Performance of Factor Models in a Simple Economy

Kerry Back* Alexander Ober†

February 21, 2025

We simulate multiple panels of firm characteristics and stock returns from the Berk et al. (1999) equilibrium model. The characteristics identified in the model are the four Fama and French (2015) characteristics plus momentum. We evaluate the performance of the Fama-French-Carhart model, the Fama-MacBeth-Rosenberg model, the instrumented principal components method proposed by Kelly et al. (2019), and the random Fourier features method proposed by Didisheim et al. (2024). We find that the last two models outperform the first two. Performance in the Didisheim et al. (2024) model is increasing in the number of factors up to at least several hundred. The Kelly et al. (2019) model performs the best.

PRELIMINARY AND INCOMPLETE – PLEASE DO NOT CITE

*Jones Graduate School of Business and School of Social Sciences, Rice University
†Jones Graduate School of Business, Rice University

The goal of this project is to perform Monte Carlo analysis of factor construction methodologies. We examine classical methodologies (Fama and French, 2015; Fama and MacBeth, 1973; Rosenberg, 1974) and two recent proposals: instrumented principal components (Kelly et al., 2019, hereafter KPS) and random Fourier features (Didisheim et al., 2024, hereafter DKKM). We compare methods based on Hansen-Jagannathan distances and Sharpe ratios. We assess statistical significance by generating a sample of independent panels. We are able to avoid some of the sampling error inherent in empirical evaluations because we can compute the true theoretical stochastic discount factor and moments of returns. We evaluate the conditional performance of the factor construction methodologies by computing the true discount factor at each date in each panel.

A prime difficulty in assessing factor models via Monte Carlo is that one must choose a data generating process, and the choice of a process may dictate the outcome. To avoid biasing the outcome through selection of the data generating process, we use an off-the-shelf equilibrium model. Berk, Green, and Naik (1999), hereafter BGN, develop a rational pricing model in which firm characteristics such as size and book-to-market have explanatory power for returns. We simulate their model. In the model, it is possible to calculate the four characteristics used in the five-factor Fama and French (2015) model and also momentum. We evaluate factor construction methods based on those five characteristics.

DKKM consider a very large number of factors, prompting the term "complexity" in their title. However, as the authors make clear, and as is well understood, the ultimate goal is to derive a single factor model, the single factor being an estimate of the stochastic discount factor (SDF). The essence of the DKKM methodology is to generate many new characteristics as sines or cosines of random linear combinations of the original characteristics. By de-meaning each generated characteristic in each cross-section, the characteristics can be interpreted as portfolios. DKKM apply penalized regression on the returns of the generated factor portfolios to form a single factor of the form $\hat{\beta}_t' f_{t+1}$, where f_{t+1} denotes the (perhaps very large) vector of generated factor returns from t to $t+1$, and $\hat{\beta}_t$ is a vector of regression coefficients at t. The regression is a ridge regression of the constant 1 on the factor returns (Hansen and Richard, 1987; Britten-Jones, 1999). This process should in principle allow the data more freedom to speak regarding what the SDF is than if we start with a small number of factors as is commonly done. One question we address in this paper is whether this principle has any effect in the relatively simple BGN economy (we find that it does).

KPS construct factors from characteristics and then reduce the number of factors by using a version of principal components analysis (PCA). Their version takes advantage of cross-sectional correlations between unobservable covariances and observable characteristics, and they call it instrumented principal components analysis (IPCA). We use the five characteristics in the BGN economy, apply IPCA to reduce the number of factors, and then run a regression (without penalization) of the constant 1 on the reduced set of factor returns.

Our findings are that the DKKM methodology outperforms the classical models in the BGN economy. Furthermore, performance is increasing in the number of factors up to several hundred factors, at which point it plateaus. Remarkably, despite the fact that the DKKM method is based on a large number of *randomly constructed* factors, the standard deviations across simulated panels of sample Sharpe ratios and Hansen-Jagannathan distances of the DKKM model are less than those of the classical models. In more complex economies, and especially when more characteristics are available to study, it seems likely that, as DKKM argue, it is beneficial to generate thousands or even hundreds of thousands of factors before attempting to consolidate them into an estimate of the SDF.

However, we also find that the KPS method with a small number of factors performs even better than the DKKM method. Our implementation of the KPS method (IPCA + OLS) is related to running principal components regression (PCR = PCA + OLS), and PCR is related to the ridge regression used by DKKM. The chief differences between IPCA + OLS on the one hand and ridge regression on the other are

(i) the difference beween running IPCA and running PCA on characteristic portfolios: IPCA exploits observed variation in firm characteristics in its construction of factors (see KPS for discussion), and

(ii) the difference between PCR and ridge: PCR applies a zero-one "shrinkage" of singular vectors, discarding small singular values and keeping large singular values without shrinkage, whereas ridge applies smooth shrinkage.

Based on the good performance of the KPS method that we document, we surmise that (i) is very important. Consequently, a hybrid of the DKKM and KPS methods that generates a large number of random factor portfolios following the DKKM recipe and then applies IPCA + OLS as a substitute for ridge regression might perform well, though we note that IPCA + OLS is more numerically intensive than ridge regression.[1]

We describe the BGN model in the next section. Section 2 describes how we assess factor models. Section 3 describes the DKKM, KPS, and classical factor models that we study. Section 4 presents our results regarding those models. Section 5 concludes.

1 Berk-Green-Naik Model

In the BGN model, firms invest optimally given an exogenous pricing kernel and random investment opportunities. The SDF at date t for pricing cash flows at $t+1$ is

$$m_{t+1} := \mathrm{e}^{-r_t - \frac{1}{2}\sigma_m^2 + \sigma_m \varepsilon_{t+1}} \,. \tag{1.1}$$

[1]DKKM show that applying PCA + ridge to random characteristic portfolios does not work as well as ridge, but applying IPCA + OLS (or IPCA + ridge) may work better.

The interest rate process is a Vasicek (1979) process:

$$r_{t+1} = r_t + \kappa(\mu - r_t) + \sigma_r \eta_{t+1} \,. \tag{1.2}$$

Here, ε and η are independent sequences of i.i.d. standard normals.

There are a fixed number of firms. Each firm begins at date 0 with zero capital. Each firm receives an investment opportunity each period. The opportunities expire if not taken in the period in which they arrive. All projects require the same amount of capital I and are fully equity financed. A project that is taken generates operating cash flows each period until it randomly dies. Free cash flow is paid out to shareholders.

The operating cash flow of each project has a time-invariant beta with respect to the SDF process shocks ε and a time-invariant idiosyncratic risk. The betas and idiosyncratic risks are drawn randomly for each firm and date from fixed distributions. A project's NPV depends on its beta and on the level of interest rates. Firms accept all positive NPV projects. Because the project arrival processes are the same across firms, all firms have the same value of growth options at any point in time. The value of growth options varies over time, because of variation in the interest rate. The value of assets in place varies across firms at each point in time due to differences in past project quality. The value of assets in place also depends on the interest rate.

The model generates the following data for each firm each period:

- book value of equity = book value of assets
- market value of equity
- net income = operating cash flow
- stock return

From these, we calculate size, book-to-market, ROE, asset growth, and momentum ($t-12$ through $t-2$ returns). BGN show that size, book-to-market, and momentum are correlated with subsequent returns.

We calibrate the model following BGN, and simulate it with a period length of one month, as do BGN. Like BGN, we discard the first 200 months to allow the economy to reach a steady state. We simulate multiple panels. Each panel consists of 1,000 firms and 720 months (after discarding the first 200). The following figures show data for a single panel. The exact data shown in these figures is not important. The figures only illustrate general features of the model.

Figure 1.1 shows a path of the interest rate process. Figures 1.2–1.5 provide information regarding the cross-section of firms at four distinct dates. Figure 1.2 shows the number of active projects across firms. The book equity of a firm equals its number of active projects multiplied by the cost of each project, so Figure 1.2 also provides information about the dispersion of book equity across firms. Figure 1.3 shows the distribution of market equity across firms. Aggregate market equity is relatively high at month 400 and relatively low at month 900 due to differences in the level and history of the interest rate. The level affects both the value of assets in place and the value of growth options, and the history affects the value of assets in place due to the effect of the interest rate on project choice. Figure 1.4 shows the distribution across firms of four firm characteristics: book-to-market, momentum, profitability, and asset growth. Figure 1.5 shows the distribution of returns. As in the actual data, the cross-sectional distribution of returns is leptokurtic and positively skewed.

To compute theoretical conditional moments in the BGN model, we need the list of all current projects for every firm – the number of projects and each project's beta and idiosyncratic risk, and we need to know the current interest rate. Past project decisions depend on past interest rates as well as project betas, so the economy is path dependent. To put it another way, the state space of the economy is very large. Nevertheless, we can compute the theoretical conditional moments. We compute the true conditional SDF each period and the true conditional Sharpe ratios of factor portfolios.

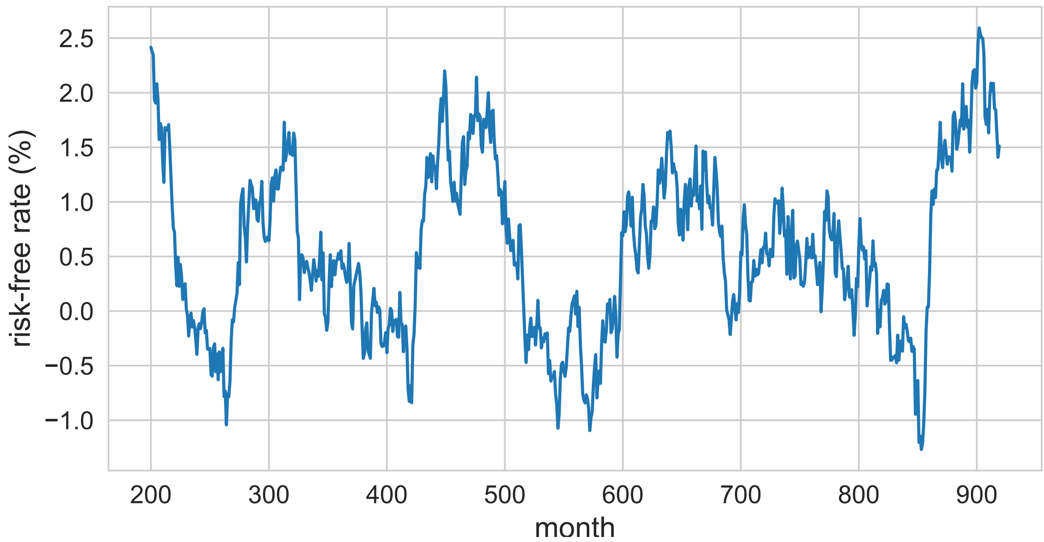

Figure 1.1: A path of the Vasicek interest rate process with the BGN calibration.

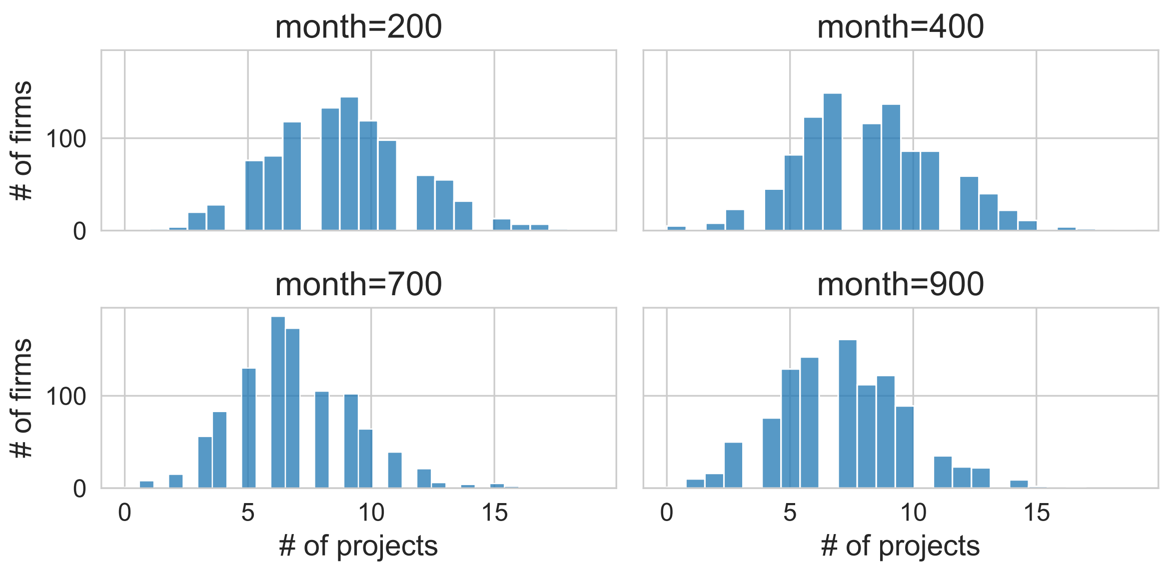

Figure 1.2

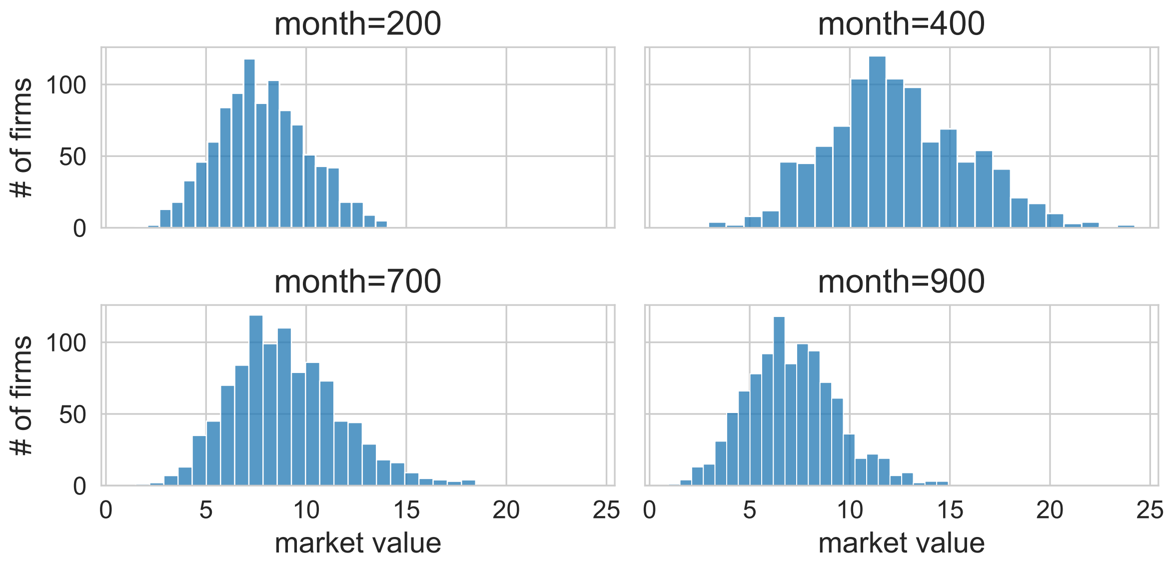

Figure 1.3

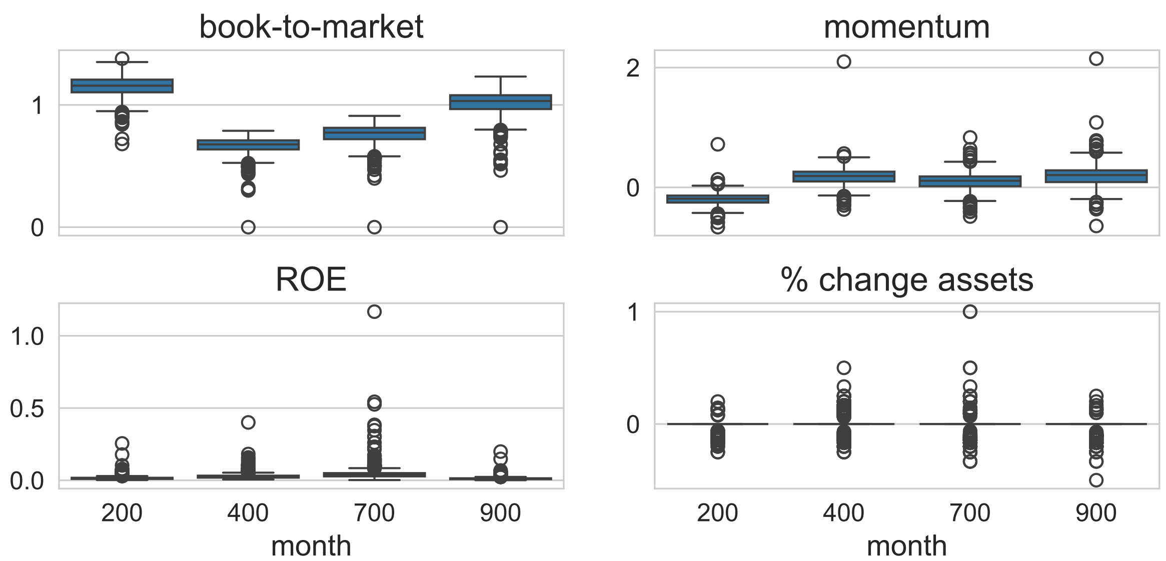

Figure 1.4

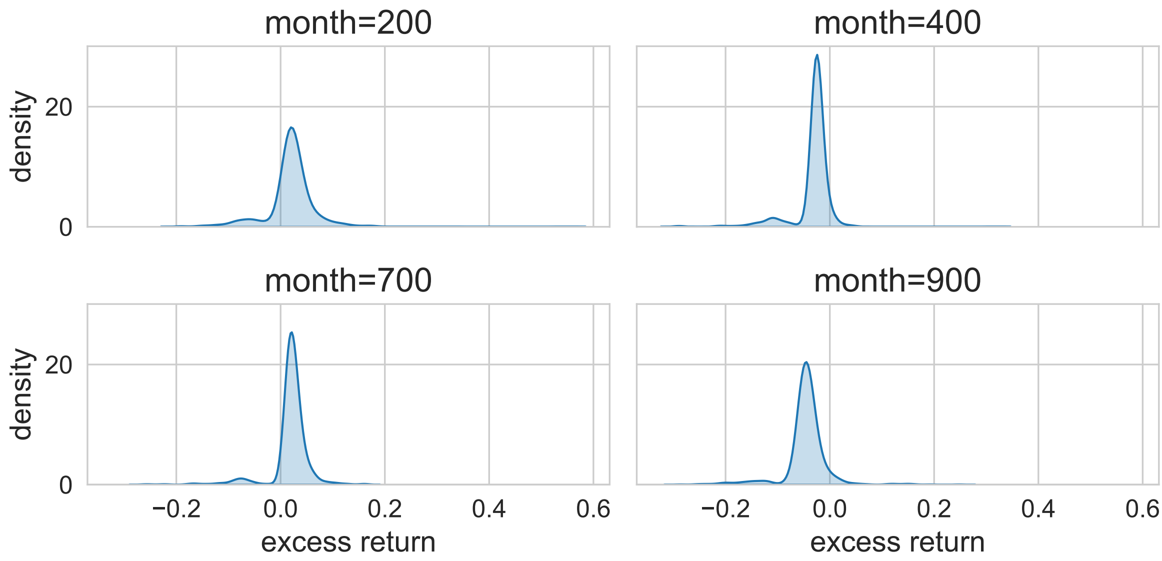

Figure 1.5

2 Assessing Factor Models

We follow DKKM closely in our evaluation of models. Hansen and Richard (1987) show that the efficient part of the mean-variance frontier is the set of returns $r_{f,t} + bz_{t+1}$ for $b \geq 0$, where $r_{f,t}$ denotes the risk-free rate from t to $t+1$, and z_{t+1} is the projection of the constant 1 on the space of excess returns from t to $t+1$. The residual $1-z_{t+1}$ in the projection is orthogonal to excess returns. The unique conditional SDF in the span of the asset returns is

$$\frac{1 - z_{t+1}}{(1 + r_{f,t})\mathsf{E}_t[(1 - z_{t+1})^2]} \,. \tag{2.1}$$

In the factor models that we study, all factors are excess returns. There is a similar representation of the mean-variance frontier spanned by each set of factors and the risk-free asset. Given a set of factors, let y_{t+1} denote the projection of the constant 1 on the set of factor portfolio returns from t to $t+1$, so $\{r_{f,t} + by_{t+1} \mid b \geq 0\}$ is the efficient part of the frontier spanned by the risk-free asset and the factor returns. The unique conditional SDF in the span of the factors and the risk-free asset for pricing the factors and the risk-free asset is

$$\frac{1 - y_{t+1}}{(1 + r_{f,t})\mathsf{E}_t[(1 - y_{t+1})^2]} . \tag{2.2}$$

In each factor model and at each date t, we estimate y_{t+1} by regressing the constant 1 on the factors (without an intercept and possibly with penalization) using the previous 360 months of returns. Britten-Jones (1999) uses this type of regression (without penalization) to compute the mean-variance frontier. DKKM use ridge regression to mitigate overfitting and to allow even more factors than time periods in the regression. Denoting the vector of regression coefficients by $\hat{\beta}_t$ and the factor returns from t to $t + 1$ by f_{t+1}, we compute $\hat{y}_{t+1} = \hat{\beta}_t' f_{t+1}$.

We calculate the square root of the mean of $(\hat{y}_{t+1} - z_{t+1})^2$ in each panel and consider it an estimate of the unconditional Hansen and Jagannathan (1997) distance (we ignore the scaling in the SDFs (2.1) and (2.2)). The Hansen-Jagannathan distance is a measure of how accurately the factor model prices assets and also a measure of how close the factor model comes to spanning the mean-variance frontier. As a second measure, we compute the mean in each panel of the theoretical conditional Sharpe ratio of $\hat{y}_{t+1}$.

3 Models

We replicate the Fama and French (2015) construction of SMB, HML, RMW, and CMA and include UMD as well as the value-weighted market excess return to form the six-factor Fama-French-Carhart (FFC) model. We also run Fama and MacBeth (1973) regressions on book-to-market, momentum, profitability, and asset growth. We standardize the portfolios implicit in the Fama-MacBeth regressions (Rosenberg, 1974; Fama, 1976) to be 100% long and 100% short and use the portfolio returns in conjunction with the equally weighted market excess return to form what we call the Fama-MacBeth-Rosenberg (FMR) model.[2]

[2]We use the equal weighted market excess return because the FMR regressions weight stocks equally. The equally weighted market excess return is the intercept in the FMR regression when the characteristics are de-meaned in each cross section. We could run weighted FMR regressions to achieve

We use the same five characteristics to implement the DKKM method. DKKM use random Fourier features to create potentially a very large number of factors. We follow their recipe to form various sets of what we call DKKM factors, ranging from a six factor model to a model with 36,000 factors.

The DKKM method begins by standardizing each characteristic in each cross-section, replacing the raw characteristic values with percentiles and then subtracting 0.5 to get ranks between -0.5 and $+0.5$. Let C_t denote the $5 \times n$ matrix of rank-standardized characteristics at date t, where n is the number of assets. From C_t, we generate an $n_f \times n$ matrix of random characteristics as follows. Let W denote a $\frac{n_f}{2} \times 5$ matrix whose entries are sampled from the standard normal distribution, and let γ denote a length $\frac{n_f}{2}$ vector whose entries are sampled uniformly from $\{0.5, 0.6, \ldots, 1\}$. We use the same W and γ for all t. Set $A_t = \gamma \odot WC_t$, where $\gamma \odot W$ denotes element by element multiplication of γ with each column of W. We compute an $n_f \times n$ matrix of random characteristics from the $\frac{n_f}{2} \times n$ matrix A_t by taking sines and cosines of the elements of A_t and stacking the sines and cosines as separate rows. We then rank standardize the rows of this matrix, replacing the raw characteristic values with percentiles and then subtracting 0.5 to get ranks between -0.5 and $+0.5$. Each row of this matrix can be interpreted as a long-short portfolio. The returns of the portfolios from t to $t+1$ are the DKKM factor realizations f_{t+1} from t to $t+1$.

We use ridge regression to form the estimate $\hat{\beta}_t' f_{t+1}$. We vary the penalty parameter in the ridge regression to create multiple estimates. To mitigate the effect of randomness in the draws of the random Fourier features, we follow DKKM by generating 20 samples of W and γ, and, for each value of α, we average the 20 estimates $\hat{\beta}_t' f_{t+1}$ to produce our final estimate $\hat{y}_{t+1}$ for that value of α.

value weighting or something between equal and value weighting, but we do not explore that.

The ridge regression is

$$\min_{\beta} \sum_{i=-359}^{0} \left(1 - \beta' f_{t+i}\right)^2 + \alpha \beta' \beta \,, \tag{3.1}$$

where α is the penalty parameter. More penalization is needed when the number of factors n_f is larger. To get a sense for how the penalty should vary with the number of factors, consider doubling the number of factors by simply replicating each factor. Then, to make $\beta'\beta$ small, we will want to split each beta evenly among the duplicate factors in each pair. This reduces the sum of squared betas by 1/2. Therefore, to maintain the same penalization, we should double α. Hence, we set $\alpha = \kappa n_f$ and vary κ. If adding more factors is more effective than simply replicating factors, then, for each value of κ, we should see performance improve as the number of factors increases. We also look at what DKKM call ridgeless regression, which can be interpreted as the limit of the ridge regression as $\alpha \to 0$ (it is OLS when the number of factors is not larger than the number of time periods). We explored ridge regression to form the estimates $\hat{\beta}_t' f_{t+1}$ for the FFC and FMR factors, but it always underperformed OLS, so in the next section we only report the OLS results for FFC and FMR.

We use the same five characteristics to implement the KPS method. We begin with the $5 \times n$ matrix C_t of rank-standardized characteristics defined above. Each row of this matrix can be interpreted as a characteristic factor portfolio. We augment this matrix with a row of constants. Call the resulting $6 \times n$ matrix $\tilde{C}_t$. The returns on the rows of $\tilde{C}_t$ from t to $t+1$ are $f^{\circ}_{t+1} := \tilde{C}_t r_{t+1}$ where r_{t+1} denotes the n-vector of excess stock returns. The IPCA step at date t defines an $n_f \times 6$ matrix A_t, where n_f now denotes the number of KPS factors, and the returns of the KPS factors are defined to be $f_{t+1} = A_t f^{\circ}_{t+1}$. We regress the constant 1 on the factor returns in rolling 360-month windows to obtain regression coefficients $\hat{\beta}_t$, and we set $\hat{y}_{t+1} = \hat{\beta}_t f_{t+1} = \hat{\beta}_t A_t f^{\circ}_{t+1}$. The six-factor KPS model does not employ dimension reduction and is the same as the FMR model run on rank-standardized characteristics.[3]

The DKKM and KPS methodologies are related to the portfolio construction methodology of Brandt, Santa-Clara, and Valkanov (2009), hereafter BSV. BSV de-mean characteristics in each cross-section as DKKM and KPS do, so they can be used as portfolio weights in a long-short portfolio, and they consider portfolios as linear combinations of characteristics. They recommend using the portfolio of this type that maximizes the past sample mean of a utility function. DKKM generate many additional characteristics as de-meaned sines or cosines of random linear combinations of the original characteristics, and KPS perform a dimension reduction of the characteristics (denoted as $A_t C_t$ above). Following DKKM, we implement both the DKKM and KPS methodologies by selecting, among the portfolios that are linear combinations of the augmented or reduced characteristics, the one that maximizes the past sample mean of the quadratic utility function $-(1-r)^2$. Here, r denotes the portfolio excess return, and we note the caveat that DKKM impose L^2 penalization in the maximization to avoid overfitting. We use this portfolio as an estimate of a mean-variance frontier portfolio and use it to estimate the SDF.

4 Results

We generate 300 panels, each consisting of 1,000 firms and 720 months. We run the regressions to estimate $\hat{\beta}_t$ to form $\hat{y}_{t+1} := \hat{\beta}_t' f_{t+1}$ in rolling 360 month windows. We compute the square root of the mean of $(\hat{y}_{t+1} - z_{t+1})^2$ as an estimate of the unconditional Hansen-Jagannathan distance and the mean of the conditional Sharpe ratio $\mathsf{E}_t[\hat{y}_{t+1}]/\text{stdev}_t(\hat{y}_{t+1})$ in each panel. We run the DKKM model with 6, 36, 360, 3600, and 36000 factors, so the ratio of the number of factors to the number of months in the ridge regression window is 0.0167, 0.1, 1, 10, and 100.

Figure 4.1 presents the means of these statistics across panels. Perfor-

[3] In our implementation of the FMR model, we do not rank-standardize the characteristics. Instead, we scale the FMR portfolios to be 100% long and 100% short.

mance improves with the number of factors in the DKKM up to 360 factors and then plateaus. Dimension reduction improves performance in the KPS model until we reach two factors. The FMR and FFC models generally underperform the DKKM and KPS models. Statistical significance of differences in performance is assessed below. In addition to the statistics for the factor models, Panel (c) of Figure 4.1 presents the mean across panels and months of the maximum achievable conditional Sharpe ratio. The difference between the mean conditional Sharpe ratio and the mean of the maximum conditional Sharpe ratio is, like the Hansen-Jagannathan distance, a measure of how far the factor model is from correctly pricing all assets.

Figure 4.2 presents the distributions of the statistics across panels for select models. It is interesting that the dispersion of the statistics is less for the DKKM model than for the FMR model. This is despite the fact that factor portfolios are randomly constructed in the DKKM model.

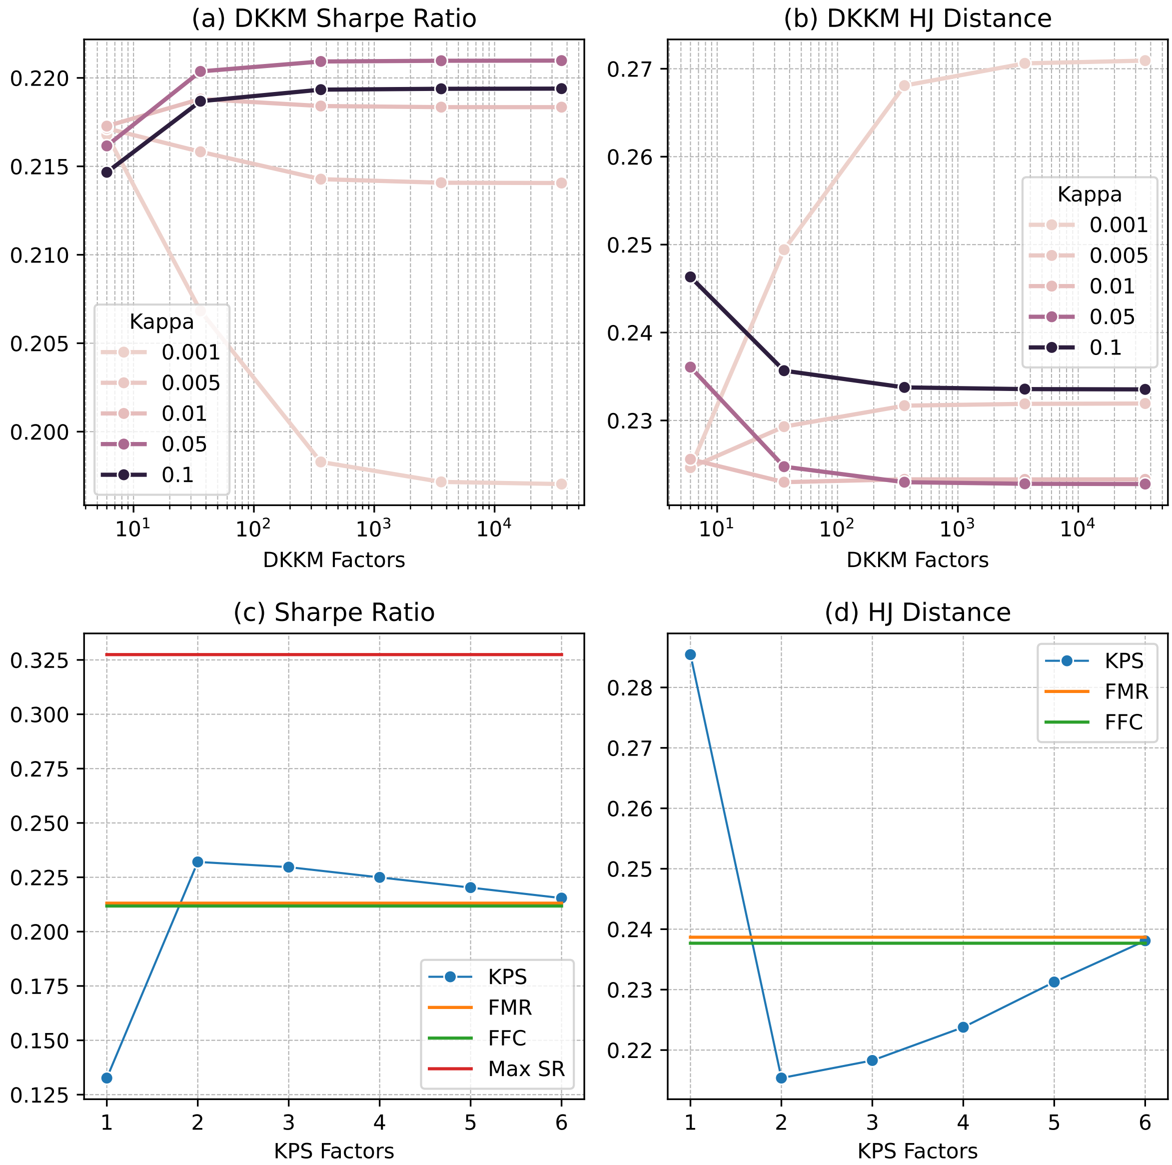

Figure 4.1: 300 panels of data are generated for the BGN economy. Panel (a) shows the mean (across months and panels) of the conditional Sharpe ratio for various numbers of factors and various penalization parameters κ. Panel (b) presents the means across panels of the square root of the within-panel mean of $(\hat{y}_{t+1} - z_{t+1})^2$ for the same models. Ridgeless regression underperforms and is omitted for reasons of scale. Panels (c) and (d) present the same statistics as Panels (a) and (b), respectively, for the FFC and FMR models and for the KPS model with various numbers of factors. Panel (c) also shows the mean across panels and months of the maximum achievable conditional Sharpe ratio.

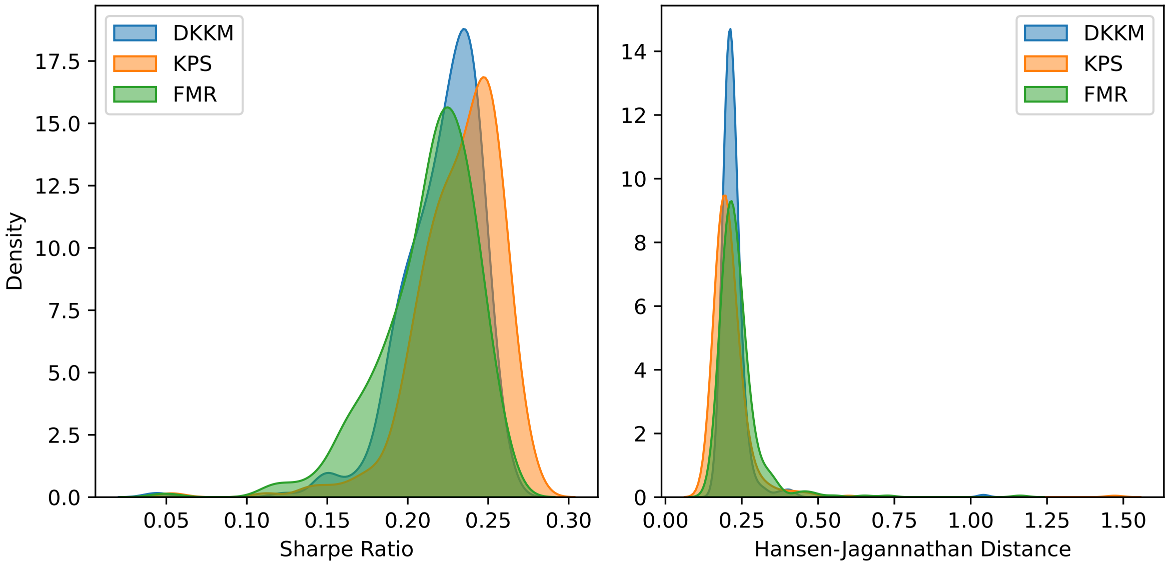

Figure 4.2: 300 panels of data are generated for the BGN economy. Panel (a) shows the distributions of the within-panel means of the conditional Sharpe ratio for the DKKM model with 360 factors and $\kappa = 0.05$, the KPS model with two factors, and the FMR model. Panel (b) presents the disributions across panels of the square root of the within-panel mean of $(\hat{y}_{t+1} - z_{t+1})^2$ for the same models.

Table 1 presents a comparison of the FMR and FFC models. The FMR model outperforms on the Sharpe ratio, and the outperformance is statistically significant. The ranking of the two models is reversed for the Hansen-Jagannathan distance estimates; however, the difference between the two models on that dimension is insignificant. Hence, we compare the DKKM and KPS models to the FMR model in what follows.

	Sharpe Ratio	HJ Distance
FMR	0.213	0.239
FFC	0.212	0.238
FMR - FFC	0.001	0.001
t-stat	1.779	0.499
p-value	0.076	0.618

Table 1: **Performance of FMR and FFC Models.** 300 panels of data are generated for the BGN economy. The mean conditional Sharpe ratio and the square root of the mean value of $(\hat{y}_{t+1} - z_{t+1})^2$ are calculated in each panel for the FMR and FFC models. The table reports t-statistics for the differences between the FMR and FMC panel statistics.

Table 2 reports t statistics for the DKKM model versus the FMR model. With sufficient penalization, and a sufficient number of factors, DKKM outperforms FMR and the outperformance is statistically significant. The t statistics plateau at 360 factors. Table 3 reports t statistics for the KPS model versus the best-performing DKKM model and the FMR model. The two-factor KPS model is clearly the best model.

(a) Sharpe Ratio

κ/Factors	6	36	360	3600	36000
0	3.61	-28.80	-117.87	-112.57	-112.54
0.001	3.78	-6.32	-14.26	-15.26	-15.37
0.005	4.07	2.90	1.24	1.02	1.01
0.01	4.12	6.13	5.69	5.61	5.61
0.05	2.77	7.34	8.02	8.08	8.09
0.1	1.34	5.30	6.04	6.11	6.12

(b) Hansen-Jagannathan Distance

κ/Factors	6	36	360	3600	36000
0	-6.89	13.54	38.74	44.47	45.99
0.001	-6.91	3.72	8.12	8.63	8.76
0.005	-6.41	-4.87	-3.19	-3.05	-3.03
0.01	-5.61	-9.01	-8.41	-8.38	-8.36
0.05	-0.87	-5.66	-6.52	-6.63	-6.64
0.1	2.28	-0.99	-1.62	-1.69	-1.70

Table 2: **Performance of the DKKM Model.** 300 panels of data are generated for the BGN economy. The mean conditional Sharpe ratio and the square root of the mean value of $(\hat{y}_{t+1} - z_{t+1})^2$ are calculated in each panel for the DKKM and FMR models. The table reports t-statistics for the differences between the DKKM and FMR panel statistics.

(a) Sharpe Ratio

Factors	1	2	3	4	5	6
vs DKKM	-47.18	23.11	13.57	5.49	-0.84	-6.34
vs FMR	-39.09	19.26	22.01	19.00	12.63	4.67

(b) Hansen-Jagannathan Distance

Factors	1	2	3	4	5	6
vs DKKM	26.92	-3.42	-2.05	0.34	3.20	5.02
vs FMR	11.46	-12.52	-12.65	-9.60	-4.44	-0.26

Table 3: **Performance of the KPS Model.** 300 panels of data are generated for the BGN economy. The mean conditional Sharpe ratio and the square root of the mean value of $(\hat{y}_{t+1} - z_{t+1})^2$ are calculated in each panel for the DKKM, KPS, and FMR models. The table reports t-statistics for the differences between the panel statistics of the KPS model compared to the DKKM model with 360 factors and $\kappa = 0.05$ (Panel (a)), and compared to the FMR model (Panel (b)).

5 Conclusion

The dynamics of the BGN economy do not have a simple state-variable representation, but the model is still a fairly simple economy with only two macro shocks each period. Despite the simplicity of the environment, the DKKM "complexity" method outperforms classical factor models. However, the instrumented principal components method of KPS performs even better. Subsequent research should explore a hybrid of the methods and should explore performance in more complex simulated economies.

References

Berk, Jonathan B., Richard C. Green, and Vasant Naik, 1999, Optimal investment, growth options, and security returns, Journal of Finance 54, 1553–1607.

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.

Britten-Jones, Mark, 1999, The sampling error in estimates of mean-variance efficient portfolio weights, Journal of Finance 54, 655–671.

Didisheim, Antoine, Shikun Ke, Bryan Kelly, and Semyon Malamud, 2024, Complexity in factor pricing models, Research Paper 23-19, Swiss Finance Institute.

Fama, Eugene F., 1976, Foundations of Finance: Portfolio Decisions and Securities Prices (Basic Books, New York).

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 James D. MacBeth, 1973, Risk, return, and equilibrium: Empirical tests, Journal of Political Economy 81, 607–636.

Hansen, Lars Peter, and Ravi Jagannathan, 1997, Assessing specification errors in stochastic discount factor models, Journal of Finance 52, 557–590.

Hansen, Lars Peter, and Scott F. Richard, 1987, The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models, Econometrica 55, 587–613.

Kelly, Brian T., Seth Pruitt, and Yinan Su, 2019, Characteristics are covariances: A unified model of risk and return, Journal of Financial Economics 134, 501–524.

Rosenberg, Barr, 1974, Extra-market components of covariance in security returns, Journal of Financial and Quantitative Analysis 9, 263–274.

Vasicek, Oldřich, 1979, An equilibrium characterization of the term structure, Journal of Financial Economics 7, 83–105.