the FF5M with high-order factors models, respectively. The dashed-red-line corresponds to the median absolute tstatistic with no controls. The t-statistics are based on Newey-West corrected standard errors.

(a) **Distribution of** $|t(\alpha)|$



(b) **Distribution of** $|t(\lambda)|$



6.3 Alternative sets of candidate factors

We start considering alternative sets of higher-order factors formed using the FF5M factors. Table A.4 presents the estimation results of the forward regression method with four alternative sets of higher-order factors. Specifically, Panel A presents estimation results using a set of high-order factors of degree 2. The latter is a lower degree than the one from the baseline analysis (i.e., degree 3). These high-order factors include all pairwise interactions of the FF5M, and the FF5M factors squared. In this case, the forward regression method selects four high-order factors, reaching an adjusted cross-sectional R-squared of 0.517 in the final step, approximately 7 pp smaller than in the baseline setup (see Table 2). Panel B presents estimation results using a set of high-order factors of higher degree, and specifically of degree up to 4. These factors, augment the high-order factors of degree 3 of the baseline setup with the FF5M to the fourth power (i.e., f_i^4), and pairwise interactions of the FF5M factors of degree 4, i.e., $f_i^2 \times f_j^2$, $f_i^3 \times f_j$, $f_i \times f_j^3$. In this case, the forward regression method selects 9 high-order factors, reaching an adjusted cross-sectional R-squared of 0.6 in the final step. The latter is a R-squared value very similar to that obtained in the baseline setup, with high-order factors up to degree 3. These results suggest that higher-order factors up to degree 3 are most relevant in explaining cross-section asset returns, corroborating our choice of using high-order factors of degree 3 in the baseline analysis.

Furthermore, Panel C and D report the results when augmenting the FF5M with only pairwise interactions of degree 2 and 3 (Panel C) and only the powers of degree 2 and 3 (Panel D). The results from these two panels reveal that incorporating pairwise interactions between FF5M is crucial. In fact, the model with interactions of degree 2 and 3 reaches an adjusted cross-sectional R-squared of 0.542, selecting 7 high-order factors. In contrast, the model with powers of degree 2 and 3 selects 3 high-order factors, for an adjusted cross-sectional R-squared of 0.467 in the final step. These results highlight that both higher-orders of FF5M and higher-order interactions of FF5M are important components of the SDF, especially higher-order interactions.

7 Conclusions

In this paper, we study the importance of higher-order factors in accounting for the stochastic discount factor and the pricing of the cross-section of asset returns. We first propose a forward selection Fama-MacBeth regression method for selecting and estimating SDF models that may be potentially high-dimensional.

The forward selection FMB regression method selects 7 higher-order factors beyond the traditional FF5M model. These higher-order factors include both higher-orders of the factors in FF5M and interactions of these factors. We show that the selected higher-order factor model significantly improves the pricing power of the benchmark factor models. Specifically, by applying this selected higher-order factor model to 484 characteristic-managed portfolios, we find that the model enhance the cross-sectional explanatory power from 31.2% based on the FF5M model to 58.7%. Furthermore, we show that the estimates of the SDF loadings of the higher-order factors are significant. These selected higher-order factors robustly price assets both in-sample and out-of-sample, substantially outperforming benchmark models.

Moreover, our findings indicate that these higher-order factors effectively subsume many alternative factors proposed in the literature, subsuming the majority of the factors typically included in the factor zoo. Overall, our paper highlights the importance of higher-orders in capturing critical dimensions of the cross-section of asset returns and provides a robust, econometrically justified framework for factor selection and model estimation in high-dimensional contexts.

Adrian, Tobias, Erkko Etula, and Tyler Muir, "Financial intermediaries and the cross-section of asset returns," *Journal of Finance*, 2014, *69* (6), 2557–2596.

Baker, Malcolm and Jeffrey Wurgler, "Investor sentiment and the cross-section of stock returns," *The journal of Finance*, 2006, *61* (4), 1645–1680.

Bali, Turan G, Stephen J Brown, and Yi Tang, "Is economic uncertainty priced in the cross-section of stock returns?," *Journal of Financial Economics*, 2017, *126* (3), 471–489.

Bansal, Ravi and Salim Viswanathan, "No arbitrage and arbitrage pricing: A new approach," *Journal of Finance*, 1993, *48* (4), 1231–1262.

Borri, Nicola, Denis Chetverikov, Yukun Liu, and Aleh Tsyvinski, "One Factor to Bind the Cross-Sectional of Returns," *Working Paper*, 2024.

Campbell, John Y, "Understanding risk and return," *Journal of Political Economy*, 1996, *104* (2), 298–345.

Carhart, Mark M, "On persistence in mutual fund performance," *Journal of Finance*, 1997, *52* (1), 57–82.

Chapman, David A, "The cyclical properties of consumption growth and the real term structure," *Journal of Monetary Economics*, 1997, *39* (2), 145–172.

Chen, Andrew Y. and Tom Zimmermann, "Open Source Cross-Sectional Asset Pricing," *Critical Finance Review*, 2022, *27* (2), 207–264.

Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins, "Double/debiased machine learning for treatment and structural parameters," 2018.

Cochrane, John, *Asset Pricing*, Princeton University Press, 2005.

Cochrane, John H, "Presidential address: Discount rates," *The Journal of Finance*, 2011, *66* (4), 1047–1108.

Connor, Gregory and Robert A Korajczyk, "Performance measurement with the arbitrage pricing theory: A new framework for analysis," *Journal of Financial Economics*, 1986, *15* (3), 373–394.

Dittmar, Robert F, "Nonlinear pricing kernels, kurtosis preference, and evidence from the cross section of equity returns," *The Journal of Finance*, 2002, *57* (1), 369–403.

Fama, Eugene F. and Kenneth R. French, "Common risk factors in the returns on stocks and bonds," *Journal of Financial Economics*, 1993, *33* (1), 3–56.

Fama, Eugene F and Kenneth R French, "Multifactor explanations of asset pricing anomalies," *Journal of Finance*, 1996, *51* (1), 55–84.

_ **and** _ , "A five-factor asset pricing model," *Journal of Financial Economics*, 2015, *116* (1), 1–22.

Feng, Guanhao, Stefano Giglio, and Dacheng Xiu, "Taming the factor zoo: A test of new factors," *Journal of Finance*, 2020, *75* (3), 1327–1370.

Freyberger, Joachim, Andreas Neuhierl, and Michael Weber, "Dissecting characteristics nonparametrically," *The Review of Financial Studies*, 2020, *33* (5), 2326–2377.

Giglio, Stefano and Dacheng Xiu, "Asset pricing with omitted factors," *Journal of Political Economy*, 2021, *129* (7), 1947–1990.

Green, Jeremiah, John RM Hand, and X Frank Zhang, "The characteristics that provide independent information about average US monthly stock returns," *The Review of Financial Studies*, 2017, *30* (12), 4389–4436.

Gu, Shihao, Bryan Kelly, and Dacheng Xiu, "Empirical asset pricing via machine learning," *Review of Financial Studies*, 2020, *33* (5), 2223–2273.

Harvey, Campbell R and Akhtar Siddique, "Conditional skewness in asset pricing tests," *Journal of Finance*, 2000, *55* (3), 1263–1295.

_ **and Yan Liu**, "Lucky factors," *Journal of Financial Economics*, 2021, *141* (2), 413–435.

_ **,** _ **, and Heqing Zhu**, "... and the cross-section of expected returns," *The Review of Financial Studies*, 2016, *29* (1), 5–68.

He, Zhiguo, Bryan Kelly, and Asaf Manela, "Intermediary asset pricing: New evidence from many asset classes," *Journal of Financial Economics*, 2017, *126* (1), 1–35.

Herskovic, Bernard, Bryan Kelly, Hanno Lustig, and Stijn Van Nieuwerburgh, "The common factor in idiosyncratic volatility: Quantitative asset pricing implications," *Journal of Financial Economics*, 2016, *119* (2), 249–283.

Huang, Dashan, Fuwei Jiang, Jun Tu, and Guofu Zhou, "Investor sentiment aligned: A powerful predictor of stock returns," *The Review of Financial Studies*, 2015, *28* (3), 791–837.

Jegadeesh, Narasimhan and Sheridan Titman, "Returns to buying winners and selling losers: Implications for stock market efficiency," *Journal of Finance*, 1993, *48* (1), 65–91.

Jensen, Theis Ingerslev, Bryan Kelly, and Lasse Heje Pedersen, "Is there a replication crisis in finance?," *Journal of Finance*, 2023, *78* (5), 2465–2518.

_ **, Hanno Lustig, and Stijn Van Nieuwerburgh**, "The cross-section and time series of stock and bond returns," *Journal of Monetary Economics*, 2017, *88*, 50–69.

Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh, "Shrinking the cross-section," *Journal of Financial Economics*, 2020, *135* (2), 271–292.

Lettau, Martin and Markus Pelger, "Factors that fit the time series and cross-section of stock returns," *Review of Financial Studies*, 2020, *33* (5), 2274–2325.

Lewellen, Jonathan, Stefan Nagel, and Jay Shanken, "A skeptical appraisal of asset pricing tests," *Journal of Financial Economics*, 2010, *96* (2), 175–194.

Manela, Asaf and Alan Moreira, "News implied volatility and disaster concerns," *Journal of Financial Economics*, 2017, *123* (1), 137–162.

McLean, R David and Jeffrey Pontiff, "Does academic research destroy stock return predictability?," *The Journal of Finance*, 2016, *71* (1), 5–32.

Nagel, Stefan, *Machine learning in asset pricing*, Princeton University Press, 2021.

Newey, Whitney K and Kenneth D West, "A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix," *Econometrica*, 1987, *55* (3), 703–708.

Pástor, L'uboš and Robert F Stambaugh, "Liquidity risk and expected stock returns," *Journal of Political Economy*, 2003, *111* (3), 642–685.

Sweeney, Richard J and Arthur D Warga, "The pricing of interest-rate risk: evidence from the stock market," *The Journal of Finance*, 1986, *41* (2), 393–410.

Online Appendix

A Proof of Theorem

B Technical Lemmas

Lemma B.1. *Let $A = (A_{i,j})_{i,j\in[n]} \in \mathbb{R}^{n\times n}$ be a symmetric matrix such that $\|A\|_2 \le C_1$, $\|A^{-1}\|_2 \le C_1$, and $\|A^{-1}\|_{\infty,1} \le C_1$ for some constant $C_1 > 0$. For some $k \in [n]$, let $B = (A_{i,j})_{i,j\in[n]\setminus k}$ be sub-matrix of A. Then B is invertible and satisfies $\|B^{-1}\|_{\infty,1} \le C$, where C is constant depending only on C_1.*

Proof. Without loss of generality, we will assume that $k = n$, so that

$$A = \begin{pmatrix} B & b \\ b^T & a \end{pmatrix},$$

where $B = (A_{i,j})_{i,j\in[n-1]} \in \mathbb{R}^{n\times n}$, the vector $b \in \mathbb{R}^{n-1}$ is the last column of A (excluding the bottom-right element), and $a = A_{n,n}$.

Because A is symmetric and invertible (and also positive definite from the smallest-eigenvalue condition), B is also positive definite and therefore invertible.

Step 2. Relating B^{-1} to A^{-1}.

We want to show that $\|B^{-1}\|_\infty$ is uniformly bounded in terms of C_1 and C_2. Consider any vector $y \in \mathbb{R}^{N-1}$. We have

$$x = B^{-1}y \quad \Longleftrightarrow \quad Bx = y.$$

We can embed this system into an N-dimensional problem:

$$A\begin{pmatrix} x \\ \alpha \end{pmatrix} = \begin{pmatrix} y \\ 0 \end{pmatrix}.$$

In block form, this means

$$\begin{pmatrix} B & b \\ b^T & a_{NN} \end{pmatrix} \begin{pmatrix} x \\ \alpha \end{pmatrix} = \begin{pmatrix} y \\ 0 \end{pmatrix}.$$

Since A is invertible, there is a unique solution

$$\begin{pmatrix} x \\ \alpha \end{pmatrix} = A^{-1} \begin{pmatrix} y \\ 0 \end{pmatrix}.$$

By the hypothesis $\|A^{-1}\|_\infty \le C_2$, we have

$$\left\| \begin{pmatrix} x \\ \alpha \end{pmatrix} \right\|_\infty \le \|A^{-1}\|_\infty \left\| \begin{pmatrix} y \\ 0 \end{pmatrix} \right\|_\infty \le C_2 \|y\|_\infty.$$

Hence $\|x\|_\infty \le C_2 \|y\|_\infty$ and $|\alpha| \le C_2 \|y\|_\infty$.

Step 3. Estimate for $\|B^{-1}y\|_\infty$.

From the block system,

$$Bx + b\,\alpha = y \quad \Longrightarrow \quad x = B^{-1}\big[y - b\,\alpha\big] = B^{-1}y - B^{-1}b\,\alpha.$$

Thus

$$B^{-1}y = x + B^{-1}b\,\alpha.$$

Taking $\|\cdot\|_\infty$ on both sides:

$$\|B^{-1}y\|_\infty \le \|x\|_\infty + \|B^{-1}b\|_\infty |\alpha| \le C_2 \|y\|_\infty + \|B^{-1}b\|_\infty \big(C_2 \|y\|_\infty\big).$$

Therefore

$$\|B^{-1}y\|_\infty \le \Big[C_2 + C_2 \|B^{-1}b\|_\infty\Big] \|y\|_\infty.$$

Since $y \in \mathbb{R}^{N-1}$ was arbitrary, we get

$$\|B^{-1}\|_\infty \;\le\; C_2 + C_2\,\|B^{-1}b\|_\infty.$$

Step 4. Bounding $\|B^{-1}b\|_\infty$.

We now estimate $\|B^{-1}b\|_\infty$ similarly. Solve

$$B\,w = b.$$

Embed into the full system

$$A\begin{pmatrix} w \\ \beta \end{pmatrix} \;=\; \begin{pmatrix} b \\ 0 \end{pmatrix}.$$

Again, the solution is

$$\begin{pmatrix} w \\ \beta \end{pmatrix} = A^{-1}\begin{pmatrix} b \\ 0 \end{pmatrix}.$$

Hence

$$\|w\|_\infty \;\le\; \|A^{-1}\|_\infty\,\|b\|_\infty \;\le\; C_2\,\|b\|_\infty.$$

But b is the last column of A (except the corner entry), so

$$\|b\|_\infty \;=\; \max_{1\le i\le N-1} |A_{i,N}| \;\le\; \max_{1\le i\le N-1} \sum_{j=1}^{N} |A_{i,j}| \;\le\; \|A\|_\infty \;\le\; C_1.$$

Thus $\|w\|_\infty \le C_2\,C_1$, and $w = B^{-1}b$, so

$$\|B^{-1}b\|_\infty \;\le\; C_2\,C_1.$$

Step 5. Final conclusion.

Putting this into our earlier bound,

$$\|B^{-1}\|_\infty \;\le\; C_2 \;+\; C_2\,\|B^{-1}b\|_\infty \;\le\; C_2 \;+\; C_2\,(C_2\,C_1) \;=\; C_2 \;+\; C_2^2\,C_1.$$

Hence $\|B^{-1}\|_\infty$ is bounded by a constant that depends only on C_1 and C_2, exactly as required. □

C Additional Figures & Tables

Figure A.1: **Distribution of Higher-Order Factors**

This figure plots the distribution of the high-order factors grouped in pairwise interactions of degree 2 (Panel A), pairwise interactions of degree 2 (Panel B), powers of degree 2 (Panel C) and powers of degree 2 (Panel D). The distribution is clipped at 4 standard deviations from the mean, so that any value below the lower bound is replaced by the lower bound and any value above the upper bound is replaced by the upper bound.

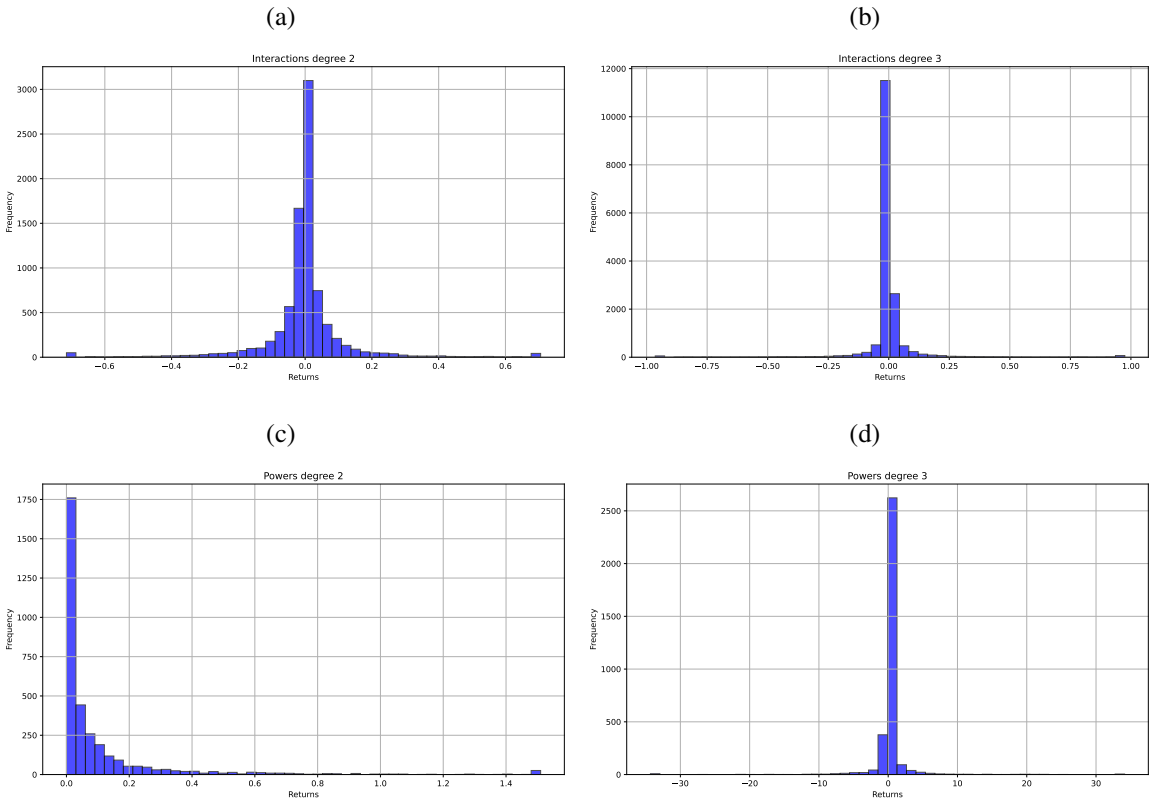


Figure A.2: R^2 of Factor Mimicking Regressions

This table reports, for each of the high-order selected factors with the forward selection FMB regression method, the adjusted R-squared of factor mimicking regressions. The factor mimicking regressions are time-series OLS regressions of high-order factor returns on the zoo factor from Jensen et al. (2023).

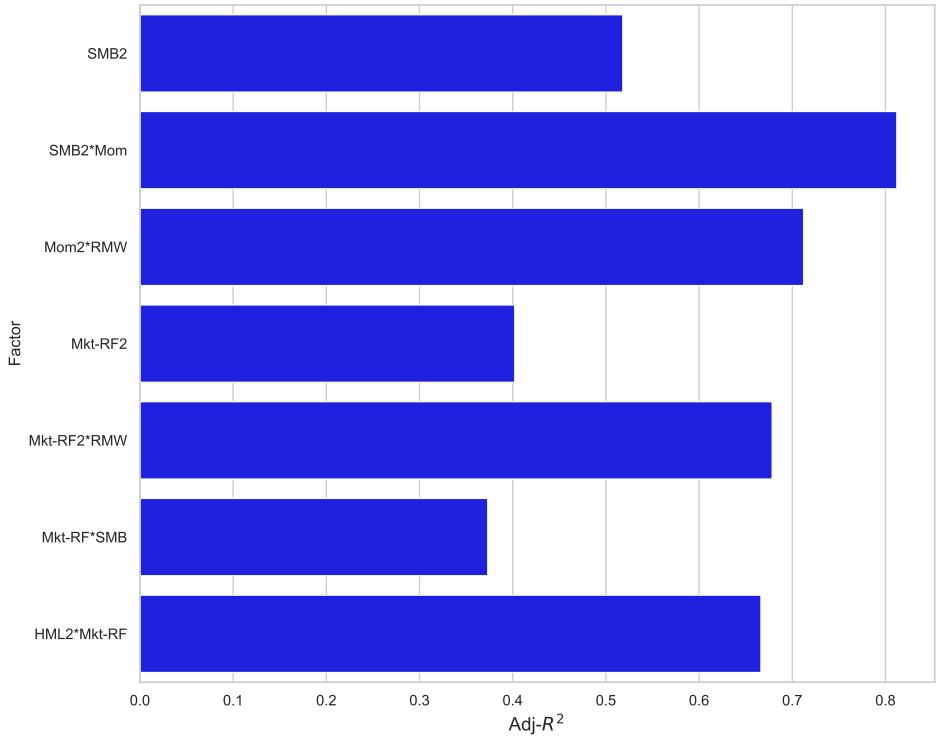


Table A.1: **Cross-Sectional Performance Restricted Models**

This table reports the results from the second step of the FMB method imposing the restriction that the risk prices of the tradable factors must equal to their sample means. The tradable factors are the FF5M factors. The high-order factors are those selected by the forward selection Fama-MacBeth regressions method (see Table 2). The table reports the adjusted cross-sectional R-squared, the estimate for the intercept (α) and associated t-statistic. The t-statistics is based on Newey-West corrected standard errors.

#	Model	Adj. R-squared	α	t-stat (α)
1	CAPM	-0.439	0.001	0.28
2	FF3	-0.371	0.001	0.334
3	FF5	-0.111	0.004	2.263
4	FF5M	0.022	0.006	2.839
5	Higher-Order	0.405	0.004	2.663

Table A.2: **Comparison of R-squareds with Random Train and Test Sample**

This table reports the in-sample and out-of-sample R-squareds in the cross-sectional regression step of the Fama-MacBeth method for different models. We randomly generate training and test samples of equal length by randomly drawing monthly returns without replacement, and estimate for each asset the covariance with respect to the factors and the SDF loadings on the training sample, and compare actual mean returns with model predictions in the test sample. For the test sample, we report R-squareds aftering re-centering the predicted values from the model. For the FF5M + Hig-Order model, we consider the high-order factors selected by the forward selection Fama-MacBeth regression method (see Table 2).

	(1) CAPM	(2) FF3	(3) FF5	(4) FF5M	(5) Higher-Order
R^2_{train}	0.06	0.167	0.274	0.316	0.498
R^2_{test}	0.063	0.04	0.076	0.085	0.14

Table A.3: **Loadings of Alternative Zoo Factors**

This table reports the fraction of zoo factors with loading significantly different from zero at the 5% confidence level on each of the FF5M factors and high-order factors selected by the forward selection Fama-MacBeth method. Significance is evaluated by t-statistics based on Newey-West corrected standard errors. The zoo factors are 159 factors constructed in Chen and Zimmermann (2022).

Factor	Frac Sig 5%	Frac Sig 5%
Mkt-RF	0.478	0.509
SMB	0.572	0.579
HML	0.459	0.421
RMW	0.572	0.553
CMA	0.428	0.491
Mom	0.535	0.421
SMB2		0.365
SMB2*Mom		0.478
Mom2*RMW		0.358
Mkt-RF2		0.302
Mkt-RF2*RMW		0.208
Mkt-RF*SMB		0.302
HML2*Mkt-RF		0.302
# zoo factors	159	159

Table A.4: **Forward Selection Fama-MacBeth Regression Method with Alternative Higher-Order Factors**

This table reports the results of the estimation of the forward selection FMB regression method for alternative sets of higher-order factors. Panel A considers the FF5M factors and all the high-order factors of degree 2 (i.e., $f_i \times f_j$ and f_j^2). Panel B considers the FF5M factors and all the high-order factors up to degree 4 (i.e., $f_i \times f_j$, $f_i^2 \times f_j$, $f_i^2 \times f_j^2$, $f_i^3 \times f_j$, f_i^2, f_i^3, and f_i^4). Panel C considers the FF5M factors and all the pairwise interactions of degree 2 and 3 (i.e., $f_i \times f_j$ and $f_i^2 \times f_j$). Panel D considers the FF5M factors all the powers of degree 2 and 3 (i.e., f_i^2 and f_i^3). For each step, we report the selected factor (h_j) adjusted cross-sectional R-squared, the estimate for the intercept (α) and associated t-statistic. The t-statistic is based on Newey-West corrected standard errors.

Step	h_j	Adj. R-squared	α	t-stat (α)
Panel A: high-order degree 2				
1	SMB2	0.41	0.003	1.821
2	HML*Mom	0.464	0.004	2.033
3	Mkt-RF2	0.503	0.005	2.738
4	SMB*Mom	0.517	0.005	2.692
Panel B: high-order degree up to 4				
1	SMB2*CMA2	0.423	0.005	2.581
2	Mkt-RF2*SMB2	0.464	0.004	2.329
3	Mom3*CMA	0.485	0.004	2.062
4	Mkt-RF2*RMW	0.498	0.004	1.931
5	HML2	0.528	0.003	1.589
6	Mkt-RF2	0.547	0.004	2.271
7	CMA2*Mom	0.57	0.004	2.13
8	CMA2*SMB	0.587	0.004	2.074
9	Mkt-RF2*Mom	0.6	0.004	2.132
Panel C: interactions degree 2 and 3				
1	SMB2*Mkt-RF	0.394	0.003	1.456
2	CMA2*SMB	0.437	0.003	1.429
3	Mkt-RF2*SMB	0.466	0.004	2.023
4	CMA2*RMW	0.487	0.004	2.084
5	HML*CMA	0.509	0.004	2.321
6	CMA2*Mkt-RF	0.527	0.004	2.248
7	CMA2*Mom	0.542	0.004	2.044
Panel D: powers degree 2 and 3				
1	SMB2	0.41	0.003	1.821
2	Mkt-RF2	0.444	0.005	2.478
3	RMW3	0.467	0.005	2.807

Table A.5: **Forward Selection Fama-MacBeth Regression Method Starting from the CAPM**

This table reports the results from the second step of the FMB method with selection by the forward selection FMB regression method starting from the CAPM. At each step, we augment the CAPM by one additional factor selected from the remaining FF5M factors and the higher-order factor (h_j) which maximizes the adjusted cross-sectional R-squared. The method stops when the improvement in the R-squared in a given step is smaller than 1 pp. The table reports the adjusted cross-sectional R-squared, the selected factor and its associated t-statistic, the estimate for the intercept (α) and associated t-statistic. The t-statistics is based on Newey-West corrected standard errors.

Step	h_j	Adj. R-squared	α	t-stat (α)
1	HML2	0.285	0.007	2.916
2	Mom2*Mkt-RF	0.323	0.005	2.153
3	SMB2*Mkt-RF	0.377	0.006	2.48
4	Mkt-RF2	0.406	0.005	2.405
5	HML*Mom	0.429	0.004	2.337
6	Mom	0.485	0.005	2.682
7	CMA2*Mom	0.499	0.006	3.087
8	SMB2*HML	0.509	0.005	2.777
9	Mom2*SMB	0.537	0.005	2.976
10	Mkt-RF3	0.552	0.005	2.722