Low-Frequency Risk Factors and Their Fundamental Drivers

Sicong Li*

This version: January 24, 2025

Abstract

There is a "zoo" of factors that capture systematic risk premia and a large number of economic variables that explain their time variation, which poses a doubly high-dimensional challenge to understanding how economic fundamentals relate to the time-varying dynamics of risk premia. I propose a method to regularize this problem by identifying low-frequency risk factors, whose risk premia are driven by latent low-frequency state variables. Empirically, one below-business-cycle-frequency factor and one business-cycle-frequency factor, whose variation concentrates on cycles longer than eight years and between 1.5 and eight years, explain the expected returns of individual stocks and characteristic-managed portfolios. The below-business-cycle-frequency factor has a high Sharpe ratio, and stocks whose current size is small compared to their long-term average load on it. Moreover, selected macroeconomic and financial variables have statistically and economically significant out-of-sample predictive power for the returns of the two low-frequency factors.

Keywords: Time-varying risk premia, cross-section of asset returns, time-series predictability, Fourier transform.
JEL classification: C14, G11, G12, G17

*CUHK Business School, The Chinese University of Hong Kong, Hong Kong SAR, China. Email: sicongli@cuhk.edu.hk. I am grateful for comments from Svetlana Bryzgalova, Victor DeMiguel, Gavin Feng, Hongye Guo (discussant), Jinyu He, Christian Julliard, Yingying Li, Paolo Zaffaroni, Yifeng Zhu (discussant) and seminar participants at Antai College at Shanghai Jiao Tong University, FISF at Fudan University, Renmin University of China, SAIF at Shanghai Jiao Tong University, School of Management at Fudan University, Shanghai University of Finance and Economics, Singapore Management University, The Chinese University of Hong Kong, The Chinese University of Hong Kong – Shenzhen, University of Delaware, University of South Carolina, 10th Hong Kong Joint Finance Research Workshop, 2024 Macau International Conference on Business Intelligence and Analytics, 5th FinTech Research Forum at Renmin University of China, and Workshop on AI in Finance and Digital Economy

1 Introduction

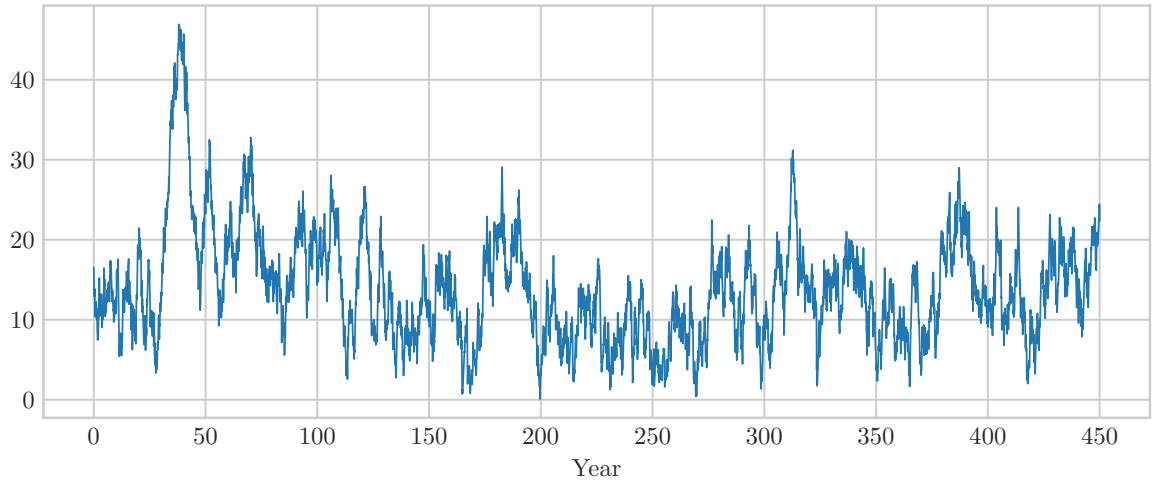
Research in the last few decades has identified more than 300 candidate risk factors that have explanatory power for the cross-section of expected returns, leading to a "zoo" of factors that capture risk premia information (Cochrane, 2011; Harvey, Liu, and Zhu, 2016). Moreover, prominent asset-pricing models imply that the systematic risk premia fluctuate over time. Examples of state variables that drive the time-varying risk premia in these models include the sensitivity function of the surplus consumption ratio (Campbell and Cochrane, 1999), conditional volatility of consumption growth (Bansal and Yaron, 2004), and squared intermediary leverage ratio (He and Krishnamurthy, 2012; He, Kelly, and Manela, 2017). In addition, the time variation of systematic risk premia is also supported by the empirical evidence of return predictability for conventional asset-pricing factors such as the market (e.g., Nelson, 1976; Shiller, 1981; Campbell, 1987; Fama and French, 1988, 1989). Understanding the time-varying dynamics of systematic risk premia and their relation to the economic fundamentals is important because it sets the stage for constructing conditional asset-pricing models. However, it is a doubly high-dimensional problem, with the systematic risk premia captured by the aforementioned high-dimensional "factor zoo" and their time variation explained by a large number of potential predictors of asset returns (see, for example, Dong, Li, Rapach, and Zhou, 2022). This manuscript proposes an innovative approach to regularize this important high-dimensional challenge by identifying *low-frequency* risk factors whose risk premia are driven by *latent low-frequency state variables*.

The motivation for my approach stems from the observation that many variables that drive the variation of risk premia in theoretical asset-pricing models and those conventionally used for empirical return prediction are persistent. For example, Figure 1 illustrates a simulated time series of the sensitivity function of the surplus consumption ratio, which drives the risk premia variation in the habit model of Campbell and Cochrane (1999). I calibrate the parameters of the series following the original paper and simulate the series for 5,400 months (450 years). This figure shows that the state variable that drives the risk premia variation in the habit model is far from independent and identically distributed (iid). Instead, it exhibits very persistent variation and thus manifests slow-moving fluctuations. In other words, a large fraction of the total variation of this variable is *low-frequency* associated with long cycles. In contrast, the innovations of asset-pricing factors (sources of risk) are conditionally uncorrelated and thus are not persistent. In a model where risk factors are the sum of the factor innovations and the state variable component that drives the factor risk premia, the difference between the persistency of the state variables and factor innovations motivates and enables me to identify low-frequency risk factors, of which a large fraction of the total variation is low-frequency, from a panel of candidate risk factors.

I propose a simple method to identify general frequency-specific risk factors from a panel of candidate risk factors, and I focus on using this method to identify those that are low-frequency. Specifically, I first use the Fourier transform to decompose the covariance matrix of candidate risk factors into different frequency-specific components (e.g., a low-frequency component and a high-frequency component). In this case, each component represents the covariance matrix of the candidate risk factors in a specific frequency interval. To identify a low-frequency factor, I use these frequency-specific covariance matrices to solve for a linear combination of the candidate factors that maximizes the ratio of the variation in the low-frequency interval to that in other frequency intervals. As a result, such a linear combination of the candidate factors, by construction, can be interpreted as a low-frequency risk factor because its variation is concentrated in the low-frequency interval.

Figure 1: Simulated series of the sensitivity function of the surplus consumption ratio

This figure illustrates a simulated time series of the sensitivity function of the surplus consumption ratio, which is the state variable that drives the time-varying systematic risk premia in the habit model of Campbell and Cochrane (1999). The series is simulated for 5,400 months (450 years) and its parameters are calibrated following the original paper.



How does identifying low-frequency risk factors regularize the high-dimensional challenge of understanding the time-varying dynamics of systematic risk premia? Note that the innovation of a risk factor is conditionally uncorrelated and thus is not persistent. Thus, intuitively, a low-frequency factor having a higher variation in the low-frequency interval than in other frequency intervals suggests that it loads on a state variable whose variation concentrates on the low-frequency interval. In other words, the low-frequency risk factor reflects the economically important low-frequency dynamics of systematic risk premia. I provide a formal argument to support this intuition. Specifically, I show that the solved linear combination of the candidate risk factors, which identifies the low-frequency factor, is the loadings of all candidate risk factors on a latent low-frequency state variable. Thus, this latent low-frequency state variable drives the risk premium variation of the low-frequency risk factor. Moreover, the low-frequency factor *fully* captures the latent state variable's explanatory power for the systematic risk premia because this latent state variable does not drive the risk premia of other factors that are cross-sectionally orthogonal to the low-frequency factor. Therefore, identifying low-frequency factors regularizes the first high-dimensional challenge associated with understanding the time-varying dynamics of systematic risk premia –"which factors' risk premia are time-varying?"– by reducing the number of factors with time-varying risk premia and thus predictable returns. In particular, a low-frequency factor aggregates a latent low-frequency state variable's explanatory power for the risk premia of all candidate risk factors in the cross-section.

Moreover, identifying low-frequency risk factors also economically regularizes the second high-dimensional challenge associated with understanding the time-varying dynamics of systematic risk premia: "Which economic variables explain the time-varying risk premia?" Intuitively, an economic variable whose variation concentrates on the high-frequency interval should not have high explanatory power for the risk premium of a low-frequency risk factor, which is driven by a low-frequency latent state variable. Thus, compared to explaining the risk premium of a conventional asset-pricing factor such as the market, the number of potential variables that explain the risk premium of a low-frequency factor is smaller. Therefore, my approach sheds light on how popular economic variables in the literature (e.g.,

variables related to the economic fundamentals) relate to the systematic risk premia. Specifically, it incorporates *economic structure in the frequency domain* into the state variables that drive the risk premia of low-frequency factors. Consequently, one should expect that economic variables whose variation concentrates on the low-frequency interval explain the risk premia of low-frequency risk factors.

Furthermore, identifying low-frequency risk factors also alleviates the signal-to-noise problem in return prediction. Specifically, asset/factor returns are notoriously hard to predict because the state variables that drive the risk premia only contribute to a small fraction of the total variation of returns, implying a low population R-squared for return prediction.[1] One conventional approach to address this problem is to study the long-term return predictability (e.g., the predictability of annual returns). In this case, the time aggregation of returns diversifies the noise originating from the uncorrelated innovations while preserving the signal of the persistent state variables. Instead of time-series aggregation, identifying low-frequency risk factors addresses the signal-to-noise problem via *cross-sectional* aggregation. In particular, identifying a low-frequency factor via a linear combination of the candidate risk factors diversifies the innovations of all candidate factors cross-sectionally and preserves the informative component of the state variables from every candidate factor. The benefit of such cross-sectional aggregation is twofold. First, it allows me to study the variation of risk premia using short-term returns rather than long-term returns, which usually involve data overlap and thus result in a cumbersome empirical analysis. Second, it allows me to study the variation of the risk premia of a cross-section of risk factors, while the time-series aggregation can only be used to study one specific factor.

Empirically, I use the Instrumented Principal Component Analysis (IPCA) of Kelly, Pruitt, and Su (2019) to obtain the candidate risk factors from which I identify the low-frequency risk factors. IPCA uses firm-specific characteristics to instrument the loadings of stocks on systematic risks and uses alternating least squares to estimate the systematic risk factors. The main reason for using IPCA to obtain the candidate risk factors is that it explicitly considers the time variation of factor loadings (instrumented by time-varying firm characteristics). Therefore, rather than the time-varying loadings on the true systematic risks, it is expected that the predictable components of the IPCA factors mainly capture the risk premia of systematic risks, whose variation is the main focus of this manuscript. Moreover, **?** show that portfolios sorted based on the past values of firm characteristics capture risk premia unexplained by those based on the present values of characteristics, implying that the persistent and transitory components of characteristics are associated with different components of risk compensation. Thus, in addition to the Basic-IPCA framework, in which only the present values of characteristics instrument the loadings of stocks on systematic risks, I also consider a frequency-specific IPCA (Freq-IPCA) framework, in which the low- and high-frequency components of the *firm characteristics* are used as different instruments for the factor loadings. As a result, Freq-IPCA ensures that the candidate IPCA factors better capture the systematic risk premia and also provides flexibility in that the persistent and transitory components of characteristics give different measures of the stocks' quantity of systematic risks. Additionally, the frequency decomposition of the firm characteristics, together with the identification of frequency-specific factors, allows me to establish an economically important link between an individual stock's frequency-specific risk compensation and the frequency-specific variation in its characteristics.

[1] For example, Ross (2005) provides an upper bound for the population R-squared of a predictive regression for asset returns. In addition, Zhou (2010) finds that the R-squared is less than 1% when predicting the excess return of the S&P 500 index using ten popular variables from the literature.

In the majority of the empirical analysis, I estimate two low-frequency factors: one below-business-cycle-frequency (B-BCF) factor whose variation concentrates on cycles longer than eight years and one business-cycle-frequency (BCF) factor whose variation concentrates on cycles between 1.5 and eight years. My empirical findings are summarized as follows. First, the returns of both factors exhibit slow-moving fluctuations, suggesting that they successfully reflect the low-frequency dynamics of systematic risk premia. Nonetheless, the average returns of the two factors are substantially different. Specifically, the in-sample unconditional Sharpe ratio of the below-business-cycle-frequency factor estimated from 37 Basic-IPCA factors is 1.193, only 23.8% lower than that of the optimal portfolio of 37 Basic-IPCA factors, 1.569. Thus, the Sharpe ratio of the below-business-cycle-frequency factor is surprisingly high, especially considering that the estimation of this factor does not use any information about the average returns of candidate risk factors. This suggests that a large fraction of the *unconditional* systematic risk premia is associated with time-series variation driven by a latent below-business-cycle-frequency state variable, and implies that the performance of most candidate risk factors with positive risk premia can be improved via timing their expected returns (by a below-business-cycle-frequency state variable). Moreover, the conditional risk premia of most candidate risk factors move in the same direction as the below-business-cycle-frequency state variable, suggesting that this latent variable can be interpreted as a fundamental "level" state variable for the risk premia variation. The business-cycle-frequency factor, on the other hand, delivers an unconditional Sharpe ratio close to zero. Correspondingly, the latent business-cycle-frequency state variable can be interpreted as a "long-short" state variable that affect the risk premia of different candidate factors from different directions. Moreover, the two low-frequency factors have low and negative correlations with a factor that mimics the market factor, implying that the market factor cannot explain the risk premia captured by these factors. Moreover, the market factor having a negative correlation to the highly-compensated below-business-cycle-frequency factor suggests that the state variable that drives up the risk premia of a large number of candidate factors simultaneously drives down the market premium. This poses a challenge to many structural asset-pricing models.

Second, parsimonious factor models that include the two low-frequency factors outperform higher-dimensional IPCA models in explaining the average returns of individual stocks and characteristic-managed portfolios. This implies that the proposed method to identify low-frequency factors can be viewed as a empirically satisfactory method to address the "factor zoo" problem by identifying a parsimonious model that prices the cross-section. In particular, although the two low-frequency factors do not explain the time-series variation of the returns of individual stocks and managed portfolios as much as the higher-dimensional IPCA models, they deliver a higher predictive R-squared for explaining the returns of these two panels of assets than the IPCA models. For instance, the predictive R-squared of the nine-factor Basic-IPCA model for explaining the monthly returns of 36 characteristic-managed portfolios is 3.23%, while that of a model that only includes the two low-frequency factors estimated from the nine Basic-IPCA factors is 3.79%. Moreover, using the Freq-IPCA framework consistently improves the performance of the IPCA and low-frequency factors. For instance, the predictive R-squared of the nine-factor Freq-IPCA model for explaining the monthly returns of 36 characteristic-managed portfolios is 3.58%, while that of a model that only includes the two low-frequency factors estimated from the nine Freq-IPCA factors is 4.54%. Thus, the Freq-IPCA framework helps the IPCA and low-frequency factors better capture the risk premia of the cross-section. Moreover, I also find that the low- and high-frequency components of firm characteristics measure the stocks' quantity of risks differently.

Specifically, compared to stocks that are persistently small, stocks whose current size is small compared to their long-term average have high loadings on the highly-compensated below-business-cycle-frequency factor. The driving force of the superior performance of the low-frequency factors is that a large fraction of the unconditional risk premia of the cross-section is associated with time-series variation empirically, and thus the low-frequency factors, which focus on capturing factors with time-varying risk premia, capture well the unconditional risk premia of the cross-section. In general, the method to identify low-frequency factors provides no guarantee to perfectly tame the "factor zoo," especially when most candidate risk factors have constant risk premia. However, it can be used jointly with other popular methods that focus on unconditional factor risk premia to construct an asset-pricing model that prices expected returns conditionally.

Third, using the three-pass regression filter of Kelly and Pruitt (2015), which is a generalization of the partial least squares (PLS; Wold, 1966) regression, I find that 19 macroeconomic variables and 24 financial variables robustly predict the one-month-ahead returns of the two low-frequency factors in- and out-of-sample, implying that these selected economic variables explain the variation of the risk premia of the two factors. Based on a set of observable predictors, PLS is a simple regression-based procedure to identify a parsimonious number of latent factors of the predictors relevant to predicting a specific time series. Therefore, it is an appropriate method to estimate the latent state variables that drive the risk premia of the two low-frequency factors. Using the one-factor PLS prediction model, the 43 economic variables yield a statistically significant out-of-sample R-squared of 6.17% for predicting the monthly returns of the below-business-cycle-frequency factor estimated from 37 Basic-IPCA factors. The significant out-of-sample R-squared further increases to 13.04% when using the Bayesian Information Criterion (BIC) to select the number of PLS factors. These figures are substantially higher than the out-of-sample prediction R-squared usually reported in the literature for a conventional asset-pricing factor, such as the market.

The business-cycle-frequency factor is less predictable, especially for that estimated from a small number of IPCA factors. However, the selected economic variables still significantly predict the monthly returns of the business-cycle-frequency factor estimated from a large number of IPCA factors, with the out-of-sample predictive R-squared being 4.24% when using the one-factor PLS prediction model and 4.32% when using the BIC criterion to select the number of PLS factors. In addition, the statistically significant predictability of the low-frequency risk factors also generates significant economic value for an investor. Specifically, the unconditional mean-variance-efficient portfolio of the two low-frequency factors based on the PLS prediction model has a statistically significant and economically sizable alpha with respect to that based on the historical-mean prediction model. In summary, my empirical results show that low-frequency state variables can time a large fraction of the unconditional systematic risk premia. Moreover, the selected economic variables have high predictive power for the low-frequency factor returns and thus explain the important low-frequency fluctuations of systematic risk premia.

As my work studies the time variation of systematic risk premia, it is thus related to the literature on the time-series predictability of asset returns.[2] The early literature focuses on the predictability of the market index (e.g., Shiller, 1981; Fama and French, 1988) and then extends to the predictability of anomaly portfolios such as value (Asness, Friedman, Krail, and Liew, 2000; Cohen, Polk, and Vuolteenaho, 2003) and momentum (Cooper, Gutierrez Jr, and Hameed, 2004). The objective of this

[2]Extensive surveys of this literature can be found in, for example, Koijen and Van Nieuwerburgh (2011) and, more recently, in Rapach and Zhou (2022).

paper is more general, focusing on the time variation of the risk premia of the cross-section of systematic risks. Thus, I highlight the predictable components of the entire cross-section of returns instead of the returns of a specific anomaly portfolio. With a similar objective, many papers study the predictability of a large cross-section of anomaly portfolios. For example, Stambaugh, Yu, and Yuan (2012) show that the anomaly portfolio returns are predictable by investor sentiment. Ehsani and Linnainmaa (2022) show that the returns of most factors are positively autocorrelated, and thus are predictable by their own past returns. More recently, Neuhierl, Randl, Reschenhofer, and Zechner (2023) study the predictability of over 300 candidate risk factors using a high-dimensional set of predictors, and find that the performance of most types of factors can be improved via factor timing.[3] Instead of studying the return predictability of each anomaly portfolio in the factor zoo, the approach proposed in this manuscript reduces the dimension of predictable factors, thus allowing me to *systematically* study the variation of the risk premia of a large cross-section of candidate risk factors through the lens of low-frequency risk factors, which aggregate low-frequency state variables' predictability for the entire cross-section of returns.

One closely related paper is Haddad, Kozak, and Santosh (2020), who regularize the return prediction problem by identifying the predictable components of the cross-section of returns. In particular, by imposing the constraint that the implied stochastic discount factor (SDF) is not too volatile, they find that the lower-order principal components (PCs) of returns should capture most of the predictability of the cross-section. Empirically, they find that the lower-order PCs of returns can be significantly predicted, with an out-of-sample R-squared of around 4%.[4]

However, there are three main differences between this manuscript and the work by Haddad et al. (2020). First, instead of imposing the restriction that the SDF is not too volatile and thus focusing on the predictability of the lower-order PCs, I let the data speak and identify factors that fully capture low-frequency state variables' predictive power for returns. In this case, I implicitly allow candidate risk factors that do not explain a substantial amount of the time-series variation of returns to have time-varying risk premia. Second, the predictability of the lower-order PCs of returns could potentially come from the predictability of the time-varying loadings of the PCs on the true systematic risks. By using the IPCA factors as the candidate risk factors, whose predictable components are expected to mainly capture risk premia, the approach in this manuscript is better suited for studying the time variation of systematic risk premia. Third, Haddad et al. (2020) use the book-to-market ratio of each PC as its own return predictor. Instead, the approach in this manuscript incorporates economic structure in the frequency domain into the state variables that drive risk premia, and thus sheds light on how popular economic variables in the literature explain the systematic risk premia.

This work also relates to the literature that uses frequency-domain techniques in asset pricing. Bandi, Perron, Tamoni, and Tebaldi (2019) apply the extended Wold representation of covariance-stationary processes and find that the component of market returns with cycles between eight and 16 years is predictable by the component of uncertainty measures with the same frequency. This im-

[3]This is consistent with my findings that the below-business-cycle-frequency factor has a high Sharpe ratio, and thus a large fraction of the unconditional systematic risk premia is associated with time-series variation driven by a latent below-business-cycle-frequency state variable.

[4]This result, combined with the insight of Kozak, Nagel, and Santosh (2018), who show that the higher-order PCs should not earn high risk premia, suggests that a large fraction of the unconditional risk premia of the cross-section is predictable. This is consistent with my finding that the highly-compensated below-business-cycle-frequency factor is also highly predictable.

plies that the long-run future market returns are predictable by the long-run past uncertainties. This manuscript shifts the focus from the predictability of a particular factor, such as the market, to the time-varying risk premia of a cross-section of candidate risk factors. Moreover, the approach in this manuscript enables me to study the variation of risk premia using the short-term one-month-ahead returns. Another application of the frequency-domain techniques in asset pricing examines how risks of different frequencies are priced. For example, Bandi, Chaudhuri, Lo, and Tamoni (2021) find that the CAPM beta within the frequency interval between 32 and 64 months can explain the expected returns of various cross-sections. Dew-Becker and Giglio (2016) study frequency-dependent risks in consumption-based models and find that risks of frequencies lower than the business-cycle frequencies are significantly priced. This manuscript studies a different asset-pricing object. In particular, instead of studying which frequency components of the market or consumption risks are priced, this manuscript takes a relatively simple stance on the priced risks. Specifically, I use a panel of candidate risk factors to span the priced risks and focus on the frequency-specific variation of the risk *compensation*.

Huang (2021) identifies frequency-specific risk factors by applying principal component analysis (PCA) to the frequency-specific components of the covariance matrix of asset returns. In essence, the objective of Huang (2021) is to identify a parsimonious (unconditional) factor model, focusing on identifying weak factors that have high loadings on frequency-specific state variables. For example, factors that do not explain a sufficient amount of the time-series variation of asset returns but have high loadings on a low-frequency state variable cannot be identified by the conventional PCA, but can be detected by applying PCA to the low-frequency covariance matrix of asset returns. Although I also use the information contained in the frequency-specific covariance matrices of the returns, there are important differences between the two manuscripts. Methodologically, instead of applying PCA to the low-frequency covariance matrix, I identify low-frequency factors by solving for a linear combination of the candidate factors that maximizes the ratio of the variation in the low-frequency interval to that in other frequency intervals. More importantly, the objective of this manuscript is different from that of Huang (2021). Specifically, instead of summarizing asset returns in a parsimonious unconditional factor model, my objective is to understand the time-varying dynamics of systematic risk premia in different frequencies and establish the relation between risk premia and state variables characterizing the economic fundamentals, which sheds light on the construction of conditional asset-pricing models.

More broadly speaking, this work also relates to studies that predict stock returns when the predictors are high-dimensional (e.g., Chinco, Clark-Joseph, and Ye, 2019; Gu, Kelly, and Xiu, 2020; Dong et al., 2022). I apply the PLS regression of Kelly and Pruitt (2013, 2015) to predict the returns of the low-frequency factors using a relatively high-dimensional set of variables. In general, the performance for predicting the returns of these two factors could potentially be improved further by applying a more flexible prediction method or using a more comprehensive set of predictors. However, this manuscript does not focus on selecting the "best" return prediction method or the most comprehensive set of return predictors. Instead, the returns of the two low-frequency factors being highly predicted by a relatively simple model further justifies that identifying low-frequency factors successfully regularizes the high-dimensional challenge of understanding how the time-varying dynamics of systematic risk premia relate to the economic fundamentals.

The remainder of the manuscript is organized as follows. Section 2 illustrates the motivation and method to estimate low-frequency factors from a panel of candidate risk factors and describes the set-up

of the empirical analysis. Section 3 presents the empirical results, and Section 4 concludes. Appendix A includes the proofs of all the main results.

2 Method

I first describe the notation used in the remainder of the manuscript. $\mathbb{E}[\cdot]$, $\mathbb{V}ar(\cdot)$, and $\mathbb{C}ov(\cdot)$ denote the unconditional expectation, variance, and covariance, respectively, and those with the subscript t denote their conditional counterparts. For illustrative purposes, Σ_m also denotes the unconditional covariance matrix of a random vector m_t, and $\Sigma_m(h) = \mathbb{E}[m_{t+h}m_t^\top] - \mathbb{E}[m_{t+h}]\mathbb{E}[m_t^\top]$ is the unconditional autocovariance matrix of m_t with lag h. σ_e^2 is the unconditional variance of a random variable e_t. $\mathbf{I}_N$ is the N-dimensional identity matrix, and $\mathbf{0}$ is a vector or matrix of zeros. By default, all vectors are column vectors.

2.1 Data-generating process

I assume that the excess stock returns follow a K-factor model:

$$\underbrace{r_{t+1}}_{N\times 1} = \underbrace{\beta_t}_{N\times K}\underbrace{f_{t+1}}_{K\times 1} + \underbrace{\varepsilon_{t+1}}_{N\times 1}, \tag{1}$$

where r_{t+1} contains the excess returns of the N stocks at time $t+1$, f_{t+1} contains the returns of K risk factors at time $t+1$, β_t contains the conditional loadings of the stocks on the risk factors at time t, and ε_{t+1} contains the idiosyncratic noises of the stock returns orthogonal to the risk factors. To model time-varying risk premia, I assume that the risk factors load on $J < K$ *unobservable* state variables:

$$\underbrace{f_{t+1}}_{K\times 1} = \underbrace{\mu_f}_{K\times 1} + \underbrace{\phi}_{K\times J}\underbrace{x_t}_{J\times 1} + \underbrace{\xi_{t+1}}_{K\times 1}, \tag{2}$$

where x_t contains the J state variables that affect the conditional mean of the risk factors, ϕ contains the loadings of the factors on the state variables, and ξ_{t+1} contains the factor innovations, which span the sources of risk and are orthogonal to x_t and ε_{t+1}. The factor innovations are conditionally uncorrelated, that is, $\mathbb{E}[\xi_{t+1}] = \mathbb{E}_t[\xi_{t+1}] = \mathbf{0}$. Without loss of generality, I assume $\mathbb{E}[x_t] = \mathbf{0}$, and thus μ_f is the unconditional mean of the factors.

Assuming that the risk factors load on the state variables as in (2) implies that the risk premia are time-varying, which is consistent with a vast literature on theoretical and empirical asset pricing. In particular, both the habit model of Campbell and Cochrane (1999) and the long-run risk model of Bansal and Yaron (2004) predict that the systematic risk premia are time-varying, driven by some (persistent) state variables. In addition, the empirical literature shows that aggregate stock returns are predictable by numerous economic variables, such as the dividend-price and earnings-price ratios.[5] Moreover, Chernov, Lochstoer, and Lundeby (2022) show that notable (unconditional) factor models fail to price the multi-horizon returns of the factors included in the models, precisely because these models do not consider the conditional dynamics of asset returns. Assuming that the factors load on the state

[5] An incomplete list includes Campbell (1987); Campbell and Shiller (1988a,b); Breen, Glosten, and Jagannathan (1989); Fama and French (1988, 1989); Kothari and Shanken (1997); Lettau and Ludvigson (2001b); Campbell and Vuolteenaho (2004); Ang and Bekaert (2007); Campbell and Diebold (2009); Cochrane (2008); Kelly and Pruitt (2013).

variables accommodates these findings because it equips the asset-pricing model with the flexibility of time-varying risk premia. Specifically, the conditional expected returns of the N stocks at time t are

$$\mathbb{E}_t\left[r_{t+1}\right] = \beta_t\left(\mu_f + \phi x_t\right), \tag{3}$$

suggesting that the factor loadings β_t and state variables x_t drive the time variation of the quantity of risk and risk compensation of the stocks, respectively.

Moreover, in the presence of time-varying risk premia, the candidate risk factors identified based on popular empirical methods, such as portfolio sorts or PCA, capture not only the sources of risk but also the sources of risk premia variation, consistent with the data-generating process (DGP) of the factors in equation (2). To see this, suppose that the stock returns are generated by the conditional CAPM of Lettau and Ludvigson (2001a). For simplicity, I assume that the stocks have constant loadings on the market risk. In this case, the data-generating process of the stock returns is

$$r_{t+1} = \beta\mu_{m,t} + \beta\xi_{m,t+1} + \varepsilon_{t+1},$$

where the risk premium of the market risk $\xi_{m,t+1}$ is $\mu_{m,t}$, which is time-varying and is driven by the state variable *cay*. Importantly, if one applies PCA to the stock returns in this world, the risk factor she identifies is proportional to $\mu_{m,t} + \xi_{m,t+1}$, the sum of the source of market risk and the component that reflects the time-varying market risk premium. As a result, the DGP of the factors in (2) is consistent with the candidate risk factors identified by popular empirical methods in the literature. In particular, these factors capture the information of the dynamics of systematic risk premia.

Understanding the time-varying dynamics of systematic risk premia empirically, however, is a high-dimensional challenge. On the one hand, systematic risks that carry risk premia are unobservable. To capture the systematic risks, researchers have proposed a "zoo" of factors that capture risk premia information and also various statistical methods, such as PCA and its modifications, to efficiently aggregate the information of the "factor zoo." On the other hand, state variables that determine the conditional mean of the risk factors are also unobservable. To capture these unobservable drivers of risk premia, various theoretical asset-pricing models and research in return prediction have proposed a large number of candidate variables to explain the risk premia variation.[6] As a result, understanding the time variation of systematic risk premia and how it relates to the economic fundamentals is an important yet difficult doubly high-dimensional challenge. The objective of this manuscript is to regularize this high-dimensional challenge.

2.2 Risk factors with time-varying risk premia

In this section, I show that identifying the loadings of the risk factors on the state variables, that is, ϕ in equation (2), regularizes the aforementioned high-dimensional challenge. In particular, identifying ϕ reduces the dimension of factors with time-varying risk premia. Section 2.3 illustrates my approach to identifying (a subspace of) ϕ.

For illustrative purposes, I assume $J = 1$ throughout this section, that is, there exists only one latent state variable driving the time-series variation of the risk premia of all candidate risk factors.[7] In

[6] Importantly, a large number of observable economic variables could have non-zero covariances with the J unobservable state variables, and thus the data-generating process of the risk factors in (2) encompasses the findings in the empirical literature that a high-dimensional set of economic variables have predictive power to factor/anomaly returns.

[7] The results and intuition of this section are unaltered when there is more than one state variable driving the risk premia variation of the risk factors.

this case, the loadings of the risk factors on the state variable are contained a K-dimensional vector ϕ. Without loss of generality, I normalize the covariance matrix of the risk factors to be an identity matrix and normalize ϕ to have unit length. In other words, I have $\Sigma_f = \mathbf{I}_K$ and $\phi^\top \phi = 1$.[8] Let Λ be a $K \times (K-1)$-dimensional matrix whose columns are orthonormal and orthogonal to ϕ, and thus $\Lambda^\top \Lambda = \mathbf{I}_{K-1}$ and $\Lambda^\top \phi = \mathbf{0}$. By rotating the risk factors by $[\phi, \Lambda]^\top$, I obtain

$$\phi^\top f_{t+1} = \phi^\top \mu_f + x_t + \phi^\top \xi_{t+1}, \tag{4}$$

$$\Lambda^\top f_{t+1} = \Lambda^\top \mu_f + \mathbf{0} + \Lambda^\top \xi_{t+1}. \tag{5}$$

Equations (4) and (5) show that, after the rotation, within the space spanned by the K risk factors, $\phi^\top f_{t+1}$ is the *only risk factor with time-varying risk premium* driven by the state variable, while $\Lambda^\top f_{t+1}$ contains the $K-1$ *risk factors with constant risk premia*.[9] Thus, identifying ϕ reduces the number of factors with time-varying risk premia from K to one. In particular, the only factor with time-varying risk premium aggregates the state variable's explanatory power for the risk premia of *all candidate factors*, and this dimension-reduction effect facilitates linking the risk premia variation of the entire cross-section to economic fundamentals.

Moreover, identifying the only factor with time-varying risk premia $\phi^\top f_{t+1}$ in (4) also alleviates the signal-to-noise problem in linking risk premia variation to economic fundamentals. Conventionally, asset/factor returns are notoriously hard to predict because the state variables that drive the risk premia variation only contribute to a small fraction of the total variation of returns. For instance, Zhou (2010) finds that the R-squared is less than 1% when predicting the excess return of the S&P 500 index using ten popular variables from the literature. By identifying ϕ, the signal of the state variable is cross-sectionally aggregated in the only factor with time-varying risk premia, and the noise originating from the innovations of the $K-1$ factors with constant risk premia is diversified away. To see this, note that

$$\frac{\text{trace}\left(\phi\phi^\top \sigma_x^2\right)}{\text{trace}\left(\Sigma_f\right)} = \frac{\sigma_x^2}{K} < \frac{\sigma_x^2}{\mathbb{V}ar\left(\phi^\top f_{t+1}\right)} = \sigma_x^2,$$

where the left- and right-hand sides of this inequality measure the fraction of the variation driven by the state variable in the total variation of all candidate risk factors f_{t+1} and in that of the only factor with time-varying risk premium $\phi^\top f_{t+1}$, respectively. Thus, this inequality shows that the signal strength of the state variable in the only factor with time-varying risk premia is higher than that in the original panel of risk factors, suggesting that identifying the only factor with time-varying risk premia alleviates the signal-to-noise problem. In addition, such cross-sectional aggregation of all candidate factors enables me to study how popular economic variables explain the risk premia of the *cross-section of risk factors*. In contrast, the conventional time-series aggregation, which studies long-term return predictability, only allows me to study how the economic variables explain the risk premium of one specific candidate factor, such as the market.

Furthermore, identifying ϕ also sheds light on how much of the unconditional systematic risk premia is associated with time-series variation. To see this, consider a special case where μ_f is proportional to ϕ, that is, the candidate factors' unconditional risk premia are proportional to their loadings on the only state variable. In this case, $\Lambda^\top \mu_f = \mathbf{0}$ and thus the $K-1$ factors with constant risk premia

[8] These assumptions are without loss of generality because standardizing the factors to have an identity covariance matrix does not change the space spanned by the factors and the state variables are latent.

[9] It is easy to show that $\mathbb{C}ov(\Lambda^\top \xi_{t+1}, \phi^\top \xi_{t+1}) = \mathbf{0}$, and thus the two sets of factors capture different sources of risk.

in equation (5) earn zero risk premia, while the only factor with time-varying risk premium captures all risk premia of the cross-section and thus is highly compensated. Thus, if μ_f is proportional to ϕ, then *all* unconditional risk premia of the cross-section are associated with time-series variation. As a result, the "scope" of factor timing is universal: for any risk factor that earns positive unconditional risk premium, its performance can be improved via factor timing.[10]

In general, the relative magnitude of the unconditional Sharpe ratio of the factors with time-varying risk premia to that of all candidate factors quantifies how much of the unconditional risk premia of the cross-section is associated with time-series variation. Specifically, the unconditional Sharpe ratio of all candidate risk factors and that of the factors with time-varying risk premia are

$$\mathrm{SR}_f = \left(\mu_f^\top \Sigma_\xi^{-1} \mu_f\right)^{1/2}, \quad \text{and}$$
$$\mathrm{SR}_{\phi^\top f} = \left[\left(\phi^\top \mu_f\right)^\top \left(\phi^\top \Sigma_\xi \phi\right)^{-1} \left(\phi^\top \mu_f\right)\right]^{1/2},$$

respectively, with $\mathrm{SR}_{\phi^\top f} \leq \mathrm{SR}_f$. Suppose these two values are close to each other. In that case, a large fraction of the unconditional risk premia of the cross-section is associated with time-series variation, implying that the scope of factor timing is large: the factors with constant risk premia earn limited risk premia, and the performance of most risk factors with positive unconditional risk premia can be improved via timing their conditional expected returns. Although a direct comparison of these two Sharpe ratios is infeasible because Σ_ξ is unobservable, the following proposition provides a feasible comparison criterion.

Proposition 1 *Assume that* $\Sigma_f = \mathbf{I}_K$ *and* $\phi^\top \phi = \mathbf{I}_J$*. Then:*

$$\frac{SR_{\phi^\top f}}{SR_f} \geq \frac{\|\phi^\top \mu_f\|_2}{\|\mu_f\|_2}. \tag{6}$$

In particular, when $\mu_f \in span(\phi)$*, the equality holds and both sides of* (6) *equal to one.*

The proof is in Appendix A.1. In particular, $\|\phi^\top \mu_f\|_2$ is the "observable" version of the unconditional Sharpe ratio (average return over standard deviation) of the factors with time-varying risk premia, and $\|\mu_f\|_2$ is that of all candidate risk factors. This proposition suggests that the ratio of $\|\phi^\top \mu_f\|_2$ to $\|\mu_f\|_2$ provides a lower bound to the ratio of the unconditional Sharpe ratio of the factors with time-varying risk premia to that of all candidate risk factors.

In summary, identifying the loadings of the factors on the state variables ϕ reduces the number of factors with time-varying risk premia and alleviates the signal-to-noise problem in linking risk premia variation to the economic fundamentals. Moreover, it enables me to study how the state variables drive the risk premia of the entire cross-section of risk factors and also sheds light on how much of the unconditional systematic risk premia is associated with time-series variation. Furthermore, Appendix B discusses how identifying the loadings of the factors on the state variables complements conventional unconditional methods to address the "factor zoo" problem when constructing a conditional asset-pricing model. In the next section, I propose a method that identifies (a subspace of) ϕ and incorporates frequency-specific economic structure into the latent state variables that drive risk premia variation.

[10] More precisely, this statement holds under the assumption that the factor innovations are homoskedastic. In this manuscript, I focus on the time-varying conditional expected returns of factors and do not explicitly model the conditional volatility of factors. Examples that study the conditional volatility include, for example, Moreira and Muir (2017).

2.3 Frequency-specific risk factors

Identifying the loadings of the risk factors on the state variables ϕ is nontrivial. In particular, system (2) is not a conventional factor model because the factor innovations ξ_{t+1} can be cross-sectionally correlated, and thus standard methods (e.g., the methods in Section 6.4 of Campbell, Lo, and MacKinlay, 1997) cannot be applied. Moreover, the number of state variables J is unobservable, and the state variables x_t are latent and do not have any economic structure. Thus, even if the factors with time-varying risk premia in (4) are identified, without any economic restrictions, their risk premia are potentially explained by a large number of candidate economic variables, posing a high-dimensional challenge for researchers to understand which economic variables drive the systematic risk premia.

In this section, I propose a nonparametric method that identifies the loadings of the risk factors on the state variables. Moreover, this method also economically regularizes the aforementioned high-dimensional challenge by incorporating frequency-specific economic structure into the state variables that drive risk premia. Specifically, note that equation (2) is equivalent to

$$f_{t+1} = \mu_f + \phi G G^\top x_t + \xi_{t+1}, \tag{7}$$

where G is a $J \times J$-dimensional rotation matrix that satisfies $GG^\top = G^\top G = \mathbf{I}_J$. In this equation, ϕG, which satisfies $(\phi G)^\top (\phi G) = \mathbf{I}_J$, contains the loadings of the factors on the latent state variables $G^\top x_t$. Following equation (4), the J factors with time-varying risk premia are

$$G^\top \phi^\top f_{t+1} = G^\top \phi^\top \mu_f + G^\top x_t + G^\top \phi^\top \xi_{t+1}. \tag{8}$$

The method proposed in this section identifies (a subspace of) ϕG, and the state variables $G^\top x_t$ that drive the systematic risk premia are equipped with frequency-specific economic structure, that is, the variation of each of these state variables concentrates on a particular frequency interval.

The motivation for my method is that the factor innovations ξ_{t+1} and state variables x_t that drive the systematic risk premia have different properties in the frequency domain. Specifically, unlike the conditionally uncorrelated factor innovations, many economic variables proposed in the literature that drive risk premia in theoretical asset-pricing models or predict asset returns empirically are persistent. In the frequency domain, this suggests that the factor innovations have equal low- and high-frequency variation, while the variation of the state variables concentrates on the low-frequency domain. Such difference between the factor innovations and state variables enables me to identify the loadings of the risk factors on the state variables.

To set the stage, I first revisit some basics of time-series analysis in the frequency domain. For illustrative purposes, in the remainder of this section, I assume that the risk factors have zero mean, that is, $\mu_f = \mathbf{0}$. All results in this section hold without this assumption. The spectral density matrix of the risk factors at frequency ω is the Fourier transform of the auto-covariance matrices of the factors:[11]

$$\mathbf{g}_f(\omega) = \sum_{k=-\infty}^{\infty} \Sigma_f(k)\, e^{-i\omega k}, \tag{9}$$

[11] For illustrative purposes, I implicitly assume that the auto-covariance matrices of f_{t+1} are absolutely summable, that is, $\sum_{k=-\infty}^{\infty} |\Sigma_f(k)| < \infty$, and thus the spectral density matrices of f_{t+1} exist. The results in this section are unaltered without this assumption.

where $i = \sqrt{-1}$, and $\omega \in [-\pi, \pi]$ denotes the frequency. Intuitively, a frequency characterizes the length of a cycle. For example, $\omega = 2\pi/12$ corresponds to a cycle of $2\pi/\omega = 12$ periods.[12] The inverse Fourier transform of the spectral density matrices recovers the auto-covariance matrices of the factors:

$$\Sigma_f(k) = \frac{1}{2\pi}\int_{-\pi}^{\pi} \mathbf{g}_f(\omega)\, e^{i\omega k} d\omega. \tag{10}$$

Therefore, the covariance matrix of the factors can be written as

$$\Sigma_f = \Sigma_f(0) = \frac{1}{2\pi}\int_{-\pi}^{\pi} \mathbf{g}_f(\omega)\, d\omega. \tag{11}$$

According to the Spectral Representation Theorem (see, for instance, Brockwell and Davis, 2009), any covariance-stationary time series X_t with zero mean can be represented as

$$X_t = \int_{-\pi}^{\pi} e^{i\omega t} dZ(\omega) = \int_{-\pi}^{\pi} [\cos(\omega t) + i\sin(\omega t)]\, dZ(\omega), \tag{12}$$

where $Z(\omega)$ is a complex-valued stochastic process defined on $[-\pi, \pi]$ with zero mean and orthogonal increments. Furthermore, $Z(\omega)$ satisfies

$$\mathbb{V}ar[Z(\omega_2) - Z(\omega_1)] = \frac{1}{2\pi}\int_{\omega_1}^{\omega_2} g_X(\omega)\, d\omega, \tag{13}$$

where $g_X(\omega)$ is the spectral density function of X_t at frequency ω. Therefore, equations (12) and (13) suggest that sine and cosine waves of frequency ω contribute to the total variation of X_t, with magnitude $\mathbb{V}ar[dZ(\omega)] = g_X(\omega)/2\pi$. Since the Spectral Representation Theorem also applies to multivariate cases, I can rewrite equation (11) to decompose the covariance matrix of the factors into a low-frequency component Σ_f^{LF} and a high-frequency component Σ_f^{HF}:

$$\Sigma_f = \underbrace{\frac{1}{2\pi}\int_{[-\omega,\omega]} \mathbf{g}_f(\lambda)\, d\lambda}_{\Sigma_f^{\text{LF}}} + \underbrace{\frac{1}{2\pi}\int_{[-\pi,-\omega)\cup(\omega,\pi]} \mathbf{g}_f(\lambda)\, d\lambda}_{\Sigma_f^{\text{HF}}}, \tag{14}$$

where the first and second terms on the right-hand side correspond to the variation of the factors contributed by waves with frequencies that lie in the low-frequency interval $[-\omega, \omega]$ and the high-frequency interval $[-\pi, -\omega) \cup (\omega, \pi]$, respectively. In other words, they are the *frequency-specific covariance matrices* of the factors. In general, such a frequency decomposition of the covariance matrix of the factors can be generalized to multiple frequency intervals.

The spectral density matrices of the factors defined in (9) are related to those of the state variables and factor innovations. Specifically, note that the factor innovations are conditionally uncorrelated. Therefore, their auto-covariance matrices satisfy $\Sigma_\xi(k) = \mathbf{0}$ for $k \neq 0$. As a result, the spectral density matrix of the factor innovations at frequency ω is

$$\mathbf{g}_\xi(\omega) = \sum_{k=-\infty}^{\infty} \Sigma_\xi(k)\, e^{-i\omega k} = \Sigma_\xi, \tag{15}$$

[12] The absolute value of the frequency is less or equal to π because one cannot identify cycles shorter than two periods of a time series.

which is constant across different frequencies. Moreover, since $\Sigma_f = \phi\Sigma_x\phi^\top + \Sigma_\xi$, by letting $\mathbf{g}_x(\omega)$ be the spectral density matrix of the state variables at frequency ω, the covariance matrix of the factors can be written as

$$\Sigma_f = \frac{1}{2\pi}\int_{-\pi}^{\pi} \phi\mathbf{g}_x\left(\omega\right)\phi^\top d\omega + \Sigma_\xi = \frac{1}{2\pi}\int_{-\pi}^{\pi}\left[\phi\mathbf{g}_x\left(\omega\right)\phi^\top + \Sigma_\xi\right]d\omega. \tag{16}$$

This equation, together with equation (11), suggests the spectral density matrix of the factors at frequency ω satisfies

$$\mathbf{g}_f\left(\omega\right) = \phi\mathbf{g}_x\left(\omega\right)\phi^\top + \Sigma_\xi. \tag{17}$$

Therefore, the low- and high-frequency covariance matrices of the factors defined in equation (14) can be written as

$$\Sigma_f^{\text{LF}} = \phi\underbrace{\left(\frac{1}{2\pi}\int_{[-\omega,\omega]}\mathbf{g}_x\left(\lambda\right)d\lambda\right)}_{\Sigma_x^{\text{LF}}}\phi^\top + \frac{\omega}{\pi}\Sigma_\xi, \quad \text{and} \tag{18}$$

$$\Sigma_f^{\text{HF}} = \phi\underbrace{\left(\frac{1}{2\pi}\int_{[-\pi,-\omega)\cup(\omega,\pi]}\mathbf{g}_x\left(\lambda\right)d\lambda\right)}_{\Sigma_x^{\text{HF}}}\phi^\top + \frac{\pi-\omega}{\pi}\Sigma_\xi, \tag{19}$$

respectively, where Σ_x^{LF} and Σ_x^{HF} correspond to the low- and high-frequency covariance matrices of the state variables.

The low- and high-frequency covariance matrices of the risk factors in equations (18) and (19) are informative in identifying the loadings of the factors on the state variables. In particular, the spectral density matrix of the factor innovations is constant across different frequencies, which implies that the factor innovations have equal variation in each frequency. If the spectral density matrix of the state variables is also constant across frequencies, then for any linear combination $\theta \in \mathbb{R}^K$ of the risk factors, the ratio of its low-frequency to high-frequency variation is

$$\frac{\theta^\top\Sigma_f^{\text{LF}}\theta}{\theta^\top\Sigma_f^{\text{HF}}\theta} = \frac{\theta^\top\phi\left(\frac{1}{2\pi}\int_{[-\omega,\omega]}\mathbf{g}_x\left(\lambda\right)d\lambda\right)\phi^\top\theta + \frac{\omega}{\pi}\theta^\top\Sigma_\xi\theta}{\theta^\top\phi\left(\frac{1}{2\pi}\int_{[-\pi,-\omega)\cup(\omega,\pi]}\mathbf{g}_x\left(\lambda\right)d\lambda\right)\phi^\top\theta + \frac{\pi-\omega}{\pi}\theta^\top\Sigma_\xi\theta} = \frac{\omega}{\pi-\omega},$$

which is the same as the ratio of the length of the low-frequency to that of the high-frequency interval.

However, the findings in the literature suggest that the spectral density matrix of the state variables that drive systematic risk premia is different across frequencies. For example, the habit model of Campbell and Cochrane (1999) suggests that the only state variable that drives the systematic risk premia is the sensitivity function of the surplus consumption ratio, which, as illustrated in Figure 1 of the introduction, is highly persistent. This suggests that the spectral density function of this state variable is higher in the low-frequency interval. Thus, compared to the factor innovations, the variation of this state variable is concentrated on the low-frequency interval, satisfying

$$\frac{\left(\sigma_x^2\right)^{\text{LF}}}{\left(\sigma_x^2\right)^{\text{HF}}} > \frac{\omega}{\pi-\omega}. \tag{20}$$

Suppose the sensitivity function of the surplus consumption ratio is the only state variable that drives the systematic risk premia. In this case, it can be shown that, among all linear combinations $\theta \in \mathbb{R}^K$ of the risk factors, the optimal θ^* that maximizes $(\theta^\top \Sigma_f^{\text{LF}} \theta)/(\theta^\top \Sigma_f^{\text{HF}} \theta)$ satisfies $\theta^* \propto \phi$. In other words, if the only state variable that drives the systematic risk premia is "low-frequency," that is, it satisfies inequality (20), then I can identify the loadings of the factors on this variable by maximizing $(\theta^\top \Sigma_f^{\text{LF}} \theta)/(\theta^\top \Sigma_f^{\text{HF}} \theta)$.

I now formalize and generalize the previous argument. Specifically, to identify the loadings of the risk factors on a low-frequency state variable, I solve the following problem:

$$\max_{c,\theta} \quad c, \tag{21}$$

$$\text{s.t.} \quad \theta^\top \left(\Sigma_f^{\text{LF}} - c \Sigma_f^{\text{HF}} \right) \theta \geq 0. \tag{22}$$

I make two remarks on this problem. First, the frequency-specific covariance matrices of the risk factors in constraint (22) can be flexibly specified for other frequency intervals. In this section, I let $[-\omega, \omega]$ be the low-frequency interval and $[-\pi, -\omega) \cup (\omega, \pi]$ be the high-frequency interval. Generally, the frequency intervals can be flexibly specified depending on the application, and their union does not need to cover the entire interval $[-\pi, \pi]$. For example, I can specify the low-frequency interval as $[-\omega_2, -\omega_1] \cup [\omega_1, \omega_2]$ and the high-frequency interval as $[-\omega_4, -\omega_3] \cup [\omega_3, \omega_4]$, where $0 < \omega_1 < \omega_2 < \omega_3 < \omega_4 < \pi$, and use the frequency-specific covariance matrices corresponding to these two intervals in (22). Moreover, solving problem (21) can identify the loadings of the factors on a general frequency-specific state variable. For instance, I can also switch Σ_f^{LF} and Σ_f^{HF} in (22) to identify the loadings of the factors on a high-frequency state variable. The main results in this section are unaltered with the aforementioned modifications.

Second, numerically solving this problem is straightforward. In particular, it is easy to show that constraint (22) is binding at the optimum, and thus c is the ratio of the low- to high-frequency variation of $\theta^\top f_{t+1}$. The optimal c^* is the maximum c that drives the highest eigenvalue of $\Sigma_f^{\text{LF}} - c\Sigma_f^{\text{HF}}$ to zero, and the optimal θ^* is the eigenvector that corresponds to the highest eigenvalue of $\Sigma_f^{\text{LF}} - c^* \Sigma_f^{\text{HF}}$.[13]

The following lemma sets the stage for the solution to problem (21).

Lemma 1 *The highest eigenvalue of*

$$W = \left(\mathbf{I}_J - \Sigma_x + \frac{\pi}{\pi - \omega} \Sigma_x^{HF} \right)^{-1/2} \left(\mathbf{I}_J - \Sigma_x + \frac{\pi}{\omega} \Sigma_x^{LF} \right) \left(\mathbf{I}_J - \Sigma_x + \frac{\pi}{\pi - \omega} \Sigma_x^{HF} \right)^{-1/2} \tag{23}$$

is greater than one if and only if there exists $\eta \in \mathbb{R}^J$ such that

$$\frac{\eta^\top \Sigma_x^{LF} \eta}{\eta^\top \Sigma_x^{HF} \eta} > \frac{\omega}{\pi - \omega}. \tag{24}$$

The proof is in Appendix A.2. Similar to inequality (20), the interpretation of inequality (24) is that the linear combination η of the state variables has higher low-frequency than high-frequency variation, and thus $\eta^\top x_t$ is a low-frequency state variable. If the panel of latent state variables x_t guarantees the existence of a low-frequency state variable, then the following proposition shows that solving problem (21) identifies the loadings of the risk factors on a low-frequency state variable.

[13] Note that θ is not uniquely identified in problem (21). In particular, if θ^* is a solution to problem (21), then $a\theta^*$ is also a solution for any $a \neq 0$.

Proposition 2 *Assume that* $\Sigma_f = \mathbf{I}_K$ *and* $\phi^\top \phi = \mathbf{I}_J$. *If there exists* $\eta \in \mathbb{R}^J$ *such that* $\eta^\top x_t$ *is a low-frequency state variable, that is, inequality* (24) *holds, then the solution to problem* (21) *is*

$$c^* = \frac{\omega}{\pi - \omega} \lambda_1^W, \quad \textit{and} \tag{25}$$

$$\theta^* \propto \phi \left(\mathbf{I}_J - \Sigma_x + \frac{\pi}{\pi - \omega} \Sigma_x^{HF} \right)^{-1/2} H_1 \equiv \phi G_1, \tag{26}$$

where $\lambda_1^W > 1$ *and* H_1 *are the highest eigenvalue of* W *in* (23) *and the corresponding eigenvector, respectively; otherwise, the solution to problem* (21) *is*

$$c^* = \frac{\omega}{\pi - \omega}, \quad \textit{and} \tag{27}$$

$$\theta^* \propto \Lambda y, \tag{28}$$

where Λ *is a* $K \times (K - J)$*-dimensional matrix whose columns are orthogonal to the columns of* ϕ*, and* y *is an arbitrary nonzero* $(K - J)$*-dimensional vector.*

The proof is in Appendix A.3. Proposition 2 shows that, if the panel of latent state variables guarantees the existence of a low-frequency state variable, then solving problem (21) identifies the loadings of the factors on the state variable

$$x_t^1 = \frac{1}{\|G_1\|_2} G_1^\top x_t. \tag{29}$$

It is straightforward to show that x_t^1 is a low-frequency state variable, and the ratio of its low-frequency to high-frequency variation is higher than c^* in (27). By normalizing $(\theta^*)^\top \theta^* = 1$, the risk factor (with time-varying risk premium) I identify by solving problem (21) is

$$(\theta^*)^\top f_{t+1} = x_t^1 + \frac{G_1^\top \phi^\top \xi_{t+1}}{\|G_1\|_2}, \tag{30}$$

which is a low-frequency risk factor because its ratio of low-frequency to high-frequency variation is $c^* > \omega/(\pi - \omega)$, and its time-varying risk premium is driven by the low-frequency state variable x_t^1.[14] Thus, the approach proposed in this section not only identifies the loadings of the factors on the state variables, but also incorporates the latent state variables with frequency-specific economic structure. In this case, one does not expect a high-frequency economic variable to highly explain the risk premium of $(\theta^*)^\top f_{t+1}$ because it is driven by a low-frequency state variable. Thus, compared to explaining the risk premium of a risk factor that loads on a state variable without the frequency-specific economic structure, the number of potential variables explaining the risk premium of the low-frequency factor $(\theta^*)^\top f_{t+1}$ is reduced. As a result, the approach proposed in this section imposes economic regularization on the high-dimensional challenge of "which economic variables explain the time-varying risk premia of the cross section?"

Moreover, problem (21) can be solved recursively to identify the loadings of the factors on other low-frequency state variables (or, in general, other state variables whose variation concentrates on other frequency intervals). For completeness, I summarize this result in the following proposition.

[14] The signal of the low-frequency state variable x_t^1 is maximized in $(\theta^*)^\top f_{t+1}$ because θ^* is orthogonal to Λ. Specifically, the variation of $(\theta^*)^\top f_{t+1}$ driven by the state variable x_t^1 is $\mathbb{V}ar(x_t^1)$, and that driven by factor innovations is $G_1^\top \phi^\top \Sigma_\xi \phi G_1 / \|G_1\|_2^2$. It can be shown that for any $\theta' \propto a\phi G_1 + b\Lambda y$ with $a, b \neq 0$ and $y \in \mathbb{R}^{K-J}$, the signal-to-noise ratio of $(\theta')^\top f_{t+1}$ is lower than $\mathbb{V}ar(x_t^1)/(G_1^\top \phi^\top \Sigma_\xi \phi G_1 / \|G_1\|_2^2)$.

Proposition 3 *Suppose that there exists $\eta \in \mathbb{R}^J$ such that $\eta^\top x_t$ is a low-frequency state variable, and one solves problem* (21) *to identify $\theta^* \propto \phi G_1$ as in* (26)*. Let $f_{t+1}^\perp$ be the residuals of regressing f_{t+1} on $(\theta^*)^\top f_{t+1}$. The following results hold:*

(a) The regression residuals $f_{t+1}^\perp$ do not load on the low-frequency state variable x_t^1 defined in (29)*.*

(b) Let $x_t^\perp = V^\top x_t$, where V is a $J \times (J-1)$ matrix whose columns are orthonormal and orthogonal to G_1. Suppose there exists $\eta' \in \mathbb{R}^{J-1}$ such that $(\eta')^\top x_t^\perp$ is a low-frequency state variable. Moreover, suppose one obtains the low- and high-frequency covariance matrices of $U^\top f_{t+1}^\perp$, where U is a $K \times (K-1)$ matrix whose columns are orthonormal and orthogonal to ϕG_1, and then solves problem (21) *based on these matrices. Then, the solution identifies the loadings of the risk factors on a low-frequency state variable spanned by $x_t^\perp$.*

The proof is in Appendix A.4. Part (a) suggests that $(\theta^*)^\top f_{t+1}$ fully captures the explanatory power of the state variable x_t^1 for the risk premia of the cross-section of risk factors. In particular, by regressing $(\theta^*)^\top f_{t+1}$ out of the risk factors, x_t^1 does not drive the risk premia of the regression residuals. Part (b) implies that I can solve problem (21) recursively to identify the loadings of the risk factors on other low-frequency state variables, provided that the remaining $J-1$ latent state variables guarantee the existence of a low-frequency state variable.

In summary, solving problem (21) identifies (a subspace of) the loadings of the factors on the state variables, and thus identifies the factors with time-varying risk premia. Moreover, the latent state variables that drive the risk premia of these factors are equipped with frequency-specific economic structure, which regularizes the high-dimensional challenge to understand which economic variables explain the systematic risk premia.

2.4 Empirical set-up

In this section, I introduce the empirical set-up of this manuscript. Specifically, Section 2.4.1 introduces the candidate risk factors from which I identify the low-frequency factors. Section 2.4.2 specifies the frequency intervals I consider in the empirical analysis and introduces the method to estimate the frequency-specific covariance matrices of the risk factors. Section 2.4.3 describes the method and observable variables for predicting the returns of the low-frequency factors.

2.4.1 Candidate risk factors

I use the restricted version of the Instrumented Principal Component Analysis (IPCA) of Kelly et al. (2019) to identify the candidate risk factors of asset returns, from which I estimate the low-frequency risk factors.[15] Specifically, IPCA models the loadings of the stocks on the risk factors as a linear function of firm-specific characteristics:

$$\beta_t = Z_t \Gamma_\beta, \tag{31}$$

[15] My approach of first selecting a set of candidate risk factors and then obtaining predictable factors from these factors is not unique to the literature. For example, Haddad et al. (2020) start from a set of pricing factors, and then use PCA to reduce this set of factors to a few dominant components, which they find highly predictable.

where Z_t is a $N \times L$ matrix that contains the L firm-specific characteristics of the N stocks at time t, and $\Gamma_\beta \in \mathbb{R}^{L\times K}$ is the mapping matrix that transforms the observable characteristics of the stocks to their loadings on the risk factors. The IPCA risk factors can be estimated jointly with the mapping matrix via alternating least squares.[16] Moreover, once I obtain the estimates of the low-frequency factors by solving problem (21), the mapping matrix that transforms the observable characteristics of individual stocks to their loadings on these low-frequency factors can be estimated by regressing the stock returns on the interaction of the low-frequency factors and characteristics.

One of the main reasons for using the IPCA factors as the candidate risk factors is that IPCA explicitly considers the time variation of factor loadings (because the characteristics are time-varying). Thus, it is expected that the predictable components of the IPCA factors mainly capture risk premia, whose variation is the main focus of this manuscript, rather than the predictable time-varying loadings on the true systematic risks. In general, other candidate factors of asset returns, such as returns of a comprehensive set of anomaly portfolios, the factors extracted from a cross-section of returns using PCA, or the risk-premium PCA factors of Lettau and Pelger (2020), can also be used as input to identify the low-frequency factors. However, compared to the IPCA factors, these candidate factors may have time-varying loadings on the true systematic risks. Thus, an economic variable's predictability of the returns of these factors might stem from that this variable drives the variation of the loadings of these factors on the true systematic risks or the interaction between the time-varying loadings and risk premia. This introduces noises in understanding how the time-varying dynamics of the systematic risk premia relate to the economic fundamentals.

Moreover, **?** show that portfolios sorted based on the past values of firm characteristics capture risk premia unexplained by those based on the present values of characteristics. This implies that the persistent and transitory components of characteristics are associated with different components of risk compensation. Thus, to ensure that the candidate risk factors obtained from IPCA capture well the systematic risk premia, in addition to the Basic-IPCA framework, in which only the present values of characteristics instrument the loadings of stocks on systematic risks, I also consider a frequency-specific IPCA (Freq-IPCA) framework, in which the low- and high-frequency components of characteristics are different instruments for the factor loadings. Specifically, the frequency-specific IPCA framework models the loadings of the stocks on the risk factors as:

$$\beta_t = Z_t^{\mathrm{LF}}\Gamma_\beta^{\mathrm{LF}} + Z_t^{\mathrm{HF}}\Gamma_\beta^{\mathrm{HF}}, \tag{32}$$

where Z_t^{LF} and Z_t^{HF} that satisfy $Z_t^{\mathrm{LF}} + Z_t^{\mathrm{HF}} = Z_t$ are the low- and high-frequency components of the L characteristics of the N stocks at time t, respectively, and $\Gamma_\beta^{\mathrm{LF}}, \Gamma_\beta^{\mathrm{HF}} \in \mathbb{R}^{L\times K}$ are the mapping matrices that transform the low- and high-frequency components of the characteristics to the loadings of the stocks on the risk factors. Thus, the Freq-IPCA framework effectively uses $2L$ characteristics to instrument the factor loadings of stocks.

The frequency decomposition of the stock characteristics in the Freq-IPCA framework provides two additional benefits. First, the low- and high-frequency components of the characteristics may instrument the covariances between the stock returns and aggregate risks differently. In this case, modeling factor loadings as a linear function of the present values of characteristics in the Basic-IPCA framework is misspecified, and the extra flexibility of the Freq-IPCA framework provides a better description

[16] I refer to Kelly et al. (2019) and Kelly, Pruitt, and Su (2020) for a more detailed discussion of the IPCA framework.

of how stocks expose to systematic risks. Second, together with the identification of the frequency-specific factors, the Freq-IPCA model enables me to establish an economically important link between an individual stock's frequency-specific risk compensation and the frequency-specific variation in its characteristics. For example, suppose only the low-frequency component of a characteristic contributes to the loadings of the stocks on a low-frequency factor. In this case, this implies that the compensation of individual stocks originating from this low-frequency factor does not vary substantially during a short period. Thus, identifying frequency-specific factors from the Freq-IPCA factors is a general framework for understanding the frequency-specific properties of the compensation of individual stocks.

I use the same data as Kelly et al. (2019) to estimate the IPCA model. In particular, the dataset is from Freyberger, Neuhierl, and Weber (2020) and includes monthly stock returns and characteristics from July 1964 to May 2014.[17] Each firm is associated with 36 characteristics in categories such as past returns (e.g., momentum and reversal) and value-related characteristics (e.g., book-to-market and earnings-to-price ratios). Table 14 in Appendix F provides a more detailed overview of the considered characteristics. In addition, I include a constant on top of the 36 characteristics so that the total number of characteristics is 37. The estimation of the IPCA model requires each stock-month observation to have non-missing return and characteristics, restricting the dataset to end up with 1,403,544 stock-month observations for 12,813 unique stocks. Following Kelly et al. (2019), I cross-sectionally transform the characteristics period-by-period. Specifically, I calculate the stocks' ranks for each characteristic, then divide the ranks by the number of non-missing observations and subtract 0.5. In this case, each characteristic is in the $[-0.5, 0.5]$ interval, and thus I focus on the ordering instead of the magnitude of the characteristics. In the estimation of the IPCA model, I normalize Γ_β such that $\Gamma_\beta^\top \Gamma_\beta = \mathbf{I}_K$ and let the in-sample average return of the estimated IPCA factors be positive.

To estimate the Freq-IPCA model, I let the low-frequency component of the characteristics at time t be the time-series average of the characteristics available in the past 36 months, and let the high-frequency component of the characteristics at time t be the difference between the present value of the characteristics and the low-frequency component of the characteristics at time t. By including a constant, the total number of characteristics used in the Freq-IPCA framework is $36 + 36 + 1 = 73$.

2.4.2 Estimating frequency-specific risk factors

I specify three frequency intervals to estimate different low-frequency risk factors. Specifically, I consider a typical business-cycle-frequency (BCF) interval that captures cycles between 1.5 and eight years (18 and 96 months). Moreover, as illustrated in the introduction, theoretical asset-pricing models such as the habit model of Campbell and Cochrane (1999) imply that state variables whose variation concentrates on extremely low frequencies drives the variation of systematic risk premia, and Dew-Becker and Giglio (2016) show that cycles longer than the business cycle are significantly priced in the equity market. Thus, I also consider a below-business-cycle-frequency (B-BCF) interval that captures cycles longer than eight years (96 months). In addition, I consider a high-frequency interval that captures cycles shorter than six months. Correspondingly, the frequency-specific covariance matrices of the risk factors in these frequency intervals are:

$$\Sigma_f^{\text{B-BCF}} = \frac{1}{2\pi} \int_{[-\frac{2\pi}{96}, \frac{2\pi}{96}]} \mathbf{g}_f(\omega)\, d\omega, \tag{33}$$

[17] Similar to Kelly et al. (2019), I only focus on months with at least 100 cross-sectional observations.

$$\Sigma_f^{\text{BCF}} = \frac{1}{2\pi} \int_{[-\frac{2\pi}{18}, -\frac{2\pi}{96}) \cup (\frac{2\pi}{96}, \frac{2\pi}{18}]} \mathbf{g}_f(\omega)\, d\omega, \quad \text{and} \tag{34}$$

$$\Sigma_f^{\text{HF}} = \frac{1}{2\pi} \int_{[-\pi, -\frac{2\pi}{6}] \cup [\frac{2\pi}{6}, \pi]} \mathbf{g}_f(\omega)\, d\omega, \tag{35}$$

respectively. The empirical estimation of these frequency-specific covariance matrices is via discrete Fourier transform and is explained in detail in Appendix C. In the main body of this manuscript, I focus on identifying the below-business-cycle- and business-cycle-frequency factors, both of which reflect the economically important low-frequency dynamics of risk premia.

One potential concern is whether using the estimated frequency covariance matrices to identify the low-frequency factors generates substantial estimation error. In Appendix D, I design a simple simulation to illustrate that, under the empirical setup of this manuscript, using the estimated frequency covariance matrices in place of their population counterparts does not generates substantial estimation error to the low-frequency factors in simulated samples that mimic the real-world data. In particular, by assuming that there is one below-business-cycle-frequency state variable driving the risk premia variation of all candidate factors, the simulation shows that the time-series correlation between the estimated and population below-business-cycle-frequency factors is higher than 0.84 in half of the simulated samples. This validates that the empirical findings of this manuscript are not driven by estimation error.

2.4.3 Method and variables to predict low-frequency factor returns

Note that an unobservable low-frequency state variable drives the variation of the risk premium of a low-frequency factor. By assuming that this latent state variable is a factor of a panel of observable economic variables, I apply the generalization of the partial least squares (PLS) regression of Kelly and Pruitt (2015) to predict the returns of the low-frequency factors using the panel of observable economic variables, thus understanding how these observable variables explain the systematic risk premia.

Specifically, I consider M observable economic variables and assume that these variables have the following factor structure:

$$z_t = b_0 + b x_t^1 + e_t, \tag{36}$$

where z_t contains the M observable economic variables, b_0 captures the level of the observable variables, $b \in \mathbb{R}^M$ contains the loadings of observable variables on x_t^1, the latent state variable that drives the risk premium of the low-frequency factor $(\theta^*)^\top f_{t+1}$ in (30), and e_t captures the variation of the observable variables uncaptured by the latent state variable x_t^1.[18] The three-pass regression filter of Kelly and Pruitt (2015) that generalizes PLS to predict the returns of the low-frequency factor consists of the following three stages. In the first stage, I regress the observable economic variables on the returns of the low-frequency factor:

$$z_t = \widehat{\kappa}_0 + \widehat{\kappa} \left[(\theta^*)^\top f_{t+1} \right] + u_t, \tag{37}$$

where $\hat{\kappa}$ captures the sensitivity of each observable economic variable to the latent state variable x_t^1 because the observable economic variables are uncorrelated with future innovations of the low-frequency

[18] In general, e_t is also allowed to have a factor structure. For illustrative purposes, I highlight the factor that drives the risk premium of the low-frequency factor.

factor. In the second stage, at each time t, I run a cross-sectional regression of the observable economic variables on their loadings on the low-frequency factor estimated in the first stage:

$$z_t = \widehat{c}_t + \widehat{\kappa}\widehat{g}_t + w_t, \tag{38}$$

where the regression coefficient $\hat{g}_t$, a linear combination of the observable economic variables, is the estimated PLS factor that predicts the returns of the low-frequency factor. In the third stage, I run the standard predictive regression of the low-frequency factor on the PLS factor $\hat{g}_t$ estimated from the second stage, and the R-squared obtained from this regression summarizes the predictive power of the observable economic variables for the returns of the low-frequency risk factor.

The three-pass regression filter combines the informative component of every observable economic variable into the PLS factor, which predicts the returns and thus explains the risk premium of the low-frequency factor. I refer to Kelly and Pruitt (2015) for a detailed discussion of the econometric development and proofs of the three-pass regression filter. Importantly, the authors show that this procedure is consistent when the number of observable economic variables and time-series observations becomes large, even when the observable economic variables load on other factors that do not drive the risk premium of the low-frequency factor. In addition, they provide the asymptotic distribution theory for the third-stage predictive coefficient under weak conditions. Moreover, the above procedure can be repeated to generate multiple PLS factors, which may deliver better performance for predicting the returns of the low-frequency factor. In the empirical analysis, to balance between goodness-of-fit and parsimony, I estimate at most two PLS factors from a panel of observable economic variables to predict the returns of the below-business-cycle- and business-cycle-frequency factors.

I use two sets of observable economic variables to predict the returns of the two low-frequency factors. The first set of variables comprises 19 macroeconomic time series including, for instance, inflation, NBER-dated recessions, and the uncertainty measures of Jurado, Ludvigson, and Ng (2015) and Ludvigson, Ma, and Ng (2021). The second set of variables comprises 24 financial indicators including, for instance, aggregate liquidity (Pastor and Stambaugh, 2003), investor sentiment (Baker and Wurgler, 2006; Huang, Jiang, Tu, and Zhou, 2015), and the variables considered by Welch and Goyal (2008) such as the dividend-price ratio and Treasury bill rate. Table 15 in Appendix F provides a more detailed description of these macroeconomic and financial variables. The macroeconomic variables span from June 1964 to April 2014, while the financial variables span from January 1970 to April 2014.

3 Empirical results

In this section, I discuss the empirical performance of the low-frequency risk factors estimated from the IPCA factors.

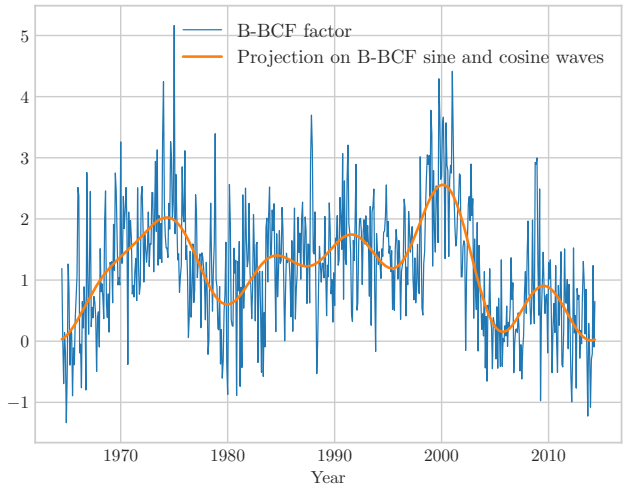
3.1 Return and composition of low-frequency factors

Note that the maximum number of IPCA factors that can be identified is the same as the number of characteristics. Thus, using the dataset described in Section 2.4.1, the maximum number of Basic-IPCA factors I identify is 37, and the maximum number of Freq-IPCA factors I identify is 73. By letting all 37 Basic-IPCA factors be the candidate risk factors, I estimate one below-business-cycle-frequency factor

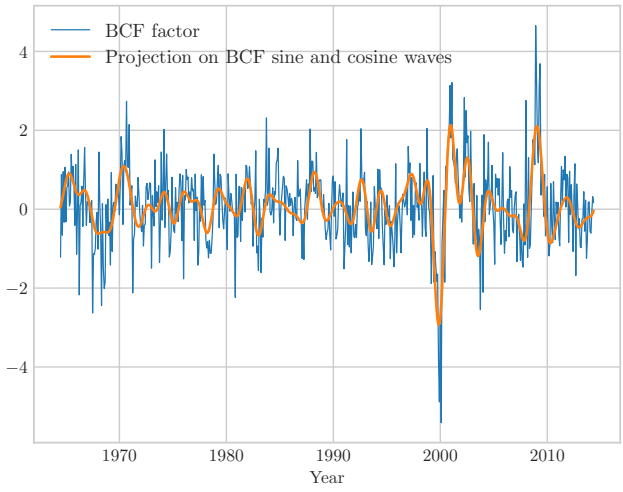
Figure 2: Low-frequency risk factors estimated from 37 Basic-IPCA factors

Panels A and B illustrate the in-sample estimates of a below-business-cycle-frequency factor and a business-cycle-frequency factor estimated from all 37 Basic-IPCA risk factors, respectively. Both factors are normalized to have positive average returns and unit standard deviation. In each panel, the blue line gives the monthly factor returns, and the orange line is the projection of the factor returns on sine and cosine waves whose frequencies lie within the corresponding low-frequency interval.

(A) Below-business-cycle-frequency factor



(B) Business-cycle-frequency factor



and one business-cycle-frequency factor.[19] Panels A and B of Figure 2 illustrate the in-sample estimates of these two factors, respectively. Both factors are normalized to have positive average returns and unit standard deviation. In each panel, the blue line gives the monthly factor returns, and the orange line is the projection of the factor returns on sine and cosine waves whose frequencies lie within the corresponding low-frequency interval.

In the empirical identification of the two low-frequency factors, I find that the optimal c^* (the ratio of low-frequency to high-frequency variation of the identified low-frequency factor) that solves problem (21) follows equation (25). This suggests that the risk premia of the cross-section of IPCA factors are associated with time-series variation, driven by below-business-cycle-frequency and business-cycle-frequency state variables. In addition, an unreported simulation study shows that this conclusion is statistically significant with a p-value lower than 0.001. Figure 2 shows that the two factors successfully reflect the low-frequency dynamics of systematic risk premia. For instance, the projection of the returns of the below-business-cycle-frequency factor on sine and cosine waves with cycles longer than eight years shows that a large fraction of the total variation of this factor is low-frequency, suggesting that a low-frequency state variable drives the conditional mean of this factor.[20] In particular, the returns of this factor gradually increased before the mid-1970s and then decreased until 1980. The returns then fluctuated mildly until the mid-1990s, reached a high level in 2000, and then decreased sharply until the

[19]Note that the identification of the low-frequency factors depends on which frequency interval is considered first. In other words, factors obtained from first estimating the below-business-cycle-frequency factors and then the business-cycle-frequency factors are different from those obtained from first estimating the business-cycle-frequency factors and then the below-business-cycle-frequency factors. In the remainder of this manuscript, I proceed with the former approach by first estimating the below-business-cycle-frequency factors and then the business-cycle-frequency factors. The results are qualitatively unchanged with the latter approach.

[20]To avoid potential ambiguity, I highlight that the state variable that drives the risk premium of the below-business-cycle-frequency factor, in general, is not the projection of the factor returns on sine and cosine waves with cycles longer than eight years (the orange line in Panel A of Figure 2).

Table 1: Sharpe ratio of low-frequency factors estimated from 37 Basic-IPCA factors

This table reports the in-sample (unconditional) monthly Sharpe ratio of the low-frequency factors estimated from all 37 Basic-IPCA factors. The first two columns report the Sharpe ratio of the below-business-cycle- and business-cycle-frequency factors, the third column that of the optimal portfolio of the two low-frequency factors, and the last column that of the optimal portfolio of all 37 Basic-IPCA factors.

	B-BCF factor	BCF factor	B-BCF+BCF factors	37 Basic-IPCA factors
Sharpe ratio	1.193	0.081	1.196	1.569

financial crisis. Similarly, the returns of the business-cycle-frequency factor also exhibit low-frequency fluctuations, which is illustrated by their projection on the business-cycle-frequency sine and cosine waves. Specifically, the return of this factor reached a low level in the dot-com bubble and a high level during the financial crisis.

Table 1 reports the in-sample (unconditional) monthly Sharpe ratio of the low-frequency risk factors estimated from all 37 Basic-IPCA factors. The Sharpe ratio of the below-business-cycle-frequency factor is 1.193, only 23.8% lower than that of the optimal portfolio of the 37 Basic-IPCA factors, 1.569. The Sharpe ratio of the below-business-cycle-frequency factor is surprisingly high, considering that the estimation of the low-frequency factors only exploits the covariance and autocovariance but not the average return information of the Basic-IPCA factors. Thus, Proposition 1 suggests that at least 76.2% of the unconditional risk premia of the cross-section is associated with time-series variation driven by a below-business-cycle-frequency state variable. Thus, the performance of most candidate risk factors with positive risk premia can be improved via timing their expected returns (by a below-business-cycle-frequency state variable).

In contrast, the Sharpe ratio of the business-cycle-frequency factor is only 0.081, and thus that of the optimal portfolio of the two low-frequency factors is 1.196, only 0.23% higher than that of the below-business-cycle-frequency factor. Thus, the state variable that drives the risk premium of the business-cycle-frequency factor does not capture a large fraction of the unconditional risk premia of the cross-section. However, the business-cycle-frequency factor delivering a low unconditional Sharpe ratio does not imply that it is uninformative for pricing the cross-section. Instead, it captures the information of candidate factors that are important for the construction of a conditional asset-pricing model, but may be overlooked by conventional methods that focus on capturing factors with positive unconditional risk premia.[21] Specifically, the business-cycle-frequency factor captures the information of candidate factors whose conditional risk premia have business-cycle-frequency variation, even if these factors could possibly earn zero unconditional risk premia. Thus, the conditional risk premia of the business-cycle-frequency factor, by construction, is driven by a business-cycle-frequency state variable. Suppose one successfully identifies this state variable, then the business-cycle-frequency factor, in general, can be exploited to enter the conditional SDF, and one can get closer to conditional mean-variance efficiency by timing the average return of this factor.

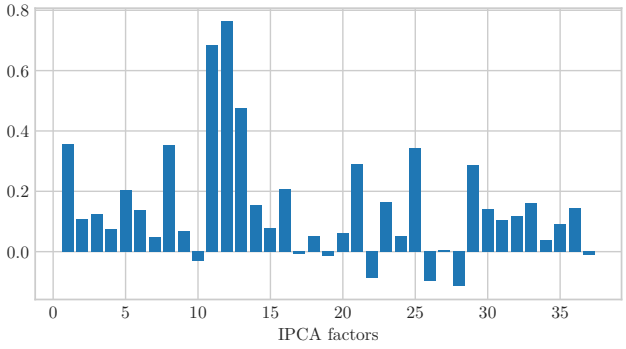
Importantly, a level factor that mimics the market does not explain the risk premia of the two low-frequency factors. Specifically, in unreported results, I find that the below-business-cycle- and business-cycle-frequency factors have low and negative correlations with the market factor (with the

[21] Appendix B provides a more detailed explanation of how the low-frequency factors capture risk factors that are informative for conditional asset-pricing.

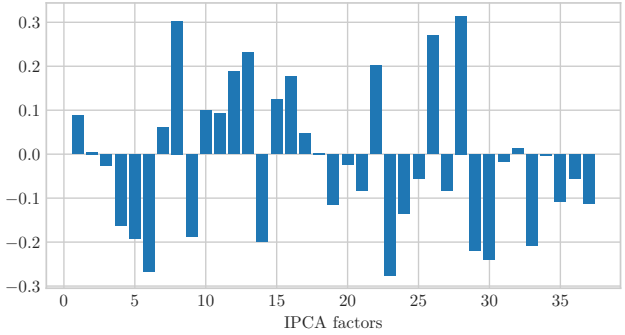
Figure 3: Composition of the low-frequency factors

Panels A and B illustrate the composition of the below-business-cycle- and business-cycle-frequency factors estimated from all 37 Basic-IPCA factors, respectively. Each Basic-IPCA factor is normalized to have a positive mean and unit standard deviation.

(A) Below-business-cycle-frequency factor



(B) Business-cycle-frequency factor



correlation coefficients being -0.103 and -0.194, respectively).[22] In addition, the alpha of the market factor with respect to the below-business-cycle-frequency factor is 1.03% per month with a significant t-statistic 3.69, and that with respect to the business-cycle-frequency factor is 0.58% per month with a significant t-statistic 3.07. Thus, the two low-frequency factors, especially the below-business-cycle-frequency factor, capture risk premia orthogonal to the market. Moreover, the limited correlation between the low-frequency factors and the market factor suggests that, similar to that the market factor is inadequate to fully capture the cross-sectional difference in asset average returns, "market timing" is inadequate to fully characterize the time-series variation of the risk premia of the cross-section. The risk premia of the sources of risk that are orthogonal to the market also have time-series variation.

To understand how the Basic-IPCA factors load on the latent low-frequency state variables that drive the risk premia of the low-frequency factors, Panels A and B of Figure 3 illustrate the weight of the linear combination of the 37 Basic-IPCA factors that result in the below-business-cycle- and business-cycle-frequency factors, respectively. Each Basic-IPCA factor is normalized to have a positive average return and unit standard deviation. Since the Basic-IPCA factors are approximately orthogonal, Proposition 2 implies that each panel of Figure 3 can be approximately viewed as the loadings of Basic-IPCA risk factors on the low-frequency state variables.

Panel A shows that most Basic-IPCA factors have positive loadings on the latent below-business-cycle-frequency state variable: only seven out of the 37 Basic-IPCA factors have negative loadings on this state variable, and the magnitude of these negative loadings is low. Since all Basic-IPCA factors are normalized to have positive average returns, this justifies the high Sharpe ratio of the below-business-cycle-frequency factor. In particular, the 11th, 12th, and 13th Basic-IPCA factors have the highest loadings on the below-business-cycle-frequency state variable, while at the same time, they also have the highest Sharpe ratio.[23] As a result, the latent below-business-cycle-frequency state variable can be approximately interpreted as a fundamental "level" state variable that moves the conditional risk premia of a cross-section of candidate risk factors from the same direction. In particular, when the

[22] I find the low-frequency factors estimated from a small number of IPCA factors also have low correlations with the market factor.

[23] In unreported results, I find the Sharpe ratio of these three factors is 0.559, 0.417, and 0.270, respectively, which are higher than all other Basic-IPCA factors.

below-business-cycle-frequency state variable reaches a high level, the conditional risk premia of most candidate risk factors with positive unconditional risk premia reach a high level.

In contrast, Panel B shows that over half of the 37 Basic-IPCA factors have negative loadings on the business-cycle-frequency state variable, which justifies the low Sharpe ratio of the business-cycle-frequency risk factor. As a result, the latent business-cycle-frequency state variable can be interpreted as a "long-short" state variable that moves the risk premia of different candidate factors from different directions. For example, suppose that the candidate risk factors are industry-specific factors. In this case, when the business-cycle-frequency state variable reaches a high level, the conditional risk premia of the stocks in one industry could become higher, while those of the stocks in another competing industry could become lower.[24]

Importantly, both panels of Figure 3 show that the loadings of the higher-order IPCA factors on the low-frequency state variables are non-negligible. For instance, Panel B shows that the loadings on the business-cycle-frequency state variable of many higher-order IPCA factors are as high (in absolute value) as those of the lower-order IPCA factors. This suggests that the risk premia of the higher-order IPCA factors, though close to zero unconditionally, are nonzero conditionally. Thus, the higher-order IPCA factors enter the conditional SDF, and a parsimonious model that only includes the lower-order IPCA factors is inadequate to fully capture risks that are priced conditionally.[25]

The performance of low-frequency factors is qualitatively unchanged if I estimate these factors from a small number of IPCA factors. Table 2 reports the in-sample (unconditional) monthly Sharpe ratio of the optimal portfolio of a small number of IPCA factors and that of the low-frequency factors estimated from a small number of IPCA factors. Panels A and B correspond to the Basic-IPCA and Freq-IPCA frameworks, respectively. In each panel, the first row reports the Sharpe ratio of the optimal portfolio of K IPCA factors, with K ranging between three and twelve. The remaining four rows report the Sharpe ratio of the optimal portfolio of (i) one below-business-cycle-frequency factor, (ii) one below-business-cycle-frequency factor and one business-cycle-frequency factor, (iii) one below-business-cycle-frequency factor and two business-cycle-frequency factors, and (iv) two below-business-cycle-frequency factors and one business-cycle-frequency factor, respectively. In each of these four rows, the low-frequency factors are estimated from the corresponding K IPCA factors.

There are two main discoveries from this table. First, under either the Basic-IPCA or Freq-IPCA framework, starting from five IPCA factors, the below-business-cycle-frequency factor delivers a high Sharpe ratio comparable to that of the optimal portfolio of the K IPCA factors. For instance, the Sharpe ratio of the optimal portfolio of nine Freq-IPCA factors is 1.455. In comparison, *one* below-business-cycle-frequency factor estimated from the nine Freq-IPCA factors has a Sharpe ratio

[24] In general, the two panels of Figure 3 reflect some of the fundamental insights of macro-finance models, while also poses additional challenges to these models. For instance, in the long-run risk model of Bansal and Yaron (2004), the risk premia of both the short-run and long-run consumption risk increase when the conditional volatility of consumption growth is high, suggesting that the conditional volatility of consumption growth is a "level" state variable that drives the comovement of the risk premia of the two different sources of risk in the same direction. However, few structural models consider that the risk premia of different sources of risk could react to a state variable from different directions. Moreover, the negative correlation between the below-business-cycle-frequency factor and the market factor implies that the market factor has a negative loading on the latent "level" state variable. Thus, when the state variable drives up the conditional risk premia of most candidate risk factors, it drives down the conditional market premium. This poses a challenge to most structural asset-pricing models.

[25] By imposing the restriction that the SDF is not too volatile, Haddad et al. (2020) show that factors that do not explain a substantial amount of time-series variation of the returns are not highly predictable. However, my results show that, empirically, some low-frequency state variables drive the risk premia of these factors.

Table 2: Sharpe ratio of IPCA factors and low-frequency factors

This table reports the in-sample (unconditional) monthly Sharpe ratio of the optimal portfolio of a small number of IPCA factors and that of the low-frequency factors estimated from a small number of IPCA factors. Panels A and B correspond to the Basic-IPCA and Freq-IPCA frameworks, respectively. In each panel, the first row reports the Sharpe ratio of the optimal portfolio of K IPCA factors with K ranging between three and twelve. The remaining four rows report the Sharpe ratio of the optimal portfolio of (i) one below-business-cycle-frequency factor, (ii) one below-business-cycle-frequency factor and one business-cycle-frequency factor, (iii) one below-business-cycle-frequency factor and two business-cycle-frequency factors, and (iv) two below-business-cycle-frequency factors and one business-cycle-frequency factor, respectively. In each of these four rows, the low-frequency factors are estimated from the corresponding K IPCA factors.

Factor type	Number of IPCA factors									
	3	4	5	6	7	8	9	10	11	12
Panel A: Basic-IPCA										
IPCA factors	0.315	0.354	1.263	1.277	1.274	1.379	1.419	1.431	1.447	1.447
B-BCF1	0.199	0.244	1.190	1.203	1.175	1.278	1.312	1.297	1.302	1.310
B-BCF1+BCF1	0.298	0.302	1.195	1.210	1.177	1.288	1.334	1.316	1.315	1.322
B-BCF1+BCF2	0.315	0.344	1.253	1.249	1.187	1.291	1.340	1.324	1.315	1.322
B-BCF2+BCF1	0.315	0.349	1.255	1.220	1.183	1.290	1.339	1.329	1.319	1.325
Panel B: Freq-IPCA										
IPCA factors	0.380	0.417	1.310	1.315	1.305	1.427	1.455	1.461	1.465	1.481
B-BCF1	0.283	0.277	1.216	1.210	1.177	1.301	1.328	1.326	1.319	1.324
B-BCF1+BCF1	0.372	0.381	1.235	1.213	1.179	1.322	1.345	1.348	1.338	1.341
B-BCF1+BCF2	0.380	0.404	1.297	1.227	1.201	1.323	1.346	1.360	1.349	1.344
B-BCF2+BCF1	0.380	0.407	1.297	1.225	1.198	1.324	1.347	1.357	1.350	1.343

of 1.328, which is only 8.76% lower. Thus, from an investment perspective, the performance of the below-business-cycle-frequency factor is close to that of the optimal portfolio of the IPCA factors. Including additional business-cycle-frequency factors or below-business-cycle-frequency factors on top of the first below-business-cycle-frequency factor only increases the Sharpe ratio marginally. Therefore, the conclusion that a large fraction of the risk premia of the cross-section is associated with the time-series variation driven by a below-business-cycle-frequency state variable still holds, even when I restrict the cross-section to include only the lower-order IPCA factors.

Second, in almost all cases, the IPCA and low-frequency factors estimated in the Freq-IPCA framework deliver higher Sharpe ratios than those estimated in the Basic-IPCA framework. In particular, although the Freq-IPCA framework estimates more parameters than the Basic-IPCA framework, and its performance improvement relative to the Basic-IPCA framework looks modest, it is important to note that the Freq-IPCA framework does not mechanically output models with higher Sharpe ratios than the Basic-IPCA models. Moreover, unreported GRS test results suggest that the Freq-IPCA factors capture risk premia unexplained by the Basic-IPCA factors.[26] Therefore, consistent with the findings of **?**, allowing the low- and high-frequency components of characteristics to instrument the stock factor loadings differently helps the IPCA factors to better capture the systematic risk premia, which are also captured by the below-business-cycle-frequency factor identified from the Freq-IPCA factors.

[26]Specifically, in unreported results, I regress the K Freq-IPCA factors on the K Basic-IPCA factors with K ranging between three and twelve. For these regressions, the GRS test strongly rejects the null hypothesis with a p-value lower than 0.001 for every K. In addition, I regress the K Basic-IPCA factors on the K Freq-IPCA factors. For these regressions, the GRS test yields a p-value greater than 0.025 for $K \geq 8$. Specifically, for $K = 8$, 9, and 12, the corresponding p-values are 0.073, 0.116, and 0.176, respectively.

3.2 Explaining returns of individual stocks and characteristic-managed portfolios

In this section, I discuss how the low-frequency factors explain the returns of individual stocks and managed portfolios of firm-specific characteristics. In particular, the previous subsection shows that a large fraction of the unconditional risk premia of the cross-section is associated with the time-series variation driven by a below-business-cycle-frequency state variable. Thus, the low-frequency factors, which, in this case, capture the large fraction of the risk premia, should have a satisfactory explanatory power to the expected returns of various cross-sections, such as individual stocks and characteristic-managed portfolios. In other words, the approached proposed in this manuscript to identify low-frequency risk factors should be viewed as a satisfactory method to address the "factor zoo" and identify a parsimonious model that prices the cross-section, if most of the risk premia of the cross-section are truly associated with low-frequency time-series variation.

Similar to Kelly et al. (2019), I use two R-squared statistics to measure the performance of a factor model.[27] The first one is the total R-squared:

$$\text{Total } R^2 = 1 - \frac{\sum_{i,t}\left(r_{i,t+1} - z_{i,t}^{\top}\Gamma_{\beta} f_{t+1}\right)^2}{\sum_{i,t} r_{i,t+1}^2}, \tag{39}$$

where $r_{i,t+1}$ is the return of stock i at month $t+1$, f_{t+1} contains the returns of the risk factors (either IPCA or low-frequency factors) at month $t+1$, $z_{i,t}$ contains the characteristics of stock i at month t (the ith row of Z_t defined in (31)), and Γ_{β} is the mapping matrix that transforms the characteristics to the factor loadings of stocks. In particular, as explained in Section 2.4.1, the mapping matrix of a model that includes the low-frequency factors and/or other pre-specified factors is obtained via regressing the stock returns on the interaction of the factors included in the model and characteristics. The total R-squared measures the fraction of return variation explained by a model's factor realizations and time-varying factor loadings (instrumented by characteristics), aggregated over all stocks and all months.

The second R-squared statistic is the predictive R-squared with constant factor risk premia:

$$\text{Predictive } R^2 \text{ with constant factor RP} = 1 - \frac{\sum_{i,t}\left(r_{i,t+1} - z_{i,t}^{\top}\Gamma_{\beta}\lambda\right)^2}{\sum_{i,t} r_{i,t+1}^2}, \tag{40}$$

where λ contains the unconditional premia (average returns) of the risk factors. The predictive R-squared measures the return variation explained by a model's description of the risk premia of stocks, assuming that the risk premia of the factors are constant. Thus, the predictive R-squared measures how well the model explains the stock expected returns.

Table 3 reports the total R-squared (in percent) for explaining the monthly individual stock returns of the IPCA models and models that include the low-frequency factors. Panels A and B report the results for the Basic-IPCA and Freq-IPCA frameworks, respectively. In each panel, the first row reports the total R-squared of the K-factor IPCA model with K ranging between three and twelve, and the second row of the model that includes one below-business-cycle-frequency factor and one business-cycle-frequency factor estimated from K IPCA factors. The last row reports the total R-squared of the model that includes the two low-frequency factors estimated from K IPCA factors and one factor

[27] To simplify the exposition, I define the two R-squared statistics for explaining individual stock returns. The two statistics can be defined similarly for explaining the returns of managed portfolios of firm characteristics.

Table 3: Total R-squared for explaining monthly individual stock returns

This table reports the total R-squared (in percent) for explaining the monthly individual stock returns of the IPCA models and models that include the low-frequency factors. Panels A and B report the results for the Basic-IPCA and Freq-IPCA frameworks, respectively. In each panel, the first row reports the R-squared of the K-factor IPCA model with K ranging between three and twelve, and the second row of the model that includes one below-business-cycle-frequency factor and one business-cycle-frequency factor estimated from K IPCA factors. The last row reports the R-squared of the model that includes the two low-frequency factors estimated from K IPCA factors and one factor that mimics the equally-weighted market factor. In the second to last column, the low-frequency factors are estimated from all available IPCA factors, that is, from all 37 Basic-IPCA factors or all 73 Freq-IPCA factors. The last column reports the R-squared of the model that only includes a factor that mimics the equally-weighted market factor.

Factor type	Number of IPCA factors											Market
	3	4	5	6	7	8	9	10	11	12	All	
Panel A: Basic-IPCA												
IPCA factors	17.42	18.05	18.56	18.88	19.12	19.32	19.45	19.57	19.69	19.79	-	
B-BCF1+BCF1	14.91	5.26	3.89	3.73	1.89	3.86	3.06	3.18	2.82	2.84	1.95	14.59
B-BCF1+BCF1+Market	17.38	16.55	16.54	16.53	15.79	16.19	16.18	16.19	16.02	16.02	15.63	
Panel B: Freq-IPCA												
IPCA factors	17.66	18.27	18.82	19.15	19.42	19.63	19.79	19.93	20.05	20.17	-	
B-BCF1+BCF1	14.71	5.24	3.40	3.38	1.85	3.64	3.03	2.99	2.84	2.91	1.76	14.66
B-BCF1+BCF1+Market	17.62	16.67	16.74	16.48	16.03	16.35	16.29	16.30	16.27	16.22	15.67	

that mimics the equally-weighted market factor.[28] In the second to last column, the low-frequency factors are estimated from all available IPCA factors, that is, from all 37 Basic-IPCA factors or all 73 Freq-IPCA factors. The last column reports the R-squared of the model that only includes a factor that mimics the equally-weighted market factor.

Table 3 shows that the IPCA models explain a large amount of the total variation of individual stock returns. Specifically, the Basic-IPCA models explain 17.42% to 19.79% of the total variation of individual stock returns when the number of IPCA factors ranges between three and twelve. The corresponding numbers for the Freq-IPCA models are even higher and range between 17.66% to 20.17%. As a comparison, Kelly et al. (2019, Table 2) show that when firm characteristics are used as the instruments for stock factor loadings, the total R-squared of the three-factor model of Fama and French (1993, FF3) that includes the market, size, and value factors is 14.2%, and it is 15.3% for the four-factor model of Carhart (1997, FFC4) that adds the momentum factor on top of FF3. For the five-factor model of Fama and French (2015, FF5) that adds the investment and profitability factors on top of FF3, the total R-squared is 14.7%, and it is 15.6% for the six-factor model of Fama and French (2018, FF6) that adds the momentum factor on top of FF5. All of these are lower than those of the Basic-IPCA and Freq-IPCA models.

The below-business-cycle- and business-cycle-frequency factors do not explain a large amount of the total variation of individual stock returns. Specifically, starting from five Basic-IPCA factors, the two low-frequency factors estimated from the Basic-IPCA factors explain less than 4% of the total variation of individual stock returns, and the two low-frequency factors estimated from all 37 Basic-

[28] The factor that mimics the equally-weighted portfolio is constructed by going long one unit of every available stock. Thus, the return of this factor at month $t+1$ is

$$X^M_{t+1} = \frac{1}{N_{t+1}} \mathbf{1}^\top r_{t+1},$$

where N_{t+1} is the number of non-missing observations of the assets at month $t+1$.

Table 4: Predictive R-squared with constant factor risk premia for explaining individual stock returns

This table reports the predictive R-squared with constant factor risk premia (in percent) for explaining the monthly individual stock returns of the IPCA models and models that include the low-frequency factors. Panels A and B report the results for the Basic-IPCA and Freq-IPCA frameworks, respectively. In each panel, the first row reports the R-squared of the K-factor IPCA model with K ranging between three and twelve, and the second row of the model that includes one below-business-cycle-frequency factor and one business-cycle-frequency factor estimated from K IPCA factors. The last row reports the R-squared of the model that includes the two low-frequency factors estimated from K IPCA factors and one factor that mimics the equally-weighted market factor. In the second to last column, the low-frequency factors are estimated from all available IPCA factors, that is, from all 37 Basic-IPCA factors or all 73 Freq-IPCA factors. The last column reports the R-squared of the model that only includes a factor that mimics the equally-weighted market factor.

Factor type	Number of IPCA factors											Market
	3	4	5	6	7	8	9	10	11	12	All	
Panel A: Basic-IPCA												
IPCA factors	0.41	0.42	0.69	0.68	0.68	0.70	0.70	0.71	0.70	0.70	-	
B-BCF1+BCF1	0.40	0.15	0.53	0.53	0.58	0.57	0.55	0.55	0.59	0.58	0.60	0.35
B-BCF1+BCF1+Market	0.42	0.37	0.69	0.69	0.68	0.71	0.72	0.72	0.71	0.72	0.71	
Panel B: Freq-IPCA												
IPCA factors	0.46	0.45	0.76	0.76	0.75	0.78	0.78	0.78	0.78	0.78	-	
B-BCF1+BCF1	0.45	0.21	0.62	0.63	0.70	0.62	0.65	0.66	0.67	0.68	0.75	0.35
B-BCF1+BCF1+Market	0.46	0.39	0.78	0.77	0.76	0.80	0.80	0.80	0.81	0.80	0.79	

IPCA factors only explain 1.95% of the total variation. The low-frequency factors estimated from the Freq-IPCA factors yield similar results, and the total R-squared of models that include these factors is, in most cases, slightly lower than that of models that include the low-frequency factors estimated from the Basic-IPCA factors. Thus, two low-frequency factors cannot explain the time-series variation of individual stock returns as much as the IPCA factors. This is not surprising because the low-frequency factors have low and negative correlations with the market factor, which is one of the most important asset-pricing factors that explains a substantial amount of the time variation of stock returns: the equally-weighted market factor explains 14.59% and 14.66% of the total variation of individual stock returns under the Basic-IPCA and Freq-IPCA frameworks, respectively.

By augmenting the low-frequency factors with the market factor, Table 3 shows that the total R-squared of the three-factor models is comparable to that of the higher-dimensional IPCA models. For instance, a three-factor model that includes the market factor and two low-frequency factors estimated from nine Basic-IPCA factors delivers a total R-squared of 16.18%, which is comparable with that of the nine-factor Basic-IPCA model (19.45%), and is higher than that of the prominent FF3, FFC4, FF5, and FF6 models. In summary, the low-frequency factors do not explain a substantial amount of the total variation of individual stock returns, mainly because they have low correlations with the market factor. By augmenting the low-frequency factors with the market factor, the total R-squared of the three-factor models is comparable to that of the IPCA models.

Importantly, although the low-frequency factors do not explain a substantial amount of the total variation of individual stock returns, they successfully capture the risk premia of individual stocks. Table 4 reports the predictive R-squared with constant factor risk premia (in percent) for explaining the monthly individual stock returns of the IPCA models and models that include the low-frequency factors. Panels A and B report the results for the Basic-IPCA and Freq-IPCA frameworks, respectively. There are two main findings from this table. First, the predictive R-squared of models that include two low-frequency factors is positive in all cases and is comparable to that of the higher-dimensional

Table 5: Total R-squared for explaining monthly returns of 36 characteristic-managed portfolios

This table reports the total R-squared (in percent) for explaining the monthly returns of portfolios managed by the 36 considered characteristics of the IPCA models and models that include the low-frequency factors. Panels A and B report the results for the Basic-IPCA and Freq-IPCA frameworks, respectively. In each panel, the first row reports the R-squared of the K-factor IPCA model with K ranging between three and twelve, and the second row of the model that includes one below-business-cycle-frequency factor and one business-cycle-frequency factor estimated from K IPCA factors. The last row reports the R-squared of the model that includes the two low-frequency factors estimated from K IPCA factors and one factor that mimics the equally-weighted market factor. In the second to last column, the low-frequency factors are estimated from all available IPCA factors, that is, from all 37 Basic-IPCA factors or all 73 Freq-IPCA factors. The last column reports the R-squared of the model that only includes a factor that mimics the equally-weighted market factor.

Factor type	Number of IPCA factors											Market
	3	4	5	6	7	8	9	10	11	12	All	
Panel A: Basic-IPCA												
IPCA factors	81.06	86.78	89.25	91.69	94.17	95.16	95.84	96.41	96.84	97.11	-	
B-BCF1+BCF1	62.58	34.02	35.86	34.81	19.76	31.84	28.20	28.90	25.56	25.69	16.25	34.62
B-BCF1+BCF1+Market	79.66	58.04	56.49	55.90	46.73	51.99	50.92	51.15	49.17	49.23	43.26	
Panel B: Freq-IPCA												
IPCA factors	80.99	86.95	89.33	92.39	94.17	94.90	95.83	96.33	96.52	96.79	-	
B-BCF1+BCF1	64.23	32.06	33.94	30.60	21.35	30.54	26.66	26.41	25.68	25.21	14.07	35.13
B-BCF1+BCF1+Market	79.61	58.00	56.85	52.99	48.42	51.76	50.44	50.74	50.45	49.87	43.02	

IPCA models. For instance, the predictive R-squared of the nine-factor Basic-IPCA model is 0.70%, while that of the model that includes the two low-frequency factors estimated from the nine Basic-IPCA factors is 0.55%. By augmenting these two low-frequency factors with a market factor, whose predictive R-squared is 0.35%, the predictive R-squared of the three-factor model becomes 0.72%, which is even higher than that of the higher-dimensional nine-factor Basic-IPCA model. Thus, a three-factor model that includes the market factor and two low-frequency risk factors subsumes a higher-dimensional IPCA model's explanatory power for the average return of individual stocks, and the low-frequency factors explain the risk premia of individual stocks uncaptured by the market factor.

Second, models with factors estimated based on the Freq-IPCA framework consistently deliver a higher predictive R-squared than those with factors estimated based on the Basic-IPCA framework. For instance, the predictive R-squared of the nine-factor Basic-IPCA model is 0.70%, while that of the nine-factor Freq-IPCA model is 0.78%. Similarly, the predictive R-squared of the model that includes the two low-frequency factors estimated from nine Basic-IPCA factors is 0.55%, while that of the model that includes the low-frequency factors estimated from nine Freq-IPCA factors is 0.65%. Considering that the variation of risk premia only constitutes a small fraction of the total variation of individual stock returns, such improvement in the predictive R-squared with a magnitude of around 0.1% is non-negligible, especially when the Freq-IPCA framework does not mechanically guarantee a higher predictive R-squared.[29] Thus, this suggests that the Freq-IPCA framework produces a better measure of the conditional loadings of individual stocks on systematic risks, and the low- and high-frequency components of firm characteristics measure the stocks' quantity of risks differently.

Table 5 reports the total R-squared for explaining the monthly returns of portfolios managed by the 36 considered characteristics of the IPCA models and models that include the low-frequency

[29]The objective of the IPCA model is to explain the *total* variation of asset returns (see equation (5) of Kelly et al., 2019). Thus, the Freq-IPCA model mechanically explains a higher fraction of the total variation of asset returns than the Basic-IPCA model. However, this does not guarantee that the Freq-IPCA model delivers a higher predictive R-squared than the Basic-IPCA model.

factors.[30] Similarly, Table 6 reports the predictive R-squared with constant factor risk premia for explaining these managed portfolio returns. In each table, Panels A and B report the results for the Basic-IPCA and Freq-IPCA frameworks, respectively, and the number of IPCA factors ranges between three and twelve. Since the long-short portfolios of the considered characteristics are known as asset-pricing "anomalies" that capture risk premia unexplained by benchmark factor models such as CAPM, examining whether the models that include low-frequency factors explain the returns of these anomaly portfolios sheds light on whether the low-frequency factors capture well the risk premia of the cross-section that benchmark models find hard to explain.

The results in Tables 5 and 6 are consistent with the previous findings in explaining monthly individual stock returns. Specifically, Table 5 shows that the IPCA models explain a large amount of the total variation of managed portfolio returns, while the below-business-cycle- and business-cycle-frequency factors have limited explanatory power for the total variation of managed portfolio returns. For instance, the total R-squared of the nine-factor Basic-IPCA model is 95.84%, while that of the two low-frequency factors estimated from the nine Basic-IPCA factors is only 28.20%. Augmenting the two low-frequency factors with the equally-weighted market factor increases the total R-squared to 50.92%, which is still substantially lower than the nine-factor Basic-IPCA model. This is unsurprising because models that include the low-frequency factors are parsimonious two- or three-factor models (depending on whether the low-frequency factors are augmented with a market factor), while the IPCA models are higher-dimensional. Moreover, IPCA can be approximately viewed as applying PCA to the managed portfolios of characteristics, suggesting that the IPCA factors are approximately designed to explain the total variation of managed portfolio returns. In contrast, the low-frequency factors are constructed as those whose risk premia are driven by low-frequency state variables, and the market factor is cross-sectionally orthogonal to the long-short managed portfolios. This justifies the low-frequency and market factors not explaining the total variation of managed portfolio returns as well as the IPCA factors.

However, the low-frequency factors capture well the risk premia of the managed portfolios. Specifically, Table 6 shows that the predictive R-squared of models that include the two low-frequency factors is even higher than that of the higher-dimensional IPCA models. For instance, the predictive R-squared of the nine-factor Basic-IPCA model is 3.23%, while that of the two low-frequency factors estimated from the nine Basic-IPCA factors is 3.79%. Such performance improvement holds for five or more IPCA factors and for both the Basic-IPCA and Freq-IPCA frameworks. Thus, a low-dimensional model that includes two low-frequency factors explains the risk premia of the managed portfolios better than a higher-dimensional IPCA model.

Moreover, models with factors estimated based on the Freq-IPCA framework consistently deliver a higher predictive R-squared for explaining the returns of managed portfolios than those with factors estimated based on the Basic-IPCA framework. For instance, the predictive R-squared of the nine-factor Freq-IPCA model is 3.58%, which is higher than that of its Basic-IPCA counterpart (3.23%). Similarly, the predictive R-squared of the model that includes the two low-frequency factors estimated from nine Freq-IPCA factors is 4.54%, which is also higher than that of its Basic-IPCA counterpart (3.79%). Thus, the Freq-IPCA framework empowers the two low-frequency factors with a 0.75% improvement in the

[30] At month $t+1$, the returns of the characteristic-managed portfolios are calculated as

$$X^C_{t+1} = \frac{1}{N_{t+1}} Z_t^\top r_{t+1},$$

where N_{t+1} is the number of non-missing stocks at month $t+1$.

Table 6: Predictive R-squared with constant factor risk premia for explaining managed portfolio returns

This table reports the predictive R-squared with constant factor risk premia (in percent) for explaining the monthly returns of portfolios managed by the 36 considered characteristics of the IPCA models and models that include the low-frequency factors. Panels A and B report the results for the Basic-IPCA and Freq-IPCA frameworks, respectively. In each panel, the first row reports the R-squared of the K-factor IPCA model with K ranging between three and twelve, and the second row of the model that includes one below-business-cycle-frequency factor and one business-cycle-frequency factor estimated from K IPCA factors. The last row reports the R-squared of the model that includes the two low-frequency factors estimated from K IPCA factors and one factor that mimics the equally-weighted market factor. In the second to last column, the low-frequency factors are estimated from all available IPCA factors, that is, from all 37 Basic-IPCA or all 73 Freq-IPCA factors. The last column reports the R-squared of the model that only includes a factor that mimics the equally-weighted market factor.

Factor type	Number of IPCA factors											Market
	3	4	5	6	7	8	9	10	11	12	All	
Panel A: Basic-IPCA												
IPCA factors	1.04	1.21	3.10	2.97	2.94	3.25	3.23	3.24	3.19	3.18	-	
B-BCF1+BCF1	1.03	< 0	3.32	3.30	3.52	3.81	3.79	3.80	3.90	3.88	3.91	0.44
B-BCF1+BCF1+Market	1.03	0.79	3.19	3.17	3.13	3.57	3.62	3.63	3.59	3.59	3.51	
Panel B: Freq-IPCA												
IPCA factors	1.30	1.41	3.56	3.43	3.36	3.66	3.58	3.61	3.59	3.54	-	
B-BCF1+BCF1	1.35	0.11	4.02	3.96	4.39	4.49	4.54	4.50	4.56	4.56	4.79	0.41
B-BCF1+BCF1+Market	1.30	0.99	3.97	3.85	3.92	4.30	4.29	4.27	4.29	4.25	4.27	

predictive R-squared for explaining the managed portfolio returns. Consistent with the previous findings in explaining individual stock returns, this suggests that the low- and high-frequency components of firm characteristics measure the stocks' quantity of risks differently, and the Freq-IPCA framework produces a better measure of the conditional loadings on systematic risks.

To further investigate how the low- and high-frequency components of firm characteristics instrument stock factor loadings, Figure 4 illustrates the IPCA mapping matrices that transform the characteristics to the stock loadings on the two low-frequency factors estimated from all 37 Basic-IPCA factors.[31] Panels A and B illustrate the elements of the matrices that map the characteristics to the loadings on the below-business-cycle- and business-cycle-frequency factors, respectively. In each panel, the blue bars graph the contributions of the present values of characteristics to factor loadings in the Basic-IPCA framework, and the orange and green bars graph the contributions of the low- and high-frequency components of characteristics to factor loadings in the Freq-IPCA framework, respectively. Panel A shows that, in the Basic-IPCA framework, size (`lme`) is negatively correlated with the loadings on the below-business-cycle-frequency factor. This justifies the size effect because the below-business-cycle-frequency factor earns a high unconditional risk premium. Specifically, small firms have higher loadings on the highly-compensated below-business-cycle-frequency factor, and thus, on average, they have higher expected returns than big firms.

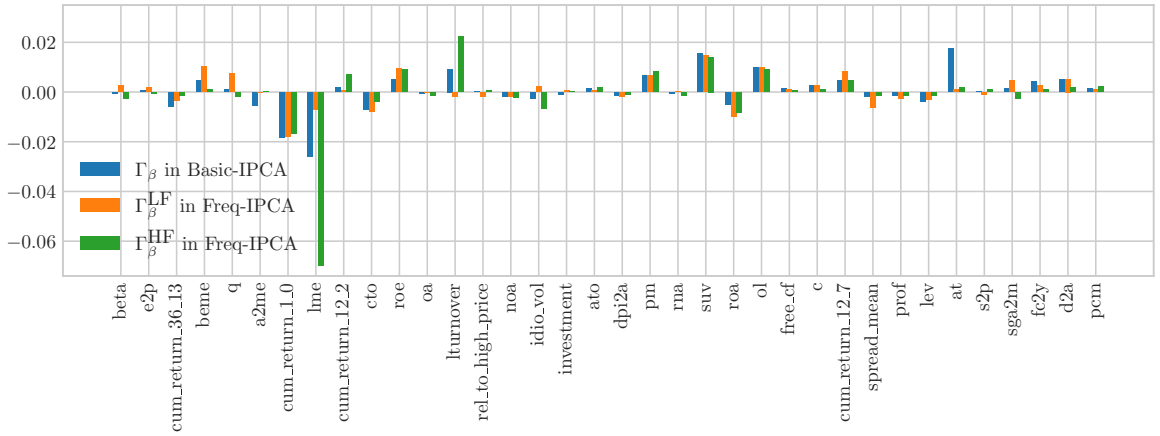
However, both Panels A and B show that the low- and high-frequency components of characteristics have different contributions to the loadings on the low-frequency factors. For instance, Panel A shows that the low- and high-frequency components of size have substantially different contributions to the loadings on the below-business-cycle-frequency factor. Specifically, the high-frequency component of size has a much stronger negative correlation with the loadings on the below-business-cycle-frequency

[31] The results are qualitatively unchanged when the two low-frequency factors are estimated from other choices of candidate risk factors estimated based on the IPCA framework.

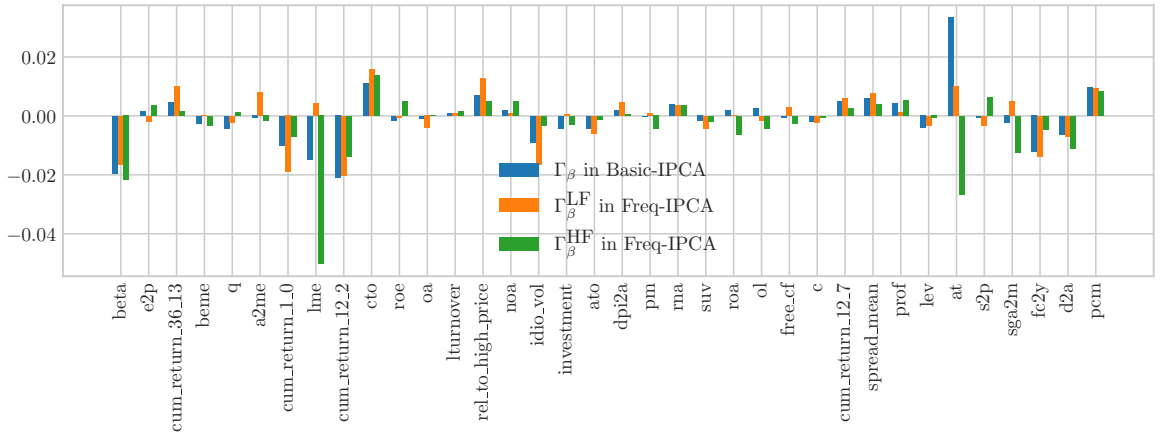
Figure 4: Mapping from characteristics to factor loadings on low-frequency factors

This figure illustrates the IPCA mapping matrices that transform the characteristics to the stock loadings on the two low-frequency factors estimated from the 37 Basic-IPCA factors. Panels A and B illustrate the elements of the matrices that map the characteristics to the loadings on the below-business-cycle- and business-cycle-frequency factors, respectively. In each panel, the blue bars graph the contributions of the present values of characteristics to factor loadings in the Basic-IPCA framework, and the orange and green bars graph the contributions of the low- and high-frequency components of characteristics to factor loadings in the Freq-IPCA framework, respectively.

(A) Below-business-cycle-frequency factor



(B) Business-cycle-frequency factor



factor, and thus it is mainly the firms whose current size is small compared to their three-year average that have high loadings on this highly-compensated factor. In contrast, the low-frequency component of the book-to-market ratio (`beme`) contributes more to the loadings on the below-business-cycle-frequency factor than the high-frequency component, suggesting that firms with a persistently high book-to-market ratio capture a higher return than firms for which the same book-to-market ratio is more transitory.[32] In summary, Figure 4 confirms that the low- and high-frequency components of characteristics instrument the stock loadings on the low-frequency factors differently.

In this section, I show that models that include the two low-frequency factors outperform higher-dimensional IPCA models in explaining the average returns of the individual stocks and characteristic-managed portfolios. As a result, the method proposed in this paper serves as a method to address the well-known "factor-zoo" problem to identify a parsimonious factor model that explains the cross-section

[32]This is consistent with the findings of **?**.

of expected returns. However, I highlight that the fundamental reason of the satisfactory performance of the low-frequency factors is that a large fraction of the unconditional risk premia of the cross-section is associated with low-frequency time-series variation empirically. By construction, the method cannot capture the information of factors with positive and constant risk premia.[33] In particular, if all candidate risk factors have constant risk premia, then the method proposed does not guarantee to identify factors that earn high risk premia. Therefore, in general, my method can be used jointly with other popular methods that focus on unconditional factor risk premia to address the "factor zoo" problem, which is explained in more detail in Appendix B.

3.3 Predicting low-frequency factors

Section 2.3 shows that latent low-frequency state variables drive the time variation of the risk premia of low-frequency factors, and the low-frequency factors aggregate the explanatory power of these state variables for systematic risk premia. In this section, I validate this conclusion empirically by showing that the one-month-ahead returns of the low-frequency factors are *highly predictable* by the 19 macroeconomic variables and 24 financial indicators described in Section 2.4.3 via the partial least squares (PLS) regression.

I start by presenting the in-sample prediction results to understand the predictability of the low-frequency factors economically. I then present the results from a pseudo-out-of-sample analysis, in which I estimate the low-frequency factors from the candidate IPCA factors using all available return data but predict the returns of the low-frequency factors at month $t+1$ recursively (via an expanding-window estimation) using the information of the economic predictor variables only up to month t. The advantage of using the full sample of return data to estimate the low-frequency factors is that it provides smoothed and efficient estimates of the low-frequency factors. In particular, to show that the low-frequency factors are predictable because they successfully aggregate the latent low-frequency state variables' explanatory power for systematic risk premia, such an advantage is important because the variation of risk premia that the low-frequency factors capture has long cycles. Estimating the low-frequency factors using the full sample of return data enables me to efficiently identify the slow-moving fluctuations in risk premia, thus facilitating the investigation of whether the economic predictor variables drive the low-frequency variation of systematic risk premia.[34] Lastly, to understand whether the returns of the low-frequency factors can be predicted in real-time, I present the results from an out-of-sample analysis, in which I use an expanding-window procedure where both the estimation and prediction of the low-frequency factors are formed without look-ahead bias.

3.3.1 In-sample prediction

Table 7 reports the in-sample prediction results for the monthly returns of the two low-frequency risk factors estimated from all 37 Basic-IPCA factors. Panels A and B present the results for the below-business-cycle- and business-cycle-frequency factors, respectively. In each panel, the first column reports the predictive coefficient of the PLS factor(s) and R-squared obtained from the third-stage PLS

[33] The existence of such factors is supported by the negative correlations between the low-frequency factors and the market factor, which earns positive risk premia.

[34] In essence, the pseudo-out-of-sample analysis addresses the question: If the data-generating process of the low-frequency factors is known, do the selected economic predictor variables predict the returns of these factors in real time?

Table 7: In-sample prediction results for low-frequency factors

This table reports the in-sample prediction results for the monthly returns of the two low-frequency risk factors estimated from all 37 Basic-IPCA factors. Panels A and B present the results for the below-business-cycle- and business-cycle-frequency factors, respectively. In each panel, the first column reports the predictive coefficient of the PLS factor(s) and R-squared obtained from the third-stage PLS regression, and the second column the statistics of the predictive coefficient of the PLS factor(s). The remaining six columns report the prediction results when using the one- and two-factor PLS prediction models and when using the 19 macroeconomic variables, 24 financial variables, and all 43 variables as predictors. For each PLS factor's predictive coefficient, the first row reports its estimated value, the second row its Newey and West (1987) t-statistic, and the third row its p-value of the t-statistic based on the asymptotic distribution of Kelly and Pruitt (2015). The fourth and fifth rows report the bias and p-value of the coefficient obtained from a parametric bootstrap procedure.

(A) Below-business-cycle-frequency factor

Parameter		Macro		Financial		Macro+financial	
		PLS1	PLS2	PLS1	PLS2	PLS1	PLS2
$\widehat{\phi}^{\text{PLS1}}$	Coeff.	0.36	0.36	0.41	0.41	0.52	0.52
		(3.82)	(4.65)	(6.33)	(6.03)	(8.47)	(11.14)
	p-value (KP)	.00	.00	.00	.00	.00	.00
	Bias (bootstrap)	0.01	0.01	0.02	0.05	0.02	0.02
	p-value (bootstrap)	.00	.00	.00	.00	.00	.00
$\widehat{\phi}^{\text{PLS2}}$	Coeff.	-	0.30	-	0.22	-	0.28
			(4.52)		(2.77)		(5.97)
	p-value (KP)	-	.00	-	.00	-	.00
	Bias (bootstrap)	-	-0.00	-	0.05	-	0.00
	p-value (bootstrap)	-	.00	-	.02	-	.00
R^2 (%)		12.76	21.67	17.55	22.40	27.53	35.42

(B) Business-cycle-frequency factor

Parameter		Macro		Financial		Macro+financial	
		PLS1	PLS2	PLS1	PLS2	PLS1	PLS2
$\widehat{\phi}^{\text{PLS1}}$	Coeff.	0.22	0.22	0.42	0.42	0.32	0.32
		(2.34)	(2.78)	(7.56)	(7.08)	(3.77)	(4.92)
	p-value (KP)	.00	.00	.00	.00	.00	.00
	Bias (bootstrap)	0.00	0.00	0.02	0.02	-0.00	-0.00
	p-value (bootstrap)	.00	.00	.00	.00	.00	.00
$\widehat{\phi}^{\text{PLS2}}$	Coeff.	-	0.16	-	0.09	-	0.23
			(2.26)		(1.33)		(4.10)
	p-value (KP)	-	.00	-	.05	-	.00
	Bias (bootstrap)	-	0.00	-	0.01	-	-0.01
	p-value (bootstrap)	-	.00	-	.13	-	.00
R^2 (%)		4.86	7.45	17.14	17.89	9.93	15.05

regression, and the second column the statistics of the predictive coefficient of the PLS factor(s). The remaining six columns report the prediction results when using the one- and two-factor PLS prediction models and when using the 19 macroeconomic variables, 24 financial variables, and all 43 variables as predictors. For each PLS factor's predictive coefficient, the first row reports its estimated value, the second row its Newey and West (1987) t-statistic, and the third row its p-value of the t-statistic based on the asymptotic distribution of Kelly and Pruitt (2015).[35]

The fourth and fifth rows report the bias and p-value of the coefficient obtained from a parametric bootstrap procedure. Specifically, I follow Haddad et al. (2020) and estimate a restricted VAR(1) system for the low-frequency factor and PLS factor(s) under the null of no predictability. I then simulate 10,000 histories with the same length as my sample from the VAR system with errors drawn with replacement from the empirical distribution. For each simulated history, I regress the low-frequency factor on the PLS factor(s). By aggregating the regression results from the 10,000 histories, I obtain the empirical bias and distribution of the coefficient of the PLS factor(s) under the null of no predictability, which enables me to construct the bootstrap p-values. In particular, this bootstrap procedure takes the persistence of the low-frequency risk factors and PLS prediction factors into account, and it also considers the correlation between the innovation of the low-frequency factors and that of the PLS factors. Thus, this procedure alleviates the issue of the potential small-sample bias of the predictive coefficient of the PLS factors (Stambaugh, 1999) and the spurious-regression issue related to the standard error of the predictive coefficient (Ferson, Sarkissian, and Simin, 2003).

[35]The Newey and West (1987) t-statistic is constructed using a 12-month window with the Bartlett kernel. The conclusion is unaltered with other band-width choices.

Table 7 shows that both the below-business-cycle- and business-cycle-frequency factors estimated from all 37 Basic-IPCA factors are highly predictable by the selected economic variables. Specifically, the R-squared of the third stage PLS regression is 12.76%, 17.55%, and 27.53% for predicting the below-business-cycle-frequency factor returns using one PLS factor of the 19 macroeconomic variables, 24 financial variables, and all 43 variables, respectively. Moreover, the tests based on the Newey and West (1987) t-statistic, the asymptotic distribution of Kelly and Pruitt (2015), and the bootstrap procedure suggest that the predictive power of each set of these economic variables is statistically significant. The aforementioned R-squared further increases to 21.67%, 22.40%, and 35.42% when I include two PLS factors in the prediction model, and the additional PLS factor is also statistically significant at the 5% level based on all considered tests. Considering that the prediction objective is the one-month-ahead return of a risk factor, the R-squared for predicting the below-business-cycle-frequency factor returns is large and is substantially higher than that for predicting the returns of a conventional asset-pricing factor, such as the market.[36]

The business-cycle-frequency factor is less predictable. Specifically, compared to the below-business-cycle-frequency factor, the macroeconomic variables have lower predictive power for the business-cycle-frequency factor, with the R-squared of the one- and two-factor PLS prediction models being 4.86% and 7.45%, respectively. However, these values are still economically sizable, and the predictive coefficients of both PLS factors are statistically significant based on all considered tests. By contrast, the financial variables have a strong predictive power for the business-cycle-frequency factor returns. Specifically, the R-squared of the one-factor PLS prediction model is 17.14%, and all considered tests suggest statistical significance. For the two-factor PLS prediction model, the R-squared further increases to 17.89%, but the predictive coefficient of the second PLS factor is only marginally significant based on the test of Kelly and Pruitt (2015) and is insignificant based on the bootstrap p-value. In general, although the business-cycle-frequency factor appears less predictable than the below-business-cycle-frequency factor, the R-squared for predicting the business-cycle-frequency factor returns is still substantially higher than that for predicting the returns of popular asset-pricing factors as reported in the literature.

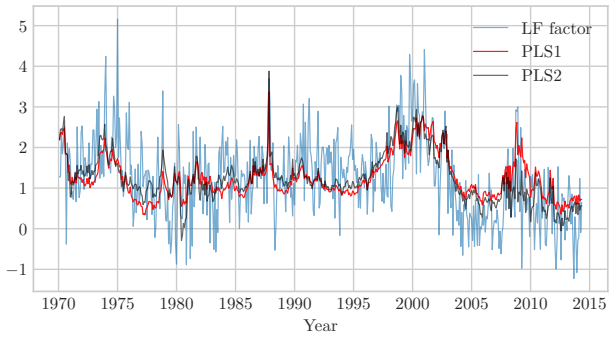
Figure 5 illustrates the returns of the two low-frequency factors estimated from all 37 Basic-IPCA factors and their in-sample predictions based on all 43 economic variables. Panels A and B correspond to the below-business-cycle- and business-cycle-frequency factors, respectively. In each panel, the blue line graphs the returns of the low-frequency factor, and the red and black lines graph the predictions of the low-frequency factor returns when using the one- and two-factor PLS prediction models, respectively. Note that the economic variables successfully capture the slow-moving fluctuations, especially of the below-business-cycle-frequency factor returns. For instance, the economic variables identify the increase of the returns of the below-business-cycle-frequency factor from 1995 to 2000, and the decrease of the returns from 2000 to 2006. For the business-cycle-frequency factor, the economic variables also identify the sharp increase of the returns before the financial crisis. Thus, the reason why the economic variables

[36] Motivated by that the PLS prediction model delivers a high predictive R-squared for the returns of the below-business-cycle-frequency factor, I also use the PLS factors of this prediction model to predict the market returns. Specifically, I find that the one PLS factor estimated from all 43 economic variables, which is constructed to predict the returns of the below-business-cycle-frequency factor, predicts the monthly returns of the market factor with an in-sample R-squared of 2.58%. This figure is quantitatively similar to that reported by Kelly and Pruitt (2013), who use a cross-section of portfolio-level book-to-market ratios and the PLS model to predict the market returns. The predictive coefficient of this PLS factor is significantly negative. This supports the finding that the state variable that drives up the risk premia of a large number of candidate factors simultaneously drives down the market premium.

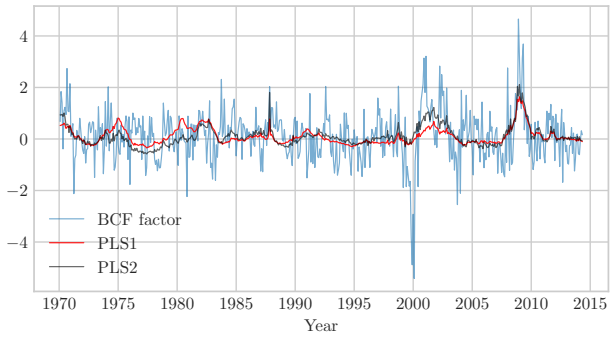
Figure 5: Low-frequency factor returns and their predictions based on all 43 economic variables

This figure illustrates the monthly returns of the two low-frequency factors estimated from all 37 Basic-IPCA factors and their in-sample predictions based on all 43 economic variables. Panels A and B correspond to the below-business-cycle- and business-cycle-frequency factors, respectively. In each panel, the blue line graphs the returns of the low-frequency factor, and the red and black lines graph the predictions of the low-frequency factor returns when using the one- and two-factor PLS prediction models, respectively.

(A) Below-business-cycle-frequency factor



(B) Business-cycle-frequency factor



have high predictive power for the low-frequency factor returns is that they successfully capture the slow-moving fluctuations of the low-frequency factor returns.

Note that a PLS factor is a linear combination of the economic variables. Thus, from every PLS prediction model, I can back out the contribution of each individual economic variable to the prediction of the returns of a low-frequency factor. Table 8 reports the contributions of the most important economic variables to the prediction of the low-frequency factor returns in a one-factor PLS prediction model, and Panels A and B report the results for the below-business-cycle- and business-cycle-frequency factors, respectively. Each panel reports the contributions of the most important economic variables among the 19 macroeconomic variables, 24 financial variables, and all 43 economic variables. Panel A suggests that the financial uncertainty measures of Jurado et al. (2015) and Ludvigson et al. (2021) provide substantial contributions to the prediction of the below-business-cycle-frequency factor returns, with the 12-month uncertainty measure being the most important predictor. In particular, the uncertainty measures positively predict the below-business-cycle-frequency factor returns, which is consistent with the intuition that higher uncertainties are associated with higher risk premia in the financial market.[37] As a result, I can interpret the below-business-cycle-frequency state variable as a

[37]The literature has documented mixed findings on the relation between the risk and return (especially for the market). For example, French, Schwert, and Stambaugh (1987); Harvey (1989); Campbell and Hentschel (1992); Ludvigson and Ng (2007) find a positive relation between market risk and return, while Campbell (1987); Breen et al. (1989); Glosten, Jagannathan, and Runkle (1993); Brandt and Kang (2004) find a negative relation. In particular, Moreira and Muir (2017) find that changes in volatility are not associated with proportional changes in expected returns. In this case, the risk price is lower when the volatility is higher, posing a challenge to macro-finance models. In contrast, the uncertainty measures positively predicting the returns of the low-frequency factors is more consistent with the prediction of these theoretical models. Two possible explanations for the difference between my results and those of Moreira and Muir (2017) are (i) the uncertainty measures and realized market volatility have different dynamics, and (ii) the low-frequency factors capture components of the cross-section different from the market. In addition, the different results may also originate from the different versions of the observed candidate factors. In particular, different types of conditional scaling of the risk factors mechanically generate different conclusions for the relation between risk and return. Specifically, suppose there exists only one systematic risk and its risk price is constant. In this case, the risk factor is

$$f_{t+1} = \mu\sigma_t^2 + \sigma_t\xi_{t+1},$$

Table 8: Most important predictors for the two low-frequency factors

This table reports the contributions of the most important economic variables to the prediction of the low-frequency factor returns in a one-factor PLS prediction model. Panels A and B report the results for the below-business-cycle- and business-cycle-frequency factors, respectively. Each panel reports the contributions of the most important economic variables among the 19 macroeconomic variables, 24 financial variables, and all 43 economic variables. A more detailed description of the economic variables is in Table 15 in Appendix F.

(A) Below-business-cycle-frequency factor

Macro		Financial		Macro+financial	
Predictor	Coefficient	Predictor	Coefficient	Predictor	Coefficient
FIN_UNC_12month	0.0908	AggLiq	-0.1186	FIN_UNC_12month	0.1093
FIN_UNC_3month	0.0876	intermed_lev_ratio_squared	-0.0788	FIN_UNC_3month	0.1072
FIN_UNC_1month	0.0871	TMS	-0.0726	FIN_UNC_1month	0.1071
BCI	-0.0531	intermed_capital_ratio	0.0696	AggLiq	-0.0973
CLI	-0.0512	DFR	-0.0694	intermed_lev_ratio_squared	-0.0663
gConsnon	-0.0334	SVAR	0.0691	TMS	-0.0606
RECESS	0.0296	DISAG	0.0647	DFR	-0.0590
gIndprod	-0.0275	EP	-0.0628	CLI	-0.0556
REAL_UNC_1month	-0.0254	PLS_SENT	0.0622	EP	-0.0522
REAL_UNC_12month	-0.0243	Real_FedFundsRate	0.0612	DFY	-0.0512

(B) Business-cycle-frequency factor

Macro		Financial		Macro+financial	
Predictor	Coefficient	Predictor	Coefficient	Predictor	Coefficient
CLI	-0.0251	AggLiq	-0.1455	AggLiq	-0.0335
gConsdur	-0.0232	SVAR	0.0965	CLI	-0.0292
gConsnon	-0.0221	PLS_SENT	0.0927	gConsnon	-0.0271
FIN_UNC_1month	0.0209	DE	0.0871	gConsdur	-0.0259
FIN_UNC_3month	0.0205	PLS_SENTorthogonalized	0.0800	SVAR	0.0240
RECESS	0.0197	NTIS	-0.0766	PLS_SENT	0.0232
FIN_UNC_12month	0.0191	DFY	0.0724	FIN_UNC_1month	0.0228
gIndprod	-0.0175	EP	-0.0683	gIndprod	-0.0227
BCI	-0.0169	intermed_capital_ratio	-0.0672	FIN_UNC_3month	0.0224
MACRO_UNC_1month	0.0166	BW_SENT	0.0464	BCI	-0.0223

fundamental "aggregate uncertainty" variable that moves the risk premia of a large number of candidate factors simultaneously from the same direction.

Panel B shows that the U.S. composite leading indicator, durable consumption growth, and non-durable consumption growth have a relatively strong predictive power to the business-cycle-frequency factor returns. The financial uncertainty measures of Jurado et al. (2015) and Ludvigson et al. (2021) also predict the business-cycle-frequency factor returns, while different from predicting the below-business-cycle-frequency factor returns, the one-month uncertainty measure is the most important predictor. However, the predictive power of these variables that reflect the fundamentals for the business-

where ξ_{t+1} is the source of risk with unit standard deviation, σ_t is the conditional volatility of the systematic risk and thus measures the quantity of risk, and μ is the price of risk. If one observes f_{t+1}, then she could conclude that the risk price is constant. However, if one observe a conditionally-scaled version of this risk factor

$$\frac{1}{\sigma_t^2} f_{t+1} = \mu + \frac{1}{\sigma_t} \xi_{t+1},$$

then her conclusion is that the risk price decreases when the volatility of the scaled factor increases. The two versions of the factor span the same component in the conditional SDF and thus deliver the same asset-pricing implication, but they generate different conclusions on the risk-return tradeoff. Empirically, it is hard to distinguish which version of the risk factor is observed (the researchers might even observe different versions at different times), which possibly explains the mixed findings in the literature on the relation between risk and return.

cycle-frequency factor returns is quantitatively lower than for the below-business-cycle-frequency factor returns. Therefore, compared to the business-cycle-frequency state variable, the below-business-cycle-frequency state variable better reflects the fundamental variation of the risk premia.

Among the financial variables, Panel A shows that the aggregate liquidity measure of Pastor and Stambaugh (2003), the squared intermediary leverage ratio of He et al. (2017), and the term spread between the long-term yield and Treasury bill strongly predict the below-business-cycle-frequency factor returns. For the business-cycle-frequency factor, the aggregate liquidity measure of Pastor and Stambaugh (2003), the variance of the S&P 500 index, and the investment sentiment measures of Huang et al. (2015) appear as the most important predictors. On average, the financial indicators, which mainly reflect the market friction, investor sentiment, and stock-market short-term behaviors, are quantitatively more important for predicting the business-cycle-frequency factor returns. These variables capture the dynamics of the business-cycle-frequency state variable, which affect the risk premia of different candidate factors from different directions. Moreover, when predicting the returns of the below-business-cycle- and business-cycle frequency factors using all economic variables, only four variables (one-month financial uncertainty, three-month financial uncertainty, aggregate liquidity, and U.S. composite leading indicator) appear as the top-ten predictors for predicting the returns of both low-frequency factors, confirming that the latent state variables that drive the risk premia of the two low-frequency factors have different time-series dynamics.

Note that the predictive sign of some variables seems inconsistent with the existing findings in the literature. For instance, the term spread is found to positively predict the market return in the literature (e.g., Fama and French, 1989), while it negatively predicts the below-business-cycle-frequency factor returns. As explained in Section 3.1, the two low-frequency factors have negative correlations with the market factor, and thus the market factor has negative loadings on the two latent low-frequency state variables. Correspondingly, a variable that positively predicts the market return may negatively predict the below-business-cycle-frequency factor returns.

3.3.2 Pseudo-out-of-sample prediction

In this section, I use a pseudo-out-of-sample analysis to investigate whether the macroeconomic and financial variables predict the returns of the low-frequency factors and how the predictability (if it exists) affects the low-frequency factors' explanatory power for the returns of individual stocks and characteristic-managed portfolios. Specifically, in the pseudo-out-of-sample analysis, I estimate the IPCA and low-frequency factors using the full sample of return data but predict their returns at month $t+1$ recursively using the information of the economic predictor variables only up to month t. In particular, I set the length of the initial prediction window as $M = 240$ months, and then add data of the next month as the prediction moves forward through time. Therefore, the out-of-sample periods of prediction models that use the macroeconomic variables span from July 1984 to May 2014, and those that use the financial variables (and both sets of variables) span from February 1990 to May 2014.

To evaluate whether the prediction of the low-frequency factor returns based on the PLS prediction model outperforms the historical-mean benchmark, I use the out-of-sample R-squared of Campbell

and Thompson (2008):

$$R^2_{\text{OOS}} = 1 - \frac{\sum_{t=M+1}^{T} \left(f_t - \widehat{f}_t\right)^2}{\sum_{t=M+1}^{T} \left(f_t - \bar{f}_t\right)^2}, \tag{41}$$

where $t = M+1, \ldots, T$ spans the out-of-sample periods, f_t contains the return of a factor at month t, and $\widehat{f}_t$ and $\bar{f}_t$ contain the out-of-sample predictions of the factor return based on the PLS prediction model and historical mean, respectively. Note that the numerator and denominator of the second term on the right-hand side of (41) give the mean squared forecast error (MSFE) for the factor returns of the PLS prediction model and the historical-mean benchmark, respectively. Therefore, if the PLS model outperforms the historical-mean benchmark, it delivers a lower mean squared forecast error than the historical-mean benchmark, and thus the out-of-sample R-squared is positive. Hence, my null hypothesis is $R^2_{\text{OOS}} \leq 0$, and my alternative hypothesis is $R^2_{\text{OOS}} > 0$. I test this hypothesis using the MSFE-adjusted statistic of Clark and West (2007), which is conveniently implemented via a simple regression:

$$\left(f_t - \bar{f}_t\right)^2 - \left(f_t - \widehat{f}_t\right)^2 + \left(\bar{f}_t - \widehat{f}_t\right)^2 = u + v_t, \tag{42}$$

where u is the intercept and v_t is the error term. The t-statistic corresponding to the OLS estimate of the intercept u is used to test my hypothesis.[38] Specifically, if the intercept u is significantly greater than zero, then the test rejects the null hypothesis of $R^2_{\text{OOS}} \leq 0$.

Table 9 reports the out-of-sample R-squared in the pseudo-out-of-sample analysis for predicting the monthly returns of the low-frequency factors estimated from the Basic-IPCA factors. Panels A and B report the results for the below-business-cycle- and business-cycle-frequency factors, respectively. Each panel includes the prediction results when using the one- and two-factor PLS prediction models and when using the 19 macroeconomic variables, 24 financial variables, and all 43 economic variables as predictors. In each panel, the first column gives the category of the predictor variables. The following ten columns give the R-squared for predicting the returns of the low-frequency factor estimated from K Basic-IPCA factors with K ranging between three and twelve, and the last column for predicting the returns of the low-frequency factor estimated from all 37 Basic-IPCA factors.

Panel A of Table 9 shows that, starting from five Basic-IPCA factors and using the one-factor PLS prediction model, the returns of the below-business-cycle-frequency factor estimated from the Basic-IPCA factors are highly predicted by every set of economic variables. Specifically, the out-of-sample R-squared exceeds 10% for predicting the returns of the below-business-cycle-frequency factor estimated from more than eight Basic-IPCA factors, and the out-of-sample R-squared is statistically significant at the 1% level in most cases. The highest R-squared, 18.49%, appears when one uses all 43 economic variables to predict the returns of the below-business-cycle-frequency factor estimated from all 37 Basic-IPCA factors.

When using the two-factor PLS model to predict the below-business-cycle-frequency factor returns, the R-squared becomes even higher and remains statistically significant. For instance, when using the macroeconomic variables to predict the returns of the below-business-cycle-frequency factor estimated from all 37 Basic-IPCA factors, the R-squared increases from 10.36% for the one-factor PLS prediction model to 19.89% for the two-factor PLS prediction model. Moreover, when using all 43

[38] I calculate the t-statistic of u using the Newey and West (1987) standard error with a 12-month window.

Table 9: Out-of-sample R-squared in pseudo-out-of-sample analysis for predicting low-frequency factors

This table reports the out-of-sample R-squared (in percent) in the pseudo-out-of-sample analysis for predicting the monthly returns of the low-frequency factors estimated from the Basic-IPCA factors. Panels A and B report the results for the below-business-cycle- and business-cycle-frequency factors, respectively. Each panel includes the prediction results when using the one- and two-factor PLS prediction models and when using the 19 macroeconomic variables, 24 financial variables, and all 43 economic variables as predictors. In each panel, the first column gives the category of the predictor variables. The following ten columns give the R-squared for predicting the returns of the low-frequency factor estimated from K Basic-IPCA factors with K ranging between three and twelve, and the last column for predicting the returns of the low-frequency factor estimated from all 37 Basic-IPCA factors. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, based on the MSFE-adjusted test of Clark and West (2007).

(A) Below-business-cycle-frequency factor

Predictor type	Number of IPCA factors to estimate low-frequency factors										
	3	4	5	6	7	8	9	10	11	12	All
One PLS factor											
Macro	< 0	< 0	4.79***	6.00***	5.14**	13.44***	17.77***	16.66***	14.86***	15.29***	10.36***
Financial	< 0	< 0	3.99***	5.57***	6.17***	11.70**	14.00***	14.06***	14.41***	13.98***	15.01***
Macro+financial	< 0	< 0	6.43***	7.49***	6.45***	14.29***	17.08***	16.65***	16.32***	16.27***	18.49***
Two PLS factors											
Macro	< 0	< 0	13.49***	14.17***	12.31***	19.86***	22.26***	21.66***	20.51***	20.73***	19.89***
Financial	< 0	1.93	6.85***	8.89***	7.38***	16.30***	22.00***	21.46***	21.54***	21.32***	23.33***
Macro+financial	< 0	< 0	13.88***	14.94***	15.01***	25.14***	29.09***	28.35***	28.17***	27.97***	32.51***

(B) Business-cycle-frequency factor

Predictor type	Number of IPCA factors to estimate low-frequency factors										
	3	4	5	6	7	8	9	10	11	12	All
One PLS factor											
Macro	< 0	9.73***	< 0	< 0	10.58*	3.65*	< 0	< 0	0.35	0.18	6.22*
Financial	< 0	< 0	< 0	< 0	10.78	8.01**	5.40**	5.45**	1.02	1.03	12.51**
Macro+financial	< 0	1.50**	< 0	< 0	15.85*	8.24*	3.36**	4.05**	2.60	2.15	13.18**
Two PLS factors											
Macro	0.01	5.86***	< 0	< 0	11.56*	3.51*	< 0	< 0	< 0	< 0	4.99*
Financial	< 0	3.09*	1.05	0.93	11.49	11.30*	7.32**	7.69**	3.62*	3.78*	17.51**
Macro+financial	< 0	2.56***	< 0	< 0	16.43*	8.97*	2.76*	3.36*	1.43	1.40	15.49**

economic variables, the R-squared is 32.51% for the two-factor PLS model and thus is substantially higher than its one-PLS-factor counterpart, 18.49%. Therefore, I conclude that the returns of the below-business-cycle-frequency factor are highly predictable by the economic variables in the pseudo-out-of-sample analysis.

Consistent with the in-sample analysis, the business-cycle-frequency factor returns are less predictable. In particular, most business-cycle-frequency factors estimated from a small number of Basic-IPCA factors cannot be significantly predicted by the selected economic variables. For instance, even though predicting the returns of the business-cycle-frequency factor estimated from seven Basic-IPCA factors yields an R-squared greater than 10% when using any set of predictor variables and when using either the one- or two-factor PLS model, the high R-squared is not statistically significant at the 5% level. In contrast, the returns of the business-cycle-frequency factor estimated from all 37 Basic-IPCA factors can be predicted more robustly. Specifically, when using the one- and two-factor PLS models, its prediction based on the financial variables yields a significant (at the 5% level) R-squared of 12.51% and 17.51%, respectively, and its prediction based on all 43 economic variables yields a significant (at

the 5% level) R-squared of 13.18% and 15.49%, respectively. One explanation of the difference between the business-cycle-frequency factor estimated from a small number of Basic-IPCA factors and that estimated from all 37 Basic-IPCA factors is the findings in Figure 3, which shows that the higher-order IPCA factors capture the information of the latent business-cycle-frequency state variable. Thus, exploiting these higher-order IPCA factors to identify the business-cycle-frequency factor increases the signal of the latent state variable that drives the risk premium of the business-cycle-frequency factor, and thus makes the factor more predictable.

Thus far, I find that the selected macroeconomic and financial variables highly predict the returns of the low-frequency factors. Consequently, it is of interest to know whether the returns of the IPCA factors are as predictable as those of the low-frequency factors. One approach to systematically compare the predictability of these two sets of factors is to compare how the predictions of the returns of these factors explain the returns of individual stocks and characteristic-managed portfolios. In particular, I evaluate the predictability of a set of factors via the predictive R-squared with varying risk premia, which is defined as

$$\text{Predictive } R^2 \text{ with varying factor RP} = 1 - \frac{\sum_{i,t}\left(r_{i,t+1} - z_{i,t}^{\top}\Gamma_{\beta}\widehat{f}_{t+1}\right)^2}{\sum_{i,t} r_{i,t+1}^2} \tag{43}$$

for the returns of individual stocks, where $\hat{f}_{t+1}$ contains the prediction of the factor returns.[39] To simplify the exposition, I only report the results where the factor returns are predicted by a one-factor PLS model.[40]

Panels A and B of Table 10 report the predictive R-squared with varying factor risk premia for explaining the monthly returns of individual stocks and portfolios managed by the 36 considered characteristics, respectively. Each panel reports the results for both the Basic-IPCA and Freq-IPCA frameworks. For each set of economic variables (macroeconomic variables, financial variables, and their combination), the first row reports the predictive R-squared of the K-factor IPCA model with K ranging between three and twelve, and the second row of the model that includes one below-business-cycle-frequency factor and one business-cycle-frequency factor estimated from K IPCA factors.

There are two main discoveries based on this table. First, both panels of Table 10 show that the prediction of the low-frequency factor returns explains the returns of individual stocks and managed portfolios better than that of the IPCA factor returns. For instance, Panel A shows that the predictive R-squared for explaining the individual stock returns of models that include two low-frequency factors estimated from no less than five IPCA factors is always positive and higher than that of their IPCA counterparts. In particular, when using the financial variables to predict the factor returns, no Basic-IPCA model delivers a positive predictive R-squared for explaining individual stock returns, and only the nine-, ten-, and twelve-factor Freq-IPCA models deliver a positive predictive R-squared. Thus, the IPCA factors cannot be robustly predicted by the selected economic variables to explain the returns of individual stocks. In contrast, the economic variables successfully predict the returns of the low-frequency factors such that these factors explain well the risk premia of individual stocks.

Panel B shows that the prediction of the IPCA factor returns explains the returns of managed portfolios better than individual stocks. In particular, both the Basic-IPCA and Freq-IPCA models

[39] This R-squared statistic can also be defined similarly for the returns of managed portfolios of characteristics.

[40] In unreported results, I find that the results where a two-factor PLS model predicts the factor returns are qualitatively the same as those reported in this section.

Table 10: Predictive R-squared with varying factor risk premia

Panels A and B report the predictive R-squared with varying factor risk premia (in percent) for explaining the monthly returns of individual stocks and portfolios managed by the 36 considered characteristics, respectively. Each panel reports the results for both the Basic-IPCA and Freq-IPCA frameworks. All factor returns are predicted by a one-factor PLS model. For each set of economic predictors (macroeconomic variables, financial variables, and their combination), the first row reports the R-squared of the K-factor IPCA model with K ranging between three and twelve, and the second row of the model that includes one below-business-cycle-frequency factor and one business-cycle-frequency factor estimated from K IPCA factors.

(A) Individual stocks

Predictor type	Factor type	Number of IPCA factors										
		3	4	5	6	7	8	9	10	11	12	All
Basic-IPCA												
Macro	IPCA	0.11	0.08	0.37	0.20	0.40	0.33	0.19	0.13	0.36	0.24	-
	B-BCF1+BCF1	0.23	< 0	0.52	0.52	0.59	0.71	0.50	0.51	0.53	0.52	0.56
Financial	IPCA	< 0	< 0	< 0	< 0	< 0	< 0	< 0	< 0	< 0	< 0	-
	B-BCF1+BCF1	< 0	< 0	0.58	0.57	0.60	0.79	0.73	0.73	0.65	0.66	0.54
Macro+financial	IPCA	< 0	< 0	< 0	< 0	< 0	< 0	< 0	< 0	< 0	< 0	-
	B-BCF1+BCF1	< 0	< 0	0.58	0.57	0.59	0.80	0.70	0.73	0.56	0.56	0.59
Freq-IPCA												
Macro	IPCA	0.21	0.03	0.54	0.53	0.59	0.52	0.59	0.66	0.30	0.57	-
	B-BCF1+BCF1	0.32	0.03	0.66	0.64	0.75	0.66	0.61	0.65	0.66	0.67	0.74
Financial	IPCA	< 0	< 0	< 0	< 0	< 0	< 0	0.14	0.18	< 0	0.33	-
	B-BCF1+BCF1	0.04	< 0	0.74	0.75	0.79	0.88	0.84	0.85	0.85	0.78	0.78
Macro+financial	IPCA	< 0	< 0	< 0	< 0	0.15	0.22	0.32	0.58	< 0	0.34	-
	B-BCF1+BCF1	< 0	< 0	0.66	0.58	0.79	0.92	0.80	0.81	0.81	0.73	0.74

(B) Managed portfolios

Predictor type	Factor type	Number of IPCA factors										
		3	4	5	6	7	8	9	10	11	12	All
Basic-IPCA												
Macro	IPCA	0.51	0.65	1.94	1.50	2.47	2.48	2.37	2.36	2.70	2.16	-
	B-BCF1+BCF1	1.20	< 0	2.27	2.25	3.27	3.40	2.47	2.61	2.88	2.84	3.09
Financial	IPCA	1.43	1.22	1.69	1.73	1.27	2.41	1.95	1.91	1.90	2.34	-
	B-BCF1+BCF1	0.05	< 0	2.49	2.48	3.13	4.12	3.86	3.93	3.13	3.22	3.15
Macro+financial	IPCA	1.30	1.51	2.26	1.84	2.25	2.69	2.72	2.53	2.73	3.07	-
	B-BCF1+BCF1	1.03	< 0	2.70	2.63	3.29	4.04	3.24	3.47	2.35	2.41	2.86
Freq-IPCA												
Macro	IPCA	0.28	0.39	2.42	2.86	2.58	2.31	2.78	3.16	1.73	2.31	-
	B-BCF1+BCF1	1.44	< 0	3.40	3.34	4.17	3.57	3.25	3.48	3.55	3.68	4.48
Financial	IPCA	1.73	1.61	2.11	1.98	1.94	3.16	2.81	2.70	1.83	2.96	-
	B-BCF1+BCF1	0.56	< 0	3.67	3.65	4.11	4.61	4.37	4.43	4.43	3.96	5.57
Macro+financial	IPCA	1.55	1.75	3.12	2.67	2.55	2.98	3.10	3.38	2.13	2.43	-
	B-BCF1+BCF1	1.06	< 0	3.57	2.89	4.34	4.53	3.73	3.85	3.92	3.46	4.74

consistently deliver a positive predictive R-squared for explaining the returns of managed portfolios. However, the predictive R-squared of the higher-dimensional IPCA models is, in most cases, lower than that of the models that only include the two low-frequency factors. For instance, when the factors are predicted by the macroeconomic variables, the predictive R-squared of the nine-factor Basic-IPCA model is 1.95%, while that of the model that includes the two low-frequency factors estimated from

the nine Basic-IPCA factors is 3.86%. Thus, the prediction of the low-frequency factor returns also outperforms that of the IPCA factor returns in explaining the risk premia of managed portfolios.

Second, the prediction of the returns of factors estimated from the Freq-IPCA framework explains the returns of individual stocks and managed portfolios better than that of factors estimated from the Basic-IPCA framework. For instance, when using all 43 economic variables to predict the factor returns, no Basic-IPCA model delivers a positive predictive R-squared for explaining individual stock returns. In contrast, the seven-, eight-, nine-, ten- and twelve-factor Freq-IPCA models deliver a positive predictive R-squared. When using the financial variables to predict the factor returns, the model that includes the two low-frequency factors estimated from all 37 Basic-IPCA factors delivers a predictive R-squared of 0.54% for explaining individual stock returns and 3.15% for explaining managed portfolio returns. As a comparison, the two corresponding predictive R-squared values of the model that includes the two low-frequency factors estimated from all 73 Freq-IPCA factors are 0.78% and 5.57%, respectively. Thus, the Freq-IPCA framework facilitates identifying the predictable components of asset returns.

In summary, the pseudo-out-of-sample results in this section show that the selected economic variables highly predict the returns of the low-frequency factors. Moreover, the IPCA factors are not as predictable as the low-frequency factors because the prediction of the latter outperforms that of the IPCA factors in explaining the returns of individual stocks and managed portfolios.

3.3.3 Real-time out-of-sample prediction

In this section, I use an out-of-sample analysis to investigate whether the selected economic variables can predict the low-frequency factor returns in real-time and whether the predictability (if it exists) generates economic value for an investor.

The out-of-sample analysis proceeds as follows. I use the first $M = 240$ months of return data to estimate the IPCA factors and then estimate the low-frequency factors from the estimated IPCA factors. Each low-frequency factor is normalized to have a positive average return and unit standard deviation. I then exploit the first 240 months of economic predictor data to estimate a prediction model for the factor returns. The parameters for estimating the IPCA and low-frequency factors are used to construct the out-of-sample one-month-ahead realization of these factors, and the parameters of the prediction model are used to construct the out-of-sample one-month-ahead prediction of the factor returns. I repeat the above procedure with added data as the factor estimation and prediction move forward through time. Importantly, this procedure guarantees that the estimation and prediction of the factors at month $t + 1$ use no information beyond month t. Similar to the pseudo-out-of-sample procedure, the out-of-sample periods of prediction models that use the macroeconomic variables span from July 1984 to May 2014, and those that use the financial variables (and both sets of variables) span from February 1990 to May 2014. To simplify the exposition, I only report results for factors estimated from the Basic-IPCA framework in this section.

The analysis in the previous sections shows that using two PLS factors in the prediction model provides additional predictability for the low-frequency factor returns on top of the one-factor PLS model. However, using an extra PLS factor requires me to estimate more parameters in the prediction model, which could cause an overfitting issue in the out-of-sample analysis. Thus, to balance

Table 11: Out-of-sample R-squared for predicting low-frequency factors

This table reports the out-of-sample R-squared (in percent) for predicting the monthly returns of the low-frequency factors estimated from the Basic-IPCA factors. Panels A and B report the results for the below-business-cycle- and business-cycle-frequency factors, respectively. Each panel includes the prediction results when using the one-factor PLS prediction model and the PLS model with the number of factors selected by BIC and when using the 19 macroeconomic variables, 24 financial variables, and all 43 economic variables as predictors. In each panel, the first column gives the category of the predictor variables. The following ten columns give the R-squared for predicting the returns of the low-frequency factor estimated from K Basic-IPCA factors with K ranging between three and twelve, and the last column for predicting the low-frequency factor estimated from all 37 Basic-IPCA factors. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, based on the MSFE-adjusted test of Clark and West (2007).

(A) Below-business-cycle-frequency factor

Predictor type	Number of IPCA factors to estimate low-frequency factors										
	3	4	5	6	7	8	9	10	11	12	All
One PLS factor											
Macro	< 0	0.69**	3.71***	4.81***	5.32***	9.49***	9.99***	8.76***	9.20***	7.66***	7.56***
Financial	< 0	2.90**	4.24***	4.10***	2.22**	2.09**	2.70*	2.74**	3.96**	1.42**	3.34**
Macro+financial	< 0	< 0	3.22***	2.39***	< 0	2.29**	3.74**	2.73***	4.27***	0.79**	6.17***
Number of PLS factors selected by BIC											
Macro	< 0	4.15***	7.86***	8.54***	8.44***	10.04***	11.06***	9.59***	9.14***	7.48***	11.48***
Financial	< 0	3.28***	5.14***	5.36***	2.12**	5.15***	9.24***	7.37***	8.79***	6.03***	7.57***
Macro+financial	< 0	2.13***	7.58***	7.08***	8.96***	8.36***	10.66***	9.21***	8.76***	5.04***	13.04***

(B) Business-cycle-frequency factor

Predictor type	Number of IPCA factors to estimate low-frequency factors										
	3	4	5	6	7	8	9	10	11	12	All
One PLS factor											
Macro	< 0	2.18	< 0	0.62	3.03**	0.91*	< 0	3.84***	4.98**	3.19**	2.04**
Financial	< 0	< 0	< 0	< 0	< 0	< 0	< 0	< 0	< 0	2.39*	1.19**
Macro+financial	< 0	< 0	< 0	< 0	1.24**	< 0	< 0	0.40*	7.03*	4.51**	4.24***
Number of PLS factors selected by BIC											
Macro	< 0	3.82*	1.83	2.63*	3.53**	2.40**	2.14**	6.22***	6.42**	5.13**	2.09**
Financial	< 0	< 0	< 0	< 0	< 0	< 0	< 0	< 0	< 0	4.12**	0.97**
Macro+financial	< 0	0.24	< 0	< 0	0.19	< 0	< 0	2.63*	9.08*	4.67*	4.32***

between goodness-of-fit and parsimony in the prediction model, in addition to estimating a one-factor PLS prediction model, I also estimate a prediction model in which the number of PLS factors is selected based on the Bayesian Information Criterion (BIC) of Kelly and Pruitt (2015).

Table 11 reports the out-of-sample R-squared for predicting the monthly returns of the low-frequency factors estimated from the Basic-IPCA factors. Panels A and B report the results for the below-business-cycle- and business-cycle-frequency factors, respectively. Each panel includes the prediction results when using the one-factor PLS prediction model and the PLS model with the number of factors selected by the BIC criterion and when using the 19 macroeconomic variables, 24 financial variables, and all 43 economic variables as predictors. In each panel, the first column gives the category of the predictor variables. The following ten columns give the R-squared for predicting the returns of the low-frequency factor estimated from K Basic-IPCA factors with K ranging between three and twelve, and the last column for predicting the returns of the low-frequency factor estimated from all 37 Basic-IPCA factors.

Consistent with the findings of the pseudo-out-of-sample analysis, Panel A of Table 11 shows that starting from five Basic-IPCA factors, the returns of the below-business-cycle-frequency factor estimated from the Basic-IPCA factors are highly predicted by the selected economic variables using the one-factor PLS model.[41] For instance, the macroeconomic variables deliver an economically sizable out-of-sample R-squared ranging between 4.81% to 9.99% for predicting the returns of the below-business-cycle-frequency factor, and the R-squared is statistically significant at the 1% level in all cases. In particular, the R-squared for predicting the returns of the below-business-cycle-frequency factor estimated from all 37 Basic-IPCA factors is 7.56%. The financial variables, in general, deliver a lower out-of-sample R-squared. For instance, using the financial variables to predict the returns of the below-business-cycle-frequency factor estimated from all 37 Basic-IPCA factors yields an R-squared of 3.34%, which is still sizable and statistically significant at the 5% level. When using all 43 economic variables to predict the returns of the below-business-cycle-frequency factor estimated from all 37 Basic-IPCA factors, the R-squared becomes 6.17% and is significant at the 1% level.

When using the PLS prediction model with the number of factors selected by the BIC criterion in place of the one-factor PLS model, the out-of-sample R-squared for predicting the returns of the below-business-cycle-frequency factor further increases in most cases. For instance, when using the macroeconomic variables to predict the returns of the below-business-cycle-frequency factor estimated from all 37 Basic-IPCA factors, the R-squared increases from 7.56% for the one-factor PLS model to 11.48% for the model with the number of PLS factors selected by the BIC criterion. Moreover, using all 43 economic variables to predict the returns of the below-business-cycle-frequency factor estimated from all 37 Basic-IPCA factors yields an R-squared of 13.04%, substantially higher than the one-PLS-factor counterpart, 6.17%.

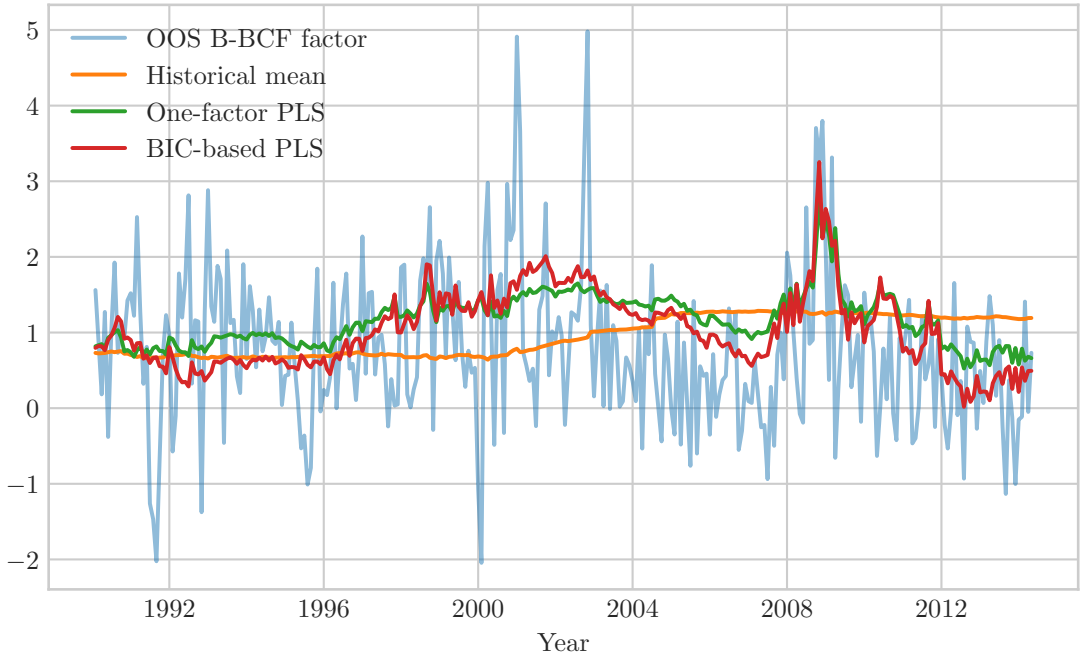
The business-cycle-frequency factor returns are less predictable by the selected economic variables compared to the below-business-cycle-frequency factor returns. In particular, when using the one-factor PLS model to predict the business-cycle-frequency factor returns, the out-of-sample R-squared is negative when a small number of IPCA factors are used to estimate the business-cycle-frequency factor. Nonetheless, the business-cycle-frequency factor estimated from all 37 Basic-IPCA factors is still significantly predictable by the selected economic variables. Specifically, when using all 43 economic variables to predict the returns of this factor, the R-squared is 4.24% for the one-factor PLS model and 4.32% for the PLS prediction model with the number of factors selected by the BIC criterion. In both cases, the R-squared is significant at the 1% level. As explained in the previous section, one explanation of the difference between the business-cycle-frequency factor estimated from a small number of Basic-IPCA factors and that estimated from all 37 Basic-IPCA factors is that the higher-order IPCA factors capture the information of the latent business-cycle-frequency state variable. Exploiting the higher-order IPCA factors to estimate the business-cycle-frequency factor increases the signal of this latent state variable, and thus makes the factor more predictable.

To provide better intuition about how the selected economic variables predict the below-business-cycle-frequency factor returns, Figure 6 illustrates the out-of-sample realizations of the monthly returns of this factor estimated from all 37 Basic-IPCA factors and their out-of-sample predictions based on

[41] In unreported results, I find that the out-of-sample Sharpe ratio of the below-business-cycle-frequency factor estimated from K Basic-IPCA factors ranges between 0.759 and 0.918 when K ranges between five and twelve. Moreover, the out-of-sample Sharpe ratio of the below-business-cycle-frequency factor estimated from all 37 Basic-IPCA factors is 0.774. Thus, consistent with the in-sample findings, the below-business-cycle-frequency factor earns a high Sharpe ratio out-of-sample.

Figure 6: Out-of-sample below-business-cycle-frequency factor and its predictions

This figure illustrates the out-of-sample realizations of the monthly returns of the below-business-cycle-frequency factor estimated from all 37 Basic-IPCA factors and their predictions based on all 43 economic variables. The blue line graphs the out-of-sample realizations of the factor returns. The orange, green, and red lines graph the predictions of the factor returns based on the historical mean, the one-factor PLS prediction model, and the PLS model with the number of factors selected by the BIC criterion, respectively.



all 43 economic variables. The blue line graphs the out-of-sample realizations of the factor returns. The orange, green, and red lines graph the predictions of the factor returns based on the historical mean, the one-factor PLS prediction model, and the PLS model with the number of factors selected by the BIC criterion, respectively. This figure shows that the out-of-sample below-business-cycle-frequency factor has slow-moving fluctuations (though less salient than its in-sample counterpart). Specifically, the returns of this factor gradually increased from 1996 to 1999, then decreased until the start of the financial crisis. The returns then sharply increased during the financial crisis and then decreased afterward. The 43 economic variables successfully capture the slow-moving fluctuations of the below-business-cycle-frequency factor returns, which are uncaptured by the historical-mean prediction. This explains the economically sizable out-of-sample R-squared for predicting the returns of the below-business-cycle-frequency factors.

Next, I study whether the statistically significant predictability of the two low-frequency factors is associated with significantly improved economic value for an investor.[42] Specifically, one conventional approach to evaluate the economic value a prediction model generates for the low-frequency factors is

[42]Even if a prediction model delivers a negative out-of-sample R-squared for predicting the returns of a factor, this does not imply that the prediction model has no predictability for the factor returns. In particular, Clark and West (2007) show that the difference between the MSPE of a parsimonious model (historical mean) and that of a larger model (PLS prediction model) has a negative center in a finite sample. Thus, the out-of-sample R-squared can be negative in a finite sample even if the larger PLS prediction model has additional predictive power for the factor returns than the parsimonious historical-mean model, and such additional predictability can potentially generates economic value for an investor. Therefore, examining whether a larger prediction model generates economic value is not only of economic interest but could also unveil the larger model's predictability masked by the negative out-of-sample R-squared.

to evaluate the performance of the following conditional mean-variance-efficient (CMVE) portfolio of the low-frequency factors (see, for example, Moreira and Muir, 2017):

$$c\widehat{\mathbb{E}}_t\left[f_{t+1}^{\text{LF}}\right]^\top\left[\widehat{\mathbb{V}ar}_t\left(f_{t+1}^{\text{LF}}\right)\right]^{-1}f_{t+1}^{\text{LF}}, \tag{44}$$

where c is a constant that governs the leverage of the portfolio, $f_{t+1}^{\text{LF}} = (f_{t+1}^{\text{B-BCF}}, f_{t+1}^{\text{BCF}})$ stacks the returns of the below-business-cycle- and business-cycle-frequency factors, and $\widehat{\mathbb{E}}_t[\cdot]$ and $\widehat{\mathbb{V}ar}_t(\cdot)$ denote the conditional mean and variance at month t implied by the prediction model.[43] To examine whether the PLS prediction model generates additional economic value for the low-frequency factors on top of the historical-mean benchmark, one first constructs the CMVE portfolios of the low-frequency factors based on the PLS prediction model and historical-mean model, and then regresses the returns of the former portfolio on those of the latter. If the regression intercept is positive, then it can be concluded that the PLS prediction model generates additional economic value for an investor on top of the historical-mean benchmark.

However, in Appendix E, I design a simple simulation to show that the aforementioned approach could generate misleading results, particularly when the factors have high unconditional risk premia. Specifically, the simulation results show that the returns of the CMVE portfolio based on the true prediction model could have a negative alpha with respect to those based on a model without predictive power. To alleviate this issue, I evaluate the economic value a prediction model generates for the low-frequency factors by exploiting the unconditional mean-variance-efficient (UMVE) portfolio of the low-frequency factors (see, for example, Ferson and Siegel, 2001; Chernov et al., 2022):

$$\frac{c\widehat{\mathbb{E}}_t\left[f_{t+1}^{\text{LF}}\right]^\top\left[\widehat{\mathbb{V}ar}_t\left(f_{t+1}^{\text{LF}}\right)\right]^{-1}}{1+\widehat{\mathbb{E}}_t\left[f_{t+1}^{\text{LF}}\right]^\top\left[\widehat{\mathbb{V}ar}_t\left(f_{t+1}^{\text{LF}}\right)\right]^{-1}\widehat{\mathbb{E}}_t\left[f_{t+1}^{\text{LF}}\right]}f_{t+1}^{\text{LF}}. \tag{45}$$

Note that the UMVE portfolio is both unconditionally efficient and conditionally efficient for every month t. To examine whether the PLS prediction model generates additional economic value for the low-frequency factors on top of the historical-mean benchmark, I first construct the UMVE portfolios of the low-frequency factors based on the PLS prediction model and historical-mean model, and then regress the returns of the former portfolio on those of the latter. I conclude that the PLS prediction model generates additional economic value for an investor on top of the historical-mean benchmark if the regression intercept is positive.[44]

Table 12 reports the intercept of these regressions. For ease of interpretation, I adjust the magnitude of c such that the UMVE portfolio of the two low-frequency factors based on the PLS prediction model (the left-hand side of the regression) is normalized to have the same unconditional standard deviation as the market. The first column gives the category of the predictor variables used in the PLS prediction model. The following ten columns give the intercept of the regression in which the low-frequency factors are estimated from K Basic-IPCA factors with K ranging between three and twelve. The last column gives the intercept of the regression in which the low-frequency factors are

[43] I compute the conditional variance of the low-frequency factors as the estimated variance of the residuals of a return-prediction model. In other words, I assume that the innovations of low-frequency factors are homoskedastic. I leave the volatility timing of the low-frequency factors to further work.

[44] The simulation results in Appendix E also show that this approach alleviates the issue associated with the approach that exploits the CMVE portfolios.

Table 12: Regression intercept: returns of UMVE portfolio of low-frequency factors based on PLS prediction with respect to those based on historical mean

This table reports the intercept (in percent) of regressing the monthly returns of the UMVE portfolio of the low-frequency factors based on the PLS prediction model on those based on the historical-mean model. The first column gives the category of the predictor variables used in the PLS prediction model. The following ten columns give the intercept of the regression in which the low-frequency factors are estimated from K Basic-IPCA factors with K ranging between three and twelve. The last column gives the intercept of the regression in which the low-frequency factors are estimated from all 37 Basic-IPCA factors. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Predictor type	Number of IPCA factors to estimate low-frequency factors										
	3	4	5	6	7	8	9	10	11	12	All
One PLS factor											
Macro	0.08	0.41***	0.26***	0.21***	0.33***	0.24***	0.22***	0.49***	0.45***	0.37***	0.11*
Financial	0.57**	0.60***	0.53***	0.25**	0.43***	0.27**	0.37***	0.63***	0.75***	0.75***	0.15
Macro+financial	0.46**	0.62***	0.59***	0.27***	0.53***	0.42***	0.35***	0.57***	0.60***	0.55***	0.26**
Number of PLS factors selected by BIC criterion											
Macro	0.12	0.54***	0.75***	0.50***	0.49***	0.44***	0.47***	0.58***	0.51***	0.50***	0.29***
Financial	0.74***	0.82***	0.50***	0.25***	0.45***	0.36***	0.48***	0.68***	0.86***	0.82***	0.14
Macro+financial	0.54**	0.65***	0.72***	0.31***	0.50***	0.49***	0.53***	0.66***	0.63***	0.56***	0.29**

estimated from all 37 Basic-IPCA factors. The table shows that the regression intercept is positive and economically sizable in all cases. For instance, by predicting the returns of the two low-frequency factors based on one PLS factor of all 43 economic variables, the UMVE portfolio achieves a monthly alpha of 0.26% with respect to its historical-mean counterpart. Moreover, the regression intercept is statistically significant at the 1% level in most cases. Therefore, I conclude that the PLS prediction model generates sizable and significant additional economic value for the low-frequency factors on top of the historical-mean benchmark.

In summary, I show strong evidence that the returns of the below-business-cycle- and business-cycle-frequency factors are highly predictable by the selected macroeconomic and financial variables. Moreover, I also find that the statistically significant predictability of the low-frequency factors is associated with sizable improved economic value for an investor. In particular, the unconditional mean-variance-efficient portfolio of the low-frequency factors based on the PLS prediction model has a significant alpha with respect to that based on the historical-mean prediction model.

4 Conclusion

Multiple sources of risks in addition to the aggregate market earn positive risk premia. In addition, prominent asset-pricing models imply that the systematic risk premia fluctuate over time, which is also supported by the empirical evidence of stock-return predictability. Thus, understanding the time-varying dynamics of systematic risk premia and how they relate to economic fundamentals is a doubly high-dimensional problem, with the systematic risk premia captured by a high-dimensional "factor zoo" and their time variation explained by a large number of potential predictors of asset returns.

This manuscript proposes an innovative approach to regularize this high-dimensional challenge by identifying *low-frequency* risk factors whose risk premia are driven by *latent low-frequency state variables*. On the one hand, identifying low-frequency factors regularizes the high-dimensional problem –"which factors' risk premia are time-varying?"– by reducing the dimension of factors whose risk premia

are driven by time-varying state variables. In particular, a low-frequency factor aggregates a latent low-frequency state variable's explanatory power for the risk premia of all candidate risk factors in the cross-section. On the other hand, identifying low-frequency risk factors also incorporates economic structure in the frequency domain into the state variables that drive systematic risk premia, and thus economically regularizes the second high-dimensional problem –"which economic variables explain the time-varying risk premia?"– because one expects that the economic variables whose variation concentrates on the low-frequency interval explain the risk premia of the low-frequency risk factors.

Empirically, I use the Instrumented Principal Component Analysis (IPCA) of Kelly et al. (2019) to obtain the candidate risk factors, from which I estimate two low-frequency risk factors: one below-business-cycle-frequency factor whose variation concentrates on cycles longer than eight years and one business-cycle-frequency factor whose variation concentrates on cycles between 1.5 and eight years. Both factors exhibit slow-moving fluctuations, suggesting that the risk premia of the cross-section have common sources of low-frequency time-series variation, and the two factors successfully reflect such low-frequency dynamics of systematic risk premia. The Sharpe ratio of the below-business-cycle-frequency factor is surprisingly high, which implies that a large fraction of the unconditional systematic risk premia of the cross-section is associated with time-series variation driven by a below-business-cycle-frequency state variable. Moreover, parsimonious models that include the two low-frequency factors outperform higher-dimensional IPCA models in explaining the average returns of individual stocks and characteristic-managed portfolios, suggesting that identifying low-frequency factors is a satisfactory method to empirically address the "factor zoo" problem. In addition, using the three-pass regression filter of Kelly and Pruitt (2015), a generalization of the partial least squares (PLS) regression, I find that 19 selected macroeconomic variables and 24 selected financial variables robustly predict the one-month-ahead returns of the two low-frequency factors in- and out-of-sample, implying that the selected economic variables highly explain the variation of the risk premia of the two factors. Moreover, the statistically significant predictability of the low-frequency risk factors also generates significant economic value for an investor. In particular, the unconditional mean-variance-efficient portfolio of the low-frequency factors based on the PLS prediction model has a statistically significant and economically sizable alpha with respect to that based on the benchmark historical-mean prediction model.

Identifying low-frequency factors regularizes the high-dimensional challenge of understanding the time-varying dynamics of systematic risk premia and how they relate to economic fundamentals. Thus, it sets the stage for constructing conditional asset-pricing models. Although this manuscript does not focus on constructing an SDF that prices stock returns conditionally, one can exploit the insight of this manuscript to achieve this objective in future research. For instance, this manuscript shows that the low-frequency factors mainly capture risk premia that are uncaptured by the market factor, which is perhaps the most important factor in asset pricing. Therefore, to construct a factor model with the objective of pricing stock returns conditionally, one can augment the low-frequency factors with more factors that capture additional information (such as the market factor) and then exploit the predictability of the low-frequency factors to approach conditional mean-variance efficiency. Moreover, this manuscript focuses on the time variation of risk premia instead of risk prices. Thus, to construct the conditional SDF, one could also apply different approaches of volatility timing on the low-frequency factors to obtain the dynamics of the risk prices that enter the SDF.

Moreover, this manuscript focuses on the performance of a few low-frequency factors whose variation concentrates on the below-business-cycle and business-cycle frequencies. To further explore the

time-varying dynamics of systematic risk premia, one can use the method proposed in this manuscript to identify additional low-frequency factors or factors whose variation concentrates on other frequency intervals. Furthermore, one can also exploit other prediction methods and/or a more comprehensive set of observable economic variables to obtain potentially better estimates of the risk premia of the low-frequency factors. In general, this manuscript provides a method to systematically study the time variation of systematic risk premia through the lens of frequency-specific factors, and thus sheds light on the construction of conditional asset-pricing models. I leave the further study of conditional factor models for future research.

References

Ang, Andrew, and Geert Bekaert, 2007, Stock return predictability: Is it there?, *Review of Financial Studies* 20, 651–707.

Asness, Clifford S., Jacques A. Friedman, Robert J. Krail, and John M. Liew, 2000, Style timing: Value versus growth, *Journal of Portfolio Management* 26, 50–60.

Baker, Malcolm, and Jeffrey Wurgler, 2006, Investor sentiment and the cross-section of stock returns, *Journal of Finance* 61, 1645–1680.

Bandi, Federico M., Shomesh E. Chaudhuri, Andrew W. Lo, and Andrea Tamoni, 2021, Spectral factor models, *Journal of Financial Economics* 142, 214–238.

Bandi, Federico M., Benoit Perron, Andrea Tamoni, and Claudio Tebaldi, 2019, The scale of predictability, *Journal of Econometrics* 208, 120–140.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.

Brandt, Michael W., and Qiang Kang, 2004, On the relationship between the conditional mean and volatility of stock returns: A latent var approach, *Journal of Financial Economics* 72, 217–257.

Breen, William, Lawrence R. Glosten, and Ravi Jagannathan, 1989, Economic significance of predictable variations in stock index returns, *Journal of Finance* 44, 1177–1189.

Brockwell, Peter J., and Richard A. Davis, 2009, *Time Series: Theory and Methods* (Springer Science & Business Media).

Campbell, John Y., 1987, Stock returns and the term structure, *Journal of Financial Economics* 18, 373–399.

Campbell, John Y., and John H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.

Campbell, John Y., and Ludger Hentschel, 1992, No news is good news: An asymmetric model of changing volatility in stock returns, *Journal of Financial Economics* 31, 281–318.

Campbell, John Y., Andrew Wen-Chuan Lo, and Archie Craig MacKinlay, 1997, *The Econometrics of Financial Markets* (Princeton University Press, Princeton, New Jersey).

Campbell, John Y., and Robert J. Shiller, 1988a, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.

Campbell, John Y., and Robert J. Shiller, 1988b, Stock prices, earnings, and expected dividends, *Journal of Finance* 43, 661–676.

Campbell, John Y., and Samuel B. Thompson, 2008, Predicting excess stock returns out of sample: Can anything beat the historical average?, *Review of Financial Studies* 21, 1509–1531.

Campbell, John Y., and Tuomo Vuolteenaho, 2004, Inflation illusion and stock prices, *American Economic Review* 94, 19–23.

Campbell, Sean D., and Francis X. Diebold, 2009, Stock returns and expected business conditions: Half a century of direct evidence, *Journal of Business & Economic Statistics* 27, 266–278.

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

Chernov, Mikhail, Lars A. Lochstoer, and Stig R.H. Lundeby, 2022, Conditional dynamics and the multihorizon risk-return trade-off, *Review of Financial Studies* 35, 1310–1347.

Chinco, Alex, Adam D. Clark-Joseph, and Mao Ye, 2019, Sparse signals in the cross-section of returns, *Journal of Finance* 74, 449–492.

Clark, Todd E., and Kenneth D. West, 2007, Approximately normal tests for equal predictive accuracy in nested models, *Journal of Econometrics* 138, 291–311.

Cochrane, John H., 2008, The dog that did not bark: A defense of return predictability, *Review of Financial Studies* 21, 1533–1575.

Cochrane, John H., 2011, Presidential address: Discount rates, *Journal of Finance* 66, 1047–1108.

Cohen, Randolph B., Christopher Polk, and Tuomo Vuolteenaho, 2003, The value spread, *Journal of Finance* 58, 609–641.

Cooper, Michael J., Roberto C. Gutierrez Jr, and Allaudeen Hameed, 2004, Market states and momentum, *Journal of Finance* 59, 1345–1365.

Dew-Becker, Ian, and Stefano Giglio, 2016, Asset pricing in the frequency domain: Theory and empirics, *Review of Financial Studies* 29, 2029–2068.

Dong, Xi, Yan Li, David E. Rapach, and Guofu Zhou, 2022, Anomalies and the expected market return, *Journal of Finance* 77, 639–681.

Ehsani, Sina, and Juhani T. Linnainmaa, 2022, Factor momentum and the momentum factor, *Journal of Finance* 77, 1877–1919.

Fama, Eugene F., and Kenneth R. French, 1988, Dividend yields and expected stock returns, *Journal of Financial Economics* 22, 3–25.

Fama, Eugene F., and Kenneth R. French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23–49.

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

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, 2018, Choosing factors, *Journal of Financial Economics* 128, 234–252.

Ferson, Wayne E., Sergei Sarkissian, and Timothy T. Simin, 2003, Spurious regressions in financial economics?, *Journal of Finance* 58, 1393–1413.

Ferson, Wayne E., and Andrew F. Siegel, 2001, The efficient use of conditioning information in portfolios, *Journal of Finance* 56, 967–982.

French, Kenneth R., G. William Schwert, and Robert F. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–29.

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

Glosten, Lawrence R., Ravi Jagannathan, and David E. Runkle, 1993, On the relation between the expected value and the volatility of the nominal excess return on stocks, *Journal of Finance* 48, 1779–1801.

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

Haddad, Valentin, Serhiy Kozak, and Shrihari Santosh, 2020, Factor timing, *Review of Financial Studies* 33, 1980–2018.

Harvey, Campbell R., 1989, Time-varying conditional covariances in tests of asset pricing models, *Journal of Financial Economics* 24, 289–317.

Harvey, Campbell R., Yan Liu, and Heqing Zhu, 2016, ... and the cross-section of expected returns, *Review of Financial Studies* 29, 5–68.

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

He, Zhiguo, and Arvind Krishnamurthy, 2012, A model of capital and crises, *Review of Economic Studies* 79, 735–777.

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

Huang, Dashan, Jiangyuan Li, and Liyao Wang, 2021, Are disagreements agreeable? Evidence from information aggregation, *Journal of Financial Economics* 141, 83–101.

Huang, Jiantao, 2021, Frequency dependent risks in the factor zoo, *Available at SSRN* 3948519.

Jurado, Kyle, Sydney C. Ludvigson, and Serena Ng, 2015, Measuring uncertainty, *American Economic Review* 105, 1177–1216.

Kelly, Bryan, and Seth Pruitt, 2013, Market expectations in the cross-section of present values, *Journal of Finance* 68, 1721–1756.

Kelly, Bryan, and Seth Pruitt, 2015, The three-pass regression filter: A new approach to forecasting using many predictors, *Journal of Econometrics* 186, 294–316.

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.

Kelly, Bryan, Seth Pruitt, and Yinan Su, 2020, Instrumented principal component analysis, *Available at SSRN* 2983919.

Koijen, Ralph S.J., and Stijn Van Nieuwerburgh, 2011, Predictability of returns and cash flows, *Annual Review of Financial Economics* 3, 467–491.

Kothari, Smitu P., and Jay Shanken, 1997, Book-to-market, dividend yield, and expected market returns: A time-series analysis, *Journal of Financial Economics* 44, 169–203.

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.

Lettau, Martin, and Sydney Ludvigson, 2001a, Resurrecting the (C)CAPM: A cross-sectional test when risk premia are time-varying, *Journal of Political Economy* 109, 1238–1287.

Lettau, Martin, and Sydney C. Ludvigson, 2001b, Consumption, aggregate wealth, and expected stock returns, *Journal of Finance* 56, 815–849.

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

Ludvigson, Sydney C., Sai Ma, and Serena Ng, 2021, Uncertainty and business cycles: Exogenous impulse or endogenous response?, *American Economic Journal: Macroeconomics* 13, 369–410.

Ludvigson, Sydney C, and Serena Ng, 2007, The empirical risk–return relation: A factor analysis approach, *Journal of financial economics* 83, 171–222.

Moreira, Alan, and Tyler Muir, 2017, Volatility-managed portfolios, *Journal of Finance* 72, 1611–1644.

Nelson, Charles R., 1976, Inflation and rates of return on common stocks, *Journal of Finance* 31, 471–483.

Neuhierl, Andreas, Otto Randl, Christoph Reschenhofer, and Josef Zechner, 2023, Timing the factor zoo, *Available at SSRN* 4376898.

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

Pastor, Lubos, and Robert F. Stambaugh, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy* 111, 642–685.

Rapach, David E., and Guofu Zhou, 2022, Asset pricing: Time-series predictability, in *Oxford Research Encyclopedia of Economics and Finance*.

Ross, Stephen, 1976, The arbitrage theory of capital asset pricing, *Journal of Economic Theory* 13, 341–360.

Ross, Stephen A., 2005, *Neoclassical Finance* (Princeton University Press, Princeton, New Jersey).

Shiller, Robert J., 1981, Do stock prices move too much to be justified by subsequent changes in dividends?, *American Economic Review* 71, 421–436.

Stambaugh, Robert F., 1999, Predictive regressions, *Journal of Financial Economics* 54, 375–421.

Stambaugh, Robert F., Jianfeng Yu, and Yu Yuan, 2012, The short of it: Investor sentiment and anomalies, *Journal of Financial Economics* 104, 288–302.

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

Wold, Herman, 1966, Estimation of principal components and related models by iterative least squares, in *P. R. Krishnaiah, ed.: Multivariate Analysis* (Academic Press, New York).

Zhou, Guofu, 2010, How much stock return predictability can we expect from an asset pricing model?, *Economics Letters* 108, 184–186.

A Proofs of all results

A.1 Proof of Proposition 1

Since x_t and ξ_{t+1} are orthogonal, I have $\mathbf{I}_K = \Sigma_f = \phi\Sigma_x\phi^\top + \Sigma_\xi$, and thus $\Sigma_\xi = \mathbf{I}_K - \phi\Sigma_x\phi^\top$. Therefore,

$$\begin{aligned}\mu_f^\top \Sigma_\xi^{-1}\mu_f &= \mu_f^\top \left(\mathbf{I}_K - \phi\Sigma_x\phi^\top\right)^{-1}\mu_f \\ &= \mu_f^\top \left[\mathbf{I}_K - \phi\left(\mathbf{I}_J - \Sigma_x^{-1}\right)^{-1}\phi^\top\right]\mu_f \\ &= \mu_f^\top\mu_f - \mu_f^\top\phi\left(\mathbf{I}_J - \Sigma_x^{-1}\right)^{-1}\phi^\top\mu_f,\end{aligned}$$

where the second equality is derived from the Woodbury identity and $\phi^\top\phi = \mathbf{I}_J$. Similarly,

$$\begin{aligned}\mu_f^\top\phi\left[\phi^\top\Sigma_\xi\phi\right]^{-1}\phi^\top\mu_f &= \mu_f^\top\phi\left(\mathbf{I}_J - \Sigma_x\right)^{-1}\phi^\top\mu_f \\ &= \mu_f^\top\phi\left[\mathbf{I}_J - \left(\mathbf{I}_J - \Sigma_x^{-1}\right)^{-1}\right]\phi^\top\mu_f \\ &= \mu_f^\top\phi\phi^\top\mu_f - \mu_f^\top\phi\left(\mathbf{I}_J - \Sigma_x^{-1}\right)^{-1}\phi^\top\mu_f.\end{aligned}$$

Let $c = -\mu_f^\top\phi\left(\mathbf{I}_J - \Sigma_x^{-1}\right)^{-1}\phi^\top\mu_f > 0$ because $\mathbf{I}_J - \Sigma_x^{-1}$ is negative definite, I have

$$\frac{\mu_f^\top\phi\left[\phi^\top\Sigma_\xi\phi\right]^{-1}\phi^\top\mu_f}{\mu_f^\top\Sigma_\xi^{-1}\mu_f} = \frac{\mu_f^\top\phi\phi^\top\mu_f + c}{\mu_f^\top\mu_f + c} \geq \frac{\mu_f^\top\phi\phi^\top\mu_f}{\mu_f^\top\mu_f}, \tag{46}$$

where the last inequality holds because $\mu_f^\top\phi\phi^\top\mu_f \leq \mu_f^\top\mu_f$. In particular, when $\mu_f \in \text{span}(\phi)$, I have $\mu_f^\top\phi\phi^\top\mu_f = \mu_f^\top\mu_f$, and thus (46) holds with equality, and each side of (46) is equal to one. This completes the proof.

A.2 Proof of Lemma 1

Note that

$$\frac{\eta^\top\Sigma_x^{\text{LF}}\eta}{\eta^\top\Sigma_x^{\text{HF}}\eta} > \frac{\omega}{\pi-\omega} \;\Leftrightarrow\; \frac{\eta^\top\frac{\pi}{\omega}\Sigma_x^{\text{LF}}\eta}{\eta^\top\frac{\pi}{\pi-\omega}\Sigma_x^{\text{HF}}\eta} > 1 \;\Leftrightarrow\; \frac{\eta^\top\left(\frac{\pi}{\omega}\Sigma_x^{\text{LF}} + \mathbf{I}_J - \Sigma_x\right)\eta}{\eta^\top\left(\frac{\pi}{\pi-\omega}\Sigma_x^{\text{HF}} + \mathbf{I}_J - \Sigma_x\right)\eta} > 1, \tag{47}$$

where the second equivalence holds because

$$\mathbf{I}_J - \Sigma_x = \mathbb{V}ar(\phi^\top f_{t+1}) - \Sigma_x = \phi^\top\Sigma_\xi\phi \tag{48}$$

is positive definite, and thus $\eta^\top(\mathbf{I}_J - \Sigma_x)\eta > 0$. Furthermore, by letting

$$\begin{aligned}v &= \left(\frac{\pi}{\pi-\omega}\Sigma_x^{\text{HF}} + \mathbf{I}_J - \Sigma_x\right)^{1/2}\eta, \quad \text{and} \\ W &= \left(\frac{\pi}{\pi-\omega}\Sigma_x^{\text{HF}} + \mathbf{I}_J - \Sigma_x\right)^{-1/2}\left(\frac{\pi}{\omega}\Sigma_x^{\text{LF}} + \mathbf{I}_J - \Sigma_x\right)\left(\frac{\pi}{\pi-\omega}\Sigma_x^{\text{HF}} + \mathbf{I}_J - \Sigma_x\right)^{-1/2},\end{aligned}$$

I can write (47) as

$$\frac{\eta^\top\Sigma_x^{\text{LF}}\eta}{\eta^\top\Sigma_x^{\text{HF}}\eta} > \frac{\omega}{\pi-\omega} \;\Leftrightarrow\; \frac{v^\top W v}{v^\top v} > 1, \tag{49}$$

which completes the proof.

A.3 Proof of Proposition 2

Since constraint (22) is binding at the optimum, I have

$$\begin{aligned} & \theta^\top \left[\frac{\omega}{\pi}\Sigma_\xi + \phi\Sigma_x^{\text{LF}}\phi^\top\right]\theta - c\theta^\top\left[\frac{\pi-\omega}{\pi}\Sigma_\xi + \phi\Sigma_x^{\text{HF}}\phi^\top\right]\theta = 0 \\ \Rightarrow \quad & c = \frac{\omega}{\pi-\omega}\frac{\theta^\top\left(\mathbf{I}_K - \phi\Sigma_x\phi^\top + \frac{\pi}{\omega}\phi\Sigma_x^{\text{LF}}\phi^\top\right)\theta}{\theta^\top\left(\mathbf{I}_K - \phi\Sigma_x\phi^\top + \frac{\pi}{\pi-\omega}\phi\Sigma_x^{\text{HF}}\phi^\top\right)\theta}, \end{aligned} \tag{50}$$

where I exploit

$$\begin{aligned} \Sigma_f^{\text{LF}} &= \phi\Sigma_x^{\text{LF}}\phi^\top + \frac{\omega}{\pi}\Sigma_\xi \quad \text{and} \\ \Sigma_f^{\text{HF}} &= \phi\Sigma_x^{\text{HF}}\phi^\top + \frac{\pi-\omega}{\pi}\Sigma_\xi \end{aligned}$$

in equations (18) and (19) and the equation $\Sigma_f = \mathbf{I}_K = \phi\Sigma_x\phi^\top + \Sigma_\xi$. Therefore, solving problem (21) is equivalent to solving for θ that maximizes c. Since the columns of $\phi \in \mathbb{R}^{K\times J}$ are orthonormal, we can write

$$\mathbf{I}_K = \Lambda\Lambda^\top + \phi\phi^\top,$$

where the columns of $\Lambda \in \mathbb{R}^{K\times(K-J)}$ are orthonormal and orthogonal to ϕ. As a result, equation (50) can be written as

$$c = \frac{\omega}{\pi-\omega}\frac{\theta^\top\begin{bmatrix}\phi, & \Lambda\end{bmatrix}\begin{bmatrix}\mathbf{I}_J - \Sigma_x + \frac{\pi}{\omega}\Sigma_x^{\text{LF}} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{K-J}\end{bmatrix}\begin{bmatrix}\phi^\top \\ \Lambda^\top\end{bmatrix}\theta}{\theta^\top\begin{bmatrix}\phi, & \Lambda\end{bmatrix}\begin{bmatrix}\mathbf{I}_J - \Sigma_x + \frac{\pi}{\pi-\omega}\Sigma_x^{\text{HF}} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{K-J}\end{bmatrix}\begin{bmatrix}\phi^\top \\ \Lambda^\top\end{bmatrix}\theta}. \tag{51}$$

Let

$$z = \begin{bmatrix}\phi, & \Lambda\end{bmatrix}\begin{bmatrix}\left(\mathbf{I}_J - \Sigma_x + \frac{\pi}{\pi-\omega}\Sigma_x^{\text{HF}}\right)^{1/2} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{K-J}\end{bmatrix}\begin{bmatrix}\phi^\top \\ \Lambda^\top\end{bmatrix}\theta,$$

and thus

$$\theta = \begin{bmatrix}\phi, & \Lambda\end{bmatrix}\begin{bmatrix}\left(\mathbf{I}_J - \Sigma_x + \frac{\pi}{\pi-\omega}\Sigma_x^{\text{HF}}\right)^{-1/2} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{K-J}\end{bmatrix}\begin{bmatrix}\phi^\top \\ \Lambda^\top\end{bmatrix}z.$$

In this case, equation (51) can be written as

$$c = \frac{\omega}{\pi-\omega}\frac{z^\top\begin{bmatrix}\phi, & \Lambda\end{bmatrix}\begin{bmatrix}W & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{K-J}\end{bmatrix}\begin{bmatrix}\phi^\top \\ \Lambda^\top\end{bmatrix}z}{z^\top z}, \tag{52}$$

where

$$W = \left(\mathbf{I}_J - \Sigma_x + \frac{\pi}{\pi-\omega}\Sigma_x^{\text{HF}}\right)^{-1/2}\left(\mathbf{I}_J - \Sigma_x + \frac{\pi}{\omega}\Sigma_x^{\text{LF}}\right)\left(\mathbf{I}_J - \Sigma_x + \frac{\pi}{\pi-\omega}\Sigma_x^{\text{HF}}\right)^{-1/2}.$$

Apply eigenvalue decomposition to W:

$$W = HD_W H^\top,$$

where $D_W \in \mathbb{R}^{J\times J}$ is a diagonal matrix whose main diagonal contains the eigenvalues of W in descending order, and $H \in \mathbb{R}^{J\times J}$ contains the orthonormal eigenvectors of the corresponding eigenvalues. In this case, equation (52) can be written as

$$c = \frac{\omega}{\pi-\omega}\frac{z^\top \begin{bmatrix}\phi H, & \Lambda\end{bmatrix}\begin{bmatrix}D_W & \mathbf{0}\\ \mathbf{0} & \mathbf{I}_{K-J}\end{bmatrix}\begin{bmatrix}H^\top\phi^\top\\ \Lambda^\top\end{bmatrix} z}{z^\top z}. \tag{53}$$

If the maximum eigenvalue of W, λ_1^W, is greater than one, then z that maximizes c is (proportional to) ϕH_1, where H_1 is the first eigenvector of W, that is, the first column of H. Therefore, I have

$$\begin{aligned}
c^* &= \frac{\omega}{\pi-\omega}\lambda_1^W,\\
\theta^* &\propto \begin{bmatrix}\phi, & \Lambda\end{bmatrix}\begin{bmatrix}\left(\mathbf{I}_J - \Sigma_x + \frac{\pi}{\pi-\omega}\Sigma_x^{\text{HF}}\right)^{-1/2} & \mathbf{0}\\ \mathbf{0} & \mathbf{I}_{K-J}\end{bmatrix}\begin{bmatrix}\phi^\top\\ \Lambda^\top\end{bmatrix}\phi H_1\\
&= \phi\left(\mathbf{I}_J - \Sigma_x + \frac{\pi}{\pi-\omega}\Sigma_x^{\text{HF}}\right)^{-1/2} H_1.
\end{aligned}$$

If λ_1^W is lower than one, then z that maximizes c is (proportional to) Λx, where $x \in \mathbb{R}^{K-J}$ is an arbitrary vector. In this case, we have

$$\begin{aligned}
c^* &= \frac{\omega}{\pi-\omega},\\
\theta^* &\propto \begin{bmatrix}\phi, & \Lambda\end{bmatrix}\begin{bmatrix}\left(\mathbf{I}_J - \Sigma_x + \frac{\pi}{\pi-\omega}\Sigma_x^{\text{HF}}\right)^{-1/2} & \mathbf{0}\\ \mathbf{0} & \mathbf{I}_{K-J}\end{bmatrix}\begin{bmatrix}\phi^\top\\ \Lambda^\top\end{bmatrix}\Lambda x = \Lambda x.
\end{aligned}$$

A.4 Proof of Proposition 3

I first prove Part (a) of the proposition. Note that

$$f_{t+1}^\perp = \left(\mathbf{I}_K - \frac{\phi G_1 G_1^\top \phi^\top}{G_1^\top G_1}\right) f_{t+1}. \tag{54}$$

Therefore, the state-variable component of $f_{t+1}^\perp$ is

$$\begin{aligned}
\left(\mathbf{I}_K - \frac{\phi G_1 G_1^\top \phi^\top}{G_1^\top G_1}\right)\phi x_t &= \phi\left(\mathbf{I}_J - \frac{G_1 G_1^\top}{G_1^\top G_1}\right)x_t\\
&= \phi V V^\top x_t \equiv \phi^\perp x_t^\perp,
\end{aligned} \tag{55}$$

where V is a $J\times(J-1)$ matrix whose columns are orthonormal and are orthogonal to G_1, $x_t^\perp = V^\top x_t$, and $\phi^\perp = \phi V$. To see that $f_{t+1}^\perp$ has zero loading on x_t^1, note that equation (55) can be written as

$$\phi V V^\top \begin{bmatrix}\frac{G_1}{\|G_1\|} & V\end{bmatrix}\begin{bmatrix}\frac{G_1^\top}{\|G_1^\top\|}\\ V^\top\end{bmatrix} x_t$$

$$= \frac{1}{\|G_1\|^2}\phi V \underbrace{V^\top G_1}_{=\mathbf{0}_{J-1}} G_1^\top x_t + \phi V V^\top x_t,$$

which illustrates that $f_{t+1}^\perp$ does not load on x_t^1.

I now prove Part (b) of the proposition. Let U be a $K \times (K-1)$ matrix whose columns are orthonormal and are orthogonal to ϕG_1. In this case, equation (54) can be written as

$$f_{t+1}^\perp = UU^\top f_{t+1}.$$

Note that the loadings of $U^\top f_{t+1}^\perp$ on $x_t^\perp$ is $U^\top \phi^\perp = U^\top \phi V$. I first show that the normalizations

$$\mathbb{V}ar(U^\top f_{t+1}^\perp) = \mathbf{I}_{K-1} \quad \text{and} \quad \left(U^\top \phi^\perp\right)^\top U^\top \phi^\perp = \mathbf{I}_{J-1}$$

hold. Specifically, since $U^\top f_{t+1}^\perp = U^\top f_{t+1}$, and thus $\mathbb{V}ar(U^\top f_{t+1}) = U^\top U = \mathbf{I}_{K-1}$ holds. In addition,

$$\begin{aligned}
\left(U^\top \phi^\perp\right)^\top U^\top \phi^\perp &= V^\top \phi^\top U U^\top \phi V \\
&= V^\top \phi^\top \left(\mathbf{I}_K - \frac{\phi G_1 G_1^\top \phi^\top}{G_1^\top G_1}\right) \phi V \\
&= V^\top V + \frac{V^\top G_1 G_1^\top V}{G_1^\top G_1} \\
&= V^\top V = \mathbf{I}_{J-1},
\end{aligned}$$

where the fourth equality holds because the columns of V are orthogonal to G_1. Therefore, the solution to problem (21) by letting the risk factors be $U^\top f_{t+1}^\perp$ follows Proposition 2. If there exists η' such that $(\eta')^\top x_t^\perp$ is a low-frequency state variable, Proposition 2 implies that the optimal θ in this case is $U^\top \phi V G_2$, where $G_2 \in \mathbb{R}^{J-1}$. In this case, I have

$$UU^\top \phi V G_2 = \phi \left(\mathbf{I}_J - \frac{G_1 G_1^\top}{G_1^\top G_1}\right) V G_2 = \phi V G_2, \tag{56}$$

which is spanned by the columns of ϕ and is orthogonal to ϕG_1. This completes the proof.

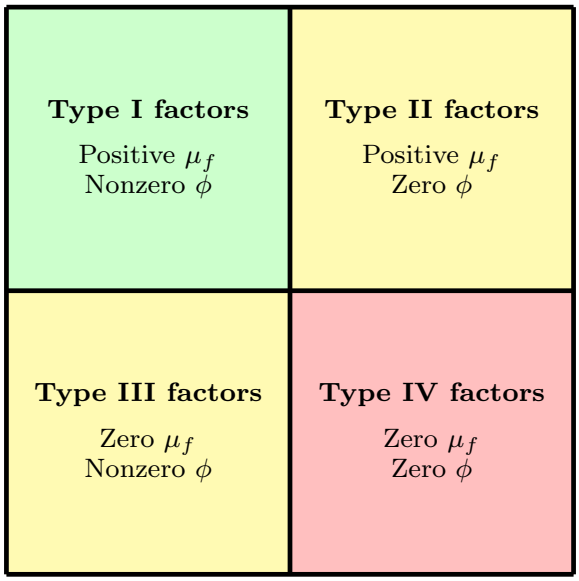
B Complement popular methods to address the "factor zoo" problem

In this section, I describe how identifying the loadings of the candidate risk factors on the state variables, that is, ϕ in equation (2), complements popular methods that focus on the unconditional factor risk premia to address the "factor zoo" problem and identify a parsimonious factor model that explains the cross-section of expected returns.

To see this, I first conceptually classify the candidate factors in the "factor zoo" into four categories, as illustrated in Figure 7. Type I factors earn positive unconditional risk premia, and their risk premia are also associated with time-series variation. Type II factors earn positive unconditional risk premia, but their risk premia are constant over time. Type III factors have zero unconditional risk premia, but their risk premia have time-series variation, suggesting that their conditional risk premia are, in general, non-zero. Type IV factors have zero unconditional and conditional risk premia. It is

Figure 7: Classification of candidate risk factors

This figure illustrates a conceptual classification of the "zoo" of candidate factors. Type I factors earn positive unconditional risk premia, and their risk premia also have time-series variation. Type II factors earn positive unconditional risk premia, but their risk premia are constant over time. Type III factors have zero unconditional risk premia, but their risk premia have time-series variation, suggesting that their conditional risk premia are, in general, non-zero. Type IV factors have zero unconditional and conditional risk premia.



easy to find that if the four types of factors have orthogonal factor innovations, an asset-pricing model that prices asset returns conditionally needs to include factors in types I, II, and III.

Popular methods that aim to address the "factor zoo" problem focus on capturing factors with positive unconditional risk premia. In other words, these methods focus on capturing factors in types I and II. For example, Kozak, Nagel, and Santosh (2020) consider a panel of characteristic-managed portfolios as the candidate risk factors, and shrink the contributions of the higher-order principal components of candidate factors to the SDF using an elastic-net estimator. The economic insight of this approach is based on the tenet of the arbitrage pricing theory (APT) of Ross (1976), in which the systematic risk factors that drive the time-series variation of the returns (lower-order principal components) should also capture the unconditional risk premia information to explain the cross-section of expected returns. The instrumented PCA (IPCA) of Kelly et al. (2019) exploits a similar insight, and their IPCA factors are exactly the principal components of characteristic-managed portfolios if the cross-sectional correlation matrix of stock characteristics is constant. Lettau and Pelger (2020) exploit the insight of APT more explicitly. In particular, their risk-premium-PCA (RP-PCA) approach generalizes PCA and explicitly increases the "signal strength" of candidate risk factors with positive unconditional risk premia. In summary, one common insight of these approaches to address the "factor zoo" problem is capturing the components of candidate risk factors with positive unconditional risk premia. Correspondingly, these methods do not aim to capture the type III factors that are informative to price asset returns conditionally.

By identifying the loadings of the candidate risk factors on the state variables, one recovers the factors with time-varying risk premia in equation (4), and thus captures the information of factors in types I and III. Therefore, identifying ϕ complements the popular methods to address the "factor zoo"

problem because it identifies the type III factors, which are hard-to-capture for the popular methods that aim to capture factors with positive unconditional risk premia. Specifically, to construct an asset-pricing model from the "zoo" of candidate factors, one method that uses the two complementing approaches simultaneously includes the following three steps:

1. Identify ϕ to obtain the type-I and type-III factors of the candidate factors;
2. Regress the factors with time-varying risk premia (identified in Step 1) out of all candidate factors and obtain the regression residuals;
3. Apply methods that focus on capturing factors with positive unconditional risk premia, such as PCA or RP-PCA, to the regression residuals to capture the type II factors.

C Estimation of frequency-specific components of covariance matrix

Given a series of factor returns $f_1, \ldots, f_T$,[45] the discrete Fourier transform (DFT) of the factor returns at frequency $\omega_l = 2\pi l/T$ (for $l = \lfloor T/2 \rfloor - T + 1, \ldots, \lfloor T/2 \rfloor$) is

$$\mathbf{d}_f(\omega_l) = \frac{1}{\sqrt{T}} \sum_{t=1}^{T} f_t e^{-i\omega_l t}. \tag{57}$$

Note that each ω_l lies in the entire frequency interval $[-\pi, \pi]$. The inverse discrete Fourier transform (IDFT) allows me to represent the sample of factor returns in their DFTs:

$$f_t = \frac{1}{\sqrt{T}} \sum_{l=\lfloor T/2 \rfloor - T+1}^{\lfloor T/2 \rfloor} \mathbf{d}_f(\omega_l) e^{i\omega_l t}. \tag{58}$$

Two important properties of the DFTs when T approaches infinity are that (i) they are uncorrelated across frequencies, and (ii) the variance of each DFT is the spectral density matrix of the factors. In other words, I have

$$\lim_{T\to\infty} \mathbb{C}ov\left(\mathbf{d}_f(\omega_l), \mathbf{d}_f(\omega_j)\right) = \mathbf{0}, \quad \text{if } l \neq j, \text{ and}$$
$$\lim_{T\to\infty} \mathbb{V}ar\left(\mathbf{d}_f(\omega_l)\right) = \mathbf{g}_f(\omega_l).$$

Therefore, equation (58) suggests that the DFTs decompose the factor returns into orthogonal frequency-specific components.

Moreover, note that

$$\mathbf{d}_f(\omega_l)\, \mathbf{d}_f^*(\omega_l) = \sum_{k=-(T-1)}^{T-1} \widehat{\Sigma}_f(k) e^{-i\omega_l k}, \tag{59}$$

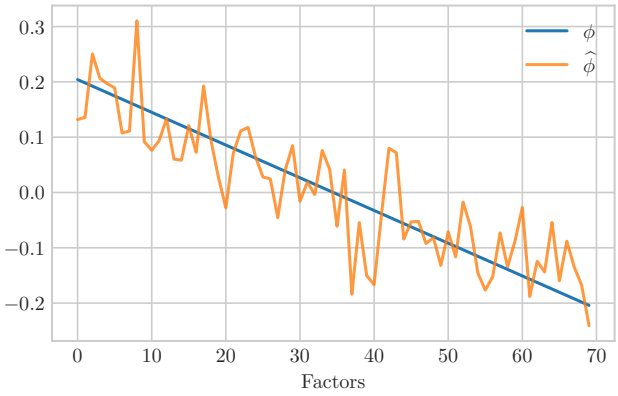
where $\widehat{\Sigma}_f(k) = T^{-1} \sum_{t=1}^{T-k} f_{t+k} f_t^\top$ for $k \geq 0$ and $\widehat{\Sigma}_f(k) = T^{-1} \sum_{t=1}^{T-k} f_t f_{t-k}^\top$ for $k < 0$, and $\mathbf{d}_f^*(\omega_l)$ is the conjugate transpose of $\mathbf{d}_f(\omega_l)$. Therefore, equation (59) shows that $\mathbf{d}_f(\omega_l)\, \mathbf{d}_f^*(\omega_l)$ is an estimator

[45] Without loss of generality, in this section, I assume that the factor returns have zero mean. This assumption is only made for illustrative purposes and can be easily relaxed with more cumbersome notations.

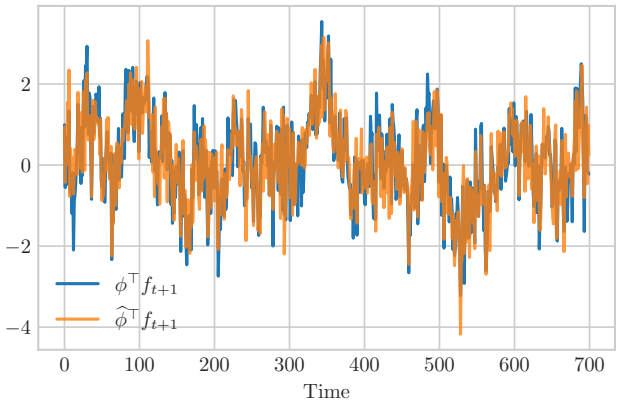
Figure 8: Simulation: loading vector and factor with time-varying risk premia

Panels A and B illustrate the loadings of the factors on the state variable and the returns of the factor with time-varying risk premia defined in equation (30), respectively. In each panel, the blue line depicts the object associated with the population loadings of the factors on the state variable, and the orange line depicts the object associated with the estimated loadings of the factors on the state variable, which are obtained by solving problem (21) using the frequency-specific covariance matrices estimated based on the simulated data.

(A) Loading vector



(B) Factor with time-varying risk premia



of the spectral density matrix $\mathbf{g}_f(\omega_l)$ of the factor returns defined in (9). Empirically, I exploit this estimator of the spectral density matrix to estimate the frequency-specific components of the covariance matrix of factor returns. Specifically, given T observations of the factor returns, I first estimate their DFTs $\widehat{\mathbf{d}}_f(\omega_l)$ using equation (57), and then estimate the low-frequency, business-cycle-frequency, and high-frequency components of the covariance matrix as

$$\begin{aligned}
\widehat{\Sigma}_f^{\text{LF}} &= \frac{1}{T}\sum\nolimits_{\omega_l \in [-\frac{2\pi}{96}, \frac{2\pi}{96}]} \widehat{\mathbf{d}}_f(\omega_l)\widehat{\mathbf{d}}_f^*(\omega_l), \\
\widehat{\Sigma}_f^{\text{BCF}} &= \frac{1}{T}\sum\nolimits_{\omega_l \in [-\frac{2\pi}{18}, -\frac{2\pi}{96}) \cup (\frac{2\pi}{96}, \frac{2\pi}{18}]} \widehat{\mathbf{d}}_f(\omega_l)\widehat{\mathbf{d}}_f^*(\omega_l), \qquad \text{and} \\
\widehat{\Sigma}_f^{\text{HF}} &= \frac{1}{T}\sum\nolimits_{\omega_l \in [-\pi, -\frac{2\pi}{6}] \cup [\frac{2\pi}{6}, \pi]} \widehat{\mathbf{d}}_f(\omega_l)\widehat{\mathbf{d}}_f^*(\omega_l),
\end{aligned}$$

which are sample counterparts of equations (33), (34), and (35), respectively.

D Simulation of Proposition 2

In this section, I design a simple simulation to illustrate the performance of my approach in Section 2.3. Specifically, this simulation illustrates that, under the empirical setup of this manuscript, the solution to problem (21) (given in Proposition 2) does not suffer from substantial estimation error in a finite sample that mimics the real-world data.

I assume that there are $K = 70$ candidate risk factors and there is $J = 1$ state variable that drives the risk premia variation of all factors. The state variable follows an AR(1) process with variance $\sigma_x^2 = 0.7$ and autoregressive parameter 0.9. The loadings of the 70 risk factors on the state variables ϕ has a linear structure centered at zero. In other words, the loading of the first factor on the state variable is d, the loading of the 70th factor is $-d$, and the loadings of the remaining 68 factors follow a monotonic linear pattern between d and $-d$. Following the main body of the manuscript, I normalize

Table 13: Correlation between population and estimated values in simulation

The first row of this table reports the quantiles of the cross-sectional correlation between the estimated and population loadings of the factors on the state variable. The second row reports the quantiles of the correlation between the factor with time-varying risk premia based on the estimated loading vector and its population counterpart. Each row reports the 0.05, 0.25, 0.50, 0.75, and 0.95 quantiles of the correlation in 1,000 simulation runs.

Quantiles	0.05	0.25	0.50	0.75	0.95
$\mathbb{C}orr\big(\phi, \widehat{\phi}\big)$	68.9%	79.0%	83.5%	87.0%	90.7%
$\mathbb{C}orr\big(\phi^\top f_{t+1}, \widehat{\phi}^\top f_{t+1}\big)$	68.7%	79.3%	84.3%	88.2%	91.7%

the covariance matrix of the factors as an identity matrix and the loading vector ϕ to have norm one. Thus, the state variable explains $\text{trace}(\sigma_x^2 \phi\phi^\top)/\text{trace}(\mathbf{I}_{70}) = 1\%$ of the total variation of the factors. I simulate the system for $T = 700$ periods.

Following the same procedure as the empirical analysis of the main body of the manuscript, I focus on the low-frequency interval $[-2\pi/96, 2\pi/96]$ in the simulation, and solve problem (21) using the frequency-specific covariance matrices estimated from the simulated data.[46] Since the only state variable in the simulation is persistent and thus is a low-frequency state variable, Proposition 2 suggests that the solution to problem (21) should identify the loadings of the factors on this state variable.

Panels A and B of Figure 8 illustrate the loadings of the factors on the state variable and the returns of the factor with time-varying risk premium defined in equation (30), respectively. In each panel, the blue line depicts the object associated with the population loadings of the factors on the state variable, and the orange line depicts the object associated with the estimated loadings of the factors on the state variable, which are obtained by solving problem (21) using the frequency-specific covariance matrices estimated based on the simulated data. Panel A shows that solving problem (21) using the estimated frequency-specific covariance matrices accurately identifies the loadings of the factors on the state variable. In particular, the estimated loadings capture well the monotonic pattern of the population loadings, and the cross-sectional correlation between the population loadings and estimated loadings is 85.7%. As a result, Panel B shows that I can accurately estimate the factor with time-varying risk premia. Specifically, the correlation between the factor with time-varying risk premia based on the estimated loading vector and its population counterpart is 88.1%.

To further investigate the performance of my approach in finite sample, I repeat this simulation for 1,000 times. The first row Table 13 reports the quantiles of the cross-sectional correlation between the estimated and population loadings of the factors on the state variable in the 1,000 simulated systems. I find that the solving problem (21) based on the estimated frequency-specific covariance matrices accurately recovers the loadings of the factors on the state variable. In particular, the cross-sectional correlation between the estimated loadings and their population counterparts is higher than 68.9% in 95% of the simulated systems. In half of the simulated systems, the cross-sectional correlation is higher than 83.5%. The high cross-sectional correlation between the estimated and population loadings translate to the high time-series correlation between the estimated and population factor with time-varying risk premium, whose quantiles are reported in the second row of Table 13. Specifically, the time-series correlation between the estimated and population factor is higher than 68.7% in 95% of the simulated systems, and is higher than 84.3% in half of the simulated systems. In summary, the

[46] Appendix C discusses the method to estimate the frequency-specific covariance matrices.

simulation in this section illustrates that the low-frequency factors estimated from a finite sample have high correlation with their population counterparts, and thus the empirical findings of this manuscript are not driven by estimation error.

E Simulation of conditional mean-variance-efficient portfolio

In this section, I design a simple simulation to illustrate that the conditional mean-variance-efficient (CMVE) portfolio based on a true prediction model could have a negative alpha with respect to that based on a model without predictive power, particularly when the average returns of the factors are high and time-varying. The simulation also shows that evaluating the performance of the unconditional mean-variance-efficient (UMVE) portfolio alleviates this issue.

To simplify the exposition, in this section, I consider one factor with time-varying average return and constant volatility. Specifically, I assume that the data-generating process of the factor is

$$f_{t+1} = \mu_t + \xi_{t+1},$$

where μ_t is the time-varying average return and $\xi_{t+1} \sim N(0, \sigma^2)$ is the innovation of the factor. The CMVE portfolio of this factor is then

$$f_{t+1}^{\text{CMVE}} = c\frac{\mu_t}{\sigma^2} f_{t+1},$$

where c is a constant that governs the leverage of the portfolio. To evaluate whether predicting the average return of the factor generates economic value, one conventional approach is to run the following regression:

$$f_{t+1}^{\text{CMVE}} = \alpha^{\text{CMVE}} + \beta^{\text{CMVE}} f_{t+1} + e_{t+1}, \tag{60}$$

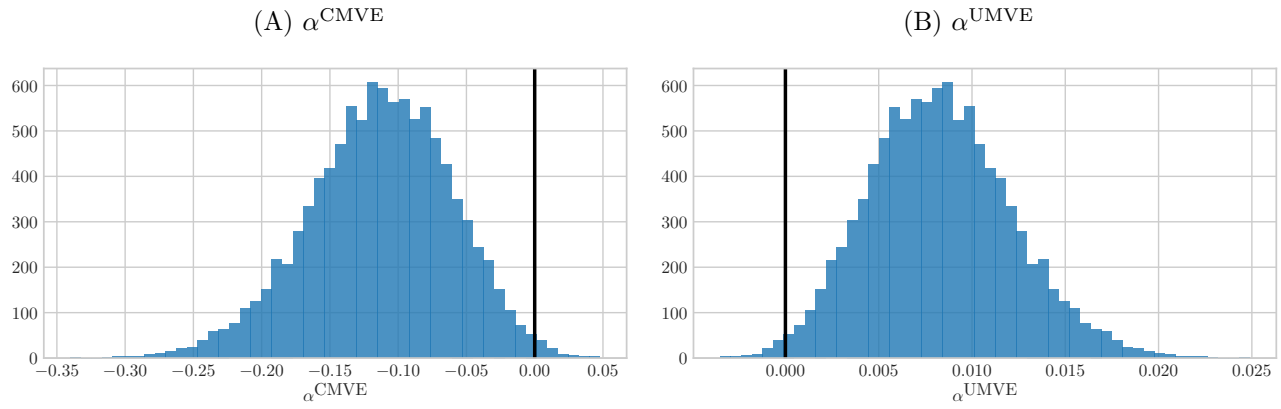
and conclude that predicting the factor returns generates additional economic value if the regression intercept is positive. Intuitively, compared to the CMVE portfolio of the factor, the buy-and-hold strategy does not exploit the predictability of the factor returns and thus is less favorable to an investor. Therefore, based on this evaluation approach, one should expect that $\alpha^{\text{CMVE}} > 0$.

I use a simulation procedure to investigate whether $\alpha^{\text{CMVE}} > 0$ holds empirically. Specifically, I simulate the factor returns for $T = 300$ months, and I assume that $\mu_t = 2$ in the first 60 months and $\mu_t = 0.8$ in the remaining 240 months. I assume that $\sigma = 1$, and thus the conditional Sharpe ratio of the factor is two in the first 60 months and 0.8 in the remaining 240 months. Using these parameters, I generate 10,000 simulated samples of the factor returns, and record the estimation of α^{CMVE} in each simulated sample.

Panel A of Figure 9 illustrates the empirical distribution of the estimation of α^{CMVE} in the 10,000 simulated samples. I find that α^{CMVE} is negative in 99.13% of the simulated samples. Thus, evaluating whether a prediction model of the factor returns generates economic value via examining whether α^{CMVE} is positive could deliver misleading results. In particular, even if one knows the true data-generating process and forms the true CMVE portfolio of the factor, the simulation results show that the true CMVE portfolio could have a negative alpha with respect to the simple buy-and-hold strategy of the factor, which does not exploit the return predictability.

Figure 9: Empirical distribution of α^{CMVE} and α^{UMVE} in 10,000 simulated samples

Panels A and B illustrate the empirical distribution of the estimation of α^{CMVE} and α^{UMVE}, which are defined in equations (60) and (61), respectively, in 10,000 simulated samples.



In the main body of the manuscript, I exploit the unconditional mean-variance-efficient (UMVE) portfolio of the factor:

$$f_{t+1}^{\text{UMVE}} = c\frac{\mu_t/\sigma^2}{1+\mu_t^2/\sigma^2}f_{t+1}.$$

To evaluate whether predicting the average return of the factor generates economic value, I run the following regression:

$$f_{t+1}^{\text{UMVE}} = \alpha^{\text{UMVE}} + \beta^{\text{UMVE}} f_{t+1} + e_{t+1}, \tag{61}$$

and conclude that predicting the factor returns generates additional economic value if α^{UMVE} is positive. Panel B of Figure 9 illustrates the empirical distribution of the estimation of α^{UMVE} in the 10,000 simulated samples, and I find that α^{UMVE} is positive in 99.13% of the simulated samples. In fact, the 87 samples that yield a negative estimation for α^{UMVE} are also those that yield a positive estimation for α^{CMVE}. Thus, compared to examining whether α^{CMVE} is positive, evaluating whether a prediction model of the factor returns generates economic value via examining whether α^{UMVE} is positive delivers less misleading results.

F Additional tables

Table 14: List of considered characteristics

This table lists the characteristics considered in the empirical analysis. The second and third columns report the acronym and definition of the variables, respectively. Section A.1 of the Online Appendix of Freyberger et al. (2020) reports the detailed definition and calculation of these characteristics.

#	Acronym	Definition
1	beta	Market beta
2	e2p	Earnings-to-price ratio: income before extraordinary items to market capitalization
3	cum_return_36_13	Cumulative return from 36 months to 13 months before
4	beme	Book-to-market ratio: book value of equity to market value of equity
5	q	Tobin's Q
6	a2me	Total assets to market capitalization
7	cum_return_1_0	Lagged one-month return
8	lme	Total market capitalization
9	cum_return_12_2	Cumulative return from 12 months to two months before
10	cto	Capital turnover: net sales to lagged total assets
11	roe	Return on equity: income before extraordinary items to lagged book value of equity
12	oa	Operating accruals: change in non-cash working capital minus depreciation scaled by lagged total assets
13	lturnover	Last month's volume over shares outstanding
14	rel_to_high_price	Price to 52 week high price
15	noa	Net operating assets: difference between operating assets and operating liabilities to lagged total assets
16	idio_vol	Idiosyncratic volatility with respect to the Fama-French three-factor model
17	investment	Percentage year-on-year growth rate in total assets
18	ato	Net sales to lagged net operating assets
19	dpi2a	Change in property, plants, and equipment and inventory to lagged total assets
20	pm	Profit margin: operating income after depreciation to sales
21	rna	Return on net operating assets: operating income after depreciation to lagged net operating assets
22	suv	Standard unexplained volume: standardized difference between actual volume and predicted volume
23	roa	Return on assets: income before extraordinary items to lagged total assets
24	ol	Operating leverage: sum of cost of goods sold and selling, general, and administrative expenses to total assets
25	free_cf	Free cash flow to book-value of equity
26	c	Cash and short-term investments to total assets
27	cum_return_12_7	Cumulative return from 12 months to seven months before
28	spread_mean	Average daily bid-ask spread in the previous months
29	prof	Gross profitability divided by the book value of equity
30	lev	Leverage: sum of long-term debt and debt in current liabilities to the sum of long-term debt, debt in current liabilities, and stockholders' equity
31	at	Total assets
32	s2p	Sale-to-price ratio: net sales to the market capitalization
33	sga2s	Selling, general and administrative expenses to net sales
34	fc2y	Fixed costs to net sales
35	d2a	Capital intensity: depreciation and amortization over total assets
36	pcm	Price-to-cost margin: difference between net sales and costs of goods sold divided by net sales

Table 15: List of macroeconomic and financial indicators

Panels A and B list the macroeconomic and financial indicators considered in the empirical analysis for predicting the returns of the low-frequency factors, respectively. In each panel, the second and third columns report the acronym and definition of the variables, respectively.

#	Acronym	Definition
Panel A: macroeconomic indicators		
1	INFL	Inflation: consumer price index for all urban consumers
2	BCI	U.S. Business confidence index from the OECD library
3	CCI	U.S. Consumer confidence index from the OECD library
4	CLI	U.S. Composite leading indicator from the OECD library
5	FIN_UNC_1month	Measure of one-month financial uncertainty (Jurado et al., 2015; Ludvigson et al., 2021)
6	FIN_UNC_3month	Measure of three-month financial uncertainty (Jurado et al., 2015; Ludvigson et al., 2021)
7	FIN_UNC_12month	Measure of 12-month financial uncertainty (Jurado et al., 2015; Ludvigson et al., 2021)
8	MACRO_UNC_1month	Measure of one-month macroeconomic uncertainty (Jurado et al., 2015; Ludvigson et al., 2021)
9	MACRO_UNC_3month	Measure of three-month macroeconomic uncertainty (Jurado et al., 2015; Ludvigson et al., 2021)
10	MACRO_UNC_12month	Measure of 12-month macroeconomic uncertainty (Jurado et al., 2015; Ludvigson et al., 2021)
11	REAL_UNC_1month	Measure of one-month real uncertainty (Jurado et al., 2015; Ludvigson et al., 2021)
12	REAL_UNC_3month	Measure of three-month real uncertainty (Jurado et al., 2015; Ludvigson et al., 2021)
13	REAL_UNC_12month	Measure of 12-month real uncertainty (Jurado et al., 2015; Ludvigson et al., 2021)
14	gIndprod	Growth of industrial production
15	gEmploy	Growth of employment
16	gConsdur	Growth of durable consumption
17	gConsnon	Growth of nondurable consumption
18	gConsserv	Growth of services consumption
19	RECESS	NBER-dated recessions
Panel B: financial indicators		
1	DP	Dividend-price ratio: log of a 12-month moving sum of dividends paid on the S&P 500 index minus the log of stock prices (S&P 500 index)
2	DY	Dividend yield: log of a 12-month moving sum of dividends paid on the S&P 500 index minus the log of lagged stock prices (S&P 500 index)
3	EP	Earnings-price ratio: log of a 12-month moving sum of earnings on the S&P 500 index minus the log of stock prices (S&P 500 index)
4	DE	Dividend-payout ratio: log of a 12-month moving sum of dividends paid on the S&P 500 index minus the log of a 12-month moving sum of earnings on the S&P 500 index
5	SVAR	Stock variance: sum of squared daily returns on the S&P 500 index
6	BM	Book-to-market ratio: ratio of book value to market value for the Dow Jones Industrial Average
7	NTIS	Net equity expansion: ratio of 12-month moving sums of net issues by NYSE stocks divided by the total end-of-year market capitalization of NYSE stocks
8	TBL	Treasury bill rate: interest rate on a three-month Treasury bill
9	LTY	Long-term yield: long-term government bond yield
10	LTR	Long-term return: return on long-term government bonds
11	TMS	Term spread: difference between the long-term yield and Treasury bill rate

Table 15 continued

#	Acronym	Definition
12	DFY	Default yield spread: difference between BAA- and AAA-rated corporate bond yields
13	DFR	Default return spread: difference between long-term corporate bond and long-term government bond returns
14	DISAG	Investor disagreement (Huang, Li, and Wang, 2021)
15	FedFundsRate	Federal funds effective rate
16	Real_FedFundsRate	Real Federal funds rate
17	AggLiq	Levels of aggregated liquidity (Pastor and Stambaugh, 2003)
18	BW_SENT	Investor sentiment index of Baker and Wurgler (2006)
19	BW_SENTorthogonalized	Investor sentiment index of Baker and Wurgler (2006), orthogonalized version
20	PLS_SENT	Aligned investor sentiment index of Huang et al. (2015)
21	PLS_SENTorthogonalized	Aligned investor sentiment index of Huang et al. (2015), orthogonalized version
22	intermediary_capital_ratio	Intermediary capital ratio of He et al. (2017)
23	intermediary_lev_ratio_squared	Squared reciprocal of the intermediary capital ratio
24	cay	Consumption-wealth ratio of Lettau and Ludvigson (2001b)