Currency Risk Premia Redux*

Federico Nucera† Lucio Sarno‡ Gabriele Zinna§

May 2022

Abstract

We study a large currency cross section using recently developed asset pricing methods which account for omitted-variable and measurement-error biases. First, we show that the implied pricing kernel includes three latent factors: a strong U.S. "Dollar" level factor, and two weak, high Sharpe ratio "Carry" and "Momentum" slope factors. The evidence for an additional "Value" factor is scant. Second, using this pricing kernel, we obtain robust estimates of the risk premia of more than 100 nontradable risk factors. Only a small fraction of these factors are priced – mostly relating to volatility, uncertainty and liquidity conditions, rather than macro variables.

Keywords: Currency risk premia, asset pricing, omitted factors, measurement error, weak factors.

JEL Classification: F31; G12; G15.

*We are indebted for constructive comments to Stefano Giglio (Editor), an anonymous Associate Editor, two anonymous Referees, Martijn Boons, Pasquale Della Corte, Antonio Di Cesare, Hsuan Fu, Angelo Ranaldo, Dacheng Xiu, and participants at presentations held at the Bank of Italy, the University of Cambridge, the 2021 "Workshop on Exchange Rates" at the Swiss National Bank. All errors are our responsibility. This research was partly carried out when Gabriele Zinna was visiting scholar at Amsterdam Business School and Tilburg University. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of Italy.

†Bank of Italy, Rome. E-mail: Federico.Nucera@bancaditalia.it

‡Cambridge Judge Business School, University of Cambridge, Trumpington Street, Cambridge CB2 1AG, and Centre for Economic Policy Research (CEPR). E-mail: l.sarno@jbs.cam.ac.uk

§Bank of Italy, Rome. E-mail: Gabriele.Zinna@bancaditalia.it

1 Introduction

In the foreign exchange (FX) market, the price of risk represents the compensation required by investors for a unit exposure to the systematic risk resulting from holding investments denominated in foreign currencies. Since the seminal paper of Lustig and Verdelhan (2007), cross-sectional asset pricing has been applied successfully to currency returns, and a growing literature continues to develop with the aim of explaining the cross section of currency returns and to provide estimates of the price of currency risk. At the same time, we have also observed a proliferation of currency investment strategies, which attract a large fraction of the over 6 trillion U.S. dollars traded in currency markets daily. It is therefore crucial, for investors and market observers alike, to uncover the sources of the underlying risk-return trade-off in this titanic market. To this end, in this paper we provide new evidence on the optimal factor model for currency returns and robust estimates of currency risk premia.[1]

Thus far, the FX literature has largely established the risk-return trade-off in terms of *tradable* risk factors. These factors represent convolutions of returns associated with currency investment strategies (e.g., carry and momentum factors) and therefore prevent a deep economic interpretation. Only a few papers focus on *nontradable* risk factors, i.e., factors representing macroeconomic and financial risks such as for example the global volatility factor of Menkhoff et al. (2012a). But this strand of the literature is evolving rapidly, so that we observe also a proliferation of FX risk factors, i.e., a "factor zoo", albeit more contained than for equities (e.g., Feng et al., 2020).[2]

When a new candidate factor is proposed, the first goal is to determine its risk premium (or price of risk). If the factor is tradable, a model-free estimate of its risk premium is readily available, being simply the time-series average of its excess return (Cochrane, 2005). By contrast, if the factor is nontradable, the task of estimating its risk premium is not trivial. A nontradable factor is by definition a non-return-based factor and, as a consequence, its mean is not informative about its price of risk. Therefore, one needs to

[1]The most popular currency strategies include carry-trade strategies based on interest rate differentials across countries (e.g., Lustig et al., 2011; Menkhoff et al., 2012a; Lettau et al., 2014), momentum strategies based on past currency returns (e.g., Menkhoff et al., 2012b, Asness et al., 2013), value strategies based on deviations from purchasing power parity (e.g., Asness et al., 2013; Kroencke et al., 2014; Menkhoff et al., 2017), global imbalances strategies based on imbalances in trade and capital flows (Della Corte et al., 2016), and macro strategies based on, for example, output gap differentials (Colacito et al., 2020).

[2]At first, the empirical asset pricing literature rested on a single factor, namely the market factor, to price the cross section of stock returns (Sharpe, 1964; Lintner, 1965). Since then, more than 300 risk factors have been claimed to explain stock returns with some statistical significance (Harvey et al., 2016), but some of these factors could just be "lucky" (Harvey and Liu, 2021).

recur to statistical methods, such as for example the standard two-pass procedure of Fama and MacBeth (1973) – FMB thereafter – to obtain an estimate of the factor risk premium.

While the two-pass FMB procedure can be easily implemented, the resulting price of risk estimates can be biased for two main reasons. First, some relevant factors entering the pricing kernel, or stochastic discount factor (SDF), could be omitted (omitted-variable bias). Second, the candidate factor could be measured with noise (measurement-error bias). Recently, Giglio and Xiu (2021) developed a three-pass procedure that helps address both sources of bias by exploiting the information contained in a reasonably large cross section of test assets. This literature, albeit very young, has already established a set of useful results for the U.S. stock market. In this paper, we build on this literature but shift the focus to currency markets. Specifically, we address the following two questions: How many (and which) factors should the optimal currency SDF comprise? Which nontradable factors, out of the plethora of factors proposed in the finance literature, are priced in the cross section of currency portfolio returns?

The FX literature has generally looked at each investment strategy in isolation, therefore resting on small cross sections of test assets. However, the use of a limited cross section of test assets may not provide a robust/valid test of an asset pricing model (Lewellen et al., 2010). In addition, the omitted-variable and measurement-error problems inherent in the estimation of the prices of risk have not been taken fully into account. For these reasons, it is fair to argue that the economic sources of the risk-return trade-off underlying popular currency investment strategies are still hotly debated. To fill this gap, we estimate the risk premia of a long list of nontradable macro-financial candidate factors from a reasonably large cross section of currency portfolios, or test assets. We do this by combining the three-pass model of Giglio and Xiu (2021) with the statistical method of Lettau and Pelger (2020a,b). The latter, as explained in detail below, is key to first extract the underlying latent factors accurately and determine the structure of the latent-factor currency SDF.

The three-pass method of Giglio and Xiu (2021) – GX thereafter – that we employ to revisit the macro-financial determinants of currency risk premia serves our purpose, as it tackles both the omitted-variable and measurement-error problems. To do so, it exploits the information contained in the panel of test-asset returns and, in particular, in the underlying latent pricing factors that are extracted from the panel of returns. In practice, this procedure projects the nontradable candidate factors onto the space of the latent pricing factors. The nontradable factors' risk-premium estimates are then simply given by linear combinations of the prices of risk of the latent pricing factors. In this way, one can remain agnostic about the set of 'true' risk factors (i.e., the controls), and yet obtain robust estimates of nontradable factors' risk premia.

It is evident, however, that the method of GX heavily relies first on estimating the latent factors, and then on determining the factor structure of the optimal SDF, i.e., the relevant pricing factors. For this reason, we amend the GX procedure by resorting to the Risk-Premium Principal Component Analysis (RP-PCA) method of Lettau and Pelger (2020a,b) – LP thereafter. In essence, RP-PCA is a generalized version of PCA, regularized by a pricing-error penalty term (named risk-premium weight or RP-weight), which "overweights" the test-asset mean returns relative to their variances. As a result, the estimated factors fit not only the time series, but also the cross section of expected returns. Strong systematic factors should be estimated more efficiently, and weak factors which possess high risk premia (Sharpe ratios) can be detected more easily. LP show that in their setting the RP-PCA estimator can be asymptotically more efficient than PCA in the sense that the SDF and factors estimated by RP-PCA are more highly correlated with the 'true' SDF and factors than those estimated by PCA. Therefore, by inspecting the properties of the factors extracted with RP-PCA, one obtains also clear indications on the structure of the optimal latent-factor currency SDF. We refer to this combined procedure that uses the methods developed separately by GX and LP as the augmented three-pass method, and we show that the use of RP-PCA enhances the three-pass model pricing performance.

In the empirical analysis, the underlying FX data consist of 49 individual currencies sampled at monthly frequency, from 1983 to 2017. We take the perspective of a U.S. investor, so that the individual currencies are expressed relative to the U.S. dollar. In the baseline analysis, the test assets consist of 46 currency portfolios, resulting from nine of the most popular currency investment strategies. Turning to the nontradable candidate risk factors, our list consists of more than 100 factors, which we categorize into three groups: financial, macro, and text-based factors. The latter factors are obtained by aggregating into an index news coverage about specific sources of uncertainty. To our knowledge, we are the first to consider such a large number of nontradable factors, capturing a wide range of macro-financial risks, and assess their implications for currency returns. Based on this extensive dataset, we uncover a number of interesting findings that help shed light (i) on the optimal latent-factor currency SDF, and (ii) on the macro-financial sources of the risk-return trade-off inherent in currency investment strategies. We present the findings in this order.

First, we show that the currency SDF consists of at least three latent pricing factors. The first factor is a strong factor, while the remaining two explain fewer portfolios, and hence are in line with a weak-factor interpretation. Yet, also these weak factors are relevant pricing factors, as they display high Sharpe ratios, and hence cannot be excluded from the SDF. Notably, the third factor is detected by RP-PCA but not by standard PCA. Hence, by neglecting the information in the portfolio means, one incurs the risk of omitting relevant factors with high Sharpe ratios, which can in turn distort the nontradable factor risk-premium estimates.

Relatedly, we find that RP-PCA changes materially the information spanned by the factors relative to PCA in a way that the estimated factors should be closer to the underlying 'true' pricing factors. For example, considering three-factor SDFs, the pricing errors drop significantly and the maximal SDF Sharpe ratio increases substantially using RP-PCA with a reasonably high RP-weight instead of PCA – note that PCA is a special case of RP-PCA with no "overweight" on the means. These differences are evident for SDFs of equal size and become even starker if one compares the respective optimal SDFs (i.e., implied by formal tests for the number of factors), which consist of three and two factors for RP-PCA and PCA, respectively. Importantly, we also document that, while the pricing accuracy improves with the RP-weights, the explained systematic variance remains essentially unchanged. Thus, in practice, there is no trade-off in choosing even very high RP-weights.[3]

Moreover, the analysis of the portfolio risk exposures reveals that the extracted, orthogonalized latent factors retain a clear economic interpretation. The first latent factor plays the role of a currency level factor, as it displays roughly equal factor loadings across currency portfolios. This factor therefore resembles the Dollar factor of Verdelhan (2018). By contrast, the remaining factors are slope factors, as we can identify investment strategies for which the corner portfolios take factor loadings of opposite signs, with almost monotonic patterns across portfolios. Put simply, these latent factors behave as spread portfolios (which are self-financed long-short investment strategies), and therefore naturally connect to specific investment

[3]We show that the gains seem to stabilize for reasonably high RP-weights. Hence, in the main analysis, we select an RP-weight of 20, which is in line with that chosen by LP for equity portfolios. However, it is important to note that the gain from using RP-PCA rather than PCA is smaller out of sample, which makes sense since RP-PCA is designed to maximize the in-sample Sharpe ratio.

strategies. In particular, the second-latent factor is a "Carry" factor, while the third factor is a (short-term) "Momentum" factor.[4] The fourth factor seems to be a "long Value short (long-term) Momentum" factor, but it is not selected by any of the statistical criteria employed, consistent with the fact that its inclusion in the SDF improves only marginally the overall model pricing performance and Sharpe ratio.[5] Therefore, this analysis ultimately shows that the currency SDF comprises three pricing factors that can be interpreted as "Dollar", "Carry", and "Momentum" factors.

Second, based on this optimal SDF, we turn to estimate the risk premia of the nontradable candidate factors. To start with, we find that the spanning regressions of the nontradable factors on the pricing latent factors deliver, on average, low R^2s. In the GX's framework, this would indicate that a large portion of nontradable risk factors is due to measurement error. The problem is particularly severe for macro factors, while some of the text-based and, especially, of the financial factors are measured more precisely. In particular, text-based and financial factors are mainly exposed to the "Carry" factor, but some of these factors (mostly financial ones) also display significant exposures to the "Momentum" factor. Interestingly, the exposures of these candidate factors to the "Carry" and "Momentum" factors generally take opposite sign. This indicates that the two strategies respond to some of the same sources of financial risk, but in opposite ways. For example, when volatility increases "Momentum" tends to perform well, while "Carry" performs poorly, in line with what is observed during periods of markets turmoil, such as during the global financial crisis. This in turn implies that, if the "Momentum" factor is omitted from the SDF, the return-based candidate factors – the original nontradable factors cleaned from measurement error and converted into return factors using the fitted value of the spanning regressions – can display different behaviors and risk premia.

Turning more specifically to the risk-premium estimates, we show that the risk premia obtained using the augmented three-pass method are substantially different from the FMB two-pass estimates.[6] In fact,

[4]Short-term and long-term momentum strategies differ in that they use as sorting signals the one-month and one-year past returns, respectively. Menkhoff et al. (2012b) show that both of these strategies are profitable and imperfectly correlated, although short-term momentum generates higher expected returns.

[5]While this factor, call-it simply "Value", displays a statistically significant mean return, the magnitudes of its return and Sharpe ratio are small in comparison with those of the "Carry" and "Momentum" factors. This helps explain why this factor takes a small weight in a four-factor SDF. We also document that the remaining latent factors retain no clear interpretation. Indeed, they are time-series factors as they have zero prices of risk, and hence take zero weights in the SDF.

[6]The FMB estimates are intentionally subject to both the omitted-variable and measurement-error problems, as no control factor other than the Dollar factor (captured by the constant) is included, in addition to the candidate factor.

the two-pass method seems to deliver higher absolute point estimates and a larger number of candidate factors with significant risk premia. This is not surprising given that inflated prices of risk are common among nontradable factors, exactly because they contain noise (Adrian et al., 2014). We document that the measurement-error problem is indeed pervasive also for a large number of our nontradable factors. Together with the omitted-variable problem, it can lead to biased risk-premium estimates and/or to erroneous selection of currency risk factors.

At the same time, thanks to the augmented three-pass method, we can also show that some nontradable factors are indeed priced in currency returns. While the list of relevant factors is shorter than using the two-pass method, it is still diverse, and mainly pertains to financial and text-based factors. Some of the nontradable factors previously uncovered by the literature turn out to be less or even not relevant, but other "novel" factors (i.e., which were not considered in previous currency research) appear to have significant risk premia, disclosing a tight link between currency and other financial markets, mainly channeled through "Carry", in line with the conjecture of Koijen et al. (2018). In particular, our findings highlight the relevance of uncertainty and volatility measures (both financial and text-based) and of liquidity factors to explain currency returns. Specifically, the global volatility factor of Menkhoff et al. (2012a) and the global Economic Policy Uncertainty (EPU) of Baker et al. (2016) are singled out, as their risk premia are large and precisely estimated. Moreover, the signs of the risk-premium estimates of the financial and text-based factors appear intuitively clear. Factors that perform poorly (well) in bad states of the world command positive (negative) currency risk premia, and hence are procyclical (countercyclical) factors, based on the three-pass estimator.

However, the results point to a substantial disconnect between currency returns and macroeconomic variables, which is disappointing as it is hardly imputable to their measurement error given that the three-pass method accounts for that. Moreover, even among the few macro factors with weakly significant premia estimates, some display risk premia with counterintuitive signs. We then show that the disconnect is not the consequence of macro factors being weak factors (i.e., factors that are relevant only for a subset of the test assets). In fact, we find similar results using the supervised principal component analysis (SPCA) estimator, recently developed by Giglio et al. (2021c) to explicitly tackle the issue of weak candidate factors in the estimation of factor risk premia. Moreover, the SPCA results show that even the few macro factors with significant risk premia can be very poorly hedged out of sample using currency portfolios. As a result, the disconnect between macro factors and currency portfolio returns is confirmed using SPCA. We show that these results hold in a number of robustness checks and additional analysis.

Finally, we verify through a simulation exercise that the augmented three-pass method works well also in finite samples that match the dimension and properties of the FX portfolio returns in our paper. The simulation is calibrated on a reduced-form SDF specification which allows for four factors that mimick the optimal SDF documented in the empirical analysis and generates simulated data that reproduce the features of our FX portfolio returns. Given this data generating process (DGP), the simulation results demonstrate that the method works well in finite samples that match our data, and also make clear that both the omitted-variable and measurement-error problems can be material in the estimation of currency risk premia, in a similar way as documented by GX for equity markets. This evidence gives us further comfort that the methods employed here are both reliable and desirable for our purposes, and that the unconditional three-pass model, if well specified, provides a satisfactory description of dynamically rebalanced FX portfolio returns.

The closest paper to ours is independent work by Chernov et al. (2021), which tackles similar objectives to the ones targeted in our paper, in a very different way. Specifically, Chernov et al. (2021) address the question of the optimal factor model for pricing currency risk, which relates to the first goal of our paper. They do so by studying directly the mean-variance efficient portfolio, and relying on the conditional projection of the SDF onto excess returns of individual currencies. Reducing the dimensionality of the problem by limiting the sample to G10 currencies, they show that this approach allows to price individual currencies and several canonical strategies (derived from carry, momentum, and value signals), both conditionally and unconditionally. On the one hand, this approach has the advantage, relative to the methods adopted in our paper, that currency pricing is carried out more directly since the estimated SDF is represented as a linear function of the unconditional mean-variance efficient portfolio. On the other hand, working directly with the mean-variance efficient portfolio can only be achieved on a set of assets that is small enough to allow reliable estimation of the covariance matrix of currency returns. Moreover, one needs to assume that the set of factors or signals that drive the conditional mean are *known*. In turn, this exposes the approach to omitted-variable problems (in addition to potential measurement-error problems), which are instead taken into account using the GX three-pass methods employed in our paper. Ultimately, we view the study of

Chernov et al. (2021) as complementary to our paper.

The remainder of the paper is organized as follows. Section 2 presents the augmented three-pass method, and Section 3 describes the FX portfolios and the nontradable candidate risk factors. Section 4 presents the baseline empirical findings, as well as robustness exercises. Section 5 studies the finite-sample performance of the estimator in simulation, whereas Section 6 deals with the weak-factor problem via SPCA. Section 7 concludes. A separate Internet Appendix briefly reviews the two-pass estimator (Section I); presents the FX investment strategies (Section II); the nontradable factors (Section III); additional empirical evidence (Sections IV and V); the simulation exercise (Section VI); and the weak-factor analysis (Section VII).

2 Asset Pricing Methods

The FMB two-pass method has long represented the workhorse model to estimate risk premia in empirical asset pricing (see a brief description of FMB in the Internet Appendix, Section I). In currency asset pricing, it is widely employed at least since the influential study of Lustig and Verdelhan (2007). Over the years, some fixes to the original two pass-procedure have been proposed, and they mainly regard the efficiency of the estimates, which relates to the use of the generated $\hat{\beta}$ covariates in the second-pass regression (e.g., Shanken, 1992; Burnside, 2011). By contrast, the omitted-variable and measurement-error problems have received less attention.

The *omitted-variable* problem arises when (some of) the relevant risk factors are omitted from the SDF. This omission biases the estimates of the risk exposures in the first pass, and the estimates of the prices of risk in the second pass. As a result, the researcher attributes the effect of the missing factors/exposures to the estimated effect of the included factors/exposures. In the first pass, the severity of the bias depends on the time-series correlation between the factors included and those omitted. In the second pass, it varies with the cross-sectional correlation of the estimated exposures and the missing exposures associated with the omitted factors. The *measurement-error* problem instead emerges even when the researcher includes all the 'true' risk factors in the SDF, but the factors are measured with noise. This problem is particularly severe in the case of nontradable (i.e., non-return-based) factors, especially those based on macroeconomic variables. The use of noisy factors may bias the first-pass estimates of the risk exposures and, as a consequence, the second-pass estimates of the prices of risk.

Both problems can manifest in many situations. A clear example is when the researcher wants to estimate the price of risk of a novel nontradable factor g_t, λ_g. In principle, the standard FMB procedure is viable but the researcher would need to (i) know the set of control factors, i.e., the set of 'true' factors entering the SDF, f_t; and (ii) use factors that are cleaned, i.e., measured without noise. By contrast, the three-pass method of GX delivers an estimate of the price of risk of the candidate factor that is not affected by (i) and (ii). To do so, the GX method exploits the information in the test assets, by projecting the candidate factor onto the space of the latent pricing factors implied in the cross section of test-asset returns. In this way, one can remain agnostic about the set of 'true' risk factors (i.e., the controls), and yet obtain an estimate of λ_g that is not affected by the omitted-variable problem. Moreover, one can easily account for the measurement error in the candidate factor.

While GX employ standard PCA to extract the latent pricing factors, one can recur to other methods to estimate the factors and still exploit in full the benefits of the three-pass method. Recently, LP developed the RP-PCA estimator. A benefit of this novel method is that the latent factors are estimated such that they fit both the time series and cross section of expected returns. Conversely, standard PCA neglects the information in the means of the portfolio returns. We therefore combine the RP-PCA method of LP with the three-pass method of GX, and this is why we call it the augmented three-pass method.

2.1 (Augmented) Three-Pass Method

Before turning to the three-pass method of GX, we first present the RP-PCA method that we use to extract the factors from the panel of currency returns and the evaluation criteria employed to shed light on the optimal currency SDF.

2.1.1 Latent Factors Estimation

To start with, we assume that K factors capture the systematic component of asset returns and the unexplained idiosyncratic component subsumes the asset-specific risks, such that

$$X_{nt} = \mathrm{F}_t \psi_n^\top + \epsilon_{nt}, \qquad n = 1, \ldots, N, \quad t = 1, \ldots, T, \tag{1}$$

where X_{nt} is the n-th test asset's time-t excess return, $F_t = [F_{1t}, \dots F_{Kt}]$ denotes the time-t $1 \times K$ vector of latent factors, ψ_n is the $1 \times K$ vector of factor loadings for test asset n, and ϵ_{nt} is the asset return's idiosyncratic component. In matrix notation, it takes the compact form $X = \mathrm{F}\psi^\top + \epsilon$, where X is a $T \times N$ matrix of returns, F is the $T \times K$ matrix of latent factors, ψ is the $N \times K$ matrix of factor loadings, and ϵ is the $T \times N$ matrix of residuals. It is then evident that, if factors and residuals are uncorrelated, the covariance matrix of the returns is given by

$$\mathrm{Var}(X) = \psi \mathrm{Var}(\mathrm{F}) \psi^\top + \mathrm{Var}(\epsilon), \tag{2}$$

which consists of a systematic part and an idiosyncratic part. Standard PCA exploits the fact that the factors relate to the largest eigenvalues of $\mathrm{Var}(X)$, which can be retrieved from the sample covariance matrix of excess returns

$$\Sigma_{\mathrm{PCA}} = \frac{1}{T} X^\top X - \overline{X}^\top \, \overline{X}, \tag{3}$$

where $\overline{X}$ denotes the sample mean of excess returns.

The estimated factor loadings $\hat{\psi}$ are proportional to the eigenvectors associated with the largest eigenvalues of Σ_{PCA}. The factors $\hat{F}_t$ are then obtained by regressing the asset returns on the factor loadings. Thus, factors identified by PCA minimize the unexplained time-series variation of the returns. Evidently, however, the information in the means of the returns is not accounted for. LP note that, in the context of asset pricing, this implies ignoring valuable information, as the role of the means is explicitly given by Ross' arbitrage pricing theory (APT).[7] Asset pricing factors should capture the information contained both in the first and second moments of test-asset returns. For this reason, LP propose to apply PCA to a covariance matrix with overweighted sample mean returns; in essence, RP-PCA is a generalized version of PCA regularized by a pricing-error penalty term, which is tantamount to applying PCA to the covariance matrix

$$\Sigma_{\mathrm{RP}} = \frac{1}{T} X^\top X + \omega \, \overline{X}^\top \, \overline{X}, \tag{4}$$

[7]Under the strong form of APT, residual risk has a risk premium of zero, which holds without loss of generality when assets are portfolios. An asset excess return is then given by its exposures to the factors times the factors' risk prices. Moreover, if the factors are excess returns, no-arbitrage implies that their means are the factors' prices of risk. Hence, the means are informative about the assets' risk premia.

where ω is the penalty term, or RP-weight. As before, the factors are constructed by regressing the returns on the factor loadings, i.e., $\hat{F} = X\hat{\psi}(\hat{\psi}^\top\hat{\psi})^{-1}$. However, the loadings $\hat{\psi}$ are now proportional to the eigenvectors associated with the largest eigenvalues of the Σ_{RP} matrix. Intuitively, in RP-PCA, the eigenvalues relate to a generalized notion of "signal strength" of a factor, while in PCA the eigenvalues are equal to the factor variances, exactly because the information in the portfolio means is neglected. That is, the matrix Σ_{RP} should converge to

$$\psi(\Sigma_F + (1+\omega)\mu_F^\top\mu_F)\psi^\top + \text{Var}(\epsilon), \tag{5}$$

where Σ_F and μ_F denote the covariance matrix and the means of F, respectively. Moreover, applying PCA to Σ_{RP} is equivalent to minimizing jointly the time-series unexplained variation and the cross-sectional pricing errors

$$\min_{\psi,F} \underbrace{\frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T}(X_{nt} - F_t\psi_n^\top)^2}_{\text{TS unexplained variation}} + \omega \underbrace{\frac{1}{N}\sum_{i=1}^{N}(\overline{X}_n - \overline{F}\psi_n^\top)^2}_{\text{CS pricing error}}, \tag{6}$$

where $\overline{F}$ is the vector of factor expected values. From Eqs. (4)-(6), it is clear that RP-PCA with $\omega = -1$ is equivalent to standard PCA as it forgoes the information in the means. Also note that RP-PCA with $\omega = 0$ corresponds to applying PCA to a correlation matrix instead of a covariance matrix. Conversely, RP-PCA with $\omega > 0$ can be interpreted as PCA applied to a matrix that "overweights" the information in the means. That is, RP-PCA combines two moment conditions, pushing up the signal-to-noise ratio and therefore leading to more efficient estimates of the factors. It selects factors that explain the time series, but at the same time penalizes factors with low Sharpe ratios. This is because factors that help price the cross section of asset returns have non-vanishing returns and higher Sharpe ratios. Thus, RP-PCA with $\omega > 0$ may help detect weak factors if they have high Sharpe ratios, exactly because the weak signal in their variances is enhanced by the information in their means. Meanwhile, it protects from selecting spurious factors (i.e., factors with vanishing loadings), as it requires the estimated factors to explain a substantial amount of time-series variation.

Evaluation Criteria. The spectrum of the estimated eigenvalues is informative about the factors' "signal strengths" and, hence, can help determine the optimal SDF. One can establish how many factors are relevant, as well as discern strong from weak factors. Statistical tests such as the ones used by LP and GX are useful in this regard.[8] Importantly, a by-product of the LP method is that factors retain a clear economic interpretation, as factors extracted using RP-PCA are return-based with unrestricted means. In fact, one can rely on several intuitive metrics to complement the evidence resulting from statistical tests. In this way, the choice of the optimal SDF is guided by both statistical and economic criteria.

A clear object of interest is the maximal Sharpe ratio from the tangency portfolio of the mean-variance frontier spanned by the linear combination of the K selected latent factors, $\hat{F} \times \hat{b}_{MV}^{\top}$, where $\hat{b}_{MV} = \mu_F {\Sigma_F}^{-1}$ is a $1 \times$ K vector; the $\hat{b}_{MV}$ entries capture the factor weights in the implied SDF, $\varphi_t = 1 - (\hat{F}_t - \mu_F)\hat{b}_{MV}^{\top}$.[9] Two further diagnostic criteria – the root-mean-square error ($\overline{RMS}_{\alpha}$) and the magnitude of the idiosyncratic variance ($\overline{\sigma}_{\epsilon}^2$) – are useful to evaluate the model performance, inform the choice of the penalty value ω, and determine which factors to include in the SDF. Such criteria are centered around the estimation of ordinary least squares (OLS) time-series regressions

$$X_{nt} = \alpha_n + \hat{F}_t \psi_n^{\top} + \epsilon_{nt}, \qquad n = 1, \ldots, N, \quad t = 1, \ldots, T, \tag{7}$$

where the intercept α_n captures the magnitude of the asset-specific pricing errors. Put simply, Eq. (7) is the OLS counterpart of the factor model of Eq. (1), but differs for two main reasons. First, it includes the intercept, while the factor model imposes no intercept and hence the residuals have means that are not necessarily zero. Second, the OLS regression (without intercept) and the factor model yield the same estimates of ψ_n only when RP-PCA uses $\omega = 0$. This is because the pricing-error term of Eq. (6) drops out, and hence the two methods minimize the same objective function.[10] Nevertheless, LP argue that the difference turns out to be negligible in the data. Thus, one can use Eq. (7) to compute $\overline{RMS}_{\alpha} = \sqrt{\hat{\alpha}\hat{\alpha}^{\top}/N}$,

[8]In order to determine the optimal number of latent factors to include in the SDF, LP use the test of Onatski (2010), whereas the GX's estimator is based on a penalty function similar to the one of Bai and Ng (2002). The Onatski (2010) test relies on the idea that the eigenvalues associated with the systematic factors diverge to infinity, while the eigenvalues associated with idiosyncratic factors cluster around a single point. Put simply, the eigenvalues of systematic factors should be separated from those of weak factors.

[9]If the estimated factors are orthogonal, Σ_F is diagonal and $\hat{b}_{MV}$ is a vector with entries $\hat{b}_{MV,k} = \mu_{F,k}/\sigma_{F,k}^2$, where $\mu_{F,k}$ and $\sigma_{F,k}^2$ denote the k-th factor's estimated mean and variance. We follow common practice and search for a small number of factors whose linear combination with constant loadings in the SDF prices assets unconditionally. In a recent paper, Chernov et al. (2021) depart from this approach with a conditional model where a single factor drives the SDF and its loading varies over time. Such factor is constructed via a conditional mean-variance efficient portfolio approach.

[10]RP-PCA with $\omega = -1$ yields the same estimates of Eq. (7) applied to demeaned X_{nt} and $\hat{F}_t$. LP show that, also for $\omega > 0$, RP-PCA loadings can be retrieved using OLS regressions. We return to this issue in the next section.

and $\overline{\sigma}_{\epsilon}^2 = \frac{1}{N}\sum_{n=1}^{N}\left[Var(\hat{\epsilon}_n)/Var(X_n)\right]$ implied by $\hat{F}$. Note that these two statistics should move in opposite direction as the penalty ω varies; *ceteris paribus*, for higher values of the penalty, the pricing error should diminish at the cost of higher variance of the idiosyncratic component. Hence, based on these statistics, one can evaluate the trade-off, and pin down the optimal value of the RP-weight, ω.

2.1.2 Candidate Factor Price of Risk Estimation

So far, we showed how to efficiently estimate the latent factors, and how to select the factors entering the optimal SDF. All of this is instrumental to apply the GX three-pass method to obtain accurate estimates of the candidate factors' risk premia, which we present next.

1. `Test-Asset Exposures to Latent Factors (`ψ`)`. The first pass consists of estimating test-asset risk exposures to latent factors. Because GX use PCA to extract the latent factors, they obtain the risk exposures through time-series OLS regressions of test-asset excess returns on the latent factors. As mentioned earlier, this is no longer exact when the factors are extracted using RP-PCA with $\omega > 0$. However, the exposures ψ implied in the factor model of Eq. (1) can still be recovered using OLS regressions. To do so, one has to transform the excess return data and the factors in such a way to incorporate the information of the pricing errors. Specifically, the time-series OLS regressions become

$$\tilde{X}_{nt} = \tilde{F}_t \psi_n^{\top} + \epsilon_{nt}, \qquad n = 1, \ldots, N, \quad t = 1, \ldots, T, \tag{8}$$

where $\tilde{X}_{nt} = X_{nt} + \tilde{\omega}\bar{X}_{nt}$, and the vector $\tilde{F}_t$ contains elements defined as $\tilde{F}_{kt} = \hat{F}_{kt} + \tilde{\omega}\bar{F}_{kt}$ for $k = 1, \ldots, K$, with $\tilde{\omega} = \sqrt{\omega + 1} - 1$. In this way, the RP-PCA risk exposures can be retrieved for any value of ω.

2. `Latent Factor Prices of Risk (`γ`)`. The second pass delivers the estimates of the prices of risk of the latent factors. The estimates are obtained by running a cross-sectional regression of average realized excess returns on the previously estimated exposures of the test assets to the latent factors,

$$\overline{X}_n = \hat{\psi}_n \gamma^{\top} + a_n, \quad n = 1, \ldots, N, \tag{9}$$

where γ is the $1 \times K$ vector of the latent factor prices of risk.[11]

`3. Candidate Factor Price of Risk` (λ_g). The last pass of the GX procedure yields the price of risk of the candidate factor g_t. First, one projects the candidate factor onto the space of the latent pricing factors, by running a time-series spanning regression of the candidate factor innovation, g_t^ι, on the demeaned latent factors, $\hat{F}_t^\iota = \hat{F}_t - \mu_F$, as follows

$$g_t^\iota = \hat{F}_t^\iota \eta^\top + u_t, \tag{10}$$

where η is the $1 \times K$ vector collecting the loadings of the candidate factor on the K latent factors. Then, using the estimated η-exposures, one implements

$$\hat{\lambda}_g = \hat{\gamma}\hat{\eta}^\top, \tag{11}$$

$$\hat{g}_t = \hat{F}_t\hat{\eta}^\top, \tag{12}$$

and obtains the *price of risk* of the candidate factor, $\hat{\lambda}_g$, and the *cleaned return-based* candidate factor, $\hat{g}_t$ (i.e., the nontradable factor after the removal of measurement error, u_t, and converted into a tradable return-based factor).

`Rotation Invariance of Risk Premia`. Before turning to the empirical analysis, it is important to introduce the rotation-invariance result shown in GX. The main result is that the risk-premium estimate of a candidate factor is rotation invariant, as its estimate does not change when the model is expressed as a function of rotated factors, $\hat{\hat{F}}_t \equiv F_t H^{-1}$, for any full-rank $k \times k$ matrix H, instead of the original factors F_t. In essence, a parameter (or quantity) is rotation invariant if it is identical in the original model or in any rotated model (Giglio and Xiu, 2021). Specifically, defining $\hat{\hat{\gamma}} \equiv \hat{\gamma}H^{-1}$ and $\hat{\hat{\eta}} \equiv \hat{\eta}H^\top$, it holds that

$$\hat{\lambda}_g = \hat{\gamma}\hat{\eta}^\top = \hat{\gamma}H^{-1}H\hat{\eta}^\top = \hat{\hat{\gamma}}\hat{\hat{\eta}}^\top. \tag{13}$$

Importantly, neither γ nor η by itself is rotation invariant, because $\hat{\hat{\gamma}} \equiv \hat{\gamma}H^{-1} \neq \hat{\gamma}$ and $\hat{\hat{\eta}} \equiv \hat{\eta}H^\top \neq \hat{\eta}$. Similarly, the risk exposures of assets to the rotated factors differ from the exposures to the original factors

[11] Note that the factors extracted using the RP-PCA method are return-based with unrestricted means. Hence, under no-arbitrage, factor prices of risk equal their means, i.e., $\gamma = \mu_F$. However, the second pass is still useful, as it allows us to determine the uncertainty around the estimates (which will be accounted for in the computation of the asymptotic standard errors of the candidate factors' prices of risk) and evaluate the fit of the latent-factor model.

($\hat{\hat{\psi}}_n \equiv \hat{\psi}_n H^\top \neq \hat{\psi}_n$). Thus, unless one knows the rotation matrix H, not all original parameters can be recovered. Yet, even without knowing H, any consistent estimator of $\hat{\hat{\gamma}}\hat{\hat{\eta}}^\top$ will consistently estimate the candidate factor risk premium, $\hat{\lambda}_g$.

3 Test Assets and Factors

In this section, we first describe the exchange rate data, and explain how excess returns are computed. We then present the currency portfolios (test assets), and the nontradable macro-financial factors (candidate risk factors).

3.1 FX Data and Excess Returns

`FX Data.` We collect spot exchange rates and one-month forward rates vis-à-vis the U.S. dollar (USD) from Barclays and Reuters, available via Datastream. We take the perspective of a U.S. investor, and define the exchange rate as units of USD per unit of foreign currency (FCU), that is, USD/FCU. Hence, an increase in the exchange rate corresponds to an appreciation of the foreign currency. The empirical analysis is based on monthly data obtained by sampling end-of-month FX rates from October 1983 to December 2017. Our sample covers 49 currencies, of which 15 are regarded as developed countries following standard definitions in prior literature (e.g., Menkhoff et al., 2012a). It is important to note that the sample size is not fixed, given that it varies over time as data for some currencies are not available from October 1983, or some currencies cease to exist due to the adoption of the euro. That is, we work with an unbalanced panel of individual currencies. We provide detailed information on the FX data in the Internet Appendix (Section II).

`FX Excess Returns.` Currency excess returns are defined as follows

$$X_{it+1} \quad = \quad \frac{S_{it+1} - F_{it}}{S_{it}}, \tag{14}$$

where, using notation local to this subsection, F_{it} is the forward exchange rate that matches the spot exchange rate S_{it} for currency i (Bekaert and Hodrick, 1993). According to Eq. (14), the excess return results from buying the foreign currency in the forward market at time t, and selling it in the spot market at time $t+1$. As a matter of convenience, throughout this paper we refer to the forward premium $fp_{it} = \frac{S_{it}-F_{it}}{S_{it}} \approx i_{it} - i_t$ as either the forward premium or interest rate differential relative to the U.S. dollar, with i_{it} and i_t denoting the foreign and U.S. interest rate, respectively. Indeed, under covered interest parity (CIP), the interest rate differential is equal to the forward premium.[12]

3.2 Test Assets

A large cross section of test assets is central to the validity of the GX three-pass method (Giglio and Xiu, 2021). Our test assets are currency portfolios rather than individual currencies. By using portfolios, we can average out idiosyncratic components of currency returns and focus only on their systematic risk (Cochrane, 2005). Moreover, portfolios dynamically include individual currencies as their returns and signals become available, resulting in a balanced panel of test assets.

We consider currency portfolios associated with widely-used trading strategies. Overall, the baseline sample consists of $N = 46$ currency portfolios that stem from nine popular investment strategies: carry (e.g., Lustig et al., 2011; Menkhoff et al., 2012a), short-term and long-term momentum (e.g., Asness et al., 2013; Menkhoff et al., 2012b), currency value (e.g., Asness et al., 2013; Kroencke et al., 2014; Menkhoff et al., 2017), net foreign assets and liabilities in domestic currencies (Della Corte et al., 2016), term spread (Bekaert et al., 2007; Lustig et al., 2019), long-term yields (Ang and Chen, 2010), and output gap (Colacito et al., 2020). In what follows, we refer to these strategies as Carry, ST and LT Mom, Value, NFA, LDC, Term, LYld, and GAP, respectively.

We provide a detailed description of each investment strategy in the Internet Appendix (Section II); here we note that these strategies differ in the signals used to allocate currencies into portfolios (e.g., interest rate differentials, past returns, etc.), but the sorting schemes are similar. In fact, all strategies are tradable and rest on single sorts (with the exception of LDC, which uses double sorts on net foreign assets and the proportion of foreign currency denomination of liabilities).[13] At time t, currencies are allocated to NP

[12] As is usual in the literature, we compute FX excess returns using forward rates rather than interest rate differentials for two main reasons. First, marginal investors (such as, e.g., hedge funds and large banks) that are responsible for the determination of exchange rates trade mostly using forward contracts (e.g., Koijen et al., 2018). Second, for many countries, forward rates are available for much longer time periods than short-term interest rates. It is reasonable, however, to exclude the months when CIP is strongly violated; in doing so, we follow Kroencke et al. (2014) and Della Corte et al. (2016), among others (see Section II, in the Internet Appendix).

[13] We refer to single sorts when a single trading signal is used to sort currencies into portfolios. Conversely, we refer to double sorts when two trading signals are sequentially used to sort currencies into portfolios.

portfolios using the past signal for the selected strategy. Then, for a generic portfolio n, the excess returns realized between time t and $t+1$, X_{nt+1}, are computed as the equally-weighted average of the individual currency excess returns allocated to that portfolio. In line with most of the FX literature, we use $NP = 5$ for single-sorted portfolios. By construction, as we move from portfolio 1 (P1) to portfolio 5 (P5), the portfolios should contain currencies with increasing riskiness. Hence, if the risk-return trade-off holds, the spread portfolio (HML) – the return difference between P5 and P1 – should give a positive return because P5 contains currencies with high risk, whereas P1 includes currencies with low risk.

3.3 Nontradable Candidate Factors

We now turn to the nontradable (or non-return-based) candidate risk factors for which we aim at estimating the price of risk. These factors feature in the last pass of the GX three-pass method (see Section 2.1.2), which is implemented separately for each candidate factor. Therefore, the choice of a candidate factor does not affect the analysis of the other factors or the optimal SDF.

To begin with, we consider a reasonably long list of macro factors. In this way, we shed light on the link between the macroeconomy and asset returns, which is a central issue in macro finance (Cochrane, 2017). While the link is clear in theory, it is hard to establish empirically. Currency returns are no exception in this regard, and the disconnect is possibly even more evident than in other financial markets. In theory, currency returns and macro fundamentals are tightly linked together (e.g., Hassan, 2013; Gabaix and Maggiori, 2015; Ready et al., 2017; Berg and Mark, 2018a). In reality, the link between the two is weak (Meese and Rogoff, 1983; Mark, 1995), or highly unstable (Rossi, 2013; Fratzscher et al., 2015). A recent finding, however, is that macro fundamentals seem to be strongly connected to the *cross section* of currency returns (e.g., Colacito et al., 2020; Dahlquist and Hasseltoft, 2020). That said, macro fundamentals are often poorly measured and are clearly nontradable factors, so that both sources of bias that we address in this paper are likely to be sizable.

We also consider another set of nontradable factors that is gaining momentum in the asset pricing literature, which pertains to financial conditions. The global financial crisis has spurred an extensive literature on volatility and liquidity risks (e.g., Menkhoff et al., 2012a; Karnaukh et al., 2015), uncertainty shocks (e.g., Bekaert et al., 2013; Dew-Becker et al., 2017), and the leverage of financial intermediaries (e.g., Adrian et al., 2014; He et al., 2017). These factors are also nontradable, and their measurement is often imprecise. Furthermore, some of these measures are global, while others focus on the U.S. market. Most measures are specific to equities and bonds, but there is by now overwhelming evidence that financial and uncertainty shocks can easily propagate across markets. We therefore attempt to capture such complexity by using multiple popular measures of (il)liquidity, volatility, and uncertainty.

In addition to macro and financial variables, we also extend the analysis to the recently developed text-based factors, which are obtained by aggregating into an index news coverage about specific sources of uncertainty. Text-based indicators based on news coverage of policy uncertainty and, more recently, of equity market volatility are becoming increasingly prominent in the literature (e.g., Baker et al., 2016; Baker et al., 2019). Their sub-categories are also particularly informative about asset returns (e.g., Giacoletti et al., 2021). We therefore estimate also the risk premia of many text-based factors. By doing so, we broaden the measurement of macro-financial risks.

Taken together, the list of nontradable candidate factors consists of a total 133 factors, which we find useful to group as financial (23), text-based (30), and macro (80). However, the following additional observations are in order. First, the set of factors is comprehensive but by no means exhaustive, mainly because some factors are not available at the monthly frequency. Second, the distinction across categories is largely adopted for convenience, not being exact for some factors, especially for those capturing multiple sources of risks. In this sense, we do not view these variables as separate and *true* risk factors (also because they are measured with error), rather as macro-financial variables that relate to the 'true' factors.[14] Finally, as is common in the asset pricing literature (Merton, 1973), we do not use the factors as such but we first convert them into innovations, capturing the unexpected changes in the factors (in the baseline analysis, we simply use the residuals from AR(1) processes as, e.g., in Menkhoff et al., 2012a). A brief overview of the candidate factors and more detailed motivations for selecting them are provided in the Appendix.

[14]For example, the latent factors of Jurado et al. (2015) are placed in the group of financial factors, but they also contain macroeconomic information. Similarly, some text-based factors measure uncertainty related to the macroeconomic environment.

4 Empirical Analysis

In this section, we present the main findings of the empirical analysis. To start with, we assess the cross section of portfolio currency returns using descriptive statistics (Section 4.1). We then present the RP-PCA estimates of the latent factors and shed light on the properties of the optimal currency SDF (Section 4.2). Next, we turn to the three-pass estimates of the risk premia of the nontradable candidate risk factors (Section 4.3). Finally, we present a number of robustness checks and additional analysis on the stability of the factor structure (Section 4.4).

4.1 Currency Portfolios

Table A3, in the Internet Appendix, presents summary statistics of the currency portfolios, i.e., our test assets, associated with the nine investment strategies described in Section 3.2. We find that 22 out 46 individual portfolios deliver statistically significant returns. Importantly, all strategies deliver spread HML portfolios (denominated as CS in Table A3 as these are cross-sectional portfolios) with positive and statistical significant average returns, with the exception of the LYld HML portfolio. Since HML portfolios are self-financed long-short strategies, they represent U.S. dollar-neutral strategies. Moreover, for these HML portfolios the average excess return is the price of risk, given that it has unit exposure by construction. Therefore, the average return of the HML portfolio is a key statistic to look at. Carry, ST Mom and GAP HML portfolios yield the highest expected excess returns (7.3, 6.9, and 6.7 percent per annum, respectively), while LYld and Term HML portfolios display the lowest excess returns (1.9 and 2.8 percent per annum, respectively).

We then resort to an intentionally simple exercise to visually illustrate the risk-return trade-off inherent in the currency portfolios. Figure 1, top panel, shows that “low-signal” portfolios (P1, P2) tend to be mostly placed in the bottom left-hand corner (low risk/ low return), whereas “high-signal” portfolios (P4, P5) in the top right-hand corner (high risk/high return). In short, with few exceptions, higher returns seem to compensate for higher risks, which is consistent with the existence of a risk-return trade-off in currency portfolios. In the bottom panel, we instead present the time series of HML portfolios’ cumulative returns, which clearly show the higher performance of Carry, ST Mom and GAP investment strategies, but also that the underlying sources of risk differ. For example, during the global financial crisis ST Mom strongly outperforms Carry.

Pair-wise correlations among HML portfolios (Table A4, in the Internet Appendix) are also of interest, as they provide a first piece of suggestive evidence on the factor structure of the optimal currency SDF. Above all, Carry appears to be a dominant strategy, as not only it yields the highest returns, but it also correlates positively with many other strategies. LYld and, to a lower extent, Term portfolios strongly correlate with Carry (76 and 54 percent, respectively), but on the backdrop of substantially lower returns than Carry. The global imbalances HML portfolios, LDC and NFA, also appear to be tightly linked to Carry. Hence, Carry singles out as a pervasive strategy. Moreover, and perhaps not surprisingly given that the signals of both strategies depend on past returns, the excess returns of ST and LT Mom portfolios are positively related, but their correlation is not particularly high (around 25 percent), possibly due to mean reversion in returns. However, in absolute terms, the LT Mom spread portfolio co-moves mostly with the Value portfolio (around –39 percent). Also, in line with the extant literature (e.g., Koijen et al., 2018), we find that momentum strategies have low correlation with Carry, suggesting that they may be driven by a separate risk factor. Similarly, the GAP strategy exhibits a high risk premium, and yet displays particularly low correlations with respect to the other strategies.

In sum, while some strategies are strongly related, others are weakly or even negatively related, suggesting that multiple sources of risk drive currency portfolio returns. It is, however, ex-ante unclear how many slope risk factors – other than Carry – are needed to capture the risk-return trade-off in the FX market. Next, we turn to assess more formally the structure of the optimal latent-factor currency SDF.

4.2 Currency Pricing Kernel

To begin with, we extract the latent factors from the panel of currency portfolio returns using different values of the RP-weight.[15] We contrast the estimates from models using PCA with those from models using RP-PCA with higher and increasing values of the RP-weights. These models can differ in terms of the detection of the factors, the factor compositions, and the order of the factors. To highlight differences across

[15]Both PCA and RP-PCA require a balanced panel of test asset returns. To fill the few missing observations in test asset returns X (GAP portfolios are available only until January 2016), we use the nuclear-norm penalized regression approach, recently employed by Giglio et al. (2021b), to which we refer for more details on the procedure. We find that the main empirical results are robust to the method used to recover X, for example if we replace the missing returns with the sample mean of returns (which we did in a previous draft of this paper).

models, we first assess the factor "signal strengths", also recurring to statistical tests. We then complement this statistical information with the model diagnostics, based on the three economic evaluation criteria presented in Section 2.1.1, to better determine the dimension and properties of the SDF. Next, we present the main results, while the more detailed analysis of the latent factor "signal strengths" and, relatedly, of the SDF are presented in the Internet Appendix (Section IV). Finally, we try to link the extracted latent pricing factors to the observable currency strategies.

`Latent Factor "Signal Strengths".` As explained in Section 2.1.1, the ability to detect a pricing factor depends on the factor's signal-to-noise ratio. In RP-PCA, the "signal strength" of a factor is captured by the associated eigenvalue of the matrix $(\frac{1}{T}X^\top X + \omega \bar{X}^\top \bar{X})$, call it Σ_{RP}. To begin the SDF analysis, it is therefore useful to inspect the behavior of the largest eigenvalues of Σ_{RP}, both plain and normalized by the idiosyncratic variance, as the latter more closely relate to the signal-to-noise ratio. At the same time, it is useful to assess what drives a factor's overall "signal strength", by simply inspecting its composition. In this way, we try to establish whether a factor (i) is strong or weak, and (ii) with high or low Sharpe ratio. We report this (and additional) evidence in Table A8, and present the main results in what follows.[16]

The first eigenvalue is large and hence is symptomatic of a systematic, strong factor. On the contrary, the remaining eigenvalues are substantially smaller, so that the associated factors are consistent with a weak-factor interpretation (i.e., factors that explain a small set of test assets). This evidence is qualitatively similar when using either PCA or RP-PCA to estimate the factors. However, by employing the latter method (and with reasonably high RP-weights), it is evident that factors' "signal strengths" are enhanced, factors are better separated from each other, and the information is aggregated in a small number of factors. In doing so, RP-PCA helps us estimate the factors more efficiently, as documented also by LP for characteristic-sorted stock portfolio returns.

In particular, we find that RP-PCA enables us to detect weak factors with high Sharpe ratios, which are missed by standard PCA. The third factor is a clear example in this regard. In fact, the O and GX tests point to a two-factor SDF when the factors are extracted using PCA, and to a three-factor SDF when

[16]In the empirical analysis, we limit the focus to the six largest eigenvalues/factors, as the remaining eigenvalues have negligible "signal strengths". Also note that, following LP, we normalize the loadings such that the factors are orthogonal with each other and have different means and variances. Specifically, we adopt the Gram-Schimdt method, which has the benefit of orthogonalizing the factors sequentially. The models based on the original and orthogonal factors are observationally equivalent. Factors are orthogonalized mainly to facilitate their economic interpretation, e.g., regarding their distinct contributions to the currency SDF.

using RP-PCA (Figure 2 shows the eigenvalues and the results of the tests for different models). This is exactly because RP-PCA overweights the sample mean returns, and hence enhances the *overall* “strength” of the factor, on the backdrop of essentially unchanged *time-series* “strength”, which neglects the effect of the factor means. Put simply, the third factor extracted via RP-PCA appears to be a weak factor but with high Sharpe ratio (or risk premium). These types of factors are particularly hard to identify, and yet have important asset pricing implications, exactly because of their large risk premia. Thus, their omission is likely to distort the candidate factor risk-premium estimates.

Using RP-PCA, also the fourth factor qualifies as a weak pricing factor, given that it has a positive and significant risk premium (see Table A8). To a larger extent, $\hat{F}_4$ behaves qualitatively like $\hat{F}_3$, but its Sharpe ratio is substantially smaller; in fact it is half of that of the third factor (despite the two having comparable time-series “strengths”). This suggests that the pricing contribution of $\hat{F}_4$ is not sufficient to increase its signal-to-noise ratio in a way that $\hat{F}_4$ is selected by the statistical tests.[17] The remaining factors have zero risk premia, and hence are time-series factors.

Therefore, based on the “signal-strength” analysis, the optimal SDF should include at least the first three pricing factors extracted via RP-PCA. Next, to complement the above evidence, we proceed with the analysis of the other evaluation criteria used to better inform the choice of the RP-weight. By doing so, we also try to further characterize the properties of the optimal latent-factor currency SDF, going beyond its dimension.

Optimal Currency SDF ($\varphi(F_K^{\omega})$). The selection of the RP-weight responds to the dual objective of achieving a model with good pricing performance and high SDF maximal Sharpe ratio, while preventing idiosyncratic variance to increase too much (e.g., Lettau and Pelger, 2020b). Therefore, in theory, there might be a trade-off such that higher (lower) RP-weights imply lower (higher) pricing errors (e.g., $\overline{RMS}_{\alpha}$) at the cost of higher (lower) idiosyncratic variance ($\overline{\sigma}_{\epsilon}$). We evaluate the trade-off for a range of RP-weights in Table 1. The results are clearcut: there is virtually no evidence of trade-off in choosing RP-PCA over PCA (in line with the signal-strength analysis of Table A8).

Specifically, the idiosyncratic variance increases only slightly with the RP-weights, while the reduction in the pricing errors is substantial (*Panel A*). However, the marginal gains in terms of pricing performance

[17]However, consistently with GX, we will show in simulation that these tests tend to underestimate the true number of factors in finite samples. Hence, the evidence resulting from these tests needs to be taken with caution.

obtained by using very large RP-weights are small. In fact, pricing-error statistics tend to stabilize for $\omega \geq 20$, and we do not see additional benefits in using RP-weights higher than 20. This choice of the RP-weight is further corroborated by the analysis of the maximal SRs (*Panel B*), which tend to level off for $\omega \geq 20$. Moreover, the SR of the three-factor SDF obtained with $\omega = 20$ is 0.45, while it drops to 0.26 using $\omega = -1$.[18] Such a wedge is almost equally due to $\hat{F}_2$ and $\hat{F}_3$ when moving from $\omega = -1$ to $\omega = 20$. By adding $\hat{F}_4$ to the SDF, the SR further increases to 0.48 using RP-PCA with $\omega = 20$, while is unchanged using PCA. The SDF-weights of $\hat{F}_4$ display a qualitatively similar pattern to those of $\hat{F}_3$, but $\hat{F}_4$'s contribution to the maximal SRs is smaller.

By looking at the SDF-weights the distinction between time-series and cross-sectional pricing factors becomes apparent. The former do not enter the SDF, while the latter take non-zero weights and contribute to price currency portfolios. Regardless of the RP-weight, two out of the six extracted factors are time-series factors. However, using RP-PCA instead of PCA, the order of the factors changes. In fact, using PCA the fourth and sixth factors are time-series factors, whereas with RP-PCA the fourth factor becomes a pricing factor, so that the fifth and sixth factors are time-series factors.[19] Nevertheless, the pricing contribution of $\hat{F}_4$ appears small compared to that of the other two weak pricing factors (i.e., $\hat{F}_2$ and $\hat{F}_3$), which helps explain why $\hat{F}_4$ is not selected by the statistical tests.

Taken together, this analysis suggests that the optimal latent-factor currency SDF should consist of at least three (and potentially four) factors, and an RP-weight of 20 seems a plausible choice. Moreover, RP-PCA appears to change materially the information spanned by the factors relative to PCA, in a way that the estimated factors should be more efficiently estimated and closer to the 'true' pricing factors, which complements the evidence uncovered by LP for equity portfolios. We test the robustness of these results along several dimensions, including the RP-weight, the SDF dimension, and out-of-sample analysis in Section 4.4. Next, we assign an economic interpretation to the pricing factors.

[18] Such difference becomes even starker if one compares the respective optimal SDFs; in fact, the SR of $\varphi(F_{1-3})$ with $\omega = 20$ is roughly three times higher than the SR of $\varphi(F_{1-2})$ with $\omega = -1$. We find similar evidence for the pricing errors. Thus, RP-PCA achieves both lower pricing errors and higher SRs than PCA. However, as discussed later, this gain is smaller when moving out of sample.

[19] Using RP-PCA, we obtain an SDF that consists only of factors with significant means. This resembles the robust SDF of Kozak et al. (2020), KNS henceforth. However, KNS extract factors using PCA and then impose sparsity, by dropping factors with means below a threshold. Moreover, factors' SDF-weights shrink toward zero relative to the mean-variance weights. Thus, both the factors and their weights differ from ours, as RP-PCA changes the construction of the factors and relies on the mean-variance weights (Lettau and Pelger, 2020b). What is common between the two, however, is that both methods "overweight" the information in the first moments.

Risk Exposures (ψ) and Factors Interpretation. Thus far, we extracted the latent factors and determined the structure of the currency SDF. Hence, we are ready to implement the three-pass method of GX. However, note that the estimates of the risk exposures and of the latent factors' prices of risk – the first-two steps of the GX method – are already subsumed into the RP-PCA factor estimation. On the one hand, the risk exposures, i.e., the currency portfolios' loadings on the latent factors, are implicit in Eq. (8). On the other hand, the point estimates of the latent factors' prices of risk are simply given by the factors' mean returns. This differs from GX as in their case factors are demeaned, and hence their risk prices are inherently model dependent and need to be estimated using the second pass of the FMB procedure. Next, we assess the portfolios' risk exposures and factor-by-factor explained variations, and in doing so relate the factors to the investment strategies.[20]

To begin with, we find that portfolios' exposures to the first factor, $\hat{F}_{1t}$, are positive and roughly equal across portfolios (not reported). This evidence is consistent with a level, strong factor interpretation. In short, the first factor resembles the "Dollar" factor (e.g., Verdelhan, 2018). Figure 3 presents the exposures (left panel) and explained variations (right panel) of the HML portfolios of the nine investment strategies to the other estimated orthogonalized factors. To present the main findings, we focus on the HML portfolios as they are arguably more interesting than individual portfolios. Moreover, in this way, we can visualize the evidence for all strategies in a clear and concise manner.[21]

The second latent factor, $\hat{F}_{2t}$, retains a clear interpretation as it mostly relates to "Carry". In fact, the Carry spread portfolio displays a strong positive exposure to $\hat{F}_{2t}$, and the associated R^2 is roughly 70 percent. Moreover, Carry portfolios' exposures to this factor increase monotonically, as we move from P1 to P5. All other HML portfolios are, to some extent, positively exposed to $\hat{F}_{2t}$. In terms of R^2s, $\hat{F}_{2t}$ mostly explains LYld, Term, LDC, and NFA spread portfolios, while it is substantially less relevant for momentum,

[20] As explained before, we center the empirical analysis around the orthogonalized latent factors. By doing this, we can also easily determine the distinct contribution of each factor in explaining portfolio returns. Moreover, while we established that the optimal SDF consists of three factors, we perform the analysis using all six latent factors (i.e., $K = 6$). In this way, we can also establish which investment strategies mostly drive the factors left out of the optimal SDF.

[21] Given that the HML portfolios are not included in the sample of test assets, their risk exposures are derived ex-post from the corner portfolio exposures of the nine investment strategies. Meanwhile, we evaluate the six factors' individual contributions to the nine HML portfolios' explained variations using Eq. (7). Specifically, we estimate 9×6 OLS time-series regressions (i.e. six regressions for each of the nine HML portfolios), as we add factors one by one; hence, we consider SDFs of increasing dimension. In the Internet Appendix, Figures A1 and A2 present, respectively, the individual portfolios exposures ($\hat{\psi}_n$) and explained variations (R_n^2) by latent factor. We omit to plot portfolios' exposures and explained variations associated with the first factor, as it becomes easier to visually detect the exposures and marginal contributions of the other factors.

value, and GAP strategies.

Conversely, $\hat{F}_{3t}$ is tightly linked to momentum investment strategies, and especially with ST Mom. In fact, ST Mom portfolios' loadings on $\hat{F}_{3t}$ increase from large negative (P1) to positive (P5) values, and the pattern across portfolios is almost monotonic. Similarly, LT Mom corner portfolios load with opposite signs on $\hat{F}_{3t}$. The GAP spread portfolio is also positively exposed to $\hat{F}_{3t}$. Notably, GAP portfolios' exposures to $\hat{F}_{3t}$ strongly resemble those of LT Mom portfolios. Hence, we uncover a novel relation between price momentum strategies and a macro strategy such as GAP.[22]

We can therefore conclude that the optimal currency SDF consists of at least a "Dollar" factor, a "Carry" factor, and a "Momentum" factor. Interestingly, while all strategies' spread portfolios are positively exposed to the first two factors, most strategies are negatively exposed to the "Momentum" factor. Thus, in line with Lustig et al. (2011) and Verdelhan (2018), we find that the currency SDF includes the "Dollar" and "Carry" factors. However, we show that an additional "Momentum" factor should also feature in the SDF. In this sense, our evidence more closely echoes that in Chernov et al. (2021), albeit uncovered using a different methodology and currency universe. Next, we turn to analyze the remaining factors.

Of particular interest is $\hat{F}_{4t}$, given that it also displays a positive risk premium, and its inclusion in the SDF increases somewhat the maximal SR and reduces the pricing errors. Moreover, it has an intuitively clear interpretation of "long Value short (long-term) Momentum" factor, or simply "Value" factor. In fact, it presents a close nexus with the Value spread portfolio. P1 and P5 Value portfolios display negative and positive loadings on $\hat{F}_{4t}$, respectively. The loadings of the middle portfolios reveal a monotonically increasing pattern. Meanwhile, it is also apparent the strong association between $\hat{F}_{4t}$ and the LT Mom spread portfolio. However, LT Mom portfolios display monotonic but decreasing exposures to $\hat{F}_{4t}$. Therefore, $\hat{F}_{4t}$ could partly be responsible for the negative correlation between (long-term) momentum and value strategies documented in Table A4 and by Asness et al. (2013) for many other asset classes.[23]

Latent Factor Prices of Risk (γ). As a first cross-validation exercise, we check that the model-free estimates of the latent factors' prices of risk (i.e., the factor means) match those obtained using the second-

[22] This differs, to some extent, from the weak association between price and economic momentum strategies (Dahlquist and Hasseltoft, 2020). At the same time, the remaining spread portfolios are mostly exposed negatively to $\hat{F}_{3t}$ (LYld and NFA, in particular).

[23] The $\hat{F}_{5t}$ and $\hat{F}_{6t}$ risk exposures, albeit relevant for some specific portfolios, display no clear patterns. Thus, it is hard to assign a precise interpretation. This is not surprising exactly because they are time-series factors.

pass of the GX method, i.e., the FMB estimates. We find that the two point estimates are equal (i.e., $\hat{\mu}_F = \hat{\gamma}$), which is reassuring as it shows that the no-arbitrage assumption is preserved (e.g., Cochrane, 2005). Specifically, for the selected SDF with $\omega = 20$, the means (or prices of risk) of the pricing factors are roughly $\hat{\mu}_{F,1}$= 18.7, $\hat{\mu}_{F,2}$= 11.3, and $\hat{\mu}_{F,3}$= 9.2 percent per annum. Based on Newey-West standard errors, $\hat{\mu}_{F,1}$ is statistically significant at the five percent level, while the remaining two at the one percent level (see Table A8). The estimates of the prices of risk of the latent factors are not particularly informative per se, but they are a crucial input in the estimation of the risk premia of the nontradable factors, λ_gs. To obtain a robust estimate of λ_g, it is important that the latent-factor model does a good job in pricing the test assets. Only if this is the case, one can argue that the price of risk estimates of the candidate factors are not affected by omitted-variable and measurement-error problems.

Test Assets' Pricing Errors. The superior pricing performance of the RP-PCA model with $\omega = 20$ relative to the PCA model already emerges in Table 1. To better appreciate the differences between these two estimation methods, Figure 4 plots realized versus model-implied average portfolios' excess returns. If a model prices perfectly the cross section of portfolios' returns, all data points lie on the 45 degree line. We find that the pricing performance of the RP-PCA model based on the optimal three-factor SDF is very accurate: the pricing errors are indeed small, and there is no tendency for the model to systematically misprice portfolio returns. By inspecting the two-factor SDF evidence, the gain from adding an extra factor clearly emerges. This reiterates the importance of including the "Momentum" latent factor to the currency SDF, in addition to the "Dollar" and "Carry" latent factors. However, for a factor model to reflect the 'true' SDF, all cross-sectional pricing errors should be zero on average (or should have zero alphas).

A natural way to proceed would be to perform standard tests of the null hypothesis that the alphas are jointly zero. However, these types of tests – such as for example the Gibbons-Ross-Shanken (GRS) test – are based on the assumption that N is constant and $T \to \infty$, while the RP-PCA estimator is derived under the assumption that N, $T \to \infty$. This implies that the covariance matrix no longer converges to the population matrix, and the tests are biased even in large samples (e.g., see Lettau and Pelger, 2020b). While some papers propose remedies to obtain consistent estimates of the covariance matrix (see Giglio et al., 2021a for a detailed review), and hence address some of the drawbacks of the GRS test, we opt for a simple multiple-testing approach. Our reasoning is as follows: the RP-PCA estimator assumes a factor model with zero alpha (or a zero-beta rate), so that each asset should have zero average pricing errors in Eq. (1), and as a result the model would span the entire asset space. Hence, we perform $N = 46$ tests of the null hypothesis $\mathcal{H}0$: $\bar{\epsilon}_n = 0$ for $n = 1, \ldots, N$, and assess how many assets have significant average pricing errors at the 5% significance level, as we vary the dimension of the SDF. The results are clear-cut. Using a two-factor model we reject the null hypothesis in 14 cases (e.g., including the ST Mom corner portfolios), while only one test asset (Value-P5) has a statistically significant average pricing error using the three-factor model. With the four-factor model, none of the assets has significant average pricing errors. In short, a four-factor model ensures that the average pricing errors are zero for all test assets, although we would argue that the three-factor model likely achieves the same outcome.[24]

Finally, the comparison between left and right panels of Figure 4 highlights the lower pricing performance of the PCA models. This is evident if one compares models' with SDFs of equal size and, even more, if one contrasts the evidence for the respective optimal SDFs (i.e., two- and three-factor SDFs for PCA and RP-PCA, respectively). Overall, this analysis validates the use of $\varphi(F_{1-3})$ based on RP-PCA to determine the prices of risk of the candidate factors, to which we turn next.

4.3 Candidate Factor Risk Premia

`Spanning Regressions.` The last pass of the GX procedure consists of projecting each of the candidate nontradable factors onto the space spanned by the estimated latent factors. We do this by estimating the following regressions

$$g^{\iota}_{jt} = a_{jk} + \hat{F}_{1:kt}\eta^{\top} + u_{jkt}, \quad j = 1, \ldots, J, \quad k = 1, \ldots, K, \quad t = 1, \ldots, T, \tag{15}$$

where g^{ι}_{jt} is the AR(1) innovation of the selected j-th nontradable factor. We again perform the regression analysis by expanding the set of latent factors, by adding one factor at a time, $\hat{F}_{1:kt}$, to single out their marginal contributions in terms of R^2s. We therefore run a total of $J \times K$ time-series regressions. GX show that the R^2s help quantify the measurement errors in the nontradable factors. Specifically, a low (high) R^2

[24]Recalling that the three-factor model only produces one significant average pricing error (with a p-value of 0.035) out of 46 assets, this can clearly arise just by chance (false discovery). For example, using a Bonferroni correction, none of the p-values is smaller than the Bonferroni corrected p-value ($5\%/N = 0.0011$), although this is known to be a conservative approach (e.g., Giglio et al., 2021b).

implies a big (small) measurement error.[25]

`Explained Variation` (R^2). Figure 5 shows the R^2s associated with each of the three latent pricing factors, grouped by type of candidate factor (we present the evidence using all six latent factors in Figure A3 in the Internet Appendix). Within each group, candidate factors are sorted by the total R^2s. We find that the measurement-error problem is pervasive, as the R^2s are generally low. At the same time, some distinct patterns across types of candidate factors emerge. A few financial factors display R^2s that exceed 10 percent; these factors mostly relate to financial and liquidity conditions, volatilities, and intermediaries' leverage. For most of these factors, the overall R^2s are driven by all three factors, albeit mainly by the "Carry" factor. There are of course a few exceptions as is clearly the case for the global financial condition index. (Its R^2 is above 35 percent, and is mostly driven by $\hat{F}_1$, although the absolute contribution of $\hat{F}_2$ is also large.)

In comparison with financial factors, the R^2s of text-based factors are generally lower. Moreover, the "Carry" factor is by far the most relevant factor, given that for many factors (especially for those with higher R^2s) it accounts for almost the entire R^2s. The U.S. EMV factors' R^2s tend to be higher than those associated with the U.S. EPU indices, with the exception of the global EPU index. Finally, turning to macro factors, they present much lower R^2s, at most in the range of 1-2 percent. The overall picture of the drivers is more mixed. In fact, for a number of macro factors, the "Dollar" and "Momentum" latent factors are the main determinants of the R^2s.

`Factor Exposures` (η). Before inspecting the estimates of the η-exposures, we note that all latent pricing factors are procyclical factors, as they command positive risk premia. In essence, procyclical factors rise in good states of the world, while dropping in bad states. It follows that a candidate factor with a positive (negative) η-exposure to a specific risk factor is procyclical (countercyclical) with respect to that source of risk. Thus, the η-exposures are economically meaningful objects. However, it is important to note that a candidate factor can present exposures of opposite signs to the individual pricing factors. Therefore, only the risk premium of the candidate factor will reveal whether the factor is either procyclical (positive

[25]Eq. (15) is the empirical counterpart to Eq. (10). It differs as we include the intercept and do not use the factors in deviation from their means. In this way, the R^2s are more meaningful, and yet the η-exposure estimates are unchanged. While the difference is negligible, when we compute the standard errors of the candidate factors' prices of risk using the formula in GX, we implement Eq. (10). Also note that, after having computed the factors' AR(1) innovations, we then standardize them such that the resulting factors have also unit variances, which helps compare the estimates across candidate factors.

risk premium), countercyclical (negative risk premium), or acyclical (zero risk premium) with respect to the state of the world.

Table 2 reports the estimates of the η-exposures of the factors. In the table, to help summarize the evidence, we focus on the candidate factors with significant prices of risk, based on the optimal SDF (we provide the evidence for all candidate factors in the Internet Appendix, Table A9). To start with, we note that financial factors tend to have exposures of the same sign to the "Dollar" and "Carry" factors (the latter are generally more precisely estimated). The signs of the η_2 estimates are consistent with the usual sources of risk inherent in currency carry strategies (e.g., liquidity and volatility risks) previously documented by the literature; however some specific risk factors are novel (e.g., otic, move). Interestingly, the exposures to the "Momentum" factor generally take opposite sign, but a smaller number of these exposures are statistically significant. Put differently, this evidence suggests that carry and momentum strategies respond to some of the same financial risk factors, but in opposite directions.[26] Turning to the text-based factors, the evidence is even more clearcut: none of the factors is exposed to the "Dollar" factor, while all factors are negatively exposed to "Carry". These factors' exposures to "Momentum" again take opposite sign with respect to "Carry", but are much smaller in absolute size, and are statistically significant only in a few cases (see Table A9).

Overall, we uncover a tight nexus between FX and other markets that is mainly channeled through the "Carry" factor, lending support to the argument of Koijen et al. (2018, p. 198) that "... [carry] could be a unifying concept that ties together many return predictors disjointly scattered across the literature from many asset classes." Finally, as expected, macro factors display only few significant η-exposures.

`Return-Based Factors` ($F\eta^\top$). To complement the above analysis, we visually inspect some examples of return-based factors. A return-based candidate factor is the original factor cleaned from measurement error, and converted into a return factor using the η-exposures, being a linear combination of the latent factors ($\hat{F}\hat{\eta}^\top$). In practice, the original nontradable factor becomes tradable by investing in the underlying currency portfolios.

Figure 6 shows selected candidate factors, transformed into return-based factors using SDFs of different dimension. In this way, we can appreciate how the return-based factors' evolutions and levels change as the

[26] Meanwhile, we observe that only few candidate factors present significant η-exposures to the other three factors, which nevertheless are excluded from the optimal SDF.

"Momentum" factor is added to the SDF. Top panels present two examples of factors (gepu and gvol) that are exposed significantly to the "Carry" factor, but not to "Momentum". Conversely, bottom panels refer to factors (icap and gliq) that are exposed significantly to both "Carry" and "Momentum" factors (Table A9), but with opposite and economically large estimates. It is evident that, for icap and gliq, by moving from the two-factor to the three-factor SDF the factor mean returns drop significantly in absolute terms, so that their risk premia eventually vanish. The in-depth analysis of the factors' risk premia is presented next.

Risk Premia. In what follows, we present the last piece of evidence resulting from the third pass of GX. That is, we report the estimates of the risk premia of the nontradable candidate factors, free from both the omitted-variable and measurement-error problems. As explained in Section 2 (and shown in Section 5 in simulation), one can also use the standard FMB two-pass procedure to obtain such estimates but, crucially, only if all relevant control factors are included in the SDF, and factors are measured without noise. The evidence reported so far clearly shows that the measurement-error problem is material. The omitted-variable problem is also likely to be important, and can therefore add to the measurement-error problem. To shed light on the severity of both problems, we first estimate factor risk premia by means of the standard FMB two-pass method.

FMB Two-Pass Method. We rely on univariate SDFs, which consist of a constant and the candidate factor at hand. Hence, for each candidate factor g_{jt}, we specify the SDF as $\varphi_t = 1 - g_{jt}^{\iota} b_j$. We intentionally omit from φ_t other potentially relevant risk factors, f_t, that could enter the SDF along with the candidate factor, g_t. In this way, the omitted-variable bias can manifest in its full strength. We find that over 90 out of the 133 candidate factors present estimates of the prices of risk that are statistically significant at least at the 10 percent level (of which more than half are the significant macro factors). Thus, the FMB estimates seem to point to a very large number of significant factors, i.e., a "factor zoo", for FX returns.

GX Three-Pass Method. Table 3 presents the factor risk-premium estimates obtained using the augmented three-pass method implemented with baseline RP-weight, $\omega = 20$. Along with the risk-premium estimates (λ_g), and the associated standard errors (se), the table reports the candidate factors' Sharpe ratios (SR), to better evaluate their economic relevance. It also presents the p-value for the Wald test that the candidate factor is weak (pval), where the null hypothesis is that the factor is weak, i.e., $\eta = 0$. We refer to

Giglio and Xiu (2021) for details on the computation of risk-premium standard errors and the weak-factor test p-value. A full analysis of the weak-factor problem is carried out later in Section 6.

Before turning to the individual factor estimates, we note that the three-pass absolute risk-premium estimates are substantially lower than the FMB ones. In Table 3, this finding holds to a large extent regardless of the specific SDF considered. Too high estimates of the prices of risk are likely to be caused especially by the measurement-error problem. In fact, when a factor is measured with noise, an attenuation bias characterizes the estimates of the portfolios' risk exposures to that factor in the first pass of FMB. This bias in turn leads to inflated prices of risk estimates in the second pass (e.g., Adrian et al., 2014). A close look at the table suggests that the problem seems to be less relevant for financial than for text-based and macro factors.

Based on the optimal SDF, $\varphi(F_{1-3})$, we find that 42 out of the 133 candidate factors have statistically significant risk premia, which is a much shorter list than that uncovered using FMB. Among these, we find that the financial factors with significant risk premia are 11 out of 23. The global volatility factor of Menkhoff et al. (2012a) singles out, as its risk premium is large and precisely estimated. The systematic FX liquidity measure of Karnaukh et al. (2015) is also priced, while the global liquidity measure of Menkhoff et al. (2012a) is not. Moreover, a number of factors relating to liquidity (noise and psliq) and volatility (move, vxo, and eqrv) conditions in the U.S. bond and equity markets turn out to have statistically significant risk premia. These factors highlight the tight link between FX returns and other markets. Interestingly, the quantity-based TIC flow measure (otic), proxying for foreign central banks' demand for US Treasuries, is positive and significant. Thus, it is a procyclical factor, possibly suggesting that foreign central banks tend to build up their reserves in good states of the world. Global financial conditions (gfc) also seem to matter for FX returns.

Turning to the text-based factors, 14 out of 30 are priced factors. In comparison with financial factors, the number of significant text-based factors drops even more substantially relative to FMB. The global EPU index of Baker et al. (2016) stands out as its risk premium is the highest (in absolute terms) and the most precisely estimated. Conversely, the U.S. EPU indices are not priced. At the same time, several EMV indicators display statistically significant negative risk premia. Some of the category-specific EMV trackers are even more precisely estimated than the overall index; namely, those relating to the macroeconomy and monetary policy (i.e., emv mout, emv mqnt, and emv mp).

To conclude the list of factors, we find that only a few macro factors are (weakly) priced in the cross section of currency returns; specifically, 17 out of 80 macro factors, stemming from six distinct macro factors measured at different frequencies. In particular, the world unemployment growth rate specified in differences versus the U.S. (unew/us) displays a negative risk premium, at many frequencies. The world industrial production factor also presents a statistically significant negative risk premium, especially when specified in differences versus the U.S. (ipw/us) at the quarterly frequency. The risk premium of world inflation (cpiw) is also negative. Consumption growth risk is negative and weakly significant (cus). However, unlike the other types of factors, several of the signs of the macro risk premia seem not to align with theoretical priors (e.g., Lustig and Verdelhan, 2007; Zviadadze, 2017).[27]

The absolute magnitude of macro risk premia is lower, on average, than that of text-based and financial factors with significant risk premia. Conversely, macro factors' SRs tend to be higher, especially those associated with unemployment. However, it is also evident that most of the macro factors are weak factors, according to the GX test (*pval*).

4.4 Robustness and Stability Analysis

To complete the baseline analysis, we perform a number of robustness checks, which are presented in detail in the Internet Appendix (Section V.1). The main results can be summarized as follows. First, we show that selecting the optimal SDF is key to obtain precise estimates of nontradable factors currency risk premia. In particular, the omission of relevant pricing factors (i.e., $\hat{F}_3$) is far more harmful than the addition of less relevant ones (i.e., $\hat{F}_4$). Second, the choice among reasonably high values of the RP-weight leads to small differences in the risk-premium estimates. Third, by including the HML portfolios to the panel of currency portfolios, some of the estimated latent factors better explain some of the high risk-premium currency strategies (e.g., GAP), but this information does not affect much the structure of the SDF, so that the method essentially selects the same relevant candidate risk factors as in the baseline results.

We then assess the stability of the SDF in both out-of-sample and in-sample time-varying settings, and provide some time-varying candidate factor risk-premium estimates (Section V.2). The in-sample results

[27]For example, considering the prospective of the U.S. investor, consumption growth risk is high in good states of the world and low in bad states. As a result, it should command a positive premium, while it turns out to be negative.

are largely confirmed out of sample. However, the performance of RP-PCA deteriorates out of sample, while PCA displays a much higher degree of stability. Therefore, there seems to be still a sizable gain in choosing RP-PCA over PCA, but this gain shrinks out of sample. Moreover, the out-of-sample evidence suggests more clearly a model with three factors, as the contribution of the fourth factor to the maximal Sharpe ratio is essentially nil. The factor structure appears rather stable over time, as indicated for example by the fact that the GX and O tests consistently point to a three-factor model throughout our recursive analysis (based on an initial window of ten years). Nevertheless, the candidate factors' risk-premium estimates do not show significant degrees of time variation, as long as the estimation window is sufficiently long, and the SDF includes at least the first three latent factors. Therefore, the unconditional three-pass model, if well specified, provides a satisfactory description of dynamically rebalanced FX portfolio returns.

5 Simulation Analysis

In this section, we study the finite-sample performance of the three-pass inference using Monte Carlo simulations. We also assess how the augmented three-pass estimator performs and compares in simulation with the two-pass estimator. Importantly, we design the simulations to capture the key features of FX returns. By doing so, we essentially tackle two key questions. First, is the three-pass method *reliable* in finite samples, with N and T equal to the dimension of our FX portfolio returns? Second, are the omitted-variable and measurement-error problems relevant for pricing currency portfolio returns, and hence is the method *desirable* for pricing currency risks?

Next we briefly describe the simulation exercise, and summarize the main findings. We present the fully-fledged simulation analysis in the Internet Appendix (Section VI).

`Simulation Method`. We set up the simulation exercise following closely Giglio and Xiu (2021), with the only relevant methodological difference due to the use of RP-PCA to extract the latent factors that drive the data generating process (DGP). However, we tailor the calibration to the specific features of the FX market. This is important because, although GX show a good performance of the three-pass estimator in simulation also for combinations of N and T that resemble the one used in our study, it is not obvious that the estimator performs the same in our case. The factor structures driving equity and FX portfolio returns may well differ, and this can in turn weigh on the estimator performance.[28] Specifically, in the simulations we consider a four-factor DGP consisting of the de-noised Dollar, Carry, ST Mom, and Value tradable factors, which most closely approximate the strong latent factor ($\hat{F}_1$) and the three weak latent factors with significant risk premia ($\hat{F}_2$, $\hat{F}_3$, and $\hat{F}_4$) documented in the empirical analysis. In essence, this DGP can be interpreted as a reduced-form model for FX returns. To remove the noise from the observed tradable factors, we use the three-pass estimator with four latent factors, extracted using RP-PCA with $\omega = 20$. These four de-noised factors should span the entire SDF (as we know from Section 4.2 that four latent factors do so and generate zero average pricing errors for all 46 assets), and hence the simulated asset returns should mirror the properties of the observed ones. This is crucial to recover the *true* risk premia via the three-pass estimator.

Next, as in GX, we assume that we may not observe all four factors but only noisy versions of them, plus a potentially spurious candidate nontradable factor. In this way, both the omitted-variable and measurement-error problems can manifest entirely. To begin with, we calibrate the candidate factor to U.S. industrial production (*ipus*) similarly to GX, which according to our three-pass estimates qualifies as a spurious factor also for FX returns. However, we then repeat the analysis replacing *ipus* with either *gvol* or *icap* (the global volatility of Menkhoff et al., 2012a, and the intermediaries' capital ratio of He et al., 2017, respectively). These two financial factors are of particular interest as they exemplify non-spurious factors whose risk premia estimates are affected to different extents by the dimension of the SDF considered (as is evident from Section 4.3, Tables 2, 3 and A9). Taken together, these three candidate factors capture the properties of the nontradable factors considered in the empirical analysis.

Finally, to evaluate the performance of the three-pass estimator, we estimate the risk premia of the noisy tradable and nontradable factors by applying the three-pass method to the simulated data. We use not only SDFs of expanding dimensions, but also with different RP-weights (this latter analysis is not in GX, as they are concerned only with the case of PCA). However, to be clear, the ultimate goal of this simulation

[28] In fact, the cross section of FX test assets is relatively small, and the underlying data are driven by fewer factors, compared to the case of equities, studied for example by GX. At least until recently, the benchmark model for FX returns has been the two-factor model of Lustig et al. (2011), consisting of a Dollar and a Carry factor. Our analysis, however, suggests that with a reasonably large N (at least much larger than the small cross sections typically used so far in the FX literature), a two-factor SDF is likely to omit relevant sources of FX risk, and that at least three and potentially four latent factors are required to achieve full spanning of the entire SDF and robust estimates of risk premia. Therefore, in larger cross sections, not only the measurement-error problem, but also the omitted-variable problem is likely to be relevant for FX returns.

exercise is not a comparison of the RP-PCA and PCA three-pass estimators. Rather, we are interested in assessing the performance of the augmented three-pass and two-pass estimators in pricing currency returns, i.e., when the DGP and the associated parameters driving the simulations match the properties of the FX portfolio returns studied in the empirical analysis. In this way, we can assess the finite-sample performance of the three-pass estimator and hence its reliability, but also shed light on the relevance of the issues of omitted factors and measurement error in the factors driving FX returns.

`Simulation Results`. The simulation analysis uncovers a number of important insights that lend support to the validity of our empirical analysis. The main results can be summarized as follows. To begin with, regarding the calibration (Table A14), we find that using the three-pass estimator applied to models including four factors, extracted with baseline $\omega = 20$, we obtain risk-premium estimates that are not statistically different from the factor averages. At the same time, the estimates are rather stable when the fifth and sixth factors are included into the SDF, consistent with full spanning of the entire SDF. Moreover, applying the three-pass estimator to *gvol*, *icap*, and *ipus* to recover their *true* premia, it clearly emerges that these factors display substantially different behaviors and premia.

`Simulation Accuracy and Factor Structure Recovery`. Turning to the simulations, we first verify the accuracy of the simulated asset returns, to ensure that the returns generated from the four-factor reduced-form model closely match the average returns, variances, and Sharpe ratios of the observed returns, as well as the cross-sectional standard deviation of observed average returns (Figure A11). This is an important check to ensure the reduced-form SDF is calibrated to generate artificial data that capture adequately the behavior of FX portfolio returns and spans the entire space of asset returns. We then show that, despite the relatively small sample size N, we can to a large extent recover in simulation the true factor structure of the return data (Tables A15 and A16). In fact, the analysis of the generalized correlations and Sharpe ratios point to an accurate recovery of the factor structure. At the same time, the simulations confirm the tendency to underestimate the true number of factors in finite samples on the basis of the statistical tests used in this paper, which is consistent with the evidence previously documented by GX. We find that this problem is attenuated, albeit not eliminated, by estimating the factors via RP-PCA, as the fourth latent pricing factor remains hard to detect. While an exact recovery of the factor structure is in itself important, the ultimate goal remains the ability of the three-pass estimator to recover the risk-premium estimates of the candidate factors in finite samples.

`Three-Pass Estimator`. We then turn to the estimation of the candidate factor risk premia using the three-pass method, which establishes two important results. First, the augmented three-pass estimator recovers the true tradable and nontradable factor risk premia in simulation (Table A17). Second, the GX central-limit result holds, which can be seen by the fact that the histograms of the bias in risk-premium estimates standardized by the asymptotic standard error match the standard normal distribution (Figures A12–A18). These results show that the augmented three-pass estimator is highly reliable also in finite samples that match the properties of our FX portfolio returns. Moreover, we note that the recovery of the true risk premia shows the rotation-invariance result – a general result shown by GX that is at the core of the three-pass estimator (see Section 2.1.2) – in our setting. This is because the asset-return DGP is driven by the de-noised tradable factors, but the risk premia are estimated via the three-pass method and, hence, as linear combinations of the latent factor risk premia. Furthermore, we find that for most factors we can recover the true risk premia even using parsimonious factor models (especially for higher RP-weights), but for some factors it is beneficial to use models that include more latent factors. Using five-factor models (thus an additional factor than in the true model), we can verify the central-limit results for all candidate factors. Thus, our simulation results are in line with those of GX and, notably, arise in a slightly different setting that matches the properties of FX portfolio returns. As a result, they lend additional support to the *reliability* of the three-pass estimator also in finite samples that match our data.

`Two-Pass Estimator`. The two-pass estimator analysis on simulated data shows that the omitted-variable problem can be material, leading to distorted risk-premium estimates (Table A18, and Figures A19-A21). The two-pass estimator recovers the true premia only if the SDF is correctly specified. In some cases, even if none of the relevant factors are omitted from the SDF, and yet some of the factors are measured with noise, we cannot retrieve the true risk premia of all factors. Therefore, both the omitted-variable and measurement-error problems manifest clearly in simulation, making the use of the three-pass estimator *desirable* for the estimation of currency investment-strategy risk premia.

Overall, the simulation results demonstrate the good performance of the three-pass estimator also in a currency setting. Moreover, they lend support to the argument made in Section 2: that is, omitting relevant pricing factors from the currency SDF, and/or measuring the factors with noise, can severely distort the two-pass risk-premium estimates. Put simply, the three-pass estimator appears to be both *reliable* and *desirable* for modeling FX portfolio returns, and hence represents a valuable method to unveil the sources of the risk-return trade-off in currency investment strategies.

6 Weak Candidate Factor Analysis

The three-pass procedure of GX tackles both the omitted-variable and measurement-error problems in the estimation of factor risk premia. However, the method as such – regardless of whether factors are extracted via PCA or RP-PCA – is not designed to explicitly address the issue of weak candidate factors. The weak-factor problem manifests if only a subset of the test assets is exposed to the candidate factor, and this can in turn disrupt the inference on the candidate factor's risk premium. Importantly, the strength of a factor is not an inherent property of the factor, as it also depends on the cross section of assets used in the analysis (Giglio et al., 2021c). Hence, a factor is not *strong* or *weak* in absolute terms, but relative to the cross section of test assets used by the researcher.

It also follows that a factor is more likely to be weak in large cross sections of test assets. At the same time, large cross sections are a prerequisite for the three-pass estimator to effectively address the omitted-variable problem, as the extracted latent factors and hence the SDF should span all relevant sources of risk. While the cross section of FX portfolios is small relative to, for example, equities, the problem of weak factors can still manifest to different degrees. For instance, it might weigh on the inference on macro factor risk premia, and thus can potentially help explain the disconnect between macro factors and currency portfolio returns, which emerges from our analysis. These considerations may also affect the risk-premium estimates of some financial and text-based factors, if some of these factors are weak. The previous analysis on the η-exposures shows, for example, that a few candidate factors have significant exposures only to a subset of the SDF, i.e., not to all pricing factors.

In light of these considerations, in what follows we assess the robustness of the candidate factors risk-premium estimates using the supervised principal component analysis (SPCA) recently proposed by Giglio et al. (2021c), GXZ henceforth. This novel method is designed to explicitly tackle the omitted-variable and measurement-error problems also accounting for the possibility that the candidate factor of interest is *weak*.

`SPCA Estimator`. Intuitively, the SPCA procedure delivers robust estimates of a weak factor's risk premium because it shrinks and adapts the assets' cross section using the information contained in the factor. By doing so, the factor can turn into a stronger factor with respect to the new, tailored cross section of assets, making the inference on the factor's risk premium more precise. Put simply, by using the factor-by-factor SPCA estimator, the SDF used to price a candidate factor is pinned down by the factor itself, through supervised selections of the relevant test assets. As a consequence, unlike the original three-pass method, the SPCA estimator does not hinge necessarily on a unique SDF common to all candidate factors, but potentially on several SDFs. This suggests that SPCA is not particularly useful to shed light on the properties of the currency SDF (the first goal of this study), but can be a useful tool to estimate the risk premia of the candidate factors (the second goal of the study), in the presence of weak factors.

The SPCA estimator relies on the two tuning parameters q and k, which depend on the candidate factor of interest. The former parameter determines how many test assets we use for factor extraction at each iteration, while the latter governs the number of iterations or, equivalently, the number of extracted latent factors. Both parameters are assumed to be known by the researcher, but GXZ show that they can be jointly determined in advance by repeating M times a $\mathcal{K}$-fold cross-validation exercise. For a given candidate factor, the exercise essentially consists of constructing a grid of out-of-sample R^2s for different combinations of q and k.[29] Notably, the R^2 denotes the fraction of the factor's variance explained out-of-sample by the return of the hedging portfolio, whereby the portfolio weights are constructed for a given combination of q and k. We then select the pair $\{q, k\}$ that yields the highest cross-validation out-of-sample R^2 for the candidate factor under consideration.

Before turning to the empirical results, two observations are in order. First, if the number of selected test assets qN equals the overall number of test assets N (i.e., there is no asset selection), then the SPCA factor risk-premium estimates are identical to the standard three-pass estimates (i.e., RP-PCA with $\omega = -1$). Therefore, for these estimates to be free from omitted-variable bias, the usual argument applies that the SDF needs to include enough latent factors to fully span the assets' space. Second, in the above sketch of the SPCA algorithm, the analysis is carried out for one candidate factor at a time, so that the asset selection is

[29]Specifically, for each m-repetition, we average over the grids obtained for the $\mathcal{K}$ possible permutations of training and testing periods. We repeat this procedure M times, and construct the final grid of out-of-sample R^2s by averaging over the M cross-validation grids. We refer to Section VII, in the Internet Appendix, for a more detailed description of the SPCA estimator and for a comprehensive analysis of the estimation results.

only driven by the factor at hand. By performing the SPCA estimation factor by factor, the risk-premium estimates are consistent and, importantly, we can determine which assets are relevant for which factors. However, factor risk premia can also be estimated via SPCA using more candidate factors simultaneously. In such joint estimation, the selection of the assets is driven simultaneously by a set that includes multiple (potentially all) candidate factors. Specifically, assets are sorted by the maximum correlation with any of the candidate factors in the set of factors considered. While both the one-by-one and the joint factor SPCA estimators are consistent, the joint estimation is required to conduct inference on the risk-premium estimates (see Giglio et al., 2021c). Intuitively, this is because, for the central-limit-theorem assumptions to hold, the candidate factors are required to have exposures to the entire SDF, which is a far more stringent requirement to be satisfied than is needed for consistency. Next, we turn first to present the risk-premium estimates obtained using SPCA, and then shed light on the identities of the most relevant assets for selected candidate factors.

`SPCA Results`. While the main focus of our analysis pertains to the estimation of nontradable factor risk premia, it is convenient to look first at the case of the *tradable* factors (i.e., the Dollar factor, and the nine HML factors associated with the currency investment strategies). This is because, for the tradable factors, we can benchmark the SPCA estimates to their sample averages, or model-free risk-premium estimates, and hence we can assess the performance of the SPCA estimator. The main results are displayed in Table A19, and can be summarized as follows. First, the factor-by-factor analysis (*Panel A*) shows that SPCA delivers risk-premium estimates that are close to the model-free premia for all tradable factors. For the HML factors, such close proximity of the two types of estimates is favored by the asset selection ($qN = 10$), coupled with SDFs including a sufficient number of latent factors ($k \approx 8$). Second, the joint analysis (*Panel B*) reveals that, using a ten-factor model, we can recover the model-free estimates of the risk premia for all tradable factors, and there is no substantial gain in using models including more latent factors. Overall, these findings suggest that SPCA is able to recover the factor averages, and a ten-factor model is a reasonable choice to carry out the risk-premium inference.

Table A20 presents the SPCA risk-premium estimates for the *nontradable* factors. In the first columns, the table shows the results of the cross-validation exercise. To start with, we use this evidence to distinguish between factors with positive and negative out-of-sample R^2s. We only report the estimation results for factors with positive R^2s, as a negative R^2 suggests that we cannot hedge that risk factor, and hence its risk-premium estimate is not informative (this is one reason why some factors appear in Table 3 but not in Table A20). Thus, the cross-validation exercise allows us to further filter out factors that cannot be hedged out of sample by the currency assets. This criterion can be seen as an additional way to further discern relevant candidate factors from non-relevant ones. This means that, while SPCA gives us the best chance to detect candidate factors with a non-zero risk premium by allowing for weak factors, it can also reduce the number of relevant factors due to this additional constraint.

While we present the detailed results in Section VII.2.1 in the Internet Appendix, here we note the following. Above all, the SPCA evidence echoes that uncovered earlier using the three-pass estimator, suggesting essentially no relationship between the macro factors and FX portfolios. But, importantly, we can now rule out that the disconnect is imputable to the fact that macro factors are weak factors in the cross section of FX portfolio returns; in fact, we find no evidence of asset selection, and the R^2s are mostly negative, especially as the models include an increasing number of latent factors, k. On the contrary, many financial and text-based factors display positive R^2s, and a few factors also display asset selection. However, for almost all candidate factors, the cross-validation exercise points to SDFs including a small number of factors (with k being typically less than three), regardless of whether there is asset selection or not. Therefore, we need to turn to the joint estimation to carry out the risk-premium inference. In doing so, we find that not all factors with positive cross-validation out-of-sample R^2s have significant risk premia (right columns of Table A20, *Joint*). This leaves us with a relatively small set of relevant nontradable factors, whose risk-premium estimates are largely consistent with our three-pass baseline estimates.[30]

`Asset Selection`. Another clear benefit of using SPCA is the asset selection, which provides valuable information by zooming into the identities of the selected assets or, in case of no selection, of the most correlated assets with the candidate factor of interest. Indeed, by doing this, we can establish a link among investment strategies (Table A21 connects tradable factors to the portfolios), and between the macro-financial risks and the investment strategies (Tables A22-A23 link relevant nontradable factors to the portfolios). To start with, we note that the procedure works particularly well. For the tradable factors,

[30]Specifically, the financial factors with significant premia are: otic, icap, gfc, gvol, psliq, move, vxo, and eqrv. Then, among the text-based factors, gepu and gepu ppp, the emv tracker and some of its subcategories (i.e., emv mout, emv inf, emv com, emv ir, and emv mp) are significant. Only the factor ipw/us(q), which measures the difference between world and U.S. quarterly growth rates in industrial production, is significant among the macro factors.

the results in Table A21 show that SPCA mostly relies on the information contained in the strategy corner portfolios to price the associated HML factor. Moreover, there is a close association between Carry and the 'usual suspects' (essentially the same strategies exposed to $\hat{F}_2$ in the three-pass analysis), while Carry is weakly related to Mom (both ST and LT), Value, and GAP HML factors; hence these strategies largely reflect distinct sources of risk. The assets selected by SPCA make sense economically also for the nontradable factors. For example, otic (a measure of foreign central banks' accumulation of U.S. Treasury securities) mostly correlates with global imbalances strategies (e.g., NFA-P5), and ipw/us appears to be tightly linked to the high-risk GAP portfolio.

The SPCA procedure also highlights some relevant nontradable factors for ST Mom, Value, and GAP strategies. This is a useful finding, as it is notoriously hard to connect these strategies with observed measures of macro-financial risks.[31] However, a larger number of nontradable factors relate to Carry, and hence to carry-related strategies. These absolute correlations provide a first means to detect the most relevant factors.

Summing up, we find that the nontradable factor risk-premium estimates are largely robust to the estimation method used. In fact, the estimates obtained with SPCA are consistent with our baseline three-pass estimates. Above all, this additional analysis confirms the disconnect between macro factors and currency portfolios. We can now also exclude that the disconnect is imputable to a weak-factor problem. Moreover, the cross-validation exercise allows us to further filter out factors that cannot be hedged out of sample by the currency assets, shrinking further the list of relevant candidate factors.[32] Finally, the asset selection (which is at the core of the SPCA estimator) proves to be a valuable tool to draw a better connection between nontradable factors and investment strategies, as well as among investment strategies.

[31]For example, focusing on the high-risk portfolios, vxo and icap seem to capture relevant sources of risk for Value, while gepu and eqrv emerge as important factors for ST Mom (see Figure A25 in the Internet Appendix).

[32]We have seen how, starting from 133 factors, the number of statistically significant risk premia reduces from over 90 using the FMB procedure to less than 50 using the three pass-method. This is due to allowing for omitted-variable and measurement-error biases. Using SPCA, where we further require a positive cross-validation out-of-sample R^2, the number of significant risk premia further drops to 17, which is only 13% of the initial list of 133 factors. It is conceivable that allowing also for multiple-hypothesis testing bias would further reduce this number. The role of this kind of bias in the context of the three-pass procedure and SPCA requires further study, along the lines of Harvey et al. (2016) and Giglio et al. (2021b), which we leave to future research.

7 Concluding Remarks

In this paper, we revisit the macro-financial sources of the risk-return trade-off inherent in currency investment strategies through the lenses of the three-pass method of Giglio and Xiu (2021), which we combine with the Risk-Premium PCA method of Lettau and Pelger (2020a,b). This approach allows us to shed light on the optimal currency SDF and to provide estimates of the risk premia of a large number of nontradable factors, while allowing for both omitted-variable and measurement-error biases.

We find that, using RP-PCA to extract the latent factors, the optimal currency pricing kernel includes at least three latent factors: a strong U.S. "Dollar" level factor, and two weak, high Sharpe ratio "Carry" and "Momentum" slope factors. We show that this pricing kernel delivers a reasonably high maximal Sharpe ratio and low pricing errors, while the explained systematic variation of the portfolios is comparable to PCA. At the same time, using standard PCA, the "Momentum" factor would be omitted from the pricing kernel, due to its low time-series "signal strength". Using RP-PCA, there is only feeble evidence in favor of a fourth weak pricing factor, which relates to "Value". However, a fourth factor can hardly be recovered in an out-of-sample setting.

Based on this optimal pricing kernel, we then show that a large portion of our long list of nontradable factors is due to noise. This helps explain why the standard two-pass FMB method can deliver inflated estimates of nontradable factor prices of risk. In particular, we find that this problem is pervasive for macro factors, while it is more contained for some financial and text-based factors. Moreover, we document that "Carry" is by far the most relevant factor for financial and text-based factors. However, the omission of the "Momentum" factor can lead to distorted risk-premium estimates. In fact, many financial factors display significant exposures to both the "Carry" and "Momentum" factors, but of opposite signs.

Overall, we find that a small fraction of nontradable – mostly financial and text-based – factors are indeed priced in currency returns. Some of the nontradable factors previously uncovered by the literature turn out to be less or even not relevant, while other factors (i.e., which were not previously related to currency returns) appear to be relevant, disclosing a tight link between FX and other markets, mainly channeled through the "Carry" factor. In particular, the results highlight the relevance of several volatility (e.g., the Global FX volatility factor of Menkhoff et al., 2012a) and uncertainty (e.g., the Global EPU index of Baker et al., 2016) measures, and of some liquidity indicators. Conversely, the results confirm a substantial disconnect between currency returns and macro factors, especially as both sources of bias are accounted for. Such disconnect manifests even if we use estimators that account for the possibility that macro factors are weak factors in the FX cross section.

Taken together, the evidence uncovered contributes to our understanding of the risk-return trade-offs inherent in currency investment strategy returns. Moreover, our results highlight the empirical relevance of achieving robust risk-premium estimates, which takes central stage in the current asset pricing research agenda. In particular, we make some progress in taming the FX "factor zoo" which, albeit in its infancy compared to other markets, is rapidly expanding. Finally, our finding that the optimal currency SDF comprises at least three factors with different strengths and risk premia could guide future theoretical research in international macro-finance with the ultimate objective of deriving currency models that rationalize such properties from first principles.

A Appendix: Nontradable Candidate Factors

Next, we provide a brief overview of the candidate factors and motivations for considering them. We present the detailed list and description of the factors in the Internet Appendix (Section III, Tables A5-A7).

Financial Factors. This category includes a wide range of financial factors that either represent key risks specific to the FX market or are notoriously relevant across asset classes. To start with, we focus on FX risk factors that feature in the FX literature, such as the global FX volatility and liquidity factors of Menkhoff et al. (2012a), and the systematic FX liquidity measure of Karnaukh et al. (2015). We then revisit the role of global risk factors that were already related to currency returns; for example, we consider the TED spread, a measure of funding liquidity risk (e.g., Brunnermeier et al., 2009); the VIX index, i.e., the so called "fear gauge", an indicator of global risk aversion (e.g., Ranaldo and Söderlind, 2010)[33]; the S&P500 monthly realized volatility; and the financial intermediaries' capital ratio, capturing intermediaries' prominent role in determining prices, given that they act as marginal investors in many asset classes (e.g., Adrian et al., 2014; Gabaix and Maggiori, 2015; He et al., 2017).[34]

Along with these factors, we consider other factors that notoriously drive risk premia of many asset classes but that – to our knowledge – have received little attention in the FX context, despite they are likely to be highly relevant also for FX returns. These factors include the liquidity measure of Pastor and Stambaugh (2003), a proxy for liquidity risk in the equity home market of the U.S. investor; the noise measure developed by Hu et al. (2013), a broad measure of liquidity conditions in the U.S. Treasury bond market that relates to the availability of arbitrage capital; the Merrill Lynch Option Volatility Estimate (MOVE) Index, a "barometer" of the U.S. Treasury market conditions, often referred to as "the VIX for Bonds"; Treasury International Capital data on official flows into U.S. Treasuries, a quantity-based factor, capturing foreign central banks' inelastic demand for Treasuries, which can affect Treasuries yields (e.g., Krishnamurthy and Vissing-Jorgensen, 2012) and currency returns (Gourinchas et al., 2020; Greenwood et al., 2020); the spread between BAA and AAA rated bond yields, a measure of credit risk in the U.S. corporate bond market, which helps price both the cross section of bond and equity returns (Fama and French, 1993); the Libor-OIS and the OIS-TBill spreads, the two components of the TED spread that

[33]As is usual in the literature (e.g., Koijen et al., 2018), we use the implied volatility of S&P100, i.e., VXO, as its sample starts earlier than VIX, and the two measures are strongly related. (The results are similar using VIX.)

[34]In the empirical analysis, we use the measure of He et al. (2017) as it is available at a monthly frequency.

are informative about tensions in the U.S. interbank market and flight to liquidity, respectively (Caballero et al., 2008); the realized volatility of crude oil prices; the global financial cycle factor of Miranda-Agrippino and Rey (2020), a broad-wide measure of global financial conditions; and the three latent macro/financial factors constructed by Jurado et al. (2015). Evidently, albeit the list is by no means exhaustive, it is long and diverse. Therefore, taken together, these factors should capture multiple sources of financial risks that pertain to a wide range of asset classes.

`Macro Factors.` This category consists of macroeconomic factors that serve as proxies for U.S. and world business cycle risks. Given that no macroeconomic factor alone is sufficient to fully capture the economic environment of a country, we take an agnostic approach and explore many widely used macroeconomic factors.[35] In fact, while there is a general consensus on how to measure inflation, this is less the case for real activity and consumption risk.[36] At the same time, the empirical analysis is conducted at a monthly frequency, which restricts somewhat the set of macro variables available.

Specifically, for the U.S., we consider the industrial production index, the Chicago Fed National Activity Index (CFNAI), the consumer price index, personal consumption expenditures, non-farm payrolls, and the unemployment rate. These macro variables are regarded as potential financial market movers and, for this reason, featured over the years in the asset pricing literature. Thus, they should capture adequately the macro risks in the home country of the representative investor. For most of these variables, however, there is no agreement on the frequency of their measurement. We do not take a stance and measure the variables as monthly, quarterly and yearly growth rates. For these frequencies, we also employ exponential moving averages. In this way, we reduce the likelihood that a macro factor is not selected as pricing factor due to its measurement.[37]

While U.S. data are also informative about global macroeconomic disturbances (Zviadadze, 2017), given

[35] Anecdotal evidence suggests that traders tend to look at simple indicators, rather than at economic indices. Moreover, our method is implemented on each indicator in isolation. Hence, we do not aggregate or extract common factors from these macro factors, rather use them one by one.

[36] In theory, consumption growth risk is arguably the most relevant macro fundamental for asset prices. However, measuring consumption growth risk is not trivial, as its definition varies with the types of agents and goods considered. For example, Malloy et al. (2009), Gonzalez-Urteaga and Rubio (2016), and Giglio and Xiu (2021) opt to use an aggregate measure of U.S. consumption growth risk, which has also the pros of being available at a monthly frequency.

[37] Regardless of the choice of the specific variable, measures of economic activity often require the estimation of unobserved equilibrium objects (like the natural rate of employment or potential output), which are notoriously hard to quantify, and are inherently model and hence also sample dependent. For this reason, we opt to use growth rates that are not contaminated by judgment and/or estimation error, being readily available to the investors.

the prominent role of the U.S. economy in the world financial system and trade (Maggiori, 2017), we also consider a range of foreign countries' macroeconomic data. Due to data limitations, we focus on a subset of the U.S. macroeconomic variables (i.e., the industrial production index, the consumer price index, and the unemployment rate). We measure foreign countries' macro indicators as before. However, we face the additional task of aggregating individual country' macro indicators into a global indicator. To do this, we follow common practice and compute GDP-weighted averages of country indicators. In the empirical analysis, we use these global indicators in levels, but also in differences relative the U.S., consistent with the fact that simple portfolios are not dollar-neutral investment strategies. We also construct measures of macro risk that capture the cross-sectional dispersion in the state of the business cycle across countries. We do so by taking each month the cross-sectional standard deviation of the individual country indicators.[38]

`Text-Based Factors.` This category comprises factors that are obtained by aggregating into an index news coverage about specific sources of uncertainty. Therefore, they are clearly nontradable factors. The U.S. Economic Policy Uncertainty (EPU) Index developed by Baker et al. (2016), which quantifies newspaper coverage of U.S. policy-related economic uncertainty, is arguably one of the most well-known text-based indicator. (Another widely used indicator of economic uncertainty is the index of Jurado et al., 2015.) Policy uncertainty is perceived by the investors as a highly undiversifiable source of risk, and should therefore command a negative risk premium (e.g., Pastor and Veronesi, 2013). That is, uncertainty-averse investors should demand extra compensation to hold assets with negative uncertainty betas, whereas they should be willing to pay high prices for assets with positive uncertainty betas (Bali et al., 2017). Relatedly, the U.S. EPU index turns out to help explain the 25 size and momentum Fama-French portfolios (Brogaard and Detzel, 2015). Hence, by now, there is compelling evidence that U.S. EPU is a risk factor for equities.

More recently, sub-indices of the U.S. EPU index based on news data and measuring different sources of policy uncertainty (e.g., fiscal, monetary, etc.) became available. More fundamentally, EPU indices have been developed for a wider set of countries and aggregated into a global version. This in turn paves the way

[38]It is important to note that we do not consider macro data in real time, for two main reasons. First, real-time data are not available for most of the foreign countries over sufficiently long time periods. Second, our objective is to explain macro factors with the currency (latent) factors, and not the other way around. Thus, the use of real time data is arguably less relevant. Finally, we intentionally omit measures of stochastic variance of macroeconomic variables – the silent feature of macroeconomic data (e.g., Zviadadze, 2017) – as also these objects are model dependent. However, some of the text-based measures, presented next, capture macroeconomic uncertainty/risk of different sources, and hence more closely relate to the second moments of macroeconomic variables.

to assess the relation between EPU indices and currency returns, which so far have been mostly related to equity returns. A notable exception is Berg and Mark (2018b) that document a link between global EPU and Carry excess returns. Mueller et al. (2017) examine the link between exchange rates and monetary policy uncertainty, both theoretically and empirically; in one of their empirical specifications, they use the EPU index.

Moreover, the same approach – based on the search and counting of selected terms from newspapers – has been used to construct a newspaper-based U.S. equity market volatility (EMV). This tracker and its sub-categories provide novel insights about the determinants of equity market volatility (Baker et al., 2019). Thus, we conjecture that EMV indices are also likely to be relevant factors for many assets, including currency returns. Therefore, we consider a battery of text-based factors, which essentially consist of EPU and EMV indices, and their sub-categories.

References

Adrian, T., E. Etula, and T. Muir (2014): "Financial Intermediaries and the Cross-Section of Asset Returns," *Journal of Finance*, 69, 2557–2596.

Ang, A. and J. Chen (2010): "Yield Curve Predictors of Foreign Exchange Returns," *Working Paper.*

Asness, C. S., T. J. Moskowitz, and L. H. Pedersen (2013): "Value and Momentum Everywhere," *Journal of Finance*, 68, 929–985.

Bai, J. and S. Ng (2002): "Determining the Number of Factors in Approximate Factor Models," *Econometrica*, 70, 191–221.

——— (2006): "Evaluating Latent and Observed Factors in Macroeconomics and Finance," *Journal of Econometrics*, 131, 507–537.

Baker, S. R., N. Bloom, and S. J. Davis (2016): "Measuring Economic Policy Uncertainty," *Quarterly Journal of Economics*, 131, 1593–1636.

Baker, S. R., N. Bloom, S. J. Davis, and K. Kost (2019): "Policy News and Stock Market Volatility," *Working Paper.*

Bali, T. G., S. J. Brown, and Y. Tang (2017): "Is Economic Uncertainty Priced in the Cross-Section of Stock Returns?" *Journal of Financial Economics*, 126, 471–489.

Bekaert, G. and R. Hodrick (1993): "On Biases in the Measurement of Foreign Exchange Risk Premiums," *Journal of International Money and Finance*, 12, 115–138.

Bekaert, G., M. Hoerova, and M. Lo Duca (2013): "Risk, Uncertainty and Monetary Policy," *Journal of Monetary Economics*, 60, 771–788.

Bekaert, G., M. Wei, and Y. Xing (2007): "Uncovered Interest Rate Parity and the Term Structure," *Journal of International Money and Finance*, 26, 1038–1069.

Benetrix, A., P. Lane, and J. C. Shambaugh (2015): "International Currency Exposures, Valuation Effects and the Global Financial Crisis," *Journal of International Economics*, 96, S98–S109.

Berg, K. and N. Mark (2018a): "Global Macro Risks in Currency Excess Returns," *Journal of Empirical Finance*, 45, 300–315.

——— (2018b): "Measures of Global Uncertainty and Carry-Trade Excess Returns," *Journal of International Money and Finance*, 88, 212–227.

Brogaard, J. and A. Detzel (2015): "The Asset-Pricing Implications of Government Economic Policy Uncertainty," *Management Science*, 61, 3–18.

Brunnermeier, M. K., S. Nagel, and L. H. Pedersen (2009): "Carry Trades and Currency Crashes," in *NBER Macroeconomics Annual 2008*, ed. by K. Rogoff, M. Woodford, and D. Acemoglu, Cambridge, MA: MIT Press, 23, 313–347.

Burnside, C. (2011): "The Cross Section of Foreign Currency Risk Premia and Consumption Growth Risk: a Comment," *American Economic Review*, 101, 3456–3476.

Caballero, R. J., E. Farhi, and P.-O. Gourinchas (2008): "Financial Crash, Commodity Prices, and Global Imbalances," *Brookings Papers on Economic Activity*, 2008, 1–55.

Chernov, M., M. Dahlquist, and L. A. Lochstoer (2021): "Pricing Currency Risks," *Working Paper.*

Cochrane, J. (2005): *Asset Pricing*, NJ: Princeton University Press.

Cochrane, J. H. (2017): "Macro-Finance," *Review of Finance*, 21, 945–985.

Colacito, R., S. Riddiough, and L. Sarno (2020): "Business Cycles and Currency Returns," *Journal of Financial Economics*, 137, 659–678.

Dahlquist, M. and H. Hasseltoft (2020): "Economic Momentum and Currency Returns," *Journal of Financial Economics*, 136, 152–167.

Della Corte, P., S. J. Riddiough, and L. Sarno (2016): "Currency Premia and Global Imbalances," *Review of Financial Studies*, 29, 2161–2193.

Dew-Becker, I., S. Giglio, A. Le, and M. Rodriguez (2017): "The Price of Variance Risk," *Journal of Financial Economics*, 123, 225–250.

Fama, E. F. and K. R. French (1993): "Common Risk Factors in the Returns of Stocks and Bonds," *Journal of Financial Economics*, 33, 3–56.

——— (2015): "A Five-Factor Asset Pricing Model," *Journal of Financial Economics*, 116, 1–22.

Fama, E. F. and J. D. MacBeth (1973): "Risk, Return, and Equilibrium: Empirical Tests," *Journal of Political Economy*, 81, 607–636.

Feng, G., S. Giglio, and D. Xiu (2020): "Taming the Factor Zoo: A Test of New Factors," *Journal of Finance*, 75, 1327–1370.

Fratzscher, M., D. Rime, L. Sarno, and G. Zinna (2015): "The Scapegoat Theory of Exchange Rates: The First Tests," *Journal of Monetary Economics*, 70, 1–21.

Gabaix, X. and M. Maggiori (2015): "International Liquidity and Exchange Rate Dynamics," *Quarterly Journal of Economics*, 130, 1369–1420.

Giacoletti, M., K. T. Laursen, and K. J. Singleton (2021): "Learning from Disagreement in the U.S. Treasury Bond Market," *Journal of Finance*, 76, 395–441.

Giglio, S., B. Kelly, and D. Xiu (2021a): "Factor Models, Machine Learning, and Asset Pricing," *Annual Review of Financial Economics, forthcoming.*

Giglio, S., Y. Liao, and D. Xiu (2021b): "Thousands of Alpha Tests," *Review of Financial Studies*, 34, 3456–3496.

Giglio, S. and D. Xiu (2021): "Asset Pricing with Omitted Factors," *Journal of Political Economy*, 129, 1947–1990.

Giglio, S., D. Xiu, and D. Zhang (2021c): "Test Assets and Weak Factors," Working Paper 29002, National Bureau of Economic Research.

Gonzalez-Urteaga, A. and G. Rubio (2016): "The Cross-Sectional Variation of Volatility Risk Premia," *Journal of Financial Economics*, 119, 353–370.

Gourinchas, P. O., W. Ray, and D. Vayanos (2020): "A Preferred-Habitat Model of Term Premia and Currency Risk," mimeo, London School of Economics.

Greenwood, R., S. Hanson, J. Stein, and A. Sunderam (2020): "A Quantity-Driven Theory of Term Premia and Exchange Rates," NBER Working Papers 27615, National Bureau of Economic Research, Inc.

Harvey, C. R. and Y. Liu (2021): "Lucky Factors," *Journal of Financial Economics*, 141, 413–435.

Harvey, C. R., Y. Liu, and H. Zhu (2016): "...and the Cross-Section of Expected Returns," *Review of Financial Studies*, 29, 5–68.

Hassan, T. A. (2013): "Country Size, Currency Unions, and International Asset Returns," *Journal of Finance*, 68, 2269–2308.

He, Z., B. Kelly, and A. Manela (2017): "Intermediary Asset Pricing: New Evidence From Many Asset Classes," *Journal of Financial Economics*, 126, 1–35.

Hu, G. X., J. Pan, and J. Wang (2013): "Noise as Information for Illiquidity," *Journal of Finance*, 68, 2341–2382.

Jurado, K., S. C. Ludvigson, and S. Ng (2015): "Measuring Uncertainty," *American Economic Review*, 105, 1177–1216.

Karnaukh, N., A. Ranaldo, and P. Söderlind (2015): "Understanding FX Liquidity," *Review of Financial Studies*, 28, 3073–3108.

Koijen, R. S., T. J. Moskowitz, L. H. Pedersen, and E. B. Vrugt (2018): "Carry," *Journal of Financial Economics*, 127, 197–225.

Kozak, S., S. Nagel, and S. Santosh (2020): "Shrinking the Cross-Section," *Journal of Financial Economics*, 135, 271–292.

Krishnamurthy, A. and A. Vissing-Jorgensen (2012): "The Aggregate Demand for Treasury Debt," *Journal of Political Economy*, 120, 233–267.

Kroencke, T., F. Schindler, and A. Schrimpf (2014): "International Diversification Benefits with Foreign Exchange Investment Styles," *Review of Finance*, 18, 1847–1883.

Lane, P. and G. Milesi-Ferretti (2004): "The External Wealth of Nations: Measures of Foreign Assets and Liabilities for Industrial and Developing Countries," *Journal of International Economics*, 55, 263–294.

——— (2007): "The External Wealth of Nations Mark II: Revised and Extended Estimates of Foreign Assets and Liabilities, 1970-2004," *Journal of International Economics*, 73, 223–250.

Lettau, M., M. Maggiori, and M. Weber (2014): "Conditional Risk Premia in Currency Markets and Other Asset Classes," *Journal of Financial Economics*, 114, 197–225.

Lettau, M. and M. Pelger (2020a): "Estimating Latent Asset-Pricing Factors," *Journal of Econometrics*, 218, 1–30.

——— (2020b): "Factors that Fit the Time Series and Cross-section of Stock Returns," *Review of Financial Studies*, 33, 2274–2325.

Lewellen, J., S. Nagel, and J. Shanken (2010): "A Skeptical Appraisal of Asset Pricing Tests," *Journal of Financial Economics*, 96, 175–194.

Lintner, J. (1965): "Security Prices, Risk, and Maximal Gains from Diversification," *Journal of Finance*, 20, 587–615.

Lustig, H., N. Roussanov, and A. Verdelhan (2011): "Common Risk Factors in Currency Markets," *Review of Financial Studies*, 24, 3731–3777.

Lustig, H., A. Stathopoulos, and A. Verdelhan (2019): "The Term Structure of Currency Carry Trade Risk Premia," *American Economic Review*, 109, 4142–77.

Lustig, H. and A. Verdelhan (2007): "The Cross Section of Foreign Currency Risk Premia and Consumption Growth Risk," *American Economic Review*, 97, 89–117.

Maggiori, M. (2017): "Financial Intermediation, International Risk Sharing, and Reserve Currencies," *American Economic Review*, 107, 3038–71.

Malloy, C. J., T. J. Moskowitz, and A. Vissing-Jorgensen (2009): “Long-Run Stockholder Consumption Risk and Asset Returns,” *Journal of Finance*, 64, 2427–2479.

Mark, N. C. (1995): “Exchange Rates and Fundamentals: Evidence on Long-Horizon Predictability,” *American Economic Review*, 85, 201–218.

Meese, R. and K. Rogoff (1983): “Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?” *Journal of International Economics*, 14, 3–24.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2012a): “Carry Trades and Global Foreign Exchange Volatility,” *Journal of Finance*, 67, 681–718.

——— (2012b): “Currency Momentum Strategies,” *Journal of Financial Economics*, 106, 681–718.

——— (2017): “Currency Value,” *Review of Financial Studies*, 30, 416–441.

Merton, R. C. (1973): “An Intertemporal Capital Asset Pricing Model,” *Econometrica*, 41, 867–887.

Miranda-Agrippino, S. and H. Rey (2020): “U.S. Monetary Policy and the Global Financial Cycle,” *Review of Economic Studies*, 87, 2754–2776.

Mueller, P., A. Tahbaz-Salehi, and A. Vedolin (2017): “Exchange Rates and Monetary Policy Uncertainty,” *Journal of Finance*, 72, 1213–1252.

Onatski, A. (2010): “Determining the Number of Factors from Empirical Distribution of Eigenvalues,” *Review of Economics and Statistics*, 92, 1004–1016.

Pastor, L. and R. Stambaugh (2003): “Liquidity Risk and Expected Stock Returns,” *Journal of Political Economy*, 111, 642–685.

Pastor, L. and P. Veronesi (2013): “Political Uncertainty and Risk Premia,” *Journal of Financial Economics*, 110, 520–545.

Ranaldo, A. and P. Söderlind (2010): “Safe Haven Currencies,” *Review of Finance*, 14, 385–407.

Ready, R., N. Roussanov, and C. Ward (2017): “Commodity Trade and the Carry Trade: A Tale of Two Countries,” *Journal of Finance*, 72, 2629–2684.

Rossi, B. (2013): "Exchange Rate Predictability," *Journal of Economic Literature*, 51, 1063–1119.

Shanken, J. (1992): "On the Estimation of Beta Pricing Models," *Review of Financial Studies*, 5, 1–34.

Sharpe, W. (1964): "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk," *Journal of Finance*, 19, 425–442.

Verdelhan, A. (2018): "The Share of Systematic Risk in Bilateral Exchange Rates," *Journal of Finance*, 73, 375–418.

Zviadadze, I. (2017): "Term Structure of Consumption Risk Premia in the Cross Section of Currency Returns," *Journal of Finance*, 72, 1529–1566.

Table 1: Latent Factor Pricing Diagnostics

The table presents model diagnostics of the first two steps of the asset pricing procedure of Giglio and Xiu (2021) applied to currency portfolios excess returns, where the pricing factors are latent and are estimated using the RP-PCA method of Lettau and Pelger (2020a,b). We report diagnostics for RP-PCA implemented without "overweight" on the means ($\omega = -1$), i.e., standard PCA, and with increasing values of the RP-weight (ω = 10, 20 and 50). We consider SDFs, $\varphi(F_{1-k})$, including an increasing number of latent factors, $k = 1, 2, \ldots, 6$. Tab A.I First pass, Panel A: Two-pass Statistics, shows the average idiosyncratic variance, $\overline{\sigma}^2_\epsilon = \frac{1}{N}\sum_{n=1}^{N}\left[Var(\hat{\epsilon}_n)/Var(X_n)\right]$, and the average root-mean-square pricing errors, $\overline{RMS}_\alpha = \sqrt{\hat{\alpha}\hat{\alpha}^\top/N}$, obtained by estimating $X_{nt} = \alpha_n + \hat{F}_t\psi_n^\top + \epsilon_{nt}$, for $n = 1..., N$ test assets, and $t = 1..., T$. Tab A.II Second pass presents the R-squared values ($R^2(\%)$), and the mean absolute errors (MAE) of the cross-sectional regression, $\overline{X}_n = \hat{\psi}_n\gamma^\top + a_n$, for $n = 1, \ldots, N$, where γ is the $1 \times K$ vector of latent factors' prices of risk. Tab B.I Components, Panel B: Sharpe Ratios, presents the maximal Sharpe ratio (SR) from the tangency portfolio of the mean-variance frontier spanned by the linear combination of the K selected latent factors, $\hat{F} \times \hat{b}_{MV}^\top$, where $\hat{b}_{MV}$ is a 1 × K vector with entries $\hat{b}_{MV,k} = \mu_{F,k}/\sigma^2_{F,k}$, with $\mu_{F,k}$ and $\sigma^2_{F,k}$ denoting the k-th factor's mean and variance. ΔSR denotes the difference in SRs between SDFs with k and $k-1$ factors. The $\hat{b}_{MV,k}$ entry represents the k-th factor's weight in the SDF, $\varphi_t = 1 - (\hat{F}_t - \mu_F)\hat{b}_{MV}^\top$. The test assets consist of the portfolios from the nine investment strategies ($N = 46$) – see Section II in the Internet Appendix – for the period 11/1983-12/2017 at monthly frequency ($T = 410$).

	Panel A: Two-pass Statistics					Panel B: Sharpe Ratios			
	A.I First pass		A.II Second pass			B.I Components			
$\omega = -1$	$\overline{\sigma}^2_\epsilon$	$\overline{RMS}_\alpha$	$R^2(\%)$	MAE		SR	ΔSR	$\hat{b}_{MV,k}$	$\mu_{F,k}$
$\varphi(F_1)$	23.41	1.73	0.17	1.38	F_1	0.10	–	0.05	17.38
$\varphi(F_{1-2})$	19.04	1.59	16.48	1.22	F_2	0.14	0.05	0.25	4.63
$\varphi(F_{1-3})$	17.07	1.30	44.94	0.94	F_3	0.26	0.12	0.75	6.25
$\varphi(F_{1-4})$	15.46	1.30	44.96	0.94	F_4	0.26	0.00	0.03	0.21
$\varphi(F_{1-5})$	14.12	0.91	72.06	0.73	F_5	0.37	0.11	1.10	6.26
$\varphi(F_{1-6})$	12.92	0.90	72.37	0.71	F_6	0.37	0.00	0.14	0.71
$\omega = 10$	$\overline{\sigma}^2_\epsilon$	$\overline{RMS}_\alpha$	$R^2(\%)$	MAE		SR	ΔSR	$\hat{b}_{MV,k}$	$\mu_{F,k}$
$\varphi(F_1)$	23.42	1.73	29.02	1.24	F_1	0.10	–	0.06	18.12
$\varphi(F_{1-2})$	19.44	1.53	77.94	0.59	F_2	0.28	0.18	0.70	10.02
$\varphi(F_{1-3})$	17.31	0.89	97.22	0.22	F_3	0.44	0.15	1.13	9.70
$\varphi(F_{1-4})$	15.79	0.71	98.49	0.17	F_4	0.46	0.03	0.63	4.05
$\varphi(F_{1-5})$	14.19	0.70	98.54	0.16	F_5	0.47	0.00	0.13	0.89
$\varphi(F_{1-6})$	12.99	0.70	98.55	0.16	F_6	0.47	0.00	0.05	0.28
$\omega = 20$	$\overline{\sigma}^2_\epsilon$	$\overline{RMS}_\alpha$	$R^2(\%)$	MAE		SR	ΔSR	$\hat{b}_{MV,k}$	$\mu_{F,k}$
$\varphi(F_1)$	23.43	1.74	59.88	1.14	F_1	0.10	–	0.06	18.65
$\varphi(F_{1-2})$	19.94	1.47	94.58	0.29	F_2	0.36	0.25	1.02	11.29
$\varphi(F_{1-3})$	17.36	0.84	99.16	0.12	F_3	0.45	0.10	0.87	9.17
$\varphi(F_{1-4})$	15.81	0.68	99.53	0.09	F_4	0.48	0.02	0.57	3.71
$\varphi(F_{1-5})$	14.21	0.68	99.53	0.09	F_5	0.48	0.00	0.07	0.51
$\varphi(F_{1-6})$	13.01	0.68	99.53	0.09	F_6	0.48	0.00	0.05	0.27
$\omega = 50$	$\overline{\sigma}^2_\epsilon$	$\overline{RMS}_\alpha$	$R^2(\%)$	MAE		SR	ΔSR	$\hat{b}_{MV,k}$	$\mu_{F,k}$
$\varphi(F_1)$	23.51	1.76	90.70	0.87	F_1	0.11	–	0.06	19.67
$\varphi(F_{1-2})$	20.36	1.37	99.26	0.11	F_2	0.40	0.29	1.25	11.85
$\varphi(F_{1-3})$	17.40	0.80	99.85	0.05	F_3	0.47	0.06	0.67	8.29
$\varphi(F_{1-4})$	15.83	0.67	99.91	0.04	F_4	0.48	0.02	0.52	3.46
$\varphi(F_{1-5})$	14.23	0.66	99.91	0.04	F_5	0.48	0.00	0.06	0.38
$\varphi(F_{1-6})$	13.02	0.66	99.91	0.04	F_6	0.48	0.00	0.05	0.26

Table 2: Exposures of Nontradable Factors to Latent Factors

The table presents the nontradable candidate factors' exposures to the latent factors (η_{F_k}) and the explained variations ($R^2_{F_{1-k}}$) obtained from the spanning regression of Eq. (15), for SDFs including an increasing number of factors, $k = 1, \ldots, K$. We report the candidate factor exposures to the first six extracted, orthogonalized latent factors (i.e., $K = 6$). The factors are extracted by applying RP-PCA with baseline weight (i.e., $\omega = 20$) to the $N = 46$ portfolios obtained from the nine investment strategies. Panels A, B, and C show the estimates for the financial, text-based, and macro factors, respectively. We present the estimates only for the candidate factors with statistically significant risk premia ($\hat{\lambda}_g$) according to the three-pass model with optimal number of factors, i.e. $\varphi(F_{1-3})$, reported later in Table 3. When a macro factor's risk premium is significant at multiple frequencies, we present only the most representative one (i.e., the frequency at which the factor's risk premium is most precisely estimated). We report all candidate factors' exposures in Table A9, in the Internet Appendix. ***,**,* denote significance, respectively, at the 1-, 5- and 10-percent levels, based on Newey-West standard errors. See nontradable factor descriptions in Tables A5-A7, in the Internet Appendix.

PANEL A: Financial Factors

	A.I: Risk Exposures						A.II: Explained Variation					
	η_{F_1}	η_{F_2}	η_{F_3}	η_{F_4}	η_{F_5}	η_{F_6}	$R^2_{F_1}$	$R^2_{F_{1-2}}$	$R^2_{F_{1-3}}$	$R^2_{F_{1-4}}$	$R^2_{F_{1-5}}$	$R^2_{F_{1-6}}$
otic	0.73**	5.73***	-0.35	-1.42	2.26	7.47**	1.17	3.67	3.68	3.77	4.01	5.99
noise	-1.36*	-11.34***	11.34**	8.06*	-0.62	-0.41	4.99	10.98	17.33	20.05	20.07	20.07
sliq	-1.29**	-10.48***	5.28	7.79*	-0.90	-6.67*	5.81	11.47	13.40	16.21	16.23	17.60
gfc	3.13***	11.71***	-7.49***	-2.96	8.80***	3.13	21.83	32.33	36.43	36.83	40.49	40.84
gvol	-1.07**	-9.95***	2.99	7.80**	-2.07	-8.34***	2.55	10.13	10.79	13.54	13.74	16.21
psliq	0.03	8.00**	-0.85	-0.88	2.42	-0.84	0.00	4.91	4.96	5.00	5.27	5.30
ted	-0.45	-6.82*	-0.44	0.38	0.72	0.27	0.47	3.94	3.95	3.96	3.98	3.98
lib ois	-1.48*	-14.14*	3.61	1.30	6.90	4.52	6.39	16.00	16.97	17.09	18.03	18.35
move	-0.98**	-10.89***	7.81***	0.12	-0.25	-1.94	2.78	8.33	11.75	11.75	11.75	11.86
vxo	-1.60***	-15.19***	10.32***	1.85	-6.91***	-1.35	6.47	19.96	26.4	26.55	28.5	28.56
eqrv	-0.65	-11.72**	1.46	3.11	-2.83	0.80	0.93	11.45	11.6	12.04	12.42	12.44

PANEL B: Text-Based Factors

	B.I: Risk Exposures						B.II: Explained Variation					
	η_{F_1}	η_{F_2}	η_{F_3}	η_{F_4}	η_{F_5}	η_{F_6}	$R^2_{F_1}$	$R^2_{F_{1-2}}$	$R^2_{F_{1-3}}$	$R^2_{F_{1-4}}$	$R^2_{F_{1-5}}$	$R^2_{F_{1-6}}$
gepu	0.00	-11.18***	3.38	3.12	-3.68	-3.52	1.12	7.36	7.64	7.99	8.26	8.56
gepu ppp	0.13	-12.01***	3.82	2.81	-3.79	-2.97	0.89	7.82	8.17	8.45	8.75	8.97
emv ov	-0.49	-9.13**	2.88	1.87	-1.51	0.31	0.48	6.73	7.35	7.51	7.61	7.61
emv mout	-0.38	-9.09**	2.19	2.74	-0.31	1.32	0.29	6.40	6.73	7.07	7.07	7.14
emv mqnt	-0.26	-7.85**	1.35	2.55	-1.12	1.09	0.14	4.78	4.92	5.21	5.27	5.31
emv inf	-0.28	-7.93**	2.92	0.09	-0.98	0.29	0.14	4.85	5.45	5.45	5.50	5.50
emv com	-0.55	-8.96**	3.64	1.21	-1.67	0.39	0.59	6.60	7.56	7.63	7.75	7.76
emv ir	-0.23	-7.97**	0.89	1.32	-1.56	1.60	0.11	4.96	5.03	5.11	5.22	5.31
emv fc	-0.72	-6.25*	2.52	2.78	2.96	-2.17	1.06	3.80	4.16	4.50	4.90	5.06
emv fp	-0.32	-7.32**	1.52	1.15	-1.58	0.34	0.21	4.29	4.48	4.54	4.65	4.65
emv tx	-0.34	-7.72**	2.09	2.10	-0.68	0.61	0.23	4.68	4.99	5.19	5.21	5.22
emv mp	-0.59	-8.69***	2.80	2.88	0.22	1.18	0.70	6.22	6.74	7.11	7.11	7.16
emv reg	-0.20	-8.56**	1.84	2.19	1.49	0.19	0.07	5.42	5.63	5.84	5.93	5.93
emv freg	-0.49	-7.54*	1.32	2.18	1.18	2.37	0.47	4.60	4.70	4.91	4.96	5.16

PANEL C: Macro Factors

	C.I: Risk Exposures						C.II: Explained Variation					
	η_{F_1}	η_{F_2}	η_{F_3}	η_{F_4}	η_{F_5}	η_{F_6}	$R^2_{F_1}$	$R^2_{F_{1-2}}$	$R^2_{F_{1-3}}$	$R^2_{F_{1-4}}$	$R^2_{F_{1-5}}$	$R^2_{F_{1-6}}$
cus(y)	-0.69	-2.24	-1.54	-0.92	-0.08	1.49	1.06	1.45	1.62	1.66	1.66	1.74
ipw(q)	0.23	-2.60	-1.80	0.53	0.95	1.87	0.12	0.63	0.87	0.88	0.92	1.05
ipw/us(q)	-0.15	-3.69**	-1.40	1.39	-0.88	1.99	0.05	1.09	1.24	1.32	1.36	1.50
cpiw(q)	0.10	-1.25	-4.07***	0.51	-0.17	-2.67	0.02	0.14	1.35	1.36	1.36	1.61
cpiw/us(ey)	-0.36	-1.77	-2.58	-1.97	2.08	-2.59	0.28	0.52	1.01	1.18	1.39	1.63
unew/us(y)	-0.27	-1.99	-3.66**	-2.53	-0.64	1.37	0.16	0.47	1.46	1.75	1.77	1.83

Table 3: Risk Premia of Nontradable Factors

The table presents the risk-premium estimates of selected nontradable factors (g_t). Panel FMB presents the risk-premium point estimates (λ_g) and Shanken standard errors (se) from the standard two-pass procedure, including the constant and the candidate factor. The remaining panels report the estimates from the augmented three-pass models of different dimensions. That is, $\varphi(F_{1-k})$ denotes the SDF including up to the k-th latent factor, whereby the factors are extracted from the panel of currency portfolio returns using the RP-PCA method with baseline weight, i.e., $\omega = 20$. The risk-premium estimates (λ_g) are reported along with the asymptotic standard errors (se) of Giglio and Xiu (2021); ***,**,* denote significance, respectively, at the 1-, 5- and 10-percent levels. For each factor and a given SDF, we also report the spanning R^2s (R2) resulting from projecting the factor onto the k latent factors entering the SDF; the Sharpe ratios (SR) associated with the projected factor, i.e., the return-based counterpart to the original nontradable factor; and the p-value (pval) of the test of GX that the g-factor is weak. In Panels A, B and C, we present financial (FIN), text-based (TXT), and macro (MAC) candidate g-factors with significant risk-premium estimates according to at least one of the SDF reported. When a macro factor is significant for multiple frequencies, we present the frequency at which the factor is most precisely estimated using $\varphi(F_{1-3})$. Factors are expressed as innovations, using the residuals from AR(1) processes, and are then standardized. The test assets consist of the portfolios from the nine investment strategies ($N = 46$). The sample period varies with the factor at hand, according to data availability over the 11/1983-12/2017 period. See nontradable factor descriptions in Tables A5-A7, in the Internet Appendix.

PANEL A:	FMB		$\varphi(F_{1-2})$					$\varphi(F_{1-3})$					$\varphi(F_{1-4})$				
FIN	λ_g	se	λ_g	se	R2	SR	pval	λ_g	se	R2	SR	pval	λ_g	se	R2	SR	pval
otic	8.36***	(2.87)	0.78***	(0.27)	3.67	0.34	0.00	0.75**	(0.32)	3.68	0.32	0.00	0.70**	(0.33)	3.77	0.30	0.01
icap	3.41**	(1.56)	1.20***	(0.34)	7.93	0.35	0.00	0.61	(0.39)	10.91	0.15	0.00	0.76*	(0.43)	11.60	0.19	0.00
noise	-3.95***	(1.33)	-1.41***	(0.43)	11.27	0.32	0.00	-0.83**	(0.35)	16.33	0.09	0.01	-0.52	(0.41)	19.51	0.03	0.01
sliq	-3.56**	(1.46)	-1.33***	(0.38)	10.33	0.31	0.00	-0.91**	(0.39)	13.13	0.15	0.01	-0.48	(0.51)	17.65	0.07	0.00
gfc	2.39**	(1.09)	1.91***	(0.45)	32.33	0.28	0.00	1.22**	(0.52)	36.43	0.17	0.00	1.11*	(0.56)	36.83	0.15	0.00
gliq	-4.28	(3.25)	-0.44**	(0.19)	1.17	0.34	0.03	-0.11	(0.23)	2.10	0.07	0.01	-0.11	(0.24)	2.10	0.06	0.01
gvol	-4.19**	(1.57)	-1.32***	(0.32)	10.13	0.35	0.00	-1.05***	(0.30)	10.79	0.27	0.00	-0.76**	(0.36)	13.54	0.17	0.00
psliq	7.64***	(2.58)	0.91**	(0.36)	4.91	0.34	0.04	0.83*	(0.43)	4.96	0.31	0.06	0.80*	(0.45)	5.00	0.30	0.10
corp	-2.03	(1.57)	-0.87**	(0.41)	7.86	0.24	0.13	-0.10	(0.32)	14.60	0.02	0.10	-0.05	(0.34)	14.73	0.03	0.11
ted	-12.27***	(3.49)	-0.89**	(0.43)	3.98	0.35	0.14	-0.97**	(0.43)	4.05	0.39	0.21	-0.97**	(0.42)	4.05	0.39	0.33
lib ois	-4.71**	(1.86)	-1.88*	(0.99)	14.32	0.33	0.09	-1.55**	(0.76)	15.23	0.24	0.12	-1.47*	(0.78)	15.31	0.23	0.16
move	-5.13***	(1.88)	-1.35***	(0.46)	7.75	0.34	0.02	-0.81**	(0.38)	11.85	0.12	0.00	-0.75**	(0.37)	11.96	0.11	0.00
vxo	-4.36***	(1.24)	-2.04***	(0.60)	20.22	0.34	0.00	-1.36**	(0.65)	26.98	0.14	0.00	-1.18*	(0.64)	27.98	0.11	0.00
eqrv	-5.74***	(1.82)	-1.44**	(0.63)	11.45	0.36	0.06	-1.31*	(0.71)	11.60	0.32	0.13	-1.19*	(0.70)	12.04	0.29	0.21
PANEL B:	FMB		$\varphi(F_{1-2})$					$\varphi(F_{1-3})$					$\varphi(F_{1-4})$				
TXT	λ_g	se	λ_g	se	R2	SR	pval	λ_g	se	R2	SR	pval	λ_g	se	R2	SR	pval
gepu	-8.12***	(2.67)	-1.52***	(0.43)	7.93	0.35	0.00	-1.41***	(0.48)	8.02	0.32	0.00	-1.18**	(0.52)	8.61	0.28	0.00
gepu ppp	-8.25***	(2.65)	-1.58***	(0.44)	8.35	0.36	0.00	-1.46***	(0.48)	8.45	0.32	0.00	-1.25**	(0.53)	8.96	0.28	0.00
epu all	-5.18**	(2.32)	-0.67*	(0.35)	2.89	0.33	0.09	-0.34	(0.40)	3.92	0.13	0.05	-0.18	(0.39)	4.93	0.05	0.01
epu mp	-7.62***	(2.65)	-0.70**	(0.33)	3.91	0.30	0.06	-0.59	(0.40)	4.02	0.25	0.07	-0.47	(0.39)	4.62	0.18	0.05
fsi tx	-4.06*	(2.27)	-0.71*	(0.38)	2.70	0.35	0.16	-0.27	(0.31)	4.44	0.10	0.07	-0.17	(0.30)	4.80	0.06	0.13
gpr a	5.07*	(2.59)	0.42**	(0.19)	1.41	0.29	0.08	0.32	(0.21)	1.48	0.22	0.14	0.17	(0.20)	2.19	0.10	0.13
emv ov	-6.79***	(2.30)	-1.15**	(0.46)	6.90	0.36	0.05	-0.91*	(0.45)	7.44	0.26	0.09	-0.84*	(0.44)	7.67	0.23	0.15
emv mout	-7.25***	(2.37)	-1.11**	(0.45)	6.55	0.35	0.05	-0.95**	(0.43)	6.80	0.29	0.11	-0.86**	(0.42)	7.17	0.25	0.18
emv mqnt	-7.77***	(2.59)	-0.95**	(0.40)	4.85	0.35	0.06	-0.84**	(0.38)	4.97	0.30	0.13	-0.75*	(0.38)	5.32	0.25	0.21
emv inf	-7.62***	(2.73)	-0.97***	(0.35)	4.99	0.35	0.02	-0.72*	(0.40)	5.52	0.24	0.02	-0.71*	(0.39)	5.52	0.24	0.04
emv com	-6.37***	(2.24)	-1.15**	(0.45)	6.82	0.36	0.04	-0.84*	(0.47)	7.68	0.23	0.05	-0.79*	(0.46)	7.79	0.21	0.07
emv ir	-9.09***	(2.94)	-0.97**	(0.44)	5.03	0.35	0.06	-0.90*	(0.51)	5.07	0.32	0.10	-0.83	(0.51)	5.23	0.29	0.16
emv fc	-7.57***	(2.73)	-0.86	(0.54)	3.86	0.34	0.17	-0.68*	(0.35)	4.13	0.25	0.28	-0.62*	(0.31)	4.30	0.22	0.42
emv fx	-5.15**	(2.19)	-0.71**	(0.29)	2.81	0.35	0.04	-0.42	(0.30)	3.54	0.17	0.02	-0.52	(0.31)	3.96	0.21	0.02
emv fp	-8.41***	(2.93)	-0.91**	(0.40)	4.38	0.35	0.08	-0.78*	(0.40)	4.55	0.29	0.17	-0.73*	(0.40)	4.66	0.26	0.27
emv tx	-7.42***	(2.63)	-0.96**	(0.43)	4.82	0.35	0.10	-0.79*	(0.43)	5.07	0.28	0.18	-0.72*	(0.43)	5.29	0.24	0.29
emv gov	-15.41***	(5.72)	-0.49*	(0.28)	1.50	0.33	0.23	-0.53	(0.32)	1.52	0.36	0.40	-0.56*	(0.33)	1.56	0.38	0.43
emv mp	-7.18***	(2.41)	-1.12***	(0.38)	6.41	0.35	0.01	-0.91**	(0.35)	6.80	0.27	0.04	-0.82**	(0.34)	7.15	0.23	0.06
emv reg	-8.27***	(2.74)	-1.01*	(0.51)	5.56	0.35	0.11	-0.89*	(0.47)	5.67	0.30	0.20	-0.83*	(0.46)	5.85	0.27	0.31
emv freg	-9.00***	(2.96)	-0.96*	(0.51)	4.70	0.35	0.20	-0.88*	(0.48)	4.75	0.32	0.35	-0.81*	(0.46)	4.95	0.29	0.45
emv tp	-7.47**	(3.34)	-0.70**	(0.31)	2.54	0.36	0.07	-0.40	(0.37)	3.38	0.16	0.02	-0.43	(0.38)	3.41	0.17	0.05
PANEL C:	FMB		$\varphi(F_{1-2})$					$\varphi(F_{1-3})$					$\varphi(F_{1-4})$				
MAC	λ_g	se	λ_g	se	R2	SR	pval	λ_g	se	R2	SR	pval	λ_g	se	R2	SR	pval
cus(y)	-13.49***	(4.14)	-0.38	(0.25)	1.45	0.26	0.12	-0.52*	(0.29)	1.62	0.34	0.15	-0.56*	(0.28)	1.66	0.36	0.25
ipw(q)	-6.27***	(1.96)	-0.25	(0.22)	0.63	0.26	0.32	-0.42*	(0.24)	0.87	0.37	0.25	-0.40	(0.26)	0.88	0.35	0.37
ipw/us(q)	-7.34***	(2.19)	-0.44**	(0.21)	1.09	0.35	0.10	-0.57**	(0.26)	1.24	0.43	0.14	-0.52*	(0.29)	1.32	0.38	0.19
cpiw(q)	-7.92***	(2.92)	-0.12	(0.21)	0.14	0.27	0.73	-0.50*	(0.28)	1.35	0.36	0.04	-0.48*	(0.28)	1.36	0.34	0.08
cpiw/us(ey)	-9.60***	(2.78)	-0.27	(0.23)	0.52	0.31	0.45	-0.50*	(0.28)	1.01	0.42	0.25	-0.58*	(0.30)	1.18	0.44	0.24
unew/us(y)	-16.91***	(4.95)	-0.27	(0.21)	0.47	0.34	0.37	-0.61**	(0.24)	1.46	0.42	0.06	-0.70***	(0.25)	1.75	0.45	0.06

Figure 1: Currency Investment Strategies

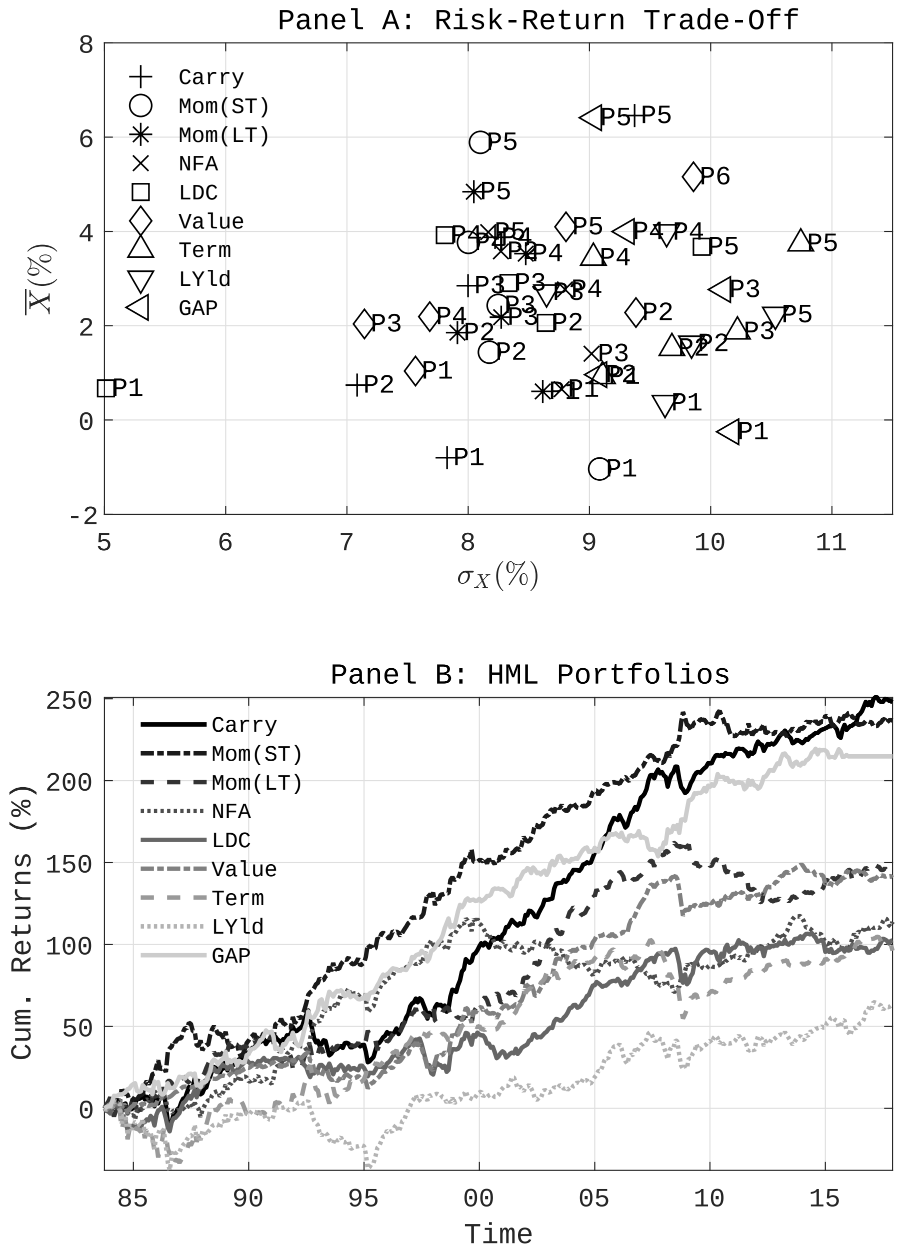

In Panel A, Risk-Return Trade-Off, the figure shows plain portfolio average excess returns ($\overline{X}(\%)$, in percent annualized) and standard deviations ($\sigma_X(\%)$, in percent annualized), proxying for risk. In Panel B, HML Portfolios, the figure plots the cumulative monthly returns of HML portfolios in percent. We consider nine investment strategies: carry (Carry), short-term momentum (Mom (ST)), long-term momentum (Mom (LT)), net foreign assets (NFA), liabilities in domestic currency (LDC), value (Value), term spreads (Term), long-term bond yields (LYld), and output gap (GAP). See Internet Appendix (Section II) for a detailed description of the strategies. The sample spans the 11/1983-12/2017 period at monthly frequency ($T = 410$).

Figure 2: Largest Normalized Eigenvalues

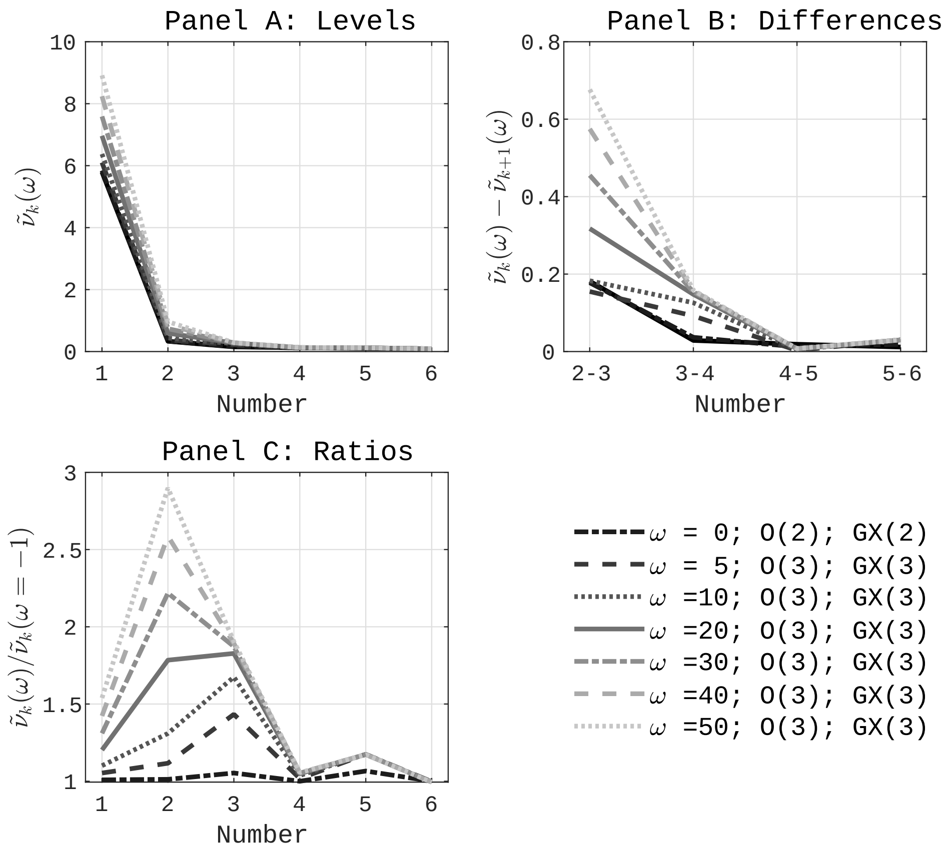

The figure shows the largest normalized eigenvalues of the matrix $\Sigma_{RP}^{(\omega)} = \frac{1}{T} X^\top X + \omega \overline{X}^\top \overline{X}$, for different values of the RP-weight (ω). The $T \times N$ matrix X collects currency portfolio excess returns from the nine investment strategies (i.e., $N = 46$), and $\overline{X}$ denote their sample averages. The eigenvalues are normalized by the average idiosyncratic variance ($\bar{\sigma}_\epsilon^2$) and hence relate more closely to factors' signal-to-noise ratios, being informative about factors' "signal strengths". Specifically, $\bar{\sigma}_\epsilon^2 = \frac{1}{N} \sum_{n=1}^{N} \sigma_{\epsilon,n}^2$, where $\sigma_{\epsilon,n}^2$ are the variances of the residuals obtained by estimating N time-series regressions, $X_{nt} = \alpha_n + \hat{F}_t^{(\omega)} \psi_n^\top + \epsilon_{nt}$, $n = 1..., N$ test assets, $t = 1..., T$ months, where $\hat{F}_t^{(\omega)}$ stacks the latent factors associated with the six largest eigenvalues of matrix $\Sigma_{RP}^{(\omega)}$, and thus vary with the RP-weight. Panel A, *Levels*, reports the normalized eigenvalues, $\tilde{\nu}_k(\omega) = \nu_k(\omega)/\sigma_\epsilon^2(\omega)$. Panel B, *Differences*, presents the difference of consecutive normalized eigenvalues, $\tilde{\nu}_k(\omega) - \tilde{\nu}_{k+1}(\omega)$, for $k = 2, \ldots, 5$. Panel C, *Ratios*, shows the eigenvalues scaled by the corresponding PCA ($\omega = -1$) eigenvalues, $\tilde{\nu}_k(\omega)/\tilde{\nu}_k(\omega = -1)$. In the legend, for a given RP-weight we present the optimal number of latent factors entering the SDF according to the Onatski (2010), O(#), and Giglio and Xiu (2021), GX(#), tests. The GX test is implemented with medium overfitting penalty value 0.5.

Figure 3: HML Portfolio Risk Exposures to Latent Factors

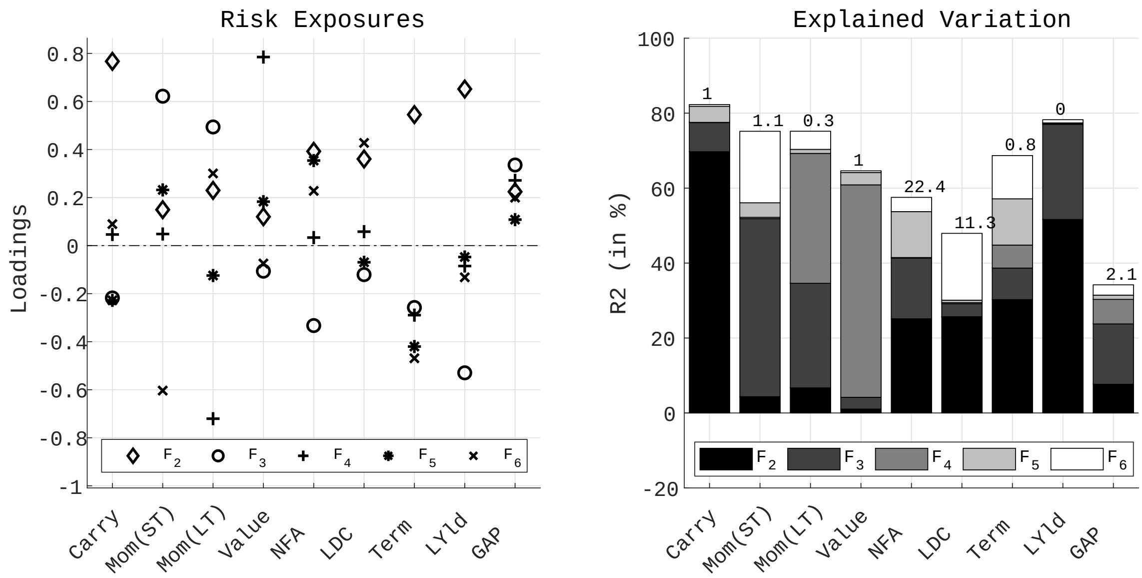

The figure shows HML portfolio loadings on the six estimated latent factors (Risk Exposures; left chart) and the associated R^2s obtained by regressing the HML portfolio returns on the latent factors (Explained Variation; right chart). The test assets' sample consists of the portfolios associated with the nine investment strategies ($N = 46$), thus HML portfolios are excluded from the estimation of latent factors (see Section II in the Internet Appendix for a description of the investment strategies). Hence, we infer HML portfolio loadings ex-post from the corner portfolio loadings of the associated strategy, as the difference between P5/6 (high) and P1 (low) loadings, obtained by means of RP-PCA with baseline penalty value, i.e., $\omega = 20$. We obtain the HML portfolios' explained variations by estimating 9×6 OLS time-series regressions of the type of Eq. (7), i.e. $X_{nt} = \alpha_n + \hat{F}_t \psi_n^\top + \epsilon_{nt}$, for $n = 1, \ldots, N$, and $t = 1, \ldots, T$. For a given HML portfolio, we run a total of 6 regressions as we include factors one by one. In this way, we can determine the marginal R^2 contributions of each factor. The numbers above the bars refer to $\hat{F}_{1t}$'s contributions. We present the $N \times K$ individual portfolio risk exposures and R^2s in Figures A1 and A2, in the Internet Appendix. The sample spans the 11/1983-12/2017 period at monthly frequency ($T = 410$).

Figure 4: Observed vs Model-Implied Portfolio Excess Returns

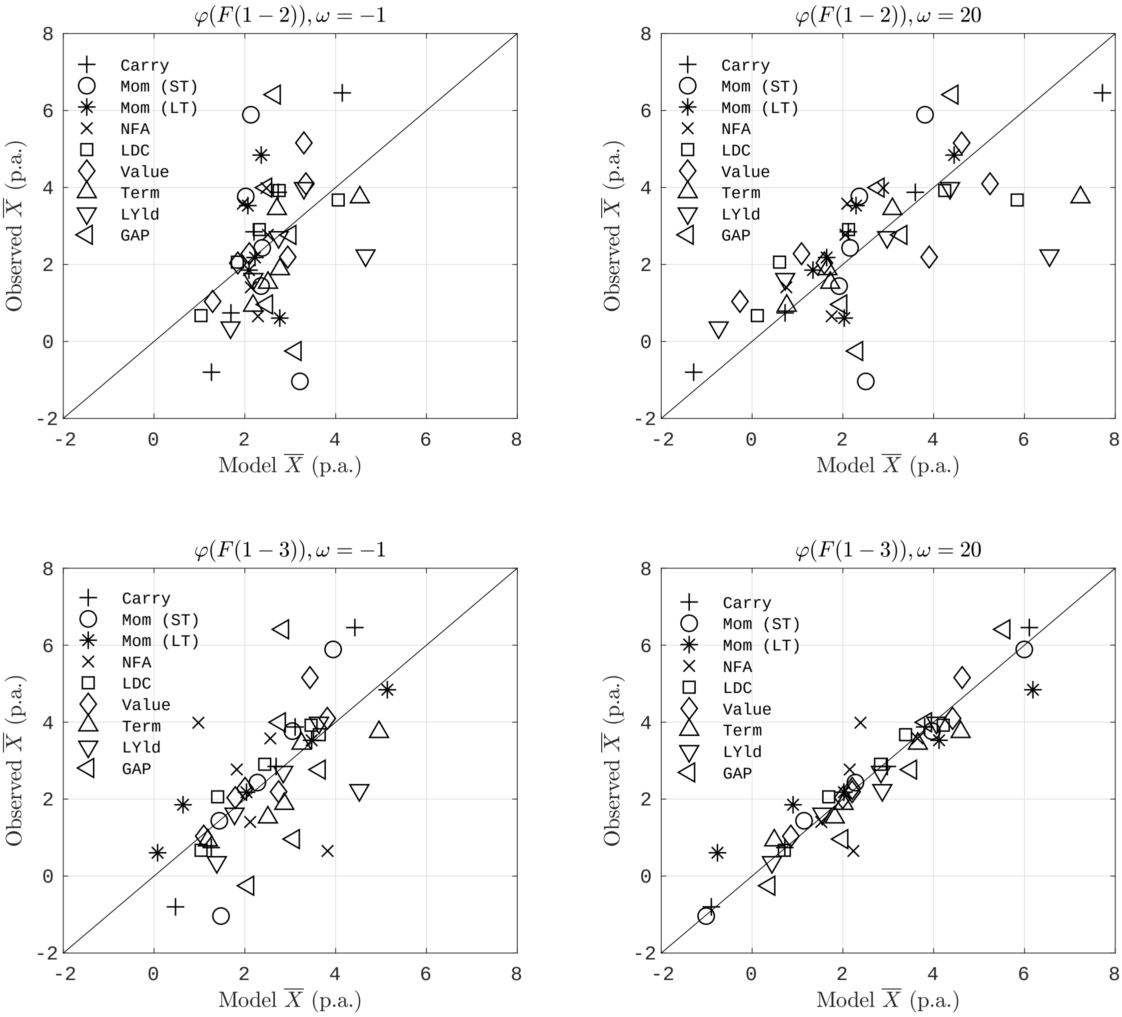

The figure shows the realized average excess returns (Observed $\overline{X}$) and the estimated, model-implied expected excess returns (Model $\overline{X}$) of the 46 currency portfolios (test assets) associated with the nine currency investment strategies, described in Section II in the Internet Appendix. For a given portfolio n, the model-implied expected excess returns are the fitted values from the second pass of the GX method. That is, the fitted return is given by $\hat{\psi}_{nk}\hat{\gamma}_k^\top$ where $\hat{\psi}_{nk}$ is the $1 \times k$ vector of the n-th portfolio's risk exposures and $\hat{\gamma}_k$ is the $1 \times k$ vector of the prices of risk of the latent factors. Left panels show the evidence for the factors extracted using PCA (i.e., RP-PCA with $\omega = -1$), while right panels using RP-PCA with $\omega = 20$ (i.e., the selected RP-weight). Top panels present the estimates for $k = 2$, while bottom panels for $k = 3$. Recall that $\varphi(F(1-2))$ is the optimal SDF for $\omega = -1$, whereas $\varphi(F(1-3))$ for $\omega = 20$. Excess returns are expressed in percent per annum (p.a). The sample period runs from 11/1983 to 12/2017 at monthly frequency ($T = 410$).

Figure 5: Nontradable Factor Explained Variations by Latent Pricing Factors

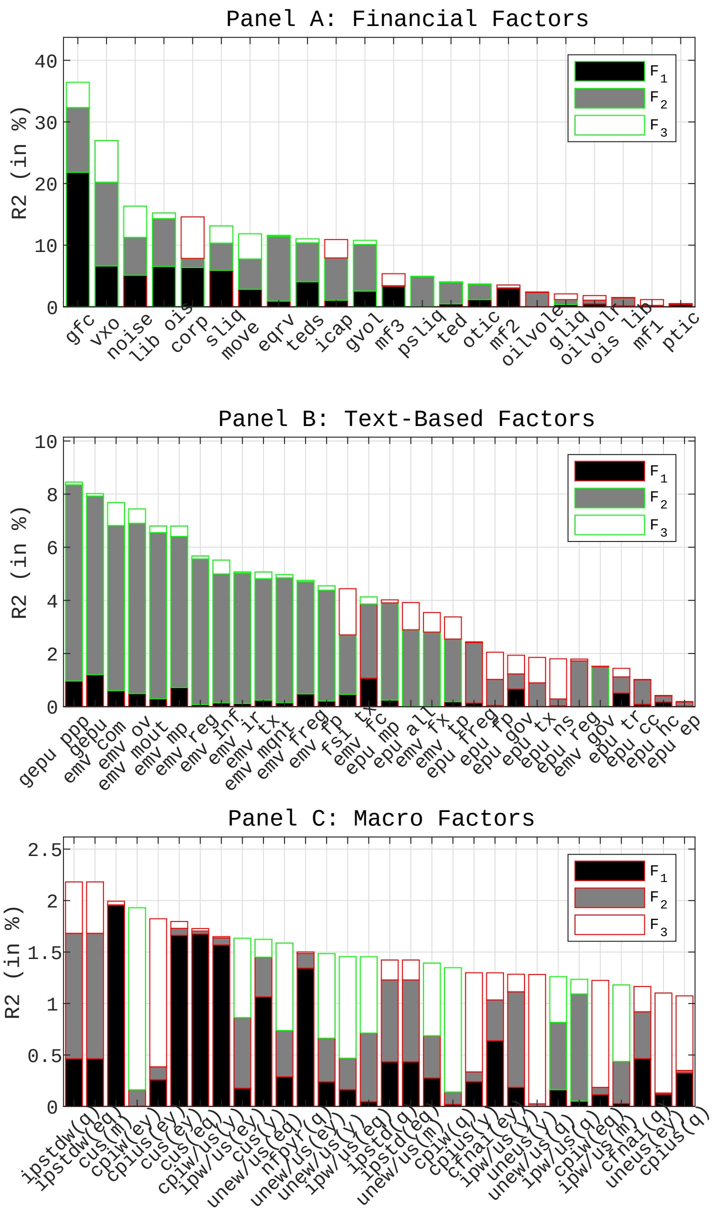

The figure shows nontradable candidate factor explained variation by latent factors. Factors are extracted by means of RP-PCA with baseline RP-weight $\omega = 20$ applied to the panel of $N = 46$ currency portfolios associated with the nine investment strategies. We limit the evidence to the first three extracted factors, given that the optimal SDF is $\varphi(F_{1-k})$ with $k = 3$ (Figure A3 in the Internet Appendix uses $k = 6$). Panels A, B, and C present the evidence for financial, text-based and macro candidate factors, respectively. In each panel, candidate factors are sorted by the spanning R^2s associated with the k-factor SDF (to help visualize the results, we only display the 30 macro factors with the highest R^2s). Bars quantify the latent factor contribution to the overall R^2. Green (red) bar edges denote (not) significant risk premia, at the 10 percent level, according to the selected SDF; we consider SDFs of expanding dimension, thus k ranges from 1 to 3. Factors are expressed as innovations, using the residuals from AR(1) processes, and are then standardized. The sample period varies with the factor at hand, according to data availability over the 11/1983-12/2017 period ($T = 410$). See factor descriptions in Tables A5-A7, in the Internet Appendix.

Figure 6: Return-Based Candidate Factors

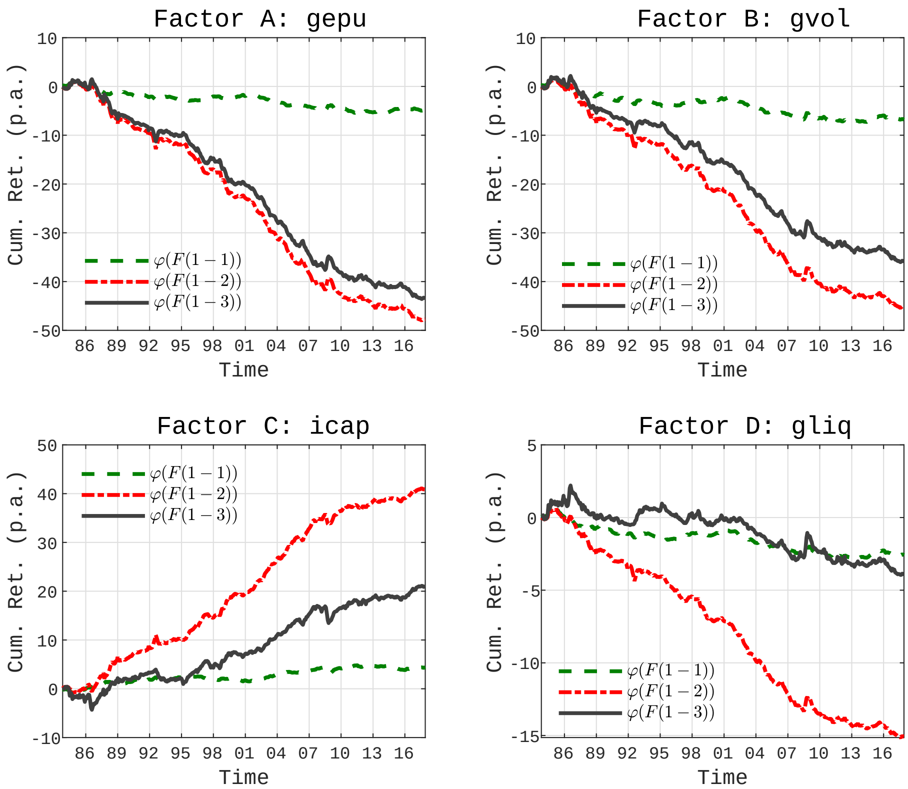

The figure shows selected return-based candidate factors. Factor A is the global economic policy uncertainty index (gepu), Factor B is the global FX volatility factor (gvol), Factor C is the financial intermediaries' capital ratio factor (icap), and Factor D is the global FX liquidity factor (gliq). For each candidate factor, we present three versions of return-based factors, by expanding the dimension of the SDF, i.e., $\varphi(F(1-k))$ with $k = 1, 2, 3$. Specifically, for $\varphi(F(1-k))$, the return-based factor is given by $\hat{F}_{1:kt}\hat{\eta}^{\top}$, where the latent factors are extracted using RP-PCA with $\omega = 20$ and the exposures are obtained by estimating the spanning regressions of Eq. (15); hence, the underlying exposures are those displayed in the first three columns of Table 2 (for *icap* and *gliq*, see Table A9 in the Internet Appendix, as their premia estimates are not statistically significant using the three-factor model). The spanning regression sample period can vary with the factor at hand, according to data availability over the 11/1983-12/2017 period. See factor descriptions in Tables A5-A7, in the Internet Appendix.

Internet Appendix
(not for publication)

Currency Risk Premia Redux

Table of Contents:

I The Two-Pass Estimator

A linear SDF is given by

$$\varphi_t = \xi[1 - (\mathrm{f}_t - \mathrm{E}[\mathrm{f}_t])\mathrm{b}^\top], \tag{I.1}$$

where ξ is the intercept, f_t is a $1 \times K$ vector of generic risk factors at time t, $\mathrm{E}[\mathrm{f}_t]$ is the vector of factor means, and b is the vector collecting the loadings of the SDF on the risk factors. Specifically, the loading b_k denotes the marginal contribution of the k-th factor to the SDF, conditional on the other factors. Then, let us denote the excess return on asset n at time t by X_{nt}, for $n = 1, \ldots, N$ test assets. Under no-arbitrage, risk-adjusted excess returns have a price of zero and satisfy the standard Euler pricing equation

$$\mathrm{E}(\varphi_t X_{nt}) = 0. \tag{I.2}$$

By combining Eqs. (I.1) and (I.2), and using some simple algebra, one can decompose expected excess returns as follows

$$\mathrm{E}[\mathrm{X}] = cov(\mathrm{X}_t, \mathrm{f}_t)\mathrm{b}^\top. \tag{I.3}$$

The linear SDF implies the following beta-pricing representation

$$\mathrm{E}[\mathrm{X}] = \underbrace{cov(\mathrm{X}_t, \mathrm{f}_t)\Sigma_f^{-1}}_{\beta} \underbrace{\Sigma_f \mathrm{b}^\top}_{\lambda^\top}, \tag{I.4}$$

whereby expected excess returns depend on the $1 \times K$ vector of risk prices (λ), and the $N \times K$ matrix of risk quantities (β). Thus, the n-th asset excess return is given by $\mathrm{E}[\mathrm{X}_n] = \beta_n \lambda^\top$. It is apparent that risk exposures are asset specific, while the prices of risk are common to all assets. That is, the price of risk of the k-th risk factor (λ_k) denotes the compensation required by the investor for a unit exposure to that factor.

It is common practice to estimate asset-specific risk exposures (β_n) and the prices of risk (λ) using the two-pass procedure of FMB. The first pass delivers estimates of the risk exposures, whereas the second pass of the prices of risk. Specifically, the first pass consists of running N OLS time-series regressions of test-asset excess returns on the vector of risk factors

$$X_{nt} = \alpha_n + f_t \beta_n^\top + \epsilon_{nt}, \quad n = 1, \ldots, N, \quad t = 1, \ldots T, \tag{I.5}$$

where the intercept α_n denotes the n-th asset risk-adjusted excess return. Then, the second pass is a cross-sectional regression of the test assets' expected returns on the previously estimated betas $\hat{\beta}_n$, including or not a constant. Empirically, it is standard to proxy expected returns, $\mathrm{E}[\mathrm{X}_n]$, with the average realized excess returns, $\overline{X}_n = \frac{1}{T}\sum_{t=1}^T X_{nt}$, so that the second-pass cross-sectional regression is given by

$$\overline{X}_n = \hat{\beta}_n \lambda^\top + a_n, \quad n = 1, \ldots, N, \tag{I.6}$$

where a_n is the n-th asset's pricing error, λ is the $1 \times K$ vector of prices of risk, and $\hat{\beta}_n \hat{\lambda}^\top$ is the n-th asset's model-implied risk premium, or excess return. The estimates of the prices of risk are given by $\hat{\lambda} = \overline{X}\hat{\beta}(\hat{\beta}^\top \hat{\beta})^{-1}$, where $\hat{\beta}$ is the matrix collecting the assets estimated risk exposures, and $\overline{X}$ is the vector of assets' average realized excess returns. Note that it might be useful to include a constant in Eq. (I.6), $\overline{X}_n = c + \hat{\beta}_n \lambda^\top + a_n$, to capture the common mispricing in test-asset expected excess returns (e.g., see Burnside, 2011).

II FX Data and Investment Strategies

`FX Data`. Table A1 reports the Datastream mnemonics for spot and forward exchange rates for the 49 currencies used in this paper.

`49 USD/FCU Currencies (1983-2017)`. Australia, Austria, Belgium, Brazil, Bulgaria, Canada, Croatia, Cyprus, Czech Republic, Denmark, Egypt, euro area, Finland, France, Germany, Greece, Hong Kong, Hungary, Iceland, India, Indonesia, Ireland, Israel, Italy, Japan, Kuwait, Malaysia, Mexico, Netherlands, New Zealand, Norway, Philippines, Poland, Portugal, Russia, Saudi Arabia, Singapore, Slovakia, Slovenia, South Africa, South Korea, Spain, Sweden, Switzerland, Taiwan, Thailand, Ukraine, the United Kingdom, and Turkey.

`15 FCU/GBP Currencies (1978-1983)`. To construct the signals needed to form the portfolios for some trading strategies, we complement our data with spot exchange rates and one-month forward rates quoted against the British Pound over the period from January 1976 to October 1983, which are also provided by Datastream. We use triangular no-arbitrage relations to retrieve exchange and forward rate quotes against the U.S. dollar (i.e., by using FCU/GBP and USD/GBP, one obtains USD/FCU). By doing this, we obtain longer time series for spot and forward rates data for a number of developed countries. The 15 FCU/GBP currencies are: Austria, Belgium, Canada, Denmark, France, Germany, Ireland, Italy, Japan, Netherlands, Norway, Portugal, Spain, Sweden, and Switzerland.

`CIP Deviations`. We follow Kroencke et al. (2014) and Della Corte et al. (2016), among others, and leave out the indicated countries for the following periods: Egypt (01/01/2011 – 30/08/2013; 03/10/2016 – 28/02/2017); Indonesia (01/12/1997 – 31/07/1998; 01/02/2001 – 31/05/2005); Malaysia (01/05/1998 – 30/06/2005); Russia (01/12/2008 – 30/01/2009; 03/11/2014 – 27/02/2015); South Africa (01/08/1985 – 30/08/1985; 01/01/2002 – 31/05/2005); Turkey (01/11/2000 – 30/11/2001); and, Ukraine (03/11/2014 – 31/12/2017).

`FX Investment Strategies`. Next, we describe the currency investment strategies that deliver the portfolios (i.e., test assets) under investigation in our empirical analysis. Table A2 provides information on strategy names, signals, and data sources. Table A3 presents the summary statistics of the currency investment strategy portfolios.

`1. Carry`. At the end of each month t, currencies are allocated to five portfolios according to their forward discounts $(S_{it} - F_{it})/S_{it}$, where S_{it} and F_{it} are the spot and forward exchange rate mid-quotes for foreign currency i, respectively (Lustig et al., 2011; Menkhoff et al., 2012a). While portfolio 1 (P1) collects the currencies with the lowest forward discounts, portfolio 5 (P5) collects currencies with the highest forward discounts. Therefore, P1 (P5) contains the currencies with the lowest (highest) interest rate differential relative to the United States, assuming that CIP holds. Portfolios are rebalanced monthly and the sample runs from November 1983 to December 2017. Data are from from Barclays Bank International (BBI), Reuters and WM/Reuters accessed via Datastream.

2. `Short-Term Momentum`. At the end of each month t, we form portfolios based on excess returns realized over the previous m months, i.e., the formation period (e.g., Asness et al., 2013; Menkhoff et al., 2012b). In Short-Term Momentum, for each foreign currency i, the sorting variable (signal) is the previous month currency excess return, X_t. While P1 contains the currencies with the lowest short-term excess returns, that is the loser currencies over the short period, P5 contains the currencies with the highest short-term excess returns, that is the winner currencies over the short period. Portfolios are rebalanced monthly and the sample runs from November 1983 to December 2017. Data are from from Barclays Bank International (BBI), Reuters and WM/Reuters accessed via Datastream.

3. `Long-Term Momentum`. The Long-Term Momentum strategy is based on a longer formation period than Short-Term Momentum (e.g., Asness et al., 2013; Menkhoff et al., 2012b). In particular, the signal is the cumulative excess return over the past 12 months, $X_{t-12:t}$. Hence, P1 contains the currencies with the lowest long-term excess returns, that is the loser currencies over the long period, whereas P5 contains the currencies with the highest long-term excess returns, that is the winner currencies over the long period. Portfolios are rebalanced monthly and the sample runs from November 1983 to December 2017. Data are from from Barclays Bank International (BBI), Reuters and WM/Reuters accessed via Datastream.

4. `Currency Value`. At the end of each month t, currencies are allocated to portfolios based on the lagged five-year real exchange rate return (e.g., Asness et al., 2013; Kroencke et al., 2014; Menkhoff et al., 2017), thus exploiting deviations from the relative purchasing power parity (PPP). P1 contains over-valued currencies, i.e., those with the highest lagged real exchange rate return, and P5 contains under-valued currencies, i.e., those with the lowest lagged real exchange rate return. In line with the literature, and differently from the other strategies, value portfolios are rebalanced every three months. The sample runs from November 1983 to December 2017. Real exchange rates are calculated by using Consumer Price Index data from IMF International Financial Statistics (the source of Taiwan's CPI data is the National Statistics).

5. `Net Foreign Assets`. Following Della Corte et al. (2016), at the end of each month t, currencies are allocated into portfolios according to the ratio between the foreign country's net foreign assets (NFA) and the country's gross domestic product (GDP), both denominated in U.S. dollar, multiplied by (-1). Hence, P1 includes creditor currencies, i.e., those with the highest NFA to GDP ratios, whereas P5 includes debtor currencies, i.e., those with the lowest NFA to GDP ratios. Portfolios are rebalanced monthly and the sample runs from November 1983 to December 2017. We thank Gian Maria Milesi-Ferretti to kindly share with us the updated version of the data on foreign assets and liabilities and GDP used in Lane and Milesi-Ferretti (2004), and Lane and Milesi-Ferretti (2007).

6. `Liabilities in Domestic Currencies`. Following Della Corte et al. (2016), we sort currencies into portfolios according to the proportion of liabilities denominated in domestic currency (LDC). At the beginning of each month t, we first allocate currencies to two portfolios according to the foreign country's lagged NFA to GDP ratio. Then, we further split each of these two portfolios into three sub-portfolios according

to the country's liabilities denominated in domestic currency (LDC). The 2×3 double-sorted portfolios are denoted by LL, LM, LH, HL, HM, and HH, where the first letter denotes the relative level of the NFA ratio, and the second letter that of LDC. As a result, we obtain 6 double-sorted LDC portfolios. P1 contains the safest currencies, i.e., those with high NFA positions coupled with a large share of their liabilities denominated in domestic currency (HH), and P6 contains the riskiest currencies, i.e., those with low NFA positions coupled with a low share of their liabilities denominated in domestic currency (LL). Portfolios are rebalanced monthly and the sample runs from November 1983 to December 2017. Data on LDC, also used in Benetrix et al. (2015), are available on Philip Lane's website.

7. `Long-Term Yields`. At the end of each month t, we sort currencies into portfolios according to the foreign country's 10-year interest rate differential relative to that of the United States ($i_{10yr} - i_{10yr}^{US}$), so that P1 includes countries with the lowest interest rates, and P5 those with the highest interest rates (e.g., Ang and Chen, 2010; Della Corte et al., 2016). Sorting on long-term interest rates allows us to capture departures from uncovered interest rate parity at the longer end of the term structure of interest rates. The portfolios are rebalanced monthly and the sample runs from November 1983 to December 2017. Data on long-term interest rates are from the OECD Monthly Monetary and Financial Statistics.

8. `Term Spread`. At the end of each month t, we sort currencies into portfolios according to the foreign country's term spread, defined as long- minus short-term rates, measured with the 10-year and 3-month rates ($i_{10yr} - i_{3mo}$), respectively (Bekaert et al., 2007; Lustig et al., 2019). We allocate to P1 countries with the highest term spread, and conversely to P5 countries with the lowest term spread. The portfolios are rebalanced every six months and the sample runs from November 1983 to December 2017. Data on long-term and short-term interest rates are from the OECD Monthly Monetary and Financial Statistics.

9. `Output Gap`. Following Colacito et al. (2020), at each month t, we sort currencies on difference between each foreign country's output gap and the US output gap, i.e., $GAP_t - GAP_t^{US}$. We allocate to P1 the currencies with the lowest output gap relative to the U.S. (low output gap currencies) whereas we place in P5 the currencies with the highest output gap relative to the U.S. (high output gap currencies). The portfolios are rebalanced every month and the sample runs from November 1983 to January 2016. Output gaps are calculated by using industrial production data from the Organisation for Economic Co-operation and Development's (OECD's) Original Release Data and Revisions Database. Output gaps are estimated using the Hodrick-Prescott filter to extract a cyclical component from the data. We thank Colacito et al. (2020) for providing the portfolio returns; given that the returns are available only until January 2016, we fill the few missing observations using the nuclear-norm penalized regression approach recently employed by Giglio et al. (2021b).

Table A1: FX Spot and Forward Rates: Data Sources

This table reports in the first column the countries analyzed in the empirical analysis. The second column reports Datastream (DS) mnemonics of the exchange rates data used for each country. Daily spot and forward exchange rates are mainly sourced from Barclays Bank International (BBI), Reuters and WM/Reuters accessed via Datastream.

Country	Spot	Forward
Australia	BBAUDSP	BBAUD1F
Austria	AUSTSC$, AUSTSCH	USATS1F, AUSTS1F
Belgium	BELGLU$, BELGLUX	USBEF1F, BELXF1F
Brazil	BRACRU$	USBRL1F
Bulgaria	BULGLV$	USBGN1F
Canada	BBCADSP, CNDOLLR	BBCAD1F, CNDOL1F
Croatia	CROATK$	USHRK1F
Cyprus	CYPRUS$	USCYP1F
Czech Rep	CZECHC$	USCZK1F
Denmark	BBDKKSP, DANISHK	BBDKK1F, DANIS1F
Egypt	EGYPTN$	USEGP1F
Euro	BBEURSP	BBEUR1F
Finland	FINMAR$	USFIM1F
France	BBFRFSP, FRENFRA	BBFRF1F, FRENF1F
Germany	BBDEMSP, DMARKER	BBDEM1F, DMARK1F
Greece	GREDRA$	USGRD1F
Hong Kong	BBHKDSP	BBHKD1F
Hungary	HUNFOR$	USHUF1F
Iceland	ICEKRO$	USISK1F
India	INDRUP$	USINR1F
Indonesia	INDORU$	USIDR1F
Ireland	BBIEPSP, IPUNTER	BBIEP1F, IPUNT1F
Israel	ISRSHE$	USILS1F
Italy	BBITLSP, ITALIRE	BBITL1F, ITALY1F
Japan	BBJPYSP, JAPAYEN	BBJPY1F, JAPYN1F
Kuwait	KUWADI$	USKWD1F
Malaysia	MALADL$	USMYR1F
Mexico	MEXPES$	USMXN1F
Netherlands	BBNLGSP, GUILDER	BBNLG1F, GUILD1F
New Zealand	BBNZDSP	BBNZD1F
Norway	BBNOKSP, NORKRON	BBNOK1F, NORKN1F
Philippines	PHILPE$	USPHP1F
Poland	POLZLO$	USPLN1F
Portugal	PORTES$, PORTESC	USPTE1F, PORTS1F
Russia	CISRUB$	USRUB1F
Saudi Arabia	SAUDRI$	USSAR1F
Singapore	BBSGDSP	BBSGD1F
Slovakia	SLOVKO$	USSKK1F
Slovenia	SLOVTO$	USSIT1F
South Africa	BBZARSP	BBZAR1F
South Korea	KORSWO$	USKRW1F
Spain	SPANPE$, SPANPES	USESP1F, SPANP1F
Sweden	BBSEKSP, SWEKRON	BBSEK1F, SWEDK1F
Switzerland	BBCHFSP, SWISSFR	BBCHF1F, SWISF1F
Taiwan	TAIWDO$	USTWD1F
Thailand	THABAH$	USTHB1F
Turkey	TURKLI$	USTRY1F
United Kingdom	BBGBPSP	BBGBP1F
Ukraine	UKRAHY$	USUAH1F
United States	USDOLLR	USDOL1F

Table A2: FX Investment Strategies: Data Sources

The table provides information about the data used for the construction of the nine currency investment strategies analyzed in the paper. First column (*Strategy*) reports the currency strategy's name, and its short name we use in the paper (`SN: xxx`). Second column (*Signal*) specifies the trading signal used by the strategy, while third column (*Source*) reports the sources of the data used for the construction of the signals.

Strategy	Signal	Source
1. Carry `SN: Carry`	$(S_t - F_t)/S_t$	Authors' calculation based on spot and forward exchange rate quotes (midquote) from Barclays Bank International (BBI), Reuters and WM/Reuters accessed via Datastream.
2. Short-term Momentum `SN: ST Mom`	X_t	Authors' calculation based on spot and forward exchange rate quotes (midquote) from Barclays Bank International (BBI), Reuters and WM/Reuters accessed via Datastream.
3. Long-term Momentum `SN: LT Mom`	$X_{t-12:t}$	Authors' calculation based on spot and forward exchange rate quotes (midquote) from Barclays Bank International (BBI), Reuters and WM/Reuters accessed via Datastream.
4. Currency Value `SN: Value`	5yr Rel. PPP Level Dev.	Real exchange rates are calculated by using Consumer Price Index data from IMF International Financial Statistics. Only for Taiwan CPI data come from National Statistics.
5. Net Foreign Assets `SN: NFA`	−NFA/GDP	Data are kindly provided to us by G.M. Ferretti, also available on IMF website as BOP/IIP Statistics. We use the latest data available to fill the missing observations for the latest years.
6. Liabilities Domestic Currency `SN: LDC`	−NFA/GDP; LDC	LDC data are available on Philip Lane's website, http://www.philiplane.org/. We use the latest data available to fill the missing observations for the latest years.
7. Long-term Yields `SN: LYld`	$i_{10yr} - i_{10yr}^{US}$	For the 10-year rates, we use the Long-term interest rates available on OECD Monthly Monetary and Financial Statistics.
8. Term Spread `SN: Term`	$-(i_{10yr} - i_{3mo}^{k})$	For the 10-year (3-month) rates, we use the Long (Short)-term interest rates available on OECD Monthly Monetary and Financial Statistics.
9. Output GAP `SN: GAP`	$GAP_t - GAP_t^{US}$	We thank Colacito et al. (2020) for providing us with currency portfolios sorted on the difference between each country's output gap and the US output gap. Output gap is defined as the logarithm of the difference between actual output and potential output.

Table A3: FX Investement Strategies: Summary Statistics

The table presents the summary statistics of the nine currency investment strategies, i.e., plain portfolios (P), HML cross-sectional (CS) and EW time series (TS) factors. We report the following excess-return statistics for the selected portfolios: mean return in p.p.a. (mean), mean p-value (pval), median (med.), standard deviation in p.p.a. (st.dev.), skewness (skew), kurtosis (kurt), Sharpe ratio annualized (SR), first-order autocorrelation coefficient (AC(1)), and number of observations (obs.).

Carry	P1	P2	P3	P4	P5	CS	TS
mean	-0.80	0.74	2.85	3.87	6.46	7.26	3.68
pval	0.57	0.56	0.06	0.01	0.00	0.00	0.00
med.	-1.40	1.41	2.59	4.63	8.55	9.32	4.84
st.dev.	7.84	7.09	8.01	8.23	9.38	8.39	7.09
skew	0.27	-0.11	-0.15	-0.42	-0.33	-0.81	-0.41
kurt	4.18	4.03	4.16	4.54	4.82	5.04	4.08
SR	-0.10	0.10	0.36	0.47	0.69	0.86	0.52
AC(1)	0.01	0.05	0.10	0.06	0.14	0.13	0.08
obs	410	410	410	410	410	410	410

Short-Term Momentum	P1	P2	P3	P4	P5	CS	TS
mean	-1.04	1.44	2.43	3.77	5.89	6.93	2.60
pval	0.52	0.34	0.12	0.01	0.00	0.00	0.03
med.	-0.09	3.41	3.65	3.92	5.34	5.99	3.20
st.dev.	9.10	8.18	8.26	8.01	8.11	9.00	7.13
skew	-0.48	-0.50	-0.40	0.15	0.23	0.28	0.07
kurt	5.72	6.51	5.59	4.11	3.95	5.26	3.79
SR	-0.11	0.18	0.29	0.47	0.73	0.77	0.37
AC(1)	0.01	0.05	0.11	0.04	0.07	-0.03	0.05
obs	410	410	410	410	410	410	410

Long-Term Momentum	P1	P2	P3	P4	P5	CS	TS
mean	0.61	1.85	2.18	3.53	4.84	4.24	2.63
pval	0.68	0.19	0.15	0.02	0.00	0.00	0.04
med.	1.23	2.98	1.19	3.92	6.21	6.12	3.57
st.dev.	8.63	7.92	8.28	8.49	8.06	8.95	7.13
skew	0.27	0.17	-0.05	-0.24	-0.68	-0.53	-0.29
kurt	5.45	4.47	4.28	4.42	6.40	4.10	3.94
SR	0.07	0.23	0.26	0.42	0.60	0.47	0.37
AC(1)	0.04	0.02	0.05	0.10	0.09	-0.05	0.08
obs	410	410	410	410	410	410	410

Value	P1	P2	P3	P4	P5	CS	TS
mean	0.65	3.58	1.41	2.77	3.99	3.33	-1.96
pval	0.69	0.02	0.39	0.08	0.01	0.01	0.16
med.	1.16	4.77	1.96	2.56	3.91	2.98	-3.31
st.dev.	8.78	8.28	9.03	8.81	8.18	7.47	7.59
skew	-0.41	-0.41	-0.14	0.03	0.24	0.11	0.42
kurt	4.87	4.39	4.25	4.16	4.03	3.95	3.98
SR	0.07	0.43	0.16	0.31	0.49	0.45	-0.26
AC(1)	0.11	0.10	0.03	0.04	-0.01	0.05	0.07
obs	410	410	410	410	410	410	410

Net Foreign Assets	P1	P2	P3	P4	P5	CS	TS
mean	0.67	2.06	2.91	3.92	3.67	3.00	1.32
pval	0.45	0.19	0.06	0.01	0.04	0.03	0.32
med.	0.20	1.83	3.00	6.04	5.44	4.77	1.53
st.dev.	5.02	8.65	8.34	7.82	9.94	7.63	7.16
skew	0.30	0.03	-0.21	-0.62	-0.30	-0.35	-0.06
kurt	4.74	3.83	4.01	4.77	5.14	4.88	3.82
SR	0.13	0.24	0.35	0.50	0.37	0.39	0.18
AC(1)	0.01	0.06	0.07	0.16	0.05	0.07	0.08
obs	410	410	410	410	410	410	410

Liabilities in Domestic Currency	P1	P2	P3	P4	P5	P6	CS
mean	1.04	2.28	2.04	2.19	4.10	5.16	4.12
pval	0.45	0.18	0.10	0.13	0.01	0.01	0.00
med.	1.70	2.32	2.35	4.29	4.30	6.98	5.67
st.dev.	7.57	9.40	7.15	7.69	8.82	9.87	6.61
skew	0.06	-0.02	-0.20	-0.25	-0.49	-0.36	-0.52
kurt	4.19	3.91	3.71	4.55	6.19	4.57	6.09
SR	0.14	0.24	0.28	0.29	0.46	0.52	0.62
AC(1)	0.06	0.07	0.02	0.10	0.04	0.12	0.15
obs	410	410	410	410	410	410	410

Term Spread	P1	P2	P3	P4	P5	CS	TS
mean	0.93	1.52	1.88	3.44	3.74	2.82	-0.58
pval	0.57	0.37	0.32	0.03	0.06	0.08	0.65
med.	1.71	2.40	3.96	3.51	4.79	5.64	-0.33
st.dev.	9.12	9.69	10.23	9.04	10.76	9.06	7.17
skew	-0.08	0.05	-0.48	-0.07	-0.75	-1.01	-0.21
kurt	4.09	3.77	4.91	4.15	6.61	6.74	3.78
SR	0.10	0.16	0.18	0.38	0.35	0.31	-0.08
AC(1)	0.04	0.04	0.08	0.02	0.08	0.03	0.05
obs	410	410	410	410	410	410	410

Long-Term Yields	P1	P2	P3	P4	P5	CS	TS
mean	0.36	1.62	2.71	3.98	2.23	1.87	3.32
pval	0.84	0.36	0.08	0.02	0.24	0.25	0.01
med.	-0.16	1.64	1.68	5.83	3.25	4.47	4.26
st.dev.	9.64	9.85	8.66	9.65	10.55	9.24	7.11
skew	0.17	-0.09	-0.30	-0.47	-0.63	-0.91	-0.23
kurt	3.30	3.92	5.08	4.58	7.00	6.37	3.96
SR	0.04	0.16	0.31	0.41	0.21	0.20	0.47
AC(1)	0.05	0.06	0.04	0.03	0.08	0.01	0.08
obs	410	410	410	410	410	410	410

Output Gap	P1	P2	P3	P4	P5	CS	TS
mean	-0.25	0.96	2.77	4.00	6.41	6.66	–
pval	0.89	0.58	0.15	0.02	0.00	0.00	–
med.	1.16	2.12	3.15	4.76	7.15	5.28	–
st.dev.	10.18	9.09	10.12	9.32	9.05	8.14	–
skew	-0.06	-0.47	-0.28	-0.27	-0.28	0.01	–
kurt	4.49	4.72	4.75	4.39	3.97	4.32	–
SR	-0.02	0.11	0.27	0.43	0.71	0.82	–
AC(1)	0.03	0.08	0.08	-0.01	0.07	-0.03	–
obs	387	387	387	387	387	387	–

Table A4: Correlations of HML Portfolios

The table presents the correlation matrix of the HML spread portfolios of the nine currency investment strategies described in Section II.

	Carry	ST Mom	LT Mom	Value	NFA	LDC	Term	LYld	GAP
Carry	1	–0.10	0.04	0.08	0.47	0.55	0.54	0.76	0.00
ST Mom	–0.10	1	0.25	–0.05	–0.15	–0.19	–0.09	–0.17	0.13
LT Mom	0.04	0.25	1	–0.39	–0.08	–0.02	0.04	–0.06	0.14
Value	0.08	–0.05	–0.39	1	0.11	–0.02	–0.11	0.14	0.09
NFA	0.47	–0.15	–0.08	0.11	1	0.60	0.25	0.49	–0.06
LDC	0.55	–0.19	–0.02	–0.02	0.60	1	0.31	0.40	–0.01
Term	0.54	–0.09	0.04	–0.11	0.25	0.31	1	0.49	–0.08
LYld	0.76	–0.17	–0.06	0.14	0.49	0.40	0.49	1	0.01
GAP	0.00	0.13	0.14	0.09	–0.06	–0.01	–0.08	0.01	1

III Nontradable Candidate Risk Factors

Table A5: Financial Factors

This table reports in the first and second columns, respectively, financial factors' names (*Name*) and short descriptions of the factors and the sources of the data used for the construction of the factors (*Description*). Finally, the third (*Start*) and fourth (*End*) columns show the period in which the factors are available. In the empirical analysis, we use innovations of the factors, i.e. residuals from AR(1) processes, instead of the raw factors.

Name	Description	Start	End
move vxo	Merrill Lynch Option Volatility Estimate Index (move), and CBOE S&P 100 Volatility Index (vxo). Source: `https://fred.stlouisfed.org/`, and Bloomberg	03/1989 11/1986	12/2017 12/2017
MF1 MF2 MF3	First three principal components extracted from a large dataset of macro and financial time series. Source: Updated Macro Factors in Bond Risk Premia database available on Sidney Ludvigson's website	11/1983	12/2017
gvol gliq	Global FX volatility (gvol) and liquidity (gliq) factors, constructed as in Menkhoff et al. (2012a). Source: Authors' calculations based on daily exchange rate quotes from Barclays Bank International (BBI), Reuters and WM/Reuters via Datastream	10/1983	12/2017
psliq	Pastor and Stambaugh equity liquidity factor (Pastor and Stambaugh, 2003). Source: Lubos Pastor's website	11/1983	12/2017
sliq	Systematic, low frequency, FX (il)liquidity factor (Karnaukh et al., 2015). Source: Angelo Ranaldo's website	02/1991	12/2017
ted	TED spread. Source: `https://fred.stlouisfed.org/`	02/1986	12/2017
noise	Market-wide liquidity measure based on the connection between the amount of arbitrage capital in the market and observed price deviations (noise) in U.S. Treasury bonds (Hu et al., 2013). Source: Jun Pan's website	02/1987	12/2016
icap	Intermediary capital risk factor based on the equity capital ratio of financial intermediaries (He et al., 2017). Source: Zhiguo He's website	11/1983	12/2017
oilvolr oilvole	Realized volatilities based on squared daily returns (oilvolr) and residuals (oilvolr) of a AR(1) process fitted to WTI oil price. Source: `https://fred.stlouisfed.org/`	11/1983	12/2017
gcf	Global financial cycle factor extracted from a dynamic factor model for a large and heterogeneous panel of risky asset prices traded around the globe (Miranda-Agrippino and Rey, 2020). Source: S. Miranda-Agrippino's website	11/1983	12/2017
otic ptic	Official (otic) and Private (ptic) net inventories in U.S. Treasuries from the Treasury International Capital (TIC) System, standardized over rolling 3-year standard deviation. Source: `https://www.treasury.gov`	11/1983	12/2017
corp	spread between BAA and AAA rated bond yields. Source: `https://fred.stlouisfed.org/`	01/1986	12/2017
lib-ois ois-tbill	Libor-OIS spread and OIS-TBill spread based on Libor, TBill, OIS data. Source: `https://fred.stlouisfed.org/`	12/2001	12/2017
eqrv	S&P500 monthly realized volatility. Authors' calculations using daily S&P500 closing prices. Source: Datastream	11/1983	12/2017

Table A6: Macro Factors

This table reports in the first and second columns, respectively, macroeconomic factors' names (*Name*) and short descriptions of the factors and the sources of the data used for the construction of the factors (*Description*). The third (*Start*) and fourth (*End*) columns show the period in which the factors are available. Each factor in the list is used in the paper for different time-formation periods. To start with, (m) is the benchmark time-formation period and indicates either a month-on-month log growth rate (for the factors from *ipus* to *cpiw/us*), or the end-of-month value (for the factors from *ipstd* to *cpistdw*). Then, we measure the variables at other frequencies using their three-month simple moving averages (q); 12-month simple moving averages (y); three-month exponential moving averages (eq); and 12-month exponential moving averages (ey). In the empirical analysis, we use innovations of the factors, i.e. residuals from AR(1) processes, instead of the raw factors.

Name	Description	Start	End
ipus	Log growth rate in U.S. industrial production index. Source: `https://fred.stlouisfed.org/`	11/1983	12/2017
cpius	Log growth rate in U.S. consumer price index. Source: `https://fred.stlouisfed.org/`	11/1983	12/2017
nfpyr	Log growth rate in U.S. non-farm payroll. Source: `https://fred.stlouisfed.org/`	11/1983	12/2017
cfnai	U.S. Chicago Fed national activity index. Source: `https://fred.stlouisfed.org/`	11/1983	12/2017
uneus	Log growth rate in U.S. harmonized unemployment rate (total). Source: `https://fred.stlouisfed.org/`	11/1983	12/2017
cus	Log growth rate in U.S. aggregate real per capita consumption expenditures on nondurable goods and services. Source: `https://fred.stlouisfed.org/`	11/1983	12/2017
ipw	Cross-country GDP-weighted average of log growth rates in industrial production. Source: `https://stats.oecd.org/`	11/1983	12/2017
ipw/us	Cross-country GDP-weighted average of log growth rates in industrial production minus log growth rate in U.S. industrial production. Source: `https://stats.oecd.org/`	11/1983	12/2017
cpiw	Cross-country GDP-weighted average of log growth rates in consumer price index. Source: `https://stats.oecd.org/`	11/1983	12/2017
cpiw/us	Cross-country GDP-weighted average of log growth rates in consumer price index minus log growth rate in U.S. consumer price index. Source: `https://stats.oecd.org/`	11/1983	12/2017
unew	Cross-country GDP-weighted average of unemployment rate log growth rates. Source: `https://stats.oecd.org/`	11/1983	12/2017
unew/us	Cross-country GDP-weighted average of log growth rates in unemployment rate minus log growth rate in U.S. unemployment rate. Source: `https://stats.oecd.org/`	11/1983	12/2017
ipstd	Cross-country standard deviation of log growth rates in industrial production. Source: `https://stats.oecd.org/`	11/1983	12/2017
ipstdw	Cross-country standard deviation of GDP-weighted log growth rates in industrial production. Source: `https://stats.oecd.org/`	11/1983	12/2017
cpistd	Cross-country standard deviation of log growth rates in consumer price index. Source: `https://stats.oecd.org/`	11/1983	12/2017
cpistdw	Cross-country standard deviation of GDP-weighted log growth rates in consumer price index. Source: `https://stats.oecd.org/`	11/1983	12/2017

Table A7: Text-Based Factors

This table reports in the first column (*Name*) text-based factors' names. The second column (*Description*) reports short descriptions of the risk factors and the sources of the data used for the construction of the factors. Finally, the third (*Start*) and fourth (*End*) columns report the period in which the factors are available. In the empirical analysis, we use innovations of the factors, i.e. residuals from AR(1) processes, instead of the raw factors.

Name	Description	Start	End
gepu gepu_ppp	Monthly indexes of Global Economic Policy Uncertainty computed as a GDP-weighted (gepu) average or as a PPP-adjusted GDP-weighted (gepu_ppp) average of national EPU indexes for 21 countries. Source: `https://www.policyuncertainty.com`	01/1997 01/1997	12/2017 12/2017
fsi_tx	newspaper-based Financial Stress Indicator for the U.S. Source: `https://www.policyuncertainty.com`	11/1983	12/2016
emv_ov emv_mout emv_mqnt emv_inf emv_com emv_ir emv_fc emv_fx emv_fp emv_tx emv_gov emv_mp emv_reg emv_freg emv_tp	Newspaper-based Equity Market Volatility (EMV) trackers that move with the CBOE Volatility Index (VIX) and with the realized volatility of returns on the S&P 500. In addition to the Overall EMV Tracker (emv_ov), other category-specific EMV trackers are also considered, e.g., Macroeconomic News and Outlook (emv_mout), Macro - Broad Quantity Indicators (emv_mqnt), Macro - Inflation Indicator (emv_inf), Commodity Markets (emv_com), Macro - Interest Rates (emv_ir), Financial Crises (emv_fc), Exchange Rates (emv_fx), Fiscal Policy (emv_fp), Taxes (emv_tx), Government Spending, Deficits, and Debt (emv_gov), Monetary Policy(emv_mp), Regulation (emv_reg), Financial Regulation (emv_freg), Trade Policy (emv_tp). Source: `https://www.policyuncertainty.com`	01/1985	12/2017
epu_all epu_mp epu_fp epu_tx epu_gov epu_hc epu_ns epu_ep epu_reg epu_freg epu_tr epu_cc	Indexes of U.S. economic policy uncertainty that hinge upon newspaper coverage frequency. In addition to the general Economic Policy Uncertainty index (epu_all), 11 categorical sub-indexes are also considered, e.g., Monetary policy (epu_mp), Fiscal Policy (epu_fp), Taxes (epu_tx), Government spending (epu_gov), Health care (epu_hc), National security (epu_ns), Entitlement programs (epu_ep), Regulation (epu_reg), Financial Regulation (epu_reg), Trade policy (epu_tr), Sovereign debt, currency crises (epu_cc). Source: `https://www.policyuncertainty.com`	01/1985	12/2017

IV Additional Baseline Analysis

In what follows, we provide a more detailed description of the results regarding the latent-factor signal strengths and the associated SDFs.

Latent Factor "Signal Strengths". Table A8 presents the "signal-strength" analysis focusing on the six largest eigenvalues and the associated factors, as we find that the remaining eigenvalues have negligible "signal strengths". First, we consider the case of PCA, i.e., RP-PCA without overweights on the means ($\omega = -1$). The first eigenvalue of Σ_{RP} (3.22) is symptomatic of a systematic, strong factor, given that it is large and substantially higher than the rest of the estimated eigenvalues. The second eigenvalue is substantially lower than the first (0.18), indicating that the associated factor is "less strong", being relevant for fewer assets. However, the second factor also stands out as its eigenvalue is well separated from the subsequent eigenvalues. In fact, the remaining eigenvalues are smaller in magnitude and roughly of comparable levels with each other. The same pattern in eigenvalues is evident when implementing RP-PCA with $\omega = 0$ (not reported), i.e., PCA applied to the correlation matrix. In contrast, for models with $\omega > 0$, the third factor becomes more clearly separated from the remaining factors. With $\omega = 10$, the value of the third eigenvalue (0.14) is twice as large as that of the fourth eigenvalue (0.07).

At the same time, higher values of the RP-weight also enhance the "signal strengths" of the first two factors.[39] These patterns are more evident for higher values of the RP-weight; all of this can be easily visualized in Figure 2. Relatedly, the sum of the first two eigenvalues extracted using RP-PCA (with $\omega > 5$) exceeds the sum of the first six eigenvalues estimated via PCA. In short, using RP-PCA, the factor "signal strengths" are enhanced. In this way, relevant factors that are not detected by PCA are instead identified by RP-PCA, as is the case of the third factor. To corroborate this interpretation, we apply the statistical tests of Onatski (2010) and Giglio and Xiu (2021) – O and GX tests, respectively – to the same Σ_{RP} matrix. Consistently, both tests detect two factors using PCA, and three factors using RP-PCA (see legend of Figure 2). Thus, we can conclude that the third factor is weak, but presumably with high SR, being therefore relevant for pricing the cross section of test assets.

To shed light on this conjecture, we assess what drives a factor's overall "signal strength", by simply inspecting its composition. Thus, we try to establish whether a factor (i) is strong or weak, and (ii) with high or low Sharpe ratio. To do this, we contrast the eigenvalues of $\Sigma^F_{RP} = \psi(\Sigma_F + (1+\omega)\mu_F^\top \mu_F)\psi^\top$ with those of $\Sigma^F_{PCA} = \psi \Sigma_F \psi^\top$.[40] The two types of eigenvalues capture, respectively, the *overall* and *time-series* "strength" of the associated factor. The time-series "strength" is informative about whether a factor is weak or strong. Instead, a comparison of the two types of eigenvalues is revealing about the factor pricing relevance, and hence about whether it is a low or high SR factor. Weak factors are factors that explain a small set of test assets, so that these factors should have low variance, i.e., low time-series "signal strength". Meanwhile, if the overall strength exceeds the time-series "strength", the factor is likely to also have high average returns (i.e., risk premia). Thus, intuitively, factors with low time-series strength and high overall strength denote weak factors with high Sharpe ratios.

Table A8 shows that factors' time-series strengths are unchanged as the RP-weight increases. Hence, for

[39]For this reason, the absolute difference between $\hat{F}_2$ and $\hat{F}_3$ is essentially the same for ω = -1 and 10 (i.e., 0.18). But, for higher RP-weights, $\hat{F}_2$ and $\hat{F}_3$ become increasingly separated (e.g., the difference is 0.32 for $\omega = 20$).

[40]Recall from Section 2.1.1 that Σ_{RP} should converge to Σ^F_{RP}.

RP-PCA with $\omega > 0$, the improvement in the factors' overall strength is entirely due to the risk-premium component. For higher values of the RP-weight, relevant factors are better separated from the remaining factors, but also between themselves (see Panel A.III Difference). Recall that, in RP-PCA, factor orderings in terms of variances and "signal strengths" may not coincide, exactly because of the effect of the means. However, for $\omega = 20$, we find that factors ordered first tend to have not only higher variances but also higher means. This is not the case using RP-PCA with low weights (e.g., $\omega = 5$). Moreover, for $\omega = 20$, leaving aside the first factor (which is strong but with relatively low Sharpe ratio), weaker factors also display descending Sharpe ratios.[41] Of particular interest are the second and third factors which, as we suspected, are weak factors with high Sharpe ratios (0.34 and 0.28, respectively; see Panel B, Table A8). The fourth factor is also weak with a positive and significant risk premium, and hence potentially represents an additional pricing factor, but its Sharpe ratio is half that of the third factor. The remaining factors have zero risk premia, and hence are time-series factors.

Thus far, we documented that RP-PCA with reasonably high RP-weights increases factor "signal strengths" and aggregates the information in a small number of factors. In doing so, it helps us estimate factors more efficiently, as documented also by LP for equities. Moreover, RP-PCA detects weak factors with high Sharpe ratios, which are missed by the standard PCA (the third latent factor is a clear example in this regard). These factors are particularly hard to identify, and yet have important asset pricing implications, exactly because of their large risk premia. In our context, their omission is likely to distort the candidate factor risk-premium estimates.

`Optimal Currency SDF` ($\varphi(F_K^\omega)$). Table 1 in the main text evaluates the trade-off for different RP-weights (ω = -1, 10, 20, 50). We consider SDFs of increasing dimension, including up to six latent factors associated with the six largest eigenvalues of matrix Σ_{RP}. In doing so, we also obtain useful clear indications on the optimal factor SDF.

To start with, we note that average idiosyncratic variance, $\overline{\sigma}_\epsilon$, increases with the RP-weight. The increase is, however, negligible for the optimal three-factor SDF, and is somewhat more pronounced for the two-factor SDF. In fact, as ω varies from -1 to 50, $\overline{\sigma}_\epsilon$ increases from 17.07 percent to 17.40 using $\varphi(F_{1-3})$, and from 19.04 percent to 20.36 using $\varphi(F_{1-2})$. As expected, $\overline{RMS}_\alpha$ moves inversely with the RP-weight; what is instead striking are the large economic gains in terms of pricing accuracy. For example, based on $\varphi(F_{1-3})$, $\overline{RMS}_\alpha$ is around 1.30 with $\omega = -1$, while it drops to 0.80 with $\omega = 50$. The cross-sectional R^2 and MAE reveal a similar pattern of the model's pricing performance. For example, the $\varphi(F_{1-3})$'s R^2 increases from 45 percent with $\omega = -1$ to slightly less than 100 percent with $\omega = 50$. This evidence, taken together, shows that in practice there is no trade-off in selecting high RP-weights in our data, which is consistent with the earlier evidence of Table A8, showing stable time-series signal strengths across RP-weights.

Note that $\overline{RMS}_\alpha$ is lower using $\varphi(F_{1-3})$ with $\omega \geq 20$ than using $\varphi(F_{1-6})$ with $\omega = -1$. Based on the R^2 and MAE criteria, we find similar evidence also for lower but positive values of the RP-weight. Thus, the RP-PCA method, implemented with a reasonably high RP-weight, achieves lower pricing errors than PCA also for more parsimonious SDFs. Moreover, it is also apparent that, based on $\varphi(F_{1-3})$, the marginal gains in terms of pricing performance obtained by using large RP-weights are small. In essence, pricing-error statistics tend to stabilize for $\omega \geq 20$, and we do not see additional benefits in using RP-weights higher than 20. The pricing contribution of $\hat{F}_4$ is, however, more evident for intermediate RP-weights than high

[41] Note that the results documented using $\omega = 20$ are basically unchanged for $\omega = 30, 40, 50$.

weights. Hence, the optimal FX pricing kernel consists of (at least) the three RP-PCA factors, extracted using $\omega = 20$.

To complete the analysis, we inspect the maximal SRs (Panel B, Table 1). To start with, we compare the SRs implied by the optimal SDFs of RP-PCA and PCA. We find that the SR of $\varphi(F_{1-3})$ with $\omega = 20$ is roughly three times higher than the SR of $\varphi(F_{1-2})$ with $\omega = -1$ (0.45 vs. 0.14). However, even for SDFs of equal dimension, the SR of RP-PCA is substantially higher than that of PCA (0.45 vs. 0.26). Such a wedge is almost equally due to $\hat{F}_2$ and $\hat{F}_3$, as both factors' means and SDF-weights ($\hat{b}_{MV}$) increase, albeit to different extents, when moving from $\omega = -1$ to $\omega = 20$. While $\hat{F}_2$'s SDF-weights and means increase monotonically with the RP-weight, those of $\hat{F}_3$ display an hump-shaped pattern. Also note that, for $\omega \geq 20$, the SRs of the three-factor SDFs display only marginal increases. Moreover, by adding $\hat{F}_4$ to the SDF, the SR further increases to 0.48 using RP-PCA with $\omega = 20$, while it is unchanged using PCA. The SDF-weights of $\hat{F}_4$ display a qualitatively similar pattern to those of $\hat{F}_3$, but $\hat{F}_4$'s contribution to the maximal SR is substantially smaller. Using four-factor SDFs, we appreciate no significance difference in the SRs of the SDFs with $\omega = 20$ and $\omega = 50$.

Overall, we can conclude that the optimal latent-factor currency SDF should include at least the first three factors and that an RP-weight of 20 seems a plausible choice.

Table A8: Latent Factor Signal Strengths

The table presents latent factors' signal strengths and statistics, whereby factors are estimated via RP-PCA using different RP-weights. Specifically, the factors are obtained from regressing the portfolio returns on the factor loadings. The factor loadings are proportional to the eigenvectors associated with the $K = 6$ largest eigenvalues of the matrix $\Sigma_{RP} = \frac{1}{T} X^\top X + \omega \bar{X}^\top \bar{X}$, where X is the $T \times N$ matrix of currency portfolio excess returns, and ω is the RP-weight (see Section 2.1.1). In Panel A: Eigenvalues, A.I Plain, we present the K largest eigenvalues (ν_K) associated with matrices, $\Sigma^F_{PCA} = \psi \Sigma_F \psi^\top$ and $\Sigma^F_{RP} = \psi(\Sigma_F + (1+\omega)\mu_F^\top \mu_F)\psi^\top$, where ψ and Σ_F are the loadings (or betas) and the annualized K-factor variance matrix, respectively; we also report the K largest eigenvalues of the annualized Σ_{RP} matrix. In A.II Normalized, eigenvalues are normalized by a constant ($\tilde{\nu}_k(\omega) = \nu_k(\omega)/\bar{\sigma}^2_\epsilon(\omega)$), and hence more directly relate to the factor signal-to-noise ratios; specifically, $\bar{\sigma}^2_\epsilon = \frac{1}{N}\sum_{n=1}^N \sigma^2_{\epsilon,n}$, where $\sigma^2_{\epsilon,n}$ is the annualized variance of the n-th portfolio's residual (i.e., the idiosyncratic variance), obtained by estimating $X_{nt} = \alpha_n + \hat{F}_t \psi_n^\top + \epsilon_{nt}$, for $t = 1..., T$, where $\hat{F}_t$ collects the six latent factors. In A.III Difference, we present the differences of consecutive normalized eigenvalues. In Panel B: Factors, B.I Statistics, we report orthogonalized factor Sharpe ratios (SR), the rank of the factor means (Rnk), and the annualized factor means (μ_F), which are starred with ***,**,* denoting significance at the 1-, 5- and 10-percent levels, respectively, based on Newey-West standard errors with optimal lag-length selection. We carry out the analysis using RP-PCA with selected RP-weights. RP-PCA with $\omega = -1$, and 0 correspond to standard PCA applied to the covariance and correlation matrices, respectively. With $\omega > 0$, RP-PCA "overweights" the factor means. The test assets consist of the currency portfolios from the nine investment strategies ($N = 46$), for the period 11/1983-12/2017 at monthly frequency ($T = 410$).

	Panel A: Eigenvalues										Panel B: Factors		
	A.I Plain			A.II Normalized			A.III Difference				B.I Statistics		
$\omega = -1$	Σ^F_{PCA}	Σ^F_{RP}	Σ_{RP}	Σ^F_{PCA}	Σ^F_{RP}	Σ_{RP}	Σ^F_{PCA}	Σ^F_{RP}	Σ_{RP}		SR	Rnk	μ_F
ν_1	3.26	3.26	3.22	5.86	5.86	5.79	5.53	5.53	5.46	F_1	0.10	1	0.17*
ν_2	0.19	0.19	0.18	0.33	0.33	0.33	0.18	0.18	0.18	F_2	0.11	4	0.05**
ν_3	0.08	0.08	0.08	0.15	0.15	0.15	0.03	0.03	0.03	F_3	0.22	3	0.06***
ν_4	0.07	0.07	0.07	0.12	0.12	0.12	0.02	0.02	0.02	F_4	0.01	6	0.00
ν_5	0.06	0.06	0.06	0.10	0.10	0.10	0.01	0.01	0.01	F_5	0.26	2	0.06***
ν_6	0.05	0.05	0.05	0.09	0.09	0.09	–	–	–	F_6	0.03	5	0.01
$\omega = 5$	Σ^F_{PCA}	Σ^F_{RP}	Σ_{RP}	Σ^F_{PCA}	Σ^F_{RP}	Σ_{RP}	Σ^F_{PCA}	Σ^F_{RP}	Σ_{RP}		SR	Rnk	μ_F
ν_1	3.26	3.44	3.40	5.84	6.18	6.10	5.51	5.80	5.73	F_1	0.10	1	0.18*
ν_2	0.19	0.21	0.21	0.33	0.37	0.37	0.18	0.16	0.16	F_2	0.19	3	0.08***
ν_3	0.08	0.12	0.12	0.15	0.22	0.22	0.03	0.09	0.09	F_3	0.35	2	0.10***
ν_4	0.07	0.07	0.07	0.12	0.12	0.12	0.03	0.00	0.00	F_4	0.11	5	0.03**
ν_5	0.05	0.07	0.07	0.10	0.12	0.12	0.00	0.03	0.03	F_5	0.15	4	0.04***
ν_6	0.05	0.05	0.05	0.09	0.09	0.09	–	–	–	F_6	0.01	6	0.00
$\omega = 10$	Σ^F_{PCA}	Σ^F_{RP}	Σ_{RP}	Σ^F_{PCA}	Σ^F_{RP}	Σ_{RP}	Σ^F_{PCA}	Σ^F_{RP}	Σ_{RP}		SR	Rnk	μ_F
ν_1	3.26	3.60	3.56	5.83	6.46	6.38	5.50	6.02	5.94	F_1	0.10	1	0.18*
ν_2	0.19	0.24	0.24	0.33	0.43	0.44	0.19	0.18	0.18	F_2	0.26	2	0.10***
ν_3	0.08	0.14	0.14	0.15	0.25	0.25	0.02	0.13	0.13	F_3	0.33	3	0.10***
ν_4	0.07	0.07	0.07	0.12	0.12	0.13	0.03	0.00	0.01	F_4	0.16	4	0.04***
ν_5	0.05	0.07	0.07	0.09	0.12	0.12	0.00	0.03	0.03	F_5	0.03	5	0.01
ν_6	0.05	0.05	0.05	0.09	0.09	0.09	–	–	–	F_6	0.01	6	0.00
$\omega = 20$	Σ^F_{PCA}	Σ^F_{RP}	Σ_{RP}	Σ^F_{PCA}	Σ^F_{RP}	Σ_{RP}	Σ^F_{PCA}	Σ^F_{RP}	Σ_{RP}		SR	Rnk	μ_F
ν_1	3.26	3.94	3.89	5.83	7.05	6.97	5.50	6.47	6.37	F_1	0.10	1	0.19**
ν_2	0.19	0.33	0.33	0.33	0.59	0.59	0.19	0.31	0.32	F_2	0.34	2	0.11***
ν_3	0.08	0.15	0.15	0.15	0.28	0.28	0.02	0.15	0.15	F_3	0.28	3	0.09***
ν_4	0.07	0.07	0.07	0.12	0.13	0.13	0.03	0.00	0.01	F_4	0.15	4	0.04***
ν_5	0.05	0.07	0.07	0.09	0.12	0.12	0.00	0.03	0.03	F_5	0.02	5	0.01
ν_6	0.05	0.05	0.05	0.09	0.09	0.09	–	–	–	F_6	0.01	6	0.00

Figure A1: Portfolio Risk Exposures to Latent Factors

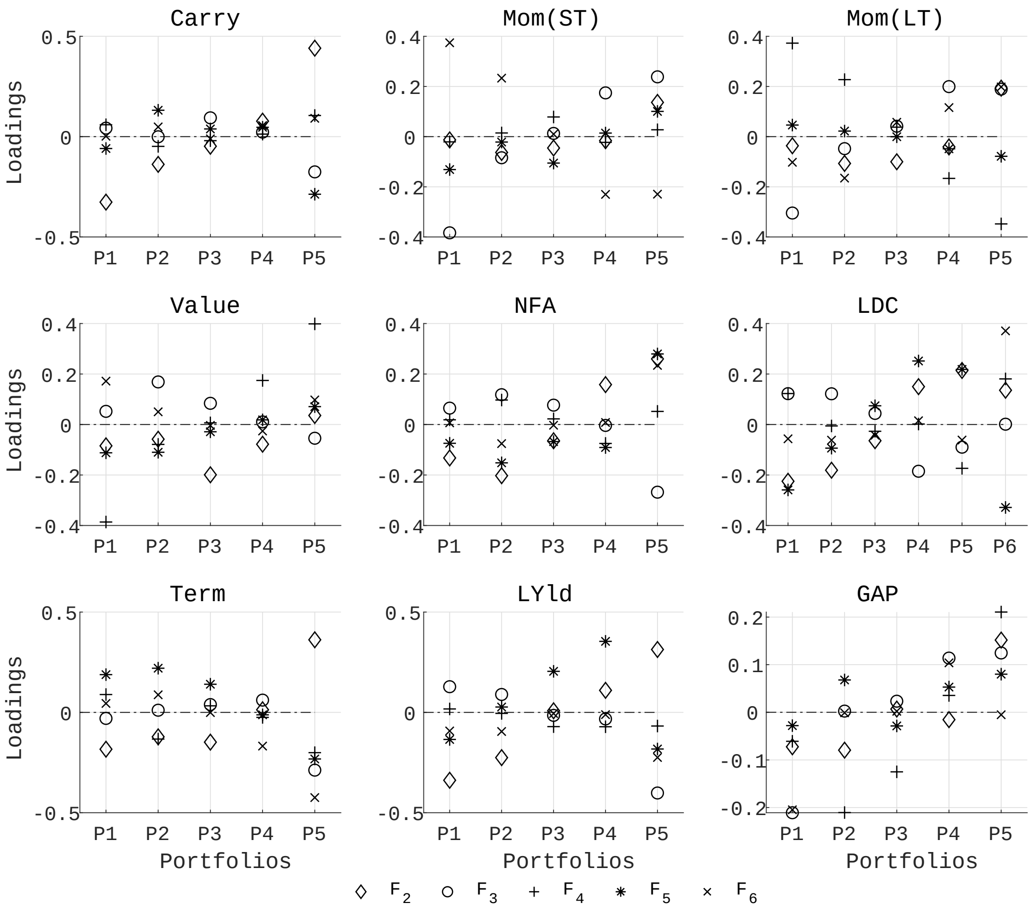

The figure shows portfolio loadings on the six estimated latent factors from the panel of currency portfolios excess returns by means of RP-PCA with baseline RP-weight, i.e., $\omega = 20$. The RP-PCA loadings are also given by the regression coefficients using transformed data, that is, by incorporating the cross-sectional error (Lettau and Pelger, 2020a,b). Specifically, define $\tilde{\omega}_{nt} = \sqrt{\omega + 1} - 1$, $\tilde{X}_{nt} = X_{nt} + \tilde{\omega}\bar{X}_{nt}$ and $\tilde{F}_{kt} = \hat{F}_{kt} + \tilde{\omega}\bar{F}_{kt}$. Then, for any value of ω, RP-PCA loadings are given by the coefficients $(\tilde{\psi}_n)$ from the OLS time-series regressions, $\tilde{X}_{nt} = \tilde{F}_t\tilde{\psi}_n^{\top} + \tilde{e}_{nt}$, $n = 1, \ldots, N$, $t = 1, \ldots, T$. The test assets' sample consists of the portfolios associated with the nine investment strategies ($N = 46$). The sample spans the 11/1983-12/2017 period at monthly frequency ($T = 410$).

Figure A2: Portfolio Explained Variations by Latent Factors

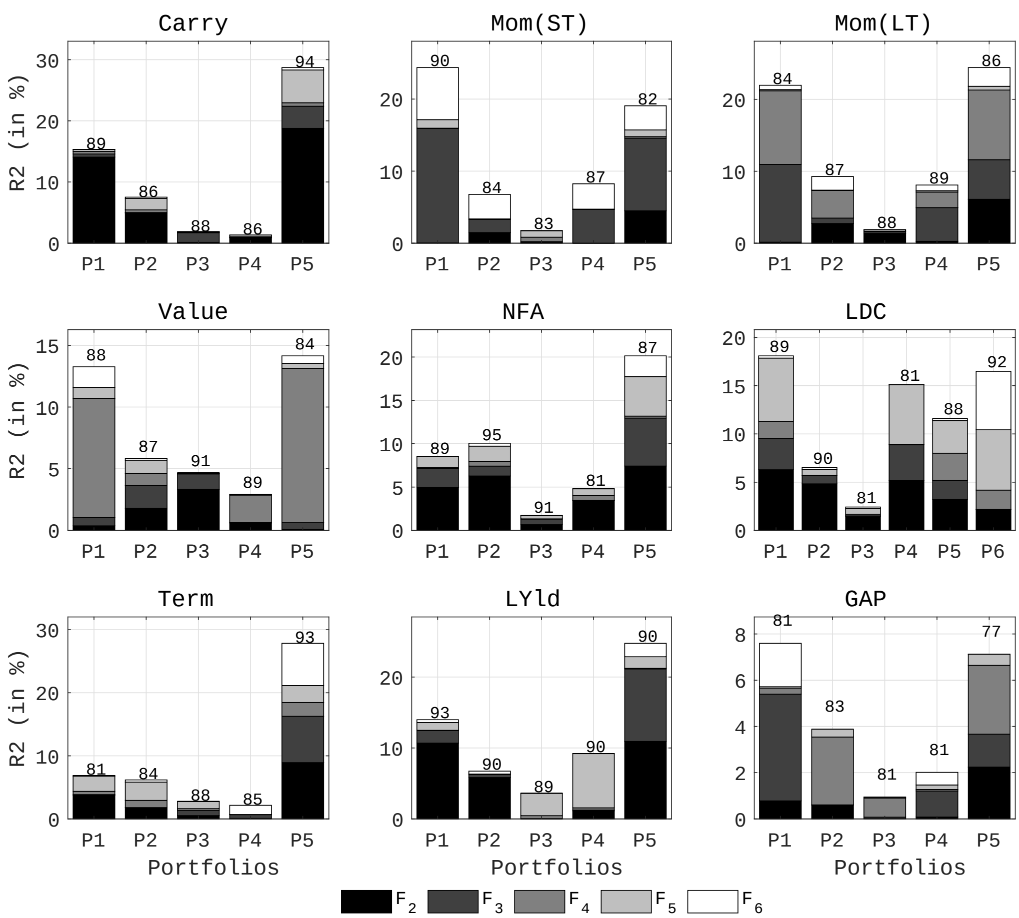

The figure shows the R^2s obtained by regressing currency portfolios' excess returns (X_{nt}) on the estimated orthogonalized latent factors ($\hat{F}_t = [\hat{F}_{1t}, \ldots, \hat{F}_{Kt}]$, with $K = 6$). Factors are estimated from the panel of test asset excess returns by means of RP-PCA with baseline RP-weight (i.e., $\omega = 20$). The estimated factors are then orthogonalized, to facilitate their economic interpretation. We estimate $N \times K$ OLS time-series regressions, $X_{nt} = \alpha_n + \hat{F}_{1:kt}\psi_n^\top + e_{nt}$, $n = 1, \ldots, N$, $t = 1, \ldots, T$, $k = 1, \ldots, K$, with $\hat{F}_{1:kt} = [\hat{F}_{1t}, \ldots, \hat{F}_{kt}]$. In this way, for each portfolio, we show factors' contributions to the overall R_n^2, omitting that of $\hat{F}_{1t}$, to better visualize the evidence of the remaining factors. We report the overall R_n^2s, which include all K factors, above the bars. The test assets' sample consists of the portfolio excess returns associated with the nine investment strategies ($N = 46$). The sample spans the 11/1983–12/2017 period at monthly frequency ($T = 410$).

Figure A3: Nontradable Factor's Explained Variations by the Six Latent Factors

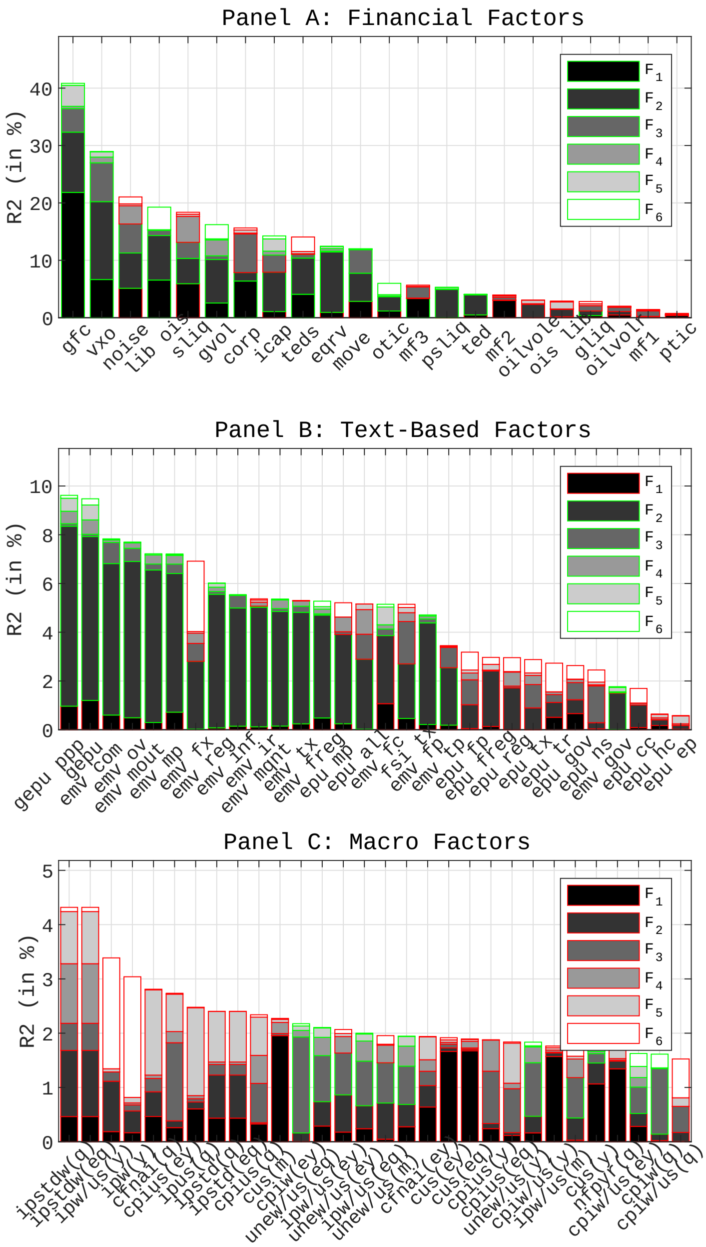

The figure shows nontradable candidate factor explained variations by the six latent factors. Factors are extracted by means of RP-PCA with baseline RP-weight (i.e., $\omega = 20$) applied to the panel of $N = 46$ currency portfolios associated with the nine investment strategies. Panels A, B, and C present the evidence for financial, text-based and macro candidate factors, respectively. In each panel, factors are sorted by the R^2s associated with the six-factor model (to help visualize the results, we only display the 30 macro factors with the highest R^2s). Green (red) K-th bar edges denote (not) significant risk premia, at the 10 percent level, based on the model including the selected number of latent factors. The test assets' sample consists of the portfolios associated with the nine investment strategies ($N = 46$). Factors are expressed as innovations, using the residuals from AR(1) processes, and are then standardized. The sample period varies with the factor at hand, according to data availability over the 11/1983-12/2017 period ($T = 410$). See factor descriptions in Tables A5-A7.

Table A9: Exposures of All Nontradable Factors to the Latent Factors

The table presents the nontradable candidate risk factors' exposures to the latent factors (η_{F_k}) and the explained variations ($R^2_{F_{1-k}}$) obtained from the spanning regression of Eq. (15), for models including an increasing number of factors, $k = 1, \ldots, K$. We report the nontradable factor exposures to the first six extracted, orthogonalized latent factors (i.e., $K = 6$). The factors are extracted by applying RP-PCA with baseline weight (i.e., $\omega = 20$) to the $N = 46$ portfolios obtained from the nine investment strategies. Panels A, B, and C show the estimates for the financial, text-based, and macro factors, respectively. To contain the space, we do not report macro factors measured using exponential moving averages. ***,**,* denote significance, respectively, at the 1-, 5- and 10-percent levels, based on Newey-West standard errors.

	PANEL A: Financial Factors											
	A.I: Risk Exposures						A.II: Explained Variation					
	η_{F_1}	η_{F_2}	η_{F_3}	η_{F_4}	η_{F_5}	η_{F_6}	$R^2_{F_1}$	$R^2_{F_{1-2}}$	$R^2_{F_{1-3}}$	$R^2_{F_{1-4}}$	$R^2_{F_{1-5}}$	$R^2_{F_{1-6}}$
otic	0.73**	5.73***	-0.35	-1.42	2.26	7.47**	1.17	3.67	3.68	3.77	4.01	5.99
otic2	0.65*	4.87***	-0.42	-2.10	2.80	8.78***	0.95	2.76	2.77	2.97	3.34	6.07
ptic	-0.44	-0.17	-1.16	-0.73	1.31	-1.81	0.43	0.43	0.53	0.55	0.63	0.75
ptic2	-0.37	-1.39	-1.15	-1.23	0.83	0.19	0.31	0.46	0.56	0.62	0.66	0.66
icap	0.69*	9.46***	-6.38***	3.89	6.72***	3.80	1.06	7.93	10.91	11.60	13.73	14.25
mf1	0.19	-1.42	3.58**	1.25	-1.81	0.63	0.08	0.23	1.17	1.24	1.40	1.41
mf2	-1.16**	-0.82	2.54	-2.38	1.51	-1.26	3.02	3.07	3.54	3.80	3.90	3.96
mf3	1.22**	1.07	-5.16**	0.25	1.35	2.45	3.34	3.42	5.37	5.37	5.46	5.67
noise	-1.36*	-11.34***	11.34**	8.06*	-0.62	-0.41	4.99	10.98	17.33	20.05	20.07	20.07
sliq	-1.29**	-10.48***	5.28	7.79*	-0.90	-6.67*	5.81	11.47	13.40	16.21	16.23	17.60
oilvole	-0.09	-5.50	0.97	-0.93	3.63	0.42	0.02	2.33	2.40	2.44	3.06	3.07
oilvolr	0.47	-2.78**	3.20	-1.35	-0.49	1.65	0.48	1.07	1.82	1.90	1.92	2.01
gfc	3.13***	11.71***	-7.49***	-2.96	8.80***	3.13	21.83	32.33	36.43	36.83	40.49	40.84
gliq	-0.41	-3.23**	3.57**	0.13	-2.60	-3.37	0.37	1.17	2.10	2.10	2.42	2.83
gvol	-1.07**	-9.95***	2.99	7.80**	-2.07	-8.34***	2.55	10.13	10.79	13.54	13.74	16.21
psliq	0.03	8.00**	-0.85	-0.88	2.42	-0.84	0.00	4.91	4.96	5.00	5.27	5.30
corp	-1.82**	-5.16**	11.02**	2.43	1.67	0.95	6.35	7.60	14.97	15.22	15.32	15.35
ted	-0.45	-6.82*	-0.44	0.38	0.72	0.27	0.47	3.94	3.95	3.96	3.98	3.98
lib ois	-1.48*	-14.14*	3.61	1.30	6.90	4.52	6.39	16.00	16.97	17.09	18.03	18.35
ois lib	-0.31	-7.43	2.17	-0.33	4.81	13.68**	0.15	1.93	2.10	2.11	2.65	5.63
move	-0.98**	-10.89***	7.81***	0.12	-0.25	-1.94	2.78	8.33	11.75	11.75	11.75	11.86
vxo	-1.60***	-15.19***	10.32***	1.85	-6.91***	-1.35	6.47	19.96	26.40	26.55	28.50	28.56
eqrv	-0.65	-11.72**	1.46	3.11	-2.83	0.80	0.93	11.45	11.60	12.04	12.42	12.44
	PANEL B: Text-Based Factors											
	B.I: Risk Exposures						B.II: Explained Variation					
	η_{F_1}	η_{F_2}	η_{F_3}	η_{F_4}	η_{F_5}	η_{F_6}	$R^2_{F_1}$	$R^2_{F_{1-2}}$	$R^2_{F_{1-3}}$	$R^2_{F_{1-4}}$	$R^2_{F_{1-5}}$	$R^2_{F_{1-6}}$
gepu	0.00	-11.18***	3.38	3.12	-3.68	-3.52	1.12	7.36	7.64	7.99	8.26	8.56
gepu ppp	0.13	-12.01***	3.82	2.81	-3.79	-2.97	0.89	7.82	8.17	8.45	8.75	8.97
fsi tx	-0.45	-5.40*	4.66**	2.29	-2.78	1.93	0.46	2.77	4.42	4.67	5.04	5.17
epu all	0.12	-6.09**	2.74	4.06**	-4.95*	-0.46	0.02	2.93	3.62	4.38	5.42	5.43
epu mp	0.35	-7.05**	0.94	3.62*	-1.25	3.25	0.25	3.95	4.02	4.62	4.70	5.08
epu fp	0.14	-3.61	3.16	2.23	-3.25	3.74	0.05	1.01	1.78	2.01	2.51	3.01
epu tx	0.04	-3.43	3.09	2.68	-3.09	3.32	0.00	0.87	1.60	1.94	2.39	2.78
epu gov	0.54	-2.83	2.58	1.55	-2.10	2.71	0.66	1.24	1.74	1.85	2.07	2.33
epu hc	-0.26	-1.69	-0.85	0.70	-2.32	0.62	0.17	0.42	0.45	0.48	0.71	0.73
epu ns	0.12	-1.64	3.80**	0.62	-3.84	-3.99*	0.03	0.26	1.45	1.47	2.04	2.61
epu ep	0.12	-1.32	-0.83	0.42	-3.41	0.26	0.02	0.19	0.21	0.22	0.73	0.73
epu reg	0.03	-4.79*	0.67	3.38	-1.02	-3.78	0.00	1.78	1.83	2.34	2.37	2.87
epu freg	0.24	-5.51	0.87	0.59	1.40	-3.05	0.14	2.40	2.44	2.46	2.56	2.89
epu tr	-0.44	-2.80	-2.38	1.22	-0.33	-5.09**	0.50	1.18	1.53	1.59	1.59	2.51
epu cc	0.22	-3.46**	-0.35	0.88	-1.30	-3.33	0.09	1.07	1.07	1.11	1.16	1.56
emv ov	-0.49	-9.13**	2.88	1.87	-1.51	0.31	0.48	6.73	7.35	7.51	7.61	7.61
emv mout	-0.38	-9.09**	2.19	2.74	-0.31	1.32	0.29	6.40	6.73	7.07	7.07	7.14
emv mqnt	-0.26	-7.85**	1.35	2.55	-1.12	1.09	0.14	4.78	4.92	5.21	5.27	5.31
emv inf	-0.28	-7.93**	2.92	0.09	-0.98	0.29	0.14	4.85	5.45	5.45	5.50	5.50

emv com	-0.55	-8.96**	3.64	1.21	-1.67	0.39	0.59	6.60	7.56	7.63	7.75	7.76
emv ir	-0.23	-7.97**	0.89	1.32	-1.56	1.60	0.11	4.96	5.03	5.11	5.22	5.31
emv fc	-0.72	-6.25*	2.52	2.78	2.96	-2.17	1.06	3.80	4.16	4.50	4.90	5.06
emv fx	-0.13	-5.90**	3.25**	-3.04	0.00	-8.69***	0.02	2.73	3.50	3.92	3.93	6.60
emv fp	-0.32	-7.32**	1.52	1.15	-1.58	0.34	0.21	4.29	4.48	4.54	4.65	4.65
emv tx	-0.34	-7.72**	2.09	2.10	-0.68	0.61	0.23	4.68	4.99	5.19	5.21	5.22
emv gov	0.06	-4.32*	-0.68	-1.57	-1.69	1.33	0.01	1.51	1.53	1.64	1.77	1.83
emv mp	-0.59	-8.69***	2.80	2.88	0.22	1.18	0.70	6.22	6.74	7.11	7.11	7.16
emv reg	-0.20	-8.56**	1.84	2.19	1.49	0.19	0.07	5.42	5.63	5.84	5.93	5.93
emv freg	-0.49	-7.54*	1.32	2.18	1.18	2.37	0.47	4.60	4.70	4.91	4.96	5.16
emv tp	-0.34	-5.56**	3.79*	-0.81	0.25	-0.77	0.18	2.42	3.38	3.41	3.42	3.44
PANEL C: Macro Factors												
	C.I: Risk Exposures						C.II: Explained Variation					
	η_{F_1}	η_{F_2}	η_{F_3}	η_{F_4}	η_{F_5}	η_{F_6}	$R^2_{F_1}$	$R^2_{F_{1-2}}$	$R^2_{F_{1-3}}$	$R^2_{F_{1-4}}$	$R^2_{F_{1-5}}$	$R^2_{F_{1-6}}$
ipus(m)	-0.33	1.98	0.70	-1.82	1.44	0.63	0.24	0.54	0.58	0.73	0.82	0.84
ipus(q)	0.52	1.29	-0.91	-1.13	5.86*	-0.51	0.60	0.73	0.79	0.85	2.47	2.48
ipus(y)	0.22	2.27	0.14	-0.18	-0.88	-1.52	0.11	0.51	0.51	0.51	0.55	0.63
cpius(m)	0.18	0.56	-2.67	-1.23	3.80	0.91	0.08	0.10	0.62	0.69	1.37	1.40
cpius(q)	0.38	0.55	-3.12*	-3.35	3.84	-1.10	0.33	0.35	1.07	1.59	2.30	2.34
cpius(y)	0.33	1.10	-3.60*	-3.53	0.24	-0.23	0.24	0.34	1.30	1.87	1.88	1.88
nfpyr(m)	0.25	-0.10	-0.51	1.01	-1.24	0.76	0.16	0.16	0.18	0.23	0.30	0.33
nfpyr(q)	0.78**	1.36	-0.45	0.83	1.84	0.57	1.34	1.49	1.50	1.53	1.69	1.70
nfpyr(y)	0.60	0.57	-0.85	-0.13	-1.72	-0.07	0.83	0.86	0.91	0.91	1.05	1.05
cfnai(m)	-0.07	0.87	-2.30	-1.90	1.52	1.10	0.01	0.07	0.46	0.62	0.73	0.77
cfnai(q)	0.46	2.44	-1.83	-1.18	5.78*	0.56	0.46	0.92	1.17	1.23	2.80	2.81
cfnai(y)	0.53	1.08	-1.73	-1.36	-1.40	0.07	0.64	0.73	0.95	1.04	1.13	1.13
uneus(m)	-0.09	0.08	-2.59	3.24	-0.90	-0.79	0.02	0.02	0.51	0.98	1.02	1.04
uneus(q)	-0.09	-1.00	0.59	1.38	-3.17	0.67	0.02	0.10	0.12	0.21	0.68	0.70
uneus(y)	-0.01	0.58	-4.13**	1.54	0.84	0.30	0.00	0.03	1.28	1.39	1.42	1.43
cus(m)	-0.94**	-0.23	-0.69	-2.12	1.11	-0.68	1.95	1.96	1.99	2.20	2.26	2.27
cus(q)	-0.31	-0.38	0.77	-0.82	0.78	-0.45	0.22	0.23	0.27	0.30	0.33	0.34
cus(y)	-0.69	-2.24	-1.54	-0.92	-0.08	1.49	1.06	1.45	1.62	1.66	1.66	1.74
ipw(m)	0.22	-1.28	-2.87	1.12	0.18	2.29	0.11	0.23	0.84	0.89	0.90	1.08
ipw(q)	0.23	-2.60	-1.80	0.53	0.95	1.87	0.12	0.63	0.87	0.88	0.92	1.05
ipw(y)	0.27	-2.29	-1.24	-0.90	-1.46	7.92***	0.17	0.57	0.68	0.72	0.82	3.04
ipw/us(m)	0.11	-2.31	-3.19	2.76	-1.05	2.16	0.03	0.44	1.18	1.52	1.58	1.74
ipw/us(q)	-0.15	-3.69**	-1.40	1.39	-0.88	1.99	0.05	1.09	1.24	1.32	1.36	1.50
ipw/us(y)	0.29	-3.47	-1.52	-0.24	-1.10	7.56**	0.19	1.11	1.28	1.29	1.34	3.39
cpiw(m)	0.34	-0.25	-2.41	-1.01	-0.26	-1.67	0.25	0.26	0.68	0.72	0.73	0.83
cpiw(q)	0.10	-1.25	-4.07***	0.51	-0.17	-2.67	0.02	0.14	1.35	1.36	1.36	1.61
cpiw(y)	-0.38	-1.23	-1.26	-2.58	1.27	2.06	0.33	0.44	0.56	0.86	0.94	1.09
cpiw/us(m)	0.08	-0.34	-0.25	-1.86	0.44	-2.69	0.01	0.02	0.03	0.18	0.19	0.45
cpiw/us(q)	-0.05	-1.47	-2.56	-0.05	1.83	-4.49*	0.01	0.17	0.65	0.65	0.81	1.53
cpiw/us(y)	-0.84***	-0.94	0.42	-0.92	0.98	0.94	1.57	1.64	1.65	1.69	1.73	1.76
ipstd(m)	0.24	1.06	-0.91	0.80	-1.91	-0.20	0.13	0.22	0.28	0.31	0.48	0.48
ipstd(q)	0.44	-3.19	1.67	1.00	-4.43**	-0.32	0.43	1.22	1.42	1.47	2.40	2.40
ipstd(y)	0.50	-2.33	0.09	2.97	0.84	1.11	0.51	0.91	0.91	1.30	1.33	1.37
ipstdw(m)	0.03	-1.00	-0.16	2.44	-1.99	-0.13	0.00	0.08	0.08	0.35	0.54	0.54
ipstdw(q)	0.46	-3.96*	2.63	2.13	-6.24***	1.34	0.46	1.68	2.19	2.40	4.25	4.31
ipstdw(y)	0.14	-0.60	-1.93	0.34	-2.75	2.35	0.03	0.07	0.28	0.28	0.63	0.83
cpistd(m)	-0.12	0.03	-0.49	2.23	2.09	-0.92	0.03	0.03	0.05	0.28	0.48	0.51
cpistd(q)	-0.24	1.02	0.09	2.92*	3.39**	-0.33	0.13	0.21	0.21	0.59	1.14	1.14
cpistd(y)	-0.22	1.71	3.11	0.77	2.18	1.89	0.08	0.34	0.96	0.98	1.19	1.32
cpistdw(m)	0.21	0.79	-1.57	0.06	-0.49	-1.98	0.10	0.15	0.33	0.33	0.34	0.48
cpistdw(q)	-0.02	-1.60	-2.59	0.98	-0.21	-2.01	0.00	0.19	0.68	0.72	0.73	0.87
cpistdw(y)	-0.12	-1.30	0.34	-1.05	0.33	0.40	0.03	0.15	0.16	0.20	0.21	0.21
unew(m)	-0.16	-1.04	1.21	3.18	2.48	1.29	0.06	0.14	0.25	0.70	0.99	1.05
unew(q)	-0.15	0.60	0.52	1.26	-0.44	0.23	0.05	0.08	0.10	0.17	0.18	0.18
unew(y)	-0.27	-1.32	0.79	0.54	1.46	0.45	0.16	0.30	0.34	0.35	0.45	0.46
unew/us(m)	-0.35	-2.31	-3.10*	-2.85	1.93	0.46	0.28	0.69	1.39	1.76	1.94	1.95
unew/us(q)	-0.27	-2.93*	-2.46	-0.95	1.70	-0.99	0.16	0.82	1.26	1.30	1.44	1.47
unew/us(y)	-0.27	-1.99	-3.66**	-2.53	-0.64	1.37	0.16	0.47	1.46	1.75	1.77	1.83

V Robustness and Stability Analysis

In this section, we first present some robustness exercises. We then investigate the stability of the factor structure, also showing the time-varying three-pass risk-premium estimates for selected factors.

V.1 Robustness Exercises

In what follows, we provide a detailed description of three main robustness exercises, which regard i) the dimension of the pricing kernel; ii) the choice of the RP-weight; and iii) the addition of HML portfolios to the sample of test assets. We then list a number of additional robustness checks.

`i) Pricing Kernel Dimension.` Along with the optimal SDF estimates, Table 3 also presents the risk-premium estimates associated with the two- and four-factor SDFs ($\varphi(F_{1-2})$ and $\varphi(F_{1-4})$ panels, respectively). In this way, we can evaluate to what extent the omission of relevant slope pricing factors (i.e., the "Momentum" factor, $\hat{F}_3$), or the inclusion of less relevant ones (i.e., the "Value" factor, $\hat{F}_4$), weigh on the risk-premium estimates and on the associated statistics.

We find that the exclusion of the "Momentum" factor from the SDF materially impacts on the risk-premium estimates. This effect is particularly evident for financial and text-based factors. In fact, for these factors, the risk-premium estimates are substantially higher (in absolute terms) when using $\varphi(F_{1-2})$ instead of $\varphi(F_{1-3})$. This is exactly because, as shown earlier in Table 2, most factors tend to have η-exposures to $\hat{F}_2$ and $\hat{F}_3$ of opposite signs. Thus, by omitting a weak factor with high SR as is $\hat{F}_3$, risk premia appear higher than what they should be. This is for example the case of the corporate credit risk (corp) and intermediaries' capital ratio (icap) factors, whose $\hat{\lambda}_g$s and SRs decrease in absolute value when $\hat{F}_3$ is added to the SDF. Text-based factors seem to follow a similar qualitative pattern. Conversely, macro factor risk-premium estimates tend to increase in absolute terms, because their η-exposures to "Carry" and "Momentum" factors are generally of the same sign.

By adding the "Value" factor to the three-factor SDF, the absolute risk-premium estimates are comparatively more stable than when detracting the "Momentum" factor. This is consistent with the fact that factor η-exposures to "Value" tend to be smaller than to "Carry" and "Momentum", coupled with a lower price of risk of "Value". Some macro factors' risk premia eventually increase, in absolute terms, when using $\varphi(F_{1-4})$, but the effect is economically small.[42]

Taken together, this analysis shows that selecting the optimal SDF is key to obtain precise estimates of nontradable factors' currency risk premia. In particular, the omission of the "Momentum" factor can materially affect the risk-premium estimates in a non-trivial way.

`ii) RP-Weights.` Next, we shed light on the effect of the RP-weights on the candidate factor risk-premium estimates. To start with, Table A10 presents the estimates for the SDF consisting of three latent factors extracted using low ($\omega = -1$, PCA), medium ($\omega = 20$, baseline), and high ($\omega = 50$, high) RP-weights. That is, we contrast the baseline estimates reviewed before with those obtained using two extreme RP-weights.

As for the PCA case, we find little difference in terms of the number of financial and text-based factors with significant risk-premium estimates relative to the baseline. However, for many of these factors, the point estimates deviate substantially. For example, moving from $\omega = 20$ to $\omega = -1$, the risk premium of

[42] Note that the effects of $\hat{F}_5$ and $\hat{F}_6$ on the candidate factors' risk-premium point estimates are essentially nil, exactly because they are time-series factors.

Global EPU (gepu) drops in absolute terms from -1.41 to -0.85, and similarly that of TED from -0.97 to -0.67; these reductions are economically meaningful. The absolute risk-premium estimates of text-based factors are generally lower using PCA than RP-PCA, while no systematic pattern is evident for financial factors. In this sense, macro factors are of particular interest. In fact, essentially none of the factors delivers statistically significant risk premia. Thus, using PCA currency returns appear even more disconnected from macro factors.

However, PCA deviates from RP-PCA also regarding the size of the currency SDF. In fact, the joint analysis of PCA factors' "signal strengths" and trade-offs suggested a two-factor SDF (see Section 4.2). As a result, to fully appreciate the differences in the risk-premium estimates between RP-PCA and PCA models, one should contrast the evidence from the three-factor SDF with $\omega = 20$ (Table 3) with that from the two-factor SDF with $\omega = -1$ (Table A11). In doing so, the differences become even starker. Specifically, using the two-factor SDF with $\omega = -1$, some financial factors' risk premia become significant (e.g., mf2, gliq, and corp), and one is no longer significant (psliq). Also, many risk premia that are statistically significant display very different magnitudes – for example, the risk premium of gvol is reduced by a half using PCA with two factors. Similarly, the risk premia for text-based factors are also very different in magnitude despite displaying significant risk premia (e.g., the estimate of gepu changes from -1.41 to -0.71 and that of emv mp from -0.91 to -0.50), and only the premia of two macro factors remain (marginally) statistically significant.

In contrast, we find that the use of an high RP-weight leads to small differences in the risk-premium estimates. Regardless of the factor types, many factors display essentially unchanged risk premia. This evidence on currency returns lends support to LP's argument that choosing too high penalties is not particularly harmful.

iii) Adding HML Portfolios to the Test Assets. It is ex-ante unclear if the inclusion of HML portfolios will weigh on the evidence uncovered so far. On the one hand, one might conjecture that adding HML portfolios is redundant, because their information should be already spanned by the corner portfolios and hence incorporated in the baseline test assets. On the other hand, the exercise is motivated as follows. First and foremost, the analysis of the R^2s of Figure 3 revealed that some strategies, although they deliver high average returns (e.g., GAP), are not explained much by the extracted factors. Second, LP show that the latent factors extracted using RP-PCA mostly load on the portfolios delivering higher absolute returns and Sharpe ratios. Thus, we add the nine HML portfolios to our sample of test assets (i.e., $N = 55$).

The main findings can be summarized as follows. We still find that there is essentially no trade-off in choosing RP-PCA (with reasonably high RP-weights) than PCA. Factors' "signal strengths" increase with the RP-weights. The key pricing factors retain the same interpretation of "Dollar", "Carry", and "Momentum" factors. However, the optimal SDF now seems to include more clearly also the fourth, "Value" pricing factor (Table A12). But the price of risk and the Sharpe ratio of the "Value" factor are low in comparison with the ones of the other pricing factors. At the same time, only few nontradable factors display significant η-exposures to "Value". As a result, the nontradable factor risk-premium estimates are essentially unchanged using either the three- or four-factor SDFs. Moreover, the nontradable factors' risk-premium estimates are largely consistent with the baseline estimates, so that the method selects a very similar list of relevant factors. This is because, by including HML portfolios, the main difference seems to be that the extracted latent factors better explain the information of the highly profitable trading strategies other than Carry (i.e., ST Mom and GAP). However, this information is mostly spanned by factors that are not relevant for pricing (i.e., $\hat{F}_5$ and $\hat{F}_6$), being excluded from the SDF.

Additional Robustness Analysis. We subject the analysis to a number of additional checks that we briefly illustrate here, before turning to the stability analysis. To start with, we consider the time-series portfolios, i.e., strategies that are either long or short U.S. dollars, of the nine investment strategies considered. These time-series portfolios are used either in place of, or in addition to the HML portfolios, finding little differences in the estimates. In comparison, the inclusion of cross-sectional HML portfolios seems to be relatively more relevant, as the impact of time-series portfolios on the currency SDF is minimal. We verify that the main results are robust to constructing cross-sectional and time-series portfolios using currency ranked-based weights, which are becoming increasingly popular in the literature (e.g., Asness et al., 2013), instead of equal weights. Turning to the candidate factors, we also computed innovations of persistent factors by taking their first differences instead of using AR(1) residuals, documenting no evident changes in their risk-premium estimates. As explained before, we measure macro factor growth rates at multiple frequencies, using either simple or exponential moving averages, and then take their innovations. However, in a robustness exercise, we also consider growth rates instead of their AR(1) innovations and still find a disconnect with currency returns.

Table A10: Risk Premia of Nontradable Factors: RP-weight Robustness

The table presents the risk-premium estimates of selected candidate nontradable risk factors (g_t). Panel FMB presents the risk-premium point estimates (λ_g) and Shanken standard errors (se) from the standard two-pass procedure, including the constant and the candidate factor. The remaining panels report the estimates from the (augmented) three-pass procedure with three-factor SDFs, $\varphi(F_{1-3})$, for $\omega = -1$ (no overweight), $\omega = 20$ (baseline RP-weight), and $\omega = 50$ (high RP-weight) in the left, mid, and right panels, respectively. The risk-premium estimates (λ_g) are reported along with the Newey-West standard errors (se), computed following Giglio and Xiu (2021); ***,**,* denote significance, respectively, at the 1-, 5- and 10-percent levels. As for the three-pass method, for each factor and a given SDF, we also report the spanning R^2s (R2) resulting from projecting the factor onto the k latent factors entering the SDF; the Sharpe ratios (SR) associated with the projected factor, i.e., the return-based counterpart to the original nontradable factor; and the p-value (pval) of the test of GX that factor g_t is weak. In Panels A, B and C, we present financial (FIN), text-based (TXT), and macro (MAC) candidate g factors that have significant risk-premium estimates according to at least one of the RP-weight reported. When a macro factor is significant for multiple frequencies, we present the most representative. Factors are expressed as innovations, using the residuals from AR(1) processes, and are then standardized. The test assets consist of the portfolios from the nine investment strategies ($N = 46$). The sample period varies with the factor at hand, according to data availability over the 11/1983-12/2017 period ($T = 410$). See nontradable factor descriptions in Tables A5-A7, in the Internet Appendix.

PANEL A:	FMB		$\omega=-1$					$\omega=20$					$\omega=50$				
FIN	λ_g	se	λ_g	se	R2	SR	pval	λ_g	se	R2	SR	pval	λ_g	se	R2	SR	pval
otic	8.36***	(2.87)	0.50**	(0.22)	3.58	0.22	0.00	0.75**	(0.32)	3.68	0.32	0.00	0.76**	(0.33)	3.66	0.33	0.00
noise	-3.95***	(1.33)	-1.12***	(0.36)	16.66	0.15	0.01	-0.83**	(0.35)	16.33	0.09	0.01	-0.79**	(0.36)	16.45	0.10	0.01
sliq	-3.56**	(1.46)	-1.02***	(0.33)	15.72	0.20	0.00	-0.91**	(0.39)	13.13	0.15	0.01	-0.87**	(0.40)	13.12	0.15	0.01
gfc	2.39**	(1.09)	1.13***	(0.40)	36.58	0.16	0.00	1.22**	(0.52)	36.43	0.17	0.00	1.22**	(0.53)	36.43	0.17	0.00
gvol	-4.19**	(1.57)	-0.97***	(0.26)	12.93	0.22	0.00	-1.05***	(0.30)	10.79	0.27	0.00	-1.03***	(0.31)	10.67	0.26	0.00
psliq	7.64***	(2.58)	0.50**	(0.23)	4.70	0.19	0.06	0.83*	(0.43)	4.96	0.31	0.06	0.84*	(0.44)	4.93	0.32	0.06
ted	-12.27***	(3.49)	-0.67*	(0.33)	3.75	0.23	0.30	-0.97**	(0.43)	4.05	0.39	0.21	-0.98**	(0.44)	4.02	0.41	0.21
lib ois	-4.71**	(1.86)	-1.11**	(0.52)	15.48	0.16	0.08	-1.55**	(0.76)	15.23	0.24	0.12	-1.52*	(0.76)	15.15	0.25	0.12
move	-5.13***	(1.88)	-0.72**	(0.30)	11.85	0.14	0.00	-0.81**	(0.38)	11.85	0.12	0.00	-0.81**	(0.39)	11.87	0.13	0.00
vxo	-4.36***	(1.24)	-1.34***	(0.46)	27.10	0.12	0.00	-1.36**	(0.65)	26.98	0.14	0.00	-1.35**	(0.66)	27.00	0.15	0.00
eqrv	-5.74***	(1.82)	-0.88**	(0.42)	11.71	0.21	0.12	-1.31*	(0.71)	11.60	0.32	0.13	-1.32*	(0.72)	11.52	0.32	0.13
PANEL B:	FMB		$\omega=-1$					$\omega=20$					$\omega=50$				
TXT	λ_g	se	λ_g	se	R2	SR	pval	λ_g	se	R2	SR	pval	λ_g	se	R2	SR	pval
gepu	-8.12***	(2.67)	-0.85***	(0.30)	8.61	0.19	0.00	-1.41***	(0.48)	8.02	0.32	0.00	-1.40***	(0.49)	7.92	0.34	0.00
gepu ppp	-8.25***	(2.65)	-0.86***	(0.31)	8.85	0.19	0.00	-1.46***	(0.48)	8.45	0.32	0.00	-1.45***	(0.50)	8.36	0.34	0.00
epu all	-5.18**	(2.32)	-0.43*	(0.24)	4.35	0.15	0.04	-0.34	(0.40)	3.92	0.13	0.05	-0.32	(0.41)	3.92	0.13	0.05
epu mp	-7.62***	(2.65)	-0.45*	(0.23)	4.50	0.16	0.06	-0.59	(0.40)	4.02	0.25	0.07	-0.59	(0.41)	4.00	0.25	0.07
epu reg	-8.86**	(3.62)	-0.42*	(0.24)	2.39	0.21	0.15	-0.45	(0.34)	1.79	0.27	0.30	-0.44	(0.34)	1.76	0.27	0.32
epu tr	-11.78***	(3.64)	-0.42**	(0.19)	1.62	0.25	0.16	-0.63	(0.38)	1.44	0.41	0.36	-0.62	(0.38)	1.39	0.42	0.37
emv ov	-6.79***	(2.30)	-0.72**	(0.32)	7.68	0.18	0.08	-0.91*	(0.45)	7.44	0.26	0.09	-0.91*	(0.46)	7.41	0.27	0.09
emv mout	-7.25***	(2.37)	-0.74**	(0.32)	7.32	0.20	0.09	-0.95**	(0.43)	6.80	0.29	0.11	-0.95**	(0.43)	6.73	0.30	0.11
emv mqnt	-7.77***	(2.59)	-0.65**	(0.29)	5.51	0.20	0.11	-0.84**	(0.38)	4.97	0.30	0.13	-0.83**	(0.39)	4.90	0.31	0.13
emv inf	-7.62***	(2.73)	-0.52**	(0.25)	5.42	0.16	0.02	-0.72*	(0.40)	5.52	0.24	0.02	-0.73*	(0.40)	5.51	0.25	0.02
emv com	-6.37***	(2.24)	-0.68**	(0.31)	7.81	0.17	0.04	-0.84*	(0.47)	7.68	0.23	0.05	-0.84*	(0.48)	7.66	0.24	0.05
emv ir	-9.09***	(2.94)	-0.62**	(0.29)	5.08	0.20	0.09	-0.90*	(0.51)	5.07	0.32	0.10	-0.90*	(0.52)	5.03	0.33	0.10
emv fc	-7.57***	(2.73)	-0.61	(0.40)	4.52	0.20	0.38	-0.68*	(0.35)	4.13	0.25	0.28	-0.68*	(0.34)	4.10	0.26	0.28
emv fp	-8.41***	(2.93)	-0.58**	(0.29)	4.71	0.19	0.17	-0.78*	(0.40)	4.55	0.29	0.17	-0.78*	(0.41)	4.51	0.30	0.17
emv tx	-7.42***	(2.63)	-0.63**	(0.30)	5.49	0.20	0.16	-0.79*	(0.43)	5.07	0.28	0.18	-0.78*	(0.44)	5.01	0.28	0.19
emv mp	-7.18***	(2.41)	-0.74**	(0.28)	7.21	0.20	0.03	-0.91**	(0.35)	6.80	0.27	0.04	-0.91**	(0.35)	6.76	0.28	0.04
emv reg	-8.27***	(2.74)	-0.65*	(0.33)	5.93	0.20	0.18	-0.89*	(0.47)	5.67	0.30	0.20	-0.90*	(0.48)	5.62	0.31	0.21
emv freg	-9.00***	(2.96)	-0.65*	(0.33)	4.98	0.21	0.29	-0.88*	(0.48)	4.75	0.32	0.35	-0.88*	(0.48)	4.71	0.33	0.36
PANEL C:	FMB		$\omega=-1$					$\omega=20$					$\omega=50$				
MAC	λ_g	se	λ_g	se	R2	SR	pval	λ_g	se	R2	SR	pval	λ_g	se	R2	SR	pval
cus(y)	-13.49***	(4.14)	-0.27	(0.20)	1.31	0.20	0.22	-0.52*	(0.29)	1.62	0.34	0.15	-0.54*	(0.29)	1.63	0.35	0.15
ipw(q)	-6.27***	(1.96)	-0.17	(0.17)	0.67	0.17	0.43	-0.42*	(0.24)	0.87	0.37	0.25	-0.43*	(0.24)	0.88	0.38	0.24
ipw/us(q)	-7.34***	(2.19)	-0.32*	(0.16)	1.14	0.25	0.16	-0.57**	(0.26)	1.24	0.43	0.14	-0.58**	(0.27)	1.23	0.44	0.14
cpiw(q)	-7.92***	(2.92)	-0.18	(0.18)	0.99	0.15	0.17	-0.50*	(0.28)	1.35	0.36	0.04	-0.51*	(0.29)	1.33	0.37	0.04
cpiw/us(ey)	-9.60***	(2.78)	-0.17	(0.18)	0.46	0.21	0.74	-0.50*	(0.28)	1.01	0.42	0.25	-0.52*	(0.29)	1.03	0.43	0.23
unew/us(y)	-16.91***	(4.95)	-0.17	(0.15)	0.46	0.21	0.55	-0.61**	(0.24)	1.46	0.42	0.06	-0.64**	(0.25)	1.52	0.44	0.05

Table A11: Risk Premia of Nontradable Factors: RP-Weight Robustness with Two-Factor SDFs

The table presents the risk-premium estimates of selected candidate nontradable factors (g_t). Panel FMB presents the risk-premium point estimates (λ_g) and Shanken standard errors (se) from the standard two-pass procedure, including the constant and the candidate factor. The remaining panels report the estimates from the (augmented) three-pass procedure of GX with two-factor SDF, $\varphi(F_{1-2})$, for $\omega = -1$ (no overweight), $\omega = 20$ (baseline RP-weight), and $\omega = 50$ (high RP-weight) in the left, mid, and right panels, respectively. The risk-premium estimates (λ_g) are reported along with the Newey-West standard errors (se), computed following Giglio and Xiu (2021); ***,**,* denote significance, respectively, at the 1-, 5- and 10-percent levels. As for the three-pass method, for each factor and a given SDF, we also report the spanning R^2s (R2) resulting from projecting the factor onto the k latent factors entering the SDF; the Sharpe ratios (SR) associated with the projected factor, i.e., the return-based counterpart to the original nontradable factor; and the p-value (pval) of the test of GX that factor g_t is weak. In Panels A, B and C, we present financial (FIN), text-based (TXT), and macro (MAC) candidate g factors that have significant risk-premium estimates according to at least one of the RP-weight reported. When a macro factor is significant for multiple frequencies, we present the most representative. Factors are expressed as innovations, using the residuals from AR(1) processes, and are then standardized. The test assets consist of the portfolios from the nine investment strategies ($N = 46$). The sample period varies with the factor at hand, according to data availability over the 11/1983-12/2017 period ($T = 410$). See nontradable factor descriptions in Tables A5-A7, in the Internet Appendix.

PANEL A:	FMB		$\omega=-1$					$\omega=20$					$\omega=50$				
FIN	λ_g	se	λ_g	se	R2	SR	pval	λ_g	se	R2	SR	pval	λ_g	se	R2	SR	pval
otic	8.36***	(2.87)	0.30*	(0.15)	2.99	0.14	0.00	0.78***	(0.27)	3.67	0.34	0.00	0.85***	(0.30)	3.55	0.38	0.00
icap	3.41**	(1.56)	0.52**	(0.25)	10.76	0.13	0.00	1.20***	(0.34)	7.93	0.35	0.00	1.22***	(0.36)	6.56	0.40	0.00
mf2	-1.22	(2.33)	-0.27*	(0.15)	3.33	0.12	0.04	-0.31	(0.21)	3.07	0.15	0.05	-0.29	(0.22)	3.04	0.14	0.04
mf3	-0.01	(2.37)	0.35**	(0.16)	4.42	0.14	0.06	0.35	(0.23)	3.42	0.16	0.09	0.28	(0.23)	3.34	0.13	0.09
noise	-3.95***	(1.33)	-1.10***	(0.37)	16.65	0.14	0.00	-1.41***	(0.43)	11.27	0.32	0.00	-1.39***	(0.43)	10.41	0.34	0.00
sliq	-3.56**	(1.46)	-0.78***	(0.25)	14.32	0.14	0.00	-1.33***	(0.38)	10.33	0.31	0.00	-1.30***	(0.38)	9.59	0.33	0.00
gfc	2.39**	(1.09)	1.04***	(0.38)	36.44	0.14	0.00	1.91***	(0.45)	32.33	0.28	0.00	1.93***	(0.48)	30.12	0.29	0.00
gliq	-4.28	(3.25)	-0.24**	(0.11)	2.06	0.14	0.00	-0.44**	(0.19)	1.17	0.34	0.03	-0.42*	(0.21)	0.90	0.36	0.06
gvol	-4.19**	(1.57)	-0.54**	(0.21)	10.21	0.14	0.00	-1.32***	(0.32)	10.13	0.35	0.00	-1.39***	(0.32)	9.07	0.38	0.00
psliq	7.64***	(2.58)	0.25	(0.16)	3.82	0.11	0.04	0.91**	(0.36)	4.91	0.34	0.04	1.00**	(0.41)	4.61	0.39	0.05
corp	-2.03	(1.57)	-0.78**	(0.32)	12.46	0.14	0.06	-0.87**	(0.41)	7.86	0.24	0.13	-0.80**	(0.39)	7.25	0.22	0.14
ted	-12.27***	(3.49)	-0.35*	(0.21)	2.56	0.14	0.24	-0.89**	(0.43)	3.98	0.35	0.14	-0.96**	(0.45)	4.02	0.40	0.13
lib ois	-4.71**	(1.86)	-1.08**	(0.52)	15.40	0.14	0.05	-1.88*	(0.99)	14.32	0.33	0.09	-1.94*	(1.04)	13.73	0.37	0.10
move	-5.13***	(1.88)	-0.72***	(0.26)	11.85	0.14	0.00	-1.35***	(0.46)	7.75	0.34	0.02	-1.34***	(0.48)	7.09	0.37	0.02
vxo	-4.36***	(1.24)	-1.42***	(0.38)	26.96	0.14	0.00	-2.04***	(0.60)	20.22	0.34	0.00	-2.05***	(0.63)	18.92	0.37	0.00
eqrv	-5.74***	(1.82)	-0.48*	(0.25)	9.33	0.13	0.07	-1.44**	(0.63)	11.45	0.36	0.06	-1.58**	(0.71)	10.71	0.40	0.07
PANEL B:	FMB		$\omega=-1$					$\omega=20$					$\omega=50$				
TXT	λ_g	se	λ_g	se	R2	SR	pval	λ_g	se	R2	SR	pval	λ_g	se	R2	SR	pval
gepu	-8.12***	(2.67)	-0.71**	(0.29)	7.61	0.13	0.00	-1.52***	(0.43)	7.93	0.35	0.00	-1.59***	(0.45)	7.64	0.40	0.00
gepu ppp	-8.25***	(2.65)	-0.73**	(0.30)	7.90	0.13	0.00	-1.58***	(0.44)	8.35	0.36	0.00	-1.65***	(0.46)	8.06	0.40	0.00
epu all	-5.18**	(2.32)	-0.29	(0.19)	4.05	0.10	0.02	-0.67*	(0.35)	2.89	0.33	0.09	-0.68*	(0.39)	2.41	0.37	0.15
epu mp	-7.62***	(2.65)	-0.21	(0.15)	3.57	0.08	0.03	-0.70**	(0.33)	3.91	0.30	0.06	-0.76**	(0.37)	3.67	0.34	0.07
fsi tx	-4.06*	(2.27)	-0.33*	(0.19)	4.46	0.13	0.05	-0.71*	(0.38)	2.70	0.35	0.16	-0.69*	(0.40)	2.09	0.39	0.24
emv ov	-6.79***	(2.30)	-0.51**	(0.24)	7.01	0.13	0.04	-1.15**	(0.46)	6.90	0.36	0.05	-1.21**	(0.50)	6.31	0.40	0.06
emv mout	-7.25***	(2.37)	-0.46*	(0.24)	6.12	0.12	0.06	-1.11**	(0.45)	6.55	0.35	0.05	-1.18**	(0.48)	6.07	0.40	0.06
emv mqnt	-7.77***	(2.59)	-0.38*	(0.22)	4.38	0.12	0.09	-0.95**	(0.40)	4.85	0.35	0.06	-1.01**	(0.43)	4.51	0.40	0.06
emv inf	-7.62***	(2.73)	-0.41**	(0.19)	5.22	0.12	0.01	-0.97***	(0.35)	4.99	0.35	0.02	-1.02**	(0.39)	4.54	0.40	0.03
emv com	-6.37***	(2.24)	-0.53**	(0.23)	7.47	0.13	0.02	-1.15**	(0.45)	6.82	0.36	0.04	-1.20**	(0.49)	6.14	0.40	0.05
emv ir	-9.09***	(2.94)	-0.36*	(0.19)	4.12	0.12	0.05	-0.97**	(0.44)	5.03	0.35	0.06	-1.04**	(0.49)	4.80	0.40	0.07
emv fx	-5.15**	(2.19)	-0.31*	(0.18)	3.49	0.11	0.02	-0.71**	(0.29)	2.81	0.35	0.04	-0.73**	(0.31)	2.44	0.40	0.06
emv fp	-8.41***	(2.93)	-0.38*	(0.21)	4.07	0.12	0.10	-0.91**	(0.40)	4.38	0.35	0.08	-0.97**	(0.43)	4.07	0.40	0.09
emv tx	-7.42***	(2.63)	-0.40*	(0.22)	4.70	0.12	0.09	-0.96**	(0.43)	4.82	0.35	0.10	-1.01**	(0.46)	4.41	0.40	0.11
emv gov	-15.41***	(5.72)	-0.13	(0.15)	0.92	0.09	0.30	-0.49*	(0.28)	1.50	0.33	0.23	-0.54*	(0.31)	1.52	0.38	0.23
emv mp	-7.18***	(2.41)	-0.50**	(0.22)	6.35	0.13	0.02	-1.12***	(0.38)	6.41	0.35	0.01	-1.18***	(0.40)	5.91	0.40	0.02
emv reg	-8.27***	(2.74)	-0.38	(0.24)	4.85	0.12	0.11	-1.01*	(0.51)	5.56	0.35	0.11	-1.08*	(0.55)	5.23	0.40	0.12
emv freg	-9.00***	(2.96)	-0.40*	(0.23)	4.01	0.13	0.22	-0.96*	(0.51)	4.70	0.35	0.20	-1.02*	(0.55)	4.48	0.40	0.20
emv tp	-7.47**	(3.34)	-0.34**	(0.17)	3.35	0.13	0.01	-0.70**	(0.31)	2.54	0.36	0.07	-0.72**	(0.34)	2.20	0.40	0.10
PANEL C:	FMB		$\omega=-1$					$\omega=20$					$\omega=50$				
MAC	λ_g	se	λ_g	se	R2	SR	pval	λ_g	se	R2	SR	pval	λ_g	se	R2	SR	pval
cpius(y)	1.06	(2.76)	0.17*	(0.09)	0.96	0.14	0.18	0.19	(0.21)	0.34	0.27	0.56	0.12	(0.23)	0.26	0.20	0.65
cfnai(q)	3.64	(2.41)	0.19*	(0.11)	1.19	0.14	0.18	0.36	(0.24)	0.92	0.31	0.32	0.36	(0.27)	0.80	0.33	0.37
ipw/us(q)	-7.34***	(2.19)	-0.10	(0.09)	0.43	0.13	0.42	-0.44**	(0.21)	1.09	0.35	0.10	-0.52**	(0.23)	1.18	0.40	0.08
unew/us(y)	-16.91***	(4.95)	-0.03	(0.10)	0.17	0.06	0.72	-0.27	(0.21)	0.47	0.34	0.37	-0.38*	(0.22)	0.67	0.39	0.19

Table A12: Latent Factor Pricing Diagnostics: Sample Including HML Portfolios

The table presents model diagnostics of the first two steps of the asset pricing procedure of Giglio and Xiu (2021) applied to currency portfolios excess returns, where the pricing factors are latent and are estimated using the RP-PCA method of Lettau and Pelger (2020a,b). We report diagnostics for RP-PCA implemented without "overweight" on the means ($\omega = -1$), i.e., standard PCA, and with increasing values of the RP-weight (ω = 10, 20 and 50). We consider stochastic discount factors, $\varphi(F_{1-k})$, including an increasing number of latent factors, $k = 1, 2, \ldots, 6$. Tab A.I First pass, Panel A: Two-pass Statistics, shows the average idiosyncratic variance, $\overline{\sigma}^2_\epsilon = \frac{1}{N}\sum_{n=1}^{N}[Var(\hat{\epsilon}_n)/Var(X_n)]$, and the average root-mean-square pricing errors, $\overline{RMS}_\alpha = \sqrt{\hat{\alpha}\hat{\alpha}^\top/N}$, obtained by estimating $X_{nt} = \alpha_n + \hat{F}_t\psi_n^\top + \epsilon_{nt}$, for $n = 1..., N$ test assets, and $t = 1..., T$ months. Tab A.II Second pass presents the R-squared values ($R^2(\%)$), and the mean absolute errors (MAE) of the cross-sectional regression, $\overline{X}_n = \hat{\psi}_n\gamma^\top + a_n$, for $n = 1, \ldots, N$, where γ is the $1 \times K$ vector of latent factors' prices of risk. Tab B.I Components, Panel B: Sharpe Ratios, presents the maximal Sharpe ratio (SR) from the tangency portfolio of the mean-variance frontier spanned by the linear combination of the K selected latent factors, $\hat{F} \times \hat{b}_{MV}^\top$, where $\hat{b}_{MV}$ is a 1 × K vector with entries $\hat{b}_{MV,k} = \mu_{F,k}/\sigma^2_{F,k}$, with $\mu_{F,k}$ and $\sigma^2_{F,k}$ denoting the k-th factor's mean and variance. ΔSR denotes the difference in SRs between SDFs with k and $k-1$ factors. The $\hat{b}_{MV,k}$ entry represents the k-th factor's weight in the SDF, $\varphi_t = 1 - (\hat{F}_t - \mu_F)\hat{b}_{MV}^\top$. The test assets consist of the plain and HML portfolios from the nine investment strategies ($N = 55$), for the period 11/1983-12/2017 at monthly frequency ($T = 410$).

	Panel A: Two-pass Statistics				Panel B: Sharpe Ratios				
	A.I First pass		A.II Second pass		B.I Components				
$\omega = -1$	$\overline{\sigma}^2_\epsilon$	$\overline{RMS}_\alpha$	$R^2(\%)$	MAE		SR	ΔSR	$\hat{b}_{MV,k}$	$\mu_{F,k}$
$\varphi(F_1)$	34.26	2.49	12.84	1.85	F_1	0.10	–	0.05	17.48
$\varphi(F_{1-2})$	25.78	2.26	16.58	1.56	F_2	0.15	0.06	0.18	7.66
$\varphi(F_{1-3})$	21.50	1.82	35.99	1.25	F_3	0.26	0.11	0.47	9.98
$\varphi(F_{1-4})$	18.60	1.11	68.22	0.87	F_4	0.39	0.12	0.74	10.71
$\varphi(F_{1-5})$	16.24	1.11	68.38	0.85	F_5	0.39	0.00	0.06	0.68
$\varphi(F_{1-6})$	14.22	0.96	74.97	0.73	F_6	0.41	0.02	0.40	4.07
$\omega = 10$	$\overline{\sigma}^2_\epsilon$	$\overline{RMS}_\alpha$	$R^2(\%)$	MAE		SR	ΔSR	$\hat{b}_{MV,k}$	$\mu_{F,k}$
$\varphi(F_1)$	34.31	2.49	1.65	1.66	F_1	0.11	–	0.06	19.48
$\varphi(F_{1-2})$	26.82	2.03	72.95	0.69	F_2	0.31	0.20	0.52	16.52
$\varphi(F_{1-3})$	22.01	1.14	95.81	0.28	F_3	0.42	0.11	0.60	13.88
$\varphi(F_{1-4})$	18.70	0.84	98.08	0.20	F_4	0.45	0.03	0.38	6.21
$\varphi(F_{1-5})$	16.34	0.84	98.09	0.20	F_5	0.45	0.00	0.03	0.41
$\varphi(F_{1-6})$	14.28	0.77	98.50	0.17	F_6	0.46	0.01	0.28	2.85
$\omega = 20$	$\overline{\sigma}^2_\epsilon$	$\overline{RMS}_\alpha$	$R^2(\%)$	MAE		SR	ΔSR	$\hat{b}_{MV,k}$	$\mu_{F,k}$
$\varphi(F_1)$	34.43	2.51	2.16	1.47	F_1	0.12	–	0.07	20.97
$\varphi(F_{1-2})$	27.68	1.86	93.26	0.35	F_2	0.37	0.25	0.67	18.01
$\varphi(F_{1-3})$	22.10	1.06	98.73	0.16	F_3	0.44	0.07	0.46	12.45
$\varphi(F_{1-4})$	18.73	0.81	99.36	0.12	F_4	0.46	0.02	0.33	5.51
$\varphi(F_{1-5})$	16.36	0.81	99.36	0.12	F_5	0.46	0.00	0.03	0.40
$\varphi(F_{1-6})$	14.30	0.75	99.50	0.10	F_6	0.47	0.01	0.27	2.73
$\omega = 50$	$\overline{\sigma}^2_\epsilon$	$\overline{RMS}_\alpha$	$R^2(\%)$	MAE		SR	ΔSR	$\hat{b}_{MV,k}$	$\mu_{F,k}$
$\varphi(F_1)$	35.01	2.63	58.20	1.02	F_1	0.14	–	0.09	23.53
$\varphi(F_{1-2})$	28.31	1.70	99.06	0.13	F_2	0.40	0.26	0.73	19.30
$\varphi(F_{1-3})$	22.18	1.01	99.77	0.07	F_3	0.45	0.05	0.36	10.99
$\varphi(F_{1-4})$	18.75	0.80	99.88	0.05	F_4	0.46	0.02	0.30	5.04
$\varphi(F_{1-5})$	16.38	0.79	99.88	0.05	F_5	0.46	0.00	0.03	0.39
$\varphi(F_{1-6})$	14.31	0.73	99.90	0.04	F_6	0.47	0.01	0.26	2.64

V.2 Stability of Factor Structure and Risk Premia

In this section we complement the in-sample (IS) analysis carried out in Sections 4.2 and 4.3 with the out-of-sample (OOS) analysis of the SDF, and provide time-varying recursive estimates of the three-pass model for some nontradable factors. The objective of the section is twofold. First, we shed light on the robustness of the in-sample estimation of the SDF in an out-of-sample setting. Second, we assess the stability of the factor structure and the risk-premium estimates over time. We tackle both objectives by varying the number of factors and the RP-weights, as in the main analysis.

V.2.1 Out-of-Sample Analysis

To start with, we address our first objective, which pertains to the estimation of the SDF in an out-of-sample setting. Similarly to the in-sample analysis, we rely on the maximal Sharpe ratios, the root-mean-square pricing errors, and the average idiosyncratic variances to assess the properties of the factor structure of returns. However, we compute these diagnostic criteria out-of-sample, as follows. For a given ω and k (for $\omega = -1, 10, 20, 50$, and $k = 1, ..., 6$), we first estimate factors and local time-varying loadings (F_t and ψ_t) containing information only up to time t. Both factors and loadings are estimated recursively with an initial expanding window of ten years to resemble the real-time behavior of a representative FX investor. We then follow closely Lettau and Pelger (2020b) in the computation of the out-of-sample diagnostics. Thus, we use the time-t loadings to predict factors ($F_{t|t+1}$) and test-asset returns ($X_{t|t+1}$) at time $t + 1$. Using the predicted returns, we obtain the out-of-sample pricing errors at time $t + 1$. Next, with the mean and variance of the pricing errors, we can easily compute the out-of-sample average pricing errors and idiosyncratic variation. Finally, to compute the maximal Sharpe ratios, we first obtain the optimal portfolio weights using information up to time t, and then recover the $t + 1$ predicted optimal portfolio return. We repeat these steps for each month t until we recover the time-series of optimal returns, from which we calculate the out-of-sample Sharpe ratios.

Table A13 reports the in-sample and out-of-sample diagnostics for the period 1993-2017. We find that the four-factor model with $\omega = 20$ (our preferred RP-weight) yields an in-sample Sharpe ratio of 0.58 (Panel A). The Sharpe ratio is mostly due to the second factor, but is also contributed to by the other two weak factors (i.e., $\hat{F}_3$ and $\hat{F}_4$) and by the first strong factor, in proportions similar to the full-sample results (Table 1). Differently from the full-sample results, we observe a small contribution of $\hat{F}_6$, while $\hat{F}_5$ remains a time-series factor with no impact on the maximal SR. The magnitudes and patterns of the average pricing errors and idiosyncratic variation estimated in-sample over the shorter sample 1993-2017 also mirror the full-sample results (again with the exception of $\hat{F}_6$ that has a more substantial impact on the pricing errors).

Panel B shows that the four-factor model with the same RP-weight of 20 yields a Sharpe ratio of 0.45 out of sample. This result is consistent with previous evidence showing that the implied Sharpe ratios are smaller out of sample than in sample using RP-PCA (Lettau and Pelger, 2020b). The out-of-sample Sharpe ratios exhibit a similar pattern to the in-sample ones, but with some differences. While the first three factors contribute in similar proportions to the maximal Sharpe ratio as the in-sample ones, the out-of-sample fourth factor differs by exerting essentially no impact on the Sharpe ratio. Moreover, the root-mean-square pricing errors and average idiosyncratic variances are higher out of sample than in sample, although their patterns are to a large degree similar. (Another difference, but arguably less important, regards $\hat{F}_6$ that does not reduce out-of-sample the pricing error as much as in-sample.) Therefore, the out-of-sample diagnostic criteria seem to suggest more clearly a model with three factors.

Turning to the results' sensitivity to the RP-weights, we find that RP-PCA outperforms PCA also in the out-of-sample setting, with no evidence of a trade-off in choosing RP-PCA over PCA. In fact, using RP-PCA we observe much lower average pricing errors than PCA, on the backdrop of essentially equal average idiosyncratic variations. Moreover, the out-of-sample implied Sharpe ratios remain substantially higher using RP-PCA than PCA.[43] However, compared to the in-sample evidence, the out-of-sample Sharpe ratios appear stable using PCA, while they deteriorate using RP-PCA. Thus, there is still a sizable gain in choosing RP-PCA over PCA, but the difference between the two shrinks out of sample.[44] Finally, we observe very similar results for RP-PCA models estimated out-of-sample using different but reasonably high RP-weights. This out-of-sample evidence on the use of different RP-weights is largely consistent with our in-sample evidence, corroborating the in-sample choice of the baseline RP-weight.

Summing up, we find that the in-sample results are largely confirmed out of sample. The currency SDF comprises at least three latent factors. Moreover, RP-PCA still outperforms PCA, in terms of both Sharpe ratios and pricing errors, and yet yields comparable average idiosyncratic variance. Hence, we document no trade-off in practice also when the SDFs are constructed in real time. However, in line with previous studies (e.g., Lettau and Pelger, 2020b; Giglio et al., 2021c), it is apparent that the performance of RP-PCA deteriorates out of sample, while PCA displays a higher degree of stability.

V.2.2 Time-Varying Analysis

In what follows, we shift the focus to the stability of both the factor structure implicit in FX portfolio returns and the nontradable factors' risk-premium estimates over time. First, we resort to the GX and O tests for the number of factors (see Sections 2.1.1 and 4.2), and the generalized correlations. Unlike the evidence in Table A13, these two additional diagnostic criteria are not constructed out of sample (but simply in-sample using an expanding window of data), yet they arguably delve directly into the stability of the factor structure over time.

We begin the analysis by conducting the GX and O tests over recursive windows with initial size of ten years. The results are clearcut: regardless of the time period, both tests select three factors using RP-PCA, and two factors using PCA (Figure A4). These results are not foregone as the consistency of the tests requires not only large N but also large T, and T is inherently shorter in a recursive estimation, especially over the first estimation windows. Hence, the time-varying evidence closely mirrors the baseline full-sample evidence, suggesting a stable number of factors driving test-asset returns over time.

Then, we proceed with the generalized correlations (GCs), a powerful tool to investigate the stability of the factor structure over time, recently employed by Lettau and Pelger (2020b) in a similar setting to ours.[45]

[43] Similar to the baseline in-sample analysis, we find that using a five-factor model the wedge in performance between the RP-PCA and PCA estimators reduces. This is because using PCA $\hat{F}_5$ is a weak factor with high Sharpe ratio. But this factor plays no role as is not selected by any of the statistical tests for the number of factors, which suggest more parsimonious SDFs (formal test results are reported in the next subsection).

[44] The results are robust if we perform the analysis using an initial expanding window of 20 years (not reported). However, we find that $\hat{F}_4$ contributes to increase the overall Sharpe ratio and reduce the pricing errors. Hence, its behavior is more in line with that documented in the in-sample analysis.

[45] Generalized correlations were first proposed by Bai and Ng (2006), to which we refer the reader for further details.

Intuitively, GCs help us determine whether two sets of factors represent the same factor model, and hence span the same vector space. In doing so, importantly, they account for the fact that a factor model is only identified up to invertible linear transformations. A high qth generalized correlation suggests that the two sets of factors have at least q common factors. In our setting, the first set denotes the factors extracted over the whole sample, while the second set consists of the factors estimated recursively over time. Specifically, we compute the generalized correlations between the total loadings (ψ^{ω}) estimated over the whole sample and the local loadings (ψ_t^{ω}) estimated recursively over expanding windows with initial length of ten years. Therefore, the GCs help us quantify the degree of stability of the factor structure over time.

Figure A5 shows the generalized correlations for selected RP-weights. We find that using $\omega = 20$ the generalized correlations of the first three factors are all above 94% and by adding the fourth factor they are above 87%. The minimum generalized correlation of the third (fourth) factor is somewhat higher (lower) using PCA ($\omega = -1$), while there are no significant differences using the other reported RP-weigths. Hence, the factor structure driving the SDF appears to be stable over time. Of course, because the correlations are estimated recursively, they converge to 100% as we approach the end of the sample. Therefore, to better zoom into the last years, we repeat the analysis using 20-year rolling windows. Above all, Figure A6 confirms the recursive evidence pointing to a stable factor structure. The correlations of the first three factors are all above 96%, and only drop slightly to 91% by including the fourth factor (using PCA they are 95% and 86%, respectively). It is also visible that the fourth generalized correlation reduces slightly over the last part of the sample. The drop is less pronounced using RP-PCA than PCA. Overall, the results point to a rather stable factor structure over time, as we observe only a mild decay of the fourth GC towards the end of the sample.

Finally, we turn to assessing the stability of the risk-premium estimates of the nontradable factors. These estimates are subject to two sources of potential instability, stemming from both the latent-factor SDF structure and the spanning regressions linking the candidate factors to the latent factors. Hence, the nontradable factors' risk premia can inherit the instability of the factor structure (which is common to all candidate factors), but they are also potentially subject to an additional layer of instability specific to the selected candidate nontradable factor. To start with, we extract the factors as before, but for brevity we only report the results using the baseline RP-weight, i.e., $\omega = 20$. Therefore, the extracted recursive factors are the same as those underlying the generalized correlations of Panel $\omega = 20$ in Figure A5. Based on these recursive estimates of the factors, we then estimate the three-pass model recursively and obtain time-varying estimates of the candidate factors' risk premia ($\lambda_g^{(1:t)}$). Figures A7-A10 present the risk-premium point estimates, together with the associated 95% confidence intervals (albeit statistical significance as such is not the main focus here), for the four selected candidate factors.[46] Moreover, for each candidate factor we show the estimates for the one-factor model and for larger models that successively add latent factors one by one (moving from top-left to bottom-right panels). Also note that, because the estimation is recursive, the risk-premium estimates will converge to the full-sample estimates of Table 3 as we reach the end of the sample. Overall, we find that the risk-premium estimates are rather stable, especially when the SDFs include at least three latent factors (the number of factors detected by the tests using RP-PCA) and the estimation window is reasonably long. Specifically, the estimates seem to stabilize around the early 2000s,

[46] We show the results for the same set of factors used in Figure 6. However, because *gepu* has many missing values for the early years of the sample, we replace it with another widely used text-based factor, i.e., equity market volatility.

which is not surprising as 10 years of data might be too few to achieve robust estimates of expected returns. For some candidate factors the impact of the global financial crisis on the estimates is more evident, but the instability is generally contained for all factors.

Summing up, the evidence obtained performing the statistical tests for the number of factors recursively closely mirrors the full-sample evidence. That is, the suggested number of factors is three using RP-PCA, and reduces to two using PCA. This evidence is largely corroborated by the generalized correlations, suggesting that the number of common factors should be at least three. Taken together, the results are consistent with a stable factor structure over time. Finally, we find that also the candidate factors' risk-premium estimates do not show significant degrees of time variation, as long as the estimation window is sufficiently long and the SDF includes at least the first three factors. Taken together, these multiple pieces of evidence seem to suggest that the unconditional three-pass model, if well specified, provides a satisfactory description of dynamically rebalanced FX portfolio returns.

Table A13: Fit of In-Sample and Out-of-Sample RP-PCA Models

The table presents maximal Sharpe ratios (SR), root-mean-square pricing errors ($\overline{RMS}_\alpha$), and unexplained idiosyncratic variation ($\overline{\sigma}^2_\epsilon$) for in-sample (Panel A) and out-of-sample (Panel B) RP-PCA models with selected RP-weights (ω) and number of factors ($\varphi(F_{1-k})$). The out-of-sample analysis is recursive with an initial, expanding window of 10 years. The test assets consist of the portfolios from the nine investment strategies ($N = 46$), for the period 11/1983-12/2017 at monthly frequency ($T = 410$). Thus, the in-sample (IS) and out-of-sample (OOS) model statistics refer to the period 12/1993-12/2017.

	Panel A: In-Sample						Panel B: Out-of-Sample					
	Panel A.I: IS RP-PCA ($\omega = -1$)						Panel B.I: OOS RP-PCA ($\omega = -1$)					
	$\varphi(F_1)$	$\varphi(F_{1-2})$	$\varphi(F_{1-3})$	$\varphi(F_{1-4})$	$\varphi(F_{1-5})$	$\varphi(F_{1-6})$	$\varphi(F_1)$	$\varphi(F_{1-2})$	$\varphi(F_{1-3})$	$\varphi(F_{1-4})$	$\varphi(F_{1-5})$	$\varphi(F_{1-6})$
SR	0.08	0.22	0.27	0.27	0.33	0.35	0.07	0.13	0.29	0.29	0.39	0.40
$\overline{RMS}_\alpha$	1.88	1.55	1.41	1.41	1.27	1.23	1.94	1.61	1.14	1.10	0.96	0.94
$\overline{\sigma}^2_\epsilon$	22.66	18.91	16.79	15.09	13.73	12.65	24.07	21.00	19.49	18.20	17.13	16.13
	Panel A.II: IS RP-PCA ($\omega = 10$)						Panel B.II: OOS RP-PCA ($\omega = 10$)					
	$\varphi(F_1)$	$\varphi(F_{1-2})$	$\varphi(F_{1-3})$	$\varphi(F_{1-4})$	$\varphi(F_{1-5})$	$\varphi(F_{1-6})$	$\varphi(F_1)$	$\varphi(F_{1-2})$	$\varphi(F_{1-3})$	$\varphi(F_{1-4})$	$\varphi(F_{1-5})$	$\varphi(F_{1-6})$
SR	0.09	0.41	0.50	0.56	0.57	0.59	0.07	0.28	0.44	0.44	0.46	0.46
$\overline{RMS}_\alpha$	1.88	1.42	1.12	0.84	0.79	0.68	1.94	1.34	0.87	0.85	0.86	0.85
$\overline{\sigma}^2_\epsilon$	22.67	19.30	17.15	15.62	13.95	12.78	24.04	21.04	19.63	18.51	17.16	16.15
	Panel A.III: IS RP-PCA ($\omega = 20$)						Panel B.III: OOS RP-PCA ($\omega = 20$)					
	$\varphi(F_1)$	$\varphi(F_{1-2})$	$\varphi(F_{1-3})$	$\varphi(F_{1-4})$	$\varphi(F_{1-5})$	$\varphi(F_{1-6})$	$\varphi(F_1)$	$\varphi(F_{1-2})$	$\varphi(F_{1-3})$	$\varphi(F_{1-4})$	$\varphi(F_{1-5})$	$\varphi(F_{1-6})$
SR	0.09	0.45	0.53	0.58	0.58	0.60	0.08	0.36	0.45	0.45	0.46	0.46
$\overline{RMS}_\alpha$	1.89	1.39	1.06	0.78	0.76	0.66	1.93	1.12	0.85	0.84	0.85	0.84
$\overline{\sigma}^2_\epsilon$	22.70	19.49	17.24	15.66	13.98	12.79	24.03	21.17	19.65	18.51	17.16	16.16
	Panel A.IV: IS RP-PCA ($\omega = 50$)						Panel B.IV: OOS RP-PCA ($\omega = 50$)					
	$\varphi(F_1)$	$\varphi(F_{1-2})$	$\varphi(F_{1-3})$	$\varphi(F_{1-4})$	$\varphi(F_{1-5})$	$\varphi(F_{1-6})$	$\varphi(F_1)$	$\varphi(F_{1-2})$	$\varphi(F_{1-3})$	$\varphi(F_{1-4})$	$\varphi(F_{1-5})$	$\varphi(F_{1-6})$
SR	0.11	0.48	0.55	0.59	0.59	0.61	0.09	0.43	0.45	0.45	0.46	0.46
$\overline{RMS}_\alpha$	1.95	1.36	1.01	0.75	0.74	0.65	1.92	0.91	0.84	0.83	0.84	0.83
$\overline{\sigma}^2_\epsilon$	22.83	19.63	17.31	15.68	13.99	12.80	23.99	21.39	19.66	18.51	17.16	16.16

Figure A4: Recursive Estimates of the Number of Factors in the SDF

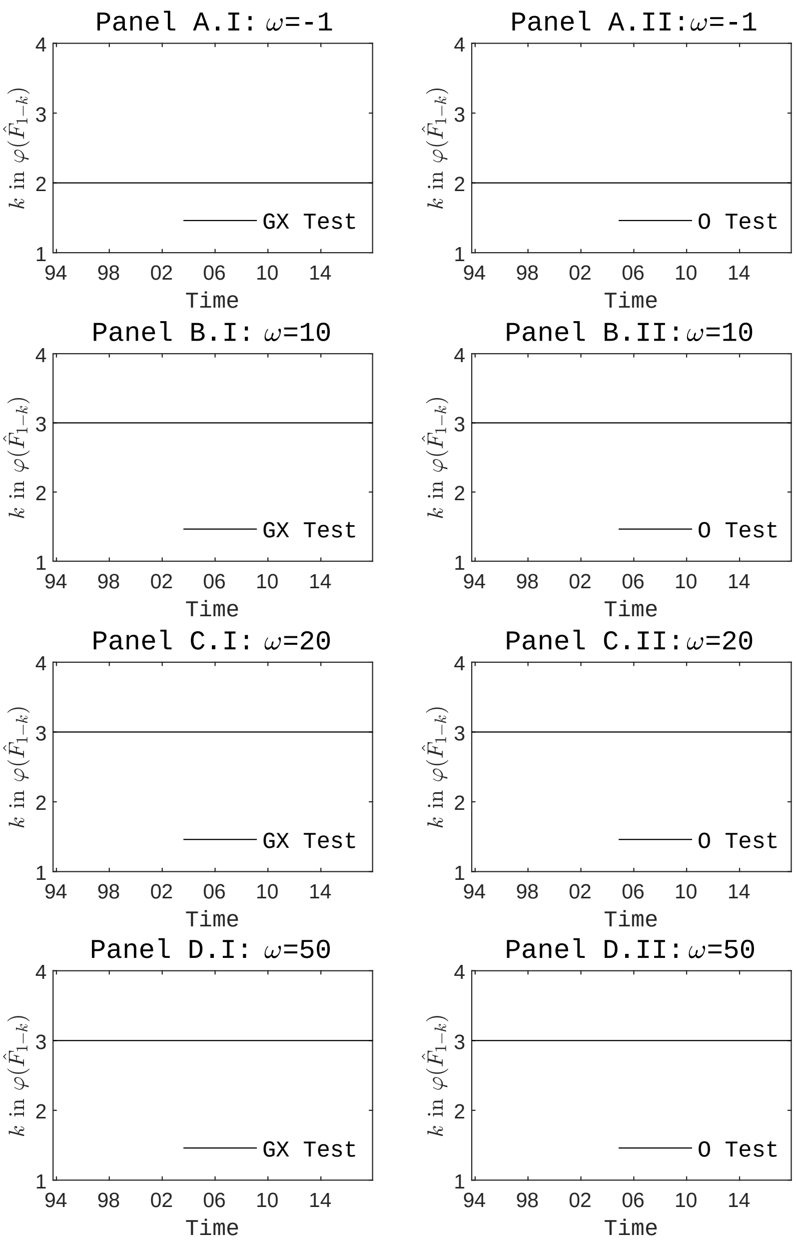

The figure shows the recursive estimates of the optimal number of factors detected by the Giglio and Xiu (2021; GX Test) and the Onatski (2010; O Test) tests for the optimal number of factors (Panels A and B, respectively). That is, the tests inform us about the optimal k in the SDF, $\varphi(\hat{F}_{1-k})$. For the selected RP-weights (with $\omega = -1$ denoting the PCA estimation), we perform the tests on the eigenvalues which are estimated recursively over expanding windows with initial length of ten years. The test assets consist of the portfolios from the nine investment strategies ($N = 46$), for the period 11/1983-12/2017 at monthly frequency ($T = 410$). Thus, model statistics refer to the period 12/1993-12/2017.

Figure A5: Generalized Correlations Between Total and Recursive Local Loadings

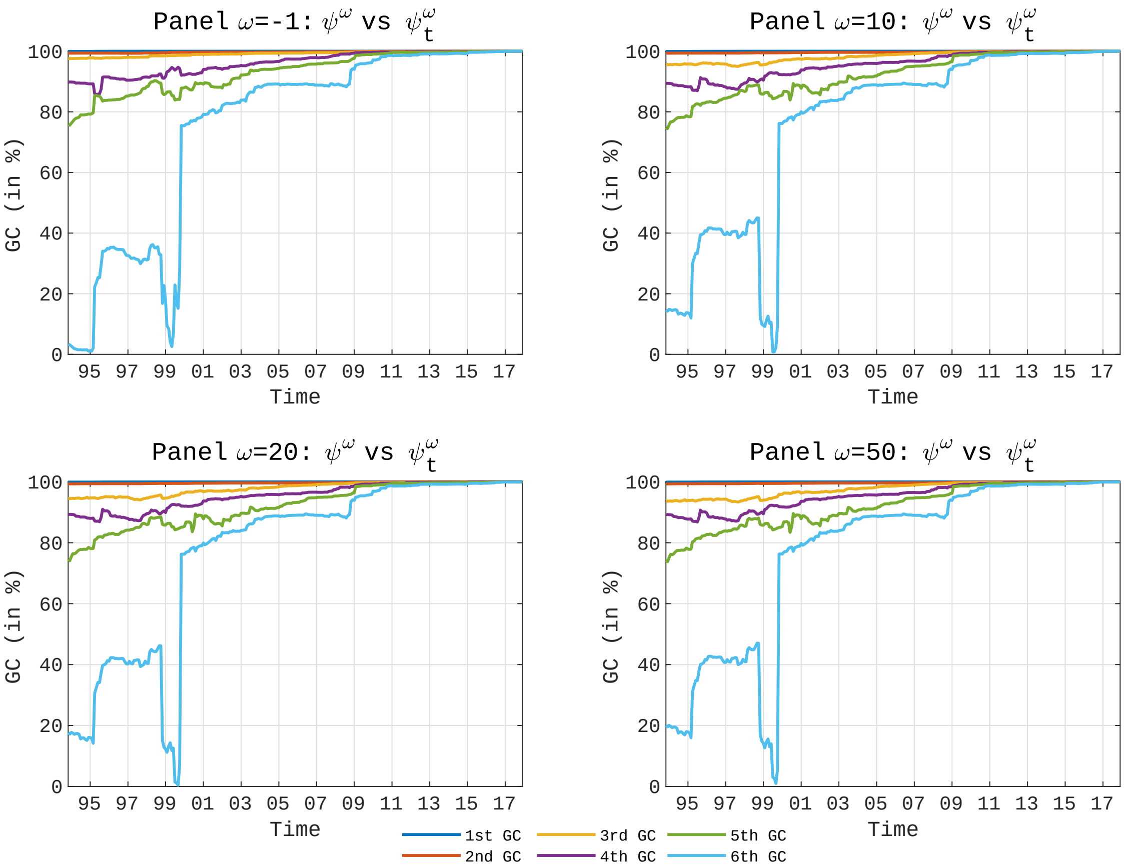

The figure shows the generalized correlations (GCs) between the total loadings (ψ^{ω}) estimated over the whole sample and the local loadings (ψ^{ω}_{t}) estimated recursively over expanding windows with initial length of ten years. We perform the analysis for the selected RP-weights, with $\omega = -1$ denoting the PCA estimation. The test assets consist of the portfolios from the nine investment strategies ($N = 46$), for the period 11/1983-12/2017 at monthly frequency ($T = 410$). Thus, model statistics refer to the period 12/1993-12/2017.

Figure A6: Generalized Correlations Between Total and Rolling Local Loadings

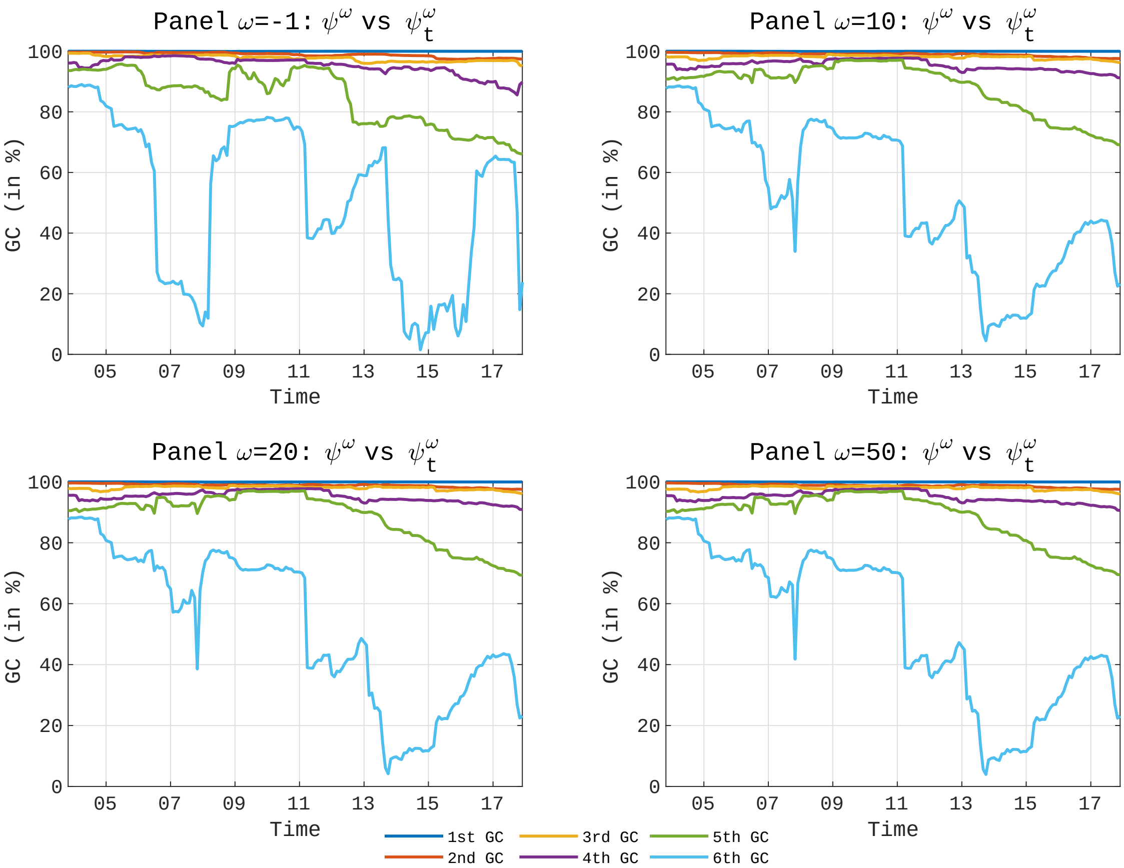

The figure shows the generalized correlations between the total loadings (ψ^{ω}) estimated over the whole sample and the local loadings (ψ^{ω}_{t}) estimated over rolling windows of 20 years. We perform the analysis for the selected RP-weights, with $\omega = -1$ denoting the PCA estimation. The test assets consist of the portfolios from the nine investment strategies ($N = 46$), for the period 11/1983-12/2017 at monthly frequency ($T = 410$). Thus, model statistics refer to the period 12/2003-12/2017.

Figure A7: Recursive Risk-Premium Estimates: Global FX Volatility

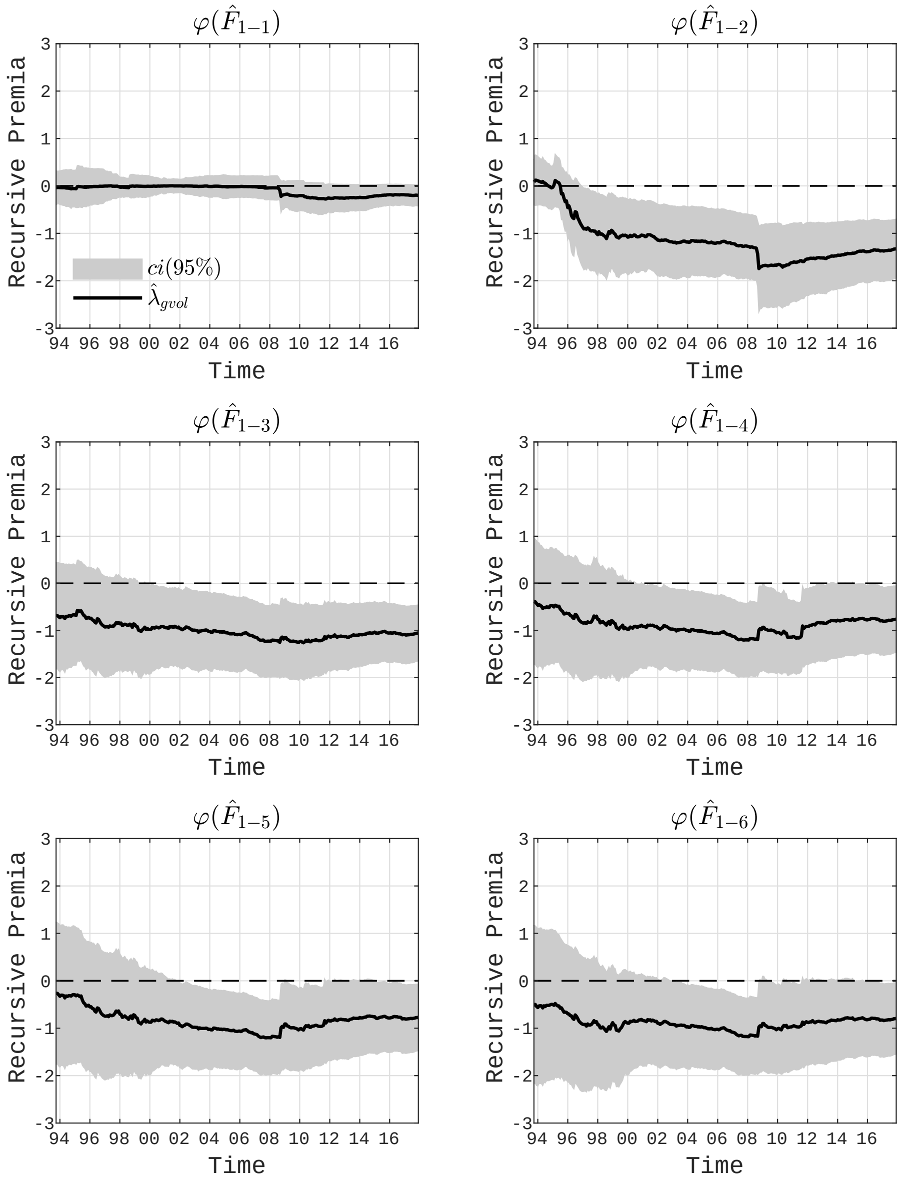

The figure shows the recursive three-pass risk-premium estimates of the global volatility ($gvol$) nontradable financial factor. The model is estimated recursively over expanding windows with initial length of 10 years, using $\omega = 20$ and SDFs including an increasing number of factors ($\varphi(\hat{F}_{1-k})$, for $k = 1, \ldots, 6$). The solid black lines show the risk-premium recursive estimates and the shaded gray areas denote the associated 95% confidence intervals. The test assets consist of the portfolios from the nine investment strategies ($N = 46$), for the period 11/1983-12/2017 at monthly frequency ($T = 410$). Thus, the recursive risk-premium estimates are displayed for the period 12/1993-12/2017.

Figure A8: Recursive Risk-Premium Estimates: Intermediaries' Capital Ratio

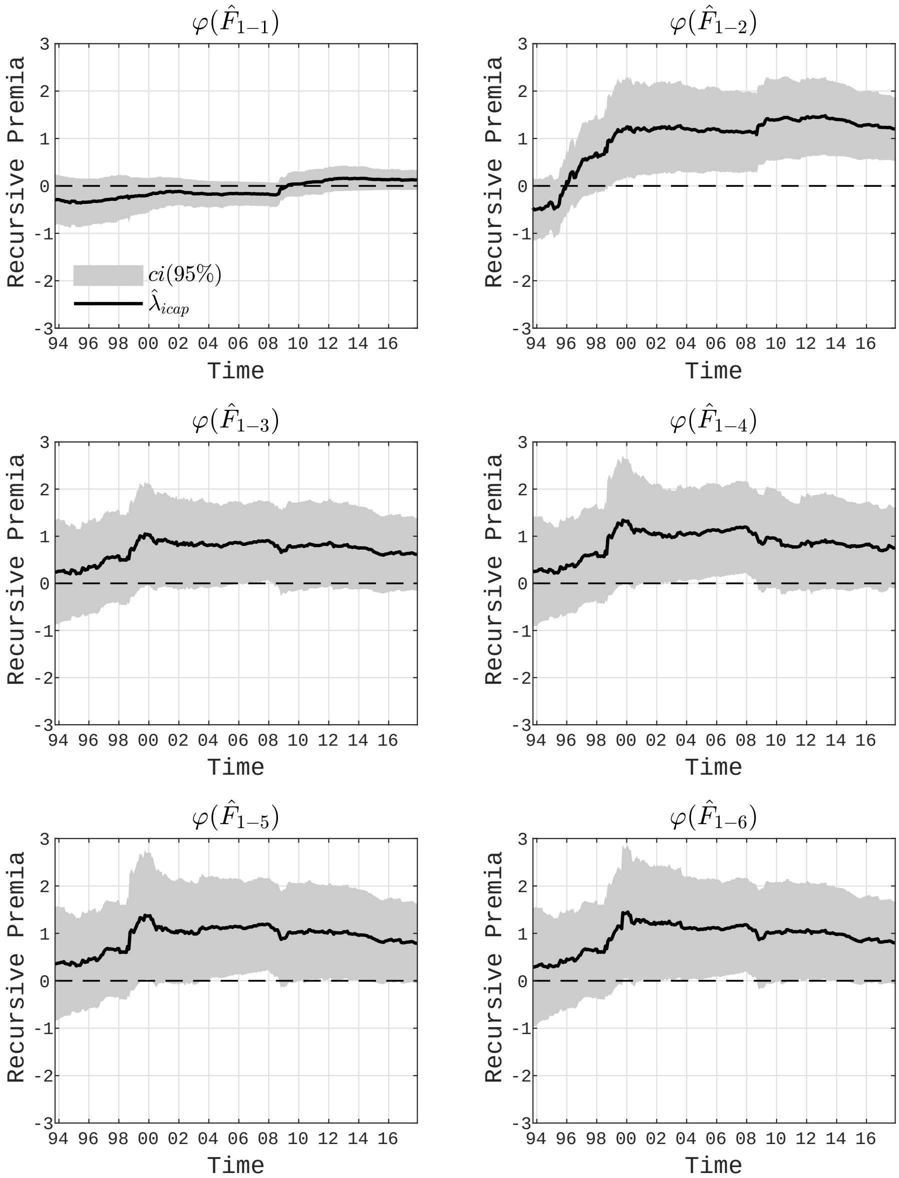

The figure shows the recursive three-pass risk-premium estimates of the intermediaries' capital ratio ($icap$) nontradable financial factor. The model is estimated recursively over expanding windows with initial length of 10 years, using $\omega = 20$ and SDFs including an increasing number of factors ($\varphi(\hat{F}_{1-k})$, for $k = 1, \ldots, 6$). The solid black lines show the risk-premium recursive estimates and the shaded gray areas denote the associated 95% confidence intervals. The test assets consist of the portfolios from the 9 investment strategies ($N = 46$), for the period 11/1983-12/2017 at monthly frequency ($T = 410$). Thus, the recursive risk-premium estimates are displayed for the period 12/1993-12/2017.

Figure A9: Recursive Risk-Premium Estimates: Global Liquidity

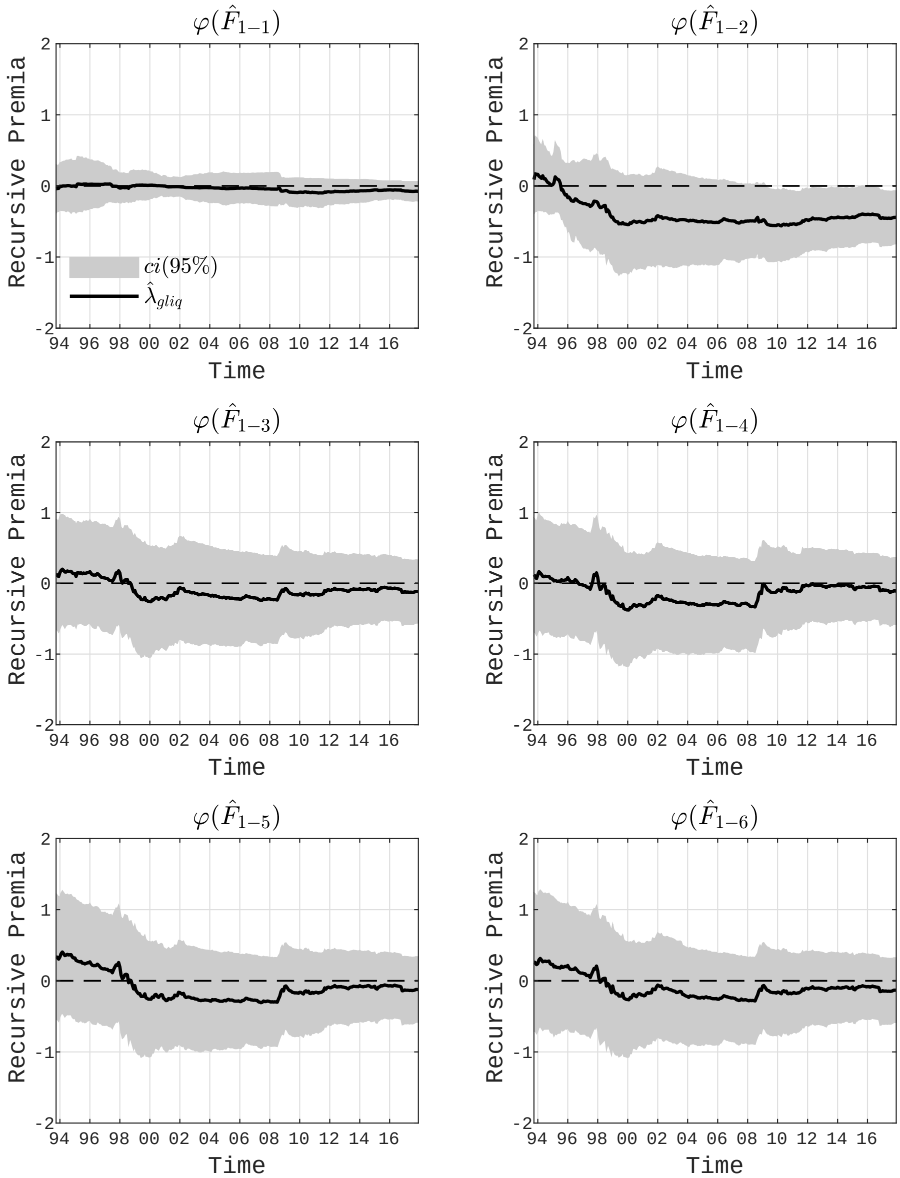

The figure shows the recursive three-pass risk-premium estimates of the global (il)liquidity ($gliq$) nontradable financial factor. The model is estimated recursively over expanding windows with initial length of 10 years, using $\omega = 20$ and SDFs including an increasing number of factors ($\varphi(\hat{F}_{1-k})$, for $k = 1, \ldots, 6$). The solid black lines show the risk-premium recursive estimates and the shaded gray areas denote the associated 95% confidence intervals. The test assets consist of the portfolios from the nine investment strategies ($N = 46$), for the period 11/1983-12/2017 at monthly frequency ($T = 410$). Thus, the recursive risk-premium estimates are displayed for the period 12/1993-12/2017.

Figure A10: Recursive Risk-Premium Estimates: Overall Equity Market Volatility

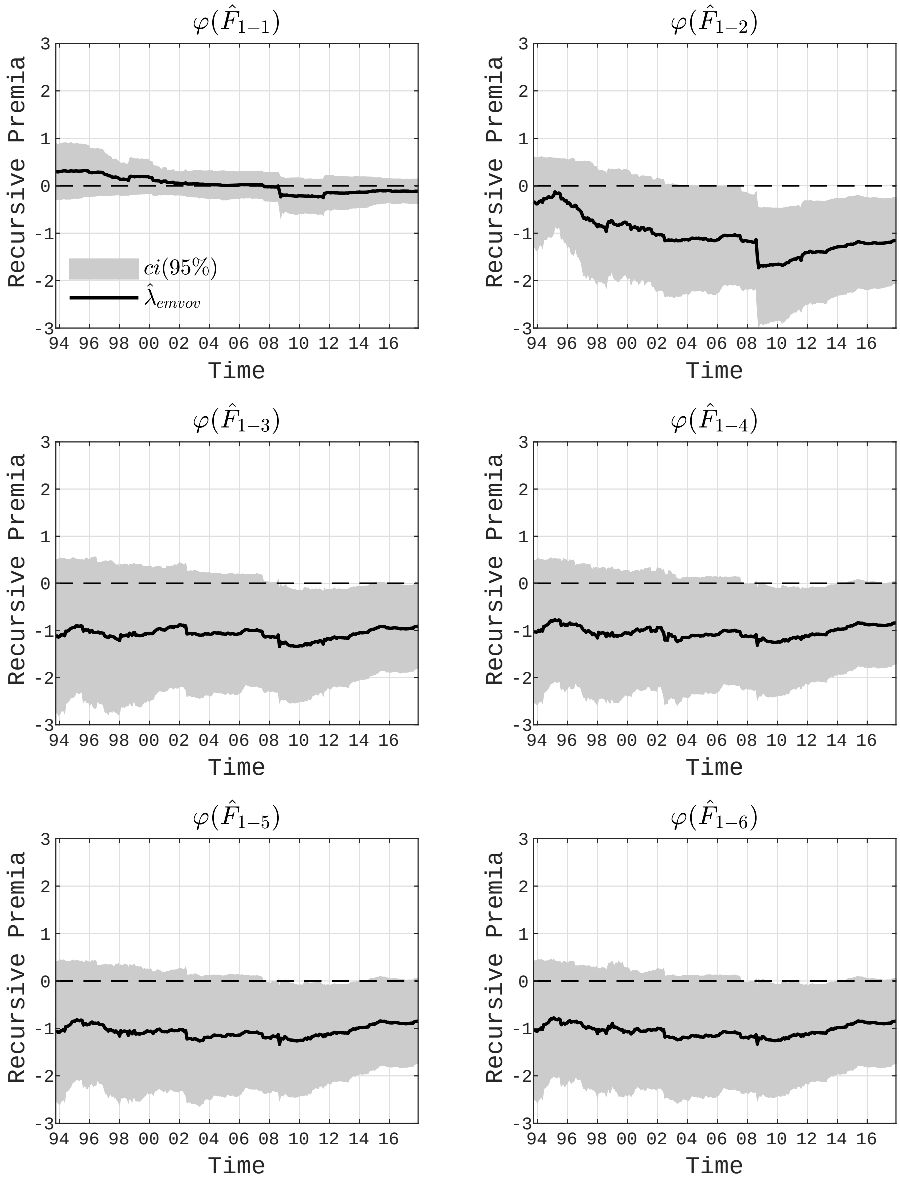

The figure shows the recursive three-pass risk-premium estimates of the overall equity market volatility (*emv ov*) nontradable text-based factor. The model is estimated recursively over expanding windows with initial length of 10 years, using $\omega = 20$ and SDFs including an increasing number of factors ($\varphi(\hat{F}_{1-k})$, for $k = 1, \ldots, 6$). The solid black lines show the risk-premium recursive estimates and the shaded gray areas denote the associated 95% confidence intervals. The test assets consist of the portfolios from the nine investment strategies ($N = 46$), for the period 11/1983-12/2017 at monthly frequency ($T = 410$). Thus, the recursive risk-premium estimates are displayed for the period 12/1993-12/2017.

VI Simulation Analysis

In this section, we study the finite-sample performance of the augmented three-pass inference using Monte Carlo simulations. We also assess how the augmented three-pass estimator performs and compares in simulation with the two-pass estimator. Importantly, we design the simulations to capture the key features of FX returns. By doing so, we essentially tackle two key questions. First, is the three-pass method *reliable* in finite samples, with N and T equal to the dimension of FX portfolio returns? Second, are the omitted-variable and measurement-error problems relevant for pricing currency portfolio returns, and hence is the method *desirable* for FX returns?

To address both questions, we set up the simulation exercise as in Giglio and Xiu (2021) (with the main difference regarding the use of RP-PCA instead of standard PCA in extracting the latent factors), but importantly we tailor the calibration to the specific features of the FX market. The cross section of FX test assets is relatively small, and the underlying data generating process (DGP) is driven by fewer factors compared to the case of equities studied, for example, by GX. Specifically, while the five-factor Fama-French model is the benchmark model for equities (see Fama and French, 2015), the benchmark model for FX returns has relied so far on much more parsimonious SDFs. At least until recently it has been the two-factor model of Lustig et al. (2011), consisting of a Dollar and a Carry factor. Regardless of the identity of the pricing factors, most of the currency SDFs employed so far in the literature have included two factors, mainly due to the small FX cross sections available. Our analysis, however, suggests that with a reasonably large N (which, albeit smaller than the typical N for equities, is much larger than the small cross sections typically used so far in the FX literature), a two-factor SDF is likely to omit relevant sources of FX risk. In fact, we showed that at least three and potentially four latent factors are required to achieve full spanning of the entire SDF, and hence robust estimates of risk premia. Therefore, in larger cross sections, not only the measurement-error problem, but also the omitted-variable problem is likely to be relevant for FX returns.

Although GX report good performance of the three-pass estimator in simulation also for combinations of N and T that resemble the one used in our study, it is not obvious that the estimator performs the same in our case. The above considerations for example suggest that the factor structures driving equity and FX portfolios returns may well differ, and this can in turn weigh on the estimator performance. Therefore, in the simulations we consider a four-factor DGP consisting of the de-noised Dollar, Carry, ST Mom, and Value tradable factors, which most closely approximate the strong latent factor ($\hat{F}_1$) and the three weak latent factors with significant risk premia ($\hat{F}_2$, $\hat{F}_3$, and $\hat{F}_4$) documented in the empirical analysis. In essence, this DGP can be interpreted as a reduced-form model for FX returns and the currency SDF. To remove the noise from the observed tradable factors, we use the three-pass estimator with four latent factors, extracted using RP-PCA with $\omega = 20$. These four de-noised factors should span the entire SDF (as we know from Section 4.2 that four latent factors guarantee this), and hence the simulated asset returns should mirror the properties of the observed ones. This is crucial to then being able to recover the *true* risk premia via the three-pass estimator.

Next, we assume that we may not observe all four factors but some noisy versions, plus a potentially spurious candidate nontradable factor. In this way, both the omitted-variable and the measurement-error problems can manifest entirely. To begin with, we calibrate the candidate factor to U.S. industrial production (*ipus*) similarly to GX, which according to the three-pass estimates qualifies as a spurious factor also for FX returns. However, we then repeat the analysis replacing *ipus* either with *gvol* or *icap* (the global volatility of Menkhoff et al., 2012a, and the intermediaries' capital ratio of He et al., 2017, respectively). These two financial factors are of particular interest as they exemplify non-spurious factors whose risk premia estimates are affected to different extents by the dimension of the SDF considered (as is evident from Section 4.3, Tables 2 and 3). Taken together, these three candidate factors capture the properties of the nontradable factors considered in the empirical analysis.

Finally, to evaluate the performance of the three-pass estimator, we estimate the risk premia of the candidate noisy tradable and nontradable factors by applying the three-pass method to the simulated data. We use not only SDFs of expanding dimensions, but also with different RP-weights (this latter analysis is not in GX, as they are clearly concerned only with the case of PCA). However, to be clear, the ultimate goal of the simulation exercise is not a comparison of the RP-PCA and PCA estimators. Rather, we are interested in assessing the performance of the three-pass and the two-pass estimators in pricing currency returns, i.e., when the DGP and the associated parameters driving the simulations match the properties of the FX portfolio returns studied in our empirical analysis. In this way, we can assess the finite-sample performance of the three-pass estimator, but also shed light on the relevance of the issues of omitted factors and measurement error in the factors for FX returns.[47]

Overall, the simulation results show a satisfactory performance of the three-pass estimator also in our setting, while the two-pass evidence lends support to the argument made in Section 2 that omitting relevant pricing factors from the currency SDF and/or measuring the factors with noise can severely distort the risk-premium estimates. Put simply, the three-pass estimator appears to be both *reliable* and *desirable* for modeling FX portfolio returns, and hence is a valuable method to unveil the sources of the risk-return trade-off in currency investment strategies.

VI.1 Calibration and Simulation Methods

In this section, we describe the details of the calibration and how we design the simulations. In what follows, F_t are the latent factors, Z_t are the observed tradable factors (Dollar, Carry, ST Mom, and Value), and G_t are the potentially spurious observed candidate nontradable factors (*gvol*, *icap*, and *ipus*). The DGP is given by

$$\begin{aligned}
X_{nt} &= \hat{\mathrm{Z}}_t \psi_{Zn}^\top + \varepsilon_{nt}, && n = 1, \dots, 46, \quad t = 1, \dots, T, && \text{(VI.1)} \\
g_{jt}^\iota &= \hat{Z}_t^\iota \eta_{Zj}^\top + u_{jt}, && j = 1, \dots, 3, \quad t = 1, \dots, T, && \text{(VI.2)} \\
\hat{g}_{jt} &= \hat{Z}_t \hat{\eta}_{Zj}^\top, && j = 1, \dots, 3, \quad t = 1, \dots, T, && \text{(VI.3)} \\
\hat{\lambda}_{g_j} &= \hat{\lambda}_Z \hat{\eta}_{Zj}^\top, && j = 1, \dots, 3. && \text{(VI.4)}
\end{aligned}$$

where $\hat{\mathrm{Z}}_t = [\hat{z}_{t1}, \hat{z}_{t2}, \hat{z}_{t3}, \hat{z}_{t4}]$ are the de-noised tradable Z-factors, $\hat{Z}_t^\iota$ denote their demeaned counterparts, $\hat{\lambda}_Z$ collects the de-noised Z-factors risk premia or sample means, and $\mathrm{G}_t = [g_{1t}^\iota, g_{2t}^\iota, g_{3t}^\iota]$ where g_{jt}^ι is the

[47] We refer the reader to GX for a fully-fledged comparison of the three-pass and two-pass estimators, e.g., for different combinations of N and T.

AR(1) innovation to the j-*th* candidate factor g_t.

First, we apply the three-pass estimator with $\omega = 20$ and $k = 4$ to the four noisy tradable Z-factors to recover their de-noised counterparts and risk premia ($\hat{Z}$ and $\hat{\lambda}_Z$). Specifically, equipped with the RP-PCA estimates of the k latent factors ($\hat{F}$), whereby $X_{nt} = \hat{\mathrm{F}}_t \psi_n^\top + \epsilon_{nt}$, we run the following four spanning regressions

$$z_{ti}^{\iota} = \hat{F}_t^{\iota} \xi_i^\top + \zeta_{ti}, \qquad i = 1, \ldots, 4, \quad t = 1, \ldots, T. \tag{VI.5}$$

We then obtain the de-noised Z-factors and their risk premia as follows

$$\hat{z}_{ti} = \hat{F}_t \hat{\xi}_i^\top, \qquad i = 1, \ldots, 4, \quad t = 1, \ldots, T, \tag{VI.6}$$

$$\hat{\lambda}_{Zi} = \hat{\gamma} \hat{\xi}_i^\top, \qquad i = 1, \ldots, 4. \tag{VI.7}$$

Second, by applying the same procedure (i.e., the third step of the estimator) to the potentially spurious nontradable candidate G-factors, we obtain the true estimates of their risk premia. This step reproduces the exposure and risk-premium estimates ($\hat{\eta}$ and $\hat{\lambda}_g$) of Section 4.3. However, while the *gvol* and *icap* risk-premium estimates are displayed in Table 3 (tab $\varphi(F_{1-4})$), the estimates of *ipus* are not, exactly because its risk-premium estimates are not significant regardless of the dimension of the SDF used.

Third, we calibrate the parameters $\hat{\lambda}_Z$, Σ_Z, $\hat{\psi}_Z$, Σ_{ψ_Z}, Σ_ε, $\hat{\eta}_Z$, σ_u and σ_ζ that drive the DGP to exactly match their counterparts in the data. Specifically, $\hat{\lambda}_Z$ collects the means of the de-noised Z-factors obtained using $\varphi(F_{1-4})$ with $\omega = 20$; Σ_Z is the covariance matrix of the de-noised Z-factors; $\hat{\psi}_Z$ collects the test assets' loadings on the de-noised Z-factors; Σ_{ψ_Z} is the covariance matrix of the factor loadings $\hat{\psi}_Z$; Σ_ε is the covariance matrix of the test-asset idiosyncratic risk; $\hat{\eta}_Z$ are the candidate G-factor exposures to the de-noised Z-factors; σ_u is the volatility of a candidate nontradable G-factor idiosyncratic risk; and σ_ζ is the volatility of a tradable Z-factor idiosyncratic risk.

Fourth, we simulate from the DGP. For the generic m-*th* Monte Carlo replication, we generate the artificial realizations of the parameters and variables $\psi_Z^{(m)}$, $\hat{Z}_t^{(m)}$, $\varepsilon_t^{(m)}$, $u_t^{(m)}$, $u_t^{(m)}$ from multivariate normals using the calibrated means and covariances. We then obtain the artificial realizations of the test assets $X_t^{(m)}$ from eq. (VI.1), the j-th noisy nontradable factor $g_{jt}^{(m)}$ from eq. (VI.2), and the i-th noisy tradable factor from $z_{it}^{(m)} = \hat{z}_{it}^{(m)} + \zeta_{it}^{(m)}$.

Fifth, we recover the objects of interest by applying the estimators to the artificial realizations. Specifically, for a given replication m, we first extract the $k = 6$ latent factors, $\hat{F}^{(m)}$, using RP-PCA with ω=-1, 10, 20, and 50. We then estimate the statistics of interest for each of the RP-weights. Namely, we apply the GX and O tests to $X^{(m)}$ to estimate the number of latent factors, $k^{(m)}$; we compute the generalized correlations, $GC^{(m)}$, between $\hat{F}$ and $\hat{F}^{(m)}$ to determine the number of common factors between the whole sample and simulated sets of factors; and we calculate the maximal Sharpe ratio, $SR^{(m)}$, for different combinations of RP-weights and number of latent factors, k in $\varphi(F_{1-k})$. Finally, we turn to estimating the Z-factor and G-factor risk premia and the associated asymptotic variances. We do this using the three-pass estimator for different combinations of ω and k. For comparison, we also estimate the factor risk premia with the two-pass (FMB) estimator, whereby we consider SDFs that include the noisy g-factor of interest and some or all of the Z-factors (i.e., the controls) measured either with or without noise. In this way, we can assess the issues of omitted variable and measurement error on simulated data.

We iterate on steps four and five for $M = 10,000$ times. We then evaluate the performance of the estimators on the M artificial realizations of the data.

VI.2 Simulation Results

In this section, we present the results of the simulations. We first report the three-pass estimates of the Z-factor and G-factor risk premia, which for $\varphi(F_{1-4})$ and $\omega = 20$ represent the true risk premia to which we benchmark the simulation estimates. We then assess the factor structure of the simulated data, and conclude the analysis by inspecting the performance of the three-pass and two-pass estimators in simulation. For the three-pass estimator, we consider models with different number of latent factors and RP-weights to estimate the risk premia of the noisy factors. For the two-pass estimator, we consider SDFs that differ in the number of omitted factors, as well as in the measurement of the factors. With the former estimator we assess the reliability of the method, while with the second we quantify the severity of the omitted-variable and measurement-error problem, and hence we try to determine the desirability of the three-pass estimator in estimating FX risk premia.

VI.2.1 Calibrating Tradable and Nontradable Factor Risk Premia

Table A14 presents the three-pass risk-premium estimates for the tradable Z-factors driving the DGP and the selected nontradable factors. To begin with, we look at the tradable factor estimates, which are an important input into the DGP and hence into the simulations. Moreover, for the tradable factors, we can benchmark the three-pass estimates with the model-free estimates (mf), i.e., the time-series sample averages of the factors. We find that we can easily recover the price of risk of the Dollar factor, regardless of the dimension of the SDF and the choice of the RP-weight. This is not surprising as the Dollar factor is a strong factor, and is mostly explained by $\hat{F}_1$. We then turn to the remaining tradable factors that are arguably more interesting as they are weak factors with high Sharpe ratios (especially Carry and ST Mom). Using RP-PCA with $\omega = 20$ (*Panel T.I*), we find that the three-pass risk-premium estimates of Carry are not different, both statistically and economically, from the mf estimate as long as the SDFs include more than two factors.[48] Similarly, the three-pass estimates of ST Mom and Value are not different from their model-free counterparts with SDFs consisting of at least three and four latent factors, respectively. Overall, this evidence is consistent with the interpretation of the latent factors provided in Section 4.2. Moreover, it is worth noting that the Z-factor risk-premium estimates are rather stable if additional factors are added to the SDF, which is also consistent with full spanning.

By contrast, even with a four-factor SDF, the PCA three-pass estimates are different from the factor averages (*Panel T.II*). By adding a fifth factor (which from Section 4.2 we know has a non-negligible risk premium in the case of PCA), the three-pass estimates move closer to the factor averages, but the gain is substantial only for ST Mom. Taken together, this evidence is useful at least for two reasons. First, it lends support to the choice of RP-PCA over PCA to extract the factors and simulate FX returns so that the artificial return realizations better mirror the properties of the observed returns. Second, these findings

[48]By including $\hat{F}_3$ in the SDF, the Sharpe ratio of de-noised Carry becomes closer to that of observed Carry. This evidence is coherent with the fact that at least three latent factors are needed to accurately span the space of FX portfolios.

provide additional support to the conjecture that for our sample of FX assets a well-specified SDF should include at least three latent factors. Therefore, to de-noise Dollar, Carry, ST Mom, and especially Value, we choose a four-factor SDF whereby the factors are estimated via RP-PCA with $\omega = 20$.

Turning to the nontradable factors, their risk-premium estimates are displayed in *Panel NT*, whereby the true risk premia are displayed in columns $\varphi(\hat{F}_{1-4})$, in *Panel NT.I: RP-PCA*. It is evident that the risk-premium estimates obtained with the RP-PCA three-pass method are rather stable already for SDFs including three factors, while the PCA estimator needs at least five factors to reach stable estimates. Moreover, the table makes clear the different characteristics and behaviors of the three candidate factors. U.S. industrial production (*ipus*) is a spurious candidate factor, as its premium is not significant regardless of the method and the dimension of the SDF used. The intermediaries' capital ratio (*icap*) of He et al. (2017) presents strong variation in the premia estimates as more latent factors are added to the SDF. Specifically, the premium is sizable and statistically significant with a two-factor SDF, turns insignificant with a three-factor SDF, and is again significant (albeit at the 10% level) with larger SDFs. It is clear that the inclusion of the third latent factor (i.e., "*Momentum*") in the SDF reduces in absolute terms the *icap* risk-premium estimates. Finally, the risk-premium estimates of the global volatility (*gvol*) of Menkhoff et al. (2012a) are significant at least at the 5% level regardless of the SDF considered (although *gvol* appears to be related mostly to $\hat{F}_2$).

Summing up, we find that using the three-pass estimator applied to models including four factors, extracted with baseline RP-weight, we achieve risk-premium estimates for the tradable factors that are not statistically different from the factor sample averages. At the same time, the estimates are rather stable when the fifth and sixth factors are included in the SDF, consistent with full spanning of the entire SDF. Finally, it clearly emerges that *ipus*, *icap*, and *gvol* represent factors with substantially different behavior and premia, which exemplify the sort of nontradable candidate risk factors considered in the baseline analysis, and hence make the simulation exercise informative.

VI.2.2 Simulation Accuracy

Next, we simulate the asset returns and the noisy factors as detailed in Section VI.1. To begin with, we assess the accuracy of the simulated asset returns by comparing the moments of the observed returns with the means of the moments of the simulated returns. Specifically, Figure A11 shows the averages and standard deviations of the test-asset returns, as well as the assets' Sharpe ratios. For the simulated returns to match the observed returns, all data points must lie on the 45 degree line. This seems to be largely the case as most assets display small deviations from the 45 degree line.

Therefore, this analysis proves the adequateness of the DGP used in simulation and also confirms that the four-factor Z-model spans almost entirely the information in the asset returns. Finally, we also note that the simulated data match closely also the cross-sectional standard deviation of the average test asset returns. In fact, the mean of the cross-sectional standard deviation of the simulated average returns is 0.0194, and that of the observed average returns is only slightly lower, 0.0173. Having established the accuracy of the simulations, we can now turn to evaluate the performance of the three-pass estimator on the simulated data.

VI.2.3 Factor Structure

Before turning to the candidate factors' risk-premium estimates, we assess the factor structure of the simulated data. First, we report in Table A15 the estimated number of factors, using the GX and O tests, as in the main analysis. In line with the GX's simulations performed with $N = 50$, we find that the estimators cannot recover the true number of factors exactly. In particular, the estimators tend to detect a lower number of factors than the true one. GX argue that, although their estimator is consistent, this 'underestimate' can happen in particular when N is small, and possibly due to arbitrary choices of the tuning parameters.

Here, based on our empirical evidence, we add that the problem of underestimating the true number of factors is less evident when using RP-PCA with reasonably high weights. In fact, the tests detect a higher number of factors when the factors are extracted using RP-PCA instead of PCA, which is consistent with the fact that RP-PCA enhances the "signal strength" of the factors and better separates factors with high Sharpe ratios from time-series factors. However, and not surprisingly, even the RP-PCA estimators do not detect $\hat{F}_4$, the "Value" factor, when extracted from observed data.[49] Hence, the results of the tests need to be interpreted with caution especially in finite samples and in the presence of weak factors with small premia. GX argue that in empirical applications it is important to select slightly more factors than indicated by the test to ensure the robustness of the estimates. Thus, in finite samples, the GX and O estimators seem to provide indications on the minimum number of factors entering the SDF. It is therefore important to complement this information with other metrics and considerations.

Table A15, Panel B, reports the generalized correlations (GC) of the true latent factors with the simulated latent factors. As stated in Section V.2, the GC was first proposed by Bai and Ng (2006) and has been employed recently by Lettau and Pelger (2020b) to assess the stability of the factor structure recovered with RP-PCA. Here, we note that the generalized correlations are natural statistics to look at also in the simulations. This is because they help us answer the question: can we recover the true factor structure in finite samples? Importantly, generalized correlations account for the fact that a factor model is only identified up to invertible linear transformations. Specifically, they measure the correlations between the true factors driving the DGP and the simulated factors after rotating them appropriately. A high qth generalized correlation suggests that the simulated factors have at least q common factors with the true factors. Therefore, generalized correlations help us quantify the degree to which a linear combination of the mth simulated factors ($\hat{F}^{(m)}$) replicate some or all of the factors in $\hat{F}$. Put simply, for a given run m they inform us on how close the true factor space and a realization of the simulated factor space are to each other.

We find that the median generalized correlations of the first four simulated factors with the true factors are high. For $\omega = 20$ they are all at least 97% and the confidence intervals are tiny, suggesting that the factors driving the observed returns and the simulated factors have at least four common factors. This in turn confirms that the simulated factors provide a good approximation of the true factors. This finding holds regardless of the RP-weigth, which is somewhat expected as correlations abstract from the role of the means of the factors.

[49] We conjecture that, while this factor is not a time-series factor, it has a much smaller risk premium compared to the other two weak factors, $\hat{F}_2$ and $\hat{F}_3$. As a result, even RP-PCA is not able to separate $\hat{F}_4$ from $\hat{F}_5$ and $\hat{F}_6$ in a sample with our N and T.

To complete this part of the analysis, we assess the maximal Sharpe ratios implied in the simulated data (Table A16). We find that, when the factors are extracted using RP-PCA, at least three factors are needed to recover the true Sharpe ratios. In fact, the two-factor model severely underestimates the Sharpe ratios, and this tendency only slightly reduces for higher RP-weights. Conversely, using more than three factors does not lead to statistically significant increases in the Sharpe ratios; this finding aligns well with the properties of the observed SDF. In fact, recall that the fourth latent factor (i.e., "Value") has a small impact on the model Sharpe ratio, and the remaining two factors are time-series factors. Their counterparts estimated on simulated data behave similarly. Finally, while we find that RP-PCA yields slightly higher point estimates using more than three factors, PCA underestimates the true Sharpe ratio for SDFs of any dimension. The absolute biases and RMSEs of PCA become comparable to those of RP-PCA only when all six latent factors are included in the SDF.

Overall, the above analysis suggests that, despite the relatively small sample size N, we can to a large extent recover in simulation the true factor structure of the return data. The simulations confirm the tendency, previously documented by GX, of the tests to underestimate the true number of factors in finite samples. We find that this problem is attenuated (albeit not eliminated) by estimating the factors via RP-PCA, as the fourth latent pricing factor remains hard to detect. However, while a good recovery of the factor structure is per se important, the ultimate object of interest is the ability of the three-pass estimator to recover the risk-premium estimates and standard errors of the candidate factors in finite samples, to which we turn next.

VI.2.4 Three-pass Estimator

In this section we assess the performance of the three-pass estimator, and hence its reliability, in finite samples. The estimation results are reported in Table A17. For each RP-weight ($\omega = -1, 10, 20, 50$) we estimate the risk premia using models that include an increasing number of factors ($k = 2, \ldots, 6$). For each combination of ω and k, we then report the biases ($Bias$, left panels), i.e., the estimate minus the true value, and the root-mean-square errors ($RMSE$, right panels) for the four tradable Z-factors and the three nontradable G-factors. Recall that the true number of latent factors is four and the RP-weight used in the factor extraction is 20.

Using $\varphi(\hat{F}_{1-4})$ and $\omega = 20$ (hence the true number of factors and RP-weight), we find that the Z-factor biases and RMSEs are small, suggesting a very good performance of the three-pass estimator also when N is relatively small. For most tradable factors, the estimator yields small biases and RMSEs already for three-factor models. In fact, we can easily recover the true Dollar factor's risk premium using a parsimonious single-factor model (not reported), while the second and third factors are particularly important to price accurately the Carry and ST Mom factors. By adding a fourth factor to the SDF, however, the gains in the pricing accuracy of the Value factor are sizable (the bias further reduces, albeit slightly, using a five-factor model).

A four-factor model with $\omega = 20$ turns out to price well also the G-factors. However, two observations are in order. First, for *gvol* and *icap*, both the biases and RMSEs further improve by including the fourth factor. Thus, although $\hat{F}_4$ has a low Sharpe ratio compared to the other two weak factors, $\hat{F}_2$ and $\hat{F}_3$, its omission can also lead to distorted risk-premium estimates. Second, for all nontradable factors, the absolute biases further reduce using a five-factor model, but the gains in terms of RMSEs are tiny or absent. Therefore, while the model already achieves an accurate pricing performance with four latent factors, adding an extra factor or two can lead to further reductions in the biases for some candidate factors, without being particularly harmful for the other factors. In contrast, it is clear that underestimating the number of relevant factors can lead to severe biases in the estimates. For example, with a two-factor model we severely overestimate (in absolute terms) the risk premia of *gvol* and *icap*. Thus, we find that including a lower number of factors than the optimal SDF is much more problematic than using a higher number of factors in the risk-premium estimation.

We then turn to the sensitivity of the estimates to the RP-weight. Starting from the comparison between PCA ($\omega = -1$) and RP-PCA with baseline RP-weight ($\omega = 20$), the difference is most apparent when examining the bias (*Panel A*). Using a four-factor model the differences in bias are substantial (much larger in absolute value for PCA than RP-PCA), and they remain high also for models with more than four factors. We observe no clear pattern in the RMSE differences using the two estimators (*Panel B*).[50] The results for RP-PCA are qualitatively identical when setting the RP-weight to either 10 or 50 as long as the SDF includes at least three factors (see *Panel II:* $\omega = 10$ and *Panel IV:* $\omega = 50$, respectively).

Next, to complete the analysis, we plot the histograms of the standardized three-pass bias estimates, using the asymptotic standard errors. In this way, we can verify the central-limit results of the estimator in the presence of small N. For the results to hold, the histograms should match the standard normal distribution. Figures A12–A15 present the results for the tradable factors. We find that, using four-factor models with $\omega = 20$, the bars essentially overlap with the standard normal distribution for all factors but Value. Specifically, the histograms of Carry and ST Mom show small deviations from the standard normal distribution for models that include at least three factors, while for the Dollar factor a one-factor model is enough. While Value benefits substantially from the inclusion of $\hat{F}_4$, the central-limit result is verified using a five-factor model. Finally, Figures A16–A18 show that the central-limit results are also confirmed for the nontradable factors. Using four-factor models, all factors' estimates display small deviations from the standard normal distributions. For *gvol* such deviations become imperceptible using a five-factor model. Moreover, in line with the above evidence, for both tradable and nontradable factors we obtain similar results when extracting the latent factors with ω equal to 10 and 50. On the contrary, for some factors (e.g., the HML factors, and *icap*) and SDFs, the PCA histograms tend to deviate substantially from the standard normal distribution.

Finally, we note that the fact that we can recover the true risk premia shows the rotation-invariance property (which is a general property that underlies the three-pass estimator; see Sections 2.1.2 and 4.2) in our sample. This is because the DGP of the asset returns is driven by the de-noised Z-factors, but the risk premia are estimated via the three-pass method and, hence, as linear combinations of the latent factors, F-factors. Put simply, the returns and the candidate factors are simulated using the reduced-form Z-factor model, and yet the risk premia are recovered using the reduced-form F-factor model.

Therefore, we conclude that the augmented three-pass estimator performs well in finite samples that

[50] As said before, here the main objective is not to compare RP-PCA with PCA, also because we use RP-PCA with $\omega = 20$ to de-noise the Z-factors that drive the DGP. However, we note that the biases and RMSEs are in general smaller if we simulate the DGP with PCA and estimate the model with RP-PCA than (as we do here) if we simulate the DGP with RP-PCA and estimate the model with PCA (results available upon request).

match the properties of our cross section of FX returns. That is, we can recover the true risk premia of both tradable and nontradable factors also when these factors are measured with noise. Hence, the simulation analysis shows that the three-pass method proves to be a *reliable* estimator also for FX risk premia. At the same time, we find that for most factors we can recover the true risk premia even using parsimonious factor models (especially for higher RP-weights), but for some factors it is beneficial to use models that include more latent factors. Using five-factor models (thus an additional factor than the true model), we can verify the central-limit results for all factors. Thus, our simulation results are in line with those of GX and, notably, are obtained in a setting that matches the properties of FX portfolio returns. As a result, they lend additional support to the validity of the (augmented) three-pass estimator.

VI.2.5 Two-pass Estimator

In what follows, we assess the finite-sample performance of the two-pass estimator. We consider SDFs of expanding dimensions by adding Dollar, Carry, ST Mom, and Value control factors one at a time along with the selected nontradable candidate risk factor (i.e., gvol, icap, or ipus). In this way, the omitted-variable problem can manifest to different degrees, as the true set of control factors should consist of all four factors.

In Panel A of Table A18 the controls are some noisy versions of the de-noised Z-factors driving the DGP, so that the measurement-error problem adds to that associated with omitting relevant factors from the SDF. It is evident that the two-pass estimator recovers the true risk premia only when the SDF includes all relevant control factors. In fact, the biases are small for all three candidate factors. On the contrary, the omission of relevant factors can severely distort the risk-premium estimates, although the biases do not necessarily reduce as more relevant controls are included. While the absolute biases of ipus monotonically reduce as more controls are added to the SDF, this is not the case for gvol or icap. For example, the inclusion of $\hat{F}_3$ widens the absolute bias for gvol, while it reduces it for icap. Nevertheless, when moving from the two- to the three-factor SDFs, the RMSEs drop for all candidate factors. Moreover, Panel B shows that the biases remain sizable even absent random noise around the controls. However, the RMSEs are (slightly) smaller for well-specified SDFs with de-noised controls. This is because the presence of *random* noise around the control factors widens the confidence intervals of the candidate factors' risk-premium estimates.

We then turn to the tradable factors' risk-premium estimates, which make even more apparent the impact of the measurement-error problem on the risk-premium estimates. On the one hand, the random noise around the tradable factors can cloud the impact of the omission of relevant factors on the tradable factor biases. On the other hand, when tradable factors are de-noised, we can appreciate fully the omitted-variable problem. In fact, Panel B shows that when factors are measured without noise we get closer to the true tradable factor risk premia as more relevant controls are added to the SDF. By contrast, this is not necessarily the case when factors are measured with *random* noise, due to the non-trivial interplay between the omission of relevant factors and the presence of noise around the factors.

Finally, Figures A19-A21 plot the histograms of the standardized two-pass bias estimates of the gvol, icap, and ipus nontradable factors, respectively, using Shanken (1992) standard errors. We find that the standardized histograms deviate substantially from the standard normal when the SDFs omit relevant factors. This result holds irrespective of whether the controls are de-noised or not. Then, the comparison of *Panels I* and *II* reveals that the standardized histograms of the tradable factors can deviate substantially from the standard normal if the controls are measured with noise, also when all relevant factors are included in the SDF. Hence, both the omitted-variable and measurement-error problems can lead to severely distorted two-pass estimates.

Summing up, the two-pass analysis on simulated data shows that the omitted-variable problem can be material, leading to distorted risk-premium estimates. The two-pass estimator delivers the true premia only if the SDF is correctly specified. In some cases, even if none of the relevant factors are omitted from the SDF, but some of the factors are measured with noise, we cannot recover the true risk premia. Therefore, both the omitted-variable and measurement-error problems manifest in simulation, making the use of the three-pass estimator *desirable* in the estimation of FX risk premia, as it performs well also in finite samples.

VI.3 Final Remarks

Taken together, the simulation analysis uncovers a number of important insights, which can be summarized as follows. To begin with, the simulated returns generated from a four-factor reduced-form model match the properties of the observed FX returns. This makes the simulation analysis reliable, but it also implies that this four-factor model spans the entire space of asset returns. Moreover, we find that RP-PCA achieves a good recovery of the true factor structure on simulated data. Similar to GX, we find a tendency of the tests to underestimate the true number of factors. Such tendency is attenuated when the factors are extracted via RP-PCA, but is not eliminated (the fourth latent factor is not detected by any of the tests). However, the analysis of the generalized correlations and Sharpe ratios point to an accurate recovery of the factor structure.

We then turn to the estimation of the candidate factor risk premia (our ultimate goal), which establishes two important results. First, the augmented three-pass estimator recovers the true factor risk premia in simulation. Second, the central-limit results are verified. These results show that the three-pass estimator is highly reliable, also in finite samples tailored to the properties of our FX asset returns. At the same time, they show the rotation invariance of the risk premia in our setting. Moreover, throughout the analysis, it emerges that omitting relevant latent factors is far more harmful than adding extra factors to the SDF. Actually, for some candidate factors it turns out to be beneficial to add an extra factor to the SDF, without any tangible adverse effect on the other candidate factors.

Finally, we find that the two-pass estimator can lead to substantially distorted estimates if some of the relevant factors are omitted from the SDF and/or are measured with noise (the interplay between these two sources of bias is material, and yet is non-trivial, so that the omitted-variable absolute bias might not reduce as more relevant factors are added to the SDF). Overall, based on the simulations, we can conclude that the three-pass estimator is both *reliable* and *desirable* in the estimation of currency risk premia.

Table A14: Three-Pass Estimates of Tradable and Nontradable Factor Risk Premia

The table presents the risk-premium estimates for selected tradable (Panel T) and nontradable (Panel NT) factors. Column *mf* shows the model-free estimates of the risk premia (λ), that is the factor's return sample average, along with the Newey-West standard errors (se) and p-values (pval), and the factor Sharpe ratio (SR). The remaining columns report the estimates from the (augmented) three-pass procedure for SDFs of different dimensions ($\varphi(\hat{F}_{1-k})$ with $k = 2, \ldots, 6$). In left panels the latent factors are estimated using the RP-PCA method with baseline RP-weight (i.e., $\omega = 20$), while in right panels using PCA (i.e., RP-PCA with $\omega = -1$) for comparison. The underlying test assets consist of the portfolios from the nine investment strategies ($N = 46$). We report the RP-PCA risk-premium estimates with the asymptotic standard errors, and p-values, and the de-noised factor Sharpe ratios. Tradable factors are Dollar, Carry, ST Mom, and Value; their de-noised version from $\varphi(\hat{F}_{1-4})$ drive the DGP in the simulations. The nontradable factors *gvol*, *icap*, and *ipus* are the AR(1) innovations to the global volatility of Menkhoff et al. (2012a), the intermediaries' capital ratio of He et al. (2017), and the U.S. industrial production, respectively (see Tables A5-A7 in the IA). The sample period is from 11/1983 to 12/2017 at monthly frequency ($T = 410$).

		Panel T: Tradable Factors									
		Panel T.I: RP-PCA ($\omega = 20$)					Panel T.II: PCA ($\omega = -1$)				
		T1.I Dollar					T1.II Dollar				
	mf	$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$	$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$
λ	0.026	0.026	0.025	0.026	0.026	0.026	0.024	0.024	0.024	0.025	0.026
(se)	0.012	0.013	0.013	0.013	0.013	0.013	0.013	0.013	0.013	0.013	0.013
<pval>	0.044	0.055	0.058	0.050	0.052	0.052	0.070	0.072	0.073	0.056	0.054
SR	0.368	0.361	0.355	0.367	0.365	0.366	0.339	0.337	0.337	0.359	0.363
		T2.I Carry					T2.II Carry				
	mf	$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$	$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$
λ	0.073	0.085	0.063	0.064	0.062	0.063	0.029	0.040	0.040	0.048	0.049
(se)	0.014	0.013	0.014	0.014	0.014	0.014	0.013	0.013	0.014	0.014	0.014
<pval>	0.000	0.000	0.000	0.000	0.000	0.000	0.031	0.005	0.006	0.001	0.001
SR	0.865	1.211	0.842	0.854	0.817	0.817	0.405	0.544	0.535	0.642	0.649
		T3.I ST Mom					T3.II ST Mom				
	mf	$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$	$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$
λ	0.069	0.019	0.079	0.082	0.083	0.082	-0.011	0.025	0.024	0.071	0.067
(se)	0.015	0.011	0.013	0.012	0.012	0.013	0.004	0.011	0.011	0.013	0.014
<pval>	0.000	0.083	0.000	0.000	0.000	0.000	0.010	0.023	0.028	0.000	0.000
SR	0.770	0.894	1.212	1.249	1.224	1.039	0.481	0.470	0.444	0.945	0.819
		T4.I Value					T4.II Value				
	mf	$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$	$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$
λ	0.033	0.006	-0.007	0.022	0.023	0.022	0.002	-0.028	-0.029	0.007	0.007
(se)	0.013	0.007	0.008	0.010	0.010	0.010	0.004	0.008	0.008	0.010	0.010
<pval>	0.014	0.377	0.421	0.034	0.034	0.038	0.630	0.001	0.001	0.468	0.485
SR	0.446	0.594	0.390	0.369	0.375	0.370	0.134	0.663	0.661	0.127	0.123
		Panel NT: Nontradable Factors									
		Panel NT.I: RP-PCA ($\omega = 20$)					Panel NT.II: PCA ($\omega = -1$)				
		NT1.I gvol					NT1.II gvol				
	mf	$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$	$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$
λ	–	-1.323	-1.048	-0.759	-0.769	-0.791	-0.540	-0.968	-0.962	-0.716	-0.773
(se)	–	0.317	0.302	0.360	0.358	0.385	0.207	0.262	0.259	0.296	0.315
<pval>	–	0.000	0.001	0.041	0.037	0.046	0.012	0.001	0.001	0.020	0.018
SR	–	1.198	0.920	0.595	0.598	0.567	0.488	0.776	0.759	0.553	0.553
		NT2.I icap					NT2.II icap				
	mf	$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$	$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$
λ	–	1.197	0.612	0.756	0.790	0.801	0.518	0.390	0.377	0.729	0.760
(se)	–	0.335	0.386	0.430	0.423	0.437	0.245	0.259	0.253	0.371	0.387
<pval>	–	0.001	0.120	0.086	0.068	0.074	0.040	0.140	0.144	0.056	0.056
SR	–	1.226	0.534	0.641	0.615	0.612	0.456	0.339	0.304	0.562	0.572
		NT3.I ipus					NT3.II ipus				
	mf	$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$	$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$
λ	–	0.163	0.227	0.159	0.166	0.168	-0.013	0.141	0.137	0.165	0.170
(se)	–	0.354	0.378	0.385	0.385	0.384	0.143	0.228	0.228	0.291	0.294
<pval>	–	0.649	0.551	0.681	0.667	0.664	0.929	0.542	0.550	0.575	0.565
SR	–	0.637	0.860	0.539	0.529	0.530	0.061	0.477	0.430	0.514	0.525

Table A15: Factor Structure in Simulation

The table presents statistics about the factor structure of the simulated data. Panel A.I and A.II report the percentiles (10*th*, 50*th*, and 90*th*) and standard deviations (SD) of the estimates of the number of factors as suggested by the tests of Onatski (2010; O), and Giglio and Xiu (2021; GX), respectively. Panel B reports the percentiles and standard deviations of the first four generalized correlations (GC). We implement the GX and O tests and compute the GC for factors extracted using RP-PCA with different RP-weights ($\omega = -1, 10, 20, 50$). The true data-generating process (DGP) has four factors, and the parameters are calibrated based on the four de-noised tradable factors (Dollar, Carry, ST Mom, and Value). Hence, the true number of factors is four. To remove noise from the factors, we applied the three-pass procedure with $\omega = 20$ and four latent factors ($\varphi(\hat{F}_{1-4})$) to the panel of observed FX portfolio returns. The results are based on $M = 10,000$ artificial Monte Carlo realizations with $N = 46$ and $T = 410$, which match the cross-section and time-series dimensions of FX portfolio returns.

	Panel A: Number of Factors							
	Panel A.I: O Test				Panel A.II: GX Test			
	10*th*	50*th*	90*th*	SD	10*th*	50*th*	90*th*	SD
ω=−1	2	2	2	0.27	2	2	2	0.23
ω=10	2	3	3	0.50	3	3	3	0.11
ω=20	3	3	3	0.40	3	3	3	0.09
ω=50	3	3	3	0.38	3	3	3	0.09
	Panel B: Generalized Correlations							
	Panel B.I: ω=−1				Panel B.II: ω=10			
	10*th*	50*th*	90*th*	SD	10*th*	50*th*	90*th*	SD
GC1	99.95	99.97	99.98	0.01	99.95	99.97	99.98	0.01
GC2	99.03	99.33	99.56	0.21	99.01	99.32	99.55	0.21
GC3	97.65	98.4	98.91	0.51	97.58	98.35	98.88	0.52
GC4	95.27	96.79	97.80	1.05	95.21	96.69	97.75	1.02
	Panel B.III: ω=20				Panel B.IV: ω=50			
	10*th*	50*th*	90*th*	SD	10*th*	50*th*	90*th*	SD
GC1	99.95	99.97	99.98	0.01	99.95	99.97	99.98	0.01
GC2	99.01	99.31	99.55	0.21	99.00	99.31	99.54	0.21
GC3	97.56	98.34	98.87	0.53	97.55	98.33	98.86	0.53
GC4	95.16	96.67	97.73	1.03	95.12	96.65	97.72	1.05

Table A16: Sharpe Ratios in Simulation

The table presents the estimates (*Avg.*), standard deviations (*SD*), biases (*Bias*), and root-mean-square errors (*RMSE*) of the maximal Sharpe ratio (SR) computed on the simulated data. The bias is given by the estimate minus the true value of the SR. We refer to Section 2.1 and Table 1 for a detailed description of the maximal SR. Here, we note that we compute the maximal SR for each artificial realization using SDFs of different dimensions ($\varphi(\hat{F}_{1-k})$ with $k = 1, \ldots, 6$), whereby the latent factors are estimated with different RP-weights ($\omega = -1, 10, 20, 50$). The true data-generating process (DGP) has four factors, and the parameters are calibrated based on the four denoised tradable factors (Dollar, Carry, ST Mom, and Value). To remove noise from the factors, we applied the three-pass procedure with RP-weight $\omega = 20$ and four latent factors ($\varphi(\hat{F}_{1-4})$) to the panel of observed FX portfolio returns. The true maximal SR is 0.476. The results are based on $M = 10,000$ artificial Monte Carlo realizations with $N = 46$ and $T = 410$, which match the cross-section and time-series dimensions of the FX portfolio returns.

	Panel A: Avg.				Panel B: SD			
	ω=-1	ω=10	ω=20	ω=50	ω=-1	ω=10	ω=20	ω=50
$\varphi(\hat{F}_1)$	0.098	0.103	0.107	0.116	0.048	0.050	0.051	0.051
$\varphi(\hat{F}_{1-2})$	0.154	0.351	0.417	0.450	0.049	0.075	0.061	0.054
$\varphi(\hat{F}_{1-3})$	0.263	0.482	0.494	0.502	0.059	0.052	0.051	0.049
$\varphi(\hat{F}_{1-4})$	0.302	0.500	0.507	0.513	0.063	0.052	0.051	0.050
$\varphi(\hat{F}_{1-5})$	0.403	0.508	0.514	0.518	0.063	0.051	0.051	0.050
$\varphi(\hat{F}_{1-6})$	0.438	0.511	0.516	0.520	0.053	0.051	0.051	0.050
	Panel C: Bias				Panel D: RMSE			
	ω=-1	ω=10	ω=20	ω=50	ω=-1	ω=10	ω=20	ω=50
$\varphi(\hat{F}_1)$	-0.378	-0.373	-0.369	-0.360	0.382	0.377	0.373	0.364
$\varphi(\hat{F}_{1-2})$	-0.322	-0.126	-0.059	-0.027	0.326	0.146	0.085	0.060
$\varphi(\hat{F}_{1-3})$	-0.213	0.006	0.018	0.025	0.221	0.053	0.054	0.055
$\varphi(\hat{F}_{1-4})$	-0.175	0.024	0.031	0.036	0.186	0.057	0.060	0.062
$\varphi(\hat{F}_{1-5})$	-0.074	0.032	0.038	0.042	0.097	0.061	0.063	0.065
$\varphi(\hat{F}_{1-6})$	-0.038	0.035	0.040	0.044	0.065	0.062	0.065	0.066

Table A17: Three-Pass Simulation Results

The table presents the bias (*Bias*) and the root-mean-square error (*RMSE*) of the risk-premium estimates using the (augmented) three-pass estimator with $k = 2, \ldots, 6$ latent factors ($\varphi(\hat{F}_{1-k})$) and selected RP-weights ($\omega = -1, 10, 20, 50$). The bias is computed as the estimate minus the true risk premium. The true data-generating process (DGP) has four factors, and the parameters are calibrated based on the four de-noised tradable factors (Dollar, Carry, ST Mom, and Value). To remove noise from the factors, we applied the three-pass procedure with $\omega = 20$ and four latent factors ($\varphi(\hat{F}_{1-4})$) to the panel of FX portfolio returns. The true (annualized) risk premia of the noisy yet observed tradable (Dollar, Carry, ST Mom, and Value) and nontradable (*gvol*, *icap*, and *ipus*) factors are reported in Table A14, columns ($\varphi(\hat{F}_{1-4})$) of Panels T.I and NT.I, respectively. The results are based on $M = 10,000$ artificial Monte Carlo realizations with $N = 46$ and $T = 410$, which match the cross-section and time-series dimensions of our baseline sample of FX portfolio returns.

	Panel A: Bias						Panel B: RMSE				
	Panel A.I: ω=-1						Panel B.I: ω=-1				
	$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$		$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$
Dol	-0.002	-0.002	-0.002	-0.001	-0.001	Dol	0.013	0.013	0.013	0.012	0.012
Carry	-0.035	-0.025	-0.022	-0.011	-0.007	Carry	0.037	0.028	0.026	0.018	0.015
ST Mom	-0.093	-0.063	-0.054	-0.027	-0.017	ST Mom	0.093	0.065	0.056	0.031	0.021
Value	-0.020	-0.047	-0.043	-0.019	-0.011	Value	0.020	0.047	0.045	0.024	0.016
gvol	0.216	-0.180	-0.185	-0.063	-0.022	gvol	0.306	0.327	0.337	0.324	0.328
icap	-0.238	-0.351	-0.318	-0.150	-0.090	icap	0.323	0.431	0.418	0.349	0.335
ipus	-0.172	-0.036	-0.022	-0.017	-0.015	ipus	0.200	0.173	0.192	0.247	0.268
	Panel A.II: ω=10						Panel B.II: ω=10				
	$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$		$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$
Dol	-0.001	-0.001	0.000	0.000	0.000	Dol	0.013	0.012	0.012	0.012	0.012
Carry	0.016	-0.001	0.000	0.000	0.000	Carry	0.026	0.014	0.013	0.013	0.013
ST Mom	-0.069	-0.003	-0.002	0.000	0.000	ST Mom	0.072	0.013	0.012	0.012	0.012
Value	-0.016	-0.028	-0.009	-0.002	-0.001	Value	0.018	0.031	0.018	0.012	0.012
gvol	-0.489	-0.294	-0.096	-0.021	-0.013	gvol	0.598	0.467	0.387	0.361	0.361
icap	0.382	-0.140	-0.046	-0.006	-0.003	icap	0.521	0.372	0.356	0.352	0.353
ipus	-0.019	0.067	0.020	0.004	0.003	ipus	0.227	0.299	0.303	0.307	0.309
	Panel A.III: ω=20						Panel B.III: ω=20				
	$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$		$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$
Dol	0.000	-0.001	0.000	0.000	0.000	Dol	0.013	0.012	0.012	0.012	0.012
Carry	0.025	0.000	0.000	0.001	0.001	Carry	0.031	0.014	0.013	0.013	0.013
ST Mom	-0.050	0.000	0.000	0.001	0.001	ST Mom	0.055	0.012	0.012	0.012	0.012
Value	-0.015	-0.024	-0.007	-0.001	0.000	Value	0.018	0.028	0.017	0.012	0.012
gvol	-0.597	-0.267	-0.084	-0.020	-0.013	gvol	0.689	0.455	0.385	0.363	0.363
icap	0.438	-0.118	-0.033	-0.001	0.002	icap	0.579	0.368	0.357	0.354	0.355
ipus	0.030	0.063	0.019	0.005	0.004	ipus	0.260	0.305	0.307	0.310	0.311
	Panel A.IV: ω=50						Panel B.IV: ω=50				
	$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$		$\varphi(\hat{F}_{1-2})$	$\varphi(\hat{F}_{1-3})$	$\varphi(\hat{F}_{1-4})$	$\varphi(\hat{F}_{1-5})$	$\varphi(\hat{F}_{1-6})$
Dol	0.000	-0.001	0.000	0.000	0.000	Dol	0.012	0.012	0.012	0.012	0.012
Carry	0.027	0.001	0.001	0.001	0.001	Carry	0.032	0.014	0.013	0.013	0.013
ST Mom	-0.038	0.002	0.001	0.002	0.002	ST Mom	0.044	0.012	0.012	0.012	0.012
Value	-0.014	-0.022	-0.006	-0.001	0.000	Value	0.018	0.026	0.016	0.012	0.012
gvol	-0.601	-0.248	-0.076	-0.019	-0.013	gvol	0.703	0.447	0.384	0.365	0.364
icap	0.421	-0.103	-0.025	0.003	0.005	icap	0.582	0.366	0.358	0.356	0.357
ipus	0.049	0.061	0.018	0.006	0.004	ipus	0.279	0.309	0.310	0.312	0.313

Table A18: Two-pass Simulation Results

The table presents the bias (*Bias*) and root-mean-square error (*RMSE*) of the risk premia estimates of selected nontradable factors (g) using the two-pass estimator. The bias is computed as the estimate minus the true risk premium. We consider SDFs of expanding dimensions whereby the four tradable control factors are added, along with the nontradable factor, one at a time. We consider the case of noisy controls ($\varphi([Z_{1-j}, g])$, for $j = 1, \ldots, 4$) in Panel A, and that of de-noised controls ($\varphi([\hat{Z}_{1-j}, g])$, for $j = 1, \ldots, 4$) in Panel B. The true data-generating process (DGP) has four factors, and the parameters are calibrated based on the four de-noised tradable factors (Dollar, Carry, ST Mom, and Value). To remove noise from the factors, we applied the three-pass procedure with $\omega = 20$ and four latent factors ($\varphi(\hat{F}_{1-4})$) to the panel of FX portfolio returns. The true (annualized) risk premia of the noisy yet observed tradable (Dollar, Carry, ST Mom, and Value) and nontradable (*gvol*, *icap*, and *ipus*) factors are provided in the 'True' column (and in column $\varphi(\hat{F}_{1-4})$ of Panels T.I and NT.I in Table A14, respectively). The results are based on $M = 10,000$ artificial Monte Carlo realizations with $N = 46$ and $T = 410$, which match the cross-section and time-series dimensions of our baseline sample of FX portfolio returns.

Panel A: Two-Pass with Noisy Controls									
		Panel AI: Bias				Panel AII: RMSE			
		Panel AI.G1: gvol				Panel AII.G1: gvol			
	True	$\varphi([Z_1, g])$	$\varphi([Z_{1-2}, g])$	$\varphi([Z_{1-3}, g])$	$\varphi([Z_{1-4}, g])$	$\varphi([Z_1, g])$	$\varphi([Z_{1-2}, g])$	$\varphi([Z_{1-3}, g])$	$\varphi([Z_{1-4}, g])$
Dol	0.0260	-0.0021	-0.0013	-0.0002	0.0001	0.0125	0.0124	0.0124	0.0124
Carry	0.0635	–	-0.0094	0.0154	0.0104	–	0.0203	0.0224	0.0195
ST Mom	0.0821	–	–	0.0321	0.0359	–	–	0.0364	0.0397
Value	0.0217	–	–	–	0.0101	–	–	–	0.0200
gvol	-0.7587	-4.1996	1.2328	3.2566	-0.0007	4.6426	4.4373	4.0290	2.6795
		Panel AI.G2: icap				Panel AII.G2: icap			
	True	$\varphi([Z_1, g])$	$\varphi([Z_{1-2}, g])$	$\varphi([Z_{1-3}, g])$	$\varphi([Z_{1-4}, g])$	$\varphi([Z_1, g])$	$\varphi([Z_{1-2}, g])$	$\varphi([Z_{1-3}, g])$	$\varphi([Z_{1-4}, g])$
Dol	0.0260	-0.0016	-0.0015	-0.0004	0.0001	0.0125	0.0125	0.0124	0.0124
Carry	0.0635	–	0.0009	0.0054	0.0103	–	0.0192	0.0175	0.0193
ST Mom	0.0821	–	–	0.0320	0.0360	–	–	0.0364	0.0398
Value	0.0217	–	–	–	0.0100	–	–	–	0.0195
icap	0.7565	3.4286	-6.6970	2.5930	0.0711	3.9838	8.3949	3.9431	2.7254
		Panel AI.G3: ipus				Panel AII.G3: ipus			
	True	$\varphi([Z_1, g])$	$\varphi([Z_{1-2}, g])$	$\varphi([Z_{1-3}, g])$	$\varphi([Z_{1-4}, g])$	$\varphi([Z_1, g])$	$\varphi([Z_{1-2}, g])$	$\varphi([Z_{1-3}, g])$	$\varphi([Z_{1-4}, g])$
Dol	0.0260	-0.0020	-0.0014	-0.0005	0.0001	0.0125	0.0125	0.0124	0.0124
Carry	0.0635	–	-0.0113	0.0095	0.0104	–	0.0201	0.0186	0.0190
ST Mom	0.0821	–	–	0.0296	0.0359	–	–	0.0344	0.0397
Value	0.0217	–	–	–	0.0101	–	–	–	0.0194
ipus	0.1592	7.0869	3.6691	-1.4835	0.0684	10.1581	8.2552	3.7566	2.9203

Panel B: Two-Pass with De-Noised Controls									
		Panel BI: Bias				Panel BII: RMSE			
		Panel BI.G1: gvol				Panel BII.G1: gvol			
	True	$\varphi([\hat{Z}_1, g])$	$\varphi([\hat{Z}_{1-2}, g])$	$\varphi([\hat{Z}_{1-3}, g])$	$\varphi([\hat{Z}_{1-4}, g])$	$\varphi([\hat{Z}_1, g])$	$\varphi([\hat{Z}_{1-2}, g])$	$\varphi([\hat{Z}_{1-3}, g])$	$\varphi([\hat{Z}_{1-4}, g])$
Dol	0.0260	-0.0022	-0.0015	-0.0004	-0.0001	0.0124	0.0123	0.0123	0.0123
Carry	0.0635	–	-0.0198	-0.0022	-0.0012	–	0.0238	0.0129	0.0127
ST Mom	0.0821	–	–	-0.0044	-0.0024	–	–	0.0122	0.0116
Value	0.0217	–	–	–	-0.0009	–	–	–	0.0100
gvol	-0.7587	-4.2043	1.2808	3.3776	-0.0090	4.6497	4.4686	4.1045	2.6131
		Panel BI.G2: icap				Panel BII.G2: icap			
	True	$\varphi([\hat{Z}_1, g])$	$\varphi([\hat{Z}_{1-2}, g])$	$\varphi([\hat{Z}_{1-3}, g])$	$\varphi([\hat{Z}_{1-4}, g])$	$\varphi([\hat{Z}_1, g])$	$\varphi([\hat{Z}_{1-2}, g])$	$\varphi([\hat{Z}_{1-3}, g])$	$\varphi([\hat{Z}_{1-4}, g])$
Dol	0.0260	-0.0018	-0.0017	-0.0006	-0.0001	0.0124	0.0124	0.0123	0.0123
Carry	0.0635	–	-0.0184	-0.0029	-0.0012	–	0.0226	0.0130	0.0127
ST Mom	0.0821	–	–	-0.0059	-0.0024	–	–	0.0130	0.0116
Value	0.0217	–	–	–	-0.0009	–	–	–	0.0100
icap	0.7565	3.4234	-6.8503	2.6437	-0.0284	3.9823	8.5134	3.9289	2.6379
		Panel BI.G3: ipus				Panel BII.G3: ipus			
	True	$\varphi([\hat{Z}_1, g])$	$\varphi([\hat{Z}_{1-2}, g])$	$\varphi([\hat{Z}_{1-3}, g])$	$\varphi([\hat{Z}_{1-4}, g])$	$\varphi([\hat{Z}_1, g])$	$\varphi([\hat{Z}_{1-2}, g])$	$\varphi([\hat{Z}_{1-3}, g])$	$\varphi([\hat{Z}_{1-4}, g])$
Dol	0.0260	-0.0022	-0.0016	-0.0007	-0.0001	0.0124	0.0124	0.0123	0.0123
Carry	0.0635	–	-0.0192	-0.0032	-0.0012	–	0.0233	0.0131	0.0127
ST Mom	0.0821	–	–	-0.0064	-0.0024	–	–	0.0132	0.0116
Value	0.0217	–	–	–	-0.0009	–	–	–	0.0100
ipus	0.1592	7.0664	3.6057	-1.6145	-0.0206	10.1536	8.2555	3.7628	2.7922

Figure A11: Observed vs. Simulated Test-Asset Return Moments

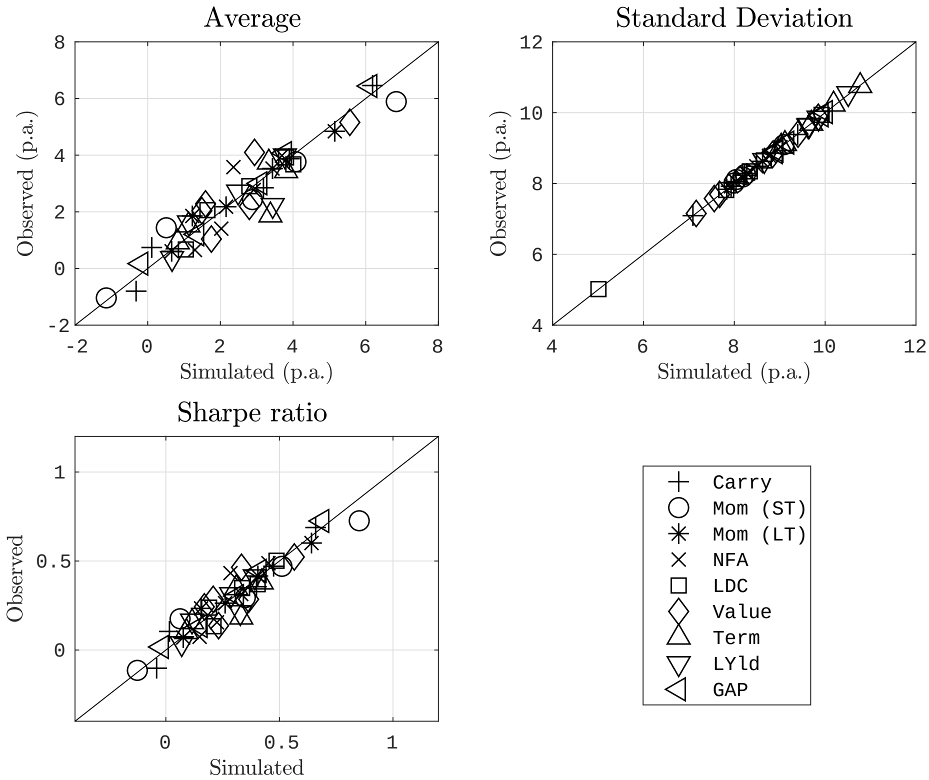

The figure shows the moments of the observed and simulated test-asset returns. In *Average*, we plot the time-series averages of the observed portfolio returns (*Observed*) against the means of the time-series averages of the simulated portfolio returns (*Simulated*; in percent, annualized); in *Standard Deviation*, we show the standard deviations of the observed returns against the means of the standard deviations of the simulated returns (in percent, annualized); in *Sharpe ratio*, we report the Sharpe ratios of the observed returns against the means of the Sharpe ratios of the simulated returns. The observed cross-sectional standard deviation of the average returns is 0.0173, and the mean of the cross-sectional standard deviation of the simulated average returns is 0.0194. The true data generating process (DGP) has four factors, and the parameters are calibrated based on the four de-noised tradable factors (Dollar, Carry, ST Mom, and Value). To de-noised the tradable factors, we used RP-PCA with $\omega = 20$ and four pricing factors. We simulate the models 10,000 times with $N = 46$ and $T = 410$ (see Section VI and Table A17).

Figure A12: Histograms of the Three-Pass Standardized Bias Estimates of the Dollar Factor in Simulation

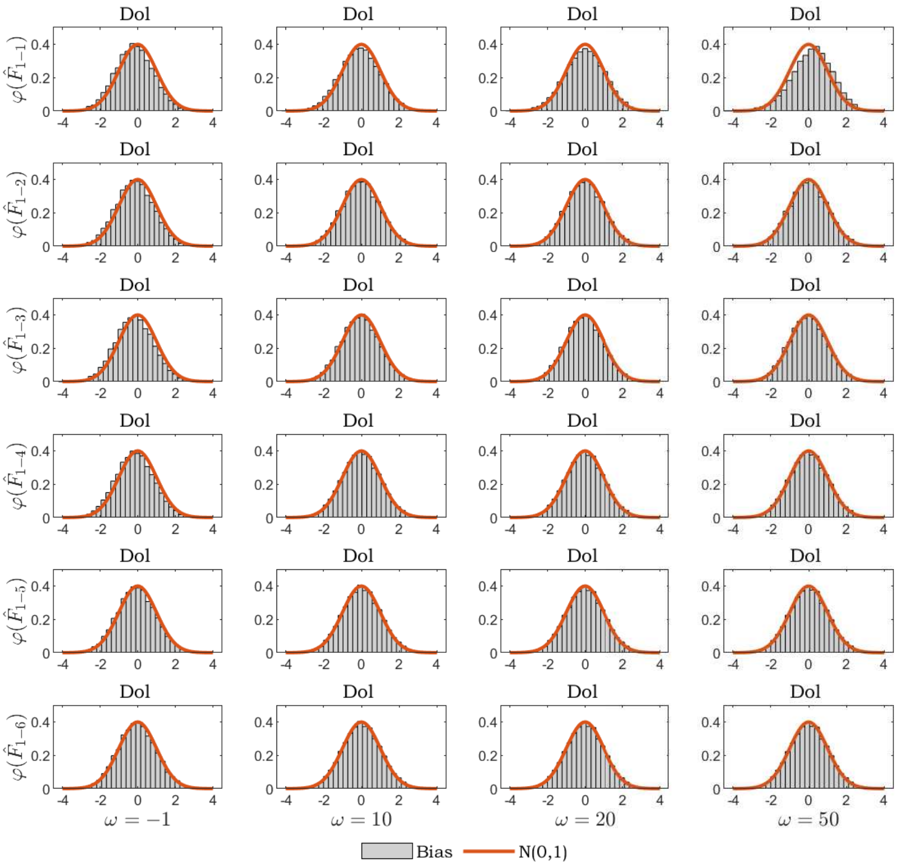

The figure shows the histograms of the three-pass standardized bias estimates of the Dollar factor using asymptotic standard errors. We implement the three-pass estimator using SDFs including an increasing number of pricing factors ($\varphi(\hat{F}_{1-k})$ for $k = 1, \ldots, 6$) and multiple RP-weights ($\omega = -1, 10, 20, 50$). The true data generating process (DGP) has four factors, and the parameters are calibrated based on the four de-noised tradable factors (Dollar, Carry, ST Mom, and Value). To de-noised the tradable factors, we used RP-PCA with $\omega = 20$ and four pricing factors. We simulate the models 10,000 times with $N = 46$ and $T = 410$ (see Section VI and Table A17).

Figure A13: Histograms of the Three-Pass Standardized Bias Estimates of the Carry Factor in Simulation

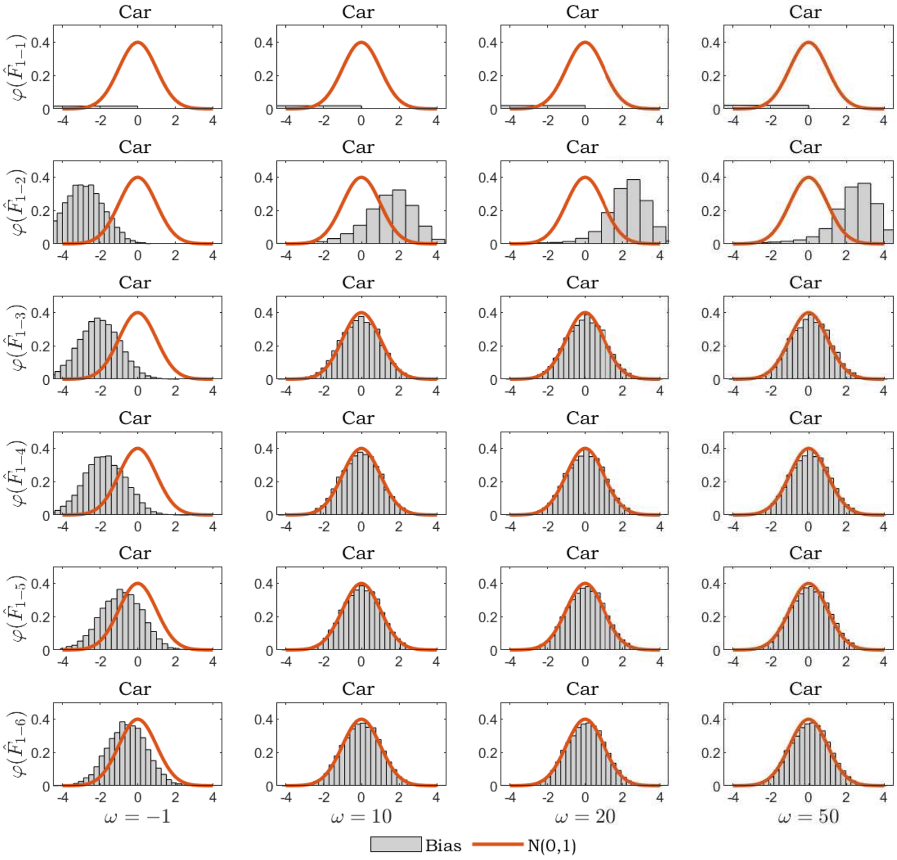

The figure shows the histograms of the three-pass standardized bias estimates of the Carry HML factor using asymptotic standard errors. We implement the three-pass estimator using SDFs including an increasing number of pricing factors ($\varphi(\hat{F}_{1-k})$ for $k = 1, \ldots, 6$) and multiple RP-weights ($\omega = -1, 10, 20, 50$). The true data generating process (DGP) has four factors, and the parameters are calibrated based on the four de-noised tradable factors (Dollar, Carry, ST Mom, and Value). To de-noised the tradable factors, we used RP-PCA with $\omega = 20$ and four pricing factors. We simulate the models 10,000 times with $N = 46$ and $T = 410$ (see Section VI and Table A17).

Figure A14: Histograms of the Three-Pass Standardized Bias Estimates of the Short-term Momentum Factor in Simulation

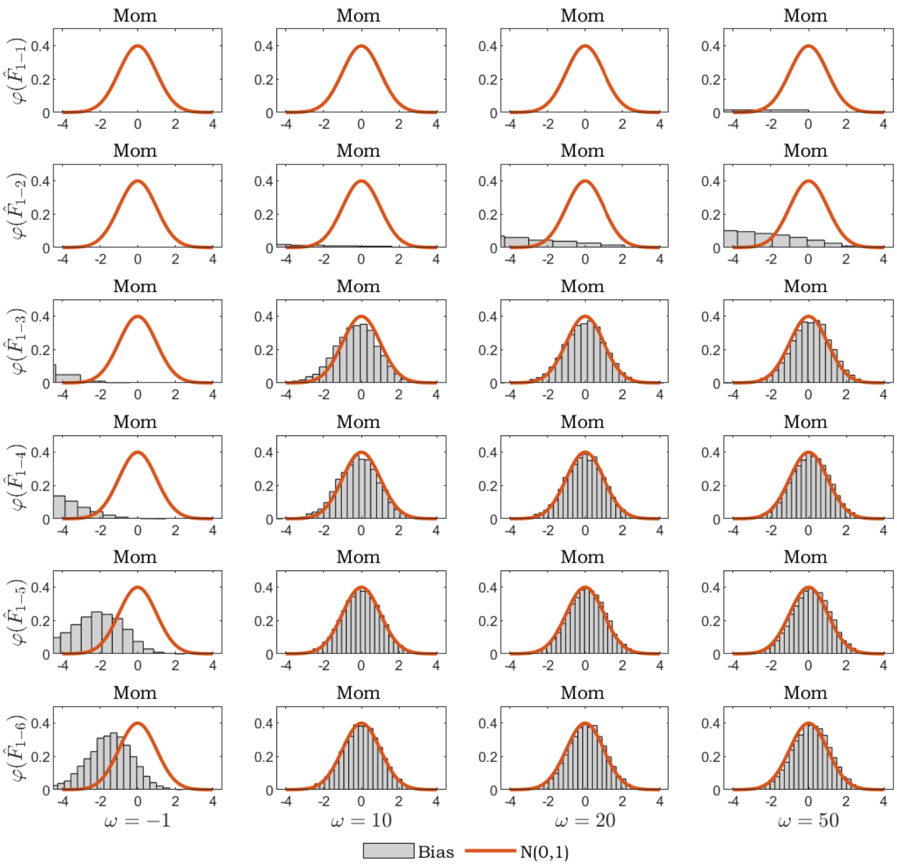

The figure shows the histograms of the three-pass standardized bias estimates of the ST Mom HML factor using asymptotic standard errors. We implement the three-pass estimator using SDFs including an increasing number of pricing factors ($\varphi(\hat{F}_{1-k})$ for $k = 1, \ldots, 6$) and multiple RP-weights ($\omega = -1, 10, 20, 50$). The true data generating process (DGP) has four factors, and the parameters are calibrated based on the four de-noised tradable factors (Dollar, Carry, ST Mom, and Value). To de-noised the tradable factors, we used RP-PCA with $\omega = 20$ and four pricing factors. We simulate the models 10,000 times with $N = 46$ and $T = 410$ (see Section VI and Table A17).

Figure A15: Histograms of the Three-Pass Standardized Bias Estimates of the Value Factor in Simulation

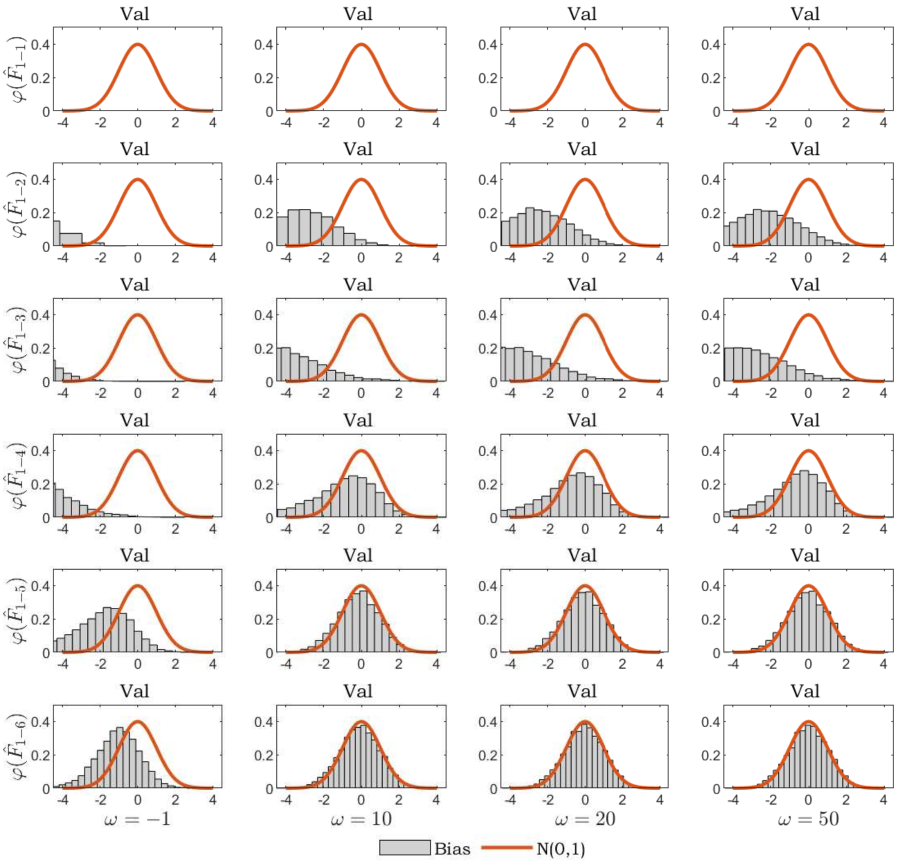

The figure shows the histograms of the three-pass standardized bias estimates of the Value HML factor using asymptotic standard errors. We implement the three-pass estimator using SDFs including an increasing number of pricing factors ($\varphi(\hat{F}_{1-k})$ for $k = 1, \ldots, 6$) and multiple RP-weights ($\omega = -1, 10, 20, 50$). The true data generating process (DGP) has four factors, and the parameters are calibrated based on the four de-noised tradable factors (Dollar, Carry, ST Mom, and Value). To de-noised the tradable factors, we used RP-PCA with $\omega = 20$ and four pricing factors. We simulate the models 10,000 times with $N = 46$ and $T = 410$ (see Section VI and Table A17).

Figure A16: Histograms of the Three-Pass Standardized Bias Estimates of the Global FX Volatility Factor in Simulation

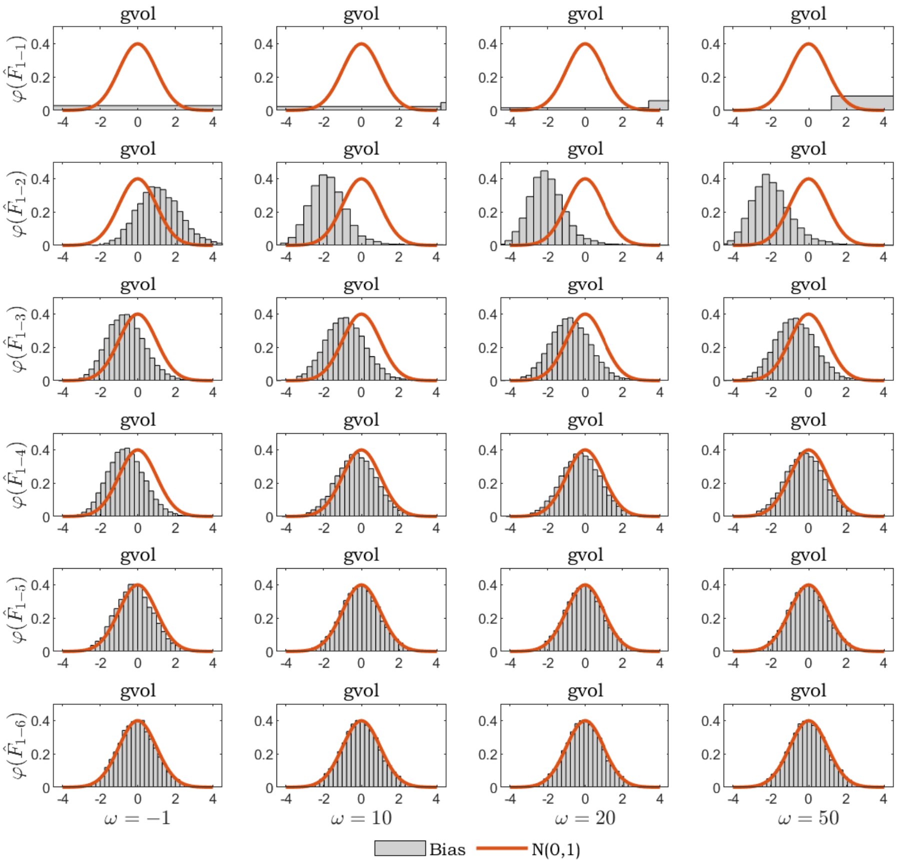

The figure shows the histograms of the three-pass standardized bias estimates of the global FX volatility factor of Menkhoff et al. (2012a) using asymptotic standard errors. We implement the three-pass estimator using SDFs including an increasing number of pricing factors ($\varphi(\hat{F}_{1-k})$ for $k = 1, \ldots, 6$) and multiple RP-weights ($\omega = -1, 10, 20, 50$). The true data generating process (DGP) has four factors, and the parameters are calibrated based on the four de-noised tradable factors (Dollar, Carry, ST Mom, and Value). To de-noised the tradable factors, we used RP-PCA with $\omega = 20$ and four pricing factors. We simulate the models 10,000 times with $N = 46$ and $T = 410$ (see Section VI and Table A17).

Figure A17: Histograms of the Three-Pass Standardized Bias Estimates of Intermediaries' Capital Ratio Factor in Simulation

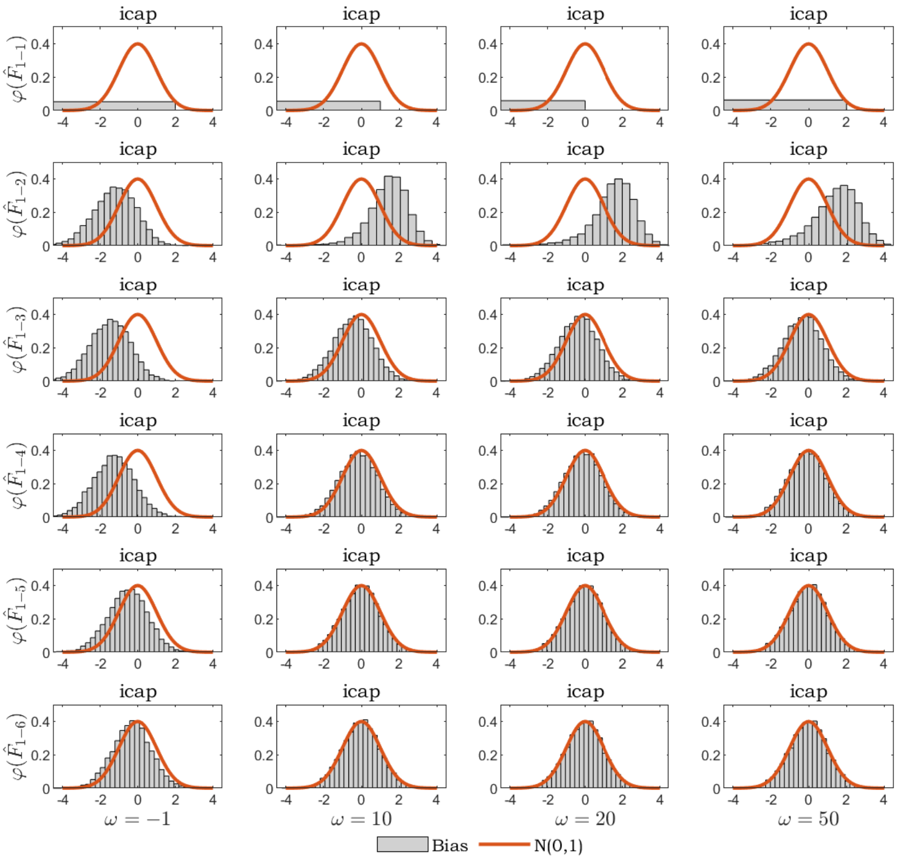

The figure shows the histograms of the three-pass standardized bias estimates of the intermediaries' capital ratio factor of He et al. (2017) using asymptotic standard errors. We implement the three-pass estimator using SDFs including an increasing number of pricing factors ($\varphi(\hat{F}_{1-k})$ for $k = 1, \dots, 6$) and multiple RP-weights ($\omega = -1, 10, 20, 50$). The true data generating process (DGP) has four factors, and the parameters are calibrated based on the four de-noised tradable factors (Dollar, Carry, ST Mom, and Value). To de-noised the tradable factors, we used RP-PCA with $\omega = 20$ and four pricing factors. We simulate the models 10,000 times with $N = 46$ and $T = 410$ (see Section VI and Table A17).

Figure A18: Histograms of the Three-Pass Standardized Bias Estimates of U.S. Industrial Production Factor in Simulation

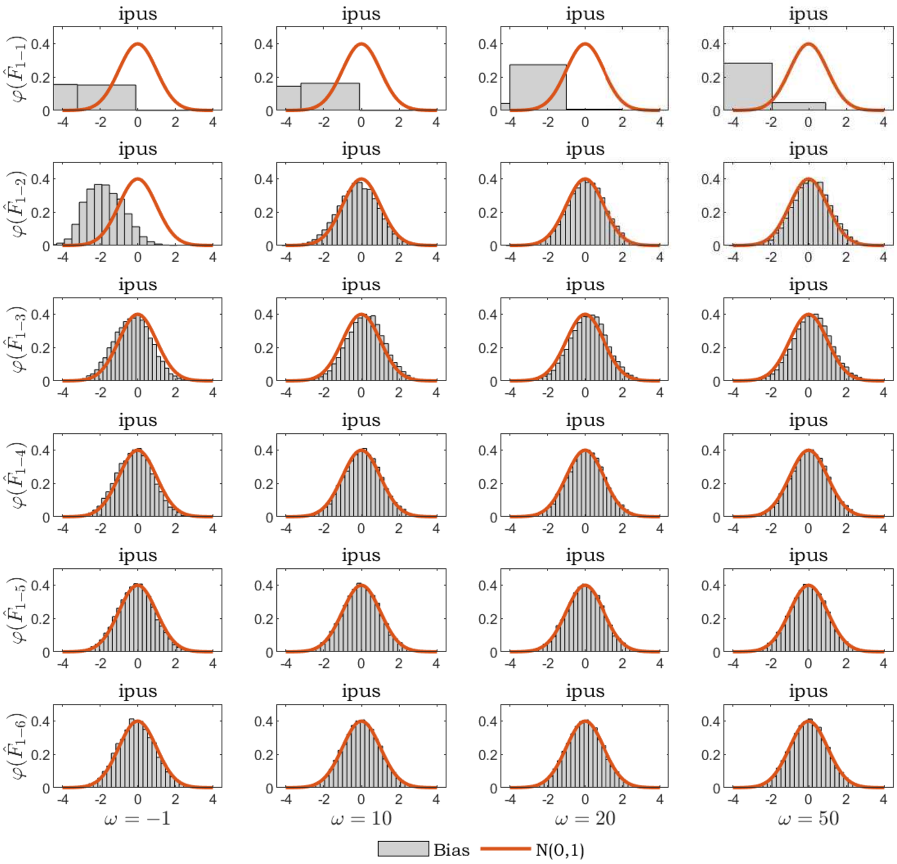

The figure shows the histograms of the three-pass standardized bias estimates of the U.S. industrial production factor using asymptotic standard errors. We implement the three-pass estimator using SDFs including an increasing number of pricing factors ($\varphi(\hat{F}_{1-k})$ for $k = 1, \ldots, 6$) and multiple RP-weights ($\omega = -1, 10, 20, 50$). The true data generating process (DGP) has four factors, and the parameters are calibrated based on the four de-noised tradable factors (Dollar, Carry, ST Mom, and Value). To de-noised the tradable factors, we used RP-PCA with $\omega = 20$ and four pricing factors. We simulate the models 10,000 times with $N = 46$ and $T = 410$ (see Section VI and Table A17).

Figure A19: Histograms of the Two-Pass Standardized Bias Estimates of the Global FX Volatility Factor in Simulation

Panel I: Noisy Controls

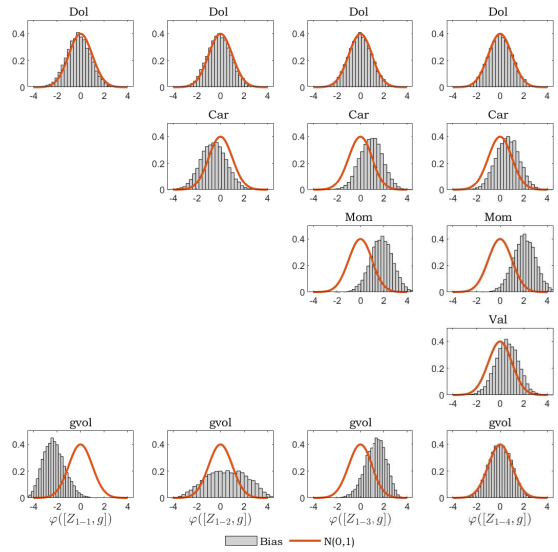

Panel II: De-Noised Controls

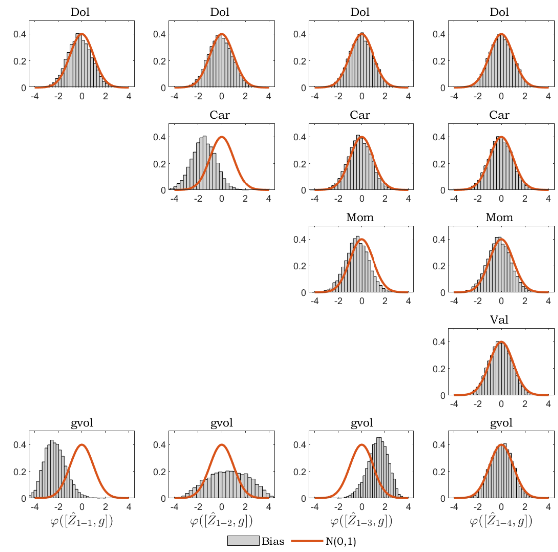

The figure shows the histograms of the two-pass standardized bias estimates of the global FX volatility factor of Menkhoff et al. (2012a) using Shanken (1992) standard errors. We consider SDFs of expanding dimension by adding Dollar, Carry, ST Mom, and Value control factors one at a time along with the candidate nontradable factor $g = gvol$. In Panel I the controls contain noise ($\varphi([Z_{1-k}, g])$), while in Panel II are de-noised ($\varphi([\hat{Z}_{1-k}, g])$). The biases are obtained as the estimates minus the true risk premia (the latter are presented in Panels AI.G1 and BI.G1 of Table A18). The true data generating process (DGP) has four de-noised factors, and we simulate the models 10,000 times with $N = 46$ and $T = 410$ (see Section VI and Table A17).

Figure A20: Histograms of the Two-Pass Standardized Bias Estimates of the Intermediaries' Capital Ratio Factor in Simulation

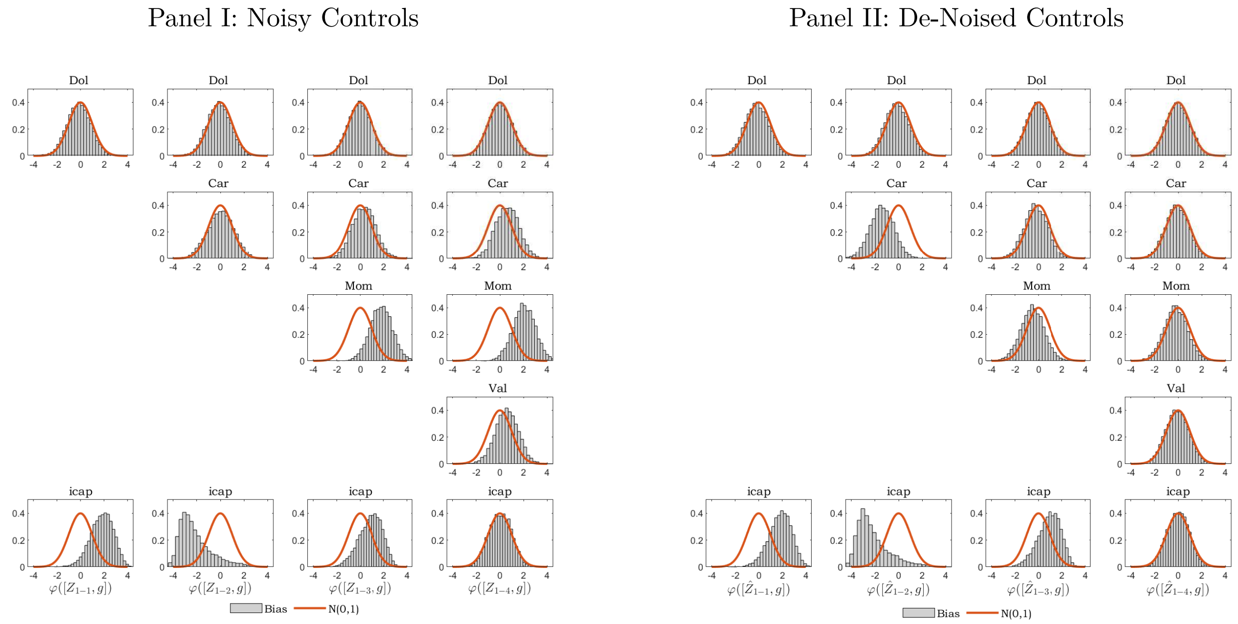

The figure shows the histograms of the two-pass standardized bias estimates of the intermediaries' capital ratio factor of He et al. (2017) using Shanken (1992) standard errors. We consider SDFs of expanding dimension by adding Dollar, Carry, ST Mom, and Value control factors one at a time along with the candidate nontradable factor $g = icap$. In Panel I the controls contain noise ($\varphi([Z_{1-k}, g])$), while in Panel II are de-noised ($\varphi([\hat{Z}_{1-k}, g])$). The biases are obtained as the estimates minus the true risk premia (the latter are presented in Panels AI.G2 and BI.G2 of Table A18). The true data generating process (DGP) has four de-noised factors, and we simulate the models 10,000 times with $N = 46$ and $T = 410$ (see Section VI and Table A17).

Figure A21: Histograms of the Two-Pass Standardized Bias Estimates of the U.S. Industrial Production Factor in Simulation

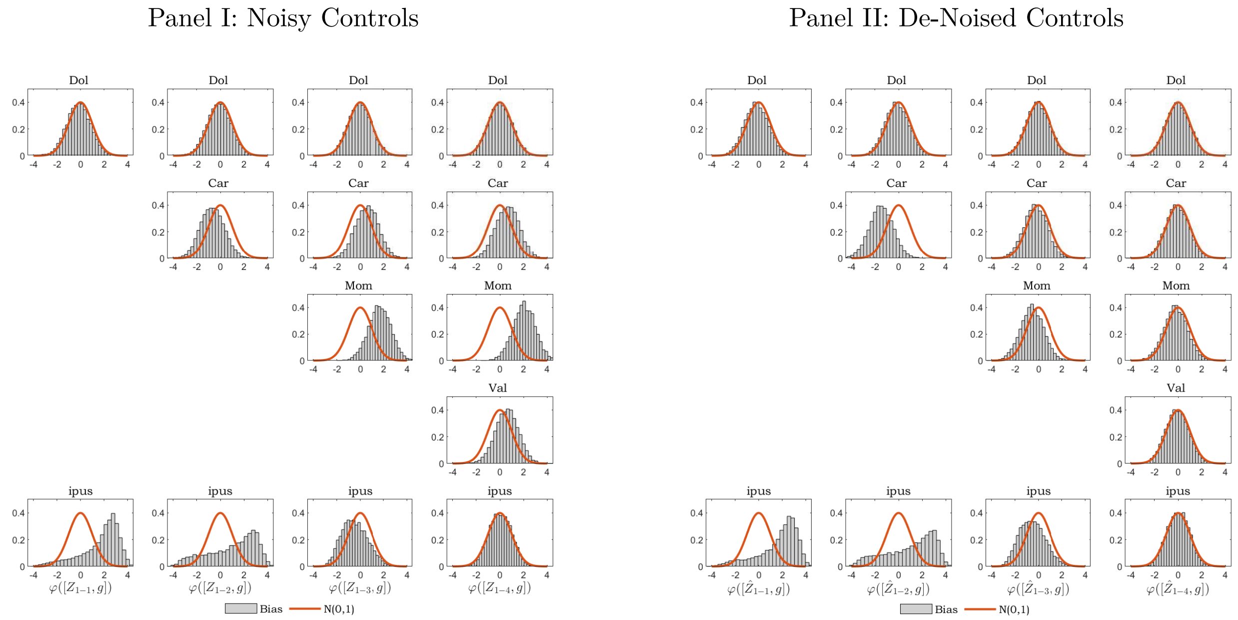

The figure shows the histograms of the two-pass standardized bias estimates of the U.S. industrial production factor using Shanken (1992) standard errors. We consider SDFs of expanding dimension by adding Dollar, Carry, ST Mom, and Value control factors one at a time along with the candidate nontradable factor $g = ipus$. In Panel I the controls contain noise ($\varphi([Z_{1-k}, g])$), while in Panel II are de-noised ($\varphi([\hat{Z}_{1-k}, g])$). The biases are obtained as the estimates minus the true risk premia (the latter are presented in Panels AI.G3 and BI.G3 of Table A18). The true data generating process (DGP) has four de-noised factors, and we simulate the models 10,000 times with $N = 46$ and $T = 410$ (see Section VI and Table A17).

VII Weak Candidate Factors

The augmented three-pass procedure tackles both the omitted-variable and measurement-error problems in the estimation of factor risk premia; see Section 2, and Section VI in the Internet Appendix. However, the method (regardless of the RP-weight used) is not designed to explicitly address the issue of weak candidate factors (Giglio et al., 2021c, GXZ henceforth). Such problem manifests if only few assets are exposed to the candidate factor, and this can in turn affect the inference on the factor risk premium. A factor is more likely to be weak in large cross sections of test assets, because the strength of a factor is not an inherent property of the factor, as it also depends on the cross section of assets used in the analysis. Hence, a factor is not strong or weak in absolute terms, but relative to the cross section of test assets. Large cross sections are a prerequisite of the three-pass estimator to effectively address the omitted-variable problem, as the extracted latent factors and hence the SDF should span all relevant sources of FX risk. Despite the cross section of FX portfolios is relatively small, the problem of weak factors could manifest to different degrees in our sample of candidate factors. For example, it might weigh on the inference on macro factor risk premia, and thus explain at least in part the disconnect that we document between macro factors and currency portfolio returns. But it could also affect the risk-premium estimates of some financial and text-based factors if these factors are also weak.

In light of these considerations, we assess the robustness of the candidate factor risk-premium estimates using the supervised principal component analysis (SPCA) recently proposed by GXZ. This novel three-pass estimator is designed to explicitly tackle the omitted-variable and measurement-error problems accounting for the possibility that the candidate factor of interest is *weak*. The SPCA procedure delivers robust estimates of a weak factor's risk premium because it shrinks and adapts the assets' cross section to match the properties of the factor. By doing so, the factor becomes strong with respect to the new tailored cross section of assets, making the inference on the factor risk premium valid. As a consequence, unlike the original GX three-pass method, the SPCA estimator does not hinge on a unique SDF common to all candidate factors (being driven by the full cross section of assets), but potentially on several SDFs. Put differently, in SPCA the SDFs are pinned down by the candidate factor properties and hence by supervised selections of the original cross section. Thus, while SPCA is not particularly useful to shed light on the properties of the currency SDF (the first goal of this study), it is highly relevant for the second goal of our study: that is, the estimation of candidate factor risk premia.

VII.1 SPCA Method

In what follows, we describe the main insights and steps of the SPCA method, while we refer the reader to GXZ for more details. Take the $T \times N$ cross section of asset excess returns X, and a generic candidate factor, either tradable or nontradable. Then, the SPCA procedure consists of iterating on the following steps: i) compute the univariate correlation of all test asset returns X with the candidate factor of interest; ii) select the top-qN assets most correlated with the factor; iii) extract the first latent factor from this subset of qN asset returns via PCA (thus the tuning parameter q determines how many assets we use to extract the latent factor); and iv) project the candidate factor and *all* the X-asset returns on the latent factor, and then take residuals. By iterating k times on steps *i-iv* (where portfolio returns and the candidate factor are replaced from the second iteration onwards by the residuals from the previous iteration projection step), one recovers the k latent factors. Equipped with the estimates of the k latent factors, one then applies the three-pass GX procedure and retrieves the risk-premium estimates of the candidate factor. The procedure is performed separately on each candidate factor.

In the above procedure, the two key tuning parameters q and k are assumed to be known by the researcher, i.e., to be predetermined. GXZ show that these parameters can be jointly determined in advance by repeating M times a $\mathcal{K}$-fold cross-validation exercise. The exercise essentially consists of first constructing a grid of out-of-sample R^2s for different combinations of the tuning parameters q and k (the R^2s are the averages from repeating the cross-validation M times), and then selecting the combination $\{q,\, k\}$ that maximizes the R^2. Specifically, in line with GXZ, we design the cross-validation exercise as follows.[51] For each repetition m (with m ranging from 1 to $M = 100$), we randomly split the sample period into $\mathcal{K} = 3$ parts of equal length, where two parts consist of the longer training period and the remaining part of the shorter testing (or out-of-sample evaluation) period. For selected combinations of the tuning parameters q and k, we build the weights of the hedging portfolio for the candidate factor by SPCA using the training data only.[52] Such weights are then used to construct the returns of the hedging portfolio over the evaluation period and then determine the fraction of the factor's variance hedged by the portfolio (the out-of-sample R^2). For each repetition m, we perform the analysis for all three possible permutations of training and testing periods. For each combination of $\{q,\, k\}$, we then average over the explained variances computed over the three testing periods and obtain the grid of R^2s for repetition m. We repeat these steps for each of the M repetitions, and then average over the M grids containing the cross-validation out-of-sample R^2s of each repetition to get the final grid. Based on this grid, we choose the pair $\{q,\, k\}$ that yields the highest cross-validation out-of-sample R^2.[53]

Before turning to the empirical results, two observations are in order. First, if the number of selected test assets qN equals the overall number of test assets N (i.e., there is no asset selection), then the SPCA factor risk-premium estimates are identical to the standard GX three-pass estimates (i.e., where the factors are estimated with PCA; see Panels $\omega = -1$ of Tables A10 and A11), for a given dimension of the SDF. Therefore, for these estimates to be free from an omitted-variable bias, the SDF needs to include enough latent factors to fully span the assets' space. We know from Section 4 that for this to be the case at least three, and potentially four, latent factors are needed using RP-PCA (more factors are needed using PCA; see also *Panel T.II* in Table A14). Second, in the above sketch of the SPCA algorithm, the analysis is carried out for one candidate factor at a time. Thus, the assets' selection is only driven by the factor at hand. By performing the SPCA estimation factor by factor, the risk-premium estimates are consistent and, importantly, we can determine which assets are relevant for which factors. However, factor risk premia can also be estimated via SPCA using more factors simultaneously. In such joint estimation, the selection of the assets is driven simultaneously by a set that includes multiple (potentially all) candidate factors. Specifically, assets are sorted by the maximum correlation with any of the factors in the set. While both the

[51]We implement the cross-validation exercise as set up in the SPCA code made available by GXZ.

[52]We require a minimum of 10 assets selected, and similarly to GXZ we work with a range of numbers of test assets (instead of working with the tuning parameter q). Specifically, we consider qN= [10, 15, 20, 25, 30, 35, 40, 46]. Meanwhile, we assess models with SDFs of expanding dimension including a maximum of 10 latent factors, so that $k = [1, \ldots, 10]$. Thus, the grid of out-of-sample R^2 has 8×10 entries.

[53]One can alternatively fix one of the two parameters, and select the second parameter that maximizes the out-of-sample R^2 for that choice of the first parameter. In this way, for a given candidate factor one can for example assess how the optimal qN and R^2 vary with the dimension of the SDF.

one-by-one and the joint factor SPCA estimators are consistent, the joint estimation is required to make inference on the premium estimates (see Giglio et al., 2021c). This is because for the central-limit-theorem assumptions to hold, the factors are required to have exposures to the entire SDF. This is a far more stringent requirement to be satisfied than is needed for consistency.

Finally, in presenting the empirical evidence below, we distinguish between factors with positive and negative cross-validation out-of-sample R^2s. We only report the estimation results for factors with positive R^2s, as a negative R^2 suggests that we cannot hedge that factor, and hence its risk-premium estimate is not informative. Thus, the cross-validation exercise allows us to further filter out factors that cannot be hedged out of sample by the currency assets. This criterion can be seen as an additional way to further discern relevant candidate factors from non-relevant ones. This means that, while SPCA gives us the best chance to detect candidate factors with a non-zero risk premium by allowing for weak factors, it can also reduce the number of relevant factors (relative to the three-pass procedure) due to this additional constraint.

VII.2 SPCA Results

Next we assess the risk-premium estimates obtained using SPCA, and then unveil the identities of the assets selected by SPCA in the estimation of the latent factors, another useful by-product of the procedure. We present the main results for both tradable and nontradable factors.

VII.2.1 Cross-Validation Analysis and Risk-Premium Estimates

While the main focus of our analysis pertains to the estimation of nontradable factor risk premia, it is convenient to look at the case of tradable factors first. This is exactly because for tradable factors we can benchmark the SPCA estimates to the model-free estimates of the factor risk premia, and hence we can precisely assess the performance of the SPCA estimator. In particular, we can determine the number of latent factors entering the SDF in the joint estimation, which is then key to carry out the inference on both tradable and nontradable factors. Thus, we first review the case of tradable factors, before turning to the nontradable factor risk-premium estimates, which we then contrast with our baseline three-pass estimates of Table 3.

Tradable Z-Factors. Table A19, Panel A, reports the estimates of the tradable factor risk premia obtained applying SPCA to each candidate factor separately. The tradable Z-factors include the Dollar level factor (the cross-sectional average of individual currency returns) and the nine HML factors of the investment strategies. The first column shows the model-free estimates (avg) for each of the tradable factors, while the subsequent blocks of columns present the results for models including an increasing number of latent factors (the tuning parameter k). For each model, the table displays the SPCA risk-premium estimates (λ), the cross-validation out-of-sample R^2s ($R2$) of the implied hedging portfolios for the tradable factor, and the number of test assets ($\#TA$) selected also in the cross-validation exercise, i.e, qN. Importantly, we find that SPCA successfully recovers the model-free risk-premium estimates for all tradable factors. However, while for the Dollar factor even a parsimonious model suffices, for most HML factors (e.g., Value) the model requires a larger number of latent factors. At times, the improvement in terms of R^2s is material also for reasonably large SDFs (e.g., for GAP when moving from a model with eighth to one with ten latent factors). By including eight latent factors, all R^2s are above 90%, with the only exceptions of Value and LDC. Moreover, the number of selected assets is 10 (i.e., the minimum allowed) for all HML factors, and does not vary with the number of latent factors. In contrast, for the Dollar factor the number of selected assets decreases with the number of latent factors (from a maximum of 46 to a minimum of 10).

Next, Panel B shows the results from the joint Z-factor estimation. Because here the asset selection is performed jointly, the models require a slightly higher number of factors to recover the model-free estimates than in the one-by-one factor analysis. But Figure A22 shows that the improvements in the model performance level off as the model includes at least 10 latent factors. Specifically, for models with more than 10 factors, the gains in the R^2s are small, and the differences in the root-mean-square errors are essentially imperceptible. Moreover, the bottom panel shows that the selected number of assets drops to 15 as the fourth latent factor is included into the SDF, and does not change as more factors are added to the SDF. Thus, taken together, this evidence suggests that a 10-factor model performs well in the joint factor analysis. Hence, we use this model to make inference on the risk-premium estimates. The standard errors show that all estimates are statistically not different from the factor averages, and meanwhile are statistically different from zero (marginally so for Term and LYld). Note that standard errors do not necessarily increase as more factors are added to the SDF. This evidence on the joint Z-factor analysis paves the way to the nontradable G-factor analysis, to which we turn next.

Nontradable G-Factors. Table A20 presents the risk-premium estimates for the nontradable factors. The table shows in the first columns (*Optimal k*) the SPCA model estimates when both the numbers of test assets and latent factors are chosen with the cross-validation exercise. To start with, we use this evidence to distinguish between factors with positive and negative R^2s. We do not report the estimation results for these latter factors, as a negative R^2 suggests that we cannot hedge that risk factor with the selected assets. If we cannot construct an hedging portfolio for a factor, its risk-premium estimate is not informative (given that the risk-premium can be interpreted as the cost of hedging insurance). We therefore focus on the remaining nontradable factors that can be at least partly hedged.

Before delving into the individual factor estimates, we look at the average R^2s by factor types (we also add the tradable factor evidence). We do so in Figure A23 where we show the average R^2s for models with different numbers of latent factors. We apply the same filter of Table A20 and only include the nontradable factors with positive R^2s. We find that, while for tradable factors the R^2s increase with the number of latent factors (that is we can better hedge the factor with more latent factors), the opposite is true for nontradable factors. But there is a clear distinction among factors of different types. The R^2s are positive for financial factors, improve substantially when moving from a one- to a two-factor model (reaching roughly 9%), and decay, albeit slowly, as more factors are added. Also for text-based factors we can appreciate a clear improvement in the R^2 when the second factor is included in the model; but the average two-factor model R^2 is much lower. Moreover, the text-based factor R^2s reduce almost linearly as the models include more factors, and eventually turn negative as the models include more than five factors. Of particular interest is the evidence on macro factors, arguably the main reason why we turned to the SPCA estimator (given our baseline results). The one-factor model R^2 is only slightly positive, while the remaining R^2s are all negative, and more so as the models include more factors. Therefore, we cannot hedge essentially any of the macro factors using FX portfolios. Put simply, this evidence confirms and refines the disconnect between FX returns and macro factors uncovered using the augmented three-pass estimator. If before we argued that, based on the three-pass estimates, the disconnect cannot be imputable either to omitted variables and/or measurement errors in the factors, here we add that it cannot be ascribed either to the fact that macro factors are weak in the cross section of FX portfolios.

While the above evidence is in itself highly informative, we then look at some of the individual nontradable factors, as there are significant differences also among factors of the same type. For most financial factors, the cross-validation exercise reveals that there is no asset selection and the 'optimal' number of latent factors is two (Panel A, Table A20).[54] But for some factors there is supervision as qN is less than N (i.e., otic, icap, noise, sliq, gfc), and for almost all these factors the optimal number of latent factors turns out to be two (except for gfc that is four). Turning to the text-based factors (Panel B), we find that six out of 13 factors display asset selection to different degrees but all yield the highest R^2 with $k = 2$ (i.e., gepu, gepu ppp, epu all, epu mp, fsi tx, and emv mp). Finally, for all macro factors (Panel C), we find no selection of test assets, and with only few exceptions a one-factor model is optimal according to the cross-validation exercise. Thus, this evidence echoes that uncovered earlier on, suggesting no relationship between the macro factors and FX assets.

To further zoom into the differences among factors, Figure A24 plots, for selected factors, the R^2 and the number of selected test assets (qN) also using SDFs of larger dimension. The factor gvol stands out as a strong factor, as the cross-validation method achieves the highest R^2 using all N assets, and enough latent factors (there is evidence of asset selection only for models including less than six latent factors). On the other hand, icap is a weaker factor, as the best performing model in the cross-validation exercise features asset selection and a contained number of latent factors. The factor gepu is also a weak factor as some supervision helps, while the model performance gradually worsens by adding latent factors and using all assets. The text-based factor emv mp behaves similarly to gepu, but the cross-validation R^2s are lower and turn negative for larger models. Finally, the two macro factors exemplify the disconnect between macro factors and currency investment strategies.

Taken together, these multiple pieces of evidence show that, for nontradable factors models with a contained number of factors perform better (this contrasts with the previous evidence on tradable factors). With the exception of gvol, there is little scope in looking at models including more than four factors. We therefore add in the subsequent columns of Table A20 the estimation results obtained with two-, three-, and four-factor models (i.e., $\varphi(F_{1-k}^{[g_j]})$, for $k = 2, 3, 4$). Consistent with Figure A24, it is evident that for most factors that display selection the R^2s tend to decrease as the models include more factors. However, as noted before, for most of these factors only the joint estimation turns out to be actually meaningful. This is because, without asset selection, the risk-premium estimates obtained using a small number of factors are distorted due to the omission of relevant factors.

Therefore, to complete the nontradable-factor analysis and carry out inference on the risk premia, we turn to the joint estimation. Here, we depart slightly from the approach used by GXZ in their empirical application, but we follow their theoretical insights to make the inference valid. Specifically, we estimate the model simultaneously on the candidate g-factor of interest and all the ten tradable Z-factors. In this way, we fulfill the key requirement that the factors need to jointly span the entire SDF. Moreover, we can make use of the evidence uncovered before on the joint estimation of tradable factors' risk premia (*Panel B*,

[54]As noted before, for these factors the 'optimal' individual factor risk-premium estimates are not of particular interest, as the SDF omits relevant sources of risks (other than Dollar and Carry).

Table A19) and set k to 10. Thus, we use the model $\varphi(F_{1-10}^{[g;Z]})$ to carry out the inference on a nontradable g-factor's risk premium. We find that not all factors with positive cross-validation out-of-sample R^2s have significant risk premia (*Joint*, Table A20). Specifically, the financial factors with significant premia are: otic, icap, gfc, gvol, psliq, move, vxo, and eqrv. Then, among the text-based factors, gepu and gepu ppp, the emv tracker and some of its subcategories (i.e., emv mout, emv inf, emv com, emv mp, and emv ir) are significant. Only ipw/us(q) is statistically significant among the macro factors.[55] Note that essentially all these factors had significant premia also using the RP-PCA three-pass estimator, and the two point estimates are generally close to each other (see Table 3).[56]

Summing up, we find that the nontradable factor risk-premium estimates are largely robust to the estimation method used. In fact, the estimates obtained with SPCA are largely consistent with our baseline three-pass estimates. Above all, this additional analysis confirms the disconnect between macro factors and currency portfolios. However, we see added value in complementing the baseline analysis with SPCA: while based on the three-pass method we argued that the disconnect cannot be attributed to either omitted variables in the SDF or measurement error in the factors, now we can further exclude that the disconnect is due to the fact that macro factors are weak in the cross section of currency portfolios. Moreover, the cross-validation exercise allows us to further filter out factors that cannot be hedged out of sample by the currency test assets. Next we turn to another useful output of SPCA which regards the assets selected (or the most relevant assets) for each nontradable factor.

VII.3 Asset Selection

In what follows we shift the focus to the assets selected by SPCA in the estimation of tradable and non-tradable factor risk premia. While our main goal is to use the asset-selection analysis to link currency portfolios to the relevant *nontradable* factors, it is first instructive to delve into the identities of the assets selected by SPCA to extract the latent factors which are then used to price the *tradable* factors. In fact, for the procedure to work well, it is reasonable to expect that SPCA exploits the information of the corner portfolios to price the HML factor of interest. At the same time, we regard the asset selection implemented on the tradable factors as an alternative valuable tool to shed light on the relationships among currency investment strategies.

Tradable Z-Factors. Table A21 reports the main findings on the assets selected by SPCA for the best performing models in the cross-validation exercise.[57] To start with, we look at the Dollar factor, and then

[55] Note that, while the point risk-premium estimates obtained with one-by-one and joint estimations are both consistent, in finite samples they can deviate from each other. For tradable factors, such differences are small. For many nontradable factors the differences are contained, while for others they are more evident. But, in line with what argued before, the comparisons of the risk premia obtained with the one-by-one and joint factor estimation is not meaningful for those factors without asset selection and SDFs including too few latent factors.

[56] Recall that we exclude from Table A20 factors that display negative out-of-sample R^2s in the cross-validation exercise. This explains why some of the candidate factors of Table 3 do not feature in Table A20, regardless of whether their premia are statistically significant or not. Thus the cross-validation R^2s can be regarded as an additional screening criterion to filter out factors that despite having positive in-sample R^2s have negative R^2s in the out-of-sample cross-validation exercise.

[57] For brevity, for a given factor of interest we only report the top-5 selected assets in terms of their absolute correlation with the factor. We also limit the focus to the first three iterations, that is, on the assets used to extract the first three factors. While we also do this to contain on space, for most factors the first three latent factors are enough to recover the bulk of the factor risk premia.

turn to the case of the HML factors (the latter are somewhat more informative as naturally connect with the currency investment strategies). From our previous analysis, it is clear that one latent factor is enough to price accurately the Dollar factor, therefore it is particularly useful to focus on the assets used by SPCA to extract the first latent factor. We find that none of the assets selected is a corner portfolio, and all the selected assets display very high correlations with the Dollar factor (being all above 90 percent). This finding is understandable because the Dollar factor plays the role of level or market factor.

On the contrary, for the HML factors the table shows that SPCA mostly relies on the information contained in the corner portfolios to extract the latent factors. In fact, for many HML factors the strategies' corner portfolios score as the top-two assets in terms of their absolute correlations with the factor. It is therefore evident that the procedure works particularly well in selecting the most relevant currency portfolios for the tradable factor at hand. In fact, the identities of the assets selected by SPCA are intuitively clear (at least based on the first three latent factors, which however account for most of the cross-validation out-of-sample R^2s). From the analysis it also emerges that Carry is an important source of risk for NFA, LDC, Term and LYld strategies (i.e., the carry-related strategies also driving $\hat{F}_2$ in the three-pass analysis), while Carry corner portfolios play little if no role (not being selected) in the extraction of the first three latent factors used to price the ST and LT Mom, Value, and GAP HML factors. Therefore, the identities of the selected assets confirm the close association between Carry and long-term interest rate, and global imbalance currency strategies. On the other hand, this evidence suggests that Carry, ST Mom, Value and GAP appear as distinct and, to a large extent, key sources of FX risks (consistently with the fact that these sources of risk feature to a different extent in the currency SDF uncovered using RP-PCA).

Nontradable G-Factors. In Tables A22 and A23, we report the top-5 assets selected by SPCA to extract the latent factors used to price the nontradable factors. We focus on the *relevant* factors, i.e., those of Table A20 with statistically significant risk premia. Moreover, like before we only show the assets used to extract the first three latent factors. However, for most nontradable factors, two is the optimal number of latent factors resulting from the cross-validation exercise. Hence, there is little scope to look beyond the third factor (for a candidate factor a column of dashes means that the latent factor is not selected by the cross-validation exercise).

It is evident that the absolute correlations are lower for nontradable than for tradable factors. But for many factors the correlations are quite high, and the assets selected make sense economically. For example, otic (a measure of foreign central banks' accumulation of U.S. Treasury securities) mostly correlates with global imbalance strategies (namely, NFA-P5). The only macro factor with significant risk premium, ipw/us(q) (measuring the difference between world and U.S. industrial production quarterly growth rates), is tightly linked to the high-risk GAP portfolio. On the other hand, GAP corner portfolios seem to be not particularly relevant for any of the other nontradable factors. In fact, most financial and text-based factors relate to Carry and carry-related strategies (e.g., LYld and NFA). This finding echoes the in-sample three-pass evidence of Table 2, showing that most factors are hedged by the second latent factor ($\hat{F}_2$), which can be interpreted as "Carry"; put differently, the SPCA analysis confirms that the tight link between currency and other asset markets is mainly channeled through Carry and hence also through carry-related strategies. Moreover, we find that for many nontradable factors the first and second extracted latent factors load respectively on high-risk and low-risk corner portfolios. In general, most of the top-ranking portfolios are corner portfolios, and this in turn helps relate the nontradable factors to specific investment strategies.

Thus, it is natural to ask, what are the most relevant nontradable factors for each currency investment strategy? To shed light on this question, Figure A25 shows for each corner portfolio its correlation with the relevant nontradable factors, i.e. those factors that display a statistically significant risk premium and a positive cross-validation out-of-sample R^2. We only report the correlation if the portfolio is selected by SPCA within the top-10 assets in the factor-by-factor analysis (thus the figure might capture factors that are not in Tables A22 and A23). More intense colors denote higher absolute correlations between the factor and the portfolio of interest, consistently with the SPCA selection criterion, while the white color means that the correlation is not within the top ten.

The key findings can be summarized as follows. First, for almost all strategies (except for Value and GAP), the high-risk corner portfolios show higher correlations with the nontradable factors than the low-risk corner portfolios; put simply, it is easier to explain P5/6 than P1 portfolios. This is particularly evident if one looks at Carry, NFA, Term, and LYld strategies. At the same, it emerges that is harder to connect ST Mom, Value, and GAP strategies than Carry and carry-related strategies to the nontradable factors. This holds if one looks at the first latent factors (*Panel A*), and even more if one looks at the assets selected in extracting the second latent factor (*Panel B*). Therefore, in comparison with the carry-strategies, the drivers of ST Mom, Value, and GAP remain less clear. However, our results show that some factors do matter also for these strategies. For example, focusing on the high-risk portfolios, vxo and icap seem to be relevant factors for Value, while gepu and eqrv single out as important risks for ST Mom.

On the other hand, as is evident by now, many sources of nontradable risks seem to explain carry-related strategies. The absolute correlations can provide a first means to try to detect the most relevant ones. However, if the SPCA method is useful in this regard, it is not explicitly designed to determine either the most relevant nontradable factors among the set of relevant ones (i.e., the factor pecking order), or which factors are subsumed by others. (The main objective remains the estimation of risk premia in the presence of weak factors.) Other methods, such as for example the one recently developed by Feng et al. (2020), are arguably better suited to address these questions.

Summing up, by means of the SPCA estimator, we detect a similar group of relevant nontradable factors to that documented in the baseline analysis, which instead uses the three-pass estimator with RP-PCA extracted factors. Above all, the two sets of results consistently point to a clear disconnect between macro factors and currency portfolio returns, at least in our sample. Thus, the SPCA estimates show that the disconnect cannot be imputable to the problem of weak factors. Moreover, the cross-validation exercise allows us to further filter out factors that cannot be hedged out of sample by the currency assets. This criterion can be seen as an additional way to further discern relevant candidate factors from non-relevant ones. Another clear benefit of using SPCA is the asset selection, which provides valuable information by zooming into the most correlated assets. In fact, by inspecting the identities of the selected assets, it emerges a clear link between the nontradable factors and the investment strategies. That is, Carry and carry-related strategies seem to react to a number of sources of nontradable risks, while we can identify a smaller number of relevant factors for ST Mom, Value, and GAP strategies. Thus, for the former strategies it might be useful in further research to try to establish a pecking order among these factors, whereas for the latter the search for other factors is still warranted to better hedge their underlying risks.

Table A19: SPCA Estimates of Tradable Factor Risk Premia

The table presents the in-sample SPCA estimates of the tradable factor risk premia. In Panel A we report in the first column the sample averages of the factors (avg). The next columns show, for models with different number of factors k, the factor-by-factor SPCA estimation results ($\varphi(F_{1-k}^{[z_j]})$). Specifically, for each choice of k, we report the risk premia estimates (λ), the cross-validation out-of-sample R^2s ($R2$) of the implied hedging portfolios (averages over evaluation periods), and the number of test assets ($\#TA$) selected by SPCA also in the cross-validation exercise (governed by the tuning parameter q). In Panel B, for convenience, we present again the factor averages. In the other columns, we show the SPCA estimates obtained including all the ten tradable Z-factors simultaneously, i.e., the joint factor analysis results ($\varphi(F_{1-k}^{[Z]})$). For each choice of k, we report the risk premia estimates (λ), the standard errors (se), and the difference between the risk premium estimate and the factor sample average (err). Finally, for the joint analysis, we also present the cross-validation out-of-sample R^2s of the implied hedging portfolios ($R2$), the root-mean-square errors (MSE), and the number of test assets selected by SPCA ($\#TA$). The test assets consist of the portfolios from the nine investment strategies ($N = 46$). The sample spans the 11/1983-12/2017 period at monthly frequency ($T = 410$).

Panel A: One-by-One Factor Analysis

		$\varphi(F_{1-2}^{[z_j]})$			$\varphi(F_{1-4}^{[z_j]})$			$\varphi(F_{1-6}^{[z_j]})$			$\varphi(F_{1-8}^{[z_j]})$			$\varphi(F_{1-10}^{[z_j]})$		
	avg	λ	R2	#TA	λ	R2	#TA	λ	R2	#TA	λ	R2	#TA	λ	R2	#TA
Dollar	2.62	2.41	0.99	46.00	2.41	1.00	15.00	2.43	1.00	15.00	2.50	1.00	10.00	2.58	1.00	10.00
Carry	7.26	3.93	0.74	10.00	5.38	0.84	10.00	6.42	0.87	10.00	6.71	0.90	10.00	7.02	0.93	10.00
Mom(ST)	6.93	1.30	0.27	10.00	4.67	0.82	10.00	6.78	0.89	10.00	7.04	0.94	10.00	7.11	0.96	10.00
Mom(LT)	4.24	3.23	0.67	10.00	3.32	0.83	10.00	4.54	0.89	10.00	4.94	0.93	10.00	4.63	0.95	10.00
Value	3.33	0.42	0.38	10.00	2.50	0.58	10.00	4.15	0.73	10.00	2.49	0.85	10.00	2.75	0.92	10.00
NFA	3.00	2.83	0.73	10.00	2.43	0.90	10.00	2.37	0.93	10.00	2.58	0.94	10.00	2.66	0.94	10.00
LDC	4.12	2.71	0.45	10.00	3.81	0.73	10.00	3.56	0.80	10.00	3.63	0.83	10.00	3.79	0.86	10.00
Term	2.82	3.24	0.46	10.00	3.51	0.81	10.00	3.01	0.90	10.00	2.53	0.93	10.00	2.40	0.96	10.00
LYld	1.87	3.30	0.83	10.00	1.82	0.90	10.00	2.56	0.93	10.00	2.17	0.94	10.00	2.42	0.95	10.00
GAP	6.27	2.30	0.13	10.00	6.78	0.69	10.00	6.36	0.85	10.00	6.48	0.92	10.00	6.38	0.96	10.00

Panel B: Joint Factor Analysis

		$\varphi(F_{1-6}^{[Z]})$			$\varphi(F_{1-8}^{[Z]})$			$\varphi(F_{1-10}^{[Z]})$			$\varphi(F_{1-11}^{[Z]})$			$\varphi(F_{1-12}^{[Z]})$		
	avg	λ	se	err	λ	se	err	λ	se	err	λ	se	err	λ	se	err
Dollar	2.62	2.49	1.30	0.14	2.51	1.30	0.11	2.44	1.30	0.18	2.42	1.30	0.20	2.42	1.30	0.20
Carry	7.26	4.89	1.39	2.37	5.46	1.42	1.80	6.20	1.50	1.05	6.08	1.56	1.18	6.17	1.57	1.08
Mom(ST)	6.93	5.90	1.47	1.02	6.73	1.53	0.20	6.68	1.51	0.25	6.62	1.51	0.31	6.63	1.51	0.30
Mom(LT)	4.24	4.44	1.44	-0.21	3.74	1.46	0.49	4.19	1.49	0.04	4.37	1.47	-0.14	4.33	1.46	-0.10
Value	3.33	1.16	1.02	2.17	2.01	1.02	1.32	2.92	1.12	0.41	3.28	1.33	0.05	3.13	1.33	0.20
NFA	3.00	4.30	1.23	-1.30	2.51	1.30	0.50	2.91	1.34	0.09	2.96	1.33	0.05	2.99	1.33	0.02
LDC	4.12	3.69	1.13	0.43	3.43	1.15	0.69	3.90	1.19	0.22	3.90	1.19	0.22	4.05	1.23	0.07
Term	2.82	0.54	1.41	2.28	1.88	1.48	0.93	2.73	1.60	0.08	2.93	1.58	-0.11	2.83	1.58	-0.02
LYld	1.87	4.01	1.48	-2.14	3.17	1.48	-1.30	2.91	1.55	-1.04	2.87	1.57	-1.00	2.90	1.57	-1.03
GAP	6.27	6.62	1.15	-0.35	7.21	1.29	-0.94	6.34	1.35	-0.07	6.27	1.35	0.00	6.26	1.35	0.01
		R2	MSE	#TA	R2	MSE	#TA	R2	MSE	#TA	R2	MSE	#TA	R2	MSE	#TA
		0.71	1.52	15.00	0.83	0.97	15.00	0.90	0.50	15.00	0.93	0.51	15.00	0.95	0.49	15.00

Table A20: SPCA Estimates of Nontradable Factor Risk Premia

The table presents the SPCA estimates of the nontradable factor risk premia (grouped in Panels A, B, and C by the types of factors). In the first-four blocks of columns we show the factor-by-factor SPCA estimation results. In the first block (Optimal k, $\varphi(F^{[g_j]}_{1-k})$), we report the SPCA risk-premium estimates (λ) using the combination of number of test assets ($\#TA$) and number of latent factors (k) that yields the highest out-of-sample R^2s ($R2$) in the cross-validation exercise. In the other blocks ($\varphi(F^{[g_j]}_{1-k})$), we report the risk premia estimates and the cross-validation out-of-sample R^2s obtained by estimating models with fixed number of latent factors (i.e., $k = 2, 3, 4$) and using the number of test assets ($\#TA$) that yields the highest out-of-sample R^2s in the cross-validation exercise for that choice k. In the last three columns (Joint, $\varphi(F^{[g_j,Z]}_{1-10})$), for a given candidate g-factor we report the risk-premium estimate (λ), the standard error (se), and the p-value ($pval$) obtained by estimating the SPCA model with 10 latent factors and including the g-factor simultaneously with the ten tradable Z-factors (see Table A19). We report the results only for the nontradable factors that have positive out-of-sample R^2s using the optimal combination of k and qN determined by the cross-validation exercise. The test assets consist of the portfolios from the nine investment strategies ($N = 46$). The sample spans the 11/1983-12/2017 period at monthly frequency ($T = 410$).

	Optimal k, $\varphi(F^{[g_j]}_{1-k})$				$\varphi(F^{[g_j]}_{1-2})$			$\varphi(F^{[g_j]}_{1-3})$			$\varphi(F^{[g_j]}_{1-4})$			Joint, $\varphi(F^{[g_j,Z]}_{1-10})$		
	λ	R2	#TA	k	λ	R2	#TA	λ	R2	#TA	λ	R2	#TA	λ	se	pval
Panel A: Financial Factors																
otic	0.51	0.02	30.00	2.00	0.51	0.02	30.00	0.50	0.01	46.00	0.72	0.01	30.00	0.78	0.34	0.03
icap	0.78	0.11	25.00	2.00	0.78	0.11	25.00	0.96	0.10	30.00	0.68	0.11	30.00	0.92	0.41	0.03
mf2	-0.20	0.02	46.00	1.00	-0.27	0.01	46.00	-0.08	0.01	46.00	-0.08	0.00	46.00	-0.22	0.29	0.44
mf3	0.10	0.02	46.00	3.00	0.35	0.02	46.00	0.10	0.02	46.00	0.10	0.01	46.00	0.10	0.26	0.70
noise	-1.08	0.08	20.00	2.00	-1.08	0.08	20.00	-0.98	0.07	40.00	-0.99	0.06	40.00	-0.64	0.43	0.15
sliq	-0.85	0.08	35.00	2.00	-0.85	0.08	35.00	-1.03	0.07	40.00	-0.95	0.08	46.00	-0.36	0.46	0.43
gfc	1.43	0.40	10.00	4.00	1.11	0.38	10.00	1.27	0.40	10.00	1.43	0.40	10.00	1.46	0.52	0.01
gliq	-0.24	0.01	46.00	2.00	-0.24	0.01	46.00	-0.19	0.00	46.00	-0.18	0.00	46.00	-0.19	0.22	0.38
gvol	-0.83	0.08	46.00	8.00	-0.73	0.07	15.00	-1.00	0.07	40.00	-0.93	0.07	40.00	-0.97	0.34	0.01
psliq	0.25	0.01	46.00	2.00	0.25	0.01	46.00	0.50	0.01	46.00	0.49	0.00	46.00	0.93	0.45	0.04
corp	-0.48	0.04	46.00	3.00	-0.76	0.04	40.00	-0.48	0.04	46.00	-0.49	0.03	46.00	-0.22	0.34	0.52
move	-0.72	0.07	46.00	2.00	-0.72	0.07	46.00	-0.72	0.06	46.00	-0.71	0.05	46.00	-0.77	0.40	0.06
vxo	-1.42	0.21	46.00	2.00	-1.42	0.21	46.00	-1.31	0.20	40.00	-1.30	0.21	46.00	-1.34	0.57	0.02
eqrv	-0.48	0.02	46.00	2.00	-0.48	0.02	46.00	-0.88	0.01	46.00	-0.88	0.00	46.00	-1.39	0.66	0.04
Panel B: Text-Based Factors																
	Optimal k, $\varphi(F^{[g_j]}_{1-k})$				$\varphi(F^{[g_j]}_{1-2})$			$\varphi(F^{[g_j]}_{1-3})$			$\varphi(F^{[g_j]}_{1-4})$			Joint, $\varphi(F^{[g_j,Z]}_{1-10})$		
	λ	R2	#TA	k	λ	R2	#TA	λ	R2	#TA	λ	R2	#TA	λ	se	pval
gepu	-0.92	0.06	35.00	2.00	-0.92	0.06	35.00	-0.85	0.06	46.00	-1.01	0.05	46.00	-1.20	0.50	0.02
gepu ppp	-0.90	0.06	35.00	2.00	-0.90	0.06	35.00	-0.86	0.06	46.00	-1.00	0.05	46.00	-1.26	0.50	0.02
epu all	-0.36	0.02	25.00	2.00	-0.36	0.02	25.00	-0.43	0.01	46.00	-0.35	0.01	46.00	-0.35	0.40	0.38
epu mp	-0.24	0.01	30.00	2.00	-0.24	0.01	30.00	-0.45	0.01	46.00	-0.44	0.00	46.00	-0.61	0.40	0.14
fsi tx	-0.45	0.01	20.00	2.00	-0.45	0.01	20.00	-0.35	0.01	46.00	-0.33	0.00	46.00	-0.38	0.37	0.31
emv ov	-0.51	0.02	46.00	2.00	-0.51	0.02	46.00	-0.72	0.01	46.00	-0.70	0.00	46.00	-0.89	0.45	0.06
emv mout	-0.46	0.01	46.00	2.00	-0.46	0.01	46.00	-0.74	0.01	46.00	-0.74	0.00	46.00	-0.89	0.43	0.05
emv inf	-0.41	0.02	46.00	2.00	-0.41	0.02	46.00	-0.52	0.01	46.00	-0.52	0.00	46.00	-0.82	0.39	0.04
emv com	-0.53	0.03	46.00	2.00	-0.53	0.03	46.00	-0.68	0.02	46.00	-0.66	0.01	46.00	-0.94	0.47	0.05
emv ir	-0.36	0.00	46.00	2.00	-0.36	0.00	46.00	-0.62	-0.01	46.00	-0.59	-0.02	46.00	-0.85	0.50	0.10
emv fx	-0.31	0.01	46.00	2.00	-0.31	0.01	46.00	-0.27	0.00	46.00	-0.33	0.00	46.00	-0.34	0.30	0.26
emv mp	-0.50	0.03	40.00	2.00	-0.50	0.03	40.00	-0.74	0.02	46.00	-0.74	0.02	46.00	-0.90	0.34	0.01
emv tp	-0.34	0.02	46.00	2.00	-0.34	0.02	46.00	-0.32	0.01	46.00	-0.34	0.00	46.00	-0.42	0.37	0.27
Panel C: Macro Factors																
	Optimal k, $\varphi(F^{[g_j]}_{1-k})$				$\varphi(F^{[g_j]}_{1-2})$			$\varphi(F^{[g_j]}_{1-3})$			$\varphi(F^{[g_j]}_{1-4})$			Joint, $\varphi(F^{[g_j,Z]}_{1-10})$		
	λ	R2	#TA	k	λ	R2	#TA	λ	R2	#TA	λ	R2	#TA	λ	se	pval
nfpyr(q)	0.13	0.01	46.00	1.00	0.18	0.00	46.00	0.18	0.00	46.00	0.18	-0.01	46.00	0.37	0.26	0.16
nfpyr(eq)	0.10	0.00	46.00	1.00	0.11	-0.01	46.00	0.06	-0.02	46.00	0.06	-0.03	46.00	0.15	0.32	0.64
cus(m)	-0.16	0.01	46.00	1.00	-0.15	0.00	46.00	-0.12	-0.01	46.00	-0.13	-0.01	46.00	-0.10	0.32	0.76
cus(eq)	-0.15	0.01	46.00	1.00	-0.16	0.00	46.00	-0.15	-0.01	46.00	-0.15	-0.01	46.00	-0.11	0.32	0.74
cus(ey)	-0.15	0.01	46.00	1.00	-0.16	0.00	46.00	-0.20	-0.01	46.00	-0.20	-0.02	46.00	-0.16	0.31	0.60
ipw/us(q)	-0.32	0.01	46.00	8.00	-0.10	-0.01	46.00	-0.32	-0.01	46.00	-0.32	-0.01	46.00	-0.60	0.26	0.03
cpiw/us(y)	-0.14	0.01	46.00	1.00	-0.18	0.01	46.00	-0.18	0.00	46.00	-0.18	0.00	46.00	-0.25	0.22	0.26
ipstdw(q)	-0.11	0.01	46.00	4.00	-0.08	0.00	46.00	-0.15	0.00	46.00	-0.11	0.01	46.00	-0.22	0.31	0.47
ipstdw(eq)	-0.11	0.01	46.00	4.00	-0.08	0.00	46.00	-0.15	0.00	46.00	-0.11	0.01	46.00	-0.22	0.31	0.47

Table A21: Asset Selections for Tradable Factors

For each tradable factor, the table reports the top-5 assets selected by SPCA in extracting the first three latent factors. While assets are ordered by their absolute correlation with the factor at hand, we show next to the name of the selected asset (*Asset*) its correlation with the candidate factor (ρ). Each panel refers to a different candidate factor and, for brevity, we report the results only for the first three iterations, i.e., the first three extracted latent factors. We refer to Table A19 for further details on the SPCA estimation and the underlying cross-validation exercise.

Dollar Factor

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
NFA-P3	94	Carry-P5	79	Carry-P1	52
Value-P3	93	NFA-P5	69	LYld-P1	49
LT Mom-P3	93	LYld-P5	59	NFA-P1	48
Value-P4	92	LDC-P6	58	LYld-P5	-43
NFA-P2	92	ST Mom-P1	54	Term-P5	-42

Carry HML Factor

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
Carry-P5	62	Carry-P1	-91	Carry-P5	81
LYld-P5	44	Carry-P5	85	Carry-P1	-65
Term-P5	38	LYld-P1	-83	NFA-P2	43
Carry-P1	-33	NFA-P1	-68	NFA-P5	-39
NFA-P4	31	NFA-P2	-67	LDC-P1	37

ST Mom HML Factor

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
ST Mom-P1	-60	ST Mom-P5	89	ST Mom-P1	-86
ST Mom-P5	44	ST Mom-P1	-84	ST Mom-P5	83
ST Mom-P2	-28	ST Mom-P4	48	ST Mom-P2	-40
LT Mom-P1	-22	Term-P3	30	LT Mom-P4	-40
LYld-P5	-18	LT Mom-P5	28	Carry-P3	-40

LT Mom HML Factor

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
LT Mom-P1	-58	LT Mom-P5	88	LT Mom-P1	-71
LT Mom-P5	49	LT Mom-P1	-87	LT Mom-P5	66
LT Mom-P2	-25	LT Mom-P4	50	Value-P2	-41
Value-P5	-21	Value-P1	43	LYld-P1	-39
LT Mom-P4	16	LT Mom-P2	-34	GAP-P2	-38

Value HML Factor

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
Value-P1	-50	Value-P5	88	Value-P5	79
Value-P5	37	Value-P1	-82	Value-P1	-75
LT Mom-P5	-27	LT Mom-P1	44	NFA-P5	-41
LT Mom-P4	-22	Value-P4	33	LDC-P5	-33
Value-P2	-21	LT Mom-P2	32	LDC-P4	-27

NFA HML Factor

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
NFA-P5	87	NFA-P5	83	NFA-P5	91
LDC-P4	64	NFA-P1	-67	NFA-P4	-54
LDC-P5	62	LDC-P1	-66	Term-P5	-46
Carry-P5	61	NFA-P2	-66	Carry-P1	41
LYld-P5	60	NFA-P3	-66	LYld-P1	35

LDC HML Factor

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
LDC-P6	64	LDC-P1	-70	LDC-P6	77
Carry-P5	55	NFA-P2	-58	LDC-P1	-48
NFA-P5	53	LYld-P1	-54	LYld-P5	-46
LYld-P5	42	LDC-P6	49	Term-P5	-38
ST Mom-P1	41	Carry-P1	-48	LDC-P5	-23

Term HML Factor

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
Term-P5	59	Term-P1	-86	Term-P5	84
Carry-P5	36	Term-P5	81	Term-P1	-75
LYld-P5	36	Carry-P2	-51	LYld-P1	44
Term-P1	-30	Value-P3	-50	NFA-P2	34
LDC-P5	20	Carry-P1	-50	LYld-P2	32

LYld HML Factor

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
LYld-P5	53	LYld-P1	-93	LYld-P5	84
Carry-P5	39	LYld-P5	93	LYld-P1	-55
LYld-P1	-38	Carry-P1	-80	LYld-P2	40
Carry-P1	-35	NFA-P2	-75	LDC-P2	34
Term-P5	34	NFA-P1	-71	NFA-P3	32

GAP HML Factor

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
GAP-P1	-53	GAP-P5	85	GAP-P5	81
GAP-P5	29	GAP-P1	-79	GAP-P1	-73
GAP-P2	-20	NFA-P3	21	GAP-P4	-29
LDC-P3	-20	Carry-P3	20	Term-P3	-25
ST Mom-P1	-20	ST Mom-P3	18	LT Mom-P4	-23

Table A22: Asset Selections for Financial Factors

For each relevant financial factor, the table reports the top-5 assets selected by SPCA in extracting the latent factors. While assets are ordered by their absolute correlation with the factor at hand, we show next to the name of the selected asset (*Asset*) its correlation with the candidate factor (ρ). Each panel refers to a different candidate factor and, for brevity, we report the results for the first three iterations, i.e., the first three extracted latent factors. For a given iteration, a column of dashes denotes that the associated latent factor is not selected by the cross-validation exercise. We do not report the results for the factors of which the risk premia estimates are not significant in Table A20. We refer to that table also for further details on the SPCA estimation and the underlying cross-validation exercise.

otic

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
NFA-P5	16	LYld-P1	-19	–	–
LDC-P5	15	LT Mom-P2	-18	–	–
Carry-P4	14	NFA-P2	-18	–	–
LDC-P4	13	NFA-P1	-17	–	–
Carry-P5	13	Carry-P1	-16	–	–

icap

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
NFA-P5	25	NFA-P2	-34	–	–
Carry-P5	23	LT Mom-P4	-30	–	–
LT Mom-P1	21	LYld-P1	-30	–	–
LDC-P4	21	NFA-P1	-29	–	–
LYld-P5	18	LYld-P2	-28	–	–

gfc

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
LDC-P5	54	LYld-P1	-38	LYld-P5	-20
LDC-P4	54	NFA-P2	-38	GAP-P2	20
NFA-P5	54	Value-P3	-35	Term-P5	-18
Carry-P5	53	Carry-P1	-34	LT Mom-P5	15
LYld-P4	52	LDC-P1	-33	LDC-P3	15

gvol

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
LT Mom-P5	-26	NFA-P2	33	NFA-P2	20
Carry-P5	-25	Carry-P1	33	LT Mom-P5	-20
NFA-P4	-24	LYld-P1	32	Carry-P1	20
NFA-P5	-23	NFA-P1	27	GAP-P2	-19
LDC-P5	-22	LDC-P1	25	LYld-P1	18

psliq

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
Carry-P5	12	Carry-P5	19	–	–
NFA-P5	10	NFA-P5	17	–	–
Carry-P1	-9	Carry-P1	-17	–	–
LT Mom-P5	8	LYld-P1	-17	–	–
ST Mom-P5	8	NFA-P2	-16	–	–

move

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
Carry-P5	-26	LYld-P1	27	–	–
ST Mom-P1	-24	Carry-P5	-22	–	–
LYld-P5	-23	NFA-P2	22	–	–
LDC-P4	-22	Carry-P1	21	–	–
LT Mom-P1	-22	LYld-P2	19	–	–

vxo

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
NFA-P5	-42	LYld-P1	42	–	–
LYld-P5	-38	NFA-P5	-38	–	–
Carry-P5	-37	NFA-P2	36	–	–
LDC-P4	-36	LYld-P5	-32	–	–
LDC-P5	-35	Carry-P1	31	–	–

eqrv

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
Carry-P5	-22	LYld-P1	29	–	–
LT Mom-P5	-21	NFA-P2	28	–	–
NFA-P5	-20	Carry-P1	27	–	–
LDC-P5	-19	Carry-P5	-24	–	–
LYld-P5	-18	LDC-P5	-23	–	–

Table A23: Asset Selections for Text-Based and Macro Factors

For each relevant text-based or macro factor, the table reports the top-5 assets selected by SPCA in extracting the latent factors. While assets are ordered by their absolute correlation with the factor at hand, we show next to the name of the selected asset (*Asset*) its correlation with the candidate factor (ρ). Each panel refers to a different candidate factor and, for brevity, we report the results for the first three iterations, i.e., the first three extracted latent factors. For a given iteration, a column of dashes denotes that the associated latent factor is not selected by the cross-validation exercise. We do not report the results for the factors of which the risk premia estimates are not significant in Table A20; we do not show the asset selection for gepu ppp as is similar to that of gepu. We refer to Table A20 also for further details on the SPCA estimation and the underlying cross-validation exercise.

gepu

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
Carry-P5	-21	LDC-P2	26	–	–
LT Mom-P5	-19	NFA-P2	25	–	–
LDC-P6	-18	LYld-P1	25	–	–
NFA-P4	-17	Carry-P1	24	–	–
NFA-P5	-16	LT Mom-P3	23	–	–

emv ov

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
Carry-P5	-17	LYld-P1	24	–	–
NFA-P5	-16	Carry-P1	21	–	–
LYld-P5	-15	NFA-P2	21	–	–
LDC-P5	-15	Carry-P5	-20	–	–
LT Mom-P5	-15	LDC-P5	-20	–	–

emv mout

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
Carry-P5	-16	LYld-P1	24	–	–
LYld-P5	-14	NFA-P2	21	–	–
NFA-P5	-14	Carry-P1	20	–	–
LT Mom-P5	-13	Carry-P5	-20	–	–
LDC-P5	-13	LYld-P5	-18	–	–

emv inf

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
Carry-P5	-15	Carry-P5	-20	–	–
NFA-P5	-13	LYld-P1	19	–	–
LYld-P5	-12	Carry-P1	18	–	–
LDC-P4	-11	NFA-P5	-17	–	–
LDC-P5	-10	NFA-P2	17	–	–

emv com

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
Carry-P5	-18	LYld-P1	25	–	–
NFA-P5	-18	Carry-P1	22	–	–
LYld-P5	-17	NFA-P2	21	–	–
LDC-P4	-15	Carry-P5	-21	–	–
LDC-P5	-15	NFA-P5	-20	–	–

emv ir

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
Carry-P5	-13	LYld-P1	20	–	–
LDC-P5	-11	Carry-P1	19	–	–
NFA-P5	-11	Carry-P5	-17	–	–
LYld-P5	-11	LDC-P5	-17	–	–
LT Mom-P5	-10	LYld-P2	16	–	–

emv mp

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
Carry-P5	-18	LYld-P1	24	–	–
LYld-P5	-17	Carry-P1	21	–	–
NFA-P5	-16	NFA-P2	20	–	–
LT Mom-P5	-16	Carry-P5	-18	–	–
Term-P5	-15	LYld-P5	-17	–	–

ipw/us(q)

Iter #1		Iter #2		Iter #3	
Asset	ρ	Asset	ρ	Asset	ρ
GAP-P5	-12	GAP-P5	-18	GAP-P5	-18
LYld-P5	-8	LT Mom-P2	17	LT Mom-P2	16
NFA-P5	-8	Term-P2	-12	Term-P2	-14
LT Mom-P5	-8	NFA-P5	-12	LYld-P1	-12
Term-P2	-7	Term-P3	12	LT Mom-P4	-11

Figure A22: Diagnostics of the Joint SPCA Estimation on Tradable Factors

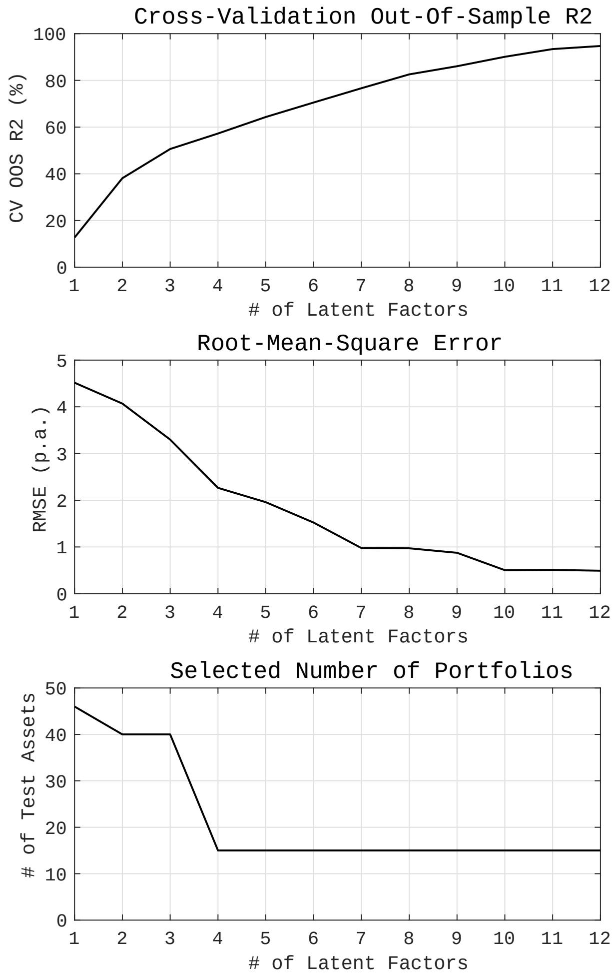

The figure shows the diagnostics for SPCA models with up to 12 latent factors estimated including all the ten tradable Z-factors simultaneously. We report in the top panel the cross-validation out-of-sample R^2s of the implied hedging portfolios, in the middle panel the root-mean-square errors (computed as the differences between the risk premium estimates and the factor sample averages), and in the bottom panel the number of test assets selected by SPCA (governed by the parameter q in GXZ). The test assets consist of the portfolios from the nine investment strategies ($N = 46$). The sample spans the 11/1983-12/2017 period at monthly frequency ($T = 410$).

Figure A23: Average Cross-Validation Out-of-Sample R^2s by Candidate Types

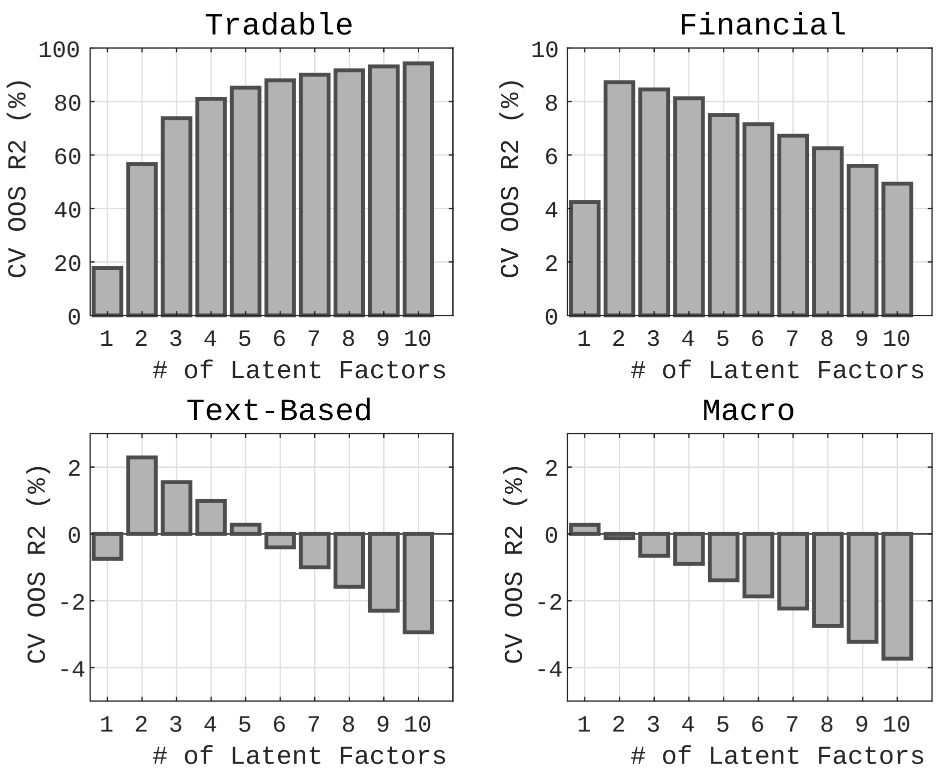

The figure shows the average cross-validation out-of-sample R^2s (*CV OOS R2 (%)*) for both tradable and nontradable factors, grouped by factor types. The individual factor R^2s are obtained by estimating the SPCA models factor by factor, using an increasing number of latent factors (k ranging from 1 to 10). We filter out candidate nontradable factors with negative R^2s based on the optimal combination of number of latent factors (k) and number of test assets (qN) selected by the cross-validation exercise. Thus, the underlying candidate nontradable factors are those displayed in Table A20. For a given k, we then average the factor R^2s by factor types. The test assets consist of the portfolios from the nine investment strategies ($N = 46$). The sample spans the 11/1983-12/2017 period at monthly frequency ($T = 410$).

Figure A24: SPCA Diagnostics of Selected Nontradable Candidate Factors

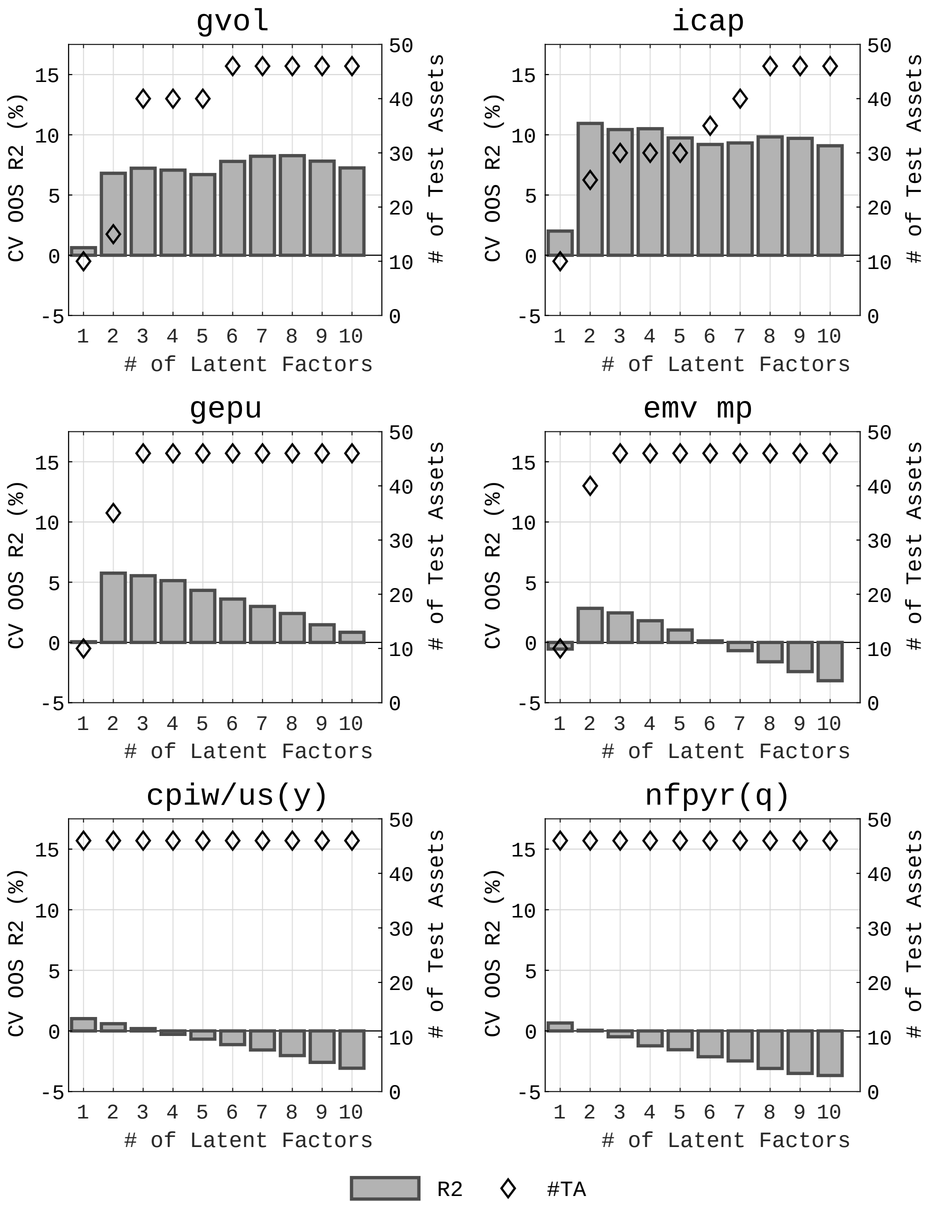

For selected nontradable candidate factors, the figure shows the cross-validation out-of-sample R^2s on the vertical axis and, on the horizontal axis, the number of test assets for SPCA models with different number of latent factors, with k ranging from 1 to 10. The SPCA models are estimated individually on each candidate factor. We select from Table A20 two financial factors (*gvol*, and *icap*), two text-based factor (*gepu*, and *emv mp*), and two macro factors ($cpiw/us(y)$, and $nfpyr(q)$). Further details on the cross-validation exercise are displayed in Table A20. The test assets consist of the portfolios from the nine investment strategies ($N = 46$). The sample spans the 11/1983-12/2017 period at monthly frequency ($T = 410$).

Figure A25: Relevant Nontradable Factors and Corner Portfolios

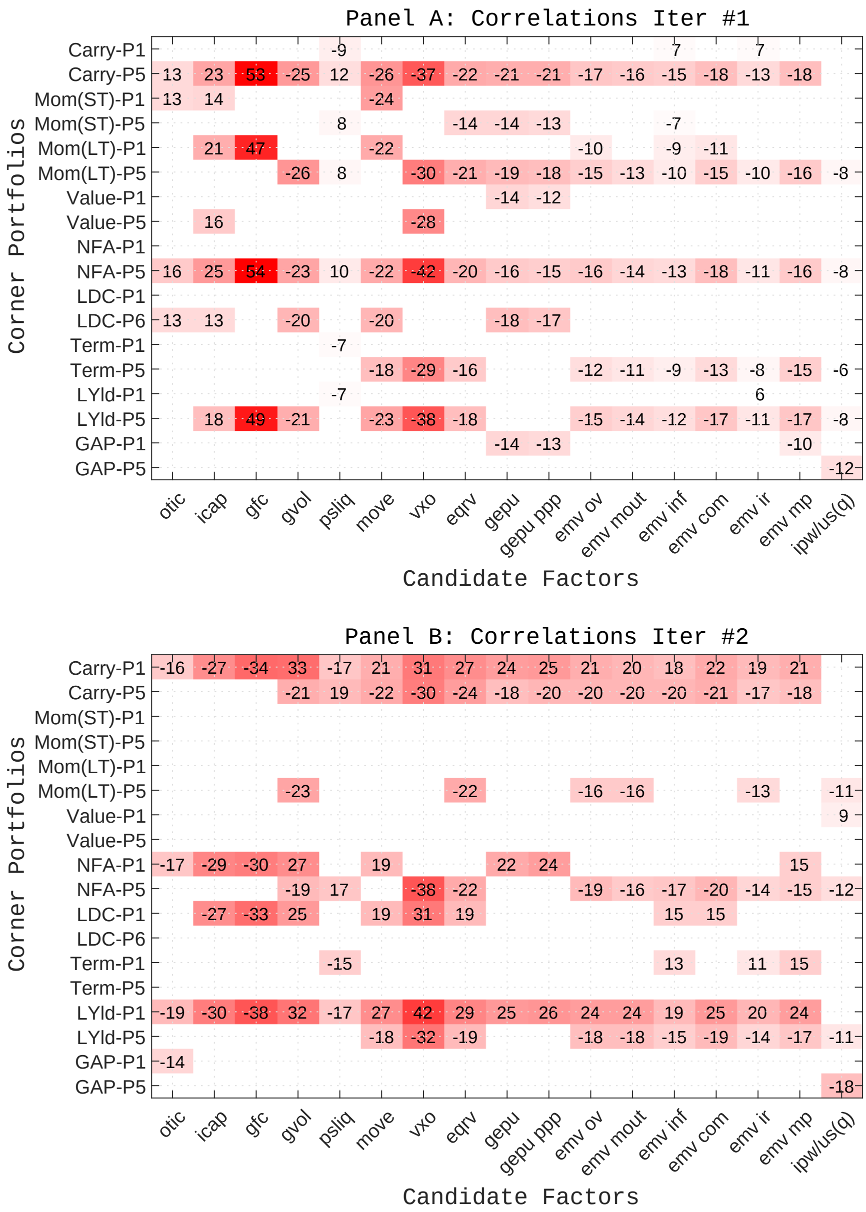

For each relevant nontradable factor, the figure shows its correlation with the investment strategies' corner portfolios, if the corner portfolio is selected by SPCA within the top-10 most correlated assets. We omit to report the middle portfolios for brevity but also because are less informative. The colors are set based on the absolute correlations, consistently with the SPCA selection criterion, while the white color means that the corner portfolio is not in the group of the top-10 most correlated assets for a given candidate factor and latent-factor extraction. We present the results for the first two iterations of SPCA, i.e., the first two extracted latent factors (for most nontradable factors the third latent factor is not selected by the cross-validation exercise, see Table A20). We only report the results for the relevant nontradable factors of Table A20, i.e. factors with significant risk premia and positive cross-validation out-of-sample R^2s. We refer to that table also for further details on the SPCA estimation and the underlying cross-validation exercise.