NBER WORKING PAPER SERIES

PRINCIPAL PORTFOLIOS

Bryan T. Kelly
Semyon Malamud
Lasse H. Pedersen

Working Paper 27388
http://www.nber.org/papers/w27388

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2020

We have received consulting income from AQR Capital Management exceeding $10,000 over the past three years. AQR Capital Management is a global investment management firm, which may or may not apply similar investment techniques or methods of analysis as described herein. The views expressed here are those of the authors and not necessarily those of AQR or the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w27388.ack

Principal Portfolios
Bryan T. Kelly, Semyon Malamud, and Lasse H. Pedersen
NBER Working Paper No. 27388
June 2020
JEL No. C1,C3,C58,C61,G11,G12,G14

ABSTRACT

We propose a new asset-pricing framework in which all securities' signals are used to predict each individual return. While the literature focuses on each security's own-signal predictability, assuming an equal strength across securities, our framework is flexible and includes cross-predictability—leading to three main results. First, we derive the optimal strategy in closed form. It consists of eigenvectors of a "prediction matrix," which we call "principal portfolios." Second, we decompose the problem into alpha and beta, yielding optimal strategies with, respectively, zero and positive factor exposure. Third, we provide a new test of asset pricing models. Empirically, principal portfolios deliver significant out-of-sample alphas to standard factors in several data sets.

Bryan T. Kelly
Yale School of Management
165 Whitney Ave.
New Haven, CT 06511
and NBER
bryan.kelly@yale.edu

Semyon Malamud
Swiss Finance Institute @ EPFL
Quartier UNIL-Dorigny, Extranef 213
CH - 1015 Lausanne
Switzerland
semyon.malamud@epfl.ch

Lasse H. Pedersen
Copenhagen Business School
Solbjerg Plads 3, A5
DK-2000 Frederiksberg
DENMARK
and NYU
Lhp001@gmail.com

1 Introduction

The starting point for much of asset pricing is a set of signals, $S_{i,t}$, that proxy for the conditional expected return for a security i at time t. In the context of an equilibrium asset pricing model $S_{i,t}$ may represent a conditional beta on compensated risk factors. Or it may be a predictor that is agnostic of equilibrium considerations, such as an asset's recent price momentum. Standard analyses, such as evaluating characteristic-sorted portfolios or asset pricing tests in the spirit of Gibbons et al. (1989), focus on own-asset predictive signals; that is, the association between $S_{i,t}$ and the return on only asset i, $R_{i,t+1}$.

We propose a new approach to analyzing asset prices through the lens of what we call the "prediction matrix." The prediction matrix, defined as $\Pi = E(R_{t+1}S_t')$, not only tracks the own-signal prediction effects ($\Pi_{i,i}$) but also all cross-predictability phenomena ($\Pi_{i,j}$) in which asset j's signal predicts asset i's return. Cross-predictability exists very generally in conditional asset pricing models, be they equilibrium in nature or purely statistical. Knowledge of the entire prediction matrix, as opposed to the typical focus on diagonal elements alone, is critical to devising optimal portfolios and understanding their risk-return tradeoff.

Our main contribution is to develop an extensive theoretical understanding of the prediction matrix and the asset pricing information it carries. The main tools of our analysis are singular value decompositions, analogous to using principal components analysis (PCA) to study variance-covariance matrices. The leading components (singular vectors) of Π are defined as those responsible for the lion's share of covariation between signals and future returns. This is where cross-predictability information becomes valuable. Like the diagonal elements, off-diagonal elements of Π are informative about the joint dynamics in signals and returns.

We refer to Π's singular vectors as "principal portfolios." They are a set of normalized portfolios ordered from those most predictable by S to those least predictable. The top principal portfolios are thus the most "timeable" portfolios, and as such they offer the highest unconditional expected returns for an investor that faces a leverage constraint (one who cannot hold infinitely large positions).

A key insight of our approach is that applying a singular value decomposition directly to Π conflates two different and opposing economic phenomena. We propose first splitting Π into a symmetric part (which is equal to its transpose and denoted Π^s) and an antisymmetric part (which is equal to minus its transpose and denoted Π^a), and applying separate matrix decompositions to Π^s and Π^a. The symmetry separation of Π,

$$\Pi = \underbrace{\frac{1}{2}(\Pi + \Pi')}_{\Pi^s} + \underbrace{\frac{1}{2}(\Pi - \Pi')}_{\Pi^a}, \tag{1}$$

is a powerful device. With eigenvalue decompositions of each part, we can take a complicated collection of predictive associations in the Π matrix and decode them into a set of well-organized facts about expected returns. These facts describe i) the nature of each predictive pattern represented in Π and ii) the strength of these patterns.

The nature of a predictive pattern is described by its classification as either symmetric or antisymmetric, which, amazingly, translate into beta and alpha. In particular, we show that eigenvectors of the symmetric matrix Π^s are optimal ways to achieve factor exposure (beta), while eigenvectors of the antisymmetric matrix Π^a are optimal factor-neutralized strategies (alpha). We refer to strategies arising from eigenvectors of the symmetric component as "principal exposure portfolios" (PEPs) and the strategies arising from the antisymmetric part as "principal alpha portfolios" (PAPs). Once classified as "exposure" versus "alpha," prediction patterns (principal portfolios) are then ordered from strongest to weakest and based on the size of their associated eigenvalues. In particular, we prove that the unconditional average returns of PEPs and PAPs are exactly proportional to their respective eigenvalues.

This decomposition has a close connection to equilibrium asset pricing. When signals are betas to the pricing kernel and there is no-arbitrage, all PAPs must deliver zero expected excess returns (because they have zero factor exposure) and all PEPs must deliver non-negative average returns (because they have positive exposure to the pricing kernel). These insights are the groundwork for a new asset pricing test based on eigenvalues of the symmetric and antisymmetric components of the prediction matrix. In rational asset pricing models, there should not be any alpha relative to the pricing kernel. When we pick signals that are supposed to be proportional to covariances with the pricing kernel (e.g., market betas), then the corresponding prediction matrix should have a zero antisymmetric part—meaning that Π should be symmetric and there should be no pure alpha portfolios. Moreover, negative eigenvalues of the symmetric part of Π correspond to strategies with negative factor exposure and positive expected returns, another form of alpha. Since rational asset pricing also rule out both forms of alpha, we get the asset pricing test that Π should be symmetric and positive definite. In other words, when signals capture exposure to the pricing kernel, all PEPs should deliver non-negative returns and all PAPs should deliver zero returns.

We also develop theoretical underpinnings for practical empirical usage of the prediction matrix from the perspective of robust statistics and machine learning. Our main theoretical results characterize the properties of principal portfolios and their role in optimal portfolios are developed in population, where Π is known. With N assets, this requires estimating N^2 parameters. Such rich parameterization can lead to noisy estimates and overfit that deteriorate the out-of-sample performance of principal portfolios. In the literature and financial practice, signals are often analyzed or traded in the form $\sum_i S_{i,t}R_{i,t+1}$, which essentially restricts the signal-based analysis to testing a single parameter equal to average own-predictability, $\sum_i E(S_{i,t}R_{i,t+1})$. While this may benefit from some robustness, restricting the analysis to a one-parameter problem is harsh—it forfeits any and all useful information about heterogeneity in own-predictability or cross-predictability in the estimated

Π matrix. Principal portfolios are ideally suited to balance the joint considerations of exploiting potentially rich information from throughout Π while controlling parameterization to reduce noise and overfit. We show that low-rank approximations of Π and its symmetry-based components Π^s and Π^a offer a means of balancing both considerations in a data-driven way in order achieve robust out-of-sample portfolio performance.

We implement the methodology empirically using three samples of U.S. equities, three samples of international equities, and a sample of futures contracts on equity indices, bonds, commodities, and currencies. As an example of a trading signal that can be used for all securities, we analyze momentum (i.e., past returns). We conduct out-of-sample analyses that, at each time period t, estimates the prediction matrix from past signals and returns (i.e., only information that is available through time t). Estimating the prediction matrix is easy: $\hat{\Pi}_t = \frac{1}{120}\sum_{\tau=t-120}^{t-1} R_{\tau+1}S'_\tau$, where we use a backward looking window of 120 time periods. Having estimated the prediction matrix, the singular value and eigenvalue decompositions of Π as well as its symmetric and antisymmetric parts immediately yield PPs, PEPs, and PAPs, and we track their out-of-sample performance. We find that the leading principal portfolios tend to deliver positive returns across all samples and large number of robustness checks, with highly significant risk-adjusted returns in a number of specifications. In the base case specification, the combination of leading PEPs and PAPs delivers more than twice the Sharpe ratio of a standard factor constructed based on the same signal.

Our paper is related to several literatures. Our asset pricing test complements other such tests, including Gibbons et al. (1989) and Hansen and Jagannathan (1991) (see Cochrane (2009) for an overview). Our method to uncover new forms of predictability complements existing methods based on regressions (see Welch and Goyal (2008) and references therein), portfolio sorts (a recent example is Fama and French (2015)), and machine learning (Gu et al. (2018)). We use momentum signals, which have been used extensively in equities (Jegadeesh and Titman (1993)), global asset markets (Asness et al. (2013), Moskowitz et al. (2012)), and, more recently, on factor returns (Arnott et al. (2019), Gupta and Kelly (2019)). Finally, we consider linear trading strategies, which have also been studied in the context of dynamic trading with transaction costs by Gârleanu and Pedersen (2013, 2016), Collin-Dufresne et al. (2015), Collin-Dufresne et al. (2019), and others. While this literature focuses on linear-quadratic programming, we instead consider eigenvalue methods.

In summary, we present a new way to uncover return predictability and test asset pricing models. We illustrate how the method works empirically with a wide range of encouraging results for out-of-sample principal portfolio performance.

2 Principal Portfolio Analysis

In this section, we lay of our principal portfolio analysis (PPA) framework. We describe the concept of linear strategies of predictive signals, show how linear strategies are intimately linked to the prediction matrix, derive optimal strategies, and introduce the notion of principal portfolios that implement optimal strategies.

Let us first introduce the setting and notation that we use throughout. The economy has N securities traded at discrete times. At each time t, each security i delivers a return in excess of the risk-free rate, $R_{i,t}$. All excess returns at time t are collected in a vector, $R_t = (R_{i,t})_{i=1}^N$ and their conditional variance-covariance matrix is $\Sigma_{R,t} = \text{Var}_t(R_{t+1})$.

For each time and security, we have a "signal" or "characteristic" $S_{i,t}$, and all signals are collected in a vector, $S_t = (S_{i,t})_{i=1}^N$. We can think of these predictive characteristics as market betas, valuation ratios, momentum scores, or other observable signals that proxy for conditional expected returns.

2.1 Linear Trading Strategies

How can an investor best exploit predictive information in an asset characteristic S? To answer this question, we work in the context of general linear trading strategies based on

S. Then, we derive an optimal linear strategy subject to leverage constraints and show the intimate connection between the optimal linear strategy and the prediction matrix Π.

A *linear strategy* based on S has portfolio weights of the form $w_t' = S_t'L$. We refer to $L \in \mathbb{R}^{N\times N}$ as the *position matrix* because each column of L translates signals into a portfolio position in each asset. For example, the first column $L_1 = (L_{i,1})_{i=1}^N$ of L translates all the signals into a position in asset 1, $S_t'L_1$. The return of a linear strategy is naturally the positions times the returns, that is,

$$R_{t+1}^{w_t} = w_t'R_{t+1} = \sum_j \underbrace{(S_t'L_j)}_{\text{position in } j} \quad \underbrace{R_{j,t+1}}_{\text{return of } j} \;=\; S_t'LR_{t+1}\,. \tag{2}$$

We see that a linear strategy generally allows the position $S_t'L_j$ in any asset j to depend on the signals of *all* assets. Interestingly, these strategies can potentially exploit both predictability using each asset's own signal as well as cross-predictability using other signals.

The large majority of return prediction patterns in the empirical literature focus on strategies that are agnostic of cross-predictability. The literature's default portfolio construction based on a characteristic S builds a simple tradable factor of the form:

$$\widetilde{F}_{t+1} \;=\; \sum_j S_{j,t}R_{j,t+1} \tag{3}$$

We refer to $\widetilde{F}_{t+1}$ as the "simple factor" henceforth. We note that the simple factor is a linear strategy with identity position matrix ($L = \mathrm{Id}$):

$$\widetilde{F}_{t+1} = \sum_i S_{i,t}R_{i,t+1} = S_t'R_{t+1} \;=\; S_t'\mathrm{Id}R_{t+1}. \tag{4}$$

Hence, our framework nests the standard framework, and allows more general strategies.

The simplicity of the simple factor-mimicking strategy makes it a helpful reference point for the strategies we advocate in this paper. It is a portfolio that relies only on own-signal predictions with no cross-prediction. Moreover, it imposes that own-signal predictions enter into the portfolio uniformly, with no regard for heterogeneity in predictive effects across assets. When a researcher reports that this type of simple factor has a positive average return, $E(\widetilde{F}_{t+1}) > 0$, it is the same as saying that the signal positively predicts own-asset returns on average.

2.2 The Prediction Matrix

A central part of our analysis makes use of what we call the **prediction matrix**:

$$\Pi = E(R_{t+1}S_t'). \tag{5}$$

Π encodes predictive information for how the signals predict all returns, based on assets' own signals as well as cross-predictability. A strategy that literally chooses an asset's position equal to its own signal $S_{i,t}$ earns a return of $R_{i,t+1}S_{i,t}$, and $\Pi_{i,i}$ is the expected value of this return. Likewise, a strategy that take a position in asset i based on the signal of another asset j earns average returns of $\Pi_{i,j}$.

If $S_{j,t}$ predicts $R_{j,t+1}$ on average across securities, then this is the same as saying that the prediction matrix has a positive trace (tr, the sum of its diagonal elements):

$$E\left(\sum_j S_{j,t}R_{j,t+1}\right) = \text{tr}(\Pi) \ > \ 0\,. \tag{6}$$

This notion of positive own-predictability on average across securities has emerged as the standard criterion by which predictive signals are measured in the empirical finance literature and is typically evaluated based on the sample average of the strategy in (3).

Average own-predictability not only abstracts from information in off-diagonal elements of Π, but also from heterogeneity in own-effects on the main diagonal. In short, strategies predicated on average own-predictability are highly constrained in the information they consider regarding the predictive content of S. Proposition 1 shows that the *entire* Π matrix is necessary (and sufficient!) for understanding the returns of more general linear strategies.

Proposition 1 (Return of Linear Strategies) *The expected excess return of a linear trading strategy $w_t' = S_t' L$ is*

$$E\left(R_{t+1}^{w_t}\right) = E\left(S_t' L R_{t+1}\right) = \mathrm{tr}(L\Pi). \tag{7}$$

An interesting linear strategy in its own right is to take positions in every asset based on the magnitude of its predictability by the signal S, whether that information comes from its own signal or from another asset's signal. This amounts to using Π itself as the position matrix ($L = \Pi'$) or using a positive multiple of Π:

Proposition 2 (Trading the Prediction Matrix) *Let M be an arbitrary positive semi-definite matrix. Then, the linear strategy with position matrix $L = M\Pi'$ has positive expected excess return:*

$$E(S_t' L R_{t+1}) \;=\; \mathrm{tr}(M\,\Pi'\Pi) \;=\; \mathrm{tr}((\Pi M^{1/2})'\,(\Pi M^{1/2})) \;\geq\; 0\,. \tag{8}$$

Moreover, the inequality is strict if and only if $M^{1/2}\Pi'$ is not identically zero.

We see that the prediction matrix plays two important roles: First, Π tells us the return of any linear strategy as seen in Proposition 1. Second, Π' is itself a return-generating linear strategy as seen in Proposition 2.

2.3 Optimal Linear Strategies

We next show that strategies based on Π not only deliver positive expected return, they actually yield *optimal* linear strategies. Our precise statement of optimality is derived from the following objective function.

$$\max_{L:\|L\|\leq 1} \; E\left(S_t' L R_{t+1}\right) \, . \tag{9}$$

The objective is to maximize the expected return of a linear strategy subject to a portfolio constraint on the position matrix L. (We naturally need a portfolio constraint, since otherwise we can increase the expected return by simply increasing position sizes, e.g., the strategy $2L$ doubles the expected return of the strategy L).

To understand the constraint that we use in (9), note first that $\|x\| \equiv \left(\sum_i x_i^2\right)^{1/2}$ is the standard Euclidean norm of a vector $x \in \mathbb{R}^N$. Second, we define the standard matrix norm as $\|L\| = \sup\{\|Lx\| : x \in \mathbb{R}^m \text{ with } \|x\| = 1\}$. It is possible to show that $\|L\| = \|L'\|$. In other words, in (9) we maximize the expected return over the set of all position matrices with matrix norm of at most one.

The economic meaning of this constraint is that we consider strategies with a bounded portfolio size. Specifically, the linear strategy has portfolio weight $S_t' L$, which has a size of $\|L' S_t\| \leq \|L'\| \, \|S_t\| \leq \|S_t\|$ when $\|L\| \leq 1$. So we consider linear strategies where the position size is always bounded by the position size of the simple strategy. Further, if S_t is normalized such that $\|S_t\| = 1$ for all signals, then the linear strategy has a position size that is similarly bounded, $\|L' S_t\| \leq 1$.[1]

We can also interpret the objective function as a robust mean-variance problem. For example, when the return variance-covariance matrix is given by $\Sigma_{R,t} = \sigma^2 \mathrm{Id}$ for some $\sigma \in \mathbb{R}$, the objective function (9) is identical to the following:

$$\max_L E(S_t' L R_{t+1}) \text{ subject to } \max_{S:\mathrm{Var}_t(S' R_{t+1}) \leq 1} \mathrm{Var}_t(S' L R_{t+1}) \leq 1 \, . \tag{10}$$

[1] Here we discuss "position size" in terms of the Euclidian norm, while the notional leverage of a position x is normally calculated as $\|x\|_1 = \sum_k |x_k|$. However, the portfolio constraint $\|L\| \leq 1$ also implies a constraint on notional leverage. Indeed, since $\|x\|_1 \leq \|x\| n^{1/2}$, notional leverage is bounded: $\|L' S_t\|_1 \leq \|L' S_t\| n^{1/2} \leq n^{1/2}$.

In words, we maximize expected return subject to a risk constraint. This risk constraint is robust in the sense that we require that the variance is bounded *regardless* of the signal realization S. This robust objective where we maximize risk with respect to S, rather than considering the risk conditional on S, is natural for a linear strategy — since the position matrix L is constant over time and should "work" for all signals. To see the equivalence of (9) and (10), note that

$$\max_{S:\mathrm{Var}_t(S'R_{t+1})\leq 1} \mathrm{Var}_t(S'LR_{t+1}) = \max_{S:S\neq 0} \frac{\mathrm{Var}_t(S'LR_{t+1})}{\mathrm{Var}_t(S'R_{t+1})} = \max_{S:S\neq 0} \frac{\sigma^2\|LS\|^2}{\sigma^2\|S\|^2} = \|L\|^2 \,. \tag{11}$$

The risk constraint says that the risk of the linear strategy should be at most as high as that of the simple factor. Another way to get the same result is to require that the risk is limited when the signals are limited, $\max_{S:\|S\|\leq 1} \mathrm{Var}_t(S'LR_{t+1}) \leq \sigma^2$.

The assumption $\Sigma_{R,t} = \sigma^2 \mathrm{Id}$ is appropriate if volatilities are similar in the cross section (or has been made similar, as we do our empirical study of futures) and if the correlation matrix is close to, or has been shrunk to, the identity—and such shrinkage can be useful in an optimization setting (Pedersen et al. (2020)). In any event, when we are have general variance-covariance matrix $\Sigma_{R,t}$, then our portfolio constraint $\|L\| \leq 1$ still serves to control both risk, leverage, and the portfolio norm.[2] Further, we show how to solve a robust mean-variance problem for general $\Sigma_{R,t}$ in Appendix A. The appendix also shows how to solve the mean-variance problem with a risk penalty driven by risk aversion (instead of the risk constraint used here).

Given the objective (9), the solution for the optimal strategy is as follows.

[2]The portfolio constraint $\|L\| \leq 1$ implies a limit on the portfolio norm by definition, a leverage limit described in Footnote 1, and the following risk limit:

$$\max_{S:\|S\|\leq 1} \sqrt{\mathrm{Var}_t(S'LR_{t+1})} \;=\; \|\Sigma_{R,t}^{1/2} L'\| \leq \|\Sigma_{R,t}^{1/2}\| \|L\| \leq \|\bar{\Sigma}\|$$

when the variance-covariance matrix is bounded, $\Sigma_{R,t} \leq \bar{\Sigma}$.

Proposition 3 *The solution to (9) is given by $L = M\Pi'$ with $M = (\Pi'\Pi)^{-1/2}$, and*

$$\max_{L:\|L\|\leq 1} E\left(S_t' L R_{t+1}\right) = \sum_{i=1}^{N} \bar{\lambda}_i \,,$$

where $\bar{\lambda}_1 \geq \cdots \geq \bar{\lambda}_N$ are the singular values of Π, i.e., the eigenvalues of $(\Pi'\Pi)^{1/2}$.

Proposition 3 shows that the full Π matrix is integral to optimal linear strategies based on the signal S_t.[3] Indeed, maximum return depends on the singular values of Π, which in general depend on all its elements.

2.4 Principal Portfolios

We next decompose the optimal solution into a collection of linear strategies that we refer to as **principal portfolios** (PP) of the signal S. Principal portfolios are the building blocks that sum to form the optimal linear strategy in Proposition 3.

The construction of PPs uses the singular value decomposition of Π. Namely, let

$$\Pi = U\,\bar{\Lambda}\,V', \tag{12}$$

where $\bar{\Lambda} = \operatorname{diag}(\bar{\lambda}_1, \cdots, \bar{\lambda}_N)$ is the diagonal matrix of singular values, and U, V are orthogonal matrices with column denoted u_k and v_k, respectively. Now, the optimal L from Proposition 3 can be rewritten as

$$(\Pi'\Pi)^{-1/2}\Pi' = V\bar{\Lambda}^{-1}V'V\bar{\Lambda}\,U' = VU' = \sum_{k=1}^{N} v_k u_k' \,.$$

We define the k^{th} principal portfolio as the linear strategy with position matrix $L_k = v_k\,(u_k)'$,

[3] In particular, the optimal strategy is of the form described in Proposition 2.

which has a return of

$$PP_{t+1}^k = S_t' \underbrace{v_k\, u_k'}_{L_k} R_{t+1} = \underbrace{S_t' v_k}_{S_t^{v_k}}\, \underbrace{u_k' R_{t+1}}_{R_t^{u_k}}\,. \tag{13}$$

We see that there are two interpretation of a principal portfolio. First, it is a simple linear strategy with a position matrix L of rank 1. Second, it is a strategy that trades the portfolio u_k (with return $R_t^{u_k}$) based on the signal coming from the portfolio v (i.e., with signal $S_t^{v_k}$). This latter interpretation plays a key role when we discuss the beta components in the next section.

The construction of principal portfolios is actually very simple. All you need to do is use your favorite program to compute the singular value decomposition of Π (a standard feature of most computing programs), take the column vectors of U and V, and you are done.

Decomposing the optimal strategy into its principal portfolios is similar to decomposing the variance into the principal components. The difference is that principal component analysis decomposes the *variance*, but principal portfolio analysis decomposes the *expected return*. Just like the variance of each principal component equals its corresponding eigenvalue, the expected return of each principal portfolio is its singular value:æ

$$E(PP_{t+1}^k) \;=\; \mathrm{tr}(\Pi\, v_k u_k') \;=\; \mathrm{tr}(U\,\bar{\Lambda}\,V' v_k u_k') \;=\; \mathrm{tr}(U\,\bar{\Lambda}\,e_k u_k') \;=\; \mathrm{tr}(\bar{\lambda}_k u_k u_k') \;=\; \bar{\lambda}_k\,. \tag{14}$$

The following proposition summarizes the results of this section.

Proposition 4 *The expected return of each principal portfolio is given by its corresponding singular value,*

$$E(PP_{t+1}^i) = \bar{\lambda}_i, \tag{15}$$

and the sum of principal portfolios is the optimal linear strategy:

$$\max_{\|L\|\leq 1} E(S_t'LR_{t+1}) = E\left(\sum_{i=1}^{N} PP_{t+1}^i\right) = \sum_{i=1}^{N} \bar{\lambda}_i. \tag{16}$$

The following example provides some intuition for this result.

Example (Signals are Expected Returns). If signals are equal to conditional expected returns, $S_{i,t} = E_t(R_{i,t+1})$, one might question the usefulness of principal portfolios. But even in this simple setting principal portfolios are insightful about the optimal strategy. In this case, the prediction matrix reduces to the unconditional second moment of S_t, denoted Σ_S,

$$\Pi = E(R_{t+1}S_t') = E(E_t(R_{t+1})S_t') = E(S_tS_t') = \Sigma_S. \tag{17}$$

Therefore, principal portfolios are given by the principal components of Σ_S. The matrix Σ_S describes the joint dynamics in conditional expected returns. Its leading principal component describes the portfolio of assets with the most variable expected return. In other words, the first principal component of Σ_S is the most "timeable" portfolio. It is the most attractive portfolio to trade for an investor facing a position size constraint and delivers the highest unconditional average profitability. The second principal component is the next most attractive, and so on. Singular values of Π relate to variability of expected returns, which explains why unconditional expected returns on principal portfolios are pinned down by the size of singular values in (15). And in this example, all principal portfolios have positive expected excess returns (assuming that Σ_S is non-degenerate), so the optimizing investor holds them all, as in (16). However, if the prediction matrix is estimated with error, it may be more robust to focus on the top PPs, as discussed in Section 5.

3 Optimal Alpha and Beta Strategies

We next derive the return of the optimal alpha and beta strategies, and show how these can be decomposed into principal portfolios, just as in the general solution in Propositions 3–4.

3.1 Alpha-Beta Symmetry Decomposition

To decompose the return into its alpha and beta components, we must first specify the beta-factor. In other words, how do we characterize the riskiness of linear strategies? To address this question, Lemma 1 introduces a factor having the special property that $S_{i,t}$ exactly describes asset i's conditional exposure to the factor.

Lemma 1 (Characteristics as Covariances) *Define the factor F_{t+1} as*

$$F_{t+1} = \left(\frac{1}{S_t' \left(\Sigma_{R,t}\right)^{-1} S_t} \left(\Sigma_{R,t}\right)^{-1} S_t \right)' R_{t+1}. \tag{18}$$

F_{t+1} is the unique tradable factor with the property that

$$S_{i,t} \;=\; \frac{\mathrm{Cov}_t(R_{i,t+1}, F_{t+1})}{\mathrm{Var}_t(F_{t+1})}. \tag{19}$$

This factor (referred to as the "latent factor" henceforth) is an economically important reference point.[4] It has a natural risk factor interpretation—it is the factor that unifies the expected return interpretation of $S_{i,t}$ and the risk exposure interpretation of $S_{i,t}$. No other factor based on S shares this property (including the literature's standard "simple factor," $\widetilde{F}$).

[4] Kelly et al. (2020a,b) propose a modeling approach and extensive empirical study of this point. Lemma 1 shows that we can always think of any signals as exposures to a factor, but it does not necessarily imply that the return predictability embodied by S is "rational" in the sense that the factor F covaries with risks that investors care about, namely the pricing kernel, and that certain eigenvalue bounds are satisfied, as discussed later.

To interpret this result, it is again helpful to consider the example in which $S_t = E_t(R_{t+1})$. In this case, F_{t+1} is the conditional tangency portfolio, and is thus the tradable representation of the pricing kernel. As a result, the expected returns and the factor exposures are equal up to a constant positive scale factor. And, being the tangency portfolio, all assets have zero alpha versus this factor in the absence of arbitrage. Importantly, while this factor is useful for interpreting some of our results, none of our results rely on actually observing F—since we don't observe it. We don't observe F because it depends on the conditional variance-covariance matrix $\Sigma_{R,t}$, which can only be estimated with noise. Instead, we develop methods that can beat the simple factor $\widetilde{F}$ without relying on observing, much less inverting, $\Sigma_{R,t}$.

The risk factor interpretation of the latent factor F positions it as the key benchmark for evaluating the performance of principal portfolios. With it, we can characterize the risk and return of signal-based linear strategies.

To state the next result, recall that any square matrix $B \in \mathbb{R}^{N\times N}$ is decomposable into its *symmetric part*, $B^s = \frac{1}{2}(B + B')$, and its *antisymmetric part*, $B^a = \frac{1}{2}(B - B')$, where $B = B^s + B^a$. The symmetric part equals its own transpose while the antisymmetric part equals minus its own transpose, and both parts have a number of interesting properties. For example, with $B^a = -B^{a\prime}$, it has zeros along the main diagonal.

Hence, any linear strategy can be seen as a sum of a symmetric and antisymmetric part, $L = L^s + L^a$. As we now show, this decomposition has a deep economic interpretation.

Proposition 5 (Alpha-Beta Symmetry Decomposition) *The conditional latent factor exposure and expected return of the strategy* $R^{w_t}_{t+1} = S_t' L R_{t+1} = S_t' L^s R_{t+1} + S_t' L^a R_{t+1}$ *is*

$$\underbrace{\frac{\mathrm{Cov}_t(R^{w_t}_{t+1},\ F_{t+1})}{\mathrm{Var}_t(F_{t+1})}}_{\textit{factor beta}} = S_t' L^s S_t \tag{20}$$

$$E(R^{w_t}_{t+1}) = \mathrm{tr}(L^s \Pi^s) + \mathrm{tr}(L^a \Pi^a)\,. \tag{21}$$

This proposition shows that the risk (beta to the latent factor) of a linear strategy $S_t'L$ is purely determined by its symmetric part; while the expected return is determined by both the symmetric or anti-symmetric parts via their interaction with the respective components of the prediction matrix, Π^s and Π^a.

This proposition has wide-ranging implications. First, an antisymmetric strategy is always factor neutral. Second, an antisymmetric strategy can nevertheless deliver positive returns as long as $\Pi^a \neq 0$. In this case, an antisymmetric strategy can deliver positive expected return with zero factor exposure, that is, pure alpha! The fact that factor exposures depend only on the symmetric component, L^s, regardless of the symmetry of Π is a direct implication of Lemma 2.

Lemma 2 *For any symmetric matrix $B \in \mathbb{R}^{N\times N}$ and any anti-symmetric matrix $A \in \mathbb{R}^{N\times N}$, we have $\operatorname{tr}(BA) = \operatorname{tr}(AB) = 0$ and $x'Ax = 0$ for all vectors $x \in \mathbb{R}^N$.*

In other words, antisymmetric matrices nullify certain matrix multiplications, which translates into factor-neutrality of trading strategies.

Proposition 5 also shows how symmetric strategies can deliver returns via the interaction with Π^s. Symmetric strategies have a beta to the factor given by $S_t'L^sS_t$, which can be positive or negative. A symmetric strategy has positive factor beta for all possible realizations of the signal vector S_t if and only if L is positive definite. So, as we analyze in more detail in the next section, eigenvalues are key to understanding both risk and return. Finally, a symmetric strategy that always has negative factor beta corresponds to a negative definite L.

As an example application of Proposition 5, consider the riskiness of the simple factor $\widetilde{F}$ in (3), which is a linear strategy with identity position matrix ($L = \mathrm{Id}$) as seen in Equation 4. Hence, this simple factor has expected return $\operatorname{tr}(L^s\Pi^s) = \operatorname{tr}(\Pi^s) = \operatorname{tr}(\Pi)$ and it always has a positive exposure to the latent factor, $\operatorname{Cov}_t(\widetilde{F}_{t+1}, F_{t+1}) = \operatorname{Var}_t(F_{t+1})S_t'S_t > 0$.

The optimal linear strategy in Proposition 3 and the corresponding principal portfolios do not distinguish whether expected returns originate from factor exposure or alpha. In the remainder of this section, we show that Π^s and Π^a lie at the heart of optimal symmetric and antisymmetric trading strategies. We derive symmetric and antisymmetric analogues of principal portfolios, and show that these are the building blocks to optimal symmetry-decomposed strategies with either pure factor exposure and no alpha, or pure alpha and no factor exposure.

Said simply, symmetry is beta, and antisymmetry is alpha. We next derive the optimal beta and alpha, respectively.

3.2 Symmetric Strategies: Principal Exposure Portfolios

As shown in equation (4), the simple factor is a simple symmetric linear strategy that trades each asset based on its own signal. The idea that symmetric strategies trade based on their own signals holds more generally. In particular, any strategy that scales the portfolio position in proportion to the signal aggregated to the portfolio level—that is, any portfolio that trades on the portfolio's own signal—is a symmetric strategy.

To see this, consider a portfolio $w \in \mathbb{R}^N$. The portfolio w has excess return $R^w_{t+1} = \sum_i w_i R_{i,t+1}$. Aggregating the underlying signals based on these weights means that the portfolio-level own signal is $S^w_t = \sum_i w_i S_{i,t}$. Trading the portfolio based on its own signal means using its signal as portfolio weight, which generates a return of

$$S^w_t R^w_{t+1} = S'_t w w' R_{t+1}. \tag{22}$$

We see that trading the portfolio based on its own signal is a linear strategy with a symmetric, positive semi-definite position matrix $L = ww'$. It's expected return is therefore

$$E\left(S^w_t R^w_{t+1}\right) = E\left(w' S_t R'_{t+1} w\right) = w'\Pi w \ = w'\Pi^s w, \tag{23}$$

which shows, in a different way from (21), that the return depends only on the symmetric part of the prediction matrix (the last equality uses Lemma 2).

All symmetric linear strategies can be represented as combinations of portfolios traded based on their own signals. This is achieved through the eigendecomposition of any symmetric position matrix L based on its eigenvalues λ_k and orthonormal eigenvectors w_k:

$$L = \sum_{k=1}^{K} \lambda_k \, w_k \, (w_k)' \,. \tag{24}$$

Furthermore, the position matrix satisfies our portfolio constraint $\|L\| \leq 1$ if $|\lambda_k| \leq 1$ for all k.

This result provides intuition for why symmetric linear strategies have factor exposure. They trade portfolios based on the portfolio's own signal. In this sense, they do what the signal prescribes, which anchors their behavior to that of the factor F. For example, if the signal $S_{i,t}$ is each security's momentum, then a symmetric linear strategy consists of trading different portfolios based on their own momentum—in the same spirit as the factor.

We next consider *optimal* symmetric linear strategies. We know from (21) that a optimal symmetric strategy maximizes $\mathrm{tr}(L\Pi^s)$, so we can use Proposition 3 with Π replaced by Π^s. The solution can be written simplæy based on the eigenvalue-decomposition

$$\Pi^s \;=\; W \Lambda^s W' = \sum_{k=1}^{N} \lambda_k^s \, w_k^s \, (w_k^s)' \,, \tag{25}$$

where $W = (w_1^s, ..., w_N^s)$ is the matrix of eigenvectors corresponding to the eigenvalues $\lambda_1^s \geq ... \geq \lambda_N^s$. We see that the optimal symmetric strategy is:

$$(\Pi^s \Pi^s)^{-1/2} \Pi^s = W |\Lambda^s|^{-1} W' \; W \Lambda^s W' = W \, \mathrm{sign}(\Lambda^s) W' = \sum_{k=1}^{N} \mathrm{sign}(\lambda_k^s) \, w_k^s \, (w_k^s)' \,. \tag{26}$$

We see that the optimal strategy naturally decomposes into N components, which we call

principal exposure portfolios (PEPs). That is, the k^{th} PEP is a linear strategy with position matrix $w_k^s(w_k^s)'$ and a return of

$$PEP_{t+1}^k = S_t^{w_k^s} R_{t+1}^{w_k^s} = S_t' w_k^s (w_k^s)' R_{t+1}. \tag{27}$$

The next result characterizes the returns of PEPs:

Proposition 6 *The expected return of each PEP is equal to its corresponding eigenvalue*

$$E(PEP_{t+1}^k) = E\left(S_t^{w_k^s} R_{t+1}^{w_k^s}\right) = E\left(S_t' w_k^s (w_k^s)' R_{t+1}\right) = \lambda_k^s, \tag{28}$$

Going long PEPs with positive eigenvalues and short those with negative is the optimal symmetric linear strategy:

$$\max_{\|L\|\leq 1,\ L=L'} E(S_t' L R_{t+1}) = \sum_{k=1}^N \text{sign}(\lambda_k^s) E(PEP_{t+1}^k) = \sum_{k=1}^N |\lambda_k^s|. \tag{29}$$

The first result shows that returns of PEPs equal their eigenvalues. The second result shows that the collection of PEPs yield the symmetric linear strategy with the highest unconditional expected return, subject to leverage constraint $\|L\| \leq 1$. This optimal performance is achieved by trading PEPs while accounting for the direction of their predictability. The optimal strategy takes long positions of size 1 in all PEPs with positive expected returns (i.e., positive eigenvalues) and short positions of size -1 in PEPs with negative expected returns.

We next consider how the PEPs relate to the simple factor $\widetilde{F}$.

Proposition 7 (Beating the Factor) *The simple factor, $\widetilde{F}$, can be decomposed as*

$$\widetilde{F}_{t+1} = \sum_{i=1}^N S_{i,t} R_{i,t+1} = \sum_{k=1}^N S_t^{w_k^s} R_{t+1}^{w_k^s} = \sum_{k=1}^N PEP_{t+1}^k. \tag{30}$$

If all eigenvalues are non-negative, $\lambda_k^s \geq 0$, then $\widetilde{F}$ the optimal symmetric strategy. Otherwise, $\widetilde{F}$ has a lower expected return than buying the subset of PEPs with positive eigenvalues, which is lower than that the optimal strategy from Proposition 6:

$$E\left(\widetilde{F}_{t+1}\right) = \sum_{k=1}^{N} \lambda_k^s \;\leq\; \sum_{k:\lambda_k^s>0} \lambda_k^s \leq \sum_{k=1}^{N} |\lambda_k^s| \,. \tag{31}$$

Interestingly, the simple factor actually equals the sum of all PEPs as seen in (30). In fact, $\widetilde{F}$ can be viewed as the sum of all possible returns of symmetric strategies, not just the PEPs. Namely, for any orthonormal basis of portfolios $B = \{b_k\}_{k=1}^N$, we have that $BB' = \mathrm{Id}$ and, hence,

$$\sum_{i=1}^{N} S_{i,t}\, R_{i,t+1} \;=\; S_t' R_{t+1} \;=\; S_t' B B' R_{t+1} \;=\; \sum_{k=1}^{N} S_t^{b_k}\, R_{t+1}^{b_k} \,. \tag{32}$$

That is, *trading the simple factor on stocks is equivalent to trading it on portfolios.*

The fact $\widetilde{F}$ equals the sum of PEPs together with Equation (28) imply that the expected excess return of the simple factor equals the sum of the eigenvalues, $E(\widetilde{F}_{t+1}) \;=\; \sum_{k=1}^{N} \lambda_k^s$. Therefore, when a researcher documents that a simple strategy $\widetilde{F}_{t+1}$ has significantly positive average returns, we learn that the sum of eigenvalues of Π^s is positive.

When all eigenvalues are non-negative, the simple factor is in fact optimal among all symmetric strategies. So, in this case, the simple strategy is not just simple — our analysis sheds new light on why it is a natural starting point.

When $E(\widetilde{F}_{t+1}) \;=\; \sum_{k=1}^{N} \lambda_k^s \;>\; 0$, the smallest eigenvalues can nevertheless be negative. Negative eigenvalues correspond to those surprising PEPs that are *negatively* predicted by their own signals.

When there exist PEPs with negative eigenvalues, we can beat the simple factor by leaving these PEPs out, buying only the PEPs that "work". Trading all the PEPs with positive eigenvalues is the optimal strategy among all linear strategies that always have positive factor exposure (i.e., among strategies with positive semi-definite L).

If we are willing to have a factor exposure that may switch sign, we can achieve an ever higher return. Indeed, negative eigenvalues also describe useful prediction patterns, just in the opposite direction. Therefore, an investor can do even better by also shorting the PEPs with negative eigenvalues, as shown in equation (31).

The between principal portfolio analysis and principal component analysis is remarkably close when we focus on the symmetric part of the prediction matrix as highlighted in Table 1. As seen in the table, PCA and PPA share five key properties. While PCA decomposes the variance into its components, PPA decomposes the expected excess return. Both have similar connections to eigenvalues, orthogonality, the trace, and optimality across orthonormal portfolios.

Example (Diagonal Prediction Matrix). Suppose there is no cross-predictability and signals are mean zero ($E(S_{j,t}) = 0$). Then $\Pi_{ij} = E(R_{i,t+1}S_{j,t}) = 0$ for all $i \neq j$. Hence, Π is symmetric, so there are no antisymmetric (zero exposure) strategies within Π. Furthermore, the PEPs are simply the unit vectors, $w_k^s = e_k$.[5] The optimal strategy is long assets with positive own-predictability and short those with negative own-predictability.

3.3 Antisymmetric Strategies: Principal Alpha Portfolios

We now turn to antisymmetric linear trading strategies. Our analysis relies on the eigendecomposition of an antisymmetric matrix, described in the next lemma.

Lemma 3 *Any antisymmetric matrix A has an even number $2K$ of non-zero eigenvalues. The non-zero eigenvalues are purely imaginary and come in complex-conjugate pairs: $i\lambda_k$ and $-i\lambda_k$. The corresponding orthonormal eigenvectors are $z_k = \frac{1}{\sqrt{2}}(x_k + iy_k)$ and the complex*

[5] Here, $e_k = (0, \ldots, 1, 0, \ldots, 0)'$, where 1 is in the k'th position.

conjugate $\bar{z}_k = \frac{1}{\sqrt{2}}(x_k - iy_k)$*, where* $x_k, y_k \in \mathbb{R}^N$ *with* $\|x_k\| = \|y_k\| = 1$*,* $x_k' y_k = 0$*, and* $x_k' x_l = x_k' y_l = y_k' y_l = 0$ *for all* $k \neq l,\ k,\ l \leq K \leq N/2$*. The corresponding eigendecomposition is given by*

$$A = \sum_{k=1}^{K} \lambda_k (x_k y_k' - y_k x_k'). \tag{33}$$

In other words, general antisymmetric matrices can be represented as a sum of building blocks that are each simple antisymmetric matrices with rank equal to two and having the form $xy' - yx'$. We refer to building blocks with form $xy' - yx'$ as *rank-2 antisymmetric strategies.*[6]

Each rank-2 building block generates a return of

$$S_t'(x_j y_j' - y_j x_j') R_{t+1} = S_t^{x_j} R_{t+1}^{y_j} - S_t^{y_j} R_{t+1}^{x_j}. \tag{34}$$

The first part of this portfolio is the return to trading the portfolio y_j based on the signal coming from the portfolio x_j. In other words, a strong signal for x_j ($S_t' x_j$), recommends scaling up the position in y_j ($y_j' R_{t+1}$), and this generates a return of $S_t^{x_j} R_{t+1}^{y_j}$. The second part is similar but flips the roles x_j and y_j and shorts the associated strategy (due to the minus sign). Thus, antisymmetric strategies are understandable as long-short strategies that trade two portfolios against each other based on the strength of each other's signal. But why does this result in zero conditional factor exposure, as guaranteed by Proposition 5? The next example helps develop intuition for the absence of factor risk in antisymmetric strategies.

[6] These are analogous to the rank-1 symmetric trading strategies, $L = ww'$, that are the basic building blocks of all symmetric trading strategies, as described in Section 3.2. An antisymmetric strategy satisfies the portfolio constraint, $\|A\| \leq 1$ as long as $|\lambda_k| \leq 1$ in (33).

Example (Beta-neutral Strategy). Consider an economy of N assets that satisfies the CAPM, save for asset 1, which has a positive alpha. That is, $E_t(R_{i,t+1}) = \alpha \mathbf{1}_{i=1} + \beta_{i,t}\theta_t$, where $\theta_t \geq 0$ is the market risk premium, $\beta_{i,t}$ is the conditional CAPM beta of stock i, and $\alpha > 0$. Suppose further that signals are defined to be the conditional betas, $S_{i,t} = \beta_{i,t}$. A standard beta-neutral strategy to exploit this scenario takes a long position in asset 1 having size equal to 1 (i.e., the size is set equal to the factor's beta on itself). The conditional beta from the long position is equal to $\beta_{1,t}$, so beta-neutrality is achieved with a position of $-\beta_{1,t} = -S_{1,t}$ in the factor. This strategy is a rank-2 antisymmetric strategy with $L = yx' - xy'$. The long position in asset 1 corresponds to $x = (1, 0, \cdots, 0)'$, and the short position in the factor corresponds to $y = (1, 1, ..., 1)'$. In other words, the beta-neutral strategy has zero symmetric component, non-zero antisymmetric component, and positive expected return, rendering it a pure alpha strategy:

$$E(S_t' L R_{t+1}) \; = \; E(\beta_t'(yx' - xy')R_{t+1}) \; = \; E\left(\sum_i \beta_{i,t} R_{1,t+1} - \beta_{1,t} \sum_i R_{i,t+1}\right) \; = \; \alpha E\left(\sum_{i=2}^{N} \beta_{i,t}\right),$$

which is positive as long as betas are positive on average. This is not the only pure alpha strategy, as a long position in asset 1 can be hedged with any other asset or combination of assets. Below, we show how to construct optimal pure alpha strategies using the eigendecomposition of Π^a.

The example illustrates that the fundamental $yx' - xy'$ structure underlying all antisymmetric strategies is closely related to the familiar approach to factor neutralization. To eliminate factor exposures, the position size in each must be equal to the factor exposure of the other, and with appropriately opposing signs.

Next, we derive *optimal* antisymmetric strategies. The first step is to apply the eigendecomposition in (33) to the antisymmetric part of the transposed prediction matrix, $(\Pi^a)'$. By Lemma 3, the matrix $(\Pi^a)'$ has $2N^a$ non-zero and purely imaginary eigenvalues, $i\lambda_k^a$ and

$-i\lambda_k^a$, for some $N^a \leq N/2$,. Their imaginary parts, $\lambda_k^a \in \mathbb{R}$, can be ordered as

$$\lambda_1^a \;\geq\; \cdots \;\geq\; \lambda_{N^a}^a \;\geq\; 0 \;\geq\; -\lambda_{N^a}^a \;\geq\; \cdots \;\geq\; -\lambda_1^a\,. \tag{35}$$

For each eigenvalue λ_j^a, we denote the corresponding real and imaginary parts of the eigenvectors by x_j and y_j, respectively.

We define the j^{th} **principal alpha portfolio (PAP)** as the linear strategy based on the j^{th} eigenvector: $L_j = x_j y_j' - y_j x_j'$. Equivalently, it has weights $(w_{j,t}^a)' = S_t'(x_j y_j' - y_j x_j')$ for $j = 1, ..., N^a$. We note that, since $N^a \leq N/2$, there exist at most $N/2$ principal alpha strategies. Moreover, Lemma 3 implies that PAPs are orthonormal.

The return of a principal alpha portfolio, like any rank-2 antisymmetric strategy, consists of two parts:

$$PAP_{t+1}^j = \; S_t'(x_j y_j' - y_j x_j')R_{t+1} \;=\; S_t^{x_j} R_{t+1}^{y_j} - S_t^{y_j} R_{t+1}^{x_j}\,. \tag{36}$$

The PAP buys portfolio y_j based on the signal coming from the portfolio x_j and simultaneously shorts portfolio x_j based on the signal from y_j.

Similar to the result for PEPs, we find that PAP expected returns are proportional to their eigenvalues and that the sum of PAPs is in fact the optimal antisymmetric linear trading strategy.

Proposition 8 *A principal alpha strategy has expected return $E(PAP_{t+1}^j) = 2\lambda_j^a$ and zero factor exposure. The sum of PAPs is the optimal antisymmetric linear strategy:*

$$\max_{\|L\|\leq 1,\; L=-L'} E(S_t' L R_{t+1}) = \sum_{k=1}^{N^a} E(PAP_{t+1}^k) = \sum_{k=1}^{N^a} 2\lambda_j^a. \tag{37}$$

The next example helps illustrate the properties of PEPs and PAPs.

Example (Constant Signals). Suppose that signals are constant over time, $S_t = S$.[7] In this case, the prediction matrix is especially simple, $\Pi = E(R_{t+1}S_t') = RS'$, where we use the short-hand notation $R := E(R_{t+1})$. We can now compute the PEPs and PAPs explicitly.

First, consider a case in which returns align with signals exactly, $R = S$. In this case, we have $\Pi = SS'$. This matrix is symmetric and has a rank of one. Hence, there is a single principal exposure portfolio with a non-zero eigenvalue, namely the eigenvector S, and no principal alpha portfolios. Therefore, this PEP is the only meaningful portfolio, and it is the same as the simple factor, S, with expected return $S'R = R'R > 0$.

Next, consider the case in which expected returns do *not* line up perfectly with the signal. Then $\Pi = RS'$ is no longer symmetric. The symmetric part is $\Pi^s = 0.5(RS' + SR')$, which has a rank of 2. Hence, Π^s has at most two non-zero eigenvalues, $\lambda_1^s = 0.5(R'S + \|R\|\,\|S\|) > 0 \geq \lambda_N^s = 0.5(R'S - \|R\|\,\|S\|)$ and the corresponding PEPs are[8]

$$w_1^s = c_1^s \left(\frac{R}{\|R\|} + \frac{S}{\|S\|} \right) , \; w_N^s = c_N^s \left(\frac{R}{\|R\|} - \frac{S}{\|S\|} \right) ,$$

where c_1^s, c_N^s are constants chosen such that $\|w_1^s\| = \|w_N^s\| = 1$. We see that the first principal exposure portfolio bets on securities with high average returns and high signals, while the last PEP bets on securities with high average returns and low signals. The negative eigenvalue PEP isolates losses due to the erroneous component of S and exploits them with a short position.

In this example, the prediction matrix also has an antisymmetric part. The strategy that trades this is $L = \Pi^{a\prime} = 0.5(SR' - RS')$. To derive the PAP, note that $\Pi^{a\prime}$ has at most two non-zero eigenvalues with purely imaginary parts $\lambda_1^a = 0.5(\|R\|\,\|S\| - R'S)^{1/2} \geq$

[7] As a concrete example, consider sorting stocks into value (book-to-market) deciles, using the decile portfolios as the baseline assets, and using a value signal defined as the decile number of each asset as the predictive signal. This is in contrast to, for example, forming assets as value-sorted portfolios, but using portfolio momentum as the trading signal. In this case signals are far from constant over time, and this is what we do empirically.

[8] These eigenvalues and eigenvectors can be verified by checking that $\Pi^s w_k^s = \lambda_k^s w_k^s$ for $k = 1, N$.

$0 \;\geq\; \lambda^a_N \;=\; -\,\lambda^a_1$ and the corresponding PAP is the linear strategy with position matrix $L = xy' - yx'$, where[9]

$$y \;=\; c^a \left(R\|S\|^2 - \; S(R'S)\right) \;,\; x = \; S/\|S\|.$$

The short part of the portfolio (x) is exactly the factor hedge. It is in place to ensure that the constraint (zero factor exposure) is satisfied. The remaining part of the problem is to find the highest average return subject to the constraint. Since the factor uses all (and only) the information in S, the remaining information that the PAP has at its disposal comes from the unconditional mean of returns. Thus the long side of the PAP (y) is determined by the information in R that is missed by S, hence the emergence in y of the difference between R and S.

Example (Betting Against Beta: PAP is the new BAB). Proceeding from the prior example, suppose that the erroneous signals S are chosen to be the expected returns in an asset pricing model, i.e. $S_j = \mathrm{Cov}(-M_t, R_{j,t})$ where M denotes the model's pricing kernel. Suppose further that there is less dispersion in true expected returns than predicted by the model—for simplicity, suppose that $\mathbf{1}'S/N = 1$ and suppose that $R = \mathbf{1}$. Then the alpha portfolio arising from the antisymmetric part of the prediction matrix has portfolio weight $w' = S'(SR' - RS') = (S'S)\mathbf{1}' - NS'$. This strategy goes long the equal-weighted portfolio (given by $R = \mathbf{1}$), while shorting the beta-weighted portfolio, S. To keep the portfolio beta-neutral, the equal-weighted portfolio (which is lower beta) is scaled up relative to the beta-weighted portfolio,[10] $S'S > N$. Hence, this strategy resembles the betting-against-beta (BAB) strategy of Frazzini and Pedersen (2014). This strategy has expected excess return of $w'R = NS'S - N^2 > 0$ and a beta of $w'S = NS'S - NS'S = 0$.

[9] c^a is determined such that $\|y\| = 1$

[10] This result follows from Cauchy-Schwarz, which yields that $N^2 = (\mathbf{1}'S)^2 \leq (\mathbf{1}'\mathbf{1})(S'S) = NS'S$, and the inequality is strict since we assume that betas vary across stocks.

In the preceding examples, signals are constant, which makes the math particularly tractable to illustrate intuitive aspects of principal portfolios. But constant signals imply that there are only static trading opportunities. In general, signals fluctuate over time, and principal portfolios use information about both static and dynamic trading opportunities. The prediction matrix can be written as a sum of its static and dynamic components:

$$\Pi = E(R_{t+1}S_t') = E(R_{t+1})E(S_t') + \text{Cov}(R_{t+1}, S_t'). \tag{38}$$

Suppose that signals do not predict future returns in the sense that $\text{Cov}(S_{i,t}, R_{j,t+1}) = 0$ for all i, j. In this case, Π simplifies to the constant signal example, $\Pi = E(R)\,E(S')$, and we have up to two PEPs and one PAP with strictly positive expected return, but these are purely based on the signals' time series average. The first term on the right side of equation (38) thus embodies information in the prediction matrix regarding "static bets."

The second term summarizes information in the prediction matrix regarding "dynamic bets." To focus purely on dynamic bets, then we can demean signals in the time series, looking at $\tilde{S}_{i,t} = S_{i,t} - E(S_{i,t})$. This redacts static information from Π and concentrates only on dynamic opportunities:

$$E(R_{t+1}\tilde{S}_t') = \text{Cov}(R_{t+1}, \tilde{S}_t') = \text{Cov}(R_{t+1}, S_t'). \tag{39}$$

Our approach allows both static and dynamic bets since both may be useful. Static bets are useful if they pick up that certain assets generally have higher returns, and if it's possible to time one's portfolio positions, then dynamic bets are profitable. We find in our empirical analysis that most of the effects we see are driven by dynamic bets.

To summarize, as the above examples illustrate, there are potentially two ways to earn alpha relative to the factor. The first stems from the observation that if Π^s has any negative eigenvalues, then shorting the corresponding PEPs yields a positive expected return with a negative factor exposure, which is alpha with respect to the factor. The second is to identify antisymmetric strategies with positive expected returns. Because an antisymmetric strategy is guaranteed to have zero factor exposure, it is also alpha to the factor.

4 Asset Pricing Tests: Positivity Bounds

We next propose a test for whether our signal S is an exposure (i.e., beta) to the true pricing kernel. Said differently, we wish to test whether the factor F corresponding to S is proportional to the true pricing kernel, $F_{t+1} \propto -M_{t+1}$ (or M's projection on the tradable space; recall that Lemma 1 shows how F is related to S). For example, we can consider signals given by betas to the market return, R^m_{t+1}, which corresponds to testing that the pricing kernel is of the form $M_{t+1} = a_t - b_t R^m_{t+1}$ for $a_t, b_t \in \mathbb{R}$ (i.e., the CAPM). Or, we can consider signals based on exposure to consumption, corresponding to testing that the pricing kernel is of the form $M_{t+1} = \beta u'(c_{t+1})/u'(c_t)$ (consumption CAPM).

Specifically, suppose that our signal $S_{i,t}$ is proportional to the exposure to the pricing kernel, $\text{Cov}_t(R_{j,t+1}, -M_{t+1})$, where we only assume proportionality (rather than equality) since we may not know the equity premium in the CAPM or the risk aversion in CCAPM. Then, signals should be closely related to expected returns. Indeed, the definition of a pricing kernel is a process M with $E_t((1 + R^f_t + R_{j,t+1})M_{t+1}) = 1$ for all assets, where R^f_t is the risk-free rate, which implies[11]

$$E_t(R_{j,t+1}) = (1 + R^f_t)\text{Cov}_t(R_{j,t+1}, -M_{t+1}) = \theta_t \, S_{j,t} \,, \tag{40}$$

where $\theta_t > 0$ is a factor of proportionality due to the risk-free rate and to our assumption that the signal S is proportional to (but not necessarily equal to) the covariance.

[11] To see this result, note that the definition of a pricing kernel applied for the risk-free asset (which has 0 excess return) yields $(1 + R^f_t)E_t(M_{t+1}) = 1$, which implies that $E_t(R_{j,t+1}M_{t+1}) = 0$ for excess returns. Therefore, $E_t(R_{j,t+1}) = (1 + R^f_t)E_t(M_{t+1})E_t(R_{j,t+1}) = (1 + R^f_t)(E_t(R_{j,t+1}M_{t+1}) - \text{Cov}_t(R_{j,t+1}, M_{t+1})) = (1 + R^f_t)\text{Cov}_t(R_{j,t+1}, -M_{t+1})$.

For example, if we are testing the CAPM, then the signal $S_{j,t}$ is typically the market beta, $\beta_{j,t} = \text{Cov}_t(R_{j,t+1}, R^m_{t+1})/\text{Var}_t(R^m_{t+1})$. In this case, the expected excess return is $E_t(R_{j,t+1}) = E_t(R^m_{t+1})\beta_{j,t}$, so here θ_t is the market risk premium, $E_t(R^m_{t+1})$. We would like to develop a test that does not require knowledge of θ_t because we may not know $E_t(R^m_{t+1})$ (or the coefficients a_t, b_t in $M_{t+1} = a_t - b_t R^m_{t+1}$).

The key insight is that, when the signal is proportional to the beta to the pricing kernel, the prediction matrix must be symmetric and positive definite—regardless of the factor of proportionality, θ. To see that, note that any off-diagonal element of the prediction matrix is

$$\Pi_{j,i} \;=\; E(S_{i,t}R_{j,t+1}) \;=\; E(S_{i,t}E_t(R_{j,t+1})) \;=\; E(\theta_t S_{i,t}S_{j,t}) \;=\; \Pi_{i,j}\,. \tag{41}$$

which shows that Π is symmetric. Further, we see that the prediction matrix is positive semi-definite since, for any $w \in \mathbb{R}^N$:

$$w'\Pi w = w'E(\theta_t S_t S_t')w = E(\theta_t[w'S_t]^2) > 0 \tag{42}$$

This finding provides new asset pricing tests as summarized here:

Proposition 9 (Positivity of Prediction Matrix) *If there exists $\theta_t \geq 0$ such that*

$$E(R_{i,t+1}|\theta_t, S_t) \;=\; \theta_t S_{i,t} \tag{43}$$

for all i, then the corresponding prediction matrix Π is symmetric and positive semi-definite, and, equivalently, all the corresponding PEPs have non-negative expected returns and all PAPs have zero expected returns. The premise (43) holds, for example, if there is no arbitrage so a pricing kernel exists, and the signal $S_{i,t}$ is proportional to exposures to the pricing kernel.

The intuition behind this result follows from our earlier portfolio theory: We know that negative eigenvalues of Π^s and a non-zero Π^a give rise to alpha strategies (Sections 3.2 and 3.3, respectively). Since alpha strategies cannot exist in a rational asset pricing model, all eigenvalues of Π^s must be positive and Π^a must be zero. In other words, Π must be symmetric and positive semi-definite.

One benefit of this approach is that we do not need to know θ_t, we just need to observe signals and returns, and then consider the positivity of the corresponding prediction matrix. Another helpful feature is that the test is unconditional, i.e., it relies on an unconditional expected value, $\Pi = E(R_{t+1}S_t')$, even if the underlying asset pricing model in conditional. Hence, while some tests require an understanding of how the risk premium varies over time or make assumptions to get from a conditional CAPM to an unconditional test, we have a test of the conditional CAPM (and other conditional models) based on an unconditional moment condition. Further, this restriction also tests cross-asset effects.

5 Robust Strategies: Shrinkage via Principal Portfolios

Our theoretical analysis up to now has taken place in population with the prediction matrix, $\Pi = E(R_{t+1}S_t')$, known. In reality, Π is unknown and must be estimated. Unfortunately, this is a highly parameterized framework; it requires estimating N^2 parameters. The standard tradable factor approach from the literature (3) essentially restricts the set of linear strategies to a single parameter problem—i.e., signals are typically assessed based only on their average own-predictability $\sum_i E(S_{i,t}R_{i,t+1})$. This approach can be viewed as a regularization device that exploits a signal while imposing many restrictions to minimize the number of parameters. But these restrictions may be unnecessarily severe. They sacrifice any and all useful information about heterogeneity in own-predictability (differences among diagonal elements of Π) or cross-predictability (off-diagonal elements).

Principal portfolios are ideally suited to balance two considerations: 1) exploiting potentially rich information from throughout the predictability matrix, and 2) controlling parameterization to reduce overfit and ensure robust out-of-sample portfolio performance. In this section, we develop robust principal portfolio trading strategies by shrinking the predictability matrix.

The analysis in Sections 2 and 3 shows that a singular value decomposition of Π (or of its symmetric and antisymmetric parts) finds orthonormal portfolios and orders them from highest expected return to lowest. This eigendecomposition has another great benefit in that it lends itself naturally to a convenient form of regularization. In particular, if we reconstitute the Π matrix by retaining only the K largest singular values and zeroing out the rest, we obtain the matrix of rank K that is as close as possible to the original Π. This idea is familiar from principal components analysis, which finds low-rank approximations to a variance-covariance matrix by zeroing out all but its largest eigenvalues.

The following proposition operationalizes the idea of robust optimal trading strategies by constraining the parameter space to position matrices with $\operatorname{rank}(L) \le K$. Here, K is a tuning parameter that can be chosen empirically.

To add further generality and another convenient tuning parameter, we introduce the Schatten p-norm for a matrix L (see, Horn and Johnson (1991)):

$$\|L\|_p \;=\; \left(\sum_{k=1}^{N} |\bar{\lambda}_k(L)|^p\right)^{1/p},$$

where $\bar{\lambda}_k(L)$ is the k-th singular value of L and $p \in [1,\infty]$. The limiting case $p = \infty$ corresponds to the standard matrix norm $\|L\| = \|L\|_\infty$, whereas $p = 2$ corresponds to the sum of squares of all elements $\|L\|_2 \;=\; (\sum_{k,l} L_{l,k}^2)^{1/2}$ (Frobenius norm). Interestingly, we show that different matrix norms correspond to different ways of weighting the principal portfolios.

We are ready to state a result that generalizes all the optimization problems that we considered so far (Propositions 3, 6, and 8).

Proposition 10 (General Solution) *Optimal portfolios subject to* $\text{rank}(L) = K$ *and* $||L||_p \leq 1$, *where* $p = [1, \infty]$ *and* q *is defined by* $1/p + 1/q = 1$, *satisfy:*

1. *The solution with no symmetry constraints depends on the top* K *singular values,* $\bar{\lambda}_k$, *of* Π*:*

$$\max_{\|L\|_p \leq 1,\ \text{rank}(L) \leq K} E(S_t' L R_{t+1}) = \left(\sum_{k=1}^{K} \bar{\lambda}_k^q \right)^{1/q}. \tag{44}$$

   The optimal L *is* $S_t' L R_{t+1} = c \sum_{k=1}^{K} \bar{\lambda}_k^{q-1} PP_{t+1}^k$, *where* $c = \left(\sum_{k=1}^{K} \bar{\lambda}_k^q \right)^{-1/p}$.

2. *The solution when restricting to symmetric strategies depends on the set* $\mathcal{K}$ *of the* K *largest absolute eigenvalues* $|\lambda_k^s|$ *of* Π^s*:*

$$\max_{\|L\|_p \leq 1,\ \text{rank}(L) \leq K,\ L = L'} E(S_t' L R_{t+1}) = \left(\sum_{k \in \mathcal{K}} |\lambda_k^s|^q \right)^{1/q}. \tag{45}$$

   The optimal L *is* $S_t' L R_{t+1} = c \sum_{\mathcal{K}} |\lambda_k^s|^{q-1} \,\text{sign}(\lambda_k^s) PEP_{t+1}^k$, *where* $c = (\sum_{\mathcal{K}} |\lambda_k^s|^q)^{-1/p}$.

3. *The solution when restricting to antisymmetric strategies depends on the eigenvalues* λ_k^a *of* Π^a*:*

$$\max_{\|L\|_p \leq 1,\ \text{rank}(L) \leq 2K,\ L = -L'} E(S_t' L R_{t+1}) = \left(2 \sum_{k=1}^{K} (\lambda_k^a)^q \right)^{1/q}. \tag{46}$$

   The optimal L *is* $S_t' L R_{t+1} = c \sum_{k=1}^{K} (\lambda_k^a)^{q-1} PAP_{t+1}^k$, *where* $c = \left(2 \sum_{k=1}^{K} (\lambda_k^a)^{q-1} \right)^{-1/p}$.

Proposition 10 shows that optimal low-dimensional trading strategies are the same as the general optimality results proven earlier, with the exception that the strategies use only the leading principal portfolios. This is true regardless of whether one considers general linear strategies (L), symmetric and hence factor-exposed strategies ($L = L'$), or antisymmetric pure alpha strategies ($L = -L'$). By truncating the strategy at the top K principal portfolios, these robust strategies replace the lesser singular values with zeros.[12] The lesser components may be dominated by noise, and are therefore likely to have poor out-of-sample performance. Zeroing them out regularizes the optimal strategy to controls overfit and its adverse out-of-sample impact. The number of principal portfolios included in a robust strategy, K, determines the extent of regularization. It serves as a hyperparameter that can be controlled by the researcher or tuned via cross-validation.

What are the implications of the more general norm $\|\cdot\|_p$ in this proposition, and what economic role does it play? Proposition 10 shows that the optimal strategy is a weighted sum of principal portfolios for any norm. This result shows that PPs are very general building blocks. The choice of norm simply affects how the principal portfolios are weighted, which also illustrates the connection between the tuning parameters p and K: The less important PPs can be "zeroed out" by the choice of K and down-weighted by the choice of p.

At that same time, the norm constraint captures the idea of constraining trading strategy leverage in the optimization problem has a natural economic motivation—risk and institutional frictions impose leverage considerations on every real-world investor. The way real-world investors try to manage their leverage concerns is dictated in part by the performance of strategies in their opportunity set. This raises an interesting practical implication of Proposition 10. The norm exponent p can be treated as a hyperparameter that can be tuned via cross-validation. An investor that tunes p along with K in essence chooses the form of leverage constraint that lends itself to robust out-of-sample trading performance.

Interestingly, when $p = 2$, part 1. of the proposition is similar to trading a version of the Π matrix that has been estimated via a reduced rank regression (see, e.g., Velu and Reinsel (1998)).[13] Further, when $p = 2$ and we do not impose a rank restrictions (i.e., we

[12]Note that singular values of a symmetric or an anti-symmetric matrix coincide with the absolute values of its eigenvalues.

[13]Reduced rank regression (RRR) seeks to minimize the mean squared error $E(\|R_{t+1} - L'S_t\|^2) = E(\|R_{t+1}\|^2) - E(S_t'LR_{t+1}) + E(S_t'LL'S_t)$ under a rank constraint on the matrix L. By direct calculation, this objective is equivalent to maximizing $\mathrm{tr}(L\Pi) - \mathrm{tr}(LL'\Sigma_S)$. Thus, reduced rank regression amounts to maximizing the expected return, $\mathrm{tr}(L\Pi)$, with a punishment term for signal variance. If $\Sigma_S = \mathrm{Id}$, the

let $K = N$), then the solution is $L = \Pi'/\|\Pi\|_2$. So, in this case, we uncover the prediction matrix itself as the optimal strategy. For $p = \infty$, i.e. $q = 1$, the solution selects components that are large in absolute value, in the spirit of lasso applied to singular values, and with no rank restriction we recover Proposition 4.

The results in Sections 2 through 4 lay out a theoretical basis for principal portfolios, and Proposition 10 prescribes a machine learning approach to implementing principal portfolios in practice. Data-driven choices for hyperparameters K and p can allow the researcher to select the level of principal portfolio model complexity best suited for constructing optimal out-of-sample strategies.

6 Empirical Results

We next present a simple empirical implementation of our method using some of the most standard data sets in finance.

6.1 Data, Signals, and Methodology

We must make a number choices to empirically implement our framework. We present a "base-case" set of choices using relatively standard methods and consider several variations around the base case for robustness. The base-case data is the 25 Fama-French portfolios, a standard data set in finance. These portfolios are constructed by double-sorting U.S. stocks by their size (measured by market capitalization) and valuation ratio (book-to-market).[14] We compute alphas based on the 5-factor model of Fama and French (2015). We use daily data from July 1963 until the end of 2019.

For robustness, we also implement our model for several other data sets. In particular, we also run the model for 25 U.S. size and operating profitability portfolios, 25 U.S. size

punishment term coincides with $\|L\|_2^2$, and hence RRR is a modification of the problem solved in Proposition 10 for $p = 2$.

[14] https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

and investment portfolios, the international counterparts of the three sets of Fama-French portfolios (i.e., developed countries excluding the U.S.), and a sample of 52 futures contracts starting in 1985. This includes contracts for 21 commodities, 17 equity indices, 5 sovereign bonds, and 9 currencies.

For each data set, we need signals and returns. Starting with signals, the base case signal is each asset's 20-day momentum (approximately one month). That is, for each asset in each sample, we compute its past 20-day cumulative return, then standardize the signal each period by converting it to a cross-sectional rank and dividing by the number of assets and subtracting the mean (mapping the signal into the [-0.5,0.5] interval).[15] We also consider other momentum signals for robustness, namely based on 40, 60, 90, 120, and 250 day past returns, following on the standard practice of considering momentum signals up to 1 year (approximately the same as 250 trading days).

Turning to returns, the base-case measure of returns is each asset's 20-day return (again, one month). Specifically, we divide the sample into non-overlapping 20-day time periods, denoted by t, and, in each time period t, we seek to predict the future 20-day returns based on the momentum signals. We cross-sectionally demean returns to focus prediction on cross section differences in returns rather than time series fluctuations in the common market component of returns. We also consider other forecast horizons, namely 1-day, 5-day, and 10-day returns. Similarly to the base case, each of the other forecast horizons correspond to a sample of non-overlapping time periods t of the same length. As advocated by Moskowitz et al. (2012), in the case of futures contracts we time-series de-volatize returns (both in the forecast target and in the momentum signal construction) using trailing 20-day return

[15] All of our theoretical results apply to cross-sectionally demeaned signals. If we start with any signal S, we can work with the cross-sectionally demeaned signal: $\tilde{S}_{j,t} = S_{j,t} - \frac{1}{N}\sum_{k=1}^{N} S_{k,t}$. The corresponding simple factor $\tilde{F}$ is dollar neutral. The eigenvalues of the prediction matrix with respect to $\tilde{S}$ and S have the same signs, except for at most two eigenvalues (see Proposition 11 in the Appendix). Further, demeaning means that we only exploit cross-sectional predictability, not time series predictability, which essentially leads to the "loss" of one eigenvalue (Proposition 12 in the Appendix).

volatility, which helps avoid a situation in which results are unduly driven by the large cross-sectional differences in volatility of raw futures contract returns.[16]

We estimate the prediction matrix as the sample counterpart of the definition $\Pi = E(R_{t+1}S_t')$ using a rolling "training window." The training window is the past 120 time periods. For example for the base case, the training period consists of the past 120 non-overlapping 20-day time periods. The estimated prediction matrix at time period t is

$$\hat{\Pi}_t = \frac{1}{120} \sum_{\tau=t-120}^{t-1} R_{\tau+1} S_\tau' \,. \tag{47}$$

Based on this empirical prediction matrix, we compute its singular vectors to form principal portfolios (PPs) and we compute and the eigenvectors of its symmetric and antisymmetric parts, giving rise to the empirical principal exposure portfolios (PEPs) and principal alpha portfolios (PAPs). We compare these to the simple factor $\widetilde{F}_t$ defined in (3). To limit the undue effects of illiquidity on our conclusions, we always add an extra 1-day buffer between the last day in the training sample and the first day in the forecast window.

6.2 The Prediction Matrix and Principal Portfolios

We first consider the PPs, PEPs, and PAPs for the base-case sample of 20-day returns using 20-day momentum signals for the 25 Fama-French size-value portfolios. Figure 1.A shows the singular values of the prediction matrix, averaged over time. Recall that, according to the theory, these singular values correspond to the expected returns of the corresponding PPs. The realized next-month returns (i.e., out of sample) of the PPs are plotted in Figure 1.D, along with their confidence bands. We find that the realized returns roughly match the shape of the ex ante singular values, with the low-numbered PPs having large eigenvalues and high realized returns. However, while this relation would be perfect on an in-sample basis (not

[16]This adjustment has tiny effects in our equity asset analysis so, in the interest of simplicity, we do not de-volatilize equity returns.

shown), we naturally see some degradation of realized returns relative to the eigenvalues when looking out of sample.

In a similar vein, Figures 1.B and 1.C show the eigenvalues of the symmetric and antisymmetric parts of the prediction matrix, respectively. Figures 1.E and 1.F report the realized returns of the corresponding PEPs and PAPs, respectively. Again we see a close relation between the ex ante predicted returns, and the out-of-sample realized ones. In this sample, only the first two PPs and first two PEPs appear to have a significant out-of-sample return, and only the first PAP return is significant.

One might wonder what these portfolios look like? We explore this in the case of PEPs and PAPs. Figure 2.A plots the weights of the eigenvector w_1 underlying the first PEP. Interestingly, this eigenvector tends to be long value versus shorting growth stocks, and simultaneously tends to be long larger stocks versus short smaller ones. Recall that PEP1 trades w_1 based its on signal, that is, PEP1 is going long or short a size-value bet based on its own momentum. Said differently, when large-value has recently outperformed, then PEP1 buys large-value, and, otherwise, it buys small-growth. To illustrate this strategy further, Figure 2.B plots the momentum, $S'w_1$, of the eigenvector. Lastly, Figure 2.C show the overall portfolio weight, $S'w_1w_1'$, averaged over time. Similarly, Figure 2, Panels D–E illustrate the PAP1 trading strategy.

6.3 PP, PEP, and PAP Returns across Forecast Horizons

Figure 3 plots the performance of the PPs, PEPs, and PAPs across several forecast horizons. In addition to the base-case specification with 20-day return periods considered above, we also consider 1-day, 5-day, and 10-day forecast horizons. In all cases, the signal is the past 20-day momentum. For simplicity, we only report the return of the sum of the top three principal portfolios (among each the PPs, PEPs, and PAPs), and the combination of the top 3 PEPs plus top 3 PAPs.[17] In each case, we compare their performance to that of the simple factor, which is just the sum-product of signals and returns. When analyzing factor performance, we use the exact same signal construction for the factor and PPs and evaluate both over the same forecast horizons, so each group of bars is an apples-to-apples comparison.

The five bars on the right in Figure 3.A show performance for the base case with 20-day returns. We see that the PEP has a similar Sharpe ratio (SR) to that of the simple factor, where SR is the average excess return divided by volatility. The PAP has a higher SR, and the combination of PEP and PAP is higher yet, more than double the SR of the simple factor. The PP strategy performs similarly to PAP, handily beating the simple factor. The best overall performance is achieved by the combination of PEPs and PAPs.

Figure 3.B plots the information ratio (IR) and its confidence interval as a measure of the risk-adjusted return of the principal portfolios. Specifically, the IR is computed by regressing the return of the PP (or PEP, PAP, or their combination) on the simple factor ($\widetilde{F}$) and the five Fama-French factors (the market MKT, the size factor SMB, the value factor HML, the profitability factor RMW, and the investment factor CMA):

$$PEP_t = \alpha + \beta^0 \widetilde{F}_t + \beta^1 MKT_t + \beta^2 SMB_t + \beta^3 HML_t + \beta^4 RMW_t + \beta^5 CMA_t + \varepsilon_t \quad (48)$$

The IR is the alpha divided by residual volatility, IR$= \alpha/\sigma(\varepsilon_t)$, which can be interpreted as the Sharpe ratio when all the factors on the right hand side are hedged out (i.e., the alpha expressed as a Sharpe ratio).

Table 2 reports the details of this regression. As seen from Table 2 (and the confidence intervals in Figure 3.B), the PEP does not have a significant alpha (or, equivalently, a significant IR), but the PAP is highly significant (t-statistic of 4.42) and so is the PP strategy

[17]When combining PEPs and PAPs, we rescale the PAP component to have the same volatility as the PEP component, then take a 50/50 combination.

and the combination of PEP and PAP. Interestingly, Table 2 also shows that PEP has a highly significant loading on the simple factor with a high R^2, while, in contrast, PAP has small and insignificant factor loadings and low R^2. These findings are consistent with the idea that PEP provides factor exposure while PAP provides uncorrelated alpha.

Finally, Figure 3 shows that the principal portfolios perform even better at shorter forecast horizons, especially the PP and PEP strategies. Indeed, at the shorter forecast horizons, even PEP earns a higher SR than the simple factor, and the risk-adjusted return as measured by the IR becomes highly significant at 1-, 5-, and 10-day forecast horizons.

6.4 Other Samples and Markets

We next implement the model in other samples. In particular, we consider three samples of U.S. stocks (the base case from before, plus two other sets of Fama-French portfolios), three sets of international stocks (i.e., global stocks outside the U.S. sorted into similar portfolios), and a set of 52 futures contracts (consisting of equity index futures, bond futures, commodity futures, and currency forwards).

Figure 4.A and B show the Sharpe ratios and information ratios for these seven data sets. In support of the model's predictive power, we see that all of the SRs and IRs are positive, and several, but not all, are statistically significant. In further support of the model, Figure 4.C shows that the ex ante eigenvalues are highly correlated to the ex post realized returns in each sample.

6.5 Robustness across Momentum Horizons and Sub-samples

Finally, we analyze the robustness of our method across momentum horizons and sub-samples. Figure 5 shows the performance of the leading PPs, PEPs, and PAPs for different look-back periods in the specification of the momentum signal. Panel A shows that the PEP performs similarly to the simple factor for all momentum horizons. However, PP, PAP, and the PEP/PAP combination deliver higher SR across all horizons. Turning to the information ratios in Panel B, we see the these are more statistically significant for the short- and medium-term momentum periods, and less significant for the longer-term momentum horizons.

Finally, Figure 6 reports the performance of the base-case strategy for each decade in the sample. We see that the performance tends to be positive across decades — so the strong overall performance is not being driven by a single decade — but, naturally, the statistical significance in each decade is reduced due to the short time window.

7 Conclusion: The Power of Principal Portfolio Analysis

We present a new method to analyze return predictability and asset pricing tests. The method provides novel intuitive portfolios that can generate effective factor exposures and alpha strategies relative to the factor. We implement our method empirically using a range of standard data sets and find significant evidence consistent with our predictions.

References

Arnott, Robert D, Mark Clements, Vitali Kalesnik, and Juhani T Linnainmaa, "Factor momentum," *Available at SSRN 3116974*, 2019.

Asness, Clifford S, Tobias J Moskowitz, and Lasse Heje Pedersen, "Value and momentum everywhere," *The Journal of Finance*, 2013, *68* (3), 929–985.

Bhatia, Rajendra, *Matrix Analysis*, Springer, 1997.

Cochrane, John H, *Asset pricing: Revised edition*, Princeton university press, 2009.

Collin-Dufresne, Pierre, Kent D Daniel, Ciamac C Moallemi, and Mehmet Sağlam, "Dynamic asset allocation with predictable returns and transaction costs," *Available at SSRN 2618910*, 2015.

___ **, Kent Daniel, and Mehmet Sağlam**, "Liquidity regimes and optimal dynamic asset allocation," *Journal of Financial Economics*, 2019.

Fama, Eugene F and Kenneth R French, "A five-factor asset pricing model," *Journal of financial economics*, 2015, *116* (1), 1–22.

Fan, Ky, "On a theorem of Weyl concerning eigenvalues of linear transformations.," *Proceeding of the National Academy of Sciences USA*, 1950, *36*, 31–35.

___ **and A. J. Hoffman**, "Some metric inequalities in the space of matrices," *Proceedings of the American Mathematical Society*, 1955, *6*, 111–116.

Frazzini, Andrea and Lasse Heje Pedersen, "Betting against beta," *Journal of Financial Economics*, 2014, *111* (1), 1 – 25.

Gârleanu, Nicolae and Lasse Heje Pedersen, "Dynamic Trading with Predictable Returns and Transaction Costs," *The Journal of Finance*, 2013, *68* (6), 2309–2340.

___ **and** ___ , "Dynamic Portfolio Choice with Frictions," *Journal of Economic Theory*, 2016, *165*, 487–516.

Gibbons, Michael R, Stephen A Ross, and Jay Shanken, "A test of the efficiency of a given portfolio," *Econometrica: Journal of the Econometric Society*, 1989, pp. 1121–1152.

Gu, Shihao, Bryan Kelly, and Dacheng Xiu, "Empirical asset pricing via machine learning," Technical Report, National Bureau of Economic Research 2018.

Gupta, Tarun and Bryan Kelly, "Factor momentum everywhere," *The Journal of Portfolio Management*, 2019, *45* (3), 13–36.

Hansen, Lars Peter and Ravi Jagannathan, "Implications of security market data for models of dynamic economies," *Journal of political economy*, 1991, *99* (2), 225–262.

Horn, Toger A. and Charles R. Johnson, *Topics in Matrix Analysis*, Cambridge, 1st Edition, 1991.

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

Kelly, Bryan, Seth Pruitt, and Yinan Su, "Characteristics are Covariances: A Unified Model of Risk and Return," *Journal of Financial Economics*, 2020.

___ , ___ , and ___ , "Instrumented Principal Component Analysis," *Working paper*, 2020.

Marshall, Albert W. and Ingram Olkin, *Inequalities: Theory of Majorization and Applications*, Academic, 1979.

Moskowitz, Tobias J, Yao Hua Ooi, and Lasse Heje Pedersen, "Time series momentum," *Journal of financial economics*, 2012, *104* (2), 228–250.

Pedersen, Lasse Heje, Abhilash Babu, and Ari Levine, "Enhanced Portfolio Optimization," *Working paper, AQR and Copenhagen Business School*, 2020.

Tao, Terrence, *Topics in Random Matrix Theory*, Springer, 2012.

Velu, Raja and Gregory C. Reinsel, *Multivariate Reduced-Rank Regression. Theory and Applications*, Springer, 1998.

Welch, Ivo and Amit Goyal, "A comprehensive look at the empirical performance of equity premium prediction," *The Review of Financial Studies*, 2008, *21* (4), 1455–1508.

A On Mean-Variance Optimization of Linear Strategies

A.1 Robust Mean-Variance Framework of Linear Strategies

Consider the following robust mean-variance objective function:

$$\max_L E(S_t' L R_{t+1}) \text{ subject to } \max_{S:\|S\|\leq 1} \text{Var}_t(S' L R_{t+1}) \leq 1 \tag{49}$$

This objective is a robust in the sense that we require that the variance is bounded regardless of the signal S. The variance term can be written as

$$\max_{S:\|S\|\leq 1} \text{Var}_t(S' L R_{t+1}) = \max_{S:\|S\|\leq 1} S' L \Sigma_{R,t} L' S = \max_{S:\|S\|\leq 1} \|\Sigma_{R,t}^{1/2} L' S\|^2 = \|\Sigma_{R,t}^{1/2} L'\|^2 \tag{50}$$

So, if the assets are normalized and uncorrelated such that $\Sigma_{R,t} = \text{Id}$, then the robust variance constraint is the same as our matrix constraint $\|L\| \leq 1$.

We also get a similar solution when $\Sigma_{R,t} = \sigma^2 \text{Id}$ for some $\sigma \in \mathbb{R}$, since, in this case, the portfolio constraint simply becomes $\|L\|\sigma \leq 1$. Hence, we just scale the position accordingly (i.e., use $L = L^*/\sigma$, where L^* is the standard solution with $\sigma = 1$). We note that, when $\Sigma_{R,t} = \sigma^2 \text{Id}$, an alternative constraint is to require that the risk of the linear strategy cannot be greater than the risk of the simple factor as we do in (10), which is, again, equivalent to $\|L\| \leq 1$.

We can also easily solve a version of the robust mean-variance problem for general $\Sigma_{R,t}$. Solving (10) directly is not convenient since it would lead us to impose $\|\Sigma_{R,t}^{1/2} L'\| \leq 1$, which is not consistent with choosing a constant position matrix L. (Recall that the idea of linear strategies is to have a constant L, but rich portfolio dynamics, $S_t' L$, driven by the signals.) Instead, we consider the transformed "synthetic assets" with returns

$$\tilde{R}_{t+1} = \Sigma_{R,t}^{-1/2} R_{t+1}\,, \tag{51}$$

signals given by

$$\tilde{S}_t = \Sigma_{R,t}^{-1/2} S_t \,, \tag{52}$$

and the corresponding prediction matrix is $\tilde{\Pi} = E(\tilde{R}_{t+1}\tilde{S}_t') = E(\Sigma_{R,t}^{-1/2}\Pi\Sigma_{R,t}^{-1/2})$. For example, if the original assets are uncorrelated (i.e., a diagonal $\Sigma_{R,t}$), then creating these synthetic assets simply means scaling the assets to have the same constant volatility (as we do in the empirical analysis of futures).

Then we consider the robust mean-variance problem for the synthetic assets:

$$\max_{\tilde{L}} \; E\left(\tilde{S}_t'\tilde{L}\tilde{R}_{t+1}\right) \;\text{ subject to }\; \max_{\tilde{S}_t} \mathrm{Var}_t(\tilde{S}_t'\tilde{L}\tilde{R}_{t+1}) = \|\tilde{L}\| \leq 1. \tag{53}$$

So we see that this is our standard problem, expressed in terms of the synthetic assets. In other words, all our results apply for the robust mean-variance problem of the synthetic assets. For example, the optimal strategy for the synthetic assets is $\tilde{L} = (\tilde{\Pi}'\tilde{\Pi})^{-1/2}\tilde{\Pi}'$, using Proposition 3. Of course, this solution can be translated back to the original assets by noting that

$$L_t = \Sigma_{R,t}^{-1/2}\tilde{L}\Sigma_{R,t}^{-1/2}, \tag{54}$$

which holds since we must have that $\tilde{S}_t'\tilde{L}\tilde{R}_{t+1} = S_t'{\Sigma_R}^{-1/2}\tilde{L}{\Sigma_R}^{-1/2}R_{t+1} = \; S_t'L_tR_{t+1}$.

A.2 Risk Aversion instead of Risk Constraint

We can also consider a robust mean-variance problem in which the investor has risk aversion γ rather than a risk constraint:

$$\max_{L}\left(E(S_t'LR_{t+1}) - \frac{\gamma}{2}\left[\max_{S:\|S\|\leq 1}\mathrm{Var}_t(S'LR_{t+1})\right]\right)\,. \tag{55}$$

We can rewrite the objective using (50) as

$$\max_L \left(E(S_t' L R_{t+1}) - \frac{\gamma}{2} \|\Sigma_{R,t}^{1/2} L\|^2 \right) \tag{56}$$

Let us first solve this portfolio problem when $\Sigma_{R,t} = \mathrm{Id}$. To find the solution, we start by finding the solution for each level of volatility:

$$L(c) = \arg \max_{L:\|L\|=c} \; E\left(S_t' L R_{t+1}\right) = c(\Pi'\Pi)^{-1/2}\Pi'. \tag{57}$$

where the last equality uses Proposition 3. Now we can solve the objective function (60) by maximizing over all possible volatilities, c:

$$\max_c \; \left(E\left(S_t' L(c) R_{t+1}\right) - \frac{\gamma}{2} c^2 \right) \; = \; \max_c \left(c \sum_{i=1}^N \bar{\lambda}_i - \frac{\gamma}{2} c^2 \right) \;\; = \;\; \frac{(\sum_{i=1}^N \bar{\lambda}_i)^2}{2\gamma}\,, \tag{58}$$

where the optimum is achieved by $c = \frac{\sum_{i=1}^N \bar{\lambda}_i}{\gamma}$, implying that the optimal strategy is

$$L = \frac{\sum_{i=1}^N \bar{\lambda}_i}{\gamma} (\Pi'\Pi)^{-1/2}\Pi'. \tag{59}$$

We see that the optimal strategy is the same as in Proposition 3, except for a scaling factor. The scaling factor naturally decreases in γ, reflecting that a more risk averse investor takes a smaller position. Similarly, the scaling factor increases in the sum of the singular values, since higher singular values imply stronger predictability, leading to a larger position.

Finally, consider the problem with a general variance-covariance matrix $\Sigma_{R,t}$. Rather than solving the objective function (60) (which could lead to a time-varying L), we consider the similar objective function for the synthetic assets:

$$\max_{\tilde{L}} \left(E(\tilde{S}_t' \tilde{L} \tilde{R}_{t+1}) - \frac{\gamma}{2} \left[\max_{\tilde{S}:\|\tilde{S}\|\le 1} \mathrm{Var}_t(\tilde{S}' \tilde{L} \tilde{R}_{t+1}) \right] \right) = \max_{\tilde{L}} \left(E(\tilde{S}_t' \tilde{L} \tilde{R}_{t+1}) - \frac{\gamma}{2} \|\tilde{L}\|^2 \right) \tag{60}$$

Since the synthetic assets have a variance-covariance matrix equal to the identity, their optimal solution is

$$\tilde{L} = \frac{\sum_{i=1}^{N} \bar{\lambda}_i(\tilde{\Pi})}{\gamma} (\tilde{\Pi}'\tilde{\Pi})^{-1/2}\tilde{\Pi}' \,, \tag{61}$$

which can be translated back to the original assets using (54).

A.3 Relation to the Standard Mean-Variance Framework

We next consider a standard mean-variance objective function, that is, we don't use the worst-case variance, but instead assume that conditional return variance is constant over time, $\Sigma_{R,t} = \Sigma_R$. In this case, we naturally recover the standard Markowitz solution. To see this, we use the notation $\Sigma_S = E(S_t S_t')$ as before and calculate:

$$\begin{aligned}
\max_L E\left(E_t(S_t' L R_{t+1}) - \frac{\gamma}{2}\mathrm{Var}_t(S_t' L R_{t+1})\right) &= \max_L E\left(S_t' L R_{t+1} - \frac{\gamma}{2} S_t' L \Sigma_R L' S_t\right) \\
&= \max_L \left(\mathrm{tr}(L\Pi) - \frac{\gamma}{2}\,\mathrm{tr}(L\Sigma_R L \Sigma_S)\right) \\
&= \max_{\tilde{L}} \left(\mathrm{tr}(\tilde{L}\tilde{\Pi}) - \frac{\gamma}{2}\,\mathrm{tr}(\tilde{L}'\tilde{L})\right) \\
&= \max_{\tilde{L}} \left(\mathrm{tr}(\tilde{L}\tilde{\Pi}) - \frac{\gamma}{2}\|\tilde{L}\|_2^2\right) \,,
\end{aligned} \tag{62}$$

where we use the change of variable $\tilde{L} = \Sigma_S^{1/2} L {\Sigma_R}^{1/2}$ and $\tilde{\Pi} = {\Sigma_R}^{-1/2}\Pi\Sigma_S^{-1/2}$. So we see that this problem has the same form as our normal objective function, except that we have another matrix norm, namely the Frobenius 2-norm, $\|\cdot\|_2$. Since we solve the problem for all p-norms in Proposition 10 (and in the proof of Propostion 3), we know that solution, which is very simple: $\tilde{L} \ = c\tilde{\Pi}'$, where c is a constant that depends on the risk aversion γ. So the solution to the mean-variance problem using the original variables is

$$L = \Sigma_S^{-1/2}\tilde{L}{\Sigma_R}^{-1/2} = c\Sigma_S^{-1/2}\tilde{\Pi}'{\Sigma_R}^{-1/2} = c\Sigma_S^{-1}\Pi{\Sigma_R}^{-1} \tag{63}$$

In other words, the optimal portfolio is

$$w_t = L'S_t = c{\Sigma_R}^{-1}\Pi\Sigma_S^{-1}S_t = \; c\;{\Sigma_R}^{-1}E(R_{t+1}|S_t) \tag{64}$$

which is the standard Markowitz tangency portfolio (scaled by c depending on risk aversion). The last equality assumes that the conditional expected return can be computed using the multivariate regression of R_{t+1} on S_t, that is, $R_{t+1} = AS_t + \varepsilon_{t+1}$, and uses that the regression coefficient is $A = \Pi\Sigma_S^{-1}$.

So while our framework can nest the standard Markowitz solution, we seek to add robustness in several ways. First, we introduce the worst-case variance (captured by the operator matrix norm). Second, we avoid having to invert two matrices. Indeed, the Markowitz would first run a regression, requiring the matrix inversion Σ_S^{-1} and then perform a portfolio optimization, requiring the matrix inversion, ${\Sigma_R}^{-1}$, which is known to be unstable in practice. Instead, we formulate a simpler objective function, leading to a solution that is simply a sum of singular vectors.

B Cross-Sectionally Demeaning the Signals

We start with a signal S, which is not cross-sectionally demeaned, and use the notation "$\sim$" (tilde) to indicate demeaning:

$$\tilde{S}_{j,t} \;=\; S_{j,t} \;-\; \frac{1}{N}\sum_{k=1}^{N} S_{k,t}$$

Similarly, the prediction matrix based on demeaned signals is

$$\tilde{\Pi} \;=\; (R_{i,t+1}\,\tilde{S}_{j,t})_{i,j=1}^{N}\,.$$

and $\tilde{\lambda}_k^s$ and $\tilde{\lambda}_k^a$ are the eigenvalues of $\tilde{\Pi}^s$ and $\tilde{\Pi}^a$, respectively.

Proposition 11 *The demeaned eigenvalues* $\tilde\lambda_k^s$ *are interlacing with the non-demeaned ones* λ_k^s *in the sense that*

$$\lambda_{k+1}^s \;\le \tilde\lambda_k^s \;\le\; \lambda_{k-1}^s \tag{65}$$

for all $k = 2, \ldots, K-1$, $\lambda_2^s \;\le\; \tilde\lambda_1^s$, *and* $\tilde\lambda_N^s \;\le \lambda_{N-1}^s$ *and similarly for the antisymmetric eigenvalues. Hence, if* Π^s *has* N^p *positive eigenvalues, then* $\tilde\Pi^s$ *has between* N^p-1 *and* N^p+1 *positive eigenvalues. Furthermore, the total performance of the cross-sectional factor*

$$E\left(\sum_{i=1}^N \tilde S_{i,t}\, R_{i,t+1}\right) \;=\; \sum_{i=1}^N \tilde\lambda_i^s$$

satisfies

$$\sum_{i=2}^N \lambda_i^s \;\le\; E\left(\sum_{i=1}^N \tilde S_{i,t}\, R_{i,t+1}\right) \;\le\; \sum_{i=1}^{N-1} \lambda_i^s\,. \tag{66}$$

In particular, if all eigenvalues λ_i^s *are positive, then cross-sectional factor performs worse than the time series factor.*

Proof. We use Weyl inequalities (Horn and Johnson (1991)): for any two symmetric or Hermitian matrices $A,\ B$,

$$\lambda_j(A) \;+\; \lambda_k(B) \;\le\; \lambda_i(A+B) \;\le\; \lambda_r(A) \;+\; \lambda_s(B) \tag{67}$$

whenever $j+k-N \ge i \ge r+s-1$.

Then, we note that, by direct calculation,

$$\tilde\Pi^s \;=\; \Pi^s \;+\; X\,,$$

where $X = 0.5(\pi\mathbf{1}' \;+\; \mathbf{1}\pi')$, and where the vector $\pi = (\pi_i) \;=\; -(E(R_{i,t+1}\frac{1}{N}\sum_{k=1}^N S_{k,t}))$.

The matrix X has rank two and at most two non-zero eigenvalues that always have opposite signs: $\lambda_1(X) \ \geq \ 0 \ \geq \ \lambda_N(X)$. Thus, by the Weyl inequalities,

$$\lambda_{i+1}(\Pi^s) \ \leq \ \lambda_N(X) \ + \ \lambda_{i+1}(\Pi^s) \ \leq \ \lambda_i(\tilde{\Pi}^s) \ \leq \ \lambda_{i-1}(\Pi^s) \ + \ \lambda_2(X) \ \leq \ \lambda_{i-1}(\Pi^s)\,.$$

The proof for the antisymmetric part is analogous.

To prove the last inequality, define the orthogonal projection $P = Id - \frac{1}{N}\mathbf{1}_{N\times N}$ Then, $\tilde{S}_t = PS_t$ and hence, by direct calculation, $\tilde{\Pi} = \Pi P$. Furthermore, since signals are demeaned,

$$\sum_{i=1}^N \tilde{S}_{i,t}\, R_{i,t+1} \ = \ \sum_{i=1}^N \tilde{S}_{i,t}\, \tilde{R}_{i,t+1}$$

where

$$\tilde{R}_{t+1} \ = \ PR_{t+1}$$

and hence

$$\begin{aligned} E\left(\sum_{i=1}^N \tilde{S}_{i,t}\, R_{i,t+1}\right) \ &= \ E\left(\sum_{i=1}^N \tilde{S}_{i,t}\, \tilde{R}_{i,t+1}\right) \ = \ \operatorname{tr}(E(\tilde{R}_{t+1}\tilde{S}_t')) \\ = \ \operatorname{tr}(P\Pi P) \ &= \ \operatorname{tr}(P\Pi^s P)\,. \end{aligned} \tag{68}$$

The eigenvalues of $P\Pi^s P$ coincide with the $N-1$ eigenvalues $\{\hat{\lambda}_k^s\}_{k=1}^{N-1}$ of $P\Pi^s P\lceil_{P\mathbb{R}^N}$ restricted onto the subspace $P\mathbb{R}^N$, plus a zero eigenvalue. By the interlacing inequalities Horn and Johnson (1991), we have $\lambda_{k+1}^s \leq \hat{\lambda}_k^s \leq \lambda_k^s$, and therefore

$$\operatorname{tr}(P\Pi^s P) \ = \ \sum_{k=1}^{N-1} \hat{\lambda}_k^s$$

satisfies the required inequalities. □

B.1 Cross-Sectionally Demeaning Portfolio Signals

Given any orthonormal tuple of portfolios $\{\pi_k\}_{k=1}^K$, we define the corresponding demeaned signals as

$$\tilde{S}_t^{\pi_k} \;=\; S_t^{\pi_k} \;-\; \frac{1}{K}\sum_{i=1}^K S_t^{\pi_i}\,.$$

Interestingly, we "lose" one eigenvalue when using demeaned factors rather than non-demeaned ones as seen in the following proposition where $\lambda_1^s \geq \cdots \geq \lambda_N^s$ are still the eigenvalues of Π^s (i.e., based on the original, non-demeaned signals).

Proposition 12 *The expected excess return of demeaned portfolios based on any orthonormal tuple of portfolios $\{\pi_k\}_{k=1}^K$ satisfies*

$$\sum_{i=1}^{K-1}\lambda_k^s \;\geq\; E(\sum_{k=1}^K \tilde{S}_t^{\pi_k} R_{t+1}^{\pi_k}) \;\geq\; \sum_{i=N-K+2}^{N}\lambda_k^s$$

and the bounds are exact.

Proof of Proposition 12. Let

$$\tilde{\pi}_k \;=\; \pi_k \;-\; \frac{1}{K}\sum_{i=1}^K \pi_i\,.$$

Then,

$$\tilde{S}_t^{\pi_k} \;=\; S_t^{\tilde{\pi}_k}$$

and, since $\sum_k S_t^{\tilde{\pi}_k} = 0$, we have

$$\sum_k S_t^{\tilde{\pi}_k} R_{t+1}^{\pi_k} \;=\; \sum_k S_t^{\tilde{\pi}_k} R_{t+1}^{\tilde{\pi}_k}\,.$$

Define the matrix

$$X \;=\; (\tilde{\pi}_1, \cdots, \tilde{\pi}_K),$$

with the columns given by $\tilde{\pi}_i$. Then,

$$E(\sum_k S_t^{\tilde{\pi}_k} R_{t+1}^{\tilde{\pi}_k}) \;=\; \mathrm{tr}(X'\Pi^s X) \;=\; \mathrm{tr}(\Pi^s X X')$$

Since, by assumption, π_k are orthonormal, we have

$$\tilde{\pi}_k'\tilde{\pi}_l \;=\; (\pi_k \;-\; \frac{1}{K}\sum_{i=1}^{K}\pi_i)'(\pi_l \;-\; \frac{1}{K}\sum_{i=1}^{K}\pi_i) \;=\; \delta_{k,l} - 1/K$$

and hence the matrix $P = X'X \in \mathbb{R}^{K\times K}$ has rank $K-1$ and eigenvalues 1 (of multiplicity $K-1$) and 0 and is therefore an orthogonal projection. Thus, we can write $X = UP$ where $U \in \mathbb{R}^{N\times K}$ is an orthogonal matrix satisfying $U'U = Id_K$. Thus,

$$\mathrm{tr}(X'\Pi^s X) \;=\; \mathrm{tr}(PU'\Pi^s UP)\,.$$

Let V be the orthogonal matrix such that $\tilde{P} = V'PV$ is the projection onto the span of the first $K-1$ standard basis vectors of $\mathbb{R}^N$. Then,

$$\mathrm{tr}(PU'\Pi^s UP) \;=\; \mathrm{tr}(PU'\Pi^s UPVV') \;=\; \mathrm{tr}(V'PVV'U'\Pi^s UVV'PV) \;=\; \mathrm{tr}(\tilde{P}\tilde{U}'\Pi^s\tilde{U}\tilde{P})$$

where $\tilde{U} = UV$. Then, $\tilde{U}$ is an arbitrary orthogonal matrix with columns $u_1, \cdots, u_k$, and

$$\mathrm{tr}(\tilde{P}\tilde{U}'\Pi^s\tilde{U}\tilde{P}) \;=\; \sum_{i=1}^{K-1} u_k'\Pi^s u_k$$

and the claim follows from the Ky Fan inequality (Fan (1950)). □

C Proofs

Proof of Proposition 1. Using the identity $\mathrm{tr}(AB) = \mathrm{tr}(AB)$ for any two square matrices A, B, we get

$$\mathrm{tr}(E(S_t' L R_{t+1})) \;=\; \mathrm{tr}(E(L R_{t+1} S_t')) \;=\; \mathrm{tr}(L\,E(R_{t+1} S_t')) \;=\; \mathrm{tr}(L\Pi) \;=\; \mathrm{tr}(\Pi L)\,.$$

□

Proof of Proposition 2. The proof follows directly from Proposition 1 and the fact that $\mathrm{tr}(X'X) \geq 0$ for any matrix X.

□

Proof of Proposition 3. We provide a proof in the case of a general Schatten p-norm considered in Proposition 10, where $p = [1, \infty]$ and q is defined by $1/p + 1/q = 1$. First, the trace of any square matrix $A = UDV'$ is less than the sum of its singular values (d_k):

$$|\,\mathrm{tr}(A)| \;= |\,\mathrm{tr}(UDV')| = |\,\mathrm{tr}(V'UD)| = |\sum_k d_k (V'U)_{k,k}| \leq \sum_k d_k = \|A\|_1\,. \tag{69}$$

since $|(V'U)_{k,k}| = |V'_{\text{column } k} U_{\text{column } k}| \leq \|V_{\text{column } k}\|\;\|U_{\text{column } k}\| = 1$. Combining this inequality with Hölder's inequality for Schatten norms (see, e.g., Bhatia (1997), Corollary IV.2.6; or Tao (2012), p. 55, Exercise 1.3.9), we get:

$$|\,\mathrm{tr}(L\Pi)| \;\leq \|\Pi L\|_1 \leq \;\|\Pi\|_q\,\|L\|_p\,, \tag{70}$$

Finally, equality is achieved if L is proportional to $(\Pi'\Pi)^{q/2-1}\Pi'$. Thus,

$$\arg\max_{\|L\|_p \leq 1} \mathrm{tr}(L\Pi) \;=\; (\Pi'\Pi)^{q/2-1}\Pi' / \|(\Pi'\Pi)^{q/2-1}\Pi'\|_p\,. \tag{71}$$

□

Proof of Proposition 4. Follows from the calculations in the main text. □

Proof of Lemma 1. Suppose that there exists a tradable factor

$$F_{t+1} \; = \; x_t' R_{t+1}$$

such that

$$S_{i,t} \; = \; \frac{\text{Cov}_t(R_{i,t+1}, F_{t+1})}{\text{Var}_t(F_{t+1})} \, .$$

We have

$$\text{Cov}_t(R_{i,t+1}, F_{t+1}) \; = \; \text{Cov}_t(R_{i,t+1}, x_t' R_{t+1}) \; = \; (\Sigma_t^R x_t)_i$$

and, hence,

$$S_t \; = \; \Sigma_t^R x_t \, / y \, ,$$

where we have defined

$$y \; = \; \text{Var}_t(F_{t+1}) \, .$$

Furthermore,

$$\text{Var}_t(F_{t+1}) \; = \; x_t' \Sigma_t^R x_t$$

Thus, we get

$$x_t \;=\; y(\Sigma_t^R)^{-1} S_t\,,$$

and we get a fixed point equation for y :

$$y \;=\; \mathrm{Var}_t(F_{t+1}) \;=\; x_t'\Sigma_t^R x_t \;=\; y^2 S_t'(\Sigma_t^R)^{-1} S_t \;\Leftrightarrow\; y \;=\; 1/S_t'(\Sigma_t^R)^{-1} S_t\,.$$

Reverting the arguments, we see that the converse is also true: the just computed portfolio x_t does satisfy $S_{i,t} \;=\; \frac{\mathrm{Cov}_t(R_{i,t+1},F_{t+1})}{\mathrm{Var}_t(F_{t+1})}$. □

Proof of Proposition 5. By Lemma 2, we have

$$\begin{aligned} &\mathrm{tr}(L\Pi) \;=\; \mathrm{tr}((L^s+L^a)(\Pi^s+\Pi^a)) \\ &=\; \mathrm{tr}(L^s\Pi^s) \;+\; \mathrm{tr}(L^s\Pi^a) \;+\; \mathrm{tr}(L^a\Pi^s) \;+\; \mathrm{tr}(L^a\Pi^a) \;=\; \mathrm{tr}(L^s\Pi^s) \;+\; \mathrm{tr}(L^a\Pi^a)\,. \end{aligned} \tag{72}$$

Finally,

$$\begin{aligned} \mathrm{Cov}_t(R_{t+1}^{w_t}, F_{t+1}) &= \mathrm{Cov}_t(w_t' R_{t+1},\; F_{t+1}) = w_t'\mathrm{Cov}_t(R_{t+1}, F_{t+1}) \\ &= \mathrm{Var}_t(F_{t+1})\, w_t' S_t \;=\; \mathrm{Var}_t(F_{t+1})\, S_t' L S_t \;=\; \mathrm{Var}_t(F_{t+1})\, S_t' L^s S_t\,, \end{aligned} \tag{73}$$

where the third identity uses the definition of F from Lemma 1 and the fact that $\mathrm{Var}_t(F_{t+1}) = 1/(S_t'\,(\Sigma_{R,t})^{-1}S_t)$, and the last identity follows because, by Lemma 2,

$$S_t' L S_t \;=\; S_t' L^s S_t \;+\; S_t' L^a S_t \;=\; S_t' L^s S_t\,. \tag{74}$$

□

Proof of Lemma 2. Since the trace of a matrix equals the trace of its transpose, we have

$$\operatorname{tr}(AB) = \operatorname{tr}((AB)') = \operatorname{tr}(B'A') = -\operatorname{tr}(BA) = -\operatorname{tr}(AB)$$

which shows that $\operatorname{tr}(AB) = \operatorname{tr}(BA) = 0$. Similarly, $x'Ax = (x'Ax)' = x'A'x = -x'Ax$, showing that $x'Ax = 0$. □

Proof of Proposition 6. To see the first result, note that the return equation (23) combined with the eigendecomposition (25) yield

$$E\left(S_t^{w_k^s} R_{t+1}^{w_k^s}\right) \;=\; (w_k^s)'\Pi^s w_k^s = (w_k^s)' \sum_{j=1}^{K} \lambda_j^s w_j^s (w_j^s)' w_k^s = \lambda_k^s\,. \tag{75}$$

The last claim follows directly from (71) for Π^s because, for $L = L'$, by Proposition 5 we have that

$$\max_{\|L\|_p \le 1,\; L=L'} \operatorname{tr}(L\Pi) \;=\; \max_{\|L\|_p \le 1,\; L=L'} \operatorname{tr}(L\Pi^s)$$

is attained by the symmetric matrix $c((\Pi^s)'\Pi^s)^{q/2-1}(\Pi^s)' \;=\; (|\Pi^s|)^{q-2}\Pi^s$. For $q = 1$, we get $|\Pi^s|^{-1}\Pi^s \;=\; \operatorname{sign}(\Pi^s)$. Here, we have used the standard functional calculus for symmetric matrices (Horn and Johnson (1991)): for any function $f(x)$ (such as $|x|$ or $\operatorname{sign}(x)$) we define $f(L) \;=\; W \operatorname{diag}(f(\lambda(L)))W'$ where $L = W \operatorname{diag}(\lambda(L))W'$ is the eigen-decomposition of a symmetric matrix L. □

Proof of Proposition 7. Next, since $W = (w_1^s, ..., w_N^s)$ forms an orthonormal basis of $\mathbb{R}^N$, we have $WW' = Id$ so

$$\widetilde{F}_{t+1} = S_t' R_{t+1} = S_t' W W' R_{t+1} = (W'S_t) \cdot (W'R_{t+1}) = \sum_{k=1}^{N} S_t^{w_k^s} R_{t+1}^{w_k^s} \tag{76}$$

Hence, the result follows from the fact that including negative eigenvalues lowers the expected return relative to the other options considered. □

Proof of Lemma 3. Equip $\mathbb{C}^N$ with the standard inner product

$$x\hat{\cdot}y = \sum_i x_i \bar{y}_i \tag{77}$$

and recall that the Hermitian adjoint of a matrix B is defined as $B^* = \bar{B}'$, where $\bar{B}$ is the complex adjoint of B. Furthermore, for any matrix B, we have

$$x\hat{\cdot}(By) = (B^*x)\hat{\cdot}y\,. \tag{78}$$

Let now A be a real anti-symmetric matrix. Consider the matrix iA. The first observation is that iA is a Hermitian matrix. Indeed, $\overline{iA}' = -iA' = iA$. Thus, iA has real eigenvalues and a basis of complex eigenvectors $\{w_k\}_{k=1}^N$. Let $\lambda \in \mathbb{R}$ be an eigenvalue of iA :

$$iAw = \lambda w \tag{79}$$

Then, taking a complex conjugate of this identity, we get

$$-iA\bar{w} = \lambda\,\bar{w} \tag{80}$$

and hence $\bar{w}$ is an eigenvector of iA with the eigenvalue $-\lambda$. Hence, all non-zero eigenvalues come in pairs.

Furthermore, $\det(A) = \det(A') = \det(-A) = (-1)^N \det(A)$ so, if N is odd, A is degenerate and has a zero eigenvalue, whereas all non-zero eigenvalues come in pairs. Let us take all nonnegative eigenvalues $\lambda_1 \geq \cdots \geq \lambda_K$ of iA and let w_k be the respective complex eigenvectors. By the above, $-\lambda_k$ is also an eigenvalue, and the respective eigenvectors are

$\bar{w}_k$. By the spectral decomposition theorem, we have

$$iA \;=\; \sum_{k=1}^{K}(\lambda_k w_k \bar{w}_k' \;-\; \lambda_k \bar{w}_k w_k') \tag{81}$$

where we have used that the orthogonal projection onto w_k is $w_k\bar{w}_k'$, where $\bar{w}_k$ is the complex conjugate vector. Now, we have

$$\begin{aligned}\lambda_k w_k \bar{w}_k' \;-\; \lambda_k \bar{w}_k w_k' &= 0.5\lambda_k(\, w_{k,1} + i w_{k,2})(w_{k,1} - i w_{k,2})' \\ &\quad - 0.5\lambda_k(w_{k,1} - i w_{k,2})(w_{k,1} + i w_{k,2})' \\ &= i\lambda_k(w_{k,2}w_{k,1}' - w_{k,1}w_{k,2}')\end{aligned} \tag{82}$$

and the claim follows.

For any Hermitian matrix (and, hence, also for iA), eigenvectors for different eigenvalues are always orthogonal. Thus, $w_{k,1} \pm i w_{k,2}$ must be orthogonal to $w_{j,1} \pm i w_{j,2}$ and hence $w_{k,1}, w_{k,2}$ are orthogonal to $w_{j,1},\; w_{j,2}$. Furthermore, $w_{k,1} \pm i w_{k,2}$ correspond to different eigenvalues $\pm\lambda_k$ and hence they also must be orthogonal:

$$0 \;=\; (w_{k,1} + i w_{k,2}) \cdot (w_{k,1} + i w_{k,2}) \;=\; \|w_{k,1}\|^2 - \|w_{k,2}\|^2 + 2i w_{k,1} \cdot w_{k,2} \tag{83}$$

and hence $\|w_{k,1}\| = \|w_{k,2}\|$ and $w_{k,1} \cdot w_{k,2} = 0$. Thus, the two vectors are $w_{k,1},\; w_{k,2}$ are also orthogonal.

Note that

$$iA(w_{k,1} + i w_{k,2}) \;=\; \lambda_k(w_{k,1} + i w_{k,2}) \tag{84}$$

is equivalent to $Aw_{k,1} \;=\; \lambda_k w_{k,2}$ and $Aw_{k,2} = -\lambda_k w_{k,1}$ implying that, in the basis $\{(w_{k,1}, w_{k,2})\}_{k=1}^K$

the matrix A is block-diagonal, composed of diagonal blocks

$$\begin{pmatrix} 0 & -\lambda_1 \\ \lambda_1 & 0 \end{pmatrix} \tag{85}$$

□

Proof of Proposition 8. By definition, $(\Pi^a)'(x_j + iy_j) = i\lambda_j^a(x_j + iy_j)$, that is $(\Pi^a)'x_j = -\lambda_j^a y_j,\ (\Pi^a)'y_j = \lambda_j^a x_j$. The expected return is

$$\begin{aligned} E(R_{t+1}^{w_{j,t}^a}) &= E(S_t'(x_j y_j' - y_j x_j')R_{t+1}) = \operatorname{tr}((x_j y_j' - y_j x_j')\Pi^a) = -\operatorname{tr}((\Pi^a)'(x_j y_j' - y_j x_j')) \\ &= \lambda_j^a \operatorname{tr}(y_j' y_j + x_j' x_j) = 2\lambda_j^a . \end{aligned}$$

The last statement follows from (71) for Π^a because, for $L = -L'$, by Proposition 5 we have that

$$\max_{\|L\|_p \le 1,\ L=-L'} \operatorname{tr}(L\Pi) = \max_{\|L\|_p \le 1,\ L=-L'} \operatorname{tr}(L\Pi^a)$$

is attained by the anti-symmetric matrix $c((\Pi^a)'\Pi^a)^{q/2-1}(\Pi^a)' = (|i\Pi^a|)^{q-2}(\Pi^a)'$. For $q = 1$, we get $|i\Pi^a|^{-1}(\Pi^a)' = \operatorname{sign}(i\Pi^a)$. Here, we have used the standard functional calculus for Hermitian matrices (Horn and Johnson (1991)): for any function $f(x)$ (such as $|x|$ or $\operatorname{sign}(x)$) we define $f(L) = W \operatorname{diag}(f(\lambda(L)))W'$ where $L = W \operatorname{diag}(\lambda(L))W'$ is the eigen-decomposition of a Hermitian matrix L. □

Proof of Proposition 9. Follows from calculations in the body of the paper. □

Proof of Proposition 10. Let $X = L\Pi$. Then, by a result of Fan and Hoffman (1955), we have $\lambda_i(X^s) \le \bar{\lambda}_i(X)$. Furthermore, since $\operatorname{rank}(L) \le K$, we also have $\operatorname{rank}(X) \le K$ and $\operatorname{rank}(X'X) = \operatorname{rank}(X) \le K$, and hence there are at most K non-zero singular values of X.

Thus,

$$\operatorname{tr}(L\Pi) \;=\; \operatorname{tr}(X) \;=\; \operatorname{tr}(X^s) \;=\; \sum_{i=1}^{N} \lambda_i(X^s) \;\le\; \sum_{i=1}^{K} \bar\lambda_i(X)\,.$$

Second, by known result about singular values of products of matrices (see, for example, Marshall and Olkin (1979), p. 248), we have

$$\sum_{i=1}^{K} \bar\lambda_i(L\Pi) \;\le\; \sum_{i=1}^{K} \bar\lambda_i(L)\,\bar\lambda_i(\Pi)$$

Third, by the Hölder inequality

$$\sum_{i=1}^{K} \bar\lambda_i(L)\,\bar\lambda_i(\Pi) \;\le\; \left(\sum_{i=1}^{K} \bar\lambda_i(L)^p\right)^{1/p} \left(\sum_{i=1}^{K} \bar\lambda_i(\Pi)^q\right)^{1/q} \;=\; \|L\|_p \left(\sum_{i=1}^{K} \bar\lambda_i(\Pi)^q\right)^{1/q}\,.$$

Thus,

$$\max_{\|L\|_p\le 1,\ \operatorname{rank}(L)\le K} \operatorname{tr}(L\Pi) \;\le\; \left(\sum_{i=1}^{K} \bar\lambda_i(\Pi)^q\right)^{1/q}\,.$$

Thus, it remains to verify that the equality holds with $L = c\sum_{k=1}^{K} \bar\lambda_k^{q-1} v_k(u_k)'$, where $c = \left(\sum_{k=1}^{K} \bar\lambda_k^q\right)^{-1/p}$. This follows directly from the identity $\operatorname{tr}(Lv_k(u_k)')$, established in (15)

Items 2 and 3 of the Proposition are in fact special cases of item 1. Indeed, for item 2, we have by Proposition 5 that

$$\begin{aligned}
\max_{\|L\|_p\le 1,\ L=L'\ \operatorname{rank}(L)\le K} \operatorname{tr}(L\Pi) \;&=\; \max_{\|L\|_p\le 1,\ L=L',\ \operatorname{rank}(L)\le K} \operatorname{tr}(L\Pi^s) \\
\le \max_{\|L\|_p\le 1,\ \operatorname{rank}(L)\le K} \operatorname{tr}(L\Pi^s) \;&=\; \left(\sum_{i=1}^{K} \bar\lambda_i(\Pi^s)^q\right)^{1/q}\,.
\end{aligned} \tag{86}$$

Furthermore, $\bar\lambda_i(\Pi^s)$ is the i-th largest absolute eigenvalue of Π^s and the equality is achieved with $L = c\sum_{\mathcal{K}} |\lambda_k^s|^{q-1}\operatorname{sign}(\lambda_k^s) w_k^s(w_k^s)'$

Similarly, for the anti-symmetric part, we have

$$
\begin{aligned}
\max_{\|L\|_p \le 1,\ L=-L' \,\mathrm{rank}(L)\le K} \mathrm{tr}(L\Pi) \;&=\; \max_{\|L\|_p \le 1,\ L=-L',\ \mathrm{rank}(L)\le K} \mathrm{tr}(L\Pi^a) \\
\le \max_{\|L\|_p \le 1,\ \mathrm{rank}(L)\le K} \mathrm{tr}(L\Pi^a) \;&=\; \left(\sum_{i=1}^{K} \bar{\lambda}_i(\Pi^a)^q\right)^{1/q} .
\end{aligned} \tag{87}
$$

Top $2K$ singular value of Π^a are just λ_i^a counted twice, and the equality is achieved with $L = c\sum_{k=1}^{K}(\lambda_k^a)^{q-1}(w_k^a(\bar{w^a}_k)' - \bar{w^a}_k(w_k^a)')$, where $w_k^a = x_k + iy_k$ are the *complex* eigenvectors of Π^a and come in complex conjugate pairs according to Lemma 3. $\square$

Table 1: Analogy between PCA and PPA

This table shows five analogies between principal component analysis (PCA) and principal portfolio analysis (PPA) for the symmetric part of the prediction matrix. For PCA (PPA): (i) the variance (expected excess returns) of each component equals its eigenvalue; (ii) different components $k \neq l$ are orthogonal; (iii) the sum variances (returns) of individual securities equals that of the components, and also equals the trace of the variance-covariance matrix (prediction matrix); (iv) the top K components maximize variance (return) for orthonormal portfolios; and (v) component $k+1$ maximizes variance (return) among all portfolios that are orthogonal to the first k ones.

	Principal Component Analysis	Principal Portfolio Analysis (Symmetric Part)
(i)	$\text{Var}(R_{t+1}^{\pi_k}) \;=\; \lambda_k(\Sigma_R)$	$E(S_t^{w_k^s} R_{t+1}^{w_k^s}) \;=\; \lambda_k(\Pi^s)$
(ii)	$\text{Cov}(R_{t+1}^{\pi_k}, R_{t+1}^{\pi_l}) \;=\; 0$	$E(S_t^{w_k^s} R_{t+1}^{w_l^s}) \;+\; E(S_t^{w_l^s} R_{t+1}^{w_k^s}) = 0$
(iii)	$\sum_k \text{Var}(R_{k,t+1}) = \sum_k \text{Var}(R_{t+1}^{\pi_k}) = \text{tr}(\Sigma_R)$	$\sum_k E(S_{k,t} R_{k,t+1}) = \sum_k E(S_t^{w_k^s} R_{t+1}^{w_k^s}) = \text{tr}(\Pi^s)$
(iv)	$(\pi_k) = \arg\max_{\text{orthon.}\{x_k\}_{k=1}^K} \sum_k \text{Var}(R_{t+1}^{x_k})$	$(w_k^s) = \arg\max_{\text{orthon.}\{x_k\}_{k=1}^K} \sum_k E(S_t^{x_k} R_{t+1}^{x_k})$
(v)	$\pi_{k+1} = \arg\max_{x \perp \{\pi_1, \cdots, \pi_k\}} \text{Var}(R_{t+1}^x)$	$w_{k+1}^s = \arg\max_{x \perp \{w_1^s, \cdots, w_k^s\}} E(S_t^x R_{t+1}^x)$

Table 2: Principal Portfolio Factor Exposures

Statistics for regressions of out-of-sample principal portfolio returns on own-predictor strategy ("Factor") and the five Fama-French factors. Portfolios are constructed from the Fama-French 25 size and value portfolios based on a 20-day momentum signal. The table reports regressions for the own-predictor strategy itself, the equal-weighted average of the top three principal portfolios ("PP 1-3"), the equal-weighted average of the top three principal exposure portfolios ("PEP 1-3"), the equal-weighted average of the top three principal alpha portfolios ("PAP 1-3"), and the equal weighted average of the top three PEP's and PAP's combined ("PEP and PAP 1-3"). In each regression, the left-hand-side portfolio is scaled to have the same full-sample volatility as the excess market return. Results are shown for a 20-day forecast horizons, and each forecast is made on an out-of-sample basis using a rolling training sample of the 120 most recent non-overlapping return observations. Sample period is 1963-2019.

Portfolio	Factor	Mkt-Rf	SMB	HML	RMW	CMA	Alpha	R^2
Factor		-0.2	0.13	-0.28	-0.26	0.36	9.35	0.08
t-statistic		-4.59	1.9	-3.43	-3.03	2.79	4.12	
PP 1-3	0.82	0.03	0.02	0.15	-0.09	-0.02	4.69	0.67
t-statistic	32.69	1.09	0.53	3.05	-1.63	-0.29	3.38	
PEP 1-3	0.94	0.01	-0.02	0.06	-0.13	-0.01	0.89	0.89
t-statistic	67.00	0.88	-0.76	1.95	-4.4	-0.16	1.14	
PAP 1-3	-0.08	0.08	0.19	0.06	0.28	0.06	10.41	0.04
t-statistic	-1.94	1.71	2.65	0.72	3.1	0.42	4.42	
PEP and PAP 1-3	0.65	0.07	0.13	0.09	0.11	0.04	8.51	0.41
t-statistic	19.53	1.93	2.31	1.32	1.58	0.35	4.62	

Figure 1: Prediction Matrix Eigenvalues

Panels A, B, and C show estimated eigenvalues of the prediction matrix and its symmetric and anti-symmetric components, respectively, averaged over training samples. Panels D, E, and F show average out-of-sample returns and ±2 standard error confidence bands for corresponding principal portfolios, principal exposure portfolios, and principal alpha portfolios, respectively. Estimates are based on predictions of 20-day returns of the Fama-French 25 size and value portfolios based on a 20-day momentum signal. Each training sample consists of 120 non-overlapping 20-day return observations. Sample period is 1963-2019.

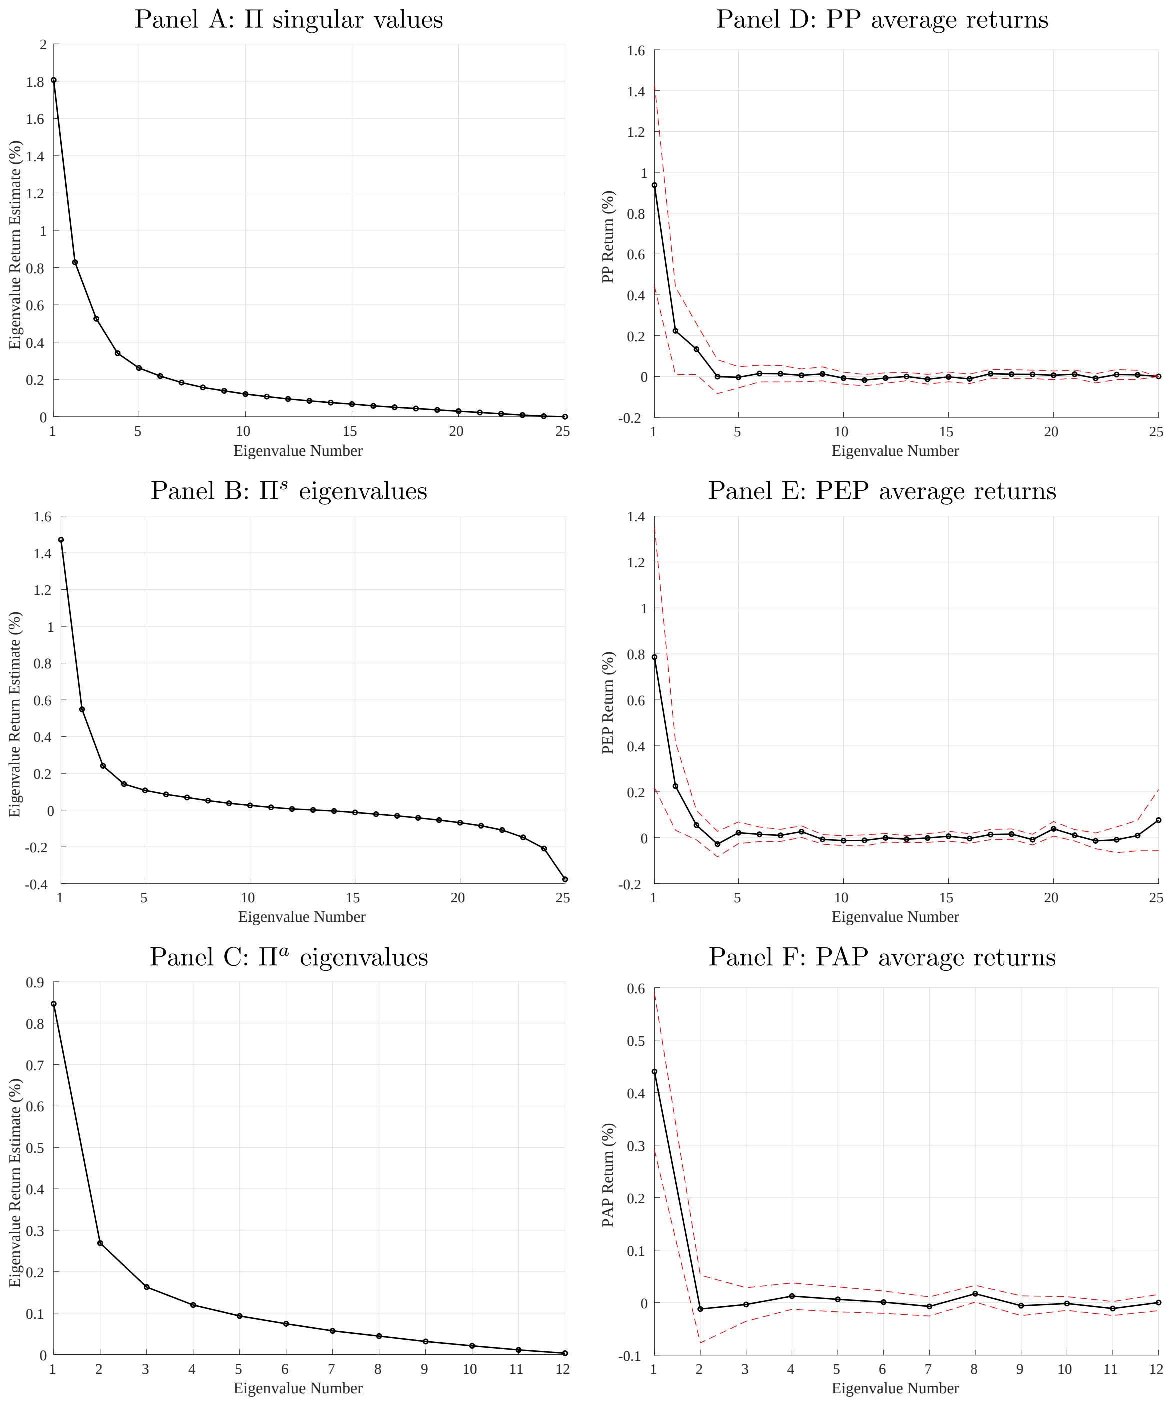

Figure 2: Portfolio Weights for Leading Principal Portfolios

Weights of the first principal exposure portfolio (Panel A) and first principal alpha portfolio (Panel B) on the 25 size and value portfolios, averaged over training samples. Portfolios are constructed based on a 20-day momentum signal and for a 20-day forecast horizon/holding period. Portfolios and estimates are made on an out-of-sample basis using a rolling training sample of the 120 most recent non-overlapping return observations. Sample period is 1963-2019.

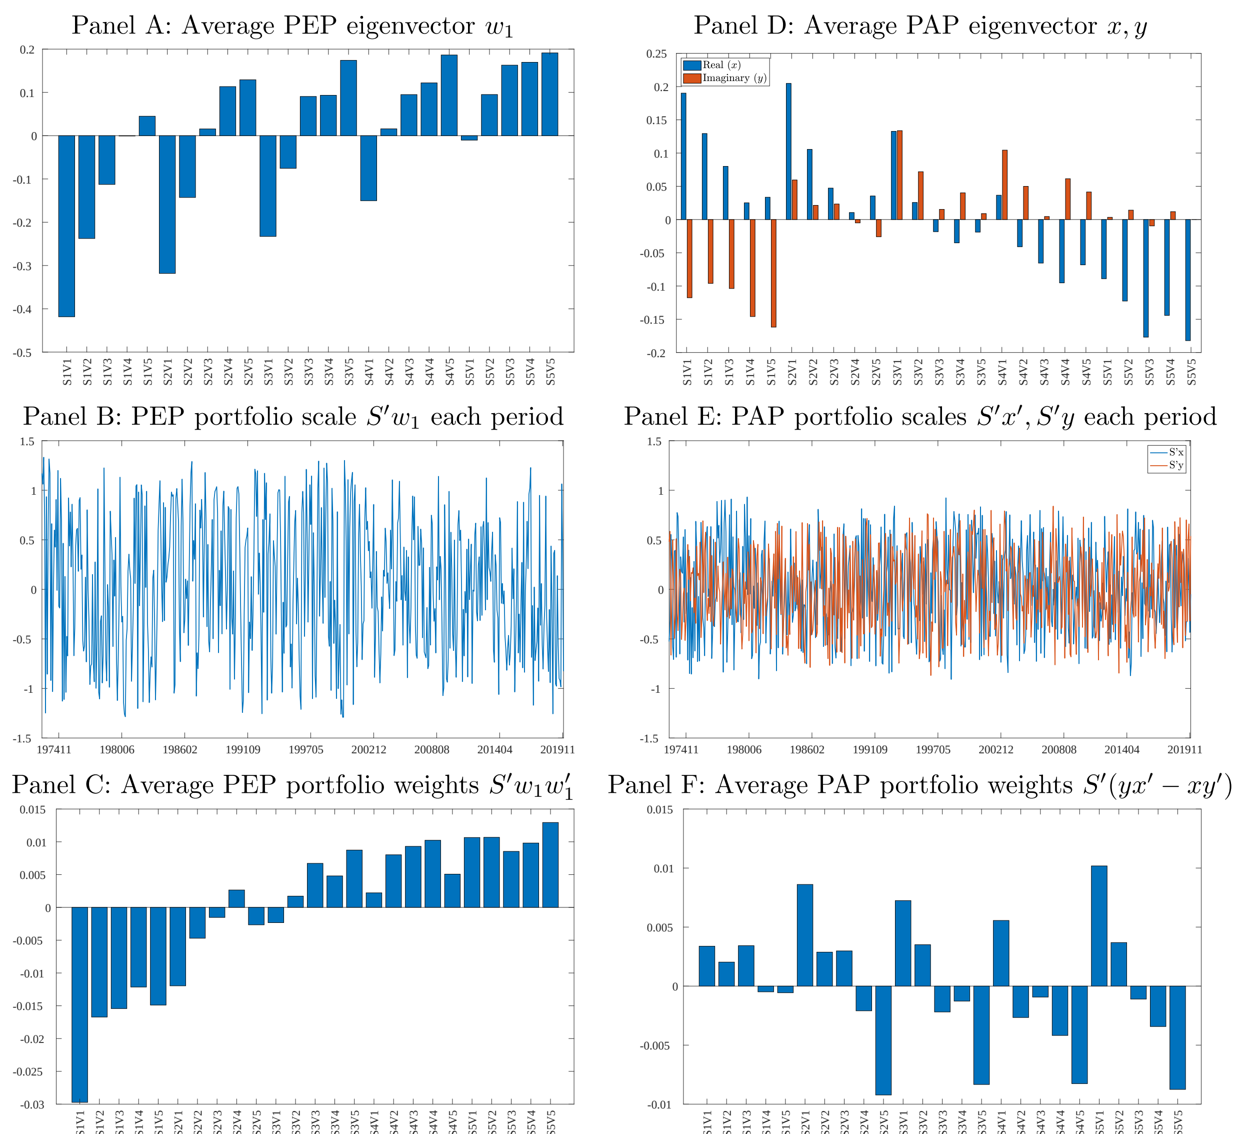

Figure 3: Principal Portfolio Performance by Forecast Horizon

Out-of-sample performance of principal portfolios in terms of annualized Sharpe ratio (Panel A) and annualized information ratio versus the own-predictor strategy and the Fama-French 5-factor model (Panel B) along with ±2 standard error band around each estimate. Portfolios are constructed from the Fama-French 25 size and value portfolios based on a 20-day momentum signal. The figure reports performance of the own-predictor strategy ("Factor"), the equal-weighted average of the top three principal portfolios ("PP 1-3"), the equal-weighted average of the top three principal exposure portfolios ("PEP 1-3"), the equal-weighted average of the top three principal alpha portfolios ("PAP 1-3"), and the equal weighted average of the top three PEP's and PAP's combined ("PEP and PAP 1-3"). Results are shown for forecast horizons (and, equivalently, holding periods) of 1, 5, 10, and 20 days. Each forecast is made on an out-of-sample basis using a rolling training sample of the 120 most recent non-overlapping return observations. Sample period is 1963-2019.

Panel A: Sharpe Ratio

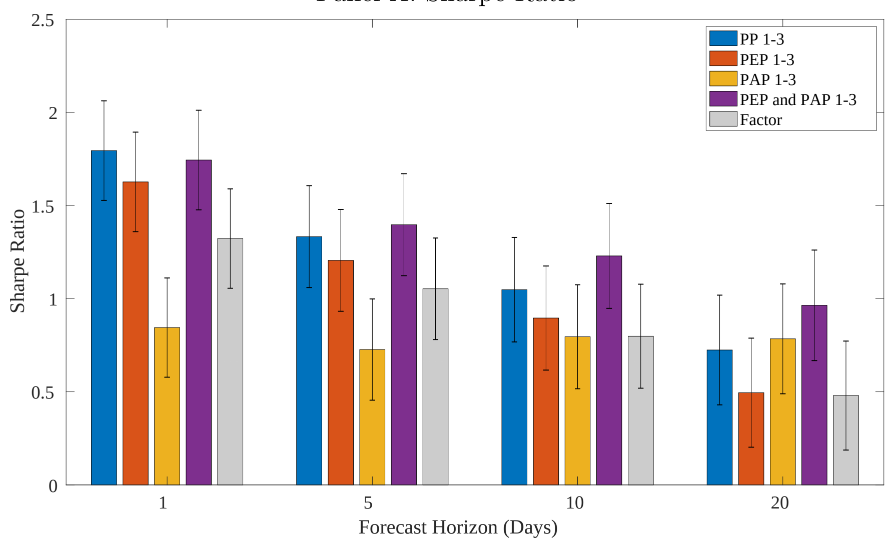

Panel B: Information Ratio

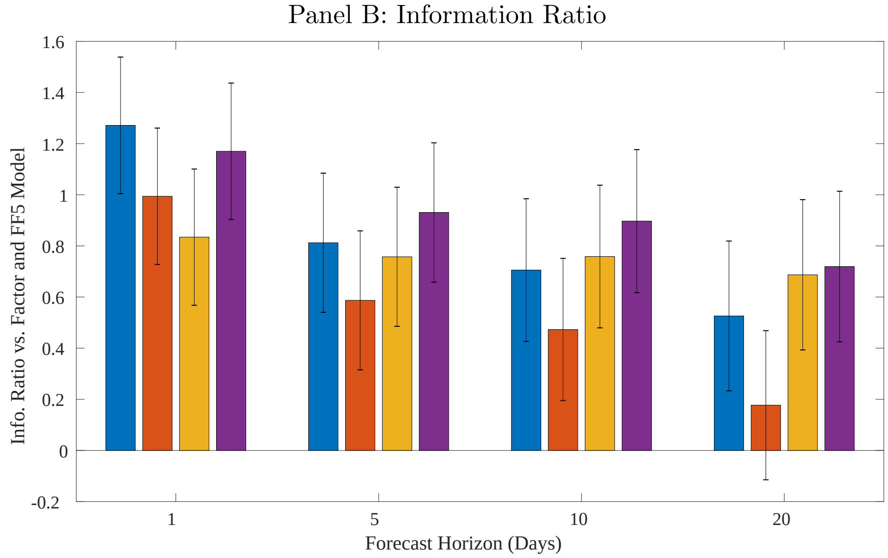

Figure 4: Principal Portfolio Performance in Other Asset Universes

Out-of-sample performance of principal portfolios in terms of annualized Sharpe ratio (Panel A) and annualized information ratio versus the own-predictor strategy and the Fama-French 5-factor model (Panel B) along with ±2 standard error band around each estimate. Portfolios are constructed based on a 20-day momentum signal from either the 25 U.S. size and value portfolios, 25 U.S. size and operating profitability portfolios, 25 U.S. size and investment portfolios, their international counterparts (developed countries excluding the U.S.), or 52 futures contracts. The figure reports performance of the own-predictor strategy ("Factor"), the equal-weighted average of the top three principal portfolios ("PP 1-3"), the equal-weighted average of the top three principal exposure portfolios ("PEP 1-3"), the equal-weighted average of the top three principal alpha portfolios ("PAP 1-3"), and the equal weighted average of the top three PEP's and PAP's combined ("PEP and PAP 1-3"). Panel C shows the correlation between out-of-sample average portfolio returns and eigenvalues of the prediction matrix. Blue bars show the correlation between PP's and singular values from the total prediction matrix, red bars show the correlation between PEP's and eigenvalues from the symmetric component, and yellow bars show the correlation between PAP's and eigenvalues from the anti-symmetric component. Results are shown for a 20-day forecast horizon/holding period. Each forecast is made on an out-of-sample basis using a rolling training sample of the 120 most recent non-overlapping return observations. Sample period is 1963-2019 for U.S. equity portfolios, 1990-2019 for international equity portfolios, and 1985-2019 for futures contracts.

Panel A: Sharpe Ratio

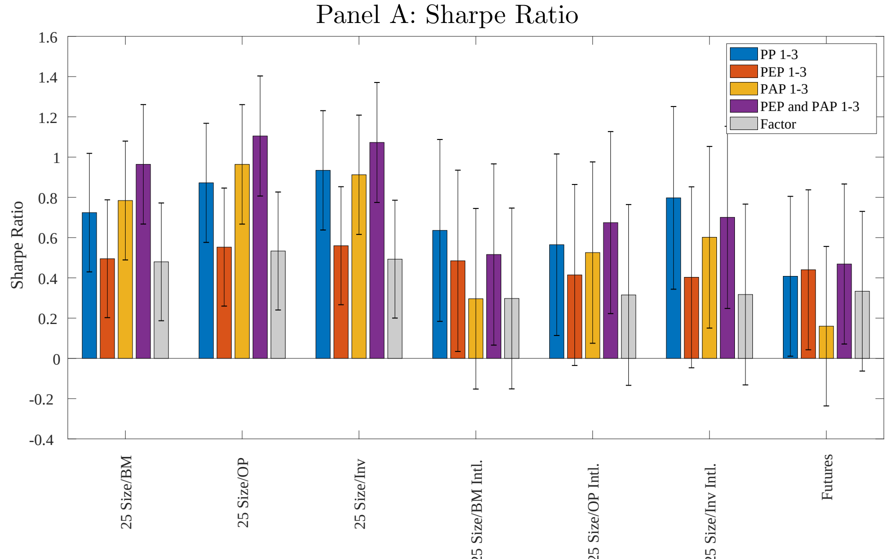

Panel B: Information Ratio

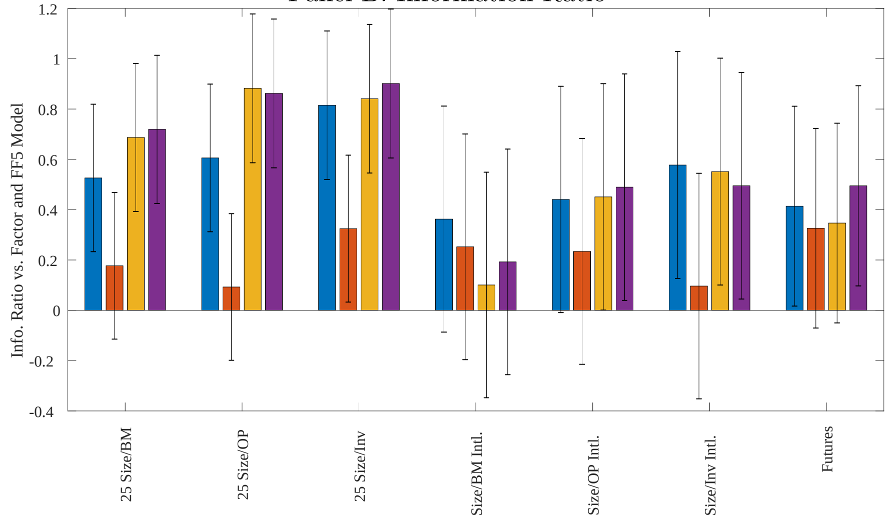

Panel C: Correlation between average portfolio returns and eigenvalues

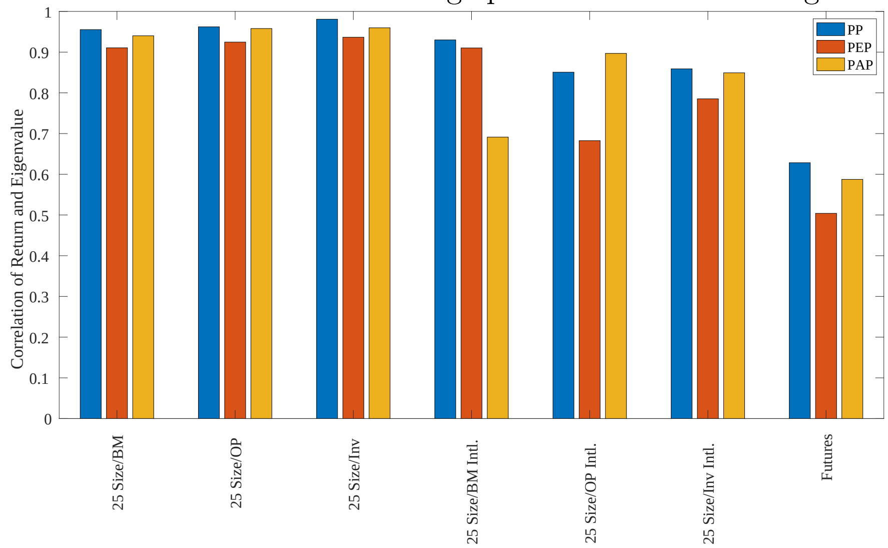

Figure 5: Principal Portfolio Performance by Momentum Lookback Window

Out-of-sample performance of principal portfolios in terms of annualized Sharpe ratio (Panel A) and annualized information ratio versus the own-predictor strategy and the Fama-French 5-factor model (Panel B) along with ±2 standard error band around each estimate. Portfolios are constructed based on a 20, 40, 60, 120 or 250-day momentum signal from 25 size and value portfolios. The figure reports performance of the own-predictor strategy ("Factor"), the equal-weighted average of the top three principal portfolios ("PP 1-3"), the equal-weighted average of the top three principal exposure portfolios ("PEP 1-3"), the equal-weighted average of the top three principal alpha portfolios ("PAP 1-3"), and the equal weighted average of the top three PEP's and PAP's combined ("PEP and PAP 1-3"). Each forecast is made on an out-of-sample basis using a rolling training sample of the 120 most recent non-overlapping return observations. Sample period is 1963-2019.

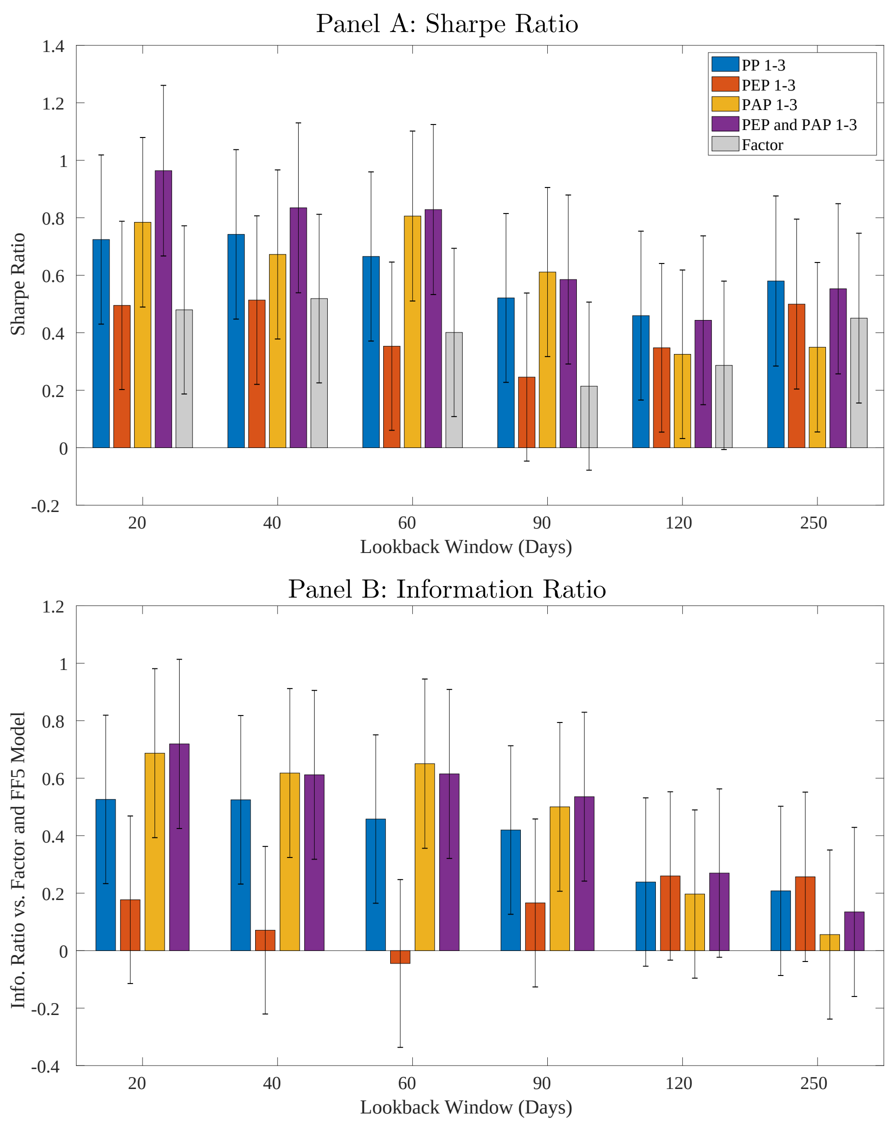

Figure 6: Principal Portfolio Performance in Subsamples

Out-of-sample performance of principal portfolios by decade in terms of annualized Sharpe ratio (Panel A) and annualized information ratio versus the own-predictor strategy and the Fama-French 5-factor model (Panel B) along with ±2 standard error band around each estimate. Portfolios are constructed based on a 20-day momentum signal from 25 size and value portfolios. The figure reports performance of the own-predictor strategy ("Factor"), the equal-weighted average of the top three principal portfolios ("PP 1-3"), the equal-weighted average of the top three principal exposure portfolios ("PEP 1-3"), the equal-weighted average of the top three principal alpha portfolios ("PAP 1-3"), and the equal weighted average of the top three PEP's and PAP's combined ("PEP and PAP 1-3"). Results are shown for a 20-day forecast horizon/holding period. Each forecast is made on an out-of-sample basis using a rolling training sample of the 120 most recent non-overlapping return observations.