

When do cross-sectional asset pricing factors span the stochastic discount factor?*

Serhiy Kozak[†]
University of Maryland

Stefan Nagel[‡]
*University of Chicago,
NBER, CEPR, and CESifo*

November 10, 2022

When expected returns are linear in asset characteristics, the stochastic discount factor (SDF) that prices individual stocks can be represented as a factor model with GLS cross-sectional regression slope factors. Factors constructed heuristically by aggregating individual stocks into characteristics-based factor portfolios using sorting, characteristics-weighting, or OLS cross-sectional regression slopes do not span this SDF unless the covariance matrix of stock returns has a specific structure. These conditions are more likely satisfied when researchers use large numbers of characteristics simultaneously. Methods to hedge unpriced components of heuristic factor returns allow partial relaxation of these conditions. We also show the conditions that must hold for dimension reduction to a number of factors smaller than the number of characteristics to be possible without having to invert a large covariance matrix. Under these conditions, instrumented and projected principal components analysis methods can be implemented as simple PCA on certain portfolio sorts.

*We are grateful for discussions with Gene Fama and thank seminar participants at Stanford University, University of Maryland, and Tianshu Lyu for excellent research assistance. Nagel gratefully acknowledges financial support from the Fama-Miller Center and the Center for Research in Security Prices at the University of Chicago Booth School of Business

[†]University of Maryland, R.H. Smith School of Business, 7621 Mowatt Lane, College Park, MD 20742, e-mail: sekozak@umd.edu

[‡]University of Chicago, Booth School of Business, 5807 South Woodlawn Avenue, Chicago, IL 60637, e-mail: stefan.nagel@chicagobooth.edu

I. INTRODUCTION

Cross-sectional asset pricing research has linked stocks’ expected returns to a large set of firm characteristics. To summarize these cross-sectional pricing patterns in a reduced-form pricing model, researchers often construct stochastic discount factor (SDF) proxies with multiple characteristics-based factors. Individual assets’ weights in these factor portfolios are functions of stock characteristics. Researchers use a variety of different heuristic approaches to specify these weights. For example, Fama and French (1993) sort stocks by characteristics and then form portfolios by applying quintile cutoffs (*sorted factors*); Kozak, Nagel, and Santosh (2020) construct portfolios with weights proportional to stocks’ centered univariate cross-sectional rank for each characteristic (*univariate factors*); Fama and French (2020) use the slopes of monthly cross-sectional OLS regressions of returns on characteristics as factor portfolio returns (*OLS factors*).¹

There is, however, only one unique SDF—or, equivalently, one mean-variance efficient portfolio return—that is spanned by returns of the individual assets (Hansen and Jagannathan 1991). Under which conditions do these different heuristic methods yield this SDF? Put differently, under which conditions is the investment opportunity set not deteriorating if one aggregates individual assets to these factor portfolios? Somewhat surprisingly, the answer to this fundamental question is not available in the literature.

Clearly, some special conditions must be met because the weights of individual assets in the mean-variance efficient portfolio weights depend on the return covariance matrix, but none of these heuristic methods use any information from the covariance matrix in factor construction. Our first objective in this paper is to work out what these conditions are.

We set estimation issues aside at first and work with population moments. We assume that conditional expected returns of N individual stocks are linear in $J \leq N$ firm charac-

1. More precisely, Fama and French (2020) use a hybrid approach where individual stocks are first sorted into a relatively large number of characteristics-based portfolios and the value-weighted returns and characteristics of these portfolios are then the input for the cross-sectional OLS regressions.

teristics collected in the $N \times J$ characteristics matrix \mathbf{X}_t . At this conceptual level, this assumption is very general since the set of characteristics could also contain nonlinear functions of some underlying basic characteristics. For instance, sorted factors can be subsumed by univariate factors if characteristics are specified as step functions. This linearity assumption only acquires substantive empirical content once a researcher has fixed a specific set of characteristics that she works with.

As a starting point, we show that the SDF that prices all individual assets can be expressed as a multifactor SDF spanned by J factors that are the slopes of cross-sectional GLS regressions of returns on lagged firm characteristics (*GLS factors*). The inverse of the conditional covariance matrix of returns serves as the GLS weighting matrix. The matrix of individual assets' conditional betas on the GLS factors is then exactly equal to \mathbf{X}_t .²

In practice, construction of these GLS factors would be difficult because it requires estimating and inverting a large conditional covariance matrix. For this reason, it is important to know whether heuristic approaches that bypass this inversion problem can deliver factors that span the SDF. Sorted factors, univariate factors, and OLS factors are all simply weighting stocks by columns of \mathbf{X}_t or a nonsingular rotation thereof. We show that these factors span the SDF if and only if the conditional covariance matrix Σ_t of individual asset returns takes the specific form

$$\Sigma_t = \mathbf{X}_t \Psi_t \mathbf{X}_t' + \mathbf{U}_t \Omega_t \mathbf{U}_t', \quad \mathbf{X}_t' \mathbf{U}_t = \mathbf{0}. \quad (1)$$

This means that there must be a clean separation among the sources of systematic risk such that loadings on up to J systematic factors are perfectly spanned by \mathbf{X}_t while loadings on the remaining ones are orthogonal to \mathbf{X}_t . When (1) holds, individual assets' betas on OLS factors are exactly equal to \mathbf{X}_t , i.e., covariances are equal to characteristics. Fama and French (2020) argue that the OLS factors can be used as asset pricing factors in time-series

2. The GLS factors are similar to the characteristics-efficient portfolios of Daniel, Mota, Rottke, and Santos (2020), but in our analysis we allow for time-varying conditional moments.

regressions with conditional betas set equal to \mathbf{X}_t , but our result shows that this is true if and only if condition (1) is satisfied.

Condition (1) is more likely to hold approximately when \mathbf{X}_t includes a large, comprehensive set of characteristics. In this case, important sources of stock return covariance can be absorbed in the first term of $\boldsymbol{\Sigma}_t$ in (1), which leaves \mathbf{U}_t and violations of $\mathbf{X}_t' \mathbf{U}_t = \mathbf{0}$ quantitatively unimportant. Additional characteristics can help even if they are unrelated to expected returns as long as they help to capture major sources of stock return covariances. But if the number of characteristics is small—as in popular low-dimensional factor models with only four or five characteristics-based factors—there is little reason to think that this small number of characteristics should be sufficient to span loadings on all major sources of covariance.

Existing empirical results are suggestive that low-dimensional factor models with heuristic factors do not satisfy condition (1). For example, Gerakos and Linnainmaa (2018) find that the HML value factor is contaminated with unpriced components; Back, Kapadia, and Ostdiek (2015) find that OLS factors have alpha with respect to the standard sorted factors of Hou, Xue, and Zhang (2015) and Fama and French (2015); Grinblatt and Saxena (2018) find that sorted factors do not price the basis portfolios from which they were constructed; Chib, Lin, Pukthuanthong, and Zeng (2021) find that the method of factor construction affects asset pricing performance. All of these findings indicate that the low-dimensional factor models do not span the SDF.

Motivated by these findings, researchers have developed heuristic methods to remove unpriced components from heuristic factors. Daniel, Mota, Rottke, and Santos (2020) (DMRS) propose a hedging approach that is meant to remove unpriced risks from the original factors. They construct hedge portfolios that have positive loadings on the original factors but zero exposure to the underlying characteristics that determine expected returns. Residualizing the original factors with respect to the hedge portfolio returns removes unpriced risks. However, it is not clear under which conditions this heuristic approach actually yields a better

approximation of the SDF. Our second objective therefore is to understand the conditions under which this hedging approach can be used to recover factors that span the SDF.

With \mathbf{H}_t representing the DMRS hedged factor portfolio weights, or a nonsingular rotation thereof, we show that the hedged factors span the SDF if the covariance matrix has the structure in (1), but with the requirement $\mathbf{X}'_t \mathbf{U}_t$ replaced with the requirement that there exists a decomposition such that

$$\mathbf{U}_t \boldsymbol{\Omega}_t \mathbf{U}'_t = \mathbf{V}_t \boldsymbol{\Gamma}_t \mathbf{V}'_t + \mathbf{E}_t \boldsymbol{\Phi}_t \mathbf{E}'_t, \quad \mathbf{X}'_t \mathbf{E}_t = \mathbf{0}, \quad (2)$$

where \mathbf{V}_t is an $N \times J$ matrix. This is a weaker condition than (1) because here columns of \mathbf{X}_t can be correlated with columns of \mathbf{U}_t , as long as this correlation comes only through the J columns of \mathbf{V}_t . Again, this condition is more likely to hold when researchers consider a large, comprehensive set of characteristics.

While DMRS consider only one round of hedging, there is no reason to stop after one round. Iteration on this approach, by hedging once more the already-hedged factor portfolios can yield further improvements. We show that with two rounds of hedging the resulting factors span the SDF if a condition like (2) holds, but in this case \mathbf{V}_t can have $2J$ instead of J columns, i.e., there can be even more sources of correlation between the columns of \mathbf{X}_t and \mathbf{U}_t than under condition (2).

The approaches we discussed so far construct J factors to capture the pricing information of J characteristics. Dimension-reduction methods aim to span the SDF with a smaller number of $K < J$ factors while again avoiding the need to invert an estimate of $\boldsymbol{\Sigma}_t$. Different approaches for dimension reduction exist in the literature, but it is not clear what the necessary conditions are for the factors constructed with these methods to span the SDF. Our third objective is therefore to establish these conditions.

We show that if and only if the conditional covariance matrix has a structure like in (1), but with \mathbf{X}_t replaced by lower-dimensional $K \leq J$ linear combinations of characteristics collected in $\mathbf{X}_t \mathbf{Q}_t$, then portfolios with weights equal to $\mathbf{X}_t \mathbf{Q}_t$, or a non-singular transfor-

mation thereof, span the SDF. We further show that under these conditions, two prominent methods of dimension reduction, the instrumented principal components method (IPCA) of Kelly, Pruitt, and Su (2019) and the projected PCA method (PPCA) of Kim, Korajczyk, and Neuhierl (2019), which both assume that loadings on latent risk factors are linear in characteristics but use different identification assumptions, can be implemented using simple PCA on OLS portfolios or univariate portfolios constructed using orthonormal characteristics, respectively.

Finally, we turn to empirical implementation. Our theoretical results are all stated in terms of population moments. However, we find that our theoretical results also characterize well the properties of factor models with empirically estimated moments. In the empirical analysis, we focus on the properties of OLS factor models constructed using the stock characteristics from Kozak, Nagel, and Santosh (2020). As one would expect based on our theoretical results, OLS factors generally do not span the SDF that prices individual stocks. We infer this from the fact that hedging the OLS factors using the DMRS method produces substantial improvements of the MVE portfolio's (maximum) Sharpe ratio attainable by the factors. As suggested by our theoretical analysis, iterating the hedging procedures produces further substantial gains in the maximum Sharpe ratio. Furthermore, while these maximum Sharpe ratio improvements are large for small-scale factor models that use only a few characteristics, they vanish when we use a large number of characteristics to construct the OLS factors. This is in line with our conclusion from the theoretical analysis that condition (1) is more likely to hold, and therefore OLS factors more likely to span the SDF that prices individual stocks, when the econometrician employs a large number of characteristics.

II. CONDITIONS FOR CHARACTERISTICS-BASED PORTFOLIOS TO SPAN THE MEAN-VARIANCE FRONTIER

We consider a cross-section of N assets with an $N \times 1$ vector of excess returns \mathbf{z}_{t+1} . Each asset features J characteristics that are observable to the econometrician, collected in the (time-varying) $N \times J$ matrix \mathbf{X}_t where $J \leq N$, $\text{rank}(\mathbf{X}_t) = J$, and the first column of \mathbf{X}_t is a vector of ones. In a number of places in our analysis we will use the residual maker matrix $\mathbf{R}_t = \mathbf{I} - \mathbf{X}_t(\mathbf{X}_t'\mathbf{X}_t)^{-1}\mathbf{X}_t'$ that generates the residuals in a projection on \mathbf{X}_t . Unless otherwise noted, we use the notation $\boldsymbol{\mu}_{y,t} = \mathbb{E}_t[\mathbf{y}_{t+1}]$, $\boldsymbol{\Sigma}_{y,t} = \text{var}_t(\mathbf{y}_{t+1})$ for the conditional moments of a random vector \mathbf{y}_{t+1} , $\boldsymbol{\Sigma}_{xy,t}$ as notation for the conditional covariance matrix of two random vectors \mathbf{x}_{t+1} and \mathbf{y}_{t+1} , and \mathbf{I}_K for a $K \times K$ identity matrix.

In what follows, all time- t conditional moments are conditioned on \mathbf{X}_t , and we denote

$$\boldsymbol{\Sigma}_t = \text{var}(\mathbf{z}_{t+1}|\mathbf{X}_t), \quad \boldsymbol{\mu}_t = \mathbb{E}[\mathbf{z}_{t+1}|\mathbf{X}_t], \quad (3)$$

and we assume that $\boldsymbol{\Sigma}_t$ is positive definite. That these conditional moments are conditioned on the characteristics observable to the econometrician is important. The set of characteristics observable to investors could be larger or smaller than what is contained in \mathbf{X}_t , without consequences for our results, as long as the law of one price holds conditional on \mathbf{X}_t .³ Therefore, it is possible that conditional on investors' information set, moments of excess returns could vary more or less than conditional on the econometrician's information. Only sources of variation linked to \mathbf{X}_t matter in our analysis.

We assume throughout that the law of one price holds and hence an SDF exists. Conditional on the econometrician's information, the maximum squared conditional Sharpe ratio that can be obtained from the N individual assets then is finite and given by $\boldsymbol{\mu}_t'\boldsymbol{\Sigma}_t^{-1}\boldsymbol{\mu}_t$. The

3. As an example that would violate this requirement, the law of one price would fail if the econometrician included elements of \mathbf{z}_{t+1} in \mathbf{X}_t . Conditional on this look-ahead information, arbitrage opportunities would seemingly exist.

SDF that uses this maximum squared conditional Sharpe ratio portfolio as risk factor,

$$M_{t+1} = 1 - \mathbf{b}'_t(\mathbf{z}_{t+1} - \boldsymbol{\mu}_t), \quad \mathbf{b}_t = \boldsymbol{\Sigma}_t^{-1} \boldsymbol{\mu}_t, \quad (4)$$

prices the N assets conditionally, i.e., $\mathbb{E}[M_{t+1} \mathbf{z}_{t+1} | \mathbf{X}_t] = 0$. This is the unique SDF in the span of excess returns. We refer to it from now on simply as the SDF.

Our analysis focuses on characteristics-based factors. These factors are generally constructed with an $N \times J$ portfolio weight matrix \mathbf{W}_t , where the weights are functions of the characteristics \mathbf{X}_t , and possibly also of $\boldsymbol{\Sigma}_t$. Using these weights, one can form J factor portfolios as

$$\mathbf{f}_{t+1} = \mathbf{W}'_t \mathbf{z}_{t+1}, \quad (5)$$

with $\boldsymbol{\mu}_{f,t} = \mathbf{W}'_t \boldsymbol{\mu}_t$ and $\boldsymbol{\Sigma}_{f,t} = \mathbf{W}'_t \boldsymbol{\Sigma}_t \mathbf{W}_t$. We assume that weights are such that $\boldsymbol{\Sigma}_{f,t}$ is positive definite.

Our aim is to understand under which conditions different specifications of the weights \mathbf{W}_t produce factors that span the conditional mean-variance frontier. Spanning the conditional mean-variance frontier is equivalent to the factors' maximum squared conditional Sharpe ratio,

$$\boldsymbol{\mu}'_{f,t} \boldsymbol{\Sigma}_{f,t}^{-1} \boldsymbol{\mu}_{f,t} = \boldsymbol{\mu}'_t \mathbf{W}_t (\mathbf{W}'_t \boldsymbol{\Sigma}_t \mathbf{W}_t)^{-1} \mathbf{W}'_t \boldsymbol{\mu}_t, \quad (6)$$

attaining the maximum squared conditional Sharpe Ratio obtainable from the individual assets. Our results below rely on the following lemma that provides conditions under which this is true.

Lemma 1 *The maximum squared conditional Sharpe ratio of the factors $\mathbf{f}_{t+1} = \mathbf{W}'_t \mathbf{z}_{t+1}$ is equal to the maximum squared conditional Sharpe Ratio of the individual assets, i.e.,*

$$\boldsymbol{\mu}'_t \boldsymbol{\Sigma}_t^{-1} \boldsymbol{\mu}_t = \boldsymbol{\mu}'_t \mathbf{W}_t (\mathbf{W}'_t \boldsymbol{\Sigma}_t \mathbf{W}_t)^{-1} \mathbf{W}'_t \boldsymbol{\mu}_t \quad (7)$$

if and only if

$$\boldsymbol{\mu}_t = \boldsymbol{\Sigma}_t \mathbf{W}_t \mathbf{b}_t \quad (8)$$

for some $J \times 1$ vector \mathbf{b}_t .

Proof. Following the proof of Lu and Schmidt (2012) Theorem 3 (A, B), express the difference of the left- and right-hand-sides of (7) as $\boldsymbol{\Delta} = \boldsymbol{\mu}'_t \boldsymbol{\Sigma}_t^{-\frac{1}{2}} \mathbf{M} \boldsymbol{\Sigma}_t^{-\frac{1}{2}} \boldsymbol{\mu}_t$, where $\mathbf{M} = \mathbf{I} - \mathbf{P}$ and \mathbf{P} are residual and projection matrices, respectively, for a projection onto the columns of $\boldsymbol{\Sigma}_t^{\frac{1}{2}} \mathbf{W}_t$. $\boldsymbol{\Delta} = \mathbf{0}$ if and only if $\boldsymbol{\Sigma}_t^{-\frac{1}{2}} \boldsymbol{\mu}_t$ is in the column space of $\boldsymbol{\Sigma}_t^{\frac{1}{2}} \mathbf{W}_t$, that is, $\boldsymbol{\Sigma}_t^{-\frac{1}{2}} \boldsymbol{\mu}_t = \boldsymbol{\Sigma}_t^{\frac{1}{2}} \mathbf{W}_t \mathbf{b}_t$ for some \mathbf{b}_t , which is equivalent to (8). ■

If the factors span the conditional mean-variance frontier, then they span the SDF that prices the individual assets:

Corollary 1 *Lemma 1 implies that if and only if equation (8) holds, an SDF can be represented in terms of the J factors:*

$$M_{t+1} = 1 - \mathbf{b}'_t (\mathbf{f}_{t+1} - \boldsymbol{\mu}_{f,t}). \quad (9)$$

This SDF perfectly prices the excess returns \mathbf{z}_{t+1} , that is, $\mathbb{E}[M_{t+1} \mathbf{z}_{t+1} | \mathbf{X}_t] = 0$. This SDF representation is equivalent to a conditional beta-pricing representation

$$\boldsymbol{\mu}_t = \boldsymbol{\beta}_t \boldsymbol{\mu}_{f,t}, \quad (10)$$

where $\boldsymbol{\beta}_t = \boldsymbol{\Sigma}_{zf,t} \boldsymbol{\Sigma}_{f,t}^{-1}$.

Equipped with this result, we can now explore under which assumptions about $\boldsymbol{\mu}_t$ and $\boldsymbol{\Sigma}_t$ various heuristic methods of factor construction that have appeared in the literature yield factors that span the SDF.

Our baseline assumption about expected returns is motivated by a large body of work that has documented cross-sectional relationships between expected return and firm characteristics:

Assumption 1 (*Linearity of expected returns in characteristics*)

$$\boldsymbol{\mu}_t = \mathbf{X}_t \boldsymbol{\phi} \tag{11}$$

with some $J \times 1$ vector $\boldsymbol{\phi}$.

At a conceptual level, the assumption that $\boldsymbol{\mu}_t$ is linear in \mathbf{X}_t is without loss of generality as \mathbf{X}_t could also include nonlinear functions of characteristics. Similarly, portfolio sorting approaches that allow expected returns to differ across but not within bins defined by characteristics can be accommodated in Assumption 1 by letting \mathbf{X}_t be a matrix of bin membership indicators. That $\boldsymbol{\phi}$ is a constant parameter vector is not restrictive either, because one could include nonlinear interactions of cross-sectional firm characteristics with time-series predictors to capture any time-variation in expected returns. In practice, though, once a researcher has chosen a specific set of characteristics to include in \mathbf{X}_t , Assumption 1 becomes a substantive assumption that restricts $\boldsymbol{\mu}_t$. Later in the paper, we discuss alternative assumptions.

Comparing Assumption 1 and equation (8), we see that $\boldsymbol{\Sigma}_t \mathbf{W}_t$ collapses to \mathbf{X}_t only in special cases when certain conditions are satisfied for \mathbf{W}_t , or certain restrictions on $\boldsymbol{\Sigma}_t$ hold. We now explore these conditions.

II.A. *The unique SDF in the span of excess returns: GLS factors and rotations thereof*

As a benchmark for understanding when and why heuristic factor models span or do not span the SDF, we first show that the SDF in (9) has a J -factor representation under Assumption 1:

Proposition 1 *Assumption 1 is equivalent to the statement that an SDF given by (9) with characteristics-based factors*

$$\mathbf{f}_{t+1} = \mathbf{S}'_t \mathbf{X}'_t \boldsymbol{\Sigma}_t^{-1} \mathbf{z}_{t+1}, \tag{12}$$

and prices of risk

$$\mathbf{b}_t = \mathbf{S}_t^{-1} \boldsymbol{\phi}, \tag{13}$$

where \mathbf{S}_t is any nonsingular $J \times J$ rotation matrix, perfectly prices the excess returns \mathbf{z}_{t+1} , that is, $\mathbb{E}[M_{t+1}\mathbf{z}_{t+1}|\mathbf{X}_t] = 0$.

Proof. Rewrite (11) as $\boldsymbol{\mu}_t = \boldsymbol{\Sigma}_t \boldsymbol{\Sigma}_t^{-1} \mathbf{X}_t \mathbf{S}_t \mathbf{S}_t^{-1} \boldsymbol{\phi}_t = \boldsymbol{\Sigma}_t \mathbf{W}_t \mathbf{b}_t$, where $\mathbf{W}_t = \boldsymbol{\Sigma}_t^{-1} \mathbf{X}_t \mathbf{S}_t$. Lemma 1 now applies. ■

Thus, when there is a linear relationship between J characteristics and conditional expected return, the SDF is spanned by J characteristics-based factors that exactly explains these conditional expected returns with zero pricing errors. Proposition 1 therefore highlights that there is no economic difference between a model that specifies expected returns directly as linear function of characteristics as in Assumption 1 and a characteristics-based factor pricing model. One can always be mapped perfectly into the other one, with equivalent pricing implications. Therefore, a horse race between direct linear prediction of \mathbf{z}_{t+1} by \mathbf{X}_t and a factor pricing model, e.g., as in as in Daniel and Titman (1997) and Davis, Fama, and French (2000) as well as many other papers, does not have economic content. If factors are constructed as in Proposition 1, there is no difference in expected returns implied by direct linear prediction and the factor model. If factors are constructed in a heuristic way that does not exactly follow the prescription of Proposition 1, then there can be a difference, but this just reflects the misspecification of the heuristic factors. The difference does not have economic content (it does not discriminate between “rational” and “behavioral” asset pricing theories, for example).

Empirical asset pricing researchers often like to work with beta-pricing specifications and, in particular, with beta-pricing specifications that can be conditioned down to deliver predictions for unconditional expected returns without elaborate estimation of time-varying conditional moments. The following example present such a case.

Example 1 Suppose $\mathbf{S}_t = (\mathbf{X}'_t \boldsymbol{\Sigma}_t^{-1} \mathbf{X}_t)^{-1}$. We then obtain an SDF with factors given by GLS cross-sectional regression slopes, $\mathbf{f}_{t+1} = (\mathbf{X}'_t \boldsymbol{\Sigma}_t^{-1} \mathbf{X}_t)^{-1} \mathbf{X}'_t \boldsymbol{\Sigma}_t^{-1} \mathbf{z}_{t+1}$. Factor risk prices are time varying, $\mathbf{b}_t = (\mathbf{X}'_t \boldsymbol{\Sigma}_t^{-1} \mathbf{X}_t) \boldsymbol{\phi}$. Factor means are constant, $\boldsymbol{\mu}_{f,t} = \boldsymbol{\phi}$. Factor betas are equal to characteristics, $\boldsymbol{\beta}_t = \mathbf{X}_t$.

The GLS slope factors in this example are the GLS counterpart to the OLS cross-sectional slope factors in Fama (1976) and Fama and French (2020). The factors in Example 1 are also similar to the “characteristic-efficient portfolios” in Daniel, Mota, Rottke, and Santos (2020), albeit here with time-varying \mathbf{X}_t and a conditional moments of excess returns. We will show later in Section IV that keeping track of time-variation in \mathbf{X}_t and conditional moments is important in empirical implementation of these factor models.

Which rotation matrix \mathbf{S}_t to pick is a matter of convenience. The next example is one in which factor covariances instead of factor betas are equal to \mathbf{X}_t :

Example 2 *Suppose $\mathbf{S}_t = \mathbf{I}_J$. Factors are $\mathbf{f}_{t+1} = \mathbf{X}'_t \boldsymbol{\Sigma}_t^{-1} \mathbf{z}_{t+1}$. Factor risk prices are constant, $\mathbf{b}_t = \boldsymbol{\phi}$. Factor means are time-varying, $\boldsymbol{\mu}_{f,t} = (\mathbf{X}'_t \boldsymbol{\Sigma}_t^{-1} \mathbf{X}_t) \boldsymbol{\phi}$. Covariances of returns and factors are equal to characteristics, $\boldsymbol{\Sigma}_{zf,t} = \mathbf{X}_t$.*

Practical implementation of the SDF in Proposition 1 is of course difficult since it involves the inversion of a large $N \times N$ conditional covariance matrix. Heuristic approaches to factor construction exist that avoid this inversion problem. We now want to find conditions that need to hold for these heuristic approaches to succeed in spanning the SDF.

II.B. Heuristic factor construction: OLS factors and rotations thereof

Many heuristic methods construct factors by taking long positions in stocks with high values of a characteristic and short positions in stocks with low values of a characteristic, with the portfolio weight matrix and factors then taking the form

$$\mathbf{W}_t = \mathbf{X}_t \mathbf{S}_t, \quad \mathbf{f}_{t+1} = \mathbf{W}'_t \mathbf{z}_{t+1}, \quad (14)$$

for some nonsingular matrix \mathbf{S}_t . For example, $\mathbf{S}_t = \mathbf{I}$ yields a univariate characteristics portfolio where weights are proportional to characteristics as, e.g., in Kozak, Nagel, and Santosh (2020). With characteristics defined as dummy variables for characteristics bins, portfolio sorts can also be represented in this way. Another example are cross-sectional regression slope

factors. Fama and French (2020) use the insight of Fama (1976) that OLS cross-sectional regression slopes are themselves portfolio returns. This is the case $\mathbf{S}_t = (\mathbf{X}'_t \mathbf{X}_t)^{-1}$.

Fama and French (2020) conjecture that the OLS factors yield an “asset pricing model that can be used in time-series applications.” In other words, they conjecture that for N assets with OLS factor betas β_t , the pricing relation $\boldsymbol{\mu}_t = \beta_t \boldsymbol{\mu}_{f,t}$ holds. However, such a pricing relationship does not generally hold for OLS factors. As we show now, this is true only if the covariance matrix takes a special form.

Proposition 2 *Suppose Assumption 1 holds and let $\mathbf{W}_t = \mathbf{X}_t \mathbf{S}_t$. Then, for any nonsingular $J \times J$ matrix \mathbf{S}_t , the maximum squared conditional Sharpe ratio of the factors $\mathbf{f}_{t+1} = \mathbf{W}'_t \mathbf{z}_{t+1}$ is equal to the maximum squared conditional Sharpe Ratio of the individual assets if and only if there exist conformable matrices $\boldsymbol{\Psi}_t$, $\boldsymbol{\Omega}_t$, and a matrix \mathbf{U}_t for which*

$$\mathbf{U}'_t \mathbf{X}_t = \mathbf{0}, \quad (15)$$

such that

$$\boldsymbol{\Sigma}_t = \mathbf{X}_t \boldsymbol{\Psi}_t \mathbf{X}'_t + \mathbf{U}_t \boldsymbol{\Omega}_t \mathbf{U}'_t. \quad (16)$$

Proof. Lu and Schmidt (2012) Theorem 1 (B, F') implies that (16) is equivalent to the statement that there exists a nonsingular \mathbf{B}_t such that $\boldsymbol{\Sigma}_t \mathbf{X}_t = \mathbf{X}_t \mathbf{B}_t$. Rewriting Assumption 1 as $\boldsymbol{\mu}_t = \mathbf{X}_t \mathbf{B}_t \mathbf{B}_t^{-1} \boldsymbol{\phi}$, we see that it is then equivalent to $\boldsymbol{\mu}_t = \boldsymbol{\Sigma}_t \mathbf{X}_t \mathbf{S}_t \mathbf{S}_t^{-1} \mathbf{B}_t^{-1} \boldsymbol{\phi} = \boldsymbol{\Sigma}_t \mathbf{W}_t \mathbf{b}_t$, where $\mathbf{b}_t = \mathbf{S}_t^{-1} \mathbf{B}_t^{-1} \boldsymbol{\phi}$. Thus, condition (8) in Lemma 1 is satisfied, which means that Lemma 1 applies. ■

Without the restriction (15), the decomposition in (16) would always exist. For instance, for any nonsingular symmetric $\boldsymbol{\Psi}_t$, we could obtain $\mathbf{U}_t \boldsymbol{\Omega}_t \mathbf{U}'_t$ from an eigendecomposition of $\boldsymbol{\Sigma}_t - \mathbf{X}_t \boldsymbol{\Psi}_t \mathbf{X}'_t$, where \mathbf{U}_t then contains the eigenvectors associated with the $N - J$ nonzero eigenvalues in the diagonal matrix $\boldsymbol{\Omega}_t$.

How can researchers wishing to use OLS factors, or rotations thereof, ensure that the condition $\mathbf{U}'_t \mathbf{X}_t = \mathbf{0}$ in (15) holds, at least approximately? Including many characteristics in \mathbf{X}_t should help. To see this, we can use the result in Lu and Schmidt (2012) that the

conditions in (15) and (16) are equivalent to J eigenvectors of Σ_t being spanned by \mathbf{X}_t . The matrix \mathbf{U}_t then contains linear combinations of the eigenvectors not spanned by \mathbf{X}_t .⁴ With only a few characteristics included in \mathbf{X}_t , it is unlikely that the J columns of \mathbf{X}_t exactly span J eigenvectors. Effectively, for each eigenvector, this is like asking whether a regression of the N elements of the eigenvector on the J variables in \mathbf{X}_t has perfect fit. Clearly, the more characteristics we add, the better the fit. In this sense, it is more likely that $\mathbf{U}'_t \mathbf{X}_t = \mathbf{0}$ holds if \mathbf{X}_t contains more characteristics.

Moreover, with a larger number of characteristics it is more likely that \mathbf{X}_t spans very well the relatively small number of eigenvectors associated with large eigenvalues, i.e., the major sources of stock return covariance. In this case, even \mathbf{X}_t does not span J eigenvectors perfectly, spanning the few important ones very well may render the violations of $\mathbf{U}'_t \mathbf{X}_t = \mathbf{0}$ quantitatively unimportant. OLS factors, or rotations thereof, may then span the SDF approximately. We investigate this further in our empirical analysis in Section V.

Importantly, for additional characteristics to be helpful in ensuring that $\mathbf{U}'_t \mathbf{X}_t = \mathbf{0}$ holds approximately, these additional characteristics do not necessarily need to contribute to variation in expected returns. If they help to span major sources of covariances, they will help OLS factors, or rotations thereof, to span the SDF, even without contribution to variation in expected returns.

The choice of rotation matrix \mathbf{S}_t is again just one of convenience. Our first example shows the choice that yields OLS cross-sectional regression slope factors:

Example 3 Suppose $\mathbf{S}_t = (\mathbf{X}'_t \mathbf{X}_t)^{-1}$ and that (16) holds. We then obtain an SDF with factors given by OLS cross-sectional regression slopes, $\mathbf{f}_{t+1} = (\mathbf{X}'_t \mathbf{X}_t)^{-1} \mathbf{X}'_t \mathbf{z}_{t+1}$. Prices of

4. If \mathbf{Q}_t and Λ_t are the matrix of eigenvectors and diagonal matrix of eigenvalues of Σ_t , respectively, and $\mathbf{Q}_t = (\mathbf{X}_t \mathbf{B}_t : \mathbf{U}_t)$ where the columns of $\mathbf{X}_t \mathbf{B}_t$, with nonsingular \mathbf{B}_t , are the J eigenvectors spanned by \mathbf{X}_t and \mathbf{U}_t are the eigenvectors not spanned by \mathbf{X}_t , then we have

$$\Sigma_t = \mathbf{X}_t \mathbf{B}_t \Lambda_{1,t} \mathbf{B}'_t \mathbf{X}'_t + \mathbf{U}_t \Lambda_{2,t} \mathbf{U}'_t \quad (17)$$

which maps into (16) with $\mathbf{B}_t \Lambda_{1,t} \mathbf{B}'_t = \Psi_t$ and $\Lambda_{2,t} = \Omega_t$. Moreover, since eigenvectors are orthogonal, $\mathbf{B}'_t \mathbf{X}'_t \mathbf{U}_t = \mathbf{0}$ and hence $\mathbf{U}'_t \mathbf{X}_t = \mathbf{0}$.

risk are $\mathbf{b}_t = \Psi_t^{-1} \mathbf{X}_t \phi$ and factor risk premia are constant, $\boldsymbol{\mu}_{f,t} = \phi$. Factor betas are equal to characteristics, $\boldsymbol{\beta}_t = \mathbf{X}_t$.

A different choice of \mathbf{S}_t produces univariate factors where weights in each characteristics portfolio depend only on one characteristic because the columns of \mathbf{X}_t serve as weights:

Example 4 Suppose $\mathbf{S}_t = \mathbf{I}_J$ and that (16) holds. We then obtain an SDF with factors given by simple univariate portfolio sorts, $\mathbf{f}_{t+1} = \mathbf{X}'_t \mathbf{z}_{t+1}$. Factor risk prices, $\mathbf{b}_t = (\mathbf{X}'_t \mathbf{X}_t)^{-1} \Psi_t^{-1} \phi$ and factor means, $\boldsymbol{\mu}_{f,t} = \mathbf{X}'_t \mathbf{X}_t \phi$, are time-varying. Factor betas are $\boldsymbol{\beta}_t = \mathbf{X}_t (\mathbf{X}'_t \mathbf{X}_t)^{-1}$.

The special case in the latter example is particularly convenient for illustrating the meaning of $\mathbf{U}'_t \mathbf{X}_t = \mathbf{0}$ in (15). For the factor model to price assets perfectly, factor covariances must span $\boldsymbol{\mu}_t = \mathbf{X}_t \phi$. This is always the case for the factors in Example 2 where factor weights are $\boldsymbol{\Sigma}_t^{-1} \mathbf{X}_t$ and hence individual assets' factor covariances are $\boldsymbol{\Sigma}_t \boldsymbol{\Sigma}_t^{-1} \mathbf{X}_t = \mathbf{X}_t$. In contrast, in the case of Example 4, individual assets' covariances with factor portfolios with factor weights \mathbf{X}_t are

$$\boldsymbol{\Sigma}_t \mathbf{X}_t = \mathbf{X}_t \Psi_t \mathbf{X}'_t \mathbf{X}_t + \mathbf{U}_t \boldsymbol{\Omega}_t \mathbf{U}'_t \mathbf{X}_t. \quad (18)$$

If $\mathbf{U}'_t \mathbf{X}_t = \mathbf{0}$, then the second component is zero and the expression hence simplifies to a term \mathbf{X}_t multiplied by a nonsingular matrix. Factor covariances therefore span $\boldsymbol{\mu}_t$. But if $\mathbf{U}'_t \mathbf{X}_t \neq \mathbf{0}$ then the second component does not disappear. As a consequence, the factor covariances are contaminated by components that are not linear in \mathbf{X}_t , and hence are unpriced as they do not earn expected return. Thus, when $\mathbf{U}'_t \mathbf{X}_t \neq \mathbf{0}$, the factors with weights \mathbf{X}_t incorporate unpriced risks, while factors that span the SDF capture only priced risks.

II.C. Hedged heuristic factors

If condition $\mathbf{U}'_t \mathbf{X}_t = \mathbf{0}$ in Proposition 2 does not hold, any factors with weights that are a nonsingular transformation of \mathbf{X}_t load on unpriced risks, i.e., risk exposure that is not compensated with higher excess returns. This prevents the factors from reaching the mean-variance frontier.

Using the GLS factors, or rotations thereof, following Proposition 1 would avoid contamination of factors with unpriced risks, but their construction requires inversion of the large covariance matrix Σ_t (that would have to be estimated in practice). For this reason, it is useful to ask whether there exist an alternative factor specifications that use *some* information about covariances to find characteristics-based factors that span the SDF, but without requiring estimation and inversion of the whole covariance matrix Σ_t . These factors will be *hedged* factors because they hedge unpriced exposures of the original factors.

We first show a result that will be helpful for checking whether a candidate hedged factor model with factor portfolio weight matrix \mathbf{H}_t spans the SDF.

Lemma 2 *Suppose Assumption 1 holds and that \mathbf{H}_t is some matrix such that $\mathbf{H}'_t \mathbf{X}_t$ has full column rank and $\mathbf{H}'_t \Sigma_t \mathbf{H}_t$ is positive definite. Then the maximum squared conditional Sharpe ratio of the factors $\mathbf{f}_{t+1} = \mathbf{H}'_t \mathbf{z}_{t+1}$ is equal to the maximum conditional squared Sharpe Ratio of the individual assets if and only if there exist a nonsingular matrix Ψ_t , and some matrices Ω_t and \mathbf{U}_t for which*

$$\mathbf{U}'_t \mathbf{H}_t = \mathbf{0}, \quad (19)$$

such that

$$\Sigma_t = \mathbf{X}_t \Psi_t \mathbf{X}'_t + \mathbf{U}_t \Omega_t \mathbf{U}'_t. \quad (20)$$

Proof. Lu and Schmidt (2012) Theorem 3 (B, F') implies that (20) is equivalent to the statement that there exists some \mathbf{B}_t such that $\mathbf{X}_t = \Sigma_t \mathbf{H}_t \mathbf{B}_t$. Rewriting Assumption 1 as $\mu_t = \Sigma_t \mathbf{H}_t \mathbf{B}_t \phi = \Sigma_t \mathbf{W}_t \mathbf{b}_t$, where $\mathbf{b}_t = \mathbf{B}_t \phi$. Thus, condition (8) is satisfied, which means that Lemma 1 applies. ■

There are two key points to note. First, the requirement that $\mathbf{H}'_t \mathbf{X}_t$ has full column rank ensures that no information about expected returns is lost when individual assets are aggregated with \mathbf{H}_t as portfolio weight matrix. Second, the requirement that $\mathbf{U}'_t \mathbf{H}_t = \mathbf{0}$ ensures that the factors do not load on unpriced risk. When both conditions hold, a similar

calculation as in (18) for the $\mathbf{S}_t = \mathbf{I}_J$ case, but now with \mathbf{H}_t as factor portfolio weights yields

$$\boldsymbol{\Sigma}_t \mathbf{H}_t = \mathbf{X}_t \boldsymbol{\Psi}_t \mathbf{X}'_t \mathbf{H}_t + \mathbf{U}_t \boldsymbol{\Omega}_t \mathbf{U}'_t \mathbf{H}_t = \mathbf{X}_t \boldsymbol{\Psi}_t \mathbf{X}'_t \mathbf{H}_t, \quad (21)$$

which means that the individual assets' covariances with these factors are perfectly linear in \mathbf{X}_t and so they span $\boldsymbol{\mu}_t$.

While Lemma 2 allows us to check whether candidate factors span the SDF, it does not show how to construct factors that satisfy these requirements. Some conditions on $\mathbf{U}_t \boldsymbol{\Omega}_t \mathbf{U}'_t$ will have to hold for the construction to be possible without using the information from the full $\boldsymbol{\Sigma}_t$ matrix. To see how additional structure on $\mathbf{U}_t \boldsymbol{\Omega}_t \mathbf{U}'_t$ can help, suppose that J columns of \mathbf{U}_t , collected in \mathbf{V}_t , are such that $\text{rank}(\mathbf{V}'_t \mathbf{X}_t) = J$, while the remaining columns, collected in \mathbf{E}_t , have $\mathbf{E}'_t \mathbf{X}_t = \mathbf{0}$ and $\mathbf{E}'_t \mathbf{V}_t = \mathbf{0}$. Moreover, suppose that $\boldsymbol{\Omega}_t$ is block-diagonal such that

$$\mathbf{U}_t \boldsymbol{\Omega}_t \mathbf{U}'_t = \mathbf{V}_t \boldsymbol{\Gamma}_t \mathbf{V}'_t + \mathbf{E}_t \boldsymbol{\Phi}_t \mathbf{E}'_t. \quad (22)$$

If we knew \mathbf{V}_t , we could then simply remove from characteristics-based factor weights $\mathbf{W}_t = \mathbf{X}_t \mathbf{S}_t$ the component that is correlated with \mathbf{U}_t by subtracting the projection of the weights on \mathbf{V}_t ,

$$\mathbf{H}_t = \mathbf{X}_t \mathbf{S}_t - \mathbf{V}_t (\mathbf{V}'_t \mathbf{V}_t)^{-1} \mathbf{V}'_t \mathbf{X}_t \mathbf{S}_t. \quad (23)$$

It is easy to verify that $\mathbf{H}'_t \mathbf{U}_t = \mathbf{0}$ and that $\mathbf{H}'_t \mathbf{X}_t$ has full column rank, i.e., the conditions in Lemma 2 hold.

We cannot directly implement this approach as \mathbf{V}_t is not directly observable. But it can be backed out from moments of \mathbf{z}_t and \mathbf{X}_t . As we show, the factor hedging method of Daniel, Mota, Rottke, and Santos (2020) (DMRS) is a feasible version of the approach above.

The goal of DMRS's procedure is to hedge the unpriced risk in heuristic factors. The first step is to construct hedging factors that go long in stocks with high loadings on the heuristic factors and short in stocks with low loadings, while holding constant the characteristics-exposure of the long and short legs of hedging factors, which ensures that they have zero

expected return according to Assumption 1. DMRS do this by sorting stocks by loadings on heuristic factors within characteristics-sorted portfolios. Here, we work with more general characteristics-based factors with weights $\mathbf{W}_t = \mathbf{X}_t \mathbf{S}_t$ and we construct a hedging portfolio that has precisely zero expected return by regressing conditional covariances of individual stocks with factors, i.e., $\boldsymbol{\Sigma}_t \mathbf{W}_t$, on \mathbf{X}_t , and then using the residuals,

$$\mathbf{W}_{h,t} = \mathbf{R}_t \boldsymbol{\Sigma}_t \mathbf{X}_t \mathbf{S}_t \quad (24)$$

as portfolio weights for hedge portfolios.

The second step is to calculate stocks' covariances with the hedge portfolio returns so that we can modify stocks' weights in the factor portfolios to remove unpriced risks:

$$\hat{\mathbf{V}}_t = \boldsymbol{\Sigma}_t \mathbf{W}_{h,t} = \mathbf{V}_t \boldsymbol{\Gamma}_t \mathbf{V}_t' \mathbf{R}_t \mathbf{V}_t \boldsymbol{\Gamma}_t \mathbf{V}_t' \mathbf{X}_t \mathbf{S}_t. \quad (25)$$

The third step is to regress the factor portfolio weights $\mathbf{W}_t = \mathbf{X}_t \mathbf{S}_t$ on $\hat{\mathbf{V}}_t$ to obtain residual factor portfolio weights that have been purged of unpriced risk exposure. Now note that $\hat{\mathbf{V}}_t$ in (25) is equal to \mathbf{V}_t post-multiplied by a nonsingular matrix. Hence regressing \mathbf{W}_t on $\hat{\mathbf{V}}_t$ produces the same residuals as regressing \mathbf{W}_t on \mathbf{V}_t . Therefore, the residuals ⁵

$$\hat{\mathbf{H}}_t = \mathbf{X}_t \mathbf{S}_t - \hat{\mathbf{V}}_t (\hat{\mathbf{V}}_t' \hat{\mathbf{V}}_t)^{-1} \hat{\mathbf{V}}_t' \mathbf{X}_t \mathbf{S}_t \quad (28)$$

are the same as the residuals in (23) and hence $\hat{\mathbf{H}}_t = \mathbf{H}_t$. In other words, the three steps

5. DMRS use a slightly different approach, but under the assumptions of Proposition 3 below, it yields the same hedged factors. They purge the heuristic factors from unpriced risks that do not earn expected return by regressing the J heuristic factors on the J hedge portfolio returns and using the J time series of residuals as the hedged factors. The $J \times J$ matrix of regression coefficients in these regressions is

$$\mathbf{K}_t = \mathbf{S}_t' \mathbf{X}_t' \mathbf{W}_{h,t} (\mathbf{W}_{h,t}' \boldsymbol{\Sigma}_t \mathbf{W}_{h,t})^{-1} \mathbf{S}_t^{-1}, \quad (26)$$

and so the hedged factors have weights

$$\hat{\mathbf{H}}_t = \mathbf{X}_t \mathbf{S}_t - \mathbf{W}_{h,t} \mathbf{K}_t'. \quad (27)$$

Substituting $\mathbf{W}_{h,t} = \mathbf{V}_t \mathbf{A}_t$, for some nonsingular matrix \mathbf{A}_t , into this expression and \mathbf{K}_t , it can be seen that this last expression is equivalent to (23).

above provide a way to construct the hedged portfolio weights in (23) from observable moments.

The following proposition states the result more formally.

Proposition 3 *If the matrices \mathbf{U}_t and $\mathbf{\Omega}_t$ in (20) are such that there exists a decomposition*

$$\mathbf{U}_t \mathbf{\Omega}_t \mathbf{U}_t' = \mathbf{V}_t \mathbf{\Gamma}_t \mathbf{V}_t' + \mathbf{E}_t \mathbf{\Phi}_t \mathbf{E}_t', \quad (29)$$

where \mathbf{V}_t is an $N \times J$ matrix of full column rank, $\mathbf{V}_t' \mathbf{X}_t$ is full rank, $\mathbf{R}_t \mathbf{V}_t$ has full column rank, $\mathbf{E}_t' \mathbf{X}_t = \mathbf{0}$, $\mathbf{E}_t' \mathbf{V}_t = \mathbf{0}$ and $\mathbf{\Gamma}_t$ is nonsingular, then the maximum squared conditional Sharpe ratio of the hedged factors $\mathbf{f}_{t+1} = \hat{\mathbf{H}}_t' \mathbf{z}_{t+1}$ with $\hat{\mathbf{H}}_t$ as defined in (28) is equal to the maximum squared conditional Sharpe Ratio of the individual assets.

Proof. Write $\hat{\mathbf{V}}_t = \mathbf{V}_t \mathbf{A}_t$ where $\mathbf{A}_t = \mathbf{\Gamma}_t \mathbf{V}_t' \mathbf{R}_t \mathbf{V}_t \mathbf{\Gamma}_t \mathbf{V}_t' \mathbf{X}_t \mathbf{S}_t$. By assumption, $\mathbf{R}_t \mathbf{V}_t$ has full column rank J , hence $\mathbf{V}_t' \mathbf{R}_t \mathbf{V}_t = \mathbf{V}_t' \mathbf{R}_t \mathbf{R}_t \mathbf{V}_t$ has full rank. Since pre- and post-multiplying this expression by full rank matrices $\mathbf{\Gamma}_t$ and $\mathbf{V}_t' \mathbf{X}_t \mathbf{S}_t$ does not change rank, it follows that \mathbf{A}_t is full rank and hence nonsingular. Then, substituting $\hat{\mathbf{V}}_t = \mathbf{V}_t \mathbf{A}_t$, with \mathbf{A}_t nonsingular into (28) yields the expression for \mathbf{H}_t in (23), i.e., $\hat{\mathbf{H}}_t = \mathbf{H}_t$. Then $\mathbf{U}_t' \hat{\mathbf{H}}_t = \mathbf{0}$ immediately follows. Therefore, by Lemma 2, the result follows. ■

The rank requirements for several matrices in Proposition 3 have an economic interpretation. That $\mathbf{V}_t' \mathbf{X}_t$ has full rank and $\mathbf{R}_t \mathbf{V}_t$ has full column rank ensures that the hedging portfolio weight vectors constructed via (25) and (28) are linearly independent. One could relax these rank requirements by building in a dimension-reduction step that removes linear dependencies in the construction of $\mathbf{W}_{h,t}$. However, for our purposes here, the benefits from greater generality of this approach would not be worth the costs of additional expositional complexity.

What do we gain from the hedging procedure? Comparing the conditions in Proposition 3 with (15) and (16) in Proposition 2, we can see that the conditions on the covariance matrix that are required to hold for the hedged factors to span the SDF are weaker than those required for the OLS factors (or nonsingular transformations thereof) to span the SDF. While Proposition 2 requires the columns of \mathbf{X}_t to be orthogonal to the columns of \mathbf{U}_t , the

conditions in Proposition 3 allow violations of this orthogonality condition as long as there are at most J linearly independent sources of such non-orthogonality as collected in the J columns of the matrix \mathbf{V}_t .

II.D. Iterated hedging

When \mathbf{V}_t has more than J columns, then the (infeasible) hedged factor construction based on the unobservable \mathbf{V}_t as in (23) still works as $\mathbf{H}'_t\mathbf{U}_t = \mathbf{0}$ still holds and $\mathbf{H}'_t\mathbf{X}_t$ still has full column rank, i.e., the conditions in Lemma 2 still hold. However, in this case the feasible hedged factor weights $\hat{\mathbf{H}}_t$ we construct in (28) are no longer equal to \mathbf{H}_t . The reason is that if we again construct $\hat{\mathbf{V}}_t$ as in (25), the J columns of $\hat{\mathbf{V}}_t$ now contain J linear combinations of the $2J$ columns in \mathbf{V}_t . Projection on $\hat{\mathbf{V}}_t$ therefore no longer produces the same residuals as a projection on \mathbf{V}_t .

However, by iterating on the hedging procedure, we can solve this problem. Repeating the hedging procedure by regressing individual stocks' conditional covariances with *hedged* factors, i.e., $\boldsymbol{\Sigma}_t\hat{\mathbf{H}}_t$, on \mathbf{X}_t and collecting the residuals $\mathbf{R}_t\boldsymbol{\Sigma}_t\hat{\mathbf{H}}_t$ analogous to (24), but here for hedged factors. Using these residuals as portfolio weights, and calculating the covariances of individual stocks with these portfolio returns, we get, in analogy to (25),

$$\hat{\mathbf{V}}_{2,t} = \mathbf{V}_t\boldsymbol{\Gamma}_t\mathbf{V}'_t\mathbf{R}_t\boldsymbol{\Sigma}_t\mathbf{V}_t\boldsymbol{\Gamma}_t\mathbf{V}'_t\hat{\mathbf{H}}_t, \quad (30)$$

where the only difference to (25) is that $\mathbf{X}_t\mathbf{S}_t$ was replaced by $\hat{\mathbf{H}}_t$. Note that $\hat{\mathbf{V}}_{2,t}$ is comprised of J linear combinations of the $2J$ columns of \mathbf{V}_t .

Under conditions that we state more formally shortly, $\hat{\mathbf{V}}_t$ and $\hat{\mathbf{V}}_{2,t}$ jointly span the same column space as \mathbf{V}_t . Therefore, the residuals from the regression of $\mathbf{X}_t\mathbf{S}_t$ on \mathbf{V}_t in (23) are the same as those from a regression of $\mathbf{X}_t\mathbf{S}_t$ on $\hat{\mathbf{V}}_t$ and $\hat{\mathbf{V}}_{2,t}$ jointly. And the latter regression can in turn be implemented in two steps, which results in an iterated hedging procedure. By the Frisch-Waugh-Lovell theorem, the residuals of a regression of $\mathbf{X}_t\mathbf{S}_t$ on $\hat{\mathbf{V}}_t$ and $\hat{\mathbf{V}}_{2,t}$ jointly are the same as the residuals of a regression of the first step residuals $\hat{\mathbf{H}}_t$

from regressing $\mathbf{X}_t \mathbf{S}_t$ on $\hat{\mathbf{V}}_t$ in (28) on the residuals from regressing $\hat{\mathbf{V}}_{2,t}$ on $\hat{\mathbf{V}}_t$. Therefore, we can construct the hedged portfolio weights as

$$\begin{aligned}\hat{\mathbf{H}}_{2,t} &= \mathbf{M}_t \mathbf{X}_t \mathbf{S}_t - \mathbf{M}_t \hat{\mathbf{V}}_{2,t} (\hat{\mathbf{V}}'_{2,t} \mathbf{M}_t \hat{\mathbf{V}}_{2,t})^{-1} \hat{\mathbf{V}}'_{2,t} \mathbf{M}_t \mathbf{X}_t \mathbf{S}_t \\ &= \hat{\mathbf{H}}_t - \mathbf{M}_t \hat{\mathbf{V}}_{2,t} (\hat{\mathbf{V}}'_{2,t} \mathbf{M}_t \hat{\mathbf{V}}_{2,t})^{-1} \hat{\mathbf{V}}'_{2,t} \mathbf{M}_t \hat{\mathbf{H}}_t,\end{aligned}\tag{31}$$

where $\mathbf{M}_t = \mathbf{I} - \hat{\mathbf{V}}_t (\hat{\mathbf{V}}'_t \hat{\mathbf{V}}_t)^{-1} \hat{\mathbf{V}}'_t$ is the residual maker matrix from regression on $\hat{\mathbf{V}}_t$, and we obtain $\hat{\mathbf{H}}_{2,t} = \mathbf{H}_t$.

The following proposition states this result formally. It looks similar to Proposition 3, but note that \mathbf{V}_t now has $2J$ columns.

Proposition 4 *If the matrices \mathbf{U}_t and $\mathbf{\Omega}_t$ in (20) are such that there exists a decomposition*

$$\mathbf{U}_t \mathbf{\Omega}_t \mathbf{U}'_t = \mathbf{V}_t \mathbf{\Gamma}_t \mathbf{V}'_t + \mathbf{E}_t \mathbf{\Phi}_t \mathbf{E}'_t,\tag{32}$$

where \mathbf{V}_t is an $N \times 2J$ matrix of full column rank, $\mathbf{R}_t \mathbf{V}_t$ has full column rank, $(\mathbf{V}'_t \mathbf{X}_t \mathbf{S}_t : \mathbf{V}'_t \hat{\mathbf{H}}_t)$ has full rank, with $\hat{\mathbf{H}}_t$ defined as in (28), $\mathbf{E}'_t \mathbf{X}_t = \mathbf{0}$, $\mathbf{E}'_t \mathbf{V}_t = \mathbf{0}$ and $\mathbf{\Gamma}_t$ is nonsingular, then the maximum squared conditional Sharpe ratio of the hedged factors $\mathbf{f}_{t+1} = \hat{\mathbf{H}}'_{2,t} \mathbf{z}_{t+1}$ with $\hat{\mathbf{H}}_{2,t}$ as defined in (31) is equal to the maximum squared Sharpe Ratio of the individual assets.

Proof. We first show $\hat{\mathbf{V}}_t$ and $\hat{\mathbf{V}}_{2,t}$ jointly span the same column space as \mathbf{V}_t . Note that we can write $(\hat{\mathbf{V}}_t : \hat{\mathbf{V}}_{2,t}) = \mathbf{V}_t \mathbf{G}_t \mathbf{A}_t$ with $\mathbf{A}_t = (\mathbf{V}'_t \mathbf{X}_t \mathbf{S}_t : \mathbf{V}'_t \hat{\mathbf{H}}_t)$ where $\mathbf{G}_t = \mathbf{\Gamma}_t \mathbf{V}'_t \mathbf{R}_t \mathbf{V}_t \mathbf{\Gamma}_t$ is a full-rank $2J \times 2J$ square matrix ($\mathbf{R}_t \mathbf{V}_t$ has full column rank, so $\mathbf{R}_t \mathbf{V}_t \mathbf{\Gamma}_t$ has rank $2J$). Premultiplying $\mathbf{R}_t \mathbf{V}_t \mathbf{\Gamma}_t$ with its own transpose then results in a matrix that is also of rank $2J$). Since \mathbf{A}_t and \mathbf{G}_t are full rank and hence invertible, we have $\mathbf{V}_t = (\hat{\mathbf{V}}_t : \hat{\mathbf{V}}_{2,t}) \mathbf{A}_t^{-1} \mathbf{G}_t^{-1}$, i.e., $\hat{\mathbf{V}}_t$ and $\hat{\mathbf{V}}_{2,t}$ jointly span the same column space as \mathbf{V}_t . Substituting this relation into (23), we obtain the residuals of a regression of $\mathbf{X}_t \mathbf{S}_t$ on $(\hat{\mathbf{V}}_t : \hat{\mathbf{V}}_{2,t})$. By the Frisch-Waugh-Lovell theorem, these residuals are in turn identical to those in the regression of $\mathbf{M}_t \mathbf{X}_t \mathbf{S}_t$ on $\mathbf{M}_t \hat{\mathbf{V}}_{2,t}$ in (31). Hence $\hat{\mathbf{H}}_{2,t} = \mathbf{H}_t$ and so $\hat{\mathbf{H}}'_{2,t} \mathbf{U}_t = \mathbf{0}$. Therefore, by Lemma 2, the result follows. ■

In analogy to the case with a single round of hedging that we discussed following Proposition 3, the rank requirements for several matrices in Proposition 4 have an economic inter-

pretation. The requirements that $(\mathbf{V}'_t \mathbf{X}_t \mathbf{S}_t : \mathbf{V}'_t \hat{\mathbf{H}}_t)$ has full rank and $\mathbf{R}_t \mathbf{V}_t$ has full column rank are both needed to ensure that iterated hedging factor portfolio weight vectors $\hat{\mathbf{H}}_{2,t}$ are linearly independent. One could again relax these rank requirements by building dimension-reduction steps that removes linear dependencies in the iterated hedging procedure.

What do we gain from iterated hedging? Comparing the conditions in Proposition 4 with those in Proposition 3, we can see that those in Proposition 4 are weaker. While the conditions in Proposition 3 allow for J linearly independent sources of such non-orthogonality of \mathbf{X}_t and the columns of \mathbf{U}_t , the conditions in Proposition 4 allow for $2J$ linearly independent sources of such non-orthogonality. In other words, iterated hedging can remove more sources of unpriced risk contamination in characteristics-based factors than a single round of hedging can.

There is no reason to necessarily stop after a second round of hedging. We do not show formal results on this, but from the logic of the hedging iteration above, it should be clear that further rounds of hedging would remove additional sources of unpriced risk contamination. When working with population moments, this should further raise the maximum squared conditional Sharpe Ratio of the hedged factors and hence get them closer to spanning the SDF. Whether this is also true in a finite sample with estimated moments is not clear. At some point, further hedging may be counterproductive and bring in estimation error contamination rather than removing unpriced risk contamination. After all, doing many iterations of the hedging procedure should be no different than constructing GLS factors by inverting an estimate of the conditional covariance matrix (which is unlikely to work well when N is not small relative to T). We investigate this further in Section V.

II.E. Summary

When conditional expected returns are linear in firm characteristics, aggregation of individual stocks into factor portfolios leads to a deterioration of the investment opportunity set, and hence a failure of the factors to span the SDF, unless the factor weights incorporate the con-

ditional covariance matrix as in GLS cross-sectional regression slope factors, or, alternatively, the conditional covariance matrix satisfies certain conditions. Methods for hedging unpriced risks in factors allow a partial relaxation of these conditions, especially if hedging procedures are applied iteratively.

III. DIMENSIONALITY REDUCTION

So far we have discussed factor models where the pricing information in J characteristics is captured by J factors in the SDF. As we show now, under certain conditions on the covariance matrix of individual stock returns, one can summarize the pricing information in J characteristics-based factors in a smaller number of $K < J$ factors. Of course, there is always a single factor that prices the individual assets (which the linear combination of J factors shown in Proposition 1), but without further assumptions, the construction of this single factor requires inversion of a large conditional covariance matrix. The point of the methods we discuss in this section is to achieve dimension reduction without having to invert or eigen-decompose the large covariance matrix of individual asset returns.

We first present general conditions on covariance matrix that need to hold such that dimension reduction is possible without loss of pricing information. Then we show that, under these conditions, various approaches that have appeared in the literature are actually equivalent or closely related.

Corollary 2 *Suppose expected returns are given by*

$$\boldsymbol{\mu}_t = \mathbf{X}_t \mathbf{Q}_t \boldsymbol{\phi}, \tag{33}$$

where \mathbf{Q}_t is a $J \times K$ matrix with $K \leq J$, and let $\mathbf{W}_t = \mathbf{X}_t \mathbf{Q}_t \mathbf{S}_t$. Then the maximum squared conditional Sharpe ratio of the factors $\mathbf{f}_{t+1} = \mathbf{W}_t' \mathbf{z}_{t+1}$, for any nonsingular $K \times K$ matrix \mathbf{S}_t , is equal to the maximum squared Sharpe Ratio of the individual assets if and only if there

exist conformable matrices $\mathbf{\Lambda}_t$, $\mathbf{\Omega}_t$, and a matrix \mathbf{U}_t for which

$$\mathbf{U}'_t \mathbf{X}_t \mathbf{Q}_t = \mathbf{0}, \quad (34)$$

such that

$$\mathbf{\Sigma}_t = \mathbf{X}_t \mathbf{Q}_t \mathbf{\Lambda}_t \mathbf{Q}'_t \mathbf{X}'_t + \mathbf{U}_t \mathbf{\Omega}_t \mathbf{U}'_t. \quad (35)$$

Proof. Directly follows from Proposition 2 by using $\mathbf{X}_t \mathbf{Q}_t$ in place of \mathbf{X}_t . ■

We have achieved dimension reduction because there are now K factors in \mathbf{f} , not J . This is made possible by the fact that the factor component of the covariance matrix related to \mathbf{X}_t is now a lower-dimensional $\mathbf{X}_t \mathbf{Q}_t$, which is $N \times K$, with $K \leq J$, rather than the larger $N \times J$ matrix \mathbf{X}_t that we had in Proposition 2. And $\mathbf{\Lambda}_t$ is a $K \times K$ matrix rather than the $J \times J$ matrix $\mathbf{\Psi}_t$ in Proposition 2.

How can we find \mathbf{Q}_t to construct the factors \mathbf{f} ? As we show now, if we make a somewhat stronger assumption than (34), namely that $\mathbf{U}'_t \mathbf{X}_t = \mathbf{0}$ we can obtain \mathbf{Q}_t through principal component analysis (PCA). Under this assumption, OLS factors, for instance, and rotations thereof span the SDF. PCA applied to OLS factors can then extract \mathbf{Q}_t . More precisely, to extract \mathbf{Q}_t as principal components, we need to add additional identification assumptions on \mathbf{Q}_t and $\mathbf{\Lambda}_t$. These assumptions pin down a specific rotation of \mathbf{Q}_t , but they do not affect the pricing implications of the factor model. With different choices of identifying assumptions, we then obtain conditional versions of two recently proposed methods of dimension-reduced factor construction.

Example 5 (IPCA) Suppose $\mathbf{U}'_t \mathbf{X}_t = \mathbf{0}$, $\mathbf{Q}'_t \mathbf{Q}_t = \mathbf{I}$ and $\mathbf{\Lambda}_t$ is diagonal with descending diagonal entries.⁶ We can then obtain \mathbf{Q}_t and $\mathbf{\Lambda}_t$ from an eigendecomposition of the conditional covariance matrix of OLS factor returns, because it factors as

$$(\mathbf{X}'_t \mathbf{X}_t)^{-1} \mathbf{X}'_t \mathbf{\Sigma}_t \mathbf{X}_t (\mathbf{X}'_t \mathbf{X}_t)^{-1} = \mathbf{Q}_t \mathbf{\Lambda}_t \mathbf{Q}'_t, \quad (36)$$

6. The last two assumptions correspond to identification assumption in Kelly, Pruitt, and Su (2019): $\mathbf{\Gamma}'_\beta \mathbf{\Gamma}_\beta = \mathbf{I}_K$ and $\text{cov}(\mathbf{f}_t)$ has only descending diagonal entries (their notation).

where \mathbf{Q}_t , given the assumptions above, becomes a matrix of eigenvectors of this covariance matrix associated with the K non-zero eigenvalues. Suppose further that $\mathbf{S}_t = (\mathbf{Q}'\mathbf{X}'_t\mathbf{X}_t\mathbf{Q})^{-1}$. Then we obtain a conditional version of the IPCA factors of Kelly, Pruitt, and Su (2019):

$$\mathbf{f}_{IPCA,t+1} = (\mathbf{Q}'_t\mathbf{X}'_t\mathbf{X}_t\mathbf{Q}_t)^{-1}\mathbf{Q}'_t\mathbf{X}'_t\mathbf{z}_{t+1}. \quad (37)$$

The expression for $\mathbf{f}_{IPCA,t+1}$ in (37) above is a conditional version of the first of two first-order conditions in Kelly, Pruitt, and Su (2019) that define the instrumented principal components analysis (IPCA) estimator. We can also show that a conditional version of their second first-order condition (their eq. 7) holds in terms of population moments. If it holds, then the right-hand side their second first-order condition should equal $\text{vec}(\mathbf{Q}_t)$ when evaluated with the factors in (37) and under the conditions of Corollary 2. Evaluating their second first-order condition, this is indeed what we obtain:

$$\begin{aligned} & (\mathbf{X}'_t\mathbf{X}_t \otimes \mathbb{E}_t[\mathbf{f}_{t+1}\mathbf{f}'_{t+1}])^{-1} \mathbb{E}_t [(\mathbf{X}'_t \otimes \mathbf{f}_{t+1}) \mathbf{z}_{t+1}] \\ &= (\mathbf{X}'_t\mathbf{X}_t \otimes \mathbb{E}_t[\mathbf{f}_{t+1}\mathbf{f}'_{t+1}])^{-1} \text{vec}(\mathbf{E}_t[\mathbf{f}_{t+1}\mathbf{z}'_{t+1}]\mathbf{X}_t) \\ &= \text{vec} \left(\mathbb{E}_t[\mathbf{f}_{t+1}\mathbf{f}'_{t+1}]^{-1} \mathbf{E}_t[\mathbf{f}_{t+1}\mathbf{z}'_{t+1}]\mathbf{X}_t (\mathbf{X}'_t\mathbf{X}_t)^{-1} \right) \\ &= \text{vec}(\mathbf{Q}_t), \end{aligned} \quad (38)$$

where for the last step we evaluated the conditional expectations using (37), (33), (35), and (34). Hence, factors constructed as in (37) with \mathbf{Q}_t obtained as eigenvectors of the OLS factor return covariance matrix in (36) solve both first-order conditions, i.e., they are indeed the IPCA factors.⁷

Kelly, Pruitt, and Su (2019) show that in the case of orthonormalized characteristics, IPCA is equivalent to PCA on returns managed portfolios with weights \mathbf{X}'_t . Our result here

7. The assumption of time-constant \mathbf{Q}_t and $\mathbf{\Lambda}_t$ can justify working with a constant \mathbf{Q} extracted from an average conditional, or approximately unconditional, covariance matrix. Working through the first-order condition in (38) expressed in terms of unconditional expectations (the population analog to the sample

shows that IPCA is more generally equivalent to PCA on managed portfolios, even in the case where characteristics are not orthonormalized, if the managed portfolios are constructed as OLS factors. In particular, applying PCA to OLS portfolios recovers \mathbf{Q}_t . By applying this matrix to univariate portfolios $\mathbf{X}'_t \mathbf{z}_{t+1}$ and further rotating them by an OLS factor (as in (37)), yields our version of the IPCA estimator.

The OLS factor population covariance matrix that we apply PCA to in (36) is singular if $K < J$ as it is a $J \times J$ matrix with only K non-zero eigenvalues. The matrices $\mathbf{\Lambda}_t$ and \mathbf{Q}_t in our notation contain only the non-zero eigenvalues and the eigenvectors associated with the non-zero eigenvalues. With an estimated covariance matrix in a finite sample, the truly zero eigenvalues would not be exactly zero but likely very small.

Example 6 (PPCA) Suppose $\mathbf{U}'_t \mathbf{X}_t = \mathbf{0}$, $\mathbf{Q}'_t \mathbf{X}'_t \mathbf{X}_t \mathbf{Q}_t = \mathbf{I}$ and $\mathbf{\Lambda}_t$ is diagonal with descending diagonal entries.⁸ We can then obtain \mathbf{Q}_t and $\mathbf{\Lambda}_t$ from an eigendecomposition of the conditional covariance matrix of univariate factor returns constructed using orthonormalized characteristics, because it factors as

$$(\mathbf{X}'_t \mathbf{X}_t)^{-\frac{1}{2}} \mathbf{X}'_t \boldsymbol{\Sigma}_t \mathbf{X}_t (\mathbf{X}'_t \mathbf{X}_t)^{-\frac{1}{2}} = (\mathbf{X}'_t \mathbf{X}_t)^{\frac{1}{2}} \mathbf{Q}_t \mathbf{\Lambda}_t \mathbf{Q}'_t (\mathbf{X}'_t \mathbf{X}_t)^{\frac{1}{2}}, \quad (39)$$

where $\mathbf{G}_t = (\mathbf{X}'_t \mathbf{X}_t)^{\frac{1}{2}} \mathbf{Q}_t$ is orthonormal by assumption and thus can be recovered as a matrix of eigenvectors of this covariance matrix associated with the K non-zero eigenvalues. We get $\mathbf{Q}_t = (\mathbf{X}'_t \mathbf{X}_t)^{-\frac{1}{2}} \mathbf{G}_t$. Suppose further that $\mathbf{S}_t = \mathbf{I}$. Then we obtain a conditional version of

averages in KPS), we the obtain $\text{vec}(\mathbf{Q})$:

$$\begin{aligned} & [\mathbb{E} (\mathbf{X}'_t \mathbf{X}_t \otimes \mathbb{E}_t [\mathbf{f}_{KPS,t+1} \mathbf{f}'_{KPS,t+1}])]^{-1} \mathbb{E} [(\mathbf{X}'_t \otimes \mathbf{f}_{KPS,t+1}) \mathbf{z}_{t+1}] \\ &= [\mathbb{E} (\mathbf{X}'_t \mathbf{X}_t) \otimes \mathbf{\Lambda}]^{-1} \text{vec} (\mathbb{E} [\mathbb{E}_t [\mathbf{f}_{KPS,t+1} \mathbf{z}'_{t+1}] \mathbf{X}_t]) \\ &= [\mathbb{E} (\mathbf{X}'_t \mathbf{X}_t) \otimes \mathbf{\Lambda}]^{-1} \text{vec} (\mathbf{\Lambda} \mathbf{Q} \mathbb{E} (\mathbf{X}'_t \mathbf{X}_t)) = \text{vec}(\mathbf{Q}). \end{aligned}$$

8. The last two assumptions correspond to identification assumption in Kim, Korajczyk, and Neuhierl (2019): $\mathbf{G}'_{\beta}(\mathbf{X}_t) \mathbf{G}_{\beta}(\mathbf{X}_t) = \boldsymbol{\Theta}'_{\beta,t} \mathbf{X}'_t \mathbf{X}_t \boldsymbol{\Theta}_{\beta,t} = \mathbf{I}_K$ and $\text{cov}(\mathbf{f}_t)$ has only descending diagonal entries (their notation).

the PPCA factors of Kim, Korajczyk, and Neuhierl (2019):

$$\mathbf{f}_{PPCA,t+1} = \mathbf{G}'_t (\mathbf{X}'_t \mathbf{X}_t)^{-\frac{1}{2}} \mathbf{X}'_t \mathbf{z}_{t+1} \quad (40)$$

$$= (\mathbf{Q}'_t \mathbf{X}'_t \mathbf{X}_t \mathbf{Q}_t)^{-1} \mathbf{Q}'_t \mathbf{X}'_t \mathbf{z}_{t+1} = \mathbf{Q}'_t \mathbf{X}'_t \mathbf{z}_{t+1}. \quad (41)$$

The expression for $\mathbf{f}_{PPCA,t+1}$ in (41) above is a conditional version of the factors in Kim, Korajczyk, and Neuhierl (2019) obtained from a cross-sectional regression of stock returns on their factor loadings $\mathbf{G}_\beta(\mathbf{X}_t)$ which we parameterize as linear here, $\mathbf{G}_\beta(\mathbf{X}_t) = \mathbf{X}_t \mathbf{Q}_t$. To see this, note that Kim, Korajczyk, and Neuhierl (2019) identify $\mathbf{G}_\beta(\mathbf{X}_t)$ via a PCA on projected returns, $\mathbf{X}_t (\mathbf{X}'_t \mathbf{X}_t)^{-1} \mathbf{X}'_t \mathbf{z}_{t+1}$. Under our assumption in (35), their covariance matrix is equal to $\mathbf{X}_t \mathbf{Q}_t \mathbf{\Lambda}_t \mathbf{Q}'_t \mathbf{X}'_t$. Because $\mathbf{X}_t \mathbf{Q}_t$ is orthonormal, their PCA solution, therefore, recovers $\mathbf{G}_\beta(\mathbf{X}_t) = \mathbf{X}_t \mathbf{Q}_t$ and their factors match ours in (41). The expression in (40) shows that we can alternatively identify these factors via a simple PCA on univariate portfolio returns (rather than projected individual stock returns) constructed using orthonormalized characteristics, to obtain \mathbf{Q}_t .

Our assumptions here are closely related to those in Kim, Korajczyk, and Neuhierl (2019). Their assumption 2 (ii) states that factor model residuals and \mathbf{X}_t are, asymptotically, cross-sectionally orthogonal. Their identifying assumptions 3 are the same as our identifying assumption in the example. Our assumption that $\mathbf{U}' \mathbf{X}_t = \mathbf{0}$ is the population version of this assumption. Lastly, our normalization assumption in Example 6 are analogous to theirs.

Overall, the results in this section show that there is a great deal of similarity in seemingly different recently proposed methods for dimension reduction. Our earlier results on the conditions required for characteristics-based factors to span the SDF provide a basis to get to these dimension-reduction in a straightforward way by applying PCA to a certain set of portfolio sorts.

IV. EXTENSIONS

Before turning to an empirical analysis, we first discuss a number of conceptual issues that come up if we want to relate our results from the previous sections to empirical data.

IV.A. Alternative assumptions about expected returns

As we discussed, our Assumption 1 that conditional expected returns are linear in characteristics is, in principle, completely general as for any given set of basis characteristics, one could define \mathbf{X}_t as including nonlinear functions and interactions of these basis characteristics. That said, once a researcher has settled on a particular set of characteristics to include in \mathbf{X}_t , the linearity assumption has economic content. For this reason, one may want to entertain alternative assumptions that link a specific characteristics matrix \mathbf{X}_t to $\boldsymbol{\mu}_t$.

For example, within a framework in which characteristics predict returns because of mispricing, our baseline Assumption 1 can be reasonable if the characteristics in \mathbf{X}_t are directly related to the magnitude of mispricing. How this could be true is easiest to see for scaled price ratios like the book-to-market ratio. If the numerator (book value) controls for differences across stocks in their fundamental scale and the remaining price variation that comes in through the denominator (market value) captures mispricing.

However, an alternative view may be that characteristics in \mathbf{X}_t capture not the magnitude of mispricing directly but rather investor sentiment, and hence sentiment investors demand for certain types of stocks. If these sentiment investors trade against mean-variance arbitrageurs, the portfolio optimization of the arbitrageurs induces cross-dependencies across expected returns and covariances that can result in equilibrium expected returns that differ from Assumption 1 (for this given \mathbf{X}_t). To illustrate, consider a CARA-normal model as in Kozak, Nagel, and Santosh (2018) where a measure $(1 - \theta)$ of rational arbitrageurs have demand $\frac{1}{a}\boldsymbol{\Sigma}_t^{-1}\boldsymbol{\mu}_t$ and a measure θ of sentiment investors have demand in excess of rational investor demand of $\mathbf{X}_t\mathbf{d}$ for some vector \mathbf{d} , i.e., $\frac{1}{a}\boldsymbol{\Sigma}_t^{-1}\boldsymbol{\mu}_t + \mathbf{X}_t\mathbf{d}$. With total asset supply of

one for each asset, collected in vector $\boldsymbol{\iota}$, market clearing implies

$$\boldsymbol{\mu}_t = a\boldsymbol{\Sigma}_t(\boldsymbol{\iota} - \theta\mathbf{X}_t\mathbf{d}) = \boldsymbol{\Sigma}_t\mathbf{X}_t\boldsymbol{\phi}, \quad (42)$$

for some vector $\boldsymbol{\phi}$, where the last equality follows because \mathbf{X}_t includes a column of ones. Thus, in this case instead of Assumption 1, we would have

Assumption 2

$$\boldsymbol{\mu}_t = \boldsymbol{\Sigma}_t\mathbf{X}_t\boldsymbol{\phi} \quad (43)$$

with some $J \times 1$ vector $\boldsymbol{\phi}$.

A closely related assumption appears in Brandt, Santa-Clara, and Valkanov (2009). They assume that mean-variance efficient portfolio weights are linear in characteristics and market portfolio weights, while here Assumption 2 implies that the weights $\boldsymbol{\Sigma}_t^{-1}\boldsymbol{\mu}_t = \mathbf{X}_t\boldsymbol{\phi}$ are linear in characteristics.

The SDF in this case is spanned by GLS factors from GLS cross-sectional regression of z_{t+1} on $\boldsymbol{\Sigma}_t\mathbf{X}_t$, or rotations of these factors. We can obtain these factors by replacing \mathbf{X}_t in Proposition 1 with $\boldsymbol{\Sigma}_t\mathbf{X}_t$ everywhere. We get factors

$$\mathbf{f}_{t+1} = \mathbf{S}'_t\mathbf{X}'_t z_{t+1}, \quad (44)$$

i.e., the GLS factors simplify to univariate factors or rotations thereof (e.g., OLS factors with $\mathbf{S}_t = (\mathbf{X}'_t\mathbf{X}_t)^{-1}$). In other words, one can construct factors that span the SDF solely based on the information in characteristics. No information about $\boldsymbol{\Sigma}_t$ is required to construct these factors! Unfortunately, as we see in the following example that summarizes the univariate factor case, conditional factor means and betas vary over time with $\boldsymbol{\Sigma}_t$, which renders empirical implementation difficult without further assumptions.

Example 7 Suppose $\mathbf{S}_t = \mathbf{I}_J$. We then obtain an SDF with factors $\mathbf{f}_{t+1} = \mathbf{X}'_t z_{t+1}$. Factor

means and covariances are time-varying, $\boldsymbol{\mu}_{f,t} = \mathbf{X}'_t \boldsymbol{\Sigma}_t \mathbf{X}_t \boldsymbol{\phi}$, $\boldsymbol{\Sigma}_{f,t} = \mathbf{X}'_t \boldsymbol{\Sigma}_t \mathbf{X}_t$, and factor risk prices are constant: $\mathbf{b}_t = \boldsymbol{\phi}$. Factor betas $\boldsymbol{\beta}_t = (\mathbf{X}'_t \boldsymbol{\Sigma}_t \mathbf{X}_t)^{-1} \boldsymbol{\Sigma}_t \mathbf{X}_t$ are varying with $\boldsymbol{\Sigma}_t$.

Under Assumption 2 dimension reduction works in the same way and under the same conditions on the covariance matrix as in Corollary 2.⁹

Whether Assumption 2 or Assumption 1 is more appropriate once a researcher has settled on a specification of \mathbf{X}_t is an empirical question. We return to our baseline Assumption 1 for the rest of this section.

IV.B. Conditioning down

To conduct empirical work, our results in terms of conditional moments are not straightforward to work with. In empirical implementation, researchers often like to work with unconditional pricing restrictions and unconditional moments as estimating conditional moments requires additional assumptions for modeling the dynamics of conditional moments.

For this purpose, it is convenient if a model implies that factors' conditional expected returns are constant and either conditional factor betas or factor prices of risk are also constant or depend only on the observable characteristics \mathbf{X}_t (and not on $\boldsymbol{\Sigma}_t$). For example, if $\boldsymbol{\beta} = \mathbf{X}_t$ and $\boldsymbol{\mu}_{f,t} = \boldsymbol{\mu}_f$, one can implement the factor model in its conditional beta pricing formulation and then condition down to

$$\mathbb{E}[\mathbf{z}_{t+1}] = \mathbb{E}[\mathbf{X}_t] \boldsymbol{\mu}_f. \quad (45)$$

Alternatively, if $\mathbf{b}_t = \boldsymbol{\phi}$ and $\boldsymbol{\mu}_{f,t} = \boldsymbol{\mu}_f$, we have an SDF

$$M_{t+1} = 1 - \boldsymbol{\phi}'(\mathbf{z}_{t+1} - \boldsymbol{\mu}_f) \quad (46)$$

9. In this case, we don't need the additional assumption about expected returns in (33) because its expected returns automatically inherit the lower-dimensional structure through their dependence on the covariance matrix in Assumption 2.

which we can rescale to

$$M_{t+1} = 1 - \frac{\phi'}{1 - \phi' \boldsymbol{\mu}_f} \mathbf{z}_{t+1} \quad (47)$$

without affecting the pricing implications for excess returns. In this formulation, one can estimate the J constant prices of risk $\mathbf{b} = \frac{\phi'}{1 - \phi' \boldsymbol{\mu}_f}$ from the J unconditional pricing restrictions $\mathbb{E}[M_{t+1} \mathbf{f}_{t+1}] = 0$ without having to model conditional moments.

Recall that our earlier results in Section II expressed factors up to a rotation by a nonsingular matrix \mathbf{S}_t . We can choose this matrix to generate factors with the desired conditioning-down properties.

Consider first the GLS factors and their rotations. The case we presented in Example 1 with $\mathbf{S}_t = (\mathbf{X}'_t \boldsymbol{\Sigma}_t^{-1} \mathbf{X}_t)^{-1}$ yields $\boldsymbol{\beta}_t = \mathbf{X}_t$ and $\boldsymbol{\mu}_{f,t} = \boldsymbol{\phi}$, so the beta-pricing formulation conditions down nicely, but there is no \mathbf{S}_t that produces both prices of risk that do not depend on $\boldsymbol{\Sigma}_t$ and factor means that do not depend on $\boldsymbol{\Sigma}_t$. As a consequence, there does not exist a version of \mathbf{S}_t that would yield an SDF that we could estimate without having to model $\boldsymbol{\Sigma}_t$.

For OLS factors and their rotations, the case in Example 3 with $\mathbf{S}_t = (\mathbf{X}'_t \mathbf{X}_t)^{-1}$ yields $\boldsymbol{\beta}_t = \mathbf{X}_t$ and $\boldsymbol{\mu}_{f,t} = \boldsymbol{\phi}$, so again the beta-pricing formulation conditions down nicely, but there is no \mathbf{S}_t that produces both prices of risk that do not depend on $\boldsymbol{\Sigma}_t$ and factor means that do not depend on $\boldsymbol{\Sigma}_t$. As a consequence, without additional assumptions, there does not exist a version of \mathbf{S}_t that would yield an SDF that we could estimate without having to model $\boldsymbol{\Sigma}_t$.

Under the alternative Assumption 2 about expected returns in Section IV.A, too, there is no specification of \mathbf{S}_t that produces, at the same time, prices of risk that do not depend on $\boldsymbol{\Sigma}_t$ and factor means that do not depend on $\boldsymbol{\Sigma}_t$.

IV.C. Orthonormalized characteristics

Empirical work often considers characteristics that are normalized in some fashion. For example, portfolio sorting procedures use only information about cross-sectional ranks of

stocks by characteristics, not the value of the characteristics themselves; other methods transform characteristics into cross-sectional ranks and use the rank-transformed characteristics as portfolio weights (Kozak, Nagel, and Santosh 2020); further alternatives include orthonormalizing characteristics such that, after orthonormalization, $\mathbf{X}'_t\mathbf{X}_t = \mathbf{I}_J$ holds. Common to these methods is that, to varying degrees, they remove time-series variation from characteristics. For example, if the original characteristics matrix includes a column of ones as first column, and characteristics are then orthonormalized using the Gram-Schmidt process, this cross-sectionally demeans all characteristics and removes time-series variation in their cross-sectional variances and correlations.

IV.D. Conditioning down with normalized characteristics

We now show that constructing factors based on such normalized characteristics can be advantageous because the requirements we discussed in Section IV.B for unconditional pricing restrictions to imply an SDF with constant factor prices of risk and constant factor means.

However, before we can discuss conditioning down the pricing relationship to unconditional moments, we first need to deal with the fact that if Assumption 1 holds for a given set of original characteristics, it does not necessarily hold for the normalized version of these characteristics. While it is, in the end, an empirical question whether it holds for the normalized or original characteristics, there are plausible reasons to think that it could hold for the former if it holds for the latter. To see why, let's focus on the case of orthonormalization and let \mathbf{C}_t be the original characteristics matrix and $\mathbf{X}_t = \mathbf{C}_t\mathbf{N}_t^{-1}$ the normalized one, with $\mathbf{N}_t = (\mathbf{C}'_t\mathbf{C}_t)^{\frac{1}{2}}$. What is needed, roughly, is that the normalized characteristics do not contain information about the time-variation in cross-sectional mean, dispersion, or correlation of characteristics that the normalization has removed. More precisely, we need that

$$\mathbb{E}[\mathbf{N}_t|\mathbf{X}_t] = \mathbf{N} \tag{48}$$

for some constant matrix \mathbf{N} . If this relationship holds, and Assumption 1 holds for the

original characteristics, i.e., $\mathbb{E}[z_{t+1}|\mathbf{C}_t] = \mathbf{C}_t\phi_C$, then,

$$\mathbb{E}[z_{t+1}|\mathbf{X}_t] = \mathbb{E}\{\mathbb{E}[z_{t+1}|\mathbf{C}_t]|\mathbf{X}_t\} = \mathbf{X}_t\phi, \quad \phi = \mathbf{N}\phi_C. \quad (49)$$

i.e., we see that the relationship between characteristics and conditional expected excess returns remains linear with constant coefficients ϕ . In this case, GLS factors constructed based on the normalized characteristics price perfectly all assets conditional on \mathbf{X}_t . The maximum squared Sharpe ratio attainable conditional on \mathbf{X}_t may be lower than conditional on \mathbf{C}_t , but all of our earlier analysis of the conditions for OLS factors to span the SDF, for factor hedging, and dimension reduction then go through based on the normalized characteristics with conditional moments conditioned on \mathbf{X}_t .

Normalization of characteristics can be useful if we wish to condition down to unconditional pricing restrictions and obtain an SDF with constant factor prices of risk and constant factor means. For instance, purging characteristics of information about time-varying cross-sectional mean, dispersion, or correlation of characteristics, removes much of the information that in characteristics that could be related to time-variation in Σ_t . As a consequence, relatively mild assumptions suffice to obtain constant factor prices of risk and constant factor means.

Based on orthonormalized characteristics, the OLS factors in Example 3 have means $\mu_{f,t} = \phi$ and prices of risk $\mathbf{b}_t = \Psi_t^{-1}\phi$. So time-variation in Ψ_t is the only remaining source of time-variation in the prices of risk. With orthonormalized characteristics, the assumption that Ψ_t is constant is a relatively weak one. Recall that all conditional moments in our analysis, including Ψ_t , are conditioned on \mathbf{X}_t . Since orthonormalization removes variation over time in the average value of characteristics, their dispersion, and their correlation, there may not be much information left in characteristics that captures time-variation in Ψ . Therefore, conditional on the normalized characteristics \mathbf{X}_t , Ψ_t could be constant even it is not constant conditional \mathbf{C}_t .

For example, consider book-to-market equity ratios. Before normalization, the average

book-to-market ratio across firms may have time-series variation that is informative about time-variation in conditional covariances Ψ_t . Orthonormalization removes this common variation. Similarly, before normalization, book-to-market ratios may have time-varying cross-sectional dispersion that is informative about time-variation in Ψ_t . Orthonormalization removes this information. There could potentially still be some information in, say, the cross-sectional ordering of firms by characteristics each period that could contain information about time-varying in Ψ_t , but it seems likely that orthonormalizing removes most of the variation in characteristics that could be informative about time-variation in Ψ_t .

If Ψ_t is indeed constant conditional on the orthonormalized characteristics, then prices of risk are constant, $\mathbf{b}_t = \mathbf{b}$, and hence the SDF

$$M_{t+1} = 1 - \mathbf{b}'(\mathbf{f}_{t+1} - \boldsymbol{\mu}_f) \tag{50}$$

can be estimated from unconditional pricing restrictions and without estimating a conditional covariance matrix. Thus, normalization combined with a relatively weak assumption about Ψ_t makes it possible to use standard estimation approaches that rely on unconditional moments.

IV.E. Testing

We close this section with a few remarks on testing. The previous analysis made clear that heuristic factor models, such as OLS factors, only span the SDF when the conditional covariance matrix satisfies certain conditions. How can we let the data tell us whether these conditions hold? Going into the sampling theory of estimation and testing is beyond the scope of this paper. Instead, we will highlight population moment conditions that reveal misspecification (and ones that do not). We focus our discussion on OLS factors.

It may seem straightforward to test an OLS factor model. Let \mathbf{f}_{t+1} denote the OLS factors from Example 3. In this case we have observable conditional betas $\boldsymbol{\beta}_t = \mathbf{X}_t$ and

constant factor means $\boldsymbol{\mu}_{f,t} = \boldsymbol{\phi}$. Therefore, it may seem that one could just evaluate whether

$$\mathbb{E}[\mathbf{z}_{t+1}] = \mathbb{E}[\mathbf{X}_t \mathbf{f}_{t+1}] \quad (51)$$

holds in the data. In fact, this is what Fama and French (2020) do in their empirical work when they evaluate an OLS factor model. However, testing the equality (51) just tests whether there is a linear relation between characteristics and expected returns as stated in Assumption 1. If Assumption 1 holds, the equality (51) is true irrespective of whether the conditions in Proposition 2 for OLS factors to span the SDF hold or not. To see this, note that $\mathbb{E}[\mathbf{X}_t \mathbf{f}_{t+1}] = \mathbb{E}[\mathbf{X}_t (\mathbf{X}'_t \mathbf{X}_t)^{-1} \mathbf{X}'_t \mathbf{z}_{t+1}] = \mathbb{E}[\mathbf{X}_t \boldsymbol{\phi}] = \mathbb{E}[\mathbf{z}_{t+1}]$ by Assumption 1. So testing the equality (51) is not a test of the OLS factor asset pricing model.

The key here is that misspecification due to the conditional covariance matrix not satisfying the conditions in Proposition 2 would show up as $\boldsymbol{\beta}_t$ deviating from \mathbf{X}_t . By assuming $\boldsymbol{\beta}_t = \mathbf{X}_t$, the approach of Fama and French (2020) assumes away any misspecification of the SDF.

One way to testing for misspecification is to construct hedged factors as in Sections II.C and II.D. If the hedged factors achieve a higher Sharpe ratio than the OLS factors, the OLS factors do not span the SDF. We implement this approach empirically in the next Section.

V. EMPIRICAL ANALYSIS

Our analysis so far provides conditions on the conditional covariance matrix of individual stock returns under which OLS factors (and rotations thereof) span the SDF, as well as conditions under which dimension-reduction via principal components analysis of OLS factor portfolios yields the same factors as IPCA. Do these conditions hold empirically for various combinations of characteristics-based factors used in the prior literature?

Directly answering this question by comparing the MVE portfolio's (maximum) Sharpe ratio of OLS factors to the maximum Sharpe ratio obtainable with GLS factors is difficult

because constructing GLS factors would require the estimation and inversion of a large conditional covariance matrix for an unbalanced panel of thousands of stocks. Instead, we use our earlier results on iterated factor hedging. If a set of OLS factors does not span the SDF because the conditional covariance matrix does not satisfy the assumptions required for SDF spanning, then hedging the factors should improve the maximum Sharpe ratio. If a set of OLS factors already spans the SDF, then factor hedging should not improve the maximum Sharpe ratio of the factors. In fact, empirically, with estimated moments that are contaminated with estimation error, factor hedging might lead to a deterioration in the Sharpe ratio.

V.A. Data and factor construction

We use rank-transformed stock characteristics from Kozak (2019) and daily stock returns from November 1973 to the end of 2020. We apply several filters to preserve characteristics with maximum data availability. In particular, we remove any characteristics for which more than 25% of the observations in the panel of firms are missing. In addition, we remove any time periods in the early part of the sample for which less than 500 firms are available. Lastly, we drop any date-stock observations for which there are missing characteristics. We collect the resulting 27 rank-transformed characteristics for each of the stocks in the monthly characteristics matrix \mathbf{X}_t .¹⁰

As we discussed in Section IV.D, normalizations such as rank-transformation remove time-varying components of characteristics. Unlike orthonormalization, rank-transformation does not remove information about time-varying correlations, but time-varying components of cross-sectional means and dispersion of characteristics are removed. On one hand, removing these components may restrict the investment opportunity set and lower the maximum Sharpe ratio that is attainable. On the other hand, the conditions necessary for means, covariances, and risk prices of OLS factors to be constant are more likely to hold. If they are constant, the unconditional maximum Sharpe ratio of the OLS factors is equal to its (constant) conditional

10. Table I provides the list of characteristics we use.

version and we can evaluate OLS factor models based on the unconditional maximum Sharpe ratio that the factors attain.

At the end of each month t , we construct OLS factors weights as $(\mathbf{X}'_t \mathbf{X}_t)^{-1} \mathbf{X}'_t$. To avoid intra-month trading, we adjust daily portfolio weights within the following month $t + 1$ to make the factors buy-and-hold during the month. Rebalancing to OLS factor weights then takes place at the end of each month $t + 1$, and so on.

Our main analysis is conducted out of sample using a split-sample approach. Specifically, we split the sample into two parts: pre-2005 and 2005–present. We estimate any time-series parameters using the earlier sample. For cross-sectional regressions which rely on rolling covariance estimates, we use the most recently available data up to that point in time.

When reporting out-of-sample (OOS) MVE portfolio Sharpe ratios, we use the sample covariance matrix of daily OLS factor returns as an estimate of the unconditional factor covariance matrix and factor means as estimates of unconditional expected excess returns on the factors, with both estimated in the pre-2005 sample of daily returns. Combining information from covariances and means, we compute MVE portfolio weights which we then fix and apply to the 2005–present sample of *monthly* stock returns. We then compute and report annualized unconditional Sharpe ratios of these series, to which we refer as OOS maximum Sharpe ratios.

V.B. Empirical performance of hedged factors

We implement the factor hedging procedure of Section II.C. We compute daily rolling covariances of individual stocks returns with the univariate factors within overlapping backward-looking 5-year windows (with a minimum requirement of 20 days of data). We then regress these daily covariances on the characteristics \mathbf{X}_t . The residuals from these regressions give us daily portfolio weights $\mathbf{W}_{h,t}$ of the hedging portfolios which we then use to calculate daily hedging factor returns. This completes the first step in the approach we outlined in Section II.C. Instead of the second and third step outlined in the main text in Section II.C that was




FIGURE I

Improvement in OOS MVE Sharpe ratios due to iterative hedging. We construct hedged factors and iterate by hedging up to five times. The plot shows improvement in average out-of-sample MVE portfolio’s Sharpe ratios constructed from hedged OLS factors relative to unhedged factors, in %, for all models with a constant and 1–25 additional factors. Improvements are averaged across many factor draws. MVE portfolio weights are estimated in the pre-2005 sample using daily returns. Sharpe ratios are evaluated in the 2005–present sample using monthly stock returns.

convenient for the theoretical analysis, we now use the approach of DMRS that we discussed in footnote 4 of Section II.C and that is equivalent under the conditions of Proposition 3. That is, we purge the ad-hoc factors from unpriced risks by regressing the daily univariate factors on the daily hedge portfolio returns. The parameters of this regression are estimated using the pre-2005 sample and then applied to the rest of the sample to construct residuals. The residuals are the hedged univariate factors. We define characteristics associated with these factors to be “hedged characteristics”. Next, we construct OLS portfolios based on these “hedged characteristics”. To construct iterated hedged factors, we repeat this procedure.

Figure I shows improvement in average out-of-sample MVE portfolio’s Sharpe ratios constructed from hedged OLS factors relative to unhedged factors, in %. We run the hedging procedure for up to five rounds of hedging. We calculate these improvements for OLS factor models with different numbers of characteristic-based factors from one to 25, in addition to the constant characteristic which is implicitly included in all models. Since there are differ-

ent possible subsets of J factors from the full 27 OLS factors, we draw, for each J , random subsets of J factors. We do this many times and Figure I shows the percentage improvement in the maximum Sharpe ratio averaged across these random subsets for each J .

As the figure shows, the benefit of hedging decreases as the number of characteristics increases. This is what we anticipated in our discussion of Proposition 2. Including a large number of characteristics makes it more likely that loadings on major sources of covariances are spanned by the columns of \mathbf{X}_t . This renders violations of the conditions of Proposition 2 quantitatively less important. As a consequence, OLS factors approximately span the SDF and factor hedging provides little additional benefit.

The benefit approaches zero and might even turn negative when J is large. Under population moments, as in our earlier theoretical analysis, hedging would never lead to a deterioration of the Sharpe ratio in sample. However, with estimated moments, estimation error contaminates the hedging procedure and hedging can then lead to a deterioration, especially out of sample.

The figure also shows that there can be a substantial benefit from iterating on the hedging procedure using the iterated hedging approach that we developed in our theoretical analysis. This benefit is particularly big if the number of factors is relatively small. For example, with $J = 1$, hedging does not seem to raise the maximum OOS Sharpe ratio on average (although it does so for some of the portfolios individually). Iterating for the total of five hedging rounds leads to an improvement of 80%, on average, that is, the maximum Sharpe ratio after five rounds of hedging almost doubles relative to the maximum Sharpe ratio obtained from the original unhedged portfolios. The marginal benefit of each additional round of hedging is declining.

Table I shows how hedging changes the out-of-sample Sharpe ratio of various specific single factor models. The first column shows unhedged MVE Sharpe ratios, the next 7 columns hedge factors iteratively up to 7 times. In each row of the table, the characteristics matrix \mathbf{X}_t includes a constant and the characteristic listed in this row. For all of these, since only

TABLE I

Maximum hedged Sharpe ratios of all one-factor models (OLS factors; with a constant).

We report maximum out-of-sample Sharpe ratios of all models which use OLS factors (first column), OLS hedged factors for $n = 1..7$ rounds, as well as approximate GLS factors (last two columns). All models include two characteristics in \mathbf{X}_t : a constant, and one of the characteristics listed in the rows. GLS factors use covariance matrix from models which use either 30 Fama-French industry portfolios (column “GLS ind.”), or the sentire set of characteristics in the table (last column). Portfolio weights are computed using daily buy-and-hold adjusted returns; performance is evaluated using monthly (annualized) stock returns. The row labeled “ER (predicted)” uses fitted values from a panel regression of returns on all characteristics as a single characteristic. The last row averages the numbers across all models.

	OLS	Hedged n times							GLS	
		1	2	3	4	5	6	7	ind.	char.
Size	0.5	0.4	0.6	0.7	0.7	0.8	0.8	0.8	0.8	1.2
Value (A)	0.5	0.3	0.1	0.1	0.2	0.3	0.4	0.4	0.7	1.0
Gross Prof.	0.8	1.0	1.1	1.2	1.3	1.3	1.4	1.4	1.6	1.5
F-score	0.9	0.6	1.0	1.2	1.3	1.4	1.5	1.5	1.1	1.5
Debt Issuance	0.8	0.9	1.0	1.1	1.1	1.2	1.2	1.2	1.1	1.4
Share Repurchases	0.6	0.1	0.5	0.9	1.1	1.2	1.3	1.4	1.0	1.3
Net Issuance (A)	0.9	0.9	0.9	1.1	1.3	1.3	1.4	1.4	1.2	1.5
Asset Growth	0.7	0.6	0.6	0.6	0.7	0.8	0.8	0.9	1.0	1.4
Asset Turnover	0.8	0.8	1.0	1.1	1.2	1.2	1.3	1.3	1.4	1.4
Gross Margins	0.6	0.4	0.7	0.8	1.0	1.1	1.2	1.2	0.9	1.4
Earnings/Price	0.6	0.1	0.4	0.6	0.8	1.0	1.1	1.1	0.9	1.4
Leverage	0.3	0.1	0.0	-0.0	0.1	0.2	0.3	0.4	0.6	0.8
ROA (A)	0.4	0.3	0.7	1.0	1.2	1.3	1.3	1.3	0.7	1.3
Return on BE (A)	0.5	0.3	0.7	0.9	1.1	1.2	1.2	1.2	0.7	1.4
Sales/Price	0.6	0.6	0.4	0.4	0.5	0.6	0.7	0.7	0.9	1.0
Industry Mom.	0.9	1.0	1.1	1.1	1.1	1.1	1.1	1.1	1.2	1.7
Momentum (12m)	0.7	0.8	1.1	1.2	1.3	1.4	1.4	1.4	1.1	1.3
Momentum-Rev.	0.5	0.3	0.4	0.5	0.7	0.7	0.8	0.9	0.7	1.1
Value (M)	0.4	0.2	0.1	0.0	0.1	0.2	0.3	0.4	0.7	1.3
Net Issuance (M)	1.0	0.9	0.9	1.0	1.2	1.3	1.4	1.5	1.3	1.8
Short-Term Rev.	0.6	0.5	0.6	0.6	0.7	0.7	0.8	0.8	0.9	1.7
Idiosyncratic Vol.	0.8	0.4	0.8	1.0	1.2	1.3	1.4	1.5	1.0	1.7
Beta Arbitrage	0.8	1.0	0.9	1.0	1.0	1.0	1.0	1.0	0.9	1.2
Industry Rel. Rev.	0.7	0.5	0.7	0.7	0.8	0.8	0.8	0.8	0.9	1.8
Price	0.5	0.3	0.5	0.7	0.9	1.0	1.0	1.1	0.8	1.3
Firm’s age	0.6	0.4	0.7	1.0	1.1	1.2	1.2	1.3	0.9	1.4
Share Volume	0.8	0.7	0.7	0.7	0.7	0.8	0.8	0.8	0.9	1.4
ER (predicted)	1.7	1.7	1.6	1.6	1.6	1.7	1.7	1.8	1.6	2.5
Average	0.7	0.6	0.7	0.8	0.9	1.0	1.1	1.1	1.0	1.4

one characteristic is used, it is highly unlikely that the conditions hold that are required by Proposition 2 for OLS factors to span the SDF. Hedging the factors should therefore improve the Sharpe ratio. Consistent with this logic, we find improvements from hedging for every characteristic.

To interpret this correctly, it is important to keep in mind that the failure of the unhedged factors to span the SDF is not a simple consequence of the fact that single factor models omits other characteristics that are informative about expected returns but are left out from the single factor model. The hedged factors do not use any information from these other characteristics either! Instead, the reason for the inferiority of the unhedged factors is that a single characteristic is not enough to satisfy the conditions in Proposition 2 for OLS factors to span the SDF that prices assets conditional on this single characteristic.

To see this more clearly, we report Sharpe ratios based on approximate GLS factors in the last two columns of the table. In both cases the covariance matrix of returns is approximated by a separate standalone OLS factor model which uses either 30 Fama and French industries as characteristics (column labeled “ind.”) or the set of all original characteristics listed in the table (column labeled “char.”). That is, the Sharpe ratios in these two columns are still based on a single-factor model, but the factor construction uses the covariance matrix of returns estimated separately from a model with a larger number of factors. The single-factor GLS models, therefore, still omit other characteristics that are informative about expected returns, but, at the same time, they approximately satisfy conditions in proposition 1 and thus achieve mean-variance efficiency (conditional on a narrow information set which includes a single characteristic). The table shows that hedging OLS portfolios moves them in the direction of the GLS portfolio. Factors hedged seven rounds perform generally better than GLS factors which use industries in the construction of the covariance matrix, but worse than GLS factors which use all characteristics in estimating to covariance matrix.

As the table shows, there is considerable heterogeneity in how much hedging improves the Sharpe ratio. Characteristics like returns on assets or book equity show dramatic improve-

ments of about 100–200% with seven rounds of hedging, while others show little out-of-sample improvement. On average, across all portfolios, maximum Sharpe ratios increase from 0.7 (unhedged factors) to 1.1 (factors after seven rounds of hedging), and to 1.4 for characteristic-based GLS factors, as can be seen in the last row of the table.

Lastly, we construct a composite characteristic which uses fitted values from a panel regression of returns on all characteristics (row labeled “ER (predicted)”). This characteristic summarizes expected return predictability of all original characteristics, but uses a single factor and thus generally does not satisfy the conditions in Proposition 2 for OLS factors to span the SDF that prices assets conditional on this single composite characteristic. As such, it is a natural candidate for hedging or GLS factor construction. The table shows that benefits of hedging for this characteristics are limited, with only a slight improvement in OOS Sharpe ration after seven rounds of hedging. The corresponding characteristic-based GLS factor, however, does produce a significantly higher Sharpe of 2.5 (an increase from 1.7), which indicates that the conditions in Proposition 2 are indeed likely to be violated. In other words, sorting stocks on fitted expected returns preserves information in means but largely discards information in covariances, which prevents the factor from reaching mean-variance efficiency.

V.C. Dimensionality reduction

Our final empirical analysis looks at dimension reduction. In Section III, we showed the conditions necessary for dimensionality reduction to be possible. We also showed a few ways how to proceed with dimensionality reduction and how these approaches are related. In this section we explore and compare these methods empirically.

Comparison of PCA on OLS factors with IPCA. In Example 5 we showed that the factors \mathbf{f}_{t+1} in (37) and (38) are a conditional version of the first of two first-order conditions in Kelly, Pruitt, and Su (2019) that define the instrumented principal components analysis

TABLE II
Spanning of IPCA factor loadings with eigenvectors of OLS portfolios.

We report R -squared of regressions of each of the 10 IPCA factor loadings onto 10 eigenvectors associated with dominant PCs extracted from the set of OLS portfolios.

PC1	PC2	PC3	PC4	PC5	PC6	PC7	PC8	PC9	PC10
1.00	0.96	0.94	0.98	0.88	0.90	0.78	0.53	0.26	0.39

(IPCA) estimator. This means that a conditional equivalent of an IPCA estimator can be constructed using PCA on managed portfolios, even in the case where characteristics are not orthonormalized, if the managed portfolios are constructed as OLS factors. In particular, applying PCA to OLS portfolios recovers \mathbf{Q}_t . By applying this matrix to univariate portfolios $\mathbf{X}'_t \mathbf{z}_{t+1}$ and further rotating them by an OLS factor (as in (37)), yields our version of the IPCA estimator.

The assumption of time-constant \mathbf{Q}_t and $\mathbf{\Lambda}_t$ can justify working with a constant \mathbf{Q} extracted from an average conditional, or approximately unconditional, covariance matrix. In practice, however, the theoretic equivalence between our analytic IPCA approach and the iterative procedure of Kelly, Pruitt, and Su (2019) might not hold exactly because this assumption and the assumptions in Corollary 2 about the covariance matrix might not be satisfied empirically. We, therefore, investigate the extent to which this equivalence holds in the data in Table II.

We start by implementing the iterative IPCA approach of Kelly, Pruitt, and Su (2019) and applying it to our data to extract 10 IPCA factors and their respective betas (the matrix Γ_β in their paper). Next, we use PCA on OLS portfolios to construct our matrix \mathbf{Q} . In Table II we regress each of the 10 IPCA eigenvectors associated with highest eigenvalues (first 10 columns of Γ_β) onto the first 10 columns of \mathbf{Q} to see how well factor loadings based on our procedure span IPCA factor loadings. The table reports R -squared of these multivariate regressions.

The table shows that while our factors do not exactly match the IPCA ones in the

TABLE III

Dimensionality reduction: benchmarking different portfolio sorts.

The table reports maximum in-sample (top panel) and out-of-sample (bottom panel) annualized Sharpe ratios of MVE portfolios constructed from N PCs (columns) of one of five portfolio sorts (rows): (i) univariate from Kozak, Nagel, and Santosh (2018) and Kozak, Nagel, and Santosh (2020), (ii) OLS portfolios from Fama and French (2020), (iii) IPCA from Kelly, Pruitt, and Su (2019) implemented as in Example 5, (iv) PPCA from Kim, Korajczyk, and Neuhierl (2019) implemented as in Example 6, and (v) univariate with orthonormalized characteristics. Out-of-sample results are based on a split sample estimation before/after 2005.

	1	2	3	4	5	6	7	8	9	10
In-sample Sharpe										
Univariate	0.4	0.4	1.0	1.1	2.1	2.2	2.2	2.4	2.5	2.5
OLS	0.6	0.8	1.4	1.6	1.6	1.6	1.6	1.7	2.0	2.8
IPCA	0.5	1.2	2.9	3.2	3.2	3.2	3.3	3.2	3.8	3.8
PPCA	0.5	0.5	1.0	1.5	3.4	3.6	3.7	3.7	3.9	4.0
Orthonorm.	0.5	0.5	0.8	1.5	3.1	3.3	3.3	3.3	3.7	3.7
Out-of-sample Sharpe										
Univariate	0.3	0.1	0.4	0.5	1.1	1.2	1.0	1.4	1.4	1.5
OLS	0.4	0.5	0.7	0.8	0.8	0.8	0.8	0.8	0.9	1.6
IPCA	0.4	0.6	1.6	1.8	1.7	1.8	1.7	1.5	2.2	2.3
PPCA	0.4	0.2	0.3	0.7	1.9	1.7	1.7	1.7	2.0	2.0
Orthonorm.	0.4	0.2	0.4	0.9	1.7	1.5	1.6	1.6	1.9	2.0

data, the first seven IPCA factor loadings are spanned by our factor loadings extremely well, with R -squared exceeding 90% for almost all of these factors. We conclude that our approach to identifying IPCA factors might be a viable analytic alternative to the costly iterative procedure in Kelly, Pruitt, and Su (2019) and to working with orthonormalized characteristics.

Comparison of dimensionality reduction methods applied to alternative sets of portfolios. In Section III we discussed that alternative assumptions lead to different techniques of how dimensionality reduction should be implemented. In particular, Example 5 showed that under Assumption (33) we should apply PCA to OLS portfolios to obtain eigenvectors and which we then apply to univariate portfolios, followed by a rotation, to construct

our IPCA factors (related to Kelly, Pruitt, and Su (2019)). Example 6 showed that under Assumption (33) we should apply PCA to univariate portfolios constructed using orthonormalized characteristics to obtain PPCA factors from Kim, Korajczyk, and Neuhierl (2019). As a benchmark, we also consider applying PCA to univariate portfolios as motivated by Kozak, Nagel, and Santosh (2018) and Kozak, Nagel, and Santosh (2020), to OLS portfolios used in Fama and French (2020), and to portfolios constructed using orthonormalized characteristics.

We now compare empirical performance of these five methods of dimensionality reduction in terms of spanning the unconditional MVE frontier. We apply PCA to monthly returns on either of these five sets of portfolios to extract a lower-dimensional latent factor structure. Table III reports maximum in-sample and out-of-sample annualized Sharpe ratios of MVE portfolios constructed from these extracted latent factors, separately for each of the five portfolio sorts. We report our results by varying the number of latent factors from 1 to 10 (shown in columns). To compute out-of-sample metrics we split the sample in 2005, estimate MVE weights in the earlier part of the sample using daily returns, and compute Sharpe ratios in the latter part using these fixed weights and monthly returns.

The table shows that our analytical versions of IPCA factors from Example 5 and PPCA factors from Example 6 perform better than PCA on simple univariate factors or OLS factors. They achieve the highest in-sample and out-of-sample Sharpe ratios across most specifications, which are roughly comparable across these two specifications.

In the last row in either panel we orthonormalize characteristics cross-sectionally. Because in this case $(\mathbf{X}'_t \mathbf{X}_t) = \mathbf{I}$, all four methods discussed above become equivalent. As discussed in Section IV.D normalization of characteristics removes time-series variation in their cross-sectional variances and correlations, but can be advantageous for conditioning down the models. The maximum Sharpe ratio attainable conditional on orthonormalized characteristics might, therefore, be lower than that of the original characteristics. Table III shows that Sharpe ratio deterioration is small in the data – Sharpe ratios attainable from orthonormalized

characteristics are roughly the same as the ones from IPCA and PPCA methods.

VI. CONCLUSION

Heuristic factor construction by sorting on firm characteristics, weighting by characteristics, or computing OLS cross-sectional regression slopes does not use information about the covariance matrix of individual stock returns. As a consequence, these heuristic factors span the SDF only if the covariance matrix satisfies certain special conditions. We work out what these conditions are and obtain a number of insights:

First, horse races between direct linear characteristics-based prediction of excess returns and heuristic characteristics-based factor models, or between different heuristic factor models, have no economic content other than exposing the shortcomings of heuristic factor construction that are rooted in their neglect of information about the covariance matrix. Results from such horse races do not lead to insights about competing economic theories of risk premia and mispricing.

Second, when the individual stock return covariance matrix satisfies conditions such that OLS cross-sectional regression slope factors span the SDF, then nonsingular rotations of OLS factors span the SDF, too, including univariate factors in which stocks are weighted by single characteristics. Choice among these different rotations is a matter of convenience, for example, to obtain suitable conditioning-down properties.

Third, the conditions on the covariance matrix that allow OLS factors, or rotations thereof, to span the SDF are more likely to hold when the number of characteristics employed by the econometrician is larger. Additional characteristics can help even if they are unrelated to expected returns as long as they help to capture important sources of stock return covariances.

Fourth, heuristic factor models that employ only a small number of characteristics can benefit from purging unpriced risks using the hedging method proposed in Daniel, Mota, Rottke, and Santos (2020). The reason why hedging works is that hedged factors can span the

SDF under weaker conditions on the covariance matrix than the unhedged factors. Hedging unpriced risks effectively incorporates some information about the covariance matrix into factor construction, but without requiring inversion of a large covariance matrix.

Fifth, iterating on these hedging procedures allows further relaxation of the conditions on the covariance matrix.

Sixth, when the relationship of expected returns and covariance matrix to characteristics has a lower-dimensional structure such that information in J characteristics can be captured by $K < J$ characteristics, then the SDF can be spanned by K factors without requiring inversion of a large covariance matrix. Under the conditions on the covariance matrix that allow the factors to span the SDF, simple PCA on OLS factors or univariate factors constructed using orthonormalized characteristics delivers the same SDF as the the IPCA method of Kelly, Pruitt, and Su (2019), up to a rotation, and as the PPCA method of Kim, Korajczyk, and Neuhierl (2019), respectively.

Overall, our results provide a foundation for the construction of reduced-form characteristics-based factors that was missing so far in the vast empirical literature on factor models in cross-sectional asset pricing.

REFERENCES

- Back, Kerry, Nishad Kapadia, and Barbara Ost diek, 2015, “Testing factor models on characteristic and covariance pure plays,” Working paper, SSRN.
- Brandt, Michael W, Pedro Santa-Clara, and Rossen Valkanov, 2009, “Parametric Portfolio Policies: Exploiting Characteristics in the Cross-Section of Equity Returns,” *Review of Financial Studies* 22, 3411–3447.
- Chib, Siddhartha, Yi Chun Lin, Kuntara Pukthuanthong, and Xiaming Zeng, 2021, “Slope Factors Outperform: Evidence from a Large Comparative Study,” Working paper, Washington University.
- Daniel, Kent, Lira Mota, Simon Rottke, and Tano Santos, 2020, “The Cross-Section of Risk and Return,” *Review Of Financial Studies* 33, 1927–1979.
- Daniel, Kent, and Sheridan Titman, 1997, “Evidence on the Characteristics of Cross Sectional Variation in Stock Returns,” *Journal of Finance* 52, 1–33.
- Davis, James, Eugene F. Fama, and Kenneth R. French, 2000, “Characteristics, Covariances, and Average Returns: 1929 to 1997,” *Journal of Finance* 55, 389–406.
- Fama, Eugene F. *Foundations of Finance* (Basic Books, New York, NY, 1976).
- Fama, Eugene F., and Kenneth R. French, 1993, “Common Risk Factors in the Returns on Stocks and Bonds,” *Journal of Financial Economics* 33, 23–49.
- Fama, Eugene F, and Kenneth R French, 2015, “A Five-Factor Asset Pricing Model,” *Journal of Financial Economics* 116, 1–22.
- Fama, Eugene F, and Kenneth R French, 2020, “Comparing Cross-Section and Time-Series Factor Models,” *Review of Financial Studies* 33, 1891–1926.
- Gerakos, Joseph, and Juhani T Linnainmaa, 2018, “Decomposing Value,” *The Review of Financial Studies* 31, 1825–1854.
- Grinblatt, Mark, and Konark Saxena, 2018, “When Factors Do Not Span Their Basis Portfolios,” *Journal of Financial and Quantitative Analysis* 53, 2335–2354.
- Hansen, Lars P., and Ravi Jagannathan, 1991, “Implications of Security Market Data for Models of Dynamic Economies,” *Journal of Political Economy* 99, 225–262.
- Hou, Kewei, Chen Xue, and Lu Zhang, 2015, “Digesting Anomalies: An Investment Approach,” *Review of Financial Studies* 28, 650–705.
- Kelly, Bryan, Seth Pruitt, and Yinan Su, 2019, “Characteristics are Covariances: A Unified Model of Risk and Return,” *Journal of Financial Economics* 134, 501–524.
- Kim, Soohun, Robert A Korajczyk, and Andreas Neuhierl, 2019, “Arbitrage Portfolios,” *Review of Financial Studies*, forthcoming.
- Kozak, Serhiy, 2019, “Kernel Trick for the Cross Section,” Working paper, University of Maryland.
- Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh, 2018, “Interpreting Factor Models,” *Journal of Finance* 73, 1183–1223.

- Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh, 2020, “Shrinking the Cross-Section,” *Journal of Financial Economics* 135, 271–292.
- Lu, Cuicui, and Peter Schmidt, 2012, “Conditions for the Numerical Equality of the OLS, GLS and Amemiya–Cragg Estimators,” *Economics Letters* 116, 538–540.