High-Dimensional Factor Models with Cross-Sectional and Time-Series Predictors

Mingjun Sun*

November 12, 2026

Abstract

I study the informational content of cross-sectional and time-series predictors in a high-dimensional setting. I use instrumented principal components analysis (IPCA) to construct SDFs from firm characteristics, their interactions, random Fourier features, and factor-timing instruments. Nonlinear transformations of characteristics improve out-of-sample Sharpe ratios and reduce pricing errors, whereas the gains from timing information are selective. Furthermore, I find that cross-sectional and time-series information can be effectively summarized by an SDF comprising a small number of IPCA factors. I provide an economic rationale for this finding: when common factors capture nondiversifiable risks, a restriction on near-arbitrage opportunities limits the expected returns of portfolios whose remaining risks can be diversified away.

Keywords: Asset pricing, factor pricing models, cross section of returns, stochastic discount factor

*Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA 15213, USA. Email: mingjuns@andrew.cmu.edu. I am indebted to my advisors, Burton Hollifield, Bryan Routledge, Patrick Blonien and Miguel Oliveira, for their invaluable guidance and support. For helpful comments, suggestions, and discussions, I thank Chester Spatt and seminar participants at Carnegie Mellon University. Presentations: Future Alpha 2026, CMU-Pitt-Penn State Finance Conference 2024, Carnegie Mellon University Finance Brownbag.

1 Introduction

A stochastic discount factor (SDF) summarizes how investors value risky payoffs. A central question in empirical asset pricing is how much information is needed to estimate it. Models based on a few characteristic portfolios provide an influential starting point (Fama and French, 1992, 2015). A small number of priced risks, however, need not imply that only a few observable characteristics are useful. Many characteristics may help measure exposure to the same risks, and their implications for expected returns may change with economic conditions.

Machine learning methods allow SDF estimation to use many characteristics and their nonlinear transformations.[1] Time-series predictors, such as valuation ratios and interest rate spreads, provide another source of information. These variables are often studied as predictors of the equity premium (Welch and Goyal, 2008; Goyal et al., 2024). They can also inform allocations across equity factors. Multiplying a factor return by a lagged predictor gives the payoff to a strategy whose factor exposure varies with the predictor (Cochrane, 2009; Haddad et al., 2020; Lehnherr et al., 2025).

I examine how nonlinear characteristic information and time-series predictors improve SDF estimation and whether their pricing benefits can be captured by a relatively small number of factors. This distinguishes the information supplied to the model from the dimension of the SDF used to exploit it.

I represent the SDF as a linear combination of portfolio returns,

$$M_{t+1} = 1 - \sum_{p=1}^{P} \lambda_p F_{p,t+1}.$$

The returns $F_{p,t+1}$ include portfolios whose weights depend on firm characteristics and portfolios that scale these positions by lagged time-series predictors. Nonlinear transformations allow the effects of characteristics to depend on one another, while timing portfolios allow exposures to vary with economic conditions. Both constructions can generate many portfolios from a modest number of underlying predictors. I compare an SDF estimated directly from all these portfolios with one based on IPCA factors extracted from stock returns using the corresponding instruments.

The main specifications use 153 characteristics, 11,934 instruments formed from those characteristics and their interactions, 1,000 random Fourier features per draw, and interac-

[1] See Kelly and Xiu (2023) and Nagel (2021).

tions of the characteristics with fourteen shared time-series predictors. Section 2.2 describes the estimator.

Characteristic interactions and Fourier features raise Sharpe ratios and reduce pricing errors relative to the baseline specification. Timing gains are less uniform. The earnings-price ratio and recent market returns help individually, but combining all fourteen shared predictors does not improve on baseline IPCA. Extensions with factor-specific predictors also identify useful individual signals, including the value spread and recent factor returns.

The factor count associated with strong IPCA pricing performance varies across specifications and can be substantially smaller than the number of instruments. With twenty IPCA factors, the baseline IPCA SDF has a Sharpe ratio of 2.83, compared with 2.28 for the SDF estimated from all 153 characteristic portfolios. Twenty IPCA factors from the interaction specification yield a Sharpe ratio of 3.61, compared with 2.55 for the benchmark based on 11,934 managed portfolios. The Fourier specification benefits from more IPCA factors per draw: the 10-draw ensemble has a Sharpe ratio of 1.95 with twenty, 4.36 with one hundred, and 4.56 with the 210 selected by validation.

I also provide an economic rationale for a parsimonious SDF based on common risks and diversification. Many instruments may help measure exposure to the same common shocks. If the factors capture nondiversifiable risks, hedging those exposures leaves risks that can be reduced by spreading positions across assets. A restriction on near-arbitrage opportunities then limits expected returns on these diversified positions: expected returns cannot remain large as risk becomes small. This restriction supports approximate pricing by a factor SDF without requiring a separate factor for every instrument. Section 4 formalizes the argument under a common factor structure with diversifiable residual risk and a uniform bound on attainable Sharpe ratios.

This paper contributes to research on cross-sectional asset pricing in high-dimensional settings. Bryzgalova et al. (2025), Kozak et al. (2020), Kozak (2021), and Didisheim et al. (2024) show how flexible portfolio construction and regularization can use information from many characteristics. Kelly et al. (2019) estimate common factors while allowing stock exposures to depend on characteristics. I expand the instruments used in IPCA and construct SDFs from the resulting IPCA factors. This approach relates to recent work on the difference between the size of a model and the information it uses for pricing.[2] My results

[2]Kelly and Malamud (2025) distinguish nominal from effective complexity; Fan et al. (2026) learn characteristic directions that matter for pricing. Kelly et al. (2025) study flexible SDFs using a transformer,

indicate that flexible characteristic information and an SDF with fewer IPCA factors can be complementary.

A second strand of research studies time-series predictability and factor timing. Early contributions include Shiller (1981), Fama and French (1988), and Fama and Bliss (1987). Welch and Goyal (2008) and Goyal et al. (2024) document the difficulty of forecasting the equity premium outside the estimation sample. Haddad et al. (2020) study the connection between factor timing and the SDF, while Lehnherr et al. (2025) show how regularization can extract timing gains from many factors and predictors. I examine timing information alongside nonlinear characteristic information within an IPCA framework. Specific predictors improve SDF estimation, but including more timing information does not consistently improve performance.

Arbitrage pricing theory and mean–variance analysis underlie the economic argument. Ross (1976) and Chamberlain and Rothschild (1983) relate a common factor structure and diversification to restrictions on expected returns. Kozak et al. (2018, 2020) explain why many characteristics can be useful even when they measure a smaller set of common risks. I use this reasoning to interpret the difference between the dimension of the instrument set and the number of factors in the estimated SDF. When common factors leave diversifiable residual risk, a bound on attainable Sharpe ratios limits expected returns unexplained by factor exposures. This conditional rationale for approximate factor pricing leaves the empirical SDF dimension to be assessed from pricing performance.

Section 2 describes portfolio construction, IPCA, and SDF estimation. Section 3 presents the empirical results. Section 4 develops an economic rationale for the difference between the number of instruments and the number of factors used in the SDF. Section 5 examines robustness, and Section 6 concludes. The appendices provide variable definitions, details of the Fourier construction, and additional figures and portfolio diagnostics.

2 An Asset Pricing Framework

I construct and estimate SDFs from managed portfolios and from IPCA factors. The two approaches represent the same underlying characteristic and timing information differently.

and Hellum et al. (2026) examine portfolio choice when the number of assets can exceed the length of the estimation sample. Dickerson et al. (2026) show that a few underlying sources of priced risk can require many observable factors to measure them.

2.1 Pricing Restrictions

At time t, let r_{t+1} denote an $N \times 1$ vector of excess returns for N stocks, and let b_t be an $N \times 1$ vector of time-varying parameters. Hansen and Jagannathan (1991) characterize the relation between asset returns and admissible SDFs. I consider an SDF m_{t+1} in the linear span of excess returns,

$$m_{t+1} = 1 - b_t' r_{t+1}$$

and solve for the $N \times 1$ vector of SDF loadings b_t that satisfies the pricing equation $E_t[m_{t+1} r_{t+1}] = 0$.

Parameterize b_t as a function of an $N \times C$ matrix of observable characteristics Z_t,

$$b_t = f(Z_t)\lambda_t$$

where $f(Z_t)$ maps Z_t to an $N \times F$ matrix, and λ_t is an $F \times 1$ parameter vector. Substituting into the SDF $m_{t+1} = 1 - b_t' r_{t+1}$ gives

$$m_{t+1} = 1 - \lambda_t' f_{t+1}$$

where $f_{t+1} = f(Z_t)' r_{t+1}$ is an $F \times 1$ vector of zero-cost factor returns. For example, the CAPM uses market equity to form the factor weights $f(Z_t)$, and the Fama–French three-factor model uses market equity and the book-to-market equity ratio to form $f(Z_t)$. Using excess returns makes the factor portfolios zero-cost. The conditional pricing equation is

$$E_t[m_{t+1} f_{t+1}] = 0$$

where the factors f_t serve as both the returns to be priced and the candidate factors in the SDF.

2.2 SDF Construction

The high-dimensional SDF uses excess returns of N stocks $r_t^{(N\times 1)}$ and an $N \times J$ characteristic rank matrix $X_{t-1}^{(N\times J)}$. For each stock i at time t, I transform the characteristics $X_{i,t-1}$ into an L-dimensional feature space $f(X_{i,t-1})$, where $f(X_{i,t-1}) : \mathbb{R}^J \to \mathbb{R}^L$. Following Kozak et al. (2020), I demean these features and divide by the sum of absolute deviations, without applying a second ranking, to form the factor portfolio weight matrix $\phi(X_{i,t-1})^{(N\times L)}$. These weights give the $L \times 1$ vector of characteristic-managed long–short portfolio returns at time t, $F_t = \phi(X_{i,t-1})' r_t$. The high-dimensional SDF is a linear combination of all constructed

factors, $M_{t+1} = 1 - \lambda' F_{t+1}$. Algorithm 1 summarizes the construction of these SDFs, called Complex SDFs in the figures.

Algorithm 1 High-Dimensional SDF Construction

1: **Input:** Start with excess returns of N stocks $r_t^{(N\times 1)}$, $N \times J$ characteristic rank matrix $X_{t-1}^{(N\times J)}$.
2: **Output:** Factor portfolio returns F_t and the Complex SDF.
3: **Step 1:** For each stock i at time t, transform the characteristics $X_{i,t-1}$ into a L-dimensional feature space $f(X_{i,t-1})$, where $f(X_{i,t-1}) : \mathbb{R}^J \to \mathbb{R}^L$.
4: **Step 2:** Demean $f(X_{i,t-1})$ and normalize by dividing by the sum of absolute deviations (without applying a second ranking), forming the factor portfolio weight matrix $\phi(X_{i,t-1})^{(N\times L)}$.
5: **Step 3:** Construct the $L\times 1$ vector of characteristic-managed long–short portfolio returns at time t as

$$F_t = \phi(X_{i,t-1})' r_t$$

6: **Step 4:** For each standardized timing predictor x_{t-1}, add the portfolios $x_{t-1}F_t$ to the baseline returns, without further normalization.
7: **Step 5:** Estimate the coefficients on all managed returns using the ridge penalty selected in Section 2.4. The SDF is

$$M_{t+1} = 1 - \hat{\lambda}' F_{t+1}.$$

IPCA factor estimation. Algorithm 2 summarizes the IPCA-based SDF construction used in the figures. I estimate the characteristic-dependent loading model of Kelly et al. (2019) at its largest supported rank, rotate the factor returns, and select the first K to construct the SDF:

$$r_{i,t+1} = z_{i,t}' \Gamma f_{t+1} + \varepsilon_{i,t+1}.$$

Here $z_{i,t}$ contains L characteristic instruments, Γ is $L \times K_{\max}$, and f_{t+1} contains $K_{\max}$ factor returns. The model has no separate characteristic intercept or observed factors. I demean the main IPCA instruments without normalizing them to unit gross exposure, unlike the managed portfolio weights. Using the past T periods, I estimate Γ and the factor returns by minimizing

$$\frac{\sum_t \|r_{t+1} - Z_t \Gamma f_{t+1}\|^2 + \tau \sum_t N_t \|\Gamma f_{t+1}\|^2}{\sum_t N_t}, \tag{1}$$

where Z_t stacks the stock instruments and N_t is the number of stocks. The stock-level penalty is zero for the main baseline, Fourier, erf, and Chebyshev specifications. The interaction and

timing specifications use a positive penalty chosen by validation and are labeled regularized IPCA. This penalty is separate from the SDF penalty described below.

The role of the rank restriction depends on the maximum fitted dimension. Set $c_t = \Gamma f_{t+1}$ and collect the training coefficients in $C = [c_1, \ldots, c_T] \in \mathbb{R}^{L \times T}$. Apart from the constant denominator, equation (1) is equivalent to

$$\min_{\text{rank}(C) \leq K_{\max}} \sum_{t=1}^{T} \left[\|r_{t+1} - Z_t c_t\|^2 + \tau N_t \|c_t\|^2 \right]. \tag{2}$$

For the expanded main monthly specifications, $K_{\max} = T = 240$ and $L \geq 240$. Every $L \times 240$ coefficient matrix already has rank at most 240, so the rank restriction is vacuous. The minimization separates across training months. For $\tau > 0$, its unique monthly solution is

$$\hat{c}_t = (Z_t' Z_t + \tau N_t I_L)^{-1} Z_t' r_{t+1}; \tag{3}$$

for $\tau = 0$, any monthly least-squares solution minimizes the objective. When Z_t is rank deficient, coefficient solutions need not be unique, and the numerical choice of solution can matter for subsequent compression. The normalization $\Gamma'\Gamma = I$ does not restore a low-rank restriction: an orthonormal basis of the column space of any such C gives a factorization, padded to $K_{\max}$ columns if necessary. The baseline uses a different numerical support cap, $K_{\max} = 149 < T$, so the $K_{\max} = T$ argument applies directly to the expanded specifications rather than to every fit in the paper.

At the saturated maximum, the factorization imposes no cross-month low-rank restriction on the coefficient vectors. Ridge shrinkage, instrument construction, and numerical support restrictions can still regularize the fit. A smaller reported K compresses the fitted model afterward; stock estimation itself is not subject to a rank-K restriction. A reported count of twenty IPCA factors therefore refers to twenty factors selected from this fit. Table 2 compares this procedure with conventional low-rank IPCA for the baseline characteristics.

I normalize the fitted loadings so that $\Gamma'\Gamma = I$ and rotate the factors by the eigenvectors of their uncentered training second moment matrix:

$$\frac{1}{T} \sum_{t=1}^{T} f_{t+1} f_{t+1}' = V D V', \qquad P_{t+1} = V' f_{t+1}, \qquad d_1 \geq d_2 \geq \cdots \geq d_{K_{\max}}.$$

The IPCA factors are ordered by their average squared training returns. For each count K, I use the first K IPCA factors, $P_t^{(K)}$, to construct $M_{t+1} = 1 - \lambda_K' P_{t+1}^{(K)}$. I estimate the largest

supported IPCA model once in each training window and select factors from that fit for each K. With orthonormal Γ, this rotation and factor selection give a leading singular-component approximation to the fitted coefficient matrix $\hat{C}$ in Euclidean norm. They generally do not minimize the original stock loss subject to $\text{rank}(C) \leq K$, because that loss weights coefficient errors using the month-specific matrices $Z_t'Z_t + \tau N_t I_L$.

For subsequent returns, I jointly extract all $K_{\max}$ factors using all fitted loading directions, apply the training rotation, and select the first K without rescaling. Throughout the paper, "IPCA factors" refers to the rotated factor returns selected by this procedure, and "IPCA SDF" to the SDF constructed from them. Let $W_t = Z_t'Z_t/N_t$ and define

$$q_t = \#\{j : \lambda_j(W_t) > 10^{-10}\lambda_{\max}(W_t)\}, \qquad K_{\max} = \min\{T, L, \min_t q_t\}.$$

The minimum is taken over all 649 required formation months, December 1969–December 2023, covering training, validation, and evaluation factor extraction. The monthly rank is 149 in December 1969–September 1971 and 153 thereafter. In those first 22 months, `eqnetis_at`, `eqnpo_me`, `eqpo_me`, and `netis_at` have no observed cross-sectional characteristic variation: their median-filled ranks become zero after demeaning, giving four zero columns. All 153 inputs remain in the model. The same cap of 149 IPCA factors is used in every annual fit, including later windows whose monthly instrument matrices have full rank. Appendix A lists every input. The expanded main specifications support 240 IPCA factors. I evaluate counts $1, 5, 10, 15, \ldots$ and the exact common maximum.

Algorithm 2 IPCA Factor and SDF Construction

1: **Input:** Stock excess returns r_{t+1} and characteristic instruments Z_t observed before those returns.
2: **Output:** K IPCA factors $P_{t+1}^{(K)}$ and their SDF.
3: **Step 1:** Form the instruments for the specification. For the main models, transform and demean characteristics without applying a second ranking or normalizing to unit gross exposure.
4: **Step 2:** Fit equation (1) once at $K_{\max}$ (149 baseline; 240 expanded). When $K_{\max} = T$, this permits separate monthly least-squares or ridge coefficient estimates, as in equations (2)–(3).
5: **Step 3:** Normalize $\Gamma'\Gamma = I$ and rotate the fitted factors in decreasing order of their uncentered training second moments.
6: **Step 4:** Use the first K IPCA factors without rescaling or re-estimating a rank-K stock model. Estimate their SDF coefficients using the validation-selected ridge penalty in Section 2.4.
7: **Step 5:** Extract all $K_{\max}$ subsequent factors jointly, rotate using the training basis, select K IPCA factors, and construct

$$M_{t+1} = 1 - \hat{\lambda}_K' P_{t+1}^{(K)}.$$

2.3 Data and Construction of Factors

The Global Factor Data of Jensen et al. (2023) provide 406 characteristics and their corresponding factor portfolios. I use the U.S. equity subset, including 153 firm-level characteristics from the factors considered in Jensen et al. (2023). Returns come from the Center for Research in Security Prices (CRSP) database.

I use fourteen shared time-series predictors. The first eight are the equity premium predictors of Welch and Goyal (2008), obtained from the Goyal et al. (2024) dataset[3]: the dividend-price ratio (dp), earnings-price ratio (ep), book-to-market ratio (bm), net equity expansion (ntis), Treasury bill rate (tbl), term spread (tms), default spread (dfy), and stock variance (svar). To these I add the correlation neglect signal (cn) of Guo and Wachter (2026), a market timing signal that exploits investors' overreaction to the predictably repetitive earnings news released across the quarterly announcement cycle, and five macroeconomic predictors following Lehnherr et al. (2025), constructed from Treasury yields, corporate bond yields, and inflation from the Federal Reserve Bank of St. Louis: the real 1-year Treasury yield (real1y), the yield spread between 5-year and 1-year Treasuries (slope5y1y),

[3]Available at Prof. Amit Goyal's website.

the 12-month change in the 1-year yield (d1yield12m), the trailing 3-month market return (mkt3m), and the Baa–Treasury credit spread (baatsy).

I construct factors from the characteristics data of Jensen et al. (2023). I exclude observations with missing market equity in month t or missing returns in month $t+1$. To limit the influence of small, illiquid stocks, I also exclude stocks classified as nanocaps in the source data. Missing characteristics receive the median rank of 0.5.

Following Kozak et al. (2020) and Didisheim et al. (2024), I transform each variable to its cross-sectional percentile rank on [0, 1]. Subtracting each rank's cross-sectional mean gives the baseline IPCA instruments. For the characteristic-managed portfolio weights, I also divide these demeaned ranks by their sum of absolute values, giving each portfolio unit gross exposure.

I expand the fixed set of raw stock characteristics into P features $f(Z_t)$. After the nonlinear transformation, I demean each feature without ranking it again. These deterministic transformations enlarge the instrument basis without adding information beyond the underlying characteristics. The demeaned features are the IPCA instruments; for the Complex SDF, I also rescale each feature to unit gross exposure. For polynomial interactions, I take any two rank-transformed characteristics, $x^i_{s,t}$ and $x^j_{s,t}$, of stock s at time t, and form their degree-2 product from the ranks on $[0, 1]$. The managed portfolio weight, $z^{ij}_{s,t}$, is

$$z^{ij}_{s,t} = \frac{\left(x^i_{s,t}x^j_{s,t} - \frac{1}{n_t}\sum_{s=1}^{n_t} x^i_{s,t}x^j_{s,t}\right)}{\sum_{s=1}^{n_t}\left|x^i_{s,t}x^j_{s,t} - \frac{1}{n_t}\sum_{s=1}^{n_t} x^i_{s,t}x^j_{s,t}\right|}.$$

The corresponding IPCA instrument is the numerator, without division by gross exposure. The interaction specification includes the 153 ranks and all 11,781 distinct products, including squares, for a total of 11,934 instruments. I also consider nonlinear transformations based on random features, as in Didisheim et al. (2024) and Kozak (2021).

Factor timing portfolios interact time-series macroeconomic predictors with the 153 original factors. At each month-end, I convert each predictor to a z-score using its trailing 20-year mean and standard deviation through that month. Let F_t denote the return of a zero-cost long–short equity factor in month t, and let X_{t-1} denote a macroeconomic predictor measured at the end of month $t-1$, standardized using that trailing mean and standard

deviation. The factor timing portfolio return in month t, G_t, is the product of the lagged standardized predictor and the contemporaneous factor return:

$$G_t = X_{t-1}F_t$$

For IPCA, I likewise interact each baseline characteristic with the predictor and include both the baseline and interacted instruments. This gives 306 instruments with one predictor and 2,295 with all fourteen. Each specification is fitted separately. Factor-specific timing also uses each portfolio's own return, volatility, or characteristic spread; Appendix C defines all six signals, their weighting and lags, and the four startup exceptions used in Table 4.

2.4 SDF Estimation

I estimate the coefficients of the SDF

$$M_{t+1} = 1 - \sum_p \lambda_p F_{p,t+1},$$

by regressing one on portfolio returns to obtain λ. Here F_{t+1} denotes all managed returns for the Complex SDF or the IPCA returns $P_{t+1}^{(K)}$. The regression approach follows Theorem 1 of Britten-Jones (1999): an OLS regression of a constant $\mathbf{1}$ on asset excess returns without an intercept gives coefficients proportional to mean–variance efficient portfolio weights when the relevant moments are nonsingular.

To limit overfitting, I use a ridge-penalized objective following Kelly and Xiu (2023) and Kozak et al. (2020):

$$\begin{aligned}\hat{\lambda}_{K,\alpha} &= \arg\min_{\lambda} \left\{ \frac{1}{T}\sum_{t=1}^{T}(1 - \lambda' P_t^{(K)})^2 + \alpha\|\lambda\|^2 \right\} \\ &= \left(\hat{S}_K + \alpha I\right)^{-1} \hat{\mu}_K, \end{aligned} \tag{4}$$

where $\hat{\mu}_K$ is the training mean and $\hat{S}_K$ is the uncentered second moment matrix of IPCA returns. Without the penalty, these coefficients are proportional to tangency portfolio weights. The penalty is expressed per observation: its equivalent in a sum of squared errors is $T\alpha$, where T is the actual number of training months or trading days. For the Complex SDF, the sum of squares penalty is $gL^{-1}\sum_t \|F_t\|^2$, where L is the number of managed portfolios. Its coefficients are estimated using all L returns.

I estimate the SDF coefficients over rolling windows, shown in Figure 16 in Appendix D. At the end of January in each year from 2000 to 2023, I use the preceding 240 months of

available returns to estimate the IPCA loadings and SDF coefficients. The January 2000 fit uses returns from February 1980 through January 2000 and is applied to returns from February 2000 through January 2001. Repeating this procedure gives evaluation returns from February 2000 through January 2024. Loadings and SDF coefficients remain fixed between annual fits, while stock characteristics and factor exposures can change monthly. The robustness comparisons also consider 60-month windows and daily training returns.

To select the penalties, I use ten annual validation periods over 1990–1999. For each validation year, the standard monthly model is estimated using the preceding 240 months of realized returns and evaluated over the next twelve months. The first fit uses January 1970–December 1989 returns; the last uses January 1979–December 1998. For each candidate penalty, I pool the validation returns across these ten years and select the penalty with the highest annualized Sharpe ratio. The selected penalties remain fixed during the evaluation period.

For each IPCA factor count, I select α from 61 logarithmically spaced values between 10^{-10} and 10^{2}. For a regularized stock fit, I fit the full $K_{\max}$ model for every $\tau \in \{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1\}$. For each candidate τ, I optimize α over the entire 61-point grid using pooled validation SDF returns. Thus,

$$\tau^* \in \arg\max_{\tau} \max_{\alpha} \widehat{\mathrm{SR}}_{1990\text{–}1999}(\tau, \alpha; K_{\max}).$$

The comparison across τ thus profiles out α. Ties favor larger τ, then larger α. I fix τ^* across IPCA factor counts and reselect α_K for each K using that count's validation returns; ties at this stage favor larger α_K. The Complex SDF uses 50 logarithmically spaced g values between 10^{-4} and 10^{2}. For Fourier features, the auxiliary managed-return search in Appendix B selects the map and benchmark g; I retain the selected g when estimating the benchmark on all selected features. The short-window comparison uses five preceding years for each validation fit; daily estimation uses trading observations within the preceding twenty calendar years. The validation years remain the same in both comparisons.

2.5 Evaluating SDF Models

I evaluate SDF performance using portfolio Sharpe ratios and pricing errors.

I construct out-of-sample portfolios implied by the estimated SDF. At time $t+1$, the SDF is $M_{t+1} = 1 - \hat{\lambda}_t' F_{t+1}$. Its implied portfolio return is

$$r_{t+1}^{\mathrm{SDF}} = \hat{\lambda}_t' F_{t+1}.$$

I report its annualized Sharpe ratio, calculated as the mean monthly return divided by its standard deviation and multiplied by $\sqrt{12}$.

I also compute the Hansen–Jagannathan specification-error distance (HJD) of Hansen and Jagannathan (1997). It aggregates squared test-asset pricing errors, weighted by the inverse uncentered second moment matrix of test returns. For test asset returns R_{t+1} realized in the same month as M_{t+1}, the squared distance is

$$\widehat{D}^2_{HJ} = \hat{\mathbb{E}}_{OOS}[M_{t+1}R_{t+1}]' \left(\hat{\mathbb{E}}_{OOS}[R_{t+1}R'_{t+1}]\right)^{+} \hat{\mathbb{E}}_{OOS}[M_{t+1}R_{t+1}],$$

where $+$ denotes the Moore–Penrose pseudoinverse and $\hat{\mathbb{E}}_{OOS}[M_{t+1}R_{t+1}]$ is the average out-of-sample pricing error.

The test assets consist of the 153 value-weighted anomaly factors constructed by Jensen et al. (2023) and available from the Global Factor Data repository. Under a correctly specified SDF, the population moment $\mathbb{E}[M_{t+1}R_{t+1}]$ is zero. Lower sample distances therefore indicate smaller pricing errors. SDF and test asset returns are matched by the month in which returns are realized.

3 Empirical Results

I compare how firm characteristics and time-series predictors contribute to SDF performance, beginning with traditional factor benchmarks, then examining characteristics and their nonlinear transformations, and finally adding timing predictors. The main IPCA curves vary the number of IPCA factors used in the SDF. I report portfolio Sharpe ratios and, for the IPCA and Complex SDF comparisons, pricing errors on the same external test assets.

Table 1 evaluates timing within four small factor sets. Its SDF portfolios use the same monthly frequency, 240-month training windows, fixed 1990–1999 validation, and February 2000–January 2024 evaluation months as the main comparisons. Each cell uses the all-return ridge estimator and penalty normalization of the Complex SDF in Section 2.4, allowing a comparison of payoff and timing information under a common design.

Table 1: Annualized Out-of-Sample Sharpe Ratios of Traditional Factor SDFs and Timing Portfolios

	Base + Timing				**Timing Only**			
Predictor	CAPM	FF3	FF5+Mom	HXZ	CAPM	FF3	FF5+Mom	HXZ
Base factors	0.27	0.28	0.46	0.71				
Equity-premium predictors								
d/p	−0.06	0.45	0.88	0.46	−0.17	0.13	0.49	0.37
e/p	0.25	0.37	0.80	0.95	−0.49	−0.02	−0.08	−0.22
b/m	−0.02	0.14	0.60	0.42	−0.33	−0.22	0.23	−0.06
ntis	0.17	0.22	0.77	0.69	−0.16	−0.21	−0.28	−0.24
tbl	0.04	0.25	0.70	0.75	−0.05	0.07	−0.02	0.15
tms	0.23	0.24	0.42	0.69	−0.06	0.05	−0.40	−0.08
dfy	0.38	0.55	0.67	0.77	0.24	0.31	0.30	0.14
svar	0.04	0.04	0.11	0.04	−0.34	0.31	0.40	−0.32
Correlation neglect signal								
cn	0.40	0.40	0.83	0.78	0.37	0.28	0.29	0.29
Macroeconomic predictors								
real1y	0.35	0.00	0.52	0.70	−0.35	−0.19	−0.01	0.13
slope5y1y	0.24	0.22	0.46	0.69	−0.03	−0.03	−0.38	−0.09
d1yield12m	0.23	0.27	0.52	0.58	−0.09	0.09	0.07	0.05
mkt3m	0.10	0.20	0.58	0.53	−0.15	0.06	0.21	−0.08
baatsy	−0.13	0.09	0.39	0.54	−0.22	0.03	0.12	0.14
Factor-specific predictors								
3m ret	0.12	0.14	0.12	0.29	−0.11	0.00	−0.03	−0.08
12m ret	0.29	0.11	0.15	0.29	−0.02	−0.13	−0.14	−0.25
vol	0.03	0.46	0.22	0.18	−0.29	0.36	−0.11	−0.18
All 17 predictors	0.45	0.62	0.65	0.54	0.33	0.51	0.53	0.39

Notes: Returns are monthly; evaluation covers February 2000–January 2024 (288 months), and each reported ratio is $\sqrt{12}\,\bar{r}/s_r$, with sample standard deviation s_r. Each cell fits a ridge regression of one on the indicated excess returns without an intercept. During evaluation, coefficients are refitted each January using 240 trailing months, including January. For each cell, g is selected from 50 log-spaced values between 10^{-4} and 10^{2} by pooled 1990–1999 validation Sharpe ratio, with 240 preceding training months per annual fold; g is then fixed during evaluation. The sum-of-squares ridge penalty is $gL^{-1}\sum_t \|F_t\|^2$ for L returns. Base + Timing combines the published factors with their predictor-scaled returns; Timing Only excludes the base returns. The fourteen shared predictors are defined in Section 2.3. The final three individual rows use each factor's own lagged 3- or 12-month return sum or 12-month return volatility, standardized over up to 240 months (minimum 24). All signals are known at portfolio formation. The final row jointly includes all seventeen predictors.

I begin with four published factor sets. CAPM uses the market excess return; the Fama–French three-factor model (FF3) adds size and value. FF5+Mom adds profitability, investment, and momentum. FF3 and FF5 use their respective published size factors. The Hou–Xue–Zhang model (HXZ) uses market, size, investment, and return-on-equity factors (Hou et al., 2015). Table 1 reports the performance of fitted SDF-implied combinations of these returns. The base SDFs have annualized Sharpe ratios of 0.27, 0.28, 0.46, and 0.71, respectively.

Individual predictors seldom improve performance consistently across all four models, and several materially reduce it. The default spread and correlation neglect signal improve all four base SDFs; for example, correlation neglect raises the FF5+Mom Sharpe ratio from 0.46 to 0.83. Other signals have mixed effects. Adding stock variance lowers the FF3 ratio from 0.28 to 0.04 and the HXZ ratio from 0.71 to 0.04. Timing each factor by its own volatility lowers FF5+Mom from 0.46 to 0.22 and HXZ from 0.71 to 0.18. These changes show that adding a predictor can substantially alter the estimated SDF even when the base factors remain available.

Combining all seventeen predictors raises the CAPM, FF3, and FF5+Mom ratios to 0.45, 0.62, and 0.65, respectively, but lowers HXZ to 0.54. The corresponding Timing Only portfolios have ratios of 0.33, 0.51, 0.53, and 0.39. Under this common protocol, timing adds useful payoffs for some traditional factor sets, with gains that depend on the model and predictor set.

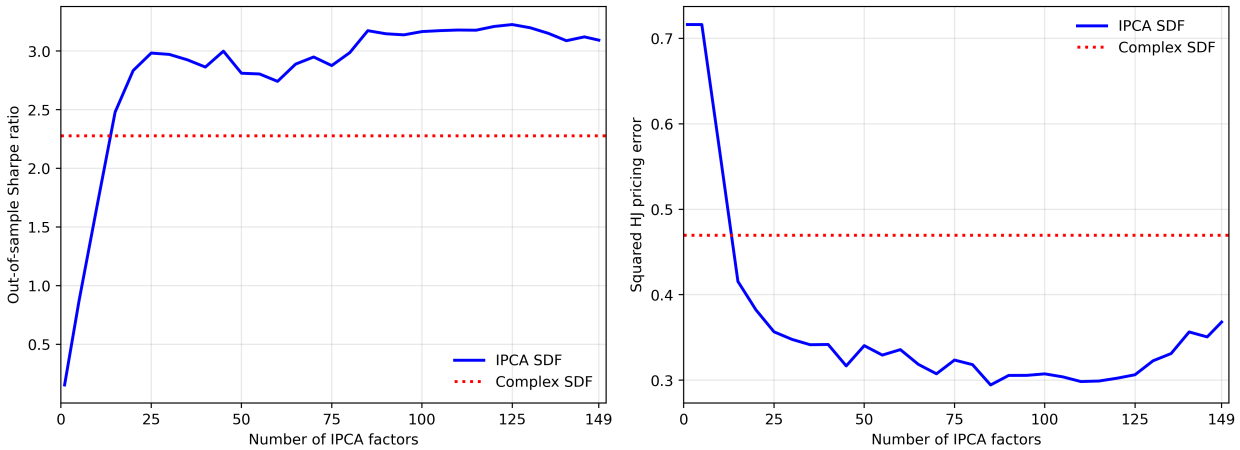


Figure 1: Out-of-Sample Sharpe Ratios and Pricing Errors by Number of IPCA Factors: 153 Characteristics. This figure compares the performance of IPCA SDFs (blue solid line) and Complex SDFs (red dotted line) using 153 characteristics. The left panel shows out-of-sample Sharpe ratios of SDF-implied portfolios, while the right panel shows pricing errors, both plotted against the number of IPCA factors used in the IPCA SDF. The out-of-sample period spans from February 2000 to January 2024. Penalties are selected using fixed 1990–1999 validation.

Figure 1 shows how pricing performance varies with the number of IPCA factors extracted using the 153 baseline characteristics. With twenty factors, the IPCA SDF has a Sharpe ratio of 2.83 and a squared Hansen–Jagannathan distance of 0.382, compared with 2.28 and 0.470 for the SDF estimated from all managed portfolios. Using all 149 IPCA factors gives

a Sharpe ratio of 3.09 and a squared distance of 0.368. Performance is not monotone in the factor count, but twenty factors improve both measures relative to the direct managed portfolio benchmark.

Comparison of estimation and compression. Table 2 holds the 153 baseline characteristics and evaluation protocol fixed and compares SDFs based on three sets of factor returns: the IPCA factors defined in Section 2.2; IPCA factors estimated separately under each rank-K restriction; and principal components from the same fixed set of 153 unit-gross-exposure managed portfolios. Both IPCA methods use $\tau = 0$. For the third method, I diagonalize the uncentered training second moment of managed returns, use unit-norm portfolio loadings, and apply those loadings to subsequent returns. No method standardizes factor returns. Each method and count selects its own SDF ridge penalty from the same grid using the same 1990–1999 validation months. The separately estimated IPCA results use the accepted stock fits at each rank K.

At $K = 5$, the Sharpe ratio is 0.87 for the IPCA SDF, 2.68 for separately fitted IPCA, and 1.55 for managed-return PCA. At $K = 20$, the corresponding values are 2.83, 3.04, and 2.26; squared HJ distances are 0.382, 0.345, and 0.451. Separately fitted IPCA has a higher Sharpe ratio at every reported interior count, although the ordering of pricing errors differs at $K = 100$. The two IPCA procedures coincide at their common full count of 149. Imposing low rank during stock fitting and compressing afterward thus yield different empirical results. Comparisons with managed-return PCA also change the portfolio construction.

Figure 2 reports results with 11,934 instruments formed from the characteristics and their interactions. With twenty factors, the IPCA Sharpe ratio is 3.61, compared with 2.83 for the baseline at the same count. Using all 240 factors raises the Sharpe ratio to 4.15 and lowers the squared Hansen–Jagannathan distance to 0.247. The corresponding managed portfolio SDF has a Sharpe ratio of 2.55 and a squared distance of 0.618. Interactions improve the use of characteristic information, and an SDF with far fewer factors than instruments outperforms the direct benchmark. Adding factors beyond the first twenty further improves performance. Both the instruments and the stock ridge penalty change in this comparison.

I next construct 1,000 instruments per draw from random Fourier features of 153 characteristics. Random features (Kozak (2021) and Didisheim et al. (2024)) map a fixed set

Table 2: Estimation versus Compression: A Matched Baseline Comparison. All methods use the same 153 characteristics, 240-month training windows, ten validation years (1990–1999), and 288 evaluation months (February 2000–January 2024). The first method uses K IPCA factors selected from a 149-factor fit; the second estimates stock exposures separately at rank K; the third retains K uncentered PCs of the common managed-return panel. Each method selects its SDF penalty separately at each K, using the same 61-point grid and validation criterion. SR is the annualized Sharpe ratio and D^2_{HJ} is the squared HJ distance on the same 153 value-weighted test assets. This comparison covers the baseline, not the expanded specifications with $K_{\max} = T = 240$.

	IPCA SDF		Separate rank-K IPCA estimation		Managed-return PCA truncation	
K	SR	D^2_{HJ}	SR	D^2_{HJ}	SR	D^2_{HJ}
5	0.87	0.716	2.68	0.395	1.55	0.598
10	1.68	0.567	2.88	0.359	1.82	0.534
20	2.83	0.382	3.04	0.345	2.26	0.451
50	2.81	0.340	3.51	0.295	2.28	0.448
100	3.17	0.307	3.39	0.321	2.23	0.484
149	3.09	0.368	3.09	0.368	2.22	0.503

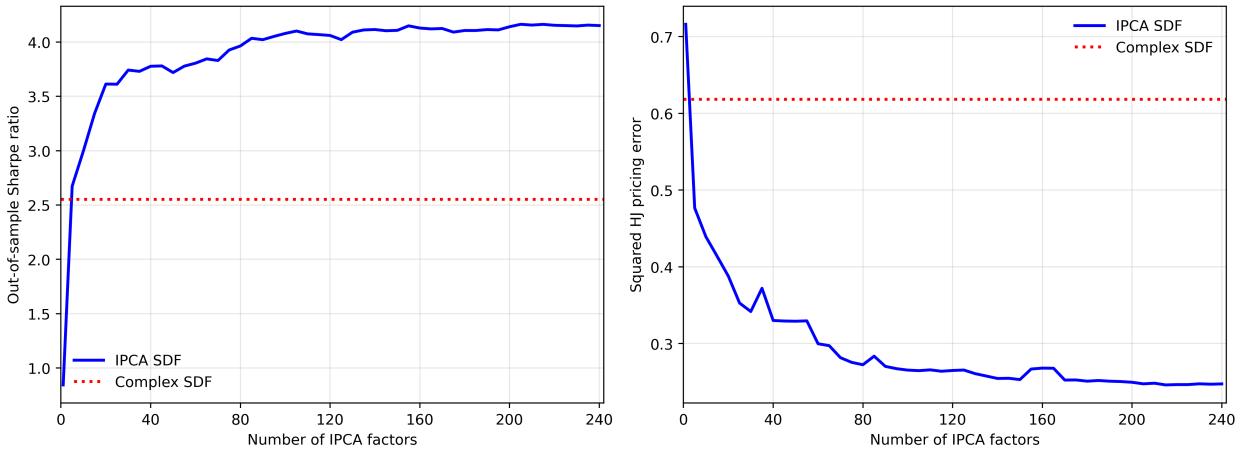


Figure 2: Out-of-Sample Sharpe Ratios and Pricing Errors by Number of IPCA Factors: 11,934 Instruments Constructed from Interactions of 153 Characteristics. This figure compares the performance of IPCA SDFs (blue solid line) and Complex SDFs (red dotted line) using 11,934 characteristic instruments. The left panel shows out-of-sample Sharpe ratios of SDF-implied portfolios, while the right panel shows pricing errors, both plotted against the number of IPCA factors used in the IPCA SDF. The out-of-sample period spans from February 2000 to January 2024. The stock penalty is $\tau = 0.1$; SDF penalties are selected using fixed 1990–1999 validation.

of raw stock characteristics into any chosen number of nonlinear features. Let $X_{i,t}$ be the 153×1 vector of characteristic ranks centered at 0.5. The transformation of $X_{i,t}$ is

$$S_{\ell,i,t} = [\sin(\gamma_\ell' X_{i,t}), \cos(\gamma_\ell' X_{i,t})], \quad \gamma_\ell = \gamma\omega_\ell, \quad \omega_\ell \sim \mathcal{N}(0, I).$$

$S_{\ell,i,t}$ applies sine and cosine to a random linear combination of $X_{i,t}$. Appendix B gives the construction details. The selected feature map uses 500 projection vectors and $\gamma = 0.1$, giving 1,000 paired features for each of ten fixed draws. The managed-return validation search in Appendix B selects the map used by both methods. I estimate IPCA separately for each draw and average the SDF portfolio returns before evaluating performance. At each IPCA factor count, the SDF penalty is selected using these averaged validation returns; the stock-fitting penalty is zero.

Figure 3 shows that the IPCA Sharpe ratio rises from 1.95 with twenty factors per draw to 4.36 with one hundred. With the 210 factors selected by validation, it is 4.56 and the squared Hansen–Jagannathan distance is 0.277, compared with 2.54 and 0.544 for the Complex SDF ensemble. Using all 240 factors gives a similar Sharpe ratio of 4.56 and a squared distance of 0.291. The Fourier specification therefore needs more factors than the baseline to achieve its gains.

Table 3 summarizes performance at the reported factor counts: 149 for the baseline specification, 240 for interactions, and 210 per draw for Fourier features. Together with Figures 1–3, it separates gains from expanding the instrument set from changes in performance as more factors enter the SDF. Interactions and Fourier features improve on baseline IPCA, but the factor counts associated with these gains differ. The smaller-count comparisons in Figures 1 and 2 show that both specifications can outperform their managed portfolio benchmarks using fewer than all available IPCA factors.

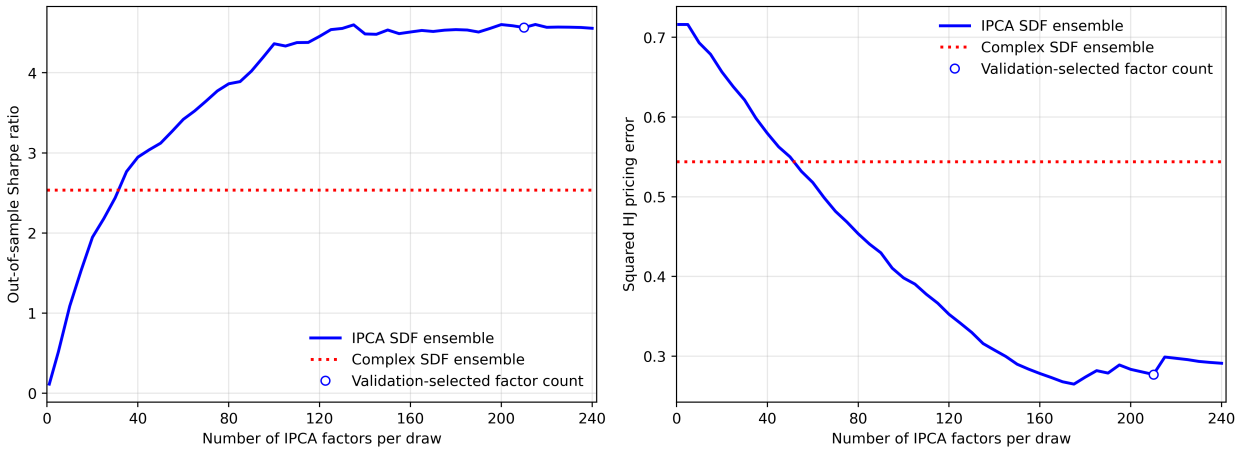


Figure 3: Out-of-Sample Sharpe Ratios and Pricing Errors by Number of IPCA Factors: 1,000 Random Fourier Features of 153 Characteristics. This figure compares the performance of IPCA SDFs (blue solid line) and Complex SDFs (red dotted line) using 1,000 random features of 153 characteristics per draw. The left panel shows out-of-sample Sharpe ratios of SDF-implied portfolios, while the right panel shows pricing errors, both plotted against the number of IPCA factors used in the IPCA SDF. The out-of-sample period spans from February 2000 to January 2024. Ten equally weighted draws each use 1,000 paired sine/cosine features. The factor count is per draw; the circle marks the validation-selected count of 210. Penalties are selected using fixed 1990–1999 validation.

Table 3: Sharpe Ratios for Three High-Dimensional Factor Models

SDF	Baseline (153)	Interactions (11,934)	Fourier (1,000)
Complex SDF	2.28	2.55	2.54
IPCA SDF	3.09	4.15	4.56

Notes: This table reports the annualized out-of-sample Sharpe ratios of SDF-implied portfolios constructed using three different high-dimensional factor models. The Complex SDF uses all managed portfolios; the IPCA SDF uses 149 IPCA factors for the baseline and 240 for interactions. The Fourier specification averages ten draws and uses its validation-selected count of 210 factors per draw. Returns span February 2000–January 2024 using a rolling 20-year estimation window, with penalties selected from 1990–1999 validation.

3.1 Alternative Nonlinear Feature Maps

To assess sensitivity to the choice of nonlinear transformation of the 153 characteristics, I evaluate two additional feature maps under the same rolling out-of-sample design. The first is a random *erf* map, $S_{\ell,i,t} = \text{erf}(g_\ell \omega_\ell' X_{i,t})$, where $X_{i,t}$ contains characteristic ranks centered at 0.5, $\omega_\ell \sim \mathcal{N}(0, I)$, and g_ℓ is drawn uniformly from $\{0.5, 0.6, 0.7, 0.8, 0.9, 1.0\}$. I draw 1,000 such features using seed 42 and hold them fixed over time. This construction is related to the arcsine, or "neural network" kernel. The second map uses a structured, deterministic basis of *Chebyshev orthogonal polynomials*, $T_2, \ldots, T_4$, applied separately to each characteristic rank

mapped onto $[-1, 1]$. It yields 459 instruments that capture smooth nonlinearities within each characteristic but no interactions across characteristics. The map excludes $T_1(u) = u$ and therefore does not contain the baseline linear basis. Including T_1 would give $4 \times 153 = 612$ columns; that nested specification is not estimated here.

Figure 4 reports out-of-sample Sharpe ratios and pricing errors by IPCA factor count for the SDF constructed from 1,000 random erf features of the 153 characteristics. The Complex SDF has an out-of-sample Sharpe ratio of 2.32. The IPCA Sharpe ratio rises from 1.38 at 20 factors to 2.80 at 100 and 3.37 at 240. At the largest count, its squared HJ distance is 0.297, compared with 0.458 for the Complex SDF. Because transformed features are demeaned without a second ranking, a strictly monotone activation can change relative feature magnitudes and portfolio weights. The gains continue well beyond the first ten factors.

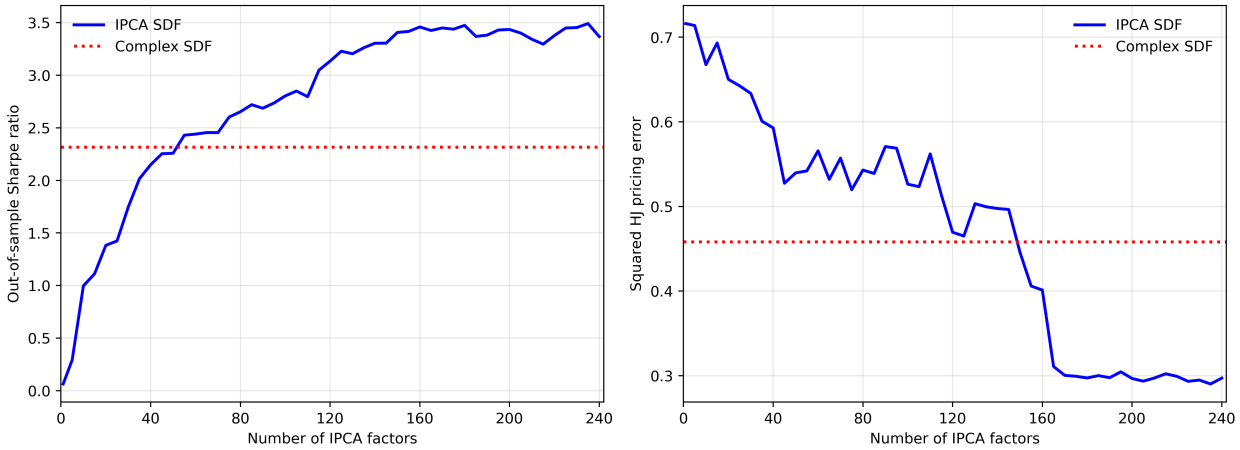


Figure 4: Out-of-Sample Sharpe Ratios and Pricing Errors by Number of IPCA Factors: 1,000 Random erf Features of 153 Characteristics. This figure compares the performance of IPCA SDFs (blue solid line) and Complex SDFs (red dotted line) constructed from 1,000 random erf features of the 153 characteristics. The left panel shows out-of-sample Sharpe ratios of SDF-implied portfolios, while the right panel shows pricing errors, both plotted against the number of IPCA factors used in the IPCA SDF. The out-of-sample period spans from February 2000 to January 2024. Penalties are selected using fixed 1990–1999 validation.

Figure 5 reports the corresponding results for the Chebyshev orthogonal polynomial basis of the 153 characteristics. The Complex SDF has an out-of-sample Sharpe ratio of 2.13. At 20 IPCA factors, the Sharpe ratio is 2.46 and the squared HJ distance is 0.452. At 240 factors, the corresponding values are 2.83 and 0.359, compared with a Complex SDF pricing error of 0.473. IPCA improves on the corresponding Complex SDF with both random and

structured nonlinearities, but performance and the number of useful IPCA factors depend on the feature map.

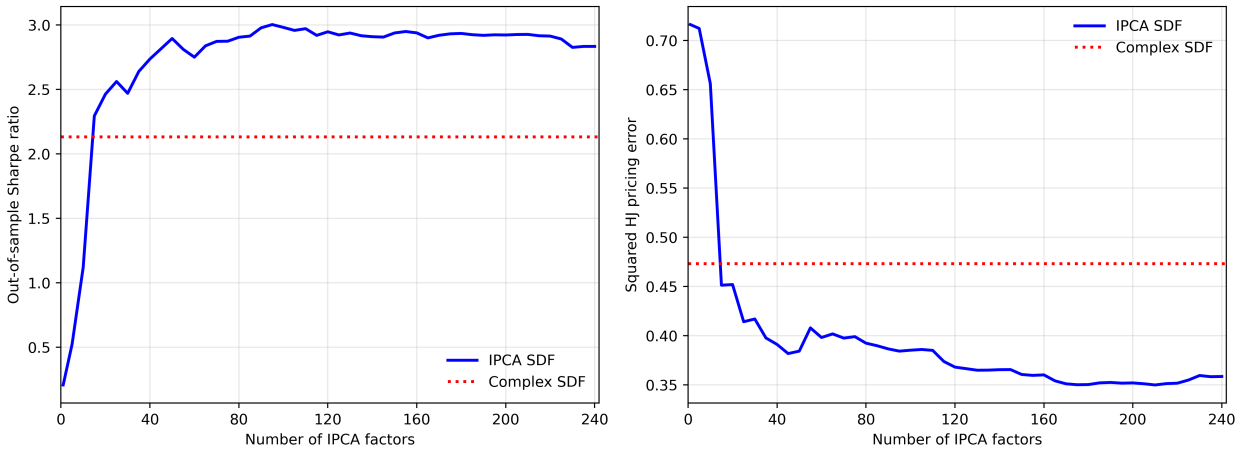


Figure 5: Out-of-Sample Sharpe Ratios and Pricing Errors by Number of IPCA Factors: 459 Instruments from an Alternative Chebyshev Feature Map of 153 Characteristics. This figure compares the performance of IPCA SDFs (blue solid line) and Complex SDFs (red dotted line) constructed from T_2, T_3, T_4 of each characteristic; the linear term T_1 is excluded. The left panel shows out-of-sample Sharpe ratios of SDF-implied portfolios, while the right panel shows pricing errors, both plotted against the number of IPCA factors used in the IPCA SDF. The out-of-sample period spans from February 2000 to January 2024. Penalties are selected using fixed 1990–1999 validation.

3.2 Factor Timing

I add time-series information through the fourteen shared predictors defined in Section 2.3. Each predictor scales the baseline managed returns and the corresponding stock characteristics using information available at portfolio formation. Figure 17 shows their standardized histories. Appendix C defines the factor-specific signals and their startup treatment.

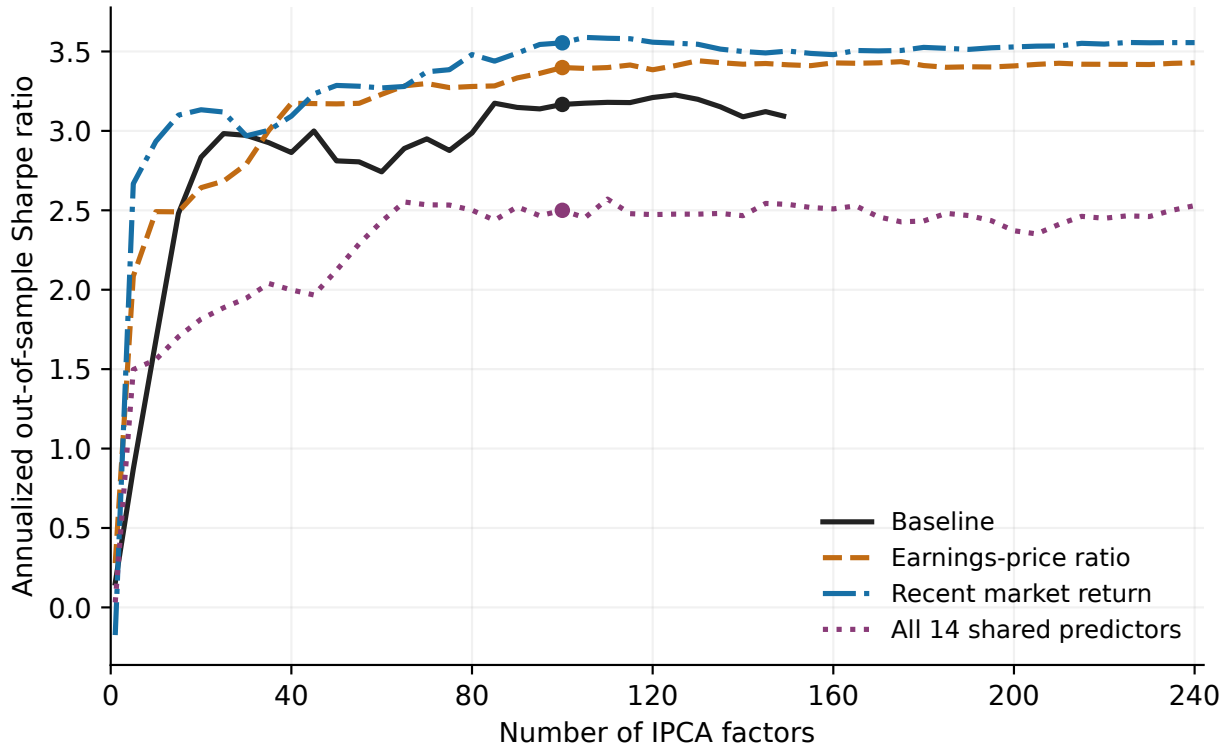


Figure 6: Central Timing Comparisons. Annualized out-of-sample Sharpe ratios for the baseline IPCA SDF, earnings-price timing, recent-market-return timing, and the joint fourteen-predictor specification. The latter three use the baseline characteristics and their timing interactions. The horizontal axis shows the number of IPCA factors, and dots mark $K = 100$. Penalties use fixed 1990–1999 validation; evaluation covers February 2000–January 2024. Complete individual-predictor comparisons with the Complex SDF appear in Appendix Figures 13–15.

Timing gains depend on the predictor. At a common count of 100 IPCA factors, adding the earnings-price ratio or the trailing 3-month market return raises the Sharpe ratio from 3.17 to 3.40 or 3.55, respectively; the other twelve individual predictors fall below the baseline at that count. Figure 6 compares these two cases with the baseline and joint timing model. The appendix reports all fourteen individual curves in larger panels with common axes.

Table 4: Out-of-Sample SDF Performance: Cross-Sectional Complexity versus Factor Timing

Instrument set	No. of instruments	Complex SDF OOS Sharpe ratio	Complex SDF Pricing error	IPCA SDF OOS Sharpe ratio	IPCA SDF Pricing error
153 base characteristics	153	2.28	0.47	3.09	0.37
Panel A: Nonlinear expansion of the 153 characteristics					
Degree-2 interactions	11,934	**2.55**	0.62	4.15	0.25
Random Fourier features	1,000	2.54	0.54	**4.56**	0.29
Random erf features	1,000	2.32	0.46	3.37	0.30
Chebyshev orthogonal polynomials	459	2.13	0.47	2.83	0.36
Panel B: Factor timing — macroeconomic predictors					
+ 1-year yield change (d1yield12m)	306	2.38	0.44	3.16	0.33
+ Trailing 3-month market return (mkt3m)	306	2.20	0.52	3.56	0.29
+ Earnings-price ratio (e/p)	306	2.07	0.49	3.43	0.30
+ Correlation neglect signal (cn)	306	2.11	0.48	3.13	0.33
+ Term spread (tms)	306	2.11	0.51	2.87	0.37
+ Treasury bill rate (tbl)	306	2.06	0.52	2.91	0.38
+ Net equity expansion (ntis)	306	2.02	0.49	2.73	0.41
+ Yield-curve slope (slope5y1y)	306	2.03	0.51	2.81	0.38
+ Dividend-price ratio (d/p)	306	1.89	0.55	2.44	0.45
+ Baa–Treasury spread (baatsy)	306	1.83	0.53	3.06	0.34
+ Real 1-year yield (real1y)	306	1.79	0.57	2.32	0.44
+ Stock variance (svar)	306	1.81	0.63	3.11	0.35
+ Book-to-market ratio (b/m)	306	1.51	0.68	2.10	0.53
+ Default spread (dfy)	306	1.30	0.79	2.81	0.39
Panel C: Factor timing — factor-specific predictors					
+ Profitability spread	306	1.93	0.51	3.07	0.36
+ 3-month factor return	306	1.91	0.58	3.53	0.28
+ 12-month factor return	306	1.72	0.62	2.92	0.36
+ Value (book-to-market) spread	306	2.08	0.50	3.70	0.28
+ Factor volatility (12-month)	306	1.70	0.52	3.28	0.33
+ Asset-growth spread	306	1.00	0.94	2.75	0.41
Panels B and C combined					
+ All timing predictors jointly	3,213	1.40	0.95	2.65	0.40

Notes: The Complex SDF uses all managed portfolios; the IPCA SDF uses all available IPCA factors: 149 for the baseline and 240 for each expanded model. Both are estimated annually over rolling 20-year windows, with penalties selected using fixed 1990–1999 validation. Sharpe ratios are annualized; pricing errors are squared Hansen–Jagannathan distances against the same 153 published value-weighted JKP anomaly factors of Jensen et al. (2023). Evaluation covers February 2000–January 2024. Panel A reports the nonlinear feature maps in Figures 2–5. For both SDFs, Fourier results average SDF-implied returns across ten random draws; the instrument count is per draw. Panels B and C augment the base characteristics with 153 interactions with a lagged, standardized predictor ($306 = 153 + 153$). Panel B uses fourteen shared macroeconomic predictors; Panel C uses six factor-specific predictors based on past factor returns, volatility, or long-minus-short characteristic spreads. Volatility uses twelve monthly factor returns. Appendix C gives the exact spread construction, signal lags, and named startup exceptions. The final row combines the base characteristics with all twenty sets of timing interactions ($3{,}213 = 153 + 20 \times 153$).

Combining all fourteen predictors does not improve on baseline IPCA. At 240 IPCA factors, the joint IPCA SDF has a Sharpe ratio of 2.53 and a squared HJ distance of 0.395, compared with 3.09 and 0.368 for the baseline with all 149 IPCA factors. It nevertheless improves on the joint Complex SDF, whose corresponding values are 1.01 and 1.793 (Figure 7).

Improving on the direct managed portfolio benchmark therefore need not imply a gain over characteristics alone.

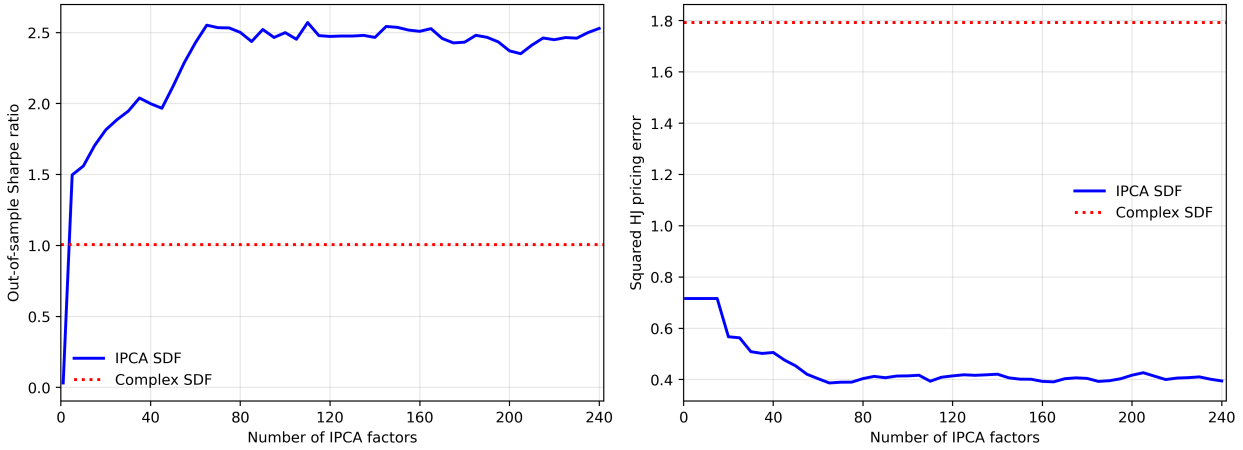


Figure 7: Out-of-Sample Sharpe Ratios and Pricing Errors versus Number of IPCA Factors: Joint Timing Specification (2,295 Portfolios). The timing portfolios scale each of the 153 base factors by each of the fourteen standardized predictors. The left panel reports the out-of-sample Sharpe ratio of the SDF-implied portfolio and the right panel the pricing error (squared Hansen–Jagannathan distance), for the IPCA SDF (blue solid line) and the Complex SDF (red dotted line), plotted against the number of IPCA factors. The predictors are the eight Welch and Goyal (2008) variables, the correlation neglect signal of Guo and Wachter (2026), and five macroeconomic predictors following Lehnherr et al. (2025). The out-of-sample period spans February 2000 to January 2024. The stock penalty is $\tau = 1$; SDF penalties are selected using fixed 1990–1999 validation.

Table 4 compares the out-of-sample performance of Complex and IPCA SDFs using the nonlinear characteristic feature maps in Figures 2–5 and factor timing instruments. Alongside the fourteen shared predictors, the table includes six factor-specific predictors based on each factor's past return, volatility, or characteristic spread. For IPCA, timing gains remain selective: the value spread and trailing 3-month factor return raise the Sharpe ratio from 3.09 for the baseline to 3.70 and 3.53, respectively. Combining all twenty timing predictors yields an IPCA Sharpe ratio of 2.65 and a squared HJ distance of 0.395, compared with 1.40 and 0.946 for the corresponding Complex SDF. Although IPCA improves on this managed portfolio benchmark, the joint timing specification does not improve on baseline IPCA, whose squared HJ distance is 0.368.

The results separate gains from additional instruments from gains from adding IPCA factors to the SDF. Characteristic interactions and Fourier features improve pricing performance, whereas timing gains depend on the predictor. Several specifications outperform

their managed portfolio benchmarks after substantial compression. Section 4 develops an economic rationale for approximate factor pricing based on common risks, diversification, and a restriction on near-arbitrage opportunities.

4 Common Risks, Diversification, and SDF Dimension

The empirical results distinguish the number of instruments describing investment opportunities from the number of IPCA factors used in the SDF. A parsimonious factor SDF has an economic rationale when common factors capture nondiversifiable risks and expected returns are bounded relative to risk. The common factor structure makes residual risk diversifiable; the restriction on near-arbitrage opportunities limits compensation for that risk. Portfolios that hedge common factor exposures then cannot offer large expected returns as diversification reduces their remaining risk. A factor SDF can therefore approximate pricing across a larger set of portfolios.

4.1 Instruments and Common Factor Exposures

IPCA uses observed instruments to describe stock exposures (Kelly et al., 2019). To interpret this parameterization economically, suppose a population common factor representation holds:

$$r_{i,t+1} = \beta_{i,t}' f_{t+1} + \varepsilon_{i,t+1}, \qquad \beta_{i,t} = \Gamma' z_{i,t}. \tag{5}$$

The dimension of $z_{i,t}$ determines how flexibly exposures can depend on observed information; the dimension of f_{t+1} determines the number of common return movements in the model. For example, exposure to a financing shock may depend jointly on leverage, profitability, and investment. Characteristic interactions can describe that exposure more accurately without introducing a separate shock for each interaction. A rich instrument set can therefore complement a factor representation of lower dimension.

Timing instruments allow exposures to vary with economic conditions. Scaling a factor return by a lagged predictor changes when its risk is borne and can generate a payoff distinct from holding the factor unconditionally (Haddad et al., 2020; Lehnherr et al., 2025). A small number of conditional shocks therefore need not imply an unconditional SDF with the same number of constant coefficients. Whether fewer IPCA factors than instruments can effectively represent the additional timing payoffs remains an empirical question.

4.2 Diversification and Expected Returns

Consider a vector R of P excess portfolio returns constructed from the available information, including characteristic and timing portfolios. Let f contain K tradable factor returns, with K small relative to P. Hold the portfolio construction rules fixed and consider population moments. Assume finite second moments and a nonsingular covariance matrix for f. Write

$$R = Bf + e, \qquad B = \mathrm{Cov}(R, f)\,\mathrm{Var}(f)^{-1}, \qquad \alpha = \mathbb{E}[e], \quad \Omega = \mathrm{Var}(e). \tag{6}$$

The residual e has zero covariance with f, but its mean is unrestricted. Thus, α measures expected returns not explained by common factor exposures.

The diversification condition restricts the covariance of these residual returns. Suppose there is a constant c, independent of P, such that

$$w'\Omega w \le c\,\|w\|^2 \quad \text{for every portfolio weight vector } w. \tag{7}$$

This approximate factor restriction allows residual correlation but bounds its contribution to portfolio risk (Chamberlain and Rothschild, 1983). Under a fixed normalization of the underlying portfolios, weights satisfying $\|w\|^2 = O(P^{-1})$ produce residual variance that approaches zero as the investment set expands.

With frictionless trading in the portfolios and factor hedges, an investor can obtain the excess payoff $u = w'R - w'Bf = w'e$. To limit near-arbitrage opportunities, suppose attainable Sharpe ratios are bounded by h, uniformly as the investment set expands. This is stronger than ruling out riskless profits alone (Cochrane and Saa-Requejo, 2000). Applied to the hedged payoff, the restriction implies

$$|w'\alpha| \le h\sqrt{w'\Omega w} \le h\sqrt{c}\,\|w\|. \tag{8}$$

Equation (8) bounds expected returns by residual risk. As diversification reduces the risk of a factor-hedged position toward zero, its expected return must also become small. Otherwise, the position would offer an increasingly large Sharpe ratio.

The same restriction yields an approximation across the original portfolios. Substituting $w = \alpha$ into equation (8), and treating $\alpha = 0$ separately, gives

$$\frac{1}{P}\sum_{j=1}^{P}\alpha_j^2 \le \frac{h^2 c}{P}. \tag{9}$$

The average squared difference between expected returns and compensation for common factor exposures becomes small as P increases, even if individual pricing errors do not all vanish.

4.3 The Corresponding Factor SDF

Let $\mu_f = \mathbb{E}[f]$ and $\Sigma_f = \mathrm{Var}(f)$. The affine factor SDF with mean one is

$$M^{(K)} = 1 - \mu_f' \Sigma_f^{-1} (f - \mu_f). \tag{10}$$

It prices the factor returns exactly, since $\mathbb{E}[M^{(K)} f] = 0$. For the larger set of portfolios,

$$\mathbb{E}[M^{(K)} R] = \alpha. \tag{11}$$

The factor component Bf is priced exactly, while $\mathrm{Cov}(e, f) = 0$ implies $\mathbb{E}[M^{(K)} e] = \mathbb{E}[e]$. The pricing errors of the factor SDF are thus the expected residual returns bounded in equations (8) and (9). Common factor exposures need not explain every expected return exactly for the SDF to approximate pricing across a large set of portfolios.[4]

For any factor representation with diversifiable residual risk, richer instruments can improve the description of common exposures without requiring a proportionate increase in the number of common factors.

4.4 Implications for the Empirical Results

Many instruments can thus contribute to pricing through a smaller set of factors. Factors with little return variation can also contain useful pricing information (Lettau and Pelger, 2020). To complement the bound on average squared errors in equation (9), I evaluate both Sharpe ratios and pricing errors on a common set of test assets as the factor count changes.

The baseline and interaction specifications outperform their managed portfolio benchmarks with relatively few IPCA factors, whereas the Fourier specification benefits substantially from additional factors. The economic argument explains why approximate pricing by fewer factors is possible when those factors leave diversifiable residual risk and near-arbitrage opportunities are restricted.

[4]This affine SDF is defined from population moments, with positivity unrestricted. Let $s_f^2 = \mu_f' \Sigma_f^{-1} \mu_f$. The population SDF with constant term one used in Section 2, without a penalty, equals $M^{(K)}/(1 + s_f^2)$. This scaling preserves exact zero pricing moments but rescales nonzero errors. Ridge estimation does not satisfy the unpenalized identities exactly. The relation between factor portfolio efficiency and SDF moments follows Hansen and Jagannathan (1991).

5 Robustness Checks

I assess robustness to the training-window length, return frequency, factor weighting scheme, and stock dataset.

5.1 Time Variation in Factors

I examine how the training-window length affects out-of-sample performance. Bessembinder et al. (2024) emphasize time variation in factors and use a 60-month training window. Figure 8 compares IPCA SDFs estimated over 60 and 240 months, with penalties selected separately for each specification by validation.

IPCA outperforms the corresponding Complex SDF with both training windows. With a 60-month window and 60 IPCA factors, the baseline Sharpe ratio is 2.80, compared with 1.77 for the Complex SDF; the interaction Sharpe ratio is 3.33, compared with 1.88. At the same count of 60 IPCA factors, a 240-month window gives a Sharpe ratio of 2.74 for the baseline and 3.81 for interactions. Using all available IPCA factors with the longer window raises these ratios to 3.09 with 149 baseline factors and 4.15 with 240 interaction factors. The benefit of a longer window thus varies with the specification and the number of IPCA factors.

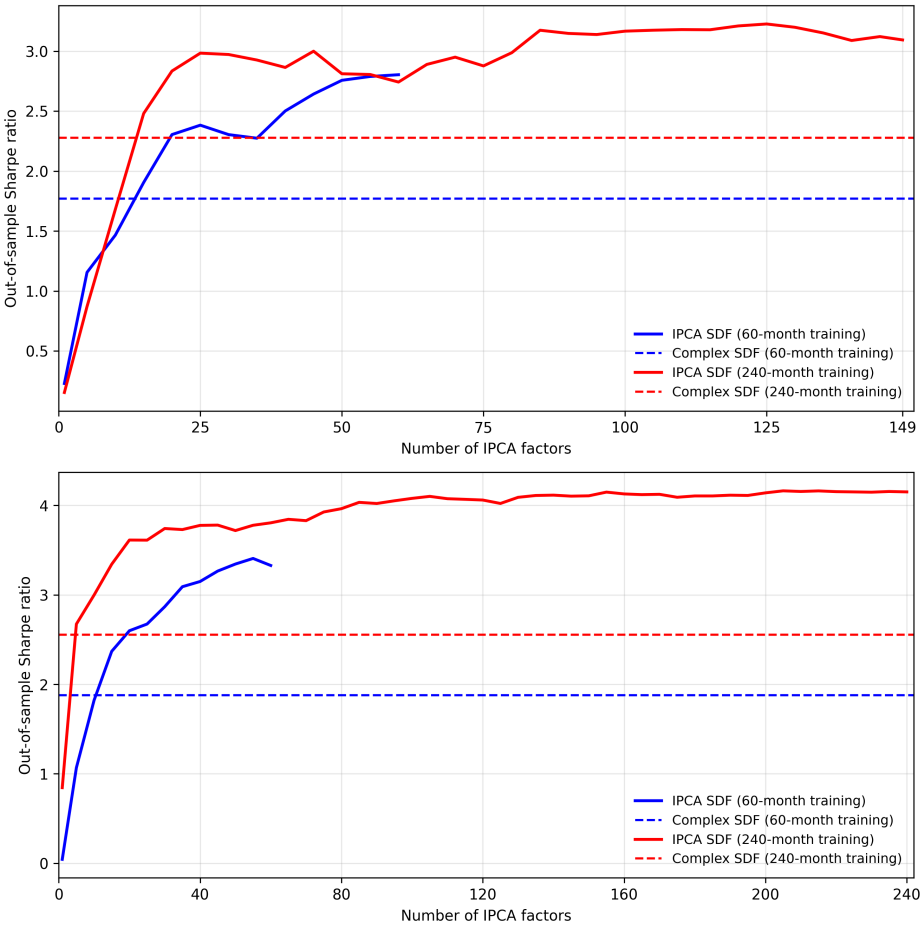


Figure 8: Out-of-Sample Sharpe Ratios across Training Windows. The two panels compare Complex and IPCA SDFs using monthly returns and, respectively, 153 baseline characteristics and 11,934 instruments formed from their linear and degree-2 terms (Jensen et al., 2023). Blue curves use 60-month training windows; red curves use 240-month windows. Solid lines show IPCA SDFs and dashed lines show Complex SDFs. The horizontal axis shows the number of IPCA factors, up to 60 for the short window and 149 baseline or 240 interaction factors for the long window. Evaluation returns span February 2000–January 2024.

5.2 Factor Return Frequency in SDF Estimation

I compare SDFs estimated using daily and monthly returns. Kozak et al. (2020) use daily factor returns to estimate second moments more precisely, whereas Bessembinder et al. (2024) and Didisheim et al. (2024) use monthly returns. The JKP stock-characteristic extract used in the main analysis is monthly; the provider also supplies daily factor returns (Jensen et al., 2026). For a comparison on matched data, I use the Common Task Framework (CTF) data of Hellum et al. (2025), which include daily and monthly returns. I evaluate both SDFs on the same monthly out-of-sample returns so that only the estimation frequency differs.

Figure 9 compares out-of-sample Sharpe ratios for daily estimation (blue lines) and monthly estimation (red lines). Each model estimates both IPCA factors and SDF coefficients at the indicated frequency and selects its own penalties. Both use 240 calendar months of training and the same 287 monthly evaluation returns from February 2000 through December 2023.

With all 153 IPCA factors, daily and monthly estimation yield similar out-of-sample Sharpe ratios: 2.07 and 2.01, respectively. Monthly estimation performs better with very few factors; at five factors, its Sharpe ratio is 1.85, compared with 0.11 for daily estimation. Daily returns therefore add little to the Sharpe ratio when all 153 IPCA factors are used.

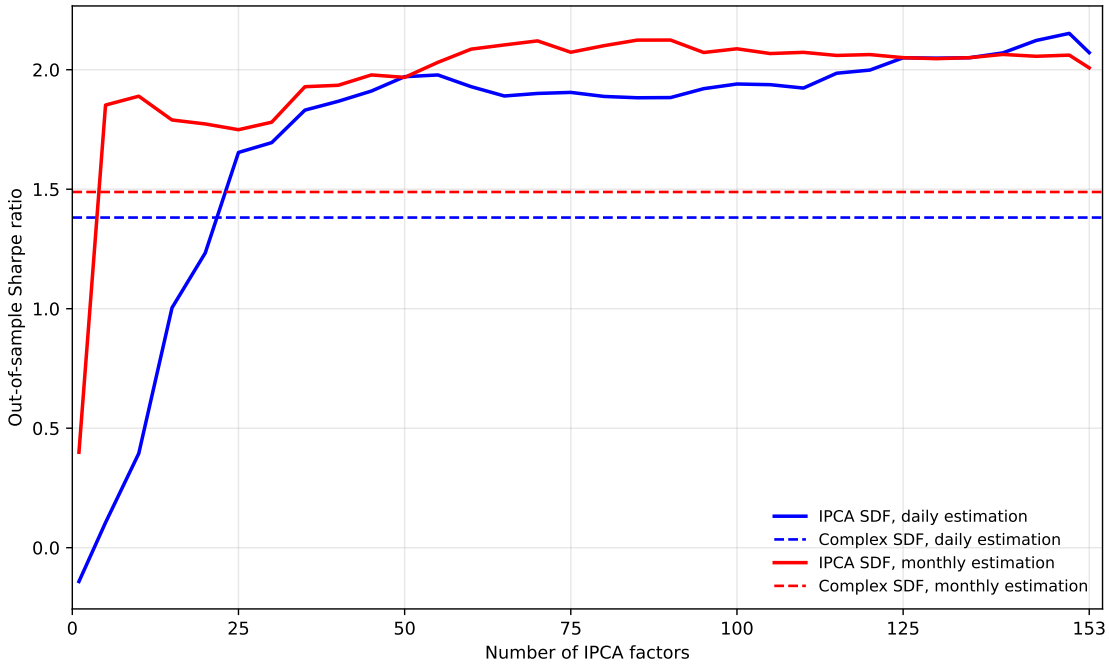


Figure 9: Out-of-Sample Sharpe Ratios by Estimation Frequency. The figure compares Complex and IPCA SDFs using the 153 characteristics of Jensen et al. (2023), measured in the Common Task Framework data of Hellum et al. (2025). IPCA factors and SDF coefficients are estimated from daily returns (blue) or monthly returns (red). Both models use 240 calendar months of training and are evaluated on the same monthly returns from February 2000–December 2023. Solid lines show IPCA SDFs and dashed lines show Complex SDFs. The horizontal axis shows the number of IPCA factors.

5.3 Alternative Factor Weighting Schemes

I examine the sensitivity of SDF performance to alternative factor weighting schemes. Bessembinder et al. (2022) show how weighting methods and the number of quantile portfo-

lios affect out-of-sample factor performance. Kozak et al. (2020) rank characteristics across stocks on $[-0.5, 0.5]$ and divide the demeaned ranks by the sum of their absolute values to form portfolio weights. Didisheim et al. (2024) construct factors from characteristic ranks on the same interval.

The main analysis in Section 3 uses demeaned characteristic instruments for IPCA and univariate weights for the benchmark portfolios. For robustness, I replace these weights with square-root-capitalization or equal-magnitude sign weights. In both alternative IPCA specifications, each instrument equals the number of stocks times its portfolio weight, with no additional centering. The three weighting schemes are:

1. Square-root-capitalization weighting: for each characteristic, I take the sign of its demeaned cross-sectional rank, multiply it by the square root of market capitalization, and normalize the resulting stock weights to unit gross exposure.

2. Equal-magnitude sign weighting: for each characteristic, I take the sign of its demeaned cross-sectional rank and normalize the resulting stock weights to unit gross exposure. Both alternatives use all eligible stocks rather than only the top and bottom deciles.

3. Univariate weighting (Kozak et al., 2018): I rank each characteristic across stocks, demean the ranks, and divide by the sum of their absolute values to form the benchmark portfolio weights.

Figure 10 reports out-of-sample Sharpe ratios under the two alternative weighting schemes. With equal-magnitude sign weights, the IPCA SDF has a Sharpe ratio of 3.08 at 20 factors and 3.14 at 153, compared with 1.95 for the Complex SDF. With square-root-capitalization weights, the IPCA Sharpe ratio is 2.13 at 20 factors and falls to 1.74 at 153, compared with 1.49 for the Complex SDF. IPCA outperforms the corresponding benchmark under both schemes, while the effect of adding IPCA factors differs across schemes.

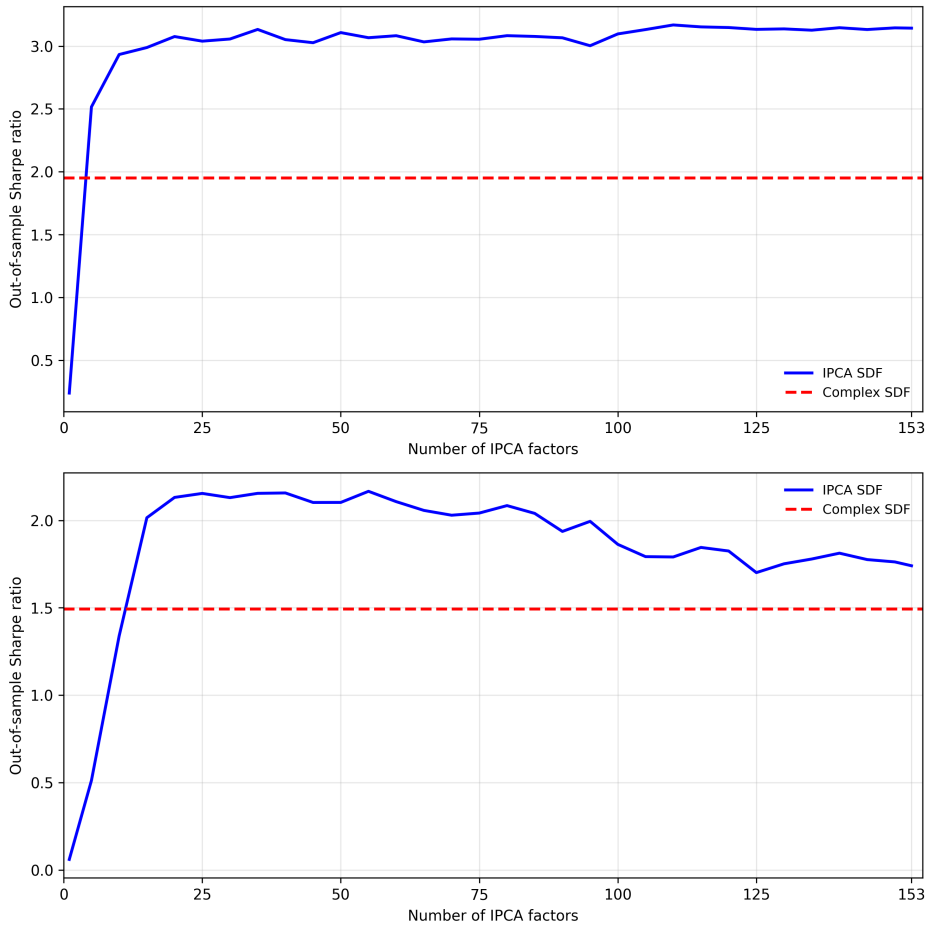


Figure 10: Out-of-Sample Sharpe Ratios across Factor Weighting Schemes. The figure uses 153 characteristics from Jensen et al. (2023). The top panel uses equal-magnitude signs of demeaned ranks; the bottom panel multiplies these signs by square-root market capitalization. Each portfolio has unit gross exposure. Blue solid lines show IPCA SDFs and red dashed lines show Complex SDFs. The horizontal axis shows the number of IPCA factors. Evaluation returns span February 2000–January 2024.

5.4 Common Task Framework Data in Hellum et al. (2025)

I repeat the baseline analysis using the Common Task Framework data of Hellum et al. (2025), which provide a common panel of stock characteristics and returns for comparing asset pricing models.

I use the end-of-month stock characteristics in `ctff_chars.parquet` to match the 153 characteristics in Jensen et al. (2023). I then compare IPCA SDFs with Complex SDFs constructed from the corresponding managed portfolios.

Figure 11 reports out-of-sample Sharpe ratios and pricing errors by IPCA factor count. At 153 factors, the IPCA Sharpe ratio is 2.01 and the squared HJ distance is 0.473, compared

with 1.49 and 0.549 for the Complex SDF. Much of the Sharpe ratio gain appears with the first five IPCA factors, which yield 1.85, although gains at larger counts are not monotone. The evaluation sample is February 2000–December 2023, as in the frequency comparison.

Sharpe ratios are lower with CTF data than with JKP data. Figure 12 in Appendix D shows that the JKP data cover a much larger universe of stocks. Broader stock coverage, along with differences in data construction and fitted IPCA models, may contribute to this performance gap.

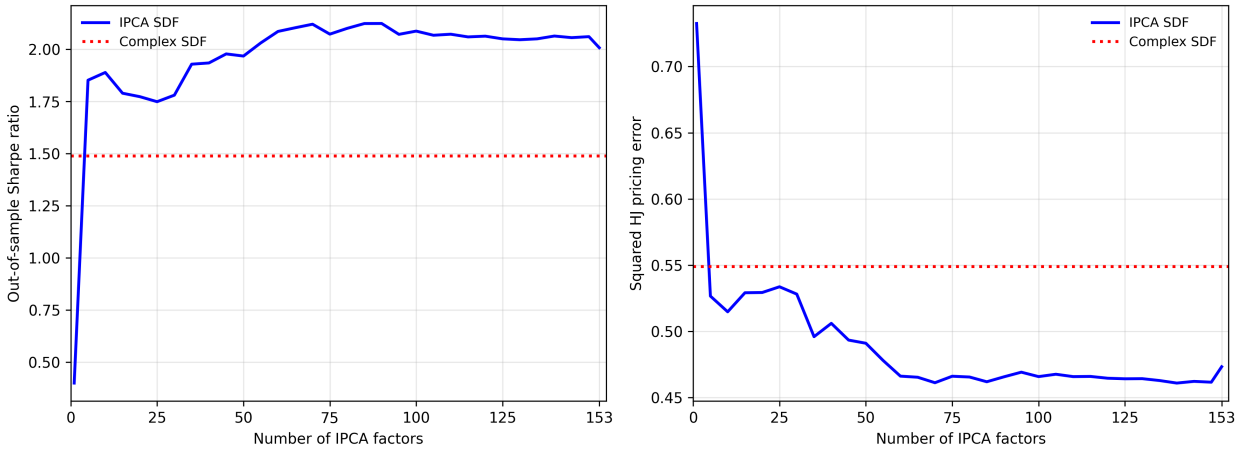


Figure 11: Out-of-Sample Sharpe Ratios and Pricing Errors by Number of IPCA Factors: Common Task Framework Data. The models use 153 characteristics. Blue solid lines show IPCA SDFs and red dotted lines show Complex SDFs. The left panel reports out-of-sample Sharpe ratios and the right panel reports squared Hansen–Jagannathan distances, both plotted against the number of IPCA factors. Evaluation returns span February 2000–December 2023. Penalties are selected using fixed 1990–1999 validation.

6 Conclusion

This paper examines whether a large set of return predictors requires an SDF with an equally large number of factors. I construct SDFs from IPCA factors extracted from stock returns using firm characteristics, their nonlinear transformations, and factor-timing instruments. Characteristic interactions and random Fourier features improve out-of-sample pricing performance, while the factor counts associated with these gains depend on the specification.

With twenty IPCA factors, the baseline and interaction specifications have annualized out-of-sample Sharpe ratios of 2.83 and 3.61, respectively. Both outperform their corresponding managed portfolio benchmarks. Random Fourier features provide further gains at higher

factor counts. Timing gains are less consistent: some individual predictors improve SDF performance, but combining many timing signals does not improve on baseline IPCA.

A rich instrument set can help measure exposures to a smaller set of common risks. Under the paper's common factor and diversification conditions, a uniform bound on attainable Sharpe ratios implies that average squared pricing errors converge to zero as the investment set expands. These conditions provide an economic rationale for approximate pricing by a factor SDF whose dimension need not increase proportionately with the number of instruments.

References

Hendrik Bessembinder, Aaron Burt, and Christopher M. Hrdlicka. Factor returns and out-of-sample alphas: Factor construction matters. *SSRN Electronic Journal*, 2022. ISSN 1556-5068. doi: https://dx.doi.org/10.2139/ssrn.4281769. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4281769.

Hendrik Bessembinder, Aaron Burt, and Christopher M. Hrdlicka. More factors matter and factors matter more than you might think: The role of time variation in factor premia. *SSRN Electronic Journal*, 2024. ISSN 1556-5068. doi: https://dx.doi.org/10.2139/ssrn.3992041. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3992041.

Mark Britten-Jones. The sampling error in estimates of mean-variance efficient portfolio weights. *Journal of Finance*, 54(2):655–671, 1999. URL https://EconPapers.repec.org/RePEc:bla:jfinan:v:54:y:1999:i:2:p:655-671.

Svetlana Bryzgalova, Markus Pelger, and Jason Zhu. Forest through the trees: Building cross-sections of stock returns. *The Journal of Finance*, 80(5):2447–2506, 2025. doi: 10.1111/jofi.13477. URL https://doi.org/10.1111/jofi.13477.

Daniel Buncic. Simplified: A closer look at the virtue of complexity in return prediction, April 2025. Available at SSRN: https://ssrn.com/abstract=5239006. Revised March 2026.

Gary Chamberlain and Michael Rothschild. Arbitrage, factor structure, and mean-variance analysis on large asset markets. *Econometrica*, 51(5):1281–1304, 1983. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/1912275.

John H. Cochrane. *Asset Pricing: Revised Edition*. Princeton University Press, Princeton, NJ, revised edition, 2009.

John H. Cochrane and Jesus Saa-Requejo. Beyond arbitrage: Good-deal asset price bounds in incomplete markets. *Journal of Political Economy*, 108(1):79–119, 2000. ISSN 00223808, 1537534X. URL http://www.jstor.org/stable/10.1086/262112.

Alexander Dickerson, Christian Julliard, and Philippe Mueller. The co-pricing factor zoo. *Journal of Financial Economics*, 182:104295, 2026. doi: 10.1016/j.jfineco.2026.104295.

Antoine Didisheim, Shikun Ke, Bryan T. Kelly, and Semyon Malamud. Complexity in Factor Pricing Models. Swiss Finance Institute Research Paper Series 23-19, Swiss Finance Institute, January 2024. URL https://ideas.repec.org/p/chf/rpseri/rp2319.html. Revised January 15, 2024.

Eugene Fama and Robert R Bliss. The information in long-maturity forward rates. *American Economic Review*, 77(4):680–92, 1987. URL https://EconPapers.repec.org/RePEc:aea:aecrev:v:77:y:1987:i:4:p:680-92.

Eugene F. Fama and Kenneth R. French. Dividend yields and expected stock returns. *Journal of Financial Economics*, 22(1):3–25, 1988. ISSN 0304-405X. doi: https://doi.org/10.1016/0304-405X(88)90020-7. URL https://www.sciencedirect.com/science/article/pii/

0304405X88900207.

Eugene F. Fama and Kenneth R. French. The cross-section of expected stock returns. *The Journal of Finance*, 47(2):427–465, 1992. ISSN 00221082, 15406261. URL http://www.jstor.org/stable/2329112.

Eugene F. Fama and Kenneth R. French. A five-factor asset pricing model. *Journal of Financial Economics*, 116(1):1–22, April 2015. ISSN 0304405X. doi: 10.1016/j.jfineco.2014.10.010. URL https://linkinghub.elsevier.com/retrieve/pii/S0304405X14002323.

Jonathan S. Fan, Bryan T. Kelly, Semyon Malamud, and Yuan Zhang. Recursive portfolio machines, April 2026. Available at SSRN: https://ssrn.com/abstract=6615198.

Amit Goyal, Ivo Welch, and Athanasse Zafirov. A comprehensive 2022 look at the empirical performance of equity premium prediction. *The Review of Financial Studies*, 37(11):3490–3557, 09 2024. ISSN 0893-9454. doi: 10.1093/rfs/hhae044. URL https://doi.org/10.1093/rfs/hhae044.

Hongye Guo and Jessica A. Wachter. Correlation neglect in asset prices. *Working paper, University of Hong Kong and The Wharton School (NBER)*, February 2026. Draft dated February 8, 2026.

Valentin Haddad, Serhiy Kozak, and Shrihari Santosh. Factor Timing. *The Review of Financial Studies*, 33(5):1980–2018, May 2020. ISSN 0893-9454, 1465-7368. doi: 10.1093/rfs/hhaa017. URL https://academic.oup.com/rfs/article/33/5/1980/5753962.

Lars Peter Hansen and Ravi Jagannathan. Implications of security market data for models of dynamic economies. *Journal of Political Economy*, 99(2):225–262, 1991. ISSN 00223808, 1537534X. URL http://www.jstor.org/stable/2937680.

Lars Peter Hansen and Ravi Jagannathan. Assessing specification errors in stochastic discount factor models. *The Journal of Finance*, 52(2):557–590, 1997. doi: 10.1111/j.1540-6261.1997.tb04813.x. URL https://doi.org/10.1111/j.1540-6261.1997.tb04813.x.

Oliver Hellum, Theis Ingerslev Jensen, Bryan T. Kelly, and Lasse Heje Pedersen. The power of the common task framework, August 2025. Available at SSRN: https://ssrn.com/abstract=5242901 or http://dx.doi.org/10.2139/ssrn.5242901.

Oliver Hellum, Theis Ingerslev Jensen, Bryan T. Kelly, and Semyon Malamud. Complex modern portfolio theory. Working Paper 35246, National Bureau of Economic Research, May 2026.

Kewei Hou, Chen Xue, and Lu Zhang. Digesting Anomalies: An Investment Approach. *Review of Financial Studies*, 28(3):650–705, March 2015. ISSN 0893-9454, 1465-7368. doi: 10.1093/rfs/hhu068. URL https://academic.oup.com/rfs/article-lookup/doi/10.1093/rfs/hhu068.

Theis Ingerslev Jensen, Bryan T. Kelly, and Lasse Heje Pedersen. Is there a replication crisis in finance? *The Journal of Finance*, 78(5):2465–2518, 2023. doi: 10.1111/jofi.13249. URL https://doi.org/10.1111/jofi.13249.

Theis Ingerslev Jensen, Bryan T. Kelly, and Lasse Heje Pedersen. *Global Factor Data Documentation*, 2026. URL https://jkpfactors-data.s3.amazonaws.com/documents/Documentation.pdf. Accessed October 19, 2026.

Bryan T. Kelly and Semyon Malamud. Understanding the virtue of complexity, 2025. Swiss Finance Institute Research Paper No. 25-96. Available at SSRN: https://ssrn.com/abstract=5346842.

Bryan T Kelly and Dacheng Xiu. Financial machine learning. Working Paper 31502, National Bureau of Economic Research, July 2023. URL http://www.nber.org/papers/w31502.

Bryan T. Kelly, Seth Pruitt, and Yinan Su. Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics*, 134(3):501–524, December 2019. ISSN 0304405X. doi: 10.1016/j.jfineco.2019.05.001. URL https://linkinghub.elsevier.com/retrieve/pii/S0304405X19301151.

Bryan T. Kelly, Boris Kuznetsov, Semyon Malamud, and Teng Andrea Xu. Artificial intelligence asset pricing models. Working Paper 33351, National Bureau of Economic Research, January 2025.

Serhiy Kozak. Kernel trick for the cross-section. *SSRN Electronic Journal*, 2021. ISSN 1556-5068. doi: https://dx.doi.org/10.2139/ssrn.3307895. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3307895#.

Serhiy Kozak, Stefan Nagel, and Shrihari Santosh. Interpreting Factor Models. *The Journal of Finance*, 73(3):1183–1223, June 2018. ISSN 0022-1082, 1540-6261. doi: 10.1111/jofi.12612. URL https://onlinelibrary.wiley.com/doi/10.1111/jofi.12612.

Serhiy Kozak, Stefan Nagel, and Shrihari Santosh. Shrinking the cross-section. *Journal of Financial Economics*, 135(2):271–292, 2020. ISSN 0304-405X. doi: https://doi.org/10.1016/j.jfineco.2019.06.008. URL https://www.sciencedirect.com/science/article/pii/S0304405X19301655.

Rob Lehnherr, Manan Mehta, and Stefan Nagel. Optimal factor timing in a high-dimensional setting. *Financial Analysts Journal*, 81(2):51–66, 2025. doi: 10.1080/0015198X.2025.2474385.

Martin Lettau and Markus Pelger. Estimating latent asset-pricing factors. *Journal of Econometrics*, 218(1):1–31, 2020. doi: 10.1016/j.jeconom.2019.08.012.

Stefan Nagel. *Machine Learning in Asset Pricing*, volume 8. Princeton University Press, 2021. ISBN 9780691218700. URL http://www.jstor.org/stable/j.ctv19fvx8r.

Stefan Nagel. Seemingly virtuous complexity in return prediction. Working Paper 34104, National Bureau of Economic Research, August 2025.

Stephen A Ross. The arbitrage theory of capital asset pricing. *Journal of Economic Theory*, 13(3):341–360, 1976. ISSN 0022-0531. doi: https://doi.org/10.1016/0022-0531(76)90046-6. URL https://www.sciencedirect.com/science/article/pii/0022053176900466.

Robert J Shiller. Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends? *American Economic Review*, 71(3):421–436, June 1981. URL https://ideas.repec.org/a/aea/aecrev/v71y1981i3p421-36.html.

Ivo Welch and Amit Goyal. A Comprehensive Look at The Empirical Performance of Equity Premium Prediction. *Review of Financial Studies*, 21(4):1455–1508, July 2008. ISSN 0893-9454, 1465-7368. doi: 10.1093/rfs/hhm014. URL https://academic.oup.com/rfs/article-lookup/doi/10.1093/rfs/hhm014.

A Variable Definitions of Firm Characteristics

This table lists all 153 model inputs, in their fixed input order, with the provider's characteristic names. Each mnemonic is an actual model column; the list contains no additional variables. Detailed accounting and market-variable formulas and original-study citations are given in Sections 6–10 of the JKP documentation (Jensen et al., 2026). All specifications use the screening, missing-rank, and transformation rules in Section 2.3.

No.	Model input	Definition
1	cowc_gr1a	Change in current operating working capital
2	oaccruals_at	Operating accruals
3	oaccruals_ni	Percent operating accruals
4	seas_16_20na	Years 16-20 lagged returns, nonannual
5	taccruals_at	Total accruals
6	taccruals_ni	Percent total accruals
7	capex_abn	Abnormal corporate investment
8	debt_gr3	Growth in book debt (3 years)
9	fnl_gr1a	Change in financial liabilities
10	ncol_gr1a	Change in noncurrent operating liabilities
11	nfna_gr1a	Change in net financial assets
12	ni_ar1	Earnings persistence
13	noa_at	Net operating assets
14	aliq_at	Liquidity of book assets
15	at_gr1	Asset Growth
16	be_gr1a	Change in common equity
17	capx_gr1	CAPEX growth (1 year)
18	capx_gr2	CAPEX growth (2 years)
19	capx_gr3	CAPEX growth (3 years)
20	coa_gr1a	Change in current operating assets
21	col_gr1a	Change in current operating liabilities
22	emp_gr1	Hiring rate
23	inv_gr1	Inventory growth
24	inv_gr1a	Inventory change
25	lnoa_gr1a	Change in long-term net operating assets
26	mispricing_mgmt	Mispricing factor: Management
27	ncoa_gr1a	Change in noncurrent operating assets
28	nncoa_gr1a	Change in net noncurrent operating assets
29	noa_gr1a	Change in net operating assets
30	ppeinv_gr1a	Change PPE and Inventory
31	ret_60_12	Long-term reversal
32	sale_gr1	Sales Growth (1 year)
33	sale_gr3	Sales Growth (3 years)
34	saleq_gr1	Sales growth (1 quarter)

Continued on next page

Firm characteristics (continued)

No.	Model input	Definition
35	seas_2_5na	Years 2-5 lagged returns, nonannual
36	age	Firm age
37	aliq_mat	Liquidity of market assets
38	at_be	Book leverage
39	bidaskhl_21d	The high-low bid-ask spread
40	cash_at	Cash-to-assets
41	netdebt_me	Net debt-to-price
42	ni_ivol	Earnings volatility
43	rd_sale	R&D-to-sales
44	rd5_at	R&D capital-to-book assets
45	tangibility	Asset tangibility
46	z_score	Altman Z-score
47	beta_60m	Market Beta
48	beta_dimson_21d	Dimson beta
49	betabab_1260d	Frazzini-Pedersen market beta
50	betadown_252d	Downside beta
51	earnings_variability	Earnings variability
52	ivol_capm_21d	Idiosyncratic volatility from the CAPM (21 days)
53	ivol_capm_252d	Idiosyncratic volatility from the CAPM (252 days)
54	ivol_ff3_21d	Idiosyncratic volatility from the Fama–French three-factor model
55	ivol_hxz4_21d	Idiosyncratic volatility from the q-factor model
56	ocfq_saleq_std	Cash flow volatility
57	rmax1_21d	Maximum daily return
58	rmax5_21d	Highest 5 days of return
59	rvol_21d	Return volatility
60	seas_6_10na	Years 6-10 lagged returns, nonannual
61	turnover_126d	Share turnover
62	zero_trades_21d	Number of zero trades with turnover as tiebreaker (1 month)
63	zero_trades_126d	Number of zero trades with turnover as tiebreaker (6 months)
64	zero_trades_252d	Number of zero trades with turnover as tiebreaker (12 months)
65	prc_highprc_252d	Current price to high price over last year
66	resff3_6_1	Residual momentum t-6 to t-1
67	resff3_12_1	Residual momentum t-12 to t-1
68	ret_3_1	Price momentum t-3 to t-1
69	ret_6_1	Price momentum t-6 to t-1
70	ret_9_1	Price momentum t-9 to t-1
71	ret_12_1	Price momentum t-12 to t-1
72	seas_1_1na	Year 1-lagged return, nonannual
73	dsale_dinv	Change sales minus change Inventory
74	dsale_drec	Change sales minus change receivables
75	dsale_dsga	Change sales minus change SG&A
76	niq_at_chg1	Change in quarterly return on assets
77	niq_be_chg1	Change in quarterly return on equity
78	niq_su	Standardized earnings surprise

Continued on next page

Firm characteristics (continued)

No.	Model input	Definition
79	ocf_at_chg1	Change in operating cash flow to assets
80	ret_12_7	Price momentum t-12 to t-7
81	sale_emp_gr1	Labor force efficiency
82	saleq_su	Standardized Revenue surprise
83	seas_1_1an	Year 1-lagged return, annual
84	tax_gr1a	Tax expense surprise
85	dolvol_var_126d	Coefficient of variation for dollar trading volume
86	ebit_bev	Return on net operating assets
87	ebit_sale	Profit margin
88	f_score	Piotroski F-score
89	ni_be	Return on equity
90	niq_be	Quarterly return on equity
91	o_score	Ohlson O-score
92	ocf_at	Operating cash flow to assets
93	ope_be	Operating profits-to-book equity
94	ope_bel1	Operating profits-to-lagged book equity
95	turnover_var_126d	Coefficient of variation for share turnover
96	at_turnover	Capital turnover
97	cop_at	Cash-based operating profits-to-book assets
98	cop_atl1	Cash-based operating profits-to-lagged book assets
99	dgp_dsale	Change gross margin minus change sales
100	gp_at	Gross profits-to-assets
101	gp_atl1	Gross profits-to-lagged assets
102	mispricing_perf	Mispricing factor: Performance
103	ni_inc8q	Number of consecutive quarters with earnings increases
104	niq_at	Quarterly return on assets
105	op_at	Operating profits-to-book assets
106	op_atl1	Operating profits-to-lagged book assets
107	opex_at	Operating leverage
108	qmj	Quality minus Junk: Composite
109	qmj_growth	Quality minus Junk: Growth
110	qmj_prof	Quality minus Junk: Profitability
111	qmj_safety	Quality minus Junk: Safety
112	sale_bev	Asset turnover
113	corr_1260d	Market correlation
114	coskew_21d	Coskewness
115	dbnetis_at	Net debt issuance
116	kz_index	Kaplan-Zingales index
117	lti_gr1a	Change in long-term investments
118	pi_nix	Taxable income-to-book income
119	seas_2_5an	Years 2-5 lagged returns, annual
120	seas_6_10an	Years 6-10 lagged returns, annual
121	seas_11_15an	Years 11-15 lagged returns, annual
122	seas_11_15na	Years 11-15 lagged returns, nonannual

Continued on next page

Firm characteristics (continued)

No.	Model input	Definition
123	seas_16_20an	Years 16-20 lagged returns, annual
124	sti_gr1a	Change in short-term investments
125	ami_126d	Amihud Measure
126	dolvol_126d	Dollar trading volume
127	market_equity	Market Equity
128	prc	Price per share
129	rd_me	R&D-to-market
130	iskew_capm_21d	Idiosyncratic skewness from the CAPM
131	iskew_ff3_21d	Idiosyncratic skewness from the Fama–French three-factor model
132	iskew_hxz4_21d	Idiosyncratic skewness from the q-factor model
133	ret_1_0	Short-term reversal
134	rmax5_rvol_21d	Highest 5 days of return scaled by volatility
135	rskew_21d	Total skewness
136	at_me	Assets-to-market
137	be_me	Book-to-market equity
138	bev_mev	Book-to-market enterprise value
139	chcsho_12m	Net stock issues
140	debt_me	Debt-to-market
141	div12m_me	Dividend yield
142	ebitda_mev	Ebitda-to-market enterprise value
143	eq_dur	Equity duration
144	eqnetis_at	Net equity issuance
145	eqnpo_12m	Equity net payout
146	eqnpo_me	Net payout yield
147	eqpo_me	Payout yield
148	fcf_me	Free cash flow-to-price
149	ival_me	Intrinsic value-to-market
150	netis_at	Net total issuance
151	ni_me	Earnings-to-price
152	ocf_me	Operating cash flow-to-market
153	sale_me	Sales-to-market

B Kernel Approximation via Random Fourier Features

Random Fourier features approximate a kernel with a finite collection of nonlinear functions of the characteristics, following Kozak (2021) and Didisheim et al. (2024). Their number need not measure how much distinct information the estimated model uses. In a market return forecasting setting, Nagel (2025) shows how a model with many random features can behave like a weighted average of past returns. Buncic (2025) emphasizes the role of the intercept restriction and the distinction between averaging strategy returns and averaging performance statistics. Kelly and Malamud (2025) distinguish the number of model parameters from

effective complexity and analyze the benefits of averaging across models. Their forecasting designs differ from the SDF construction here. I average portfolio returns across draws before calculating performance and report the IPCA factor count per draw. Regressing one on portfolio returns estimates an SDF; it is distinct from forecasting stock returns with an intercept restricted to zero.

The approximation follows Bochner's representation of a shift-invariant kernel as the Fourier transform of a probability distribution. The RBF kernel is

$$k(x, y) = \exp\left(-\frac{\gamma^2}{2}\|x - y\|^2\right),$$

whose corresponding Fourier distribution is Gaussian. To construct a finite-dimensional feature map, I sample independent random vectors $\omega_j \sim \mathcal{N}(0, I_{153})$ and apply the frequency scale γ. The resulting transformation maps the input characteristics x into a higher-dimensional space:

$$z(x) = \frac{1}{\sqrt{Q}}\left[\sin(\gamma\omega_1' x), \ldots, \sin(\gamma\omega_Q' x), \cos(\gamma\omega_1' x), \ldots, \cos(\gamma\omega_Q' x)\right]$$

The kernel is then approximated by

$$k(x, y) \approx z(x)^\top z(y)$$

Here Q is the number of projection vectors, so the feature count is $L = 2Q$. The empirical instruments omit the common $1/\sqrt{Q}$ kernel normalization and demean the sine and cosine features across stocks. The Complex SDF additionally normalizes each feature to unit gross exposure.

Candidate maps and random draws. The map search uses $Q \in \{250, 500, 1000, 2000\}$, corresponding to $L \in \{500, 1000, 2000, 4000\}$, and $\gamma \in \{0.1, 0.25, 0.5, 1.0\}$. The ten equally weighted draws use seeds

$$42,\ 1042,\ 2042,\ 3042,\ 4042,\ 5042,\ 6042,\ 7042,\ 8042,\ 9042.$$

Directions are shared across bandwidths and nested across projection counts within each seed. NumPy's default random generator draws an initial 153×500 standard-normal block. After 500 categorical draws from $\{0.5, 0.6, 0.7, 0.8, 0.9, 1.0\}$, additional 153×500 normal blocks are appended as needed. The categorical draws fix the random stream independently of γ. Each map uses the first Q columns of the resulting direction matrix and ranks centered at 0.5.

Auxiliary managed-return SDF search. I select the feature map using managed portfolio returns before estimating IPCA. For each map and draw, I compute the training mean $\hat{\mu}_F$ and uncentered second moment $\hat{S}_F$, then construct an orthonormal basis for

$$\mathcal{K}_q(\hat{S}_F, \hat{\mu}_F) = \text{span}\{\hat{\mu}_F, \hat{S}_F\hat{\mu}_F, \ldots, \hat{S}_F^{q-1}\hat{\mu}_F\}, \quad q \in \{1, 2, 5, 10, 20, 40, 80, 120\}.$$

I regress one on the returns projected into this basis, without an intercept. The sum-of-squares penalty is $gL^{-1}\sum_t \|F_t\|^2$, using the energy of all L managed returns; g ranges over 50 log-spaced values from 10^{-4} to 10^2. The auxiliary count q is a Krylov dimension, distinct from an IPCA factor count or a variance-PC count.

For every (Q, γ, q, g) candidate, I fit each draw separately in the ten annual 1990–1999 validation folds, with 240 preceding training months per fold. I average the ten SDF portfolio returns within each validation month, pool the resulting 120 monthly returns, and maximize their annualized Sharpe ratio. Training spectral calculations retain eigenvalues above 10^{-13} times the largest training eigenvalue; candidates must be estimable in every fold and draw. This gives $4 \times 4 \times 8 \times 50 = 6{,}400$ eligible candidates. Equal scores favor fewer projections, lower γ, smaller q, and then larger g. The selected tuple is $Q = 500$, $\gamma = 0.1$, $q = 120$, and $g = 10^{-4}$. The search selects the feature map for pricing performance using managed returns, rather than the IPCA stock-fitting objective, before comparing estimators on that map. The Fourier Complex benchmark then uses all 1,000 selected features per draw with this fixed g, rather than only the auxiliary Krylov coordinates.

Subsequent IPCA estimation. I hold the selected map and all ten seeds fixed and fit IPCA separately at $K_{\max} = 240$ for each draw and training window, with stock penalty $\tau = 0$. For each IPCA factor count K, I average the ten validation SDF portfolio returns for each of the 61 α candidates in Section 2.4, then select α_K by pooled validation Sharpe ratio. The validation-selected number of IPCA factors is $K = 210$ per draw. Both α_K and this count are fixed before evaluation. The feature count 1,000, auxiliary dimension 120, and IPCA factor count 210 refer to distinct steps; none is the dimension of the combined 10-draw portfolio.

C Factor-Specific Timing Predictors

Let $u_{ij,t}$ be stock i's percentile rank for characteristic j, using the stock screen and median-rank imputation in Section 2.3. The unit-gross portfolio weights and realized return are

$$w_{ij,t} = \frac{u_{ij,t} - \bar{u}_{j,t}}{\sum_i |u_{ij,t} - \bar{u}_{j,t}|}, \qquad F_{j,t+1} = \sum_i w_{ij,t} r_{i,t+1}.$$

If the denominator is zero, all weights for that characteristic-month are zero. For a spread characteristic p, use its raw characteristic level $x^p_{i,t}$, fill missing levels with the cross-sectional median (zero if the entire cross-section is missing), and winsorize the filled levels at their cross-sectional 1st and 99th percentiles. Denote the result by $\tilde{x}^p_{i,t}$. The levels are the provider's month-end characteristics with their accounting availability lags; no additional monthly shift is applied. The implemented spread is

$$s_{p,j,t} = \sum_i w_{ij,t} \tilde{x}^p_{i,t}, \qquad p \in \{\texttt{be_me}, \texttt{ope_be}, \texttt{at_gr1}\}. \tag{12}$$

These are book equity divided by market equity (value), operating profits divided by book equity (profitability), and 1-year asset growth (growth). The weights depend on the characteristic j defining each base portfolio; the levels depend on p. When both legs exist, each leg has absolute weight one half, so this unit-gross spread is half the difference between separately normalized long- and short-leg means. No additional market-capitalization weighting is applied.

For each portfolio and predictor, I standardize the raw signal $h_{j,t}$ using its own trailing history, including the value known at formation month t:

$$x_{j,t} = \frac{h_{j,t} - \bar{h}^{(240)}_{j,t}}{s^{(240)}_{h,j,t}}.$$

I use up to 240 observations, require at least 24, and compute the sample standard deviation. For spread timing, $h_{j,t} = s_{p,j,t}$. The other three raw signals are

$$h^{(3)}_{j,t} = \sum_{\ell=0}^{2} F_{j,t-\ell}, \qquad h^{(12)}_{j,t} = \sum_{\ell=0}^{11} F_{j,t-\ell}, \qquad h^{(v)}_{j,t} = \mathrm{sd}(F_{j,t}, \ldots, F_{j,t-11}).$$

Own-return signals use sums of monthly returns, and volatility uses twelve monthly returns without annualization. The timing payoff is $x_{j,t} F_{j,t+1}$; stock instruments pair each demeaned baseline characteristic j with its own $x_{j,t}$. All signal information is measured by formation

month t, before the return in $t+1$. Predictor histories start in January 1950 and use the same screened stock data as the main factors.

The only neutral startup replacements concern `eqnetis_at` (net equity issuance), `eqnpo_me` (net payout yield), `eqpo_me` (payout yield), and `netis_at` (net total issuance). Their three spread predictors are undefined in formation months December 1969–September 1971: 88 cells per spread. Their own-return and volatility predictors are undefined through October 1971: 92 cells per predictor. Only these undefined standardized signal cells are set to zero, so the associated timing interaction is neutral. Observed signals are retained, and any other undefined cell causes the construction to stop. These are formation-month dates; their payoffs occur one month later.

D Additional Figures

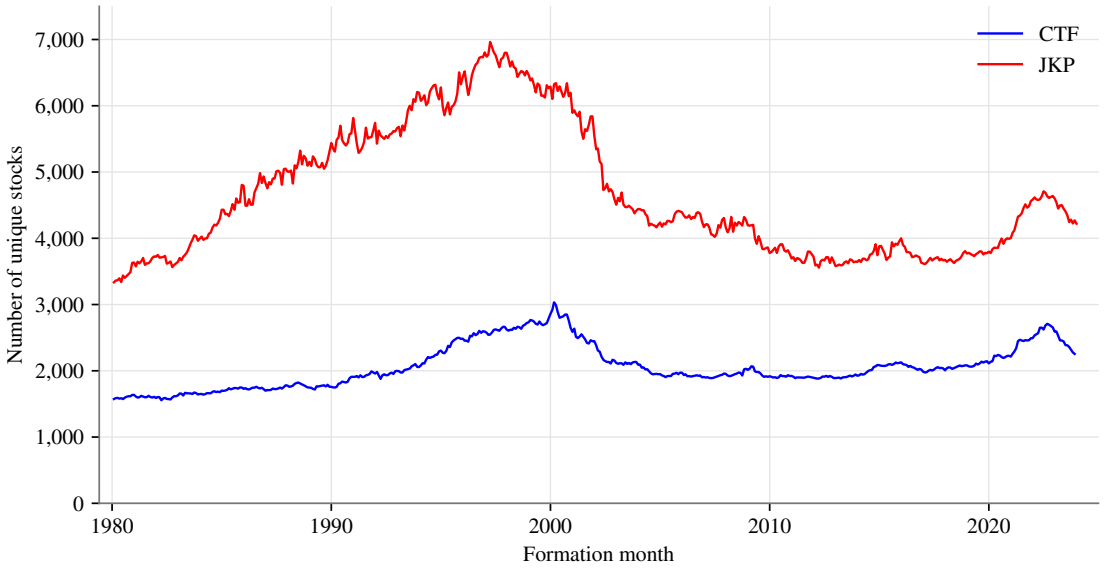


Figure 12: Stock Coverage in the JKP and Common Task Framework (CTF) Estimation Samples. Counts use the common JKP stock identifier `id` at formation month t, retaining stocks with observed excess returns in month $t+1$. The JKP sample retains the mega, large, small, and micro size groups; the CTF sample uses the supplied estimation universe without an additional size screen. Missing characteristics do not exclude a stock. The plotted formation months run from January 1980 through December 2023 for JKP and November 2023 for CTF, corresponding to final returns in January 2024 and December 2023, respectively. The JKP counts match the saved monthly IPCA sample sizes. CTF is introduced by Hellum et al. (2025); JKP Global Factor Data are described by Jensen et al. (2023).

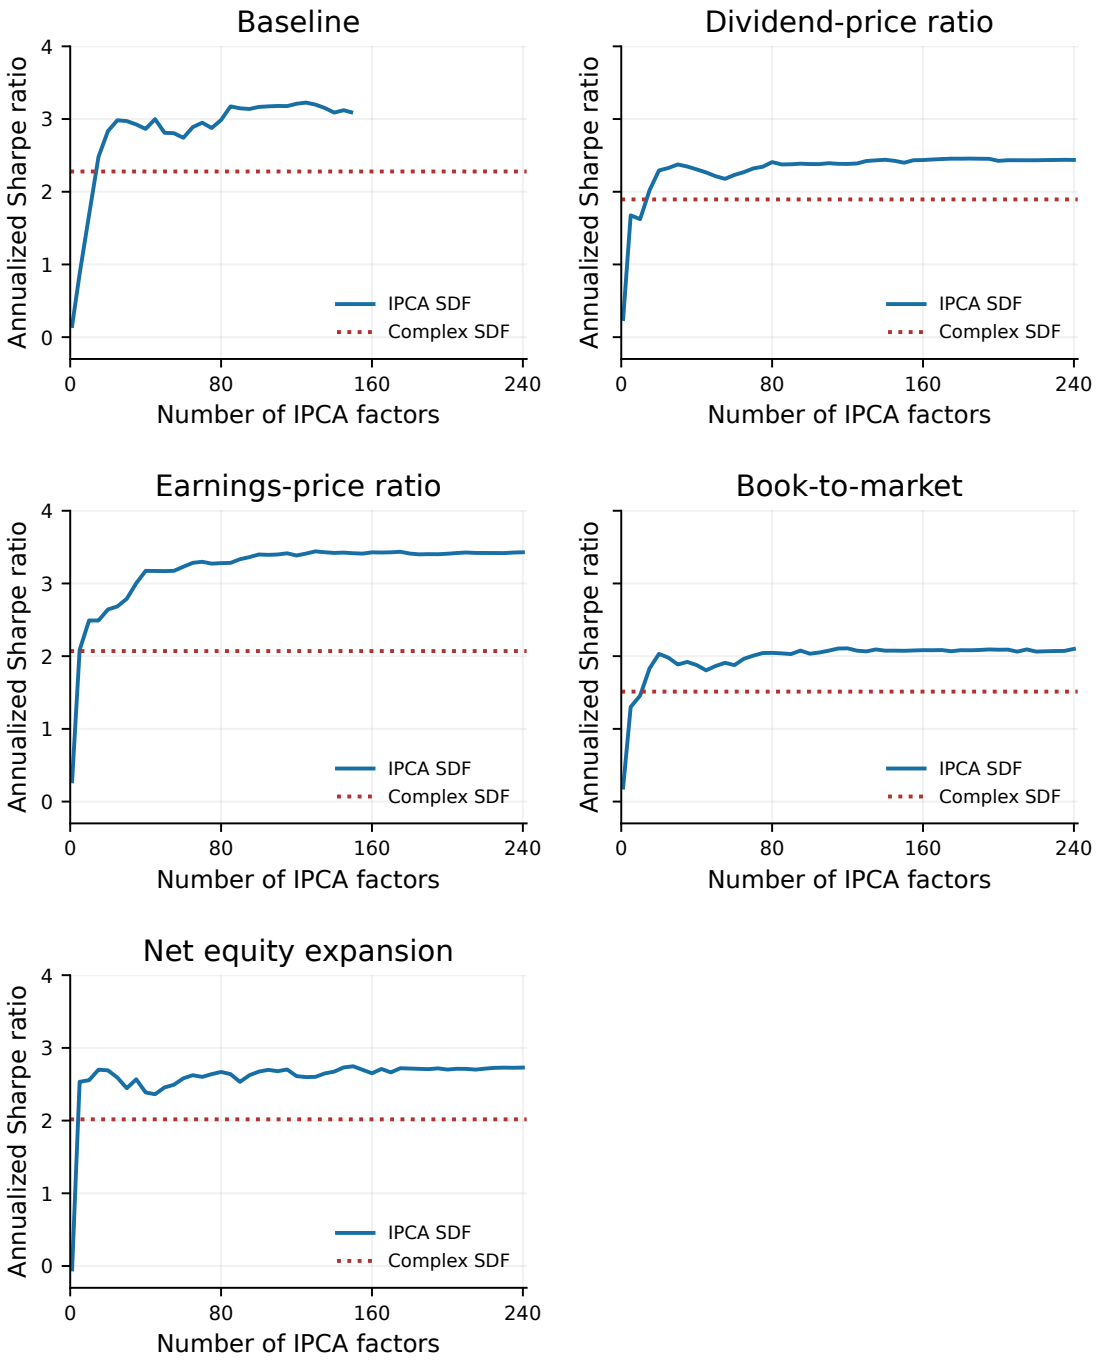


Figure 13: Individual Timing Comparisons: Baseline, dp, ep, bm, and ntis. Each panel compares the IPCA SDF (blue solid) with its all-managed-return Complex SDF benchmark (red dotted). All panels use common axes. The baseline cap is 149; each timing cap is 240. Predictors are defined in Section 2.3. Penalties use fixed 1990–1999 validation, and evaluation covers February 2000–January 2024.

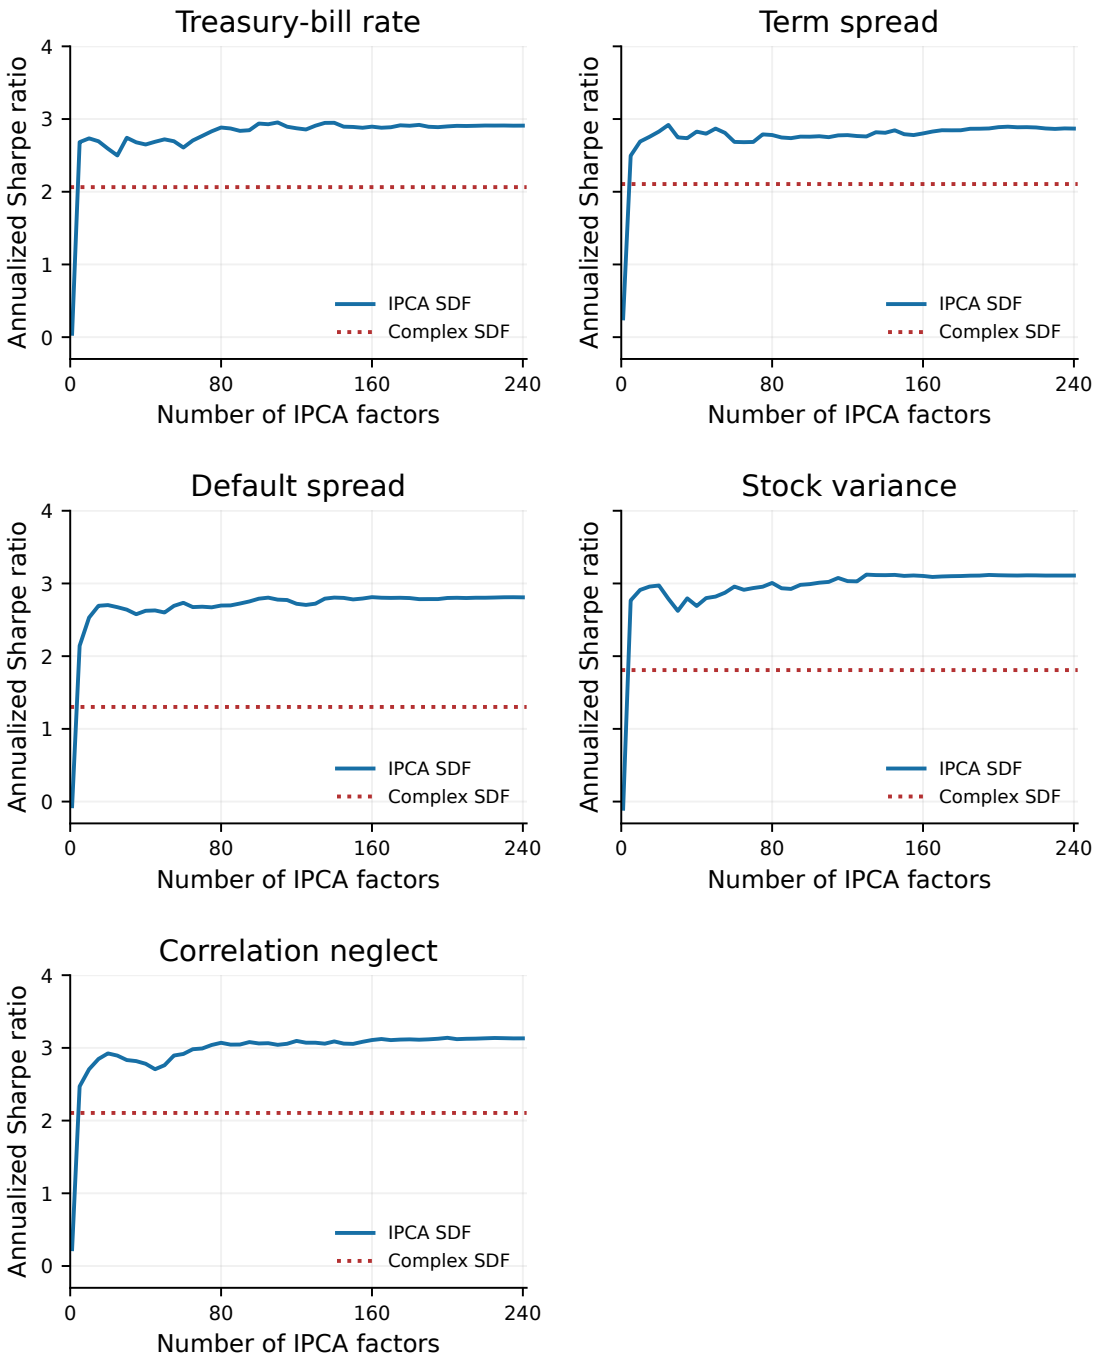


Figure 14: Individual Timing Comparisons: tbl, tms, dfy, svar, and cn. Each panel compares the IPCA SDF (blue solid) with its all-managed-return Complex SDF benchmark (red dotted). All panels use common axes. The baseline cap is 149; each timing cap is 240. Predictors are defined in Section 2.3. Penalties use fixed 1990–1999 validation, and evaluation covers February 2000–January 2024.

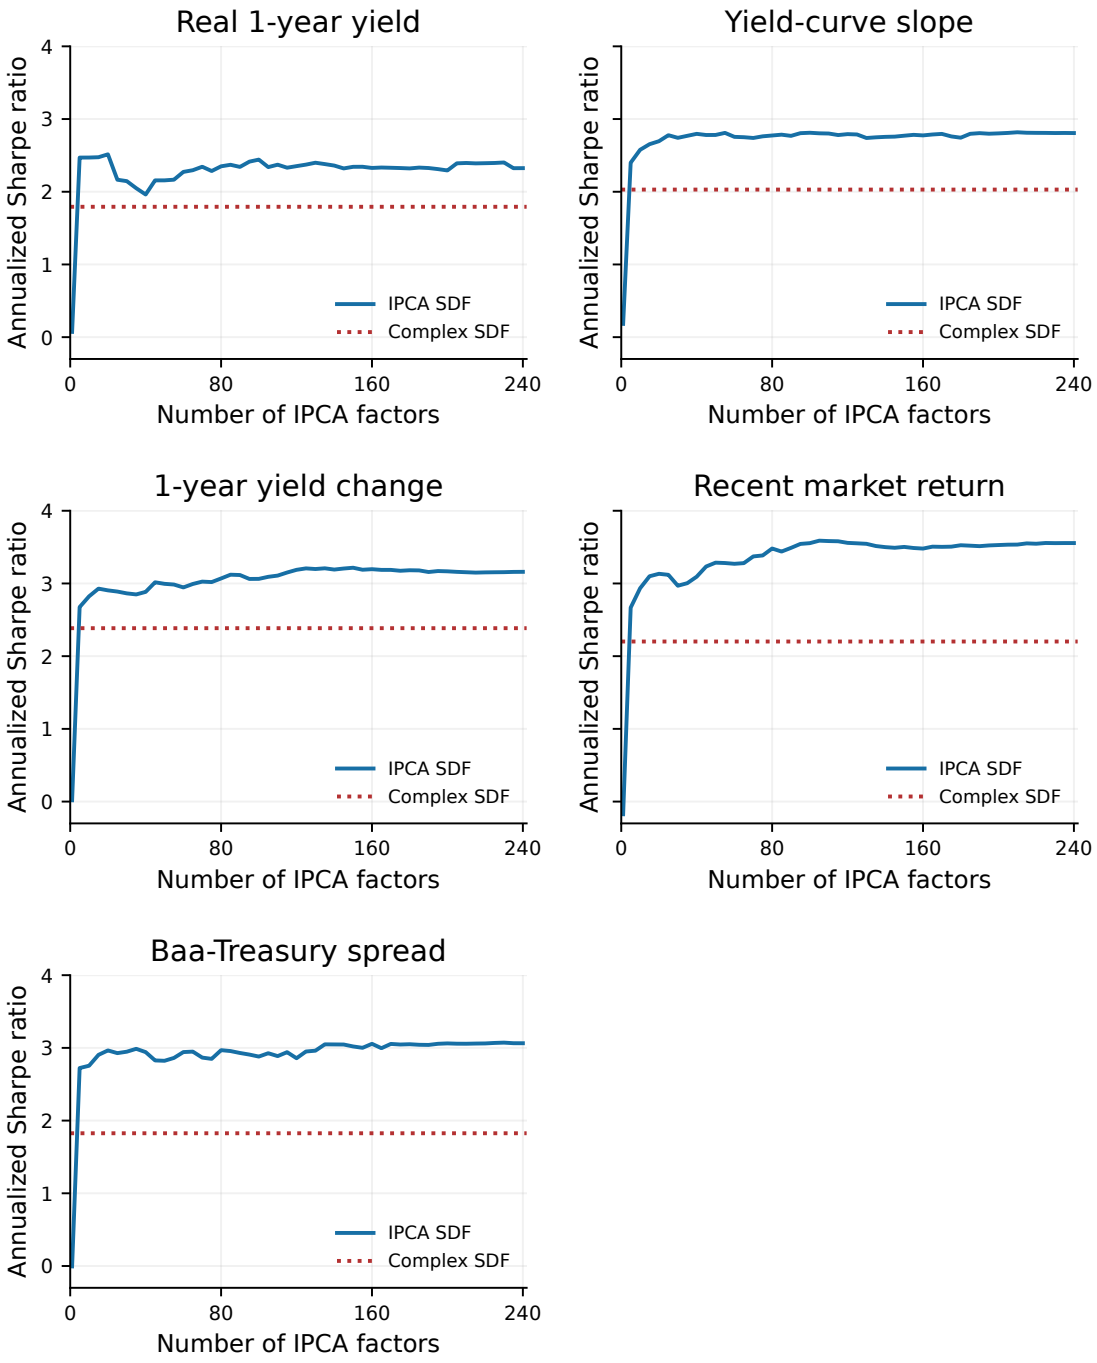
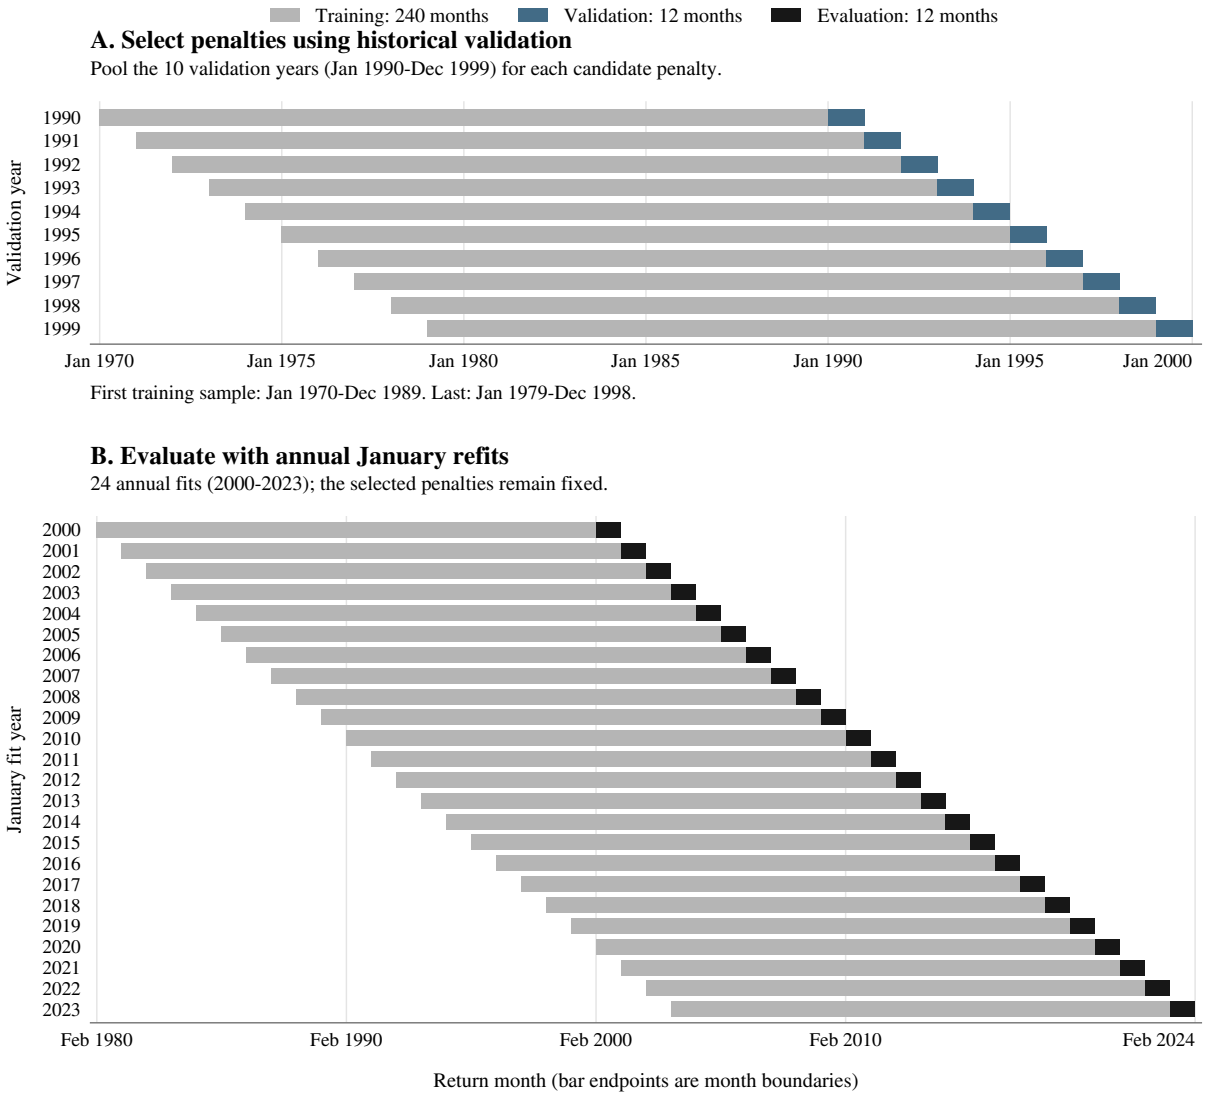


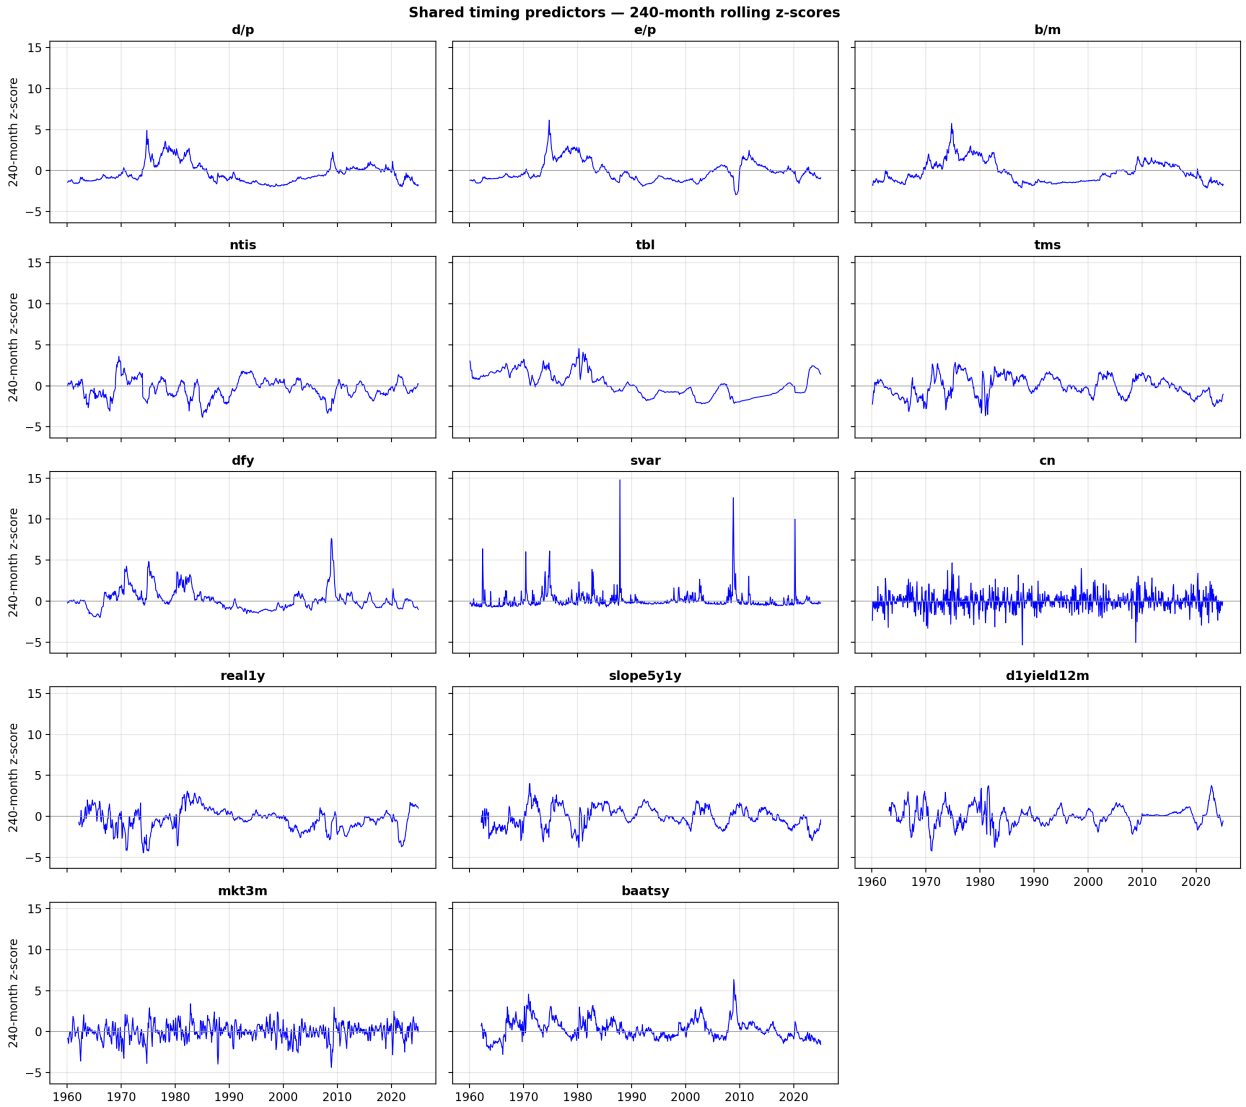
Figure 15: Individual Timing Comparisons: real1y, slope5y1y, d1yield12m, mkt3m, and baatsy. Each panel compares the IPCA SDF (blue solid) with its all-managed-return Complex SDF benchmark (red dotted). All panels use common axes. The baseline cap is 149; each timing cap is 240. Predictors are defined in Section 2.3. Penalties use fixed 1990–1999 validation, and evaluation covers February 2000–January 2024.

Figure 16: Historical Validation and Rolling Estimation Windows



Note: The panels show return months for the main monthly specifications. Panel A selects penalties using pooled 1990–1999 validation returns, with 240 preceding training months per fold. Panel B holds these penalties fixed and reestimates IPCA loadings, the factor rotation, and SDF coefficients each January. The first evaluation fit trains on February 1980–January 2000 and evaluates February 2000–January 2001; the last trains on February 2003–January 2023 and evaluates February 2023–January 2024. Fitted parameters remain fixed for twelve evaluation months; characteristics and exposures can change monthly. Historical validation months can subsequently enter training windows. Bar endpoints are the start of the month following the final return.

Figure 17: Shared Timing Predictors and Their 20-Year Rolling Z-Scores



Note: Each panel plots the 240-month (20-year) rolling z-score of one of the fourteen shared timing predictors—the standardized form in which the predictor scales the base factors—over 1960 to 2024, and all panels share a common vertical scale. The predictors are the eight Welch and Goyal (2008) variables (dp, ep, bm, ntis, tbl, tms, dfy, svar), the correlation neglect signal of Guo and Wachter (2026) (cn), and five macroeconomic predictors following Lehnherr et al. (2025) (real1y, slope5y1y, d1yield12m, mkt3m, baatsy). Standardizing each series highlights the periods in which a predictor deviates most from its trailing two-decade norm.