Benchmark Currency Stochastic Discount Factors *

Piotr Orłowski
HEC Montréal

Valeri Sokolovski
HEC Montréal

Erik Sverdrup
Stanford University

December 30, 2021

Abstract

We examine the pricing performance of out-of-sample pricing factors in the broad cross-section of currency returns. To this end, we develop a methodology for estimating empirical minimum-dispersion stochastic discount factors (SDFs) under constraints on maximum position leverage. Under leverage constraints compatible with those observed in the currency markets, our empirical SDFs deliver smaller out-of-sample pricing errors than existing factor models, and are priced in individual currency and hedge fund cross-sections. After transaction costs, an investable SDF portfolio delivers a Sharpe ratio of around 0.8 and positively skewed returns. These empirical SDFs offer tractable benchmarks for candidate currency pricing models.

Keywords: FX risk premium, FX hedge funds, leverage, machine learning, SDF.
JEL Classification: F31, G12, G15

*For their helpful comments and suggestions, we thank Caio Almeida, Gurdip Bakshi, Jaroslav Borovička, Nicole Branger, Magnus Dahlquist, Kent Daniel, Esther Eiling, Bob Hodrick, Alexandre Jeanneret, Pascal François, René Garcia, Richard Levich, Matti Suominen, Simon van Norden, Steven Riddiough, Kirsten Smart, Michael Weber, Josef Zechner, Irina Zviadadze, and the conference and seminar participants at the 2019 Vienna Symposium on Foreign Exchange Markets and HEC Montréal. Sokolovski gratefully acknowledges the financial support from les Fonds de recherche du Québec (FRQ). Orłowski: piotr.orlowski@hec.ca; Sokolovski: valeri.sokolovski@hec.ca; Sverdrup: erikcs@stanford.edu.

1 Introduction

A large literature documents persistent foreign exchange (FX) risk premiums. Specifically, there are at least three distinct, widely-used, and historically profitable currency strategies: carry, momentum, and value.[1] For the profitability of these strategies to persist in equilibrium, it ought to be commensurate compensation for risk. In addition, if the FX market is not segmented, a good model should explain the returns of all basic currency strategies. Yet, there is no consensus in the literature on a benchmark model to explain these premiums (see, e.g., Burnside, Eichenbaum, and Rebelo (2011b), and Daniel, Hodrick, and Lu (2017); and surveys by Burnside (2012), and Hassan and Zhang (2020)).

A promising solution is to extract currency pricing factors directly from historic returns as it is always possible to construct an SDF from a given set of assets that would perfectly price it (Hansen and Jagannathan (1991), HJ hereafter, and Almeida and Garcia (2017)). In practice, however, in broad cross-sections such factors often exhibit economically unrealistic properties like extreme volatility. Associated with these factors are in-sample optimal portfolios, which typically perform poorly out-of-sample. One distinctive characteristic of these portfolios is the presence of large portfolio weights that imply unfeasible levels of leverage. In this paper, we show that imposing realistic (i.e., in line with institutional practices) leverage restrictions on factor-associated portfolios allows us to uncover pricing factors with desirable properties and superior out-of-sample pricing performance.

In particular, we develop a new methodology for estimating empirical SDFs under constraints on maximum position leverage.[2] Such constraints are economically motivated as (i) they relate to margin constraints as featured in asset pricing models (e.g., Gârleanu and Pedersen (2013)) and (ii) maximum leverage levels are often institutionally mandated in the FX market (e.g., the Commodity Futures Trading Commission prescribes that retail investors are restricted to a maximum leverage of 1:50 in their currency trades, while the Japanese Financial Services Agency restricts retail and corporate FX traders' maximum leverage to 1:25). We use our approach to estimate empirical out-of-sample pricing factors

[1]Currency carry, or the carry trade, exploits average return differences between high- and low-interest-rate currencies (see, e.g., Burnside, Eichenbaum, Kleshchelski, and Rebelo (2011a) and Lustig, Roussanov, and Verdelhan (2011)). Currency momentum involves buying recently appreciated currencies and selling depreciated ones (see, e.g., Okunev and White (2003) and Menkhoff, Sarno, Schmeling, and Schrimpf (2012b)). Currency value is implemented by buying relatively undervalued currencies and selling overvalued ones (see, e.g., Menkhoff, Sarno, Schmeling, and Schrimpf (2016)). Practitioners also mainly focus on these three strategies (Pojarliev and Levich (2008)).

[2]Our methodology allows us to estimate an empirical SDF for a given leverage limit such as 1:25 or 1:50. Hence, we typically refer to empirical SDFs (plural) rather than an (individual) empirical SDF. However, as we show later in the paper, the empirical SDFs and their performance are very similar across a wide range of leverage limits.

and demonstrate their superior pricing performance on a large cross-section of currency returns. Importantly, factor parameters and risk prices are estimated on a rolling basis on different samples, in line with the suggestion of Martin and Nagel (2019) that out-of-sample tests are those that are economically relevant. In a similar spirit, we conduct supplementary asset pricing tests on cross-sections of assets that were not used for empirical SDF estimation (e.g., FX hedge funds), finding consistent and statistically significant risk price estimates.

Our main empirical result is summarized by Figure 1, which plots (for the cross-section of currency carry, momentum, and value portfolios) the predicted mean returns implied by our out-of-sample empirical SDF with a 1:50 leverage constraint along the horizontal axis and realized mean returns along the vertical axis. Our pricing factor jointly prices this cross-section of currency returns, delivering an economically meaningful and highly statistically significant risk price estimate of -0.207 and a close model fit ($\mathrm{R}^2 = 0.91$), which compares favorably to the best-performing alternative factor models found in the literature. For instance, on the same cross-section, the FX-based model of Lustig et al. (2011) and the financial-intermediary model of He, Kelly, and Manela (2017) deliver R^2s of only 0.55 and 0.45, respectively. Moreover, the mean absolute pricing error produced by either of these two models is around twice as large as that of our pricing factor. It is worth noting that our constrained out-of-sample empirical SDF's pricing performance is largely insensitive to the choice of leverage constraints, provided the maximum leverage levels are economically meaningful.[3] Our out-of-sample empirical SDFs present an ambitious and parsimonious benchmark for candidate currency risk models.

As a base, we use the sample version of the minimum-dispersion SDF, as studied by Stutzer (1996), Almeida and Garcia (2017), and Ghosh, Julliard, and Taylor (2019).[4] There are several advantages of minimum-dispersion SDFs. First, like the HJ SDF, they can be recast as optimal portfolio problems, which have a convenient interpretation and are easy to solve numerically. Second, unlike the HJ SDF, they account for higher moments of returns, rendering them better suited for the analysis of skewed and leptokurtic data such as currency returns, especially if higher moments command a risk premium (see, e.g., Brunnermeier, Nagel, and Pedersen (2009)). We add to the minimum-dispersion SDF a

[3]For example, for leverage levels between 1:15 and 1:100, R^2 ranges from 0.80 to 0.90.

[4]Almeida and Garcia (2017) develop a general framework for the construction of empirical SDFs that minimize alternative risk (dispersion) measures that are able to account for the higher moments. For each choice of dispersion measure (e.g., entropy), the empirical minimum-dispersion SDF has the lowest level of dispersion and perfectly prices the set of assets from which it is constructed. The HJ SDF is a special case in the Almeida and Garcia (2017) SDF family, which also encompasses the approaches of Stutzer (1996) and Ghosh et al. (2019).

tractable constraint on the sum of the absolute values of portfolio weights, which can be interpreted as a maximum leverage constraint. The constraint is economically motivated, as leverage limits are ubiquitous in practice and predominately fall within a fairly narrow range of between 1:2 and 1:100 (we provide supporting evidence in Section 2). Thus, given the optimal portfolio formulation of the SDF problem it makes economic sense to limit the maximum leverage that can be deployed by a marginal currency investor.[5]

However, from the in-sample fit perspective, a binding leverage constraint inevitably introduces mispricing. The central question is whether imposing leverage constraints at levels typically observed in practice induces economically and statistically significant mispricing. We address this question in two stages. First, in the leverage-constrained SDF construction, we ensure that our SDFs mimic the key features of the unconstrained minimum-dispersion SDF within their estimation samples (i.e., are "closest" to a fully-pricing SDF), thus explicitly aiming to minimize any mispricing. Second, we evaluate at which maximum leverage levels are our leverage-constrained empirical SDFs statistically indistinguishable from an SDF that perfectly prices the cross-section of currency portfolio returns. We find that, in the full sample, it is the case at leverage levels above 1:45, which roughly coincides with the typical observed maximum leverage limits. Therefore, we conclude that the leverage constraints that are in line with institutional setting of the FX market are a valuable guideline for creating benchmark FX SDFs.

Utilizing our methodology, we construct out-of-sample empirical SDFs for different maximum leverage levels. The procedure involves, at each point in time, estimating an SDF on a rolling window, and then using the estimated parameters and subsequent period's observed returns to calculate that period's out-of-sample SDF realization. Our construction is purely out-of-sample, conditioning only on the information available at the time of the empirical SDF estimation. Before examining empirical SDFs estimated on the full cross-section of currency returns, we first, as a sanity check, examine empirical SDFs estimated solely on the five carry portfolios. This approach allows us to test our methodology on a well-studied cross-section of currency returns, which existing factors are already able to explain, before applying it to a wider cross-section where there are no clear priors. Reassuringly, we find that the estimated portfolio weights for our out-of-sample empirical SDFs, estimated only on the carry portfolios, are consistent with economic intuition (e.g., the two extreme carry portfolios typically have the largest weights and enter the SDFs with opposite signs, consistent with the patterns reported by Lustig et al. (2011)).

[5] We emphasize that our methodology and setting are not suited for studying the impact of leverage constraint changes on asset prices, a question better handled with granular data (see, e.g., Lu and Qin (2021)).

Having ascertained that our methodology performs as expected on carry portfolios, we take it to a richer cross-section of currency returns. Our estimated empirical SDFs display reasonable time-series properties: their average annualized standard deviation is roughly 150%, and they peak around known currency crises. In contrast, we find that the unconstrained empirical out-of-sample minimum-dispersion SDF displays extreme time-series properties (its annualized volatility is approximately 1140%).[6] Importantly, our empirical SDFs' most prominent spikes are associated with episodes of adverse currency shocks, and not business cycle ones, a finding consistent with the disconnect between exchange rates and macroeconomic variables documented in the literature (e.g., Meese and Rogoff (1983), Itskhoki and Mukhin (2021)). We find that a large fraction of the variation in our out-of-sample empirical SDFs appears different from the existing factors. We also observe that estimated optimal portfolio weights do not correspond to static high-minus-low strategies typically adopted in pricing factor construction in the literature. This suggests that the high-minus-low factor formation may be disregarding important pricing information.

Equipped with estimates of our empirical SDFs, we conduct standard asset pricing tests. In the restricted cross-section of only the carry portfolios, we find that our empirical out-of-sample SDFs perform marginally better than the existing factors, but any performance differences are slight, due to the fact that existing factors, namely the high-minus-low factor of Lustig et al. (2011); the FX volatility factor of Menkhoff, Sarno, Schmeling, and Schrimpf (2012a); and the intermediary factor of He et al. (2017), already adequately explain the cross-section of carry returns. However, in the full cross-section of currency portfolio returns, we find that none of the existing factors can effectively explain it. In contrast, our out-of-sample empirical SDFs consistently price the cross-section and deliver an excellent fit, as illustrated in Figure 1. In a larger cross-section, the unconstrained out-of-sample empirical SDF tends to repeatedly overfit within the samples it is recursively estimated on, thus displaying extreme volatility, and consequently performing poorly in asset pricing tests. These findings highlight that our economically-motivated restrictions are important for extracting effective pricing factors.

We also evaluate our out-of-sample empirical SDFs (estimated on the cross section of currency portfolios) on cross-sections of individual currencies and FX hedge funds. We find that our empirical SDFs are priced in the large cross-section of individual currencies, and even in the smaller cross-section of developed country currencies. Similarly, we find that our empirical SDFs are priced in the cross-section of FX hedge funds (sourced from

[6] It is important to note that we cannot simply rescale the unconstrained empirical SDF to reduce its ex-post volatility, as both the unconstrained and constrained empirical SDFs have to respect the condition, $E\left[M_t R_f\right] = 1$, where M_t is the candidate SDF and R_f is the gross risk-free return.

two different databases), and their risk price estimates are similar to the ones estimated on the cross-section of carry, momentum, and value portfolios. In contrast, we do not find statistically significant risk prices for any of the existing factors when tested on this cross-section. FX hedge funds are essentially dynamically-managed currency strategies, however, we, as researchers, have no control over their portfolio compositions. This feature renders the FX hedge fund cross-section an ideal out-of-sample laboratory for the evaluation of pricing factors. Thus, finding that our empirical SDFs perform well in this setting suggests that they are indeed capturing key elements of currency risk premiums.

Given the empirical out-of-sample SDFs' superior pricing performance, they could be informative about the optimal currency portfolio. We indeed find that investable versions of our SDFs have desirable return properties. For example, between 1991 and 2020, an investable SDF portfolio of our out-of-sample empirical SDF with 1:50 leverage constraint delivers an annual Sharpe ratio of around 0.8 after transaction costs, and returns that are positively skewed. Moreover, this investable SDF portfolio has the lowest maximum drawdown compared to the basic currency strategies. Our investable SDF portfolios' performance are similar to optimal currency portfolios proposed by, for example, Berge, Jordà, and Taylor (2010) and Barroso and Santa-Clara (2015). It is worth noting, however, that the post-2008 period (the time when most currency strategies performed rather poorly) makes up a relatively large proportion of our sample.

Lastly, we conduct a battery of robustness tests. Specifically, we consider different estimation windows and different frequencies for updating SDF weights. The estimated risk price remains virtually unchanged, and the pricing performance is similar across specifications. It is particularly worth noting that the asset pricing performance remains strong even when SDF weights are only updated every 24 months, highlighting the out-of-sample viability of our methodology. We also consider an alternative specification for the constraint in our estimation procedure. We find that we can generate similar pricing results using a quadratic ridge constraint (that implicitly penalizes large, concentrated positions) in place of our leverage constraint. However, we find similar results only when we map our benchmark leverage limits onto the choice parameter governing ridge penalties (e.g., a ridge parameter of 484 produces portfolio weights that satisfy the 1:50 leverage limit in most of the sample). This exercise demonstrates the robustness of our procedure as we converge on similar empirical SDF estimates even when starting from a different optimization problem. It also highlights the importance for regularized machine learning methods to have a clear economic interpretation to better extract meaningful estimates from a profusion of potential specifications, as is also empathized by Nagel (2021).

Related literature Our work relates most to the large literature on the risk-based explanation for the profitability of currency trading strategies, particularly the cross-sectional studies.[7] Burnside, Eichenbaum, and Rebelo (2007) and Burnside et al. (2011a) report the profitability of simple currency trading strategies, specifically the carry trade, and suggest that neither the standard risk factors nor the "peso-problem" explain the carry trade's historic profitability. Brunnermeier et al. (2009) also document the carry trade's profitability, suggesting that it may be related to crash risk, however they do not conduct formal asset pricing tests. Subsequent studies, mainly examining FX options, find that crash risk accounts for up to a third of currency risk premiums (Jurek (2014), Farhi, Fraiberger, Gabaix, Ranciere, and Verdelhan (2015), and Chernov, Graveline, and Zviadadze (2018)). Relatedly, Della Corte, Ramadorai, and Sarno (2016) and Della Corte, Kozhan, and Neuberger (2021) find the existence of currency volatility premiums. If higher moments are important for pricing, it would be reflected in our empirical SDF, as our methodology is tailored to extracting pricing factors with nonlinear properties. The superior pricing performance of our empirical SDF suggests that non-linearity may play a role.

There are a number of studies offering explanations for the cross-section of currency portfolio returns. Lustig et al. (2011) propose an empirical two-factor model, comprised of the dollar and the high-minus-low carry factor, finding that it explains the average monthly carry portfolio returns. Menkhoff et al. (2012a) build on that work, showing global currency volatility innovations to be the key factor for explaining the currency cross-section. Relatedly, there is evidence that aggregate FX liquidity may also influence expected currency returns (e.g., Mancini, Ranaldo, and Wrampelmeyer (2013), Karnaukh, Ranaldo, and Söderlind (2015)). Lettau, Maggiori, and Weber (2014) find that a conditional downside risk capital asset pricing model can explain the cross-section of carry portfolios. In turn, He et al. (2017) find that their intermediary capital risk factor is priced in the cross-section of currency portfolio returns. Most existing studies focus predominantly on carry portfolios. It has been shown, however, that currency momentum and value strategies present additional, distinct currency risk premiums (Menkhoff et al. (2012b), Asness, Moskowitz, and Pedersen (2013), Menkhoff et al. (2016)). Hence, we jointly examine all the primary currency trading strategies, finding that the well-established factors are unable to explain the wide cross-section of currency returns. In contrast, our empirical out-of-sample SDFs price the cross-section, suggesting that adequate SDF projections are accessible to investors, but existing factors are not able to fully capture them. We also differ significantly from the aforementioned studies in our methodological approach.

[7] The FX literature is too vast for us to cover fully. We thus discuss the most related works, and refer to the literature surveys for additional references (see, e.g., Burnside (2012), Hassan and Zhang (2020)).

Methodologically, our study relates to the large literature on extracting SDFs from the data, an idea pioneered by Hansen and Jagannathan (1991), using the minimum-variance SDF, and generalized by Almeida and Garcia (2017) to other risk measures. Most studies using these techniques focus on empirical SDFs' properties, such as their variance bounds, rather than focusing on their out-of-sample asset-pricing performance. In this strand of the literature, our work is closest to Ghosh et al. (2019), who estimate unconstrained minimum-dispersion SDFs and evaluate their out-of-sample pricing performance on a cross-section of stock portfolio returns. We focus exclusively on currency returns and show that unconstrained empirical SDFs perform poorly in shorter time series, as found in FX research. Our methodological contribution of introducing a leverage constraint to SDF estimation renders this approach better suited to currencies, and connects our work to recent papers that show that leverage constrains play a key role in explaining certain asset pricing puzzles (see, e.g., Jylhä (2018), and Lu and Qin (2021)). The idea of imposing economically-plausible restrictions in SDF estimation also relates to the work of Korsaye, Quaini, and Trojani (2019) that proposes a general theory of estimating regularized SDFs with an economic interpretation of the regularization device; the work of Kozak, Nagel, and Santosh (2020) in which they impose economically motivated priors when estimating SDF coefficients using equity return data; and to the work of Haddad, Kozak, and Santosh (2020), in which they use a bound on the average squared Sharpe ratio of their optimal dynamic trading strategy as a regularization device to extract a better-fitting empirical SDF in the cross-section of stocks.

In analyzing the investment performance of the traded version of our empirical SDFs, we relate to the work on optimal currency portfolios (e.g., Berge et al. (2010), Jordà and Taylor (2012), Barroso and Santa-Clara (2015), Bekaert and Panayotov (2020), Maurer, To, and Tran (2020), and Opie and Riddiough (2020)). However, we differ from these studies in that currency portfolio optimization is not our main goal, but is simply an outcome of finding an empirical SDF with good pricing power.

In a broad sense, our work relates to the international asset pricing literature analyzing FX risk from the perspective of multiple international investors who trade assets across borders. Maurer, Tô, and Tran (2019) construct projections of country-specific SDFs on their common components under the assumption of complete and integrated capital markets and examine their properties and ability to price equity risk. Similarly, Sandulescu, Trojani, and Vedolin (2021) and Korsaye, Trojani, and Vedolin (2020) study the implications of frictions in international capital flows for the properties of the resulting extracted SDFs. These papers, however, do not speak directly to explaining the return cross-section of active FX strategies, while our work does. In this strand of the literature, our work is closest to Chernov, Dahlquist, and Lochstoer (2020), who construct a conditional, linear projection of the SDF (using the approach of Hansen and Richard (1987)) on the space of G10 currency returns. Our studies are, however, complementary, as Chernov et al. (2020) focus exclusively on the time series dimension, while our approach is mostly cross-sectional.

Lastly, our study relates to the rapidly growing literature that applies machine learning to asset pricing problems (see Nagel (2021) for a textbook treatment). Our work is closest in spirit to Chen, Pelger, and Zhu (2020), who argue that it is important to use economically-motivated restrictions when adapting machine learning tools to asset pricing questions.

2 Institutional background of maximum leverage constraints

In this section, we briefly review the evidence on the maximum leverage used in the FX market. Our goal is, first, to ascertain that leverage limits are important, and, second, to identify a typical range of leverage limits used in practice. While the U.S. Securities and Exchange Commission (SEC) has only permitted a 1:2 leverage on long positions in U.S. stocks since 1934, leverage in the FX market has typically been much higher. Nevertheless, the relative high leverage afforded to currency traders is still limited. Since 2010, the U.S. Commodity Futures Trading Commission (CFTC) prescribes that retail investors are restricted to a maximum leverage of 1:50 on all major currency pairs and 1:20 on others. While retail leverage could exceed 1:50 before the CFTC imposed the regulatory limit, Heimer and Simsek (2019) report that before 2010 it was the case for only around 16% of FX traders. Similarly, the Japanese Financial Services Agency restricts both retail and corporate FX traders' maximum leverage to only 1:25. Ang, Gorovyy, and Van Inwegen (2011) report maximum leverage limits for different markets: the maximum leverage in FX futures and FX swaps, the two markets most likely to be used by professional traders, is limited to 1:25 and 1:100, respectively.

Finally, even considering the potential leverage of sophisticated financial intermediaries like investment banks and hedge funds, who may be playing an important role in the FX market (e.g., Gabaix and Maggiori (2015)), we find that leverage limits are still relevant. We examine the implied leverage levels (defined as total assets over market equity) of 29 key global investment banks. The Internet Appendix plots the time series of investment bank leverage. We find that, over the period 1999–2020, the average (median) implied leverage is around 1:21 (1:16), and the 10^{th} and 90^{th} percentiles are 1:9 and 1:39, respectively.[8]

[8] These reported statistics are the time-series averages of the cross-sectional statistics.

Leverage levels do vary over time, nevertheless, in this cross-section of investment banks, the 90$^{\text{th}}$ percentile of leverage levels has exceeded the 1:100 level only once, in the third quarter of 2008 (at that time, the average and median leverage levels were 1:57 and 1:43, respectively, and the 90$^{\text{th}}$ percentile was 1:125). Since then, however, Basel III banking regulation introduced a 3% minimum bank capital requirement, which implies a maximum leverage constraint of around 1:33 for large financial institutions (see, e.g., Boyarchenko, Eisenbach, Gupta, Shachar, and Van Tassel (2020)). Similarly, the literature finds that observed hedge fund leverage is fairly limited. Using a proprietary database, Ang et al. (2011) report that average hedge fund leverage is only around 1:2, but could reach 1:8 for some strategies. Using comprehensive regulatory data, Barth, Hammond, and Monin (2020) corroborate the results of Ang et al. (2011), finding that average hedge fund balance sheet leverage is around 1:2. Barth et al. (2020) also examine leverage levels implied by notional exposure, finding the 90$^{\text{th}}$ percentile of global macro hedge funds' leverage levels to be around 1:16, which, although higher than the average, is still well below, for example, the maximum regulatory leverage levels for FX futures.[9]

In sum, irrespective of how one measures it, maximum available leverage in the FX market typically lies between 1:10 and 1:50, and very rarely exceeds 1:100. It is thus reasonable to consider this range of leverage limits as the most economically meaningful and use it to guide our analysis.

3 Theoretical background and methodology

In this section, we discuss the construction of data-driven SDFs under constraints on leverage. First, we briefly review the theory of constructing unconstrained SDFs. We discuss their out-of-sample pricing properties and argue that the results call for introducing the aforementioned constraints. Lastly, we show how to include the constraints while monitoring the constrained SDFs' in- and out-of-sample pricing properties.

3.1 Non-parametric stochastic discount factors

Our starting point is the exponential tilting SDF previously studied by Stutzer (1996); Branger (2004); Almeida and Garcia (2017), and Ghosh et al. (2019). The exponential

[9] It is worth noting that the lion's share of hedge fund leverage stems from their prime brokers, which are typically large investment banks that are subject to strict maximum leverage requirements (Barth et al. (2020), Boyarchenko et al. (2020)), hence maximum hedge fund leverage is unlikely to be significantly larger than their observed leverage.

tilting SDF is a member of a larger family of minimum-dispersion SDFs, defined by Almeida and Garcia (2017). We briefly describe its construction.

Consider the following problem. let R_t be an $N \times 1$ vector of gross returns on risky assets, and R_f the gross risk-free rate. In general, there exists an infinite number of SDFs M_t that price all the assets. We focus on an SDF with a particular extremal property: an SDF $M_t^\star$ such that

$$M_t^\star = \underset{\{M_t\}}{\arg\min} \quad E\left[M_t \ln M_t\right] \tag{1}$$

subject to

$$E\left[M_t(R_t - R_f)\right] = 0 \tag{2}$$

$$E\left[M_t R_f\right] = 1\,, \tag{3}$$

where R_f is the risk-free rate of return. The minimization criterion in (1) is known as the Kullback-Leibler (KL, hereafter) discrepancy, and it is a measure of the variability of positive random variables, which is similar to variance but takes higher moments into account. Solving the problem (1)-(3) yields a non-parametric estimate of the true SDF, which is a function of the underlying asset returns R_f and R_t, and which is always positive.[10] However, directly finding $M_t^\star$ by solving problem (1)-(3) is difficult because $M_t^\star$ is an infinite-dimensional object. Stutzer (1996) and Almeida and Garcia (2017) demonstrate how convex duality relationships stated by Borwein and Lewis (1991) can be used in order to reformulate problem (1)-(3) as a finite-dimensional optimal portfolio problem.[11]

Specifically, Stutzer and Almeida and Garcia show that the KL discrepancy-minimizing SDF $M_t^\star$ takes the form

$$M_t^\star = \frac{e^{-\theta_0^\star R_f - \sum_{j=1}^N \theta_j^\star (R_{jt} - R_f)}}{R_f E\left[e^{-\theta_0^\star R_f - \sum_{j=1}^N \theta_j^\star (R_{jt} - R_f)}\right]}\,, \tag{4}$$

where the $\theta_j^\star$ are optimal wealth allocations in the portfolio choice problem

$$\theta^\star = \underset{\theta \in \mathbb{R}^{N+1}}{\arg\max} \quad -\theta_0 - E\left[e^{-\theta_0 R_f - \sum_{j=1}^N \theta_j^\star (R_{jt} - R_f)}\right]\,. \tag{5}$$

[10]Among all possible M_t that price the assets, the SDF $M_t^\star$ exhibits the lowest KL discrepancy. It is nonetheless possible to find, e.g., an SDF M_t^{HJ} that exhibits the lowest variance (the HJ SDF). The KL discrepancy of M_t^{HJ} will be, by construction, not lower than the discrepancy of $M_t^\star$, with equality holding if, and only if, the joint distribution of the returns on risky assets is log-normal. Both SDFs will price the risky and riskless assets, but they will have different time-series properties.

[11]Ghosh et al. (2019) study problem (1) without constraint (3).

3.2 Finite-sample version of the KL SDF and its properties

The duality between the SDF construction problem (1) and its portfolio formulation (5) also holds in finite samples, as long as the number of observation periods T is sufficiently greater than the number of assets N. To see that, note that for $\theta_T^\star$ that solves the in-sample portfolio problem,

$$\theta_T^\star = \underset{\theta \in \mathbb{R}^{N+1}}{\arg\max} \quad -\theta_0 - \frac{1}{T}\sum_{t=1}^{T} e^{-\theta_0 R_f - \sum_{j=1}^{N} \theta_j^\star (R_{jt} - R_f)}\,, \tag{6}$$

the first-order conditions are equivalent to the sample analogues of the asset-pricing moment conditions

$$\frac{1}{T}\sum_{t=1}^{T}(R_{jt} - R_f)e^{-\theta_{T0}^\star R_f - \sum_{j=1}^{N} \theta_{Tj}^\star (R_{jt} - R_f)} =: \frac{1}{T}\sum_{t=1}^{T}(R_{jt} - R_f) M_t^{T,\star} = 0\,, \tag{7}$$

i.e., $M_t^{T,\star}$, constructed by plugging $\theta_T^\star$ into equation (4) and replacing the expectation therein with a sample average, exactly prices the assets in-sample.

However, even if the true optimal portfolio allocations $\theta^\star$ were known, the SDF $M_t^\star$ would display some degree of pricing error on a sample of test asset returns purely due to random noise in SDF and return realizations. When optimal allocations $\theta_T^\star$ are estimated, they are additionally affected by estimation error, which adds to the observed pricing error when the SDF is evaluated out of sample. Furthermore, when the sample length T is small relative to the size of the cross-section N, the estimated allocations $\theta_T^\star$ are likely to suffer from overfitting when estimation errors are large and strongly correlated, which manifests itself in the extreme nature of the estimated allocations.[12] These large allocations then translate directly into extreme time-series properties of the SDF which are economically implausible. When taken out of sample, the compounding of these errors leads to very poor pricing performance.

We present the unconstrained out-of-sample SDF estimated at a monthly frequency (i.e., with an out-of-sample window of 1 month) on a rolling window of 180 months in Figure IA.2 in the Internet Appendix. Indeed, we find that its time-series properties are

[12]Overfitting is a situation where a model fits the data well (e.g., HJ's SDF delivers zero in-sample pricing errors), but, as a negative side effect of obtaining the perfect or near-perfect in-sample fit, the model's predictions (e.g., the SDF's time-series) are extremely volatile (see, e.g., Ghojogh and Crowley (2019)). In portfolio selection applications, overfitting is typically exacerbated if assets share a dominant common factor, like currency portfolios (Lustig et al. (2011)), or if the cross-section of assets is large (see, e.g., Jagannathan and Ma (2003) for a discussion).

undesirable. First, the SDF peaks at a value of above 50, which implies a 50-fold rise in the marginal utility of the marginal investor between two periods, a value which is beyond the predictions of standard parametrizations of utility functions combined with observed fluctuations in investor wealth and consumption. Second, the SDF's annualized standard deviation is 1140% per year. While this is not the HJ minimum-variance SDF and, hence, its variance is not the maximum attainable Sharpe ratio, its value, nevertheless, indicates that the investor following the optimal trading strategy should be able to obtain an extremely high Sharpe ratio, beyond anything typically reported in practice. Third, and most importantly, the leverage embedded in the optimal portfolio that drives the SDF is never lower than 1:170 and peaks at levels higher than 1:1000.

In order to address these difficulties in the constrution of data-driven pricing factors, we turn the evidence on the leverage limits in FX markets. In particular, we introduce a leverage-constrained version of the SDF implied by the portfolio selection problem and we examine whether the constrained SDF has good in- and out-of-sample pricing performance. Our proposed solution draws inspiration from the regularized/constrained portfolio optimization literature (see, e.g., Jagannathan and Ma (2003); DeMiguel, Garlappi, Nogales, and Uppal (2009)). Moreover, in addition to being economically motivated, our solution also follows in the footsteps of the machine learning and the econometrics literatures, where a natural remedy to increases in problem dimensionality is the use of regularization (which induces a trade-off between in-sample model fit and model complexity). Simpler models often perform significantly better when confronted with new data.

3.3 Leverage-constrained empirical SDF

Consider a reformulation of the in-sample portfolio optimization problem (6) which adds a leverage constraint on the risky asset allocations in the vector θ:

$$\theta_T^L := \underset{\theta \in \mathbb{R}^{N+1}}{\arg\max} \quad -\theta_0 - \frac{1}{T}\sum_{t=1}^{T} e^{-\theta R_f - \sum_{j=1}^{N} \theta j^\star (R_{jt} - r_f)} \tag{8}$$

subject to

$$\sum_{j=1}^{N} |\theta_j| \leq L\,. \tag{9}$$

This constraint is related to the one in Gârleanu and Pedersen (2013).[13]

[13]The constraint in the Gârleanu and Pedersen (2013) model is as follows: $\sum_{j=1}^{N} m_t^j\,|\theta_j| + \eta^u \leq 1$, where m_t^j is each asset j's margin at time t, and the η^u is the fraction of wealth invested in uncollateralized loans.

Whenever the constraint (9) is binding, the wealth allocations θ_j are restricted. As a consequence, the candidate in-sample SDF defined as

$$M_t^{T,C}(\theta) = e^{-\theta_{T,0}^C R_f - \sum_{j=1}^N \theta_{T,j}^C (R_{j,t} - R_f)} \tag{10}$$

must exhibit more plausible time-series variation than the unconstrained SDF. However, at the same time, the value of the objective function in the constrained portfolio allocation problem (8)-(9) must be lower than in the unconstrained problem (6). This observation implies, via convex duality relationships, that the candidate in-sample SDF $M_t^{T,L}$ does not meet the sample analogue of the pricing moment condition (7) and delivers non-zero in-sample pricing errors. The magnitude of the pricing errors must be evaluated in order to determine whether imposing leverage constraints based on the presented institutional evidence is compatible with reasonable SDF properties and good pricing performance.

3.3.1 Kullback-Leibler SDF distance

In order to evaluate the magnitude of these pricing errors, we use the Kullback-Leibler (KL) SDF distance introduced by Almeida and Garcia (2012). The KL distance is closely related to the HJ distance, however the KL distance accounts for higher-order moments of the candidate and exact in-sample SDFs. The KL SDF distance is calculated by finding an additive correction Π_t to the candidate SDF $M_t^{T,L}$ that eliminates in-sample pricing errors and has the lowest KL discrepancy among all such corrections. For a candidate SDF $M_t^{T,L}$, the KL SDF distance takes the form

$$\delta_{KL}(\theta, L) = \inf_{\Pi} \frac{1}{T} \sum_{t=1}^{T} (\Pi_t + 1) \ln (\Pi_t + 1) \tag{11}$$

subject to

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^{T} (M_t^{T,L}(\theta) + \Pi_t)(R_t - R_f) &= 0 \\ \frac{1}{T} \sum_{t=1}^{T} (M_t^{T,L}(\theta) + \Pi_t) R_f &= 1\,. \end{aligned} \tag{12}$$

In this problem, the magnitude of the pricing errors of the candidate SDF $M_t^{T,L}$ is evaluated by finding a correction Π_t to the candidate SDF $M_t^{T,L}$ that results in an exactly-pricing SDF $M_t^{T,L} + \Pi_t$, which is closest to $M_t^{T,L}$ in the sense of the KL distance. Similarly to the duality between problems (1)-(3) and (5), Almeida and Garcia (2012) show that the distance δ_{KL} and the correction Π_t in problem (11)-(12) can be found through its dual:

$$\delta_{KL}(\theta, L) = \max_{\lambda} -\lambda_0 - \frac{1}{T}\sum_{t=1}^{T}\left[e^{-\lambda_0 R_f - \sum_{j=1}^{N}\lambda_j(R_{jt}-R_f)} + \left(M_t^{T,L}(\theta) - 1\right)\left(-\lambda_0 R_f - \sum_{j=1}^{N}\lambda_j(R_{jt} - R_f)\right)\right]. \tag{13}$$

Thus, for any candidate SDF $M_t^{T,L}$, we can determine how large a correction is necessary to ensure exact in-sample pricing. Naturally, the correction is smaller for $M_t^{T,L}$ with good in-sample pricing than for one allowing for large in-sample pricing errors.

3.3.2 Minimum-distance estimation of the leverage-constrained SDF

Once the in-sample pricing performance of a constrained candidate SDF $M_t^{T,L}$ can be appropriately evaluated, the logical consequence is to think about the best candidate SDF $M_t^{T,L}$ among all those which satisfy a given leverage constraint. We consider the family of candidate SDFs $M^{T,L}(\theta; L)$ of the form (10) for a given leverage constraint parameter L as in equation (9), with the requirement that the SDFs exactly price the risk-free asset. In this family, indexed by the portfolio allocation vector θ_T^L, the leverage constraint (9), if binding, introduces regularization which results in non-zero in-sample pricing errors. The best candidate SDF $M_t^{T,L}$ must thus be the one that requires the smallest correction Π_t in order to meet the in-sample moment condition while satisfying the leverage constraint. Solving this problem amounts to minimizing, over θ, the distance between the candidate SDF family and the set of SDFs that price the assets. This leads us to the following formulation of the problem.

Definition 1. *For a maximum leverage level $L \geq 0$, the leverage-constrained SDF $M^{T,C}(\widehat{\theta}^C; L)$ is defined as*

$$M^{T,C}(\widehat{\theta}^C; L) = e^{-\widehat{\theta}_0^C R_f - \sum_{j=1}^{N}\widehat{\theta}_j^C\left(R_{jt} - R_f\right)}, \tag{14}$$

where $\widehat{\theta}^C$ *solves the problem*[14]

$$\widehat{\theta}^C = \arg\min_{\theta} \, \delta_{KL}(\theta, L) \tag{15}$$

subject to

$$\sum_{j=1}^{N} |\theta_j| \leq L \tag{16}$$

$$\frac{1}{T}\sum_{t=1}^{T} M_t^{T,C} R_f = 1 \tag{17}$$

The correction Π_t *to* $M_t^{T,L}$ *which nullifies in-sample pricing errors is*

$$\Pi_t = e^{-\lambda_0 R_f - \sum_{j=1}^{N} \lambda_j (R_{jt} - r_f)} - 1 \, . \tag{18}$$

In Definition 1 above, the KL distance criterion was chosen for the following reasons. First, consider $L = 0$. Then $\widehat{\theta}^C = 0$, and the problem of finding the SDF correction Π_t reduces to the optimization problem (6). Second, considering L that is not binding in problem (8)-(9), $M_T^L(\theta, L)$ exactly prices the primitive assets R_t. Thus $\Pi_t \equiv 0$, $\delta_{KL} \equiv 0$ and we recover exact in-sample fit.

The remaining question is how to assess whether the in-sample mispricing of $M^{T,C}$ is significant, or, in other words, whether $M^{T,C}$ is in fact distinguishable from an SDF that exactly prices the assets. In order to answer this question, we further build on the results in Almeida and Garcia (2012), who demonstrate in Theorem 4 of their paper that the vector of optimal portfolio positions θ_C^T and the corresponding vector $\lambda(M^{T,C}(\theta_C^T))$ of the SDF correction Π_t asymptotically follow a joint Gaussian distribution. Based on this result, we propose a test for the distinguishability of $M^{T,C}$ from the set of exactly pricing SDFs, which is based on the properties of the correction Π_t. More specifically, we aim to test the following null hypothesis:

$$\forall \, t \quad \Pi_t = 0 \iff \lambda(M^{T,C}) = 0 \, . \tag{19}$$

Note that there are two sources of noise in $\widehat{\lambda}$. First, for a given candidate SDF $M^{T,C}$, the sampling variation in returns induces variation in $\widehat{\lambda}(M^{T,C})$. Second, the sampling variation in returns also induces variation in $\lambda(M^{T,C})$, which then further augments the variability

[14]Note that, in this optimization problem, we incorporate the leverage limit explicitly as a constraint rather than converting it to a Lagrangean (penalization) term, as is habitual in introducing regularization. This approach renders the introduction of the leverage constraint easier to interpret.

of $\widehat{\lambda}$. In the test, we focus on the pure sampling variability of $\widehat{\lambda}$, treating $M^{T,C}$ as given.[15] The magnitude of $\lambda(M^{T,C})$ informs us of a given SDF candidate's distance from the set of SDFs that indeed exactly price the returns. Focusing on $\widehat{\lambda}$'s variability yields an estimate of the uncertainty of this set of SDFs' distance from a given candidate. Thus, we state the following proposition based on Theorem 4 in Almeida and Garcia (2012).

Proposition 2. *For a given candidate SDF $M^{T,C}$, the associated vector $\widehat{\lambda}(M^{T,C})$ of estimated loadings of the SDF correction Π_t follows a Gaussian distribution,*

$$\widehat{\lambda}(M^{T,C}) \sim N\left(\lambda(M^{T,C}), \Sigma\right) . \tag{20}$$

In the above equation, $\Sigma = H_\lambda^{-1} S_\lambda H_\lambda^{-1}$, where S_λ is the variance-covariance matrix of the first-order condition of the optimization problem (13) *and H_λ is the Hessian of the optimization objective of problem* (13)*.*

The immediate consequence of this result is that a feasible test of the hypothesis stated in equation (19) can be constructed. The test statistic follows a χ^2 distribution with $N+1$ degrees of freedom under the null $\lambda(M^{T,C}) = 0$:

$$\widehat{\lambda}(M^{T,C})'\widehat{\Sigma}^{-1}\widehat{\lambda}(M^{T,C}) \sim \chi^2(N+1)\,, \tag{21}$$

where $\widehat{\Sigma}$ is an HAC estimate of the optimization score vector's variance, Σ, defined above.

3.3.3 Discussion of the leverage constraint

If one regards θ_j as portfolio allocations, the constraint (9) has a clear interpretation in terms of a limitation on the investor's total position. The FX market strategies we consider are zero-cost forward trading strategies, which are typically subject to margin requirements. Typically, the literature considers a margin requirement such that for every \$1 of collateral, the investor can buy forward \$1 worth of foreign currency and, at the same time, sell forward an amount of foreign currency which will yield \$1 of proceeds. In such a strategy, the sum of absolute values of portfolio weights is equal to 2. Thus the strategy employs 1:2 leverage. With this interpretation in mind, the portfolio weight constraint (9) can be directly seen as a leverage constraint. Setting $L = 1$ allows for no leverage, while setting $L \geq 1$ allows

[15]In other words, we examine the distribution of $\widehat{\lambda}$ conditional on the estimated $M^{T,C}$. In this way we focus on the precision of estimating the location of the set of the SDFs that exactly price the test assets in-sample.

for levered positions, with $L = 2$ corresponding to the leverage of a standard FX carry, momentum, or value strategy, as considered in the literature.

The functional form of the constraint is such that, whenever it is binding, the coefficients θ_j are shrunk towards zero, and are set exactly to zero if the constraint is tight enough. Thus, if L is sufficiently low, the constrained in-sample SDF candidate M_T^L is spanned by fewer than N assets. As L goes to zero, the assets are consecutively dropped until the SDF candidate becomes identically equal to R_f^{-1}. In the literature on penalized regression, this type of constraint is known as the LASSO penalty and it induces sparsity when binding (e.g., Freyberger, Neuhierl, and Weber (2020) use the adaptive group LASSO to study which characteristics provide incremental information about the cross-section of stock returns).

4 Data, return calculation, portfolio construction, and descriptive statistics

In this section, we discuss our main data sources, describe the calculation of currency returns and portfolio construction, and present key descriptive statistics of currency portfolio returns.

4.1 Currency data

We source daily spot and one-month forward exchange rates relative to the U.S. Dollar from WM/Reuters via Datastream. Monthly data are obtained by sampling end-of-month exchange rates. The sample period runs from February 1976 to October 2020. Our sample includes 44 developed and emerging market currencies, namely those of Australia, Austria, Belgium, Brazil, Bulgaria, Canada, Croatia, Cyprus, Czechia, Denmark, euro area, Finland, France, Germany, Greece, Hungary, Iceland, India, Indonesia, Ireland, Israel, Italy, Japan, Malaysia, Mexico, Netherlands, New Zealand, Norway, Philippines, Poland, Portugal, Russia, Singapore, Slovakia, Slovenia, South Africa, South Korea, Spain, Sweden, Switzerland, Taiwan, Thailand, Turkey, and the United Kingdom. We filter these data following Lustig et al. (2011), and Dahlquist and Hasseltoft (2020).

4.2 Other data

We collect Consumer Price Index (CPI) data from the Organisation for Economic Co-Operation and Development (OECD). The OECD provides homogeneous CPI data for its 27 member countries, the eurozone, and a number of other countries. The intermediary factor of He et al. (2017) and the FX liquidity factor of Karnaukh et al. (2015) are from the authors' websites. The risk-free rate (one-month treasury yield), the return on the U.S. stock market (monthly return of the Center for Research in Security Prices value-weighted market portfolio in excess of the risk-free rate), and Fama and French (1993) and Carhart (1997) factor data are from Kenneth R. French's Data Library. We source the CBOE Volatility Index (VIX) and the TED spread (the difference between the 3-month U.S. Libor rate and the three-month U.S. treasury rate) series from the Federal Reserve Bank of St. Louis' economic database. The J.P. Morgan Global FX Volatility Index (VXY) is from Bloomberg. Finally, the Barclay Currency Traders Index is from BarclayHedge.

4.3 Currency excess returns and currency portfolios

We examine currency returns from the perspective of an U.S. investor. We express exchange rates as USD per foreign currency unit and denote spot and forward exchange rates as S and F, respectively. Therefore, if t is the end of a given month, and $t+1$ is the end of the following month, abstracting from transaction costs, the excess one-month return (in USD) of a long position in currency j is

$$R_{j,t+1} = \frac{S_{j,t+1} - F_{j,t}}{F_{j,t}}. \tag{22}$$

This is a zero-cost strategy, and, when covered interest rate parity (CIP) holds, it is proportional to borrowing \$1 in the U.S. money market and investing the proceeds in the foreign-currency money market for one period before converting back to USD at the spot rate. We use one-month forwards throughout our empirical analysis. Therefore, all returns are monthly and there are no inherited positions from month to month. This avoids path dependency considerations when we include transaction costs. Taking the bid-ask spread into account, the excess one-month return of a long or short position in currency j is

$$R^{\text{long}}_{j,t+1} = \frac{S^{\text{bid}}_{j,t+1} - F^{\text{ask}}_{j,t}}{F^{\text{ask}}_{j,t}} \qquad R^{\text{short}}_{j,t+1} = \frac{F^{\text{bid}}_{j,t} - S^{\text{ask}}_{j,t+1}}{F^{\text{bid}}_{j,t}}, \tag{23}$$

whereby bid and ask quotes are denoted in the superscript.

Closely following the literature, we separately construct five carry, momentum, and value portfolios of currencies. We focus our analysis on currency portfolios for similar reasons outlined in Lustig and Verdelhan (2007). In particular, it allows us to reduce the noise inherent in individual currency changes, to zoom in on key characteristics driving expected returns (namely interest rates, past returns, and real exchange rate changes), and to assess economic significance in a straightforward way.[16] Moreover, using portfolios as test assets in the cross-sectional asset pricing analysis improves the precision and stability of estimated factor loadings. Nevertheless, in additional tests, we also use individual currencies as test assets.

For carry portfolios, at the end of each period t, we allocate all currencies in our sample to five portfolios based on their forward discounts (i.e., F_t/S_t). If the CIP holds, as it has historically (see, e.g., Akram, Rime, and Sarno (2008)), sorting on forward discounts is equivalent to sorting on interest rates as $\ln(F_{j,t}/S_{j,t}) \approx i_t^{\mathrm{US}} - i_t$, where i_t^{US} and i_t are the U.S. and currency j's interest rates, respectively. Portfolios are ranked from low to high foreign interest rates (forward discount); portfolio 1 contains the currencies with the lowest interest rates, and portfolio 5 contains the currencies with the highest interest rates (our approach mimics that of Lustig et al. (2011) and Menkhoff et al. (2012a)).

Independently of the carry portfolios, we construct five currency momentum portfolios. To do so, at the end of each period t, we allocate all currencies in our sample to five portfolios based on each currency's lagged returns over the previous l months (the formation period). Portfolios are ranked from low to high past return; portfolio 1 contains the currencies with the lowest past return (so-called "losers"), and portfolio 5 contains the currencies with the highest past return (so-called "winners"). We focus on the one-month formation period ($l = 1$) and one-month holding period, as Menkhoff et al. (2012b) show that currency momentum returns are similar across formation periods between one to nine months and that shorter holding periods perform better. We do not skip the most recent month when calculating currency momentum returns, as would have been customary in equity momentum strategies, because Asness et al. (2013) find that momentum returns for currencies are stronger without skipping. Our results are not sensitive to these choices in portfolio construction.

Finally, for currency value portfolios, again independently of the other strategies, we allocate the currencies in our sample into five portfolios based on the cumulative real currency depreciation in the previous five years. In particular, the value signal at period t is measured as a log change of real exchange rate between period $t-1$ and $t-61$, and the real exchange rate of country j at t is defined as $q_{j,t} = 1/S_{j,t} \times \mathrm{CPI}_t^{\mathrm{US}} / \mathrm{CPI}_{j,t}$, where $\mathrm{CPI}_t^{\mathrm{US}}$ and $\mathrm{CPI}_{j,t}$ is the CPI of the U.S. and country j, respectively.[17] This is similar to

[16] Relatedly, machine learning methods, which our approach could be categorized as, are typically sensitive to the quality of the input data. It is best practice to thoroughly filter the data (see, e.g., Gu, Kelly, and Xiu (2020)). We believe focusing on portfolios is both a relatively innocuous and widely accepted approach.

[17] Lagging CPI to ensure the CPI is known does not materially alter our results. Moreover, given the increasing ability of market participants to "now-cast", i.e., forecast low frequency data like CPI using

currency value definitions of Barroso and Santa-Clara (2015), Asness et al. (2013), and Menkhoff et al. (2016)). Note that that the real exchange rate is defined such that a higher $q_{j,t}$ means a lower valuation level of currency j relative to USD. At the end of each period t, we sort currencies into five portfolios based on each currency's value signal; portfolio 1 contains the most relatively overvalued currencies, and portfolio 5 contains the most relatively undervalued currencies.[18] All portfolios for all the strategies are rebalanced at the end of each month, and currencies are equally weighted within each portfolio.

4.4 Currency portfolio descriptive statistics

Table 1 provides an overview of the currency portfolio returns, calculated using mid-quotes, over the full sample period. For both the carry and momentum portfolios, average returns monotonically increase when moving from portfolio 1 to portfolio 5. For the value portfolios, returns also increase monotonically as we move from value portfolio 1 to 5, except for portfolio 3, which appears to have the lowest returns (0.75% per annum) among the value portfolios.[19] All the currency portfolios exhibit pronounced deviations from normality. Returns of all the portfolios display excess kurtosis, with value portfolio 5 having the highest sample kurtosis of 5.51. As has been extensively documented in the literature (e.g., Brunnermeier et al. (2009)), carry portfolios exhibit negative skewness, which becomes more negative, although not monotonically, as we move from portfolio 1 to portfolio 5 (-0.09 and -0.41, respectively). The low skewness of these returns is suggestive of crash risk. There are no clear patterns in skewness, however, among the other portfolios. Notably, value portfolio 5 displays large, positive skewness (0.65).

For each strategy, we also report summary statistics for the returns of a long-short portfolio (denoted as H/L) that is long in portfolio 5 and short in portfolio 1. It is worth noting that the H/L portfolios are implicitly leveraged since they involve a bet size of 2 USD. For the carry strategy, the H/L portfolio corresponds to the HML_{FX} factor of Lustig et al. (2011), which we denote as Carry hereafter. The average returns of H/L carry, momentum, and value portfolios are 7.98%, 5.07 %, and 4.70%, respectively (all are statistically significant at the 1% level). The average returns, coupled with relatively low standard deviations, translate to noteworthy annual Sharpe ratios of 0.93, 0.59, and

higher frequency data releases (see, e.g., Giannone, Reichlin, and Small (2008)), it is not clear whether such an adjustment is necessary.

[18]We consider a smaller sample of 35 currencies for the construction of value portfolios due to data availability. Similarly, Barroso and Santa-Clara (2015) and Menkhoff et al. (2016) consider only the currencies of 27 and 22 countries, respectively.

[19]Menkhoff et al. (2016) document the same pattern, thus it is likely simply an artifact of using a relatively smaller sample of currencies for the construction of value portfolios.

0.52, respectively (before transaction costs). These compare favorably to the U.S. stock market's Sharpe ratio over the same period, which was 0.56. Hence, it has been historically profitable to go long high-interest-rate currencies, recent currency winners, and undervalued currencies while shorting low-interest rate currencies, recent currency losers, and relatively overvalued currencies.

Panel B of Table 1 presents the correlation matrix of returns. The puzzling pattern that originally spawned active research on the subject is that none of the H/L portfolios' returns appear strongly correlated with the U.S. stock market, MKT (the highest correlation of 0.24 is between Carry and MKT), or other classic equity risk factors (namely, Fama and French (1993) value and Carhart (1997) momentum factors). Interestingly, currency momentum is uncorrelated with stock market momentum, suggesting that currency momentum is a separate phenomenon (this pattern is also documented by Menkhoff et al. (2012b)). In other words, if high historic Sharpe ratios are the result of exposure to risk, it is not immediately clear which risk (see Burnside (2012) for a discussion). Importantly, the three currency strategies are not strongly correlated with each other, suggesting that they capture distinct components of currency risk premium.

In the next section, we examine whether our empirical SDFs can jointly explain all the currency portfolio returns and how their performance compares to existing factors proposed in the literature.

5 Empirical SDFs

In this section, we evaluate empirical SDFs estimated under different leverage constraints. Specifically, we examine their in-sample fit, time-series properties, and out-of-sample cross-sectional asset pricing power. Importantly, we always proceed in two steps. We first examine the empirical SDFs estimated solely on the five carry portfolios (and evaluated exclusively on the cross section of five carry portfolios in asset pricing tests). We then examine the empirical SDFs estimated on the full cross section of currency portfolio returns. This approach allows us to test our new methodology on a well-understood and extensively studied cross-section of returns (i.e., the carry portfolios). Our analysis of the carry portfolios thus serves as a "proof of method", or sanity check, conducted before our main analysis – studying different currency risk premiums jointly.

5.1 In-sample fit of constrained empirical SDFs

As we explain in Section 3, introducing leverage constraints in the empirical SDF estimation introduces in-sample mispricing. To validate the notion that the maximum leverage constraints observed in the FX market are relevant for constructing benchmark SDFs, we need to ascertain that the resulting degree of in-sample mispricing is small. In this subsection, we evaluate the degree of in-sample mispricing under leverage constraints for the full sample of FX strategy returns, with a focus on the range of leverage constraints that could plausibly impact FX investors.

First, for each level of maximum leverage L, we calculate the KL distance δ_{KL}, as in equation (13), of all leverage-constrained candidate SDFs from the set of SDFs that perfectly price the assets in-sample. Second, we evaluate whether this distance is statistically significant, i.e., whether a leverage-constrained SDF candidate offers sufficiently good pricing performance. The test statistic is based on the idea that for a given leverage-constrained SDF, for its pricing performance to be satisfactory, it cannot be exceedingly distant from an SDF that perfectly prices the assets in-sample. Put differently, the correction in equation (18) should not be statistically different from zero.

Specifically, we construct the leverage-constrained candidate empirical SDFs using the full samples of five carry portfolios, and five carry, momentum, and value portfolios, respectively. In both candidate groups, we vary the leverage parameter L from zero to a level at which the constraint is not binding, which we find is $L = 45$ for carry portfolios and $L = 195$ for the full set of currency portfolios.[20] In the left panels of Figures 2a and 2b we plot δ_{KL} as a function of maximum leverage L for both the carry and the full set of portfolios, respectively. We observe that, in both cases, the distance (i) decreases quickly as L increases from zero, and (ii) approaches zero as L ceases to bind, albeit at a decreasing rate. The natural interpretation is that for $L = 0$ we obtain the set of SDFs' distance from a constant. In turn for large L, the candidate SDF's in-sample pricing errors converge to zero. However, relaxing the leverage constraint offers lower gains if L is already high.

To evaluate whether a leverage-constrained candidate SDF offers satisfactory pricing performance, we calculate the test statistic of equation (21), described in Section 3.3, of the null hypothesis that the leverage-constrained SDF lies within the set of admissible SDFs. For each L, we calculate the p-value of the test of the null hypothesis that the SDF candidate's correction is equal to zero. We plot these p-values as a function of maximum

[20]For clarity, in the plots we restrict L to the region between zero and 100. For a given cross-section of assets, the level at which the constraint ceases to bind varies with sample length and we focus on the full-length sample in order to demonstrate that even in this scenario the constraint remains relevant.

leverage (L) in the right-hand panels of Figures 2a and 2b for the five carry and the five carry, momentum, and value portfolios, respectively. In both plots, the dashed horizontal lines indicate the 10% confidence level, whereas the vertical lines indicate the tightest leverage constraint at which the constrained SDF candidate's pricing errors are indistinguishable from zero. A plotted p-value above the confidence level indicates that one cannot reject the hypothesis that a candidate SDF's pricing errors are zero, hence its in-sample pricing performance is statistically indistinguishable from an SDF that prices the assets perfectly in-sample. For carry portfolios, a leverage constraint above 26 is sufficiently slack to obtain a well-performing SDF. For the full set of portfolios, relaxing the leverage constraint to 45 (and above) allows a candidate empirical SDF to be considered satisfactory.

Our analysis implies that imposing leverage constraints, which are broadly compatible with those observed in the FX market, on empirical, data-driven SDFs results in satisfactory in-sample pricing performance. The out-of-sample performance of such factors is an open question which we address in the following section. We focus our discussion on the empirical SDF estimated under the 1:50 maximum leverage constraint ($\widehat{\mathrm{SDF}}^{L_{50}}$), in order to ease the exposition. We underline, however, that the exact choice of the leverage constraint does not have a material impact on the results as long as the constraint lies within the range mandated by the institutional framework of the FX market. Therefore, we additionally report results for empirical SDFs estimated under maximum leverage constraints ranging from 2 to 100.

5.2 Time series properties of out-of-sample empirical SDFs

We estimate our out-of-sample empirical SDFs on a rolling basis using a 180-month window. This means that our out-of-sample evaluation period starts in March 1991 and ends in October 2020. We rebalance our empirical SDF estimates monthly: for each month $t+1$, the empirical SDF is constructed using the parameters estimated on the data for the period $t-180$ to t.

Panel A of Figure 3 plots the monthly time series of the out-of-sample empirical SDF with a 1:50 leverage constraint estimated only on the five carry portfolios. The empirical SDF displays reasonable time series variation and seems to spike up around known currency crises. For example, we observe spikes around the European Exchange Rate Mechanism (ERM) crisis, the Mexican Peso crisis in 1994, the Asian financial crisis, Russia's sovereign debt crisis in the late 1990s, and the Turkish currency crisis around 2018. The time-series properties of empirical SDFs with alternative leverage constraints are similar and are thus not plotted.

Panel B of Figure 3 plots the estimated portfolio weights in the empirical SDF displayed in Panel A. We can clearly observe that the carry portfolio 1 (i.e., the lowest-interest rate portfolio) is consistently assigned the most negative weights. We also observe that carry portfolios 4 and 5 (i.e., the high-interest rate portfolios) are assigned the highest weights. These patterns are in line with the findings of Lustig et al. (2011), who show that the high-minus-low portfolio is closely related to the second principal component of carry portfolio returns, and that this principal component is key for describing the cross-section of carry portfolio average returns. It is worth noting, however, that the empirical SDF is not exclusively comprised of the two extreme portfolios, suggesting that the ad hoc high-minus-low factor formation may be disregarding important pricing information.

Panel A of Figure 4 plots the monthly time series of the out-of-sample empirical SDF with a 1:50 leverage constraint estimated on the five currency carry, momentum, and value portfolios (our benchmark specification). The empirical SDF displays similar time series variation to the empirical SDF constructed solely from the carry portfolios, but is less volatile.[21] For instance, the spike around the Lehman Brothers collapse in September 2008 appears less pronounced in the SDF constructed from the full cross-section of currency portfolios. A point worth noting is that this empirical SDF's most distinctive feature is that its prominent spikes are associated with episodes of adverse currency shocks, and not general business-cycle ones. For example, the empirical SDF is particularly elevated during the ERM and the Asian financial crises, episodes characterized by extreme currency movements. In contrast, the empirical SDF is relatively low during the 2001 U.S. recession, and even around the dot-com stock market crash.

Our out-of-sample empirical SDF series broadly resembles the global FX volatility series, which is also highly elevated around the two aforementioned currency crises and is low in the early 2000s (the Internet Appendix plots the two series side by side). Although visually it appears that the two series display similar underlying trends, the unconditional correlation of monthly FX volatility innovation and the empirical SDF is just 0.30 (the Internet Appendix presents the factor correlation matrix). Hence, a substantial fraction of the variation in our empirical SDF appears different from existing factors like FX volatility. In the next sub-section, we investigate whether these differences translate to superior asset pricing performance.

Lastly, we examine the estimated portfolio weights in the empirical SDF displayed in Panel A of Figure 4. Panel B plots the time-series of weights for all 15 currency portfolios.

[21] The average correlation of empirical SDFs (estimated on the full set of currency portfolios) with their counterparts estimated only on the carry portfolios is approximately 0.63.

Again, we observe that the carry portfolio 1 is consistently assigned the most negative weights. Momentum portfolio 1 is also assigned negative weights, but the weights decrease in absolute terms after the great financial crisis and are around zero in the last five years of our sample period (i.e., 2015-2020). Value portfolio 1, on the other hand, is almost never assigned a negative weight, with its average weight being around zero. On the long side, we observe large, positive weights assigned to carry portfolios 4 and 5, momentum portfolios 4 and 5, and value portfolio 5, which is in line with economic intuition. However, the relative magnitudes of their weights fluctuate over time. The main insight that, we believe, can be gleaned from this plot is that long-short portfolios are noisy proxies for pricing factors, a view in line with Bryzgalova, Pelger, and Zhu (2020), who argue that traditional long-short factors are often only weakly related to the true SDF.

5.3 Cross sectional asset pricing tests

In this subsection, we consider formal cross-sectional asset pricing tests. Our methodology is standard. We employ a two-pass regression to estimate the risk price of each of the candidate pricing factors in the cross-section of currency portfolios. We estimate the factor risk prices in the context of specific factor models. In particular, for each asset i, we estimate betas, β_i^j, from time-series regressions of asset i's excess returns, $R_{i,t+1}$, on J_μ factors, F_{t+1}^j, of each model μ.

$$R_{i,t+1} = a_i + \sum_{j \in J_\mu} \beta_i^j F_{t+1}^j + \varepsilon_{t+1} \tag{24}$$

We then, separately for each model μ, run a cross-sectional regression of average excess asset returns, $\bar{R}_i$, on the estimated betas, $\hat{\beta}_i^j$, to estimate the factor risk price, λ^j.

$$\bar{R}_i = \gamma^\mu + \sum_{j \in J_\mu} \lambda^j \hat{\beta}_i^j + \alpha_i^\mu, \tag{25}$$

where γ^μ is the constant and α_i^μ is the residual from a cross-sectional regression which can be interpreted as the pricing error. If model μ prices a given cross-section of assets, we should find statistically and economically significant λ^j for at least one of the factors of J_μ, and pricing errors that are, on average, zero.[22]

Statistical inference on risk prices and pricing errors is conducted using generalized method of moments (GMM) standard errors (with the Newey and West (1987) kernel and

[22]We do not impose the restriction of $\gamma^\mu = 0$ in our main analysis because our out-of-sample empirical SDFs are not traded factors.

a 12-month bandwidth), which account for cross-asset correlation in the residuals and for the betas' estimation error (see Cochrane (2009)). To evaluate model fit we consider the cross-sectional adjusted R^2 (defined as $1 - Var(\widehat{\alpha})/Var(\bar{R})$, where $Var(\widehat{\alpha})$ and $Var(\bar{R})$ are the cross-sectional variance of $\widehat{\alpha}_i$s and $\bar{R}_i$s, respectively); the mean absolute pricing error (MAPE) measured in percent per month (i.e., the mean absolute residual in the cross-sectional regression multiplied by 100); and the χ^2 test statistic, $\widehat{\alpha}' \left[\widehat{V}_{\widehat{\alpha}}\right]^{-1} \widehat{\alpha}$, of the null hypothesis that all the pricing errors, $\widehat{\alpha}$, are jointly zero ($\widehat{V}_{\widehat{\alpha}}$ is the GMM variance-covariance matrix estimate). For the cross-sectional tests to be valid, it is imperative that we have a reasonable spread in first-stage beta estimates. Therefore, for each model, we also test the null hypothesis that first-stage betas are jointly zero. We then report the difference in estimated factor betas of the portfolios with the largest and the smallest beta estimates and their corresponding statistical significance.

Our main focus is on the out-of-sample empirical SDF, $\widehat{SDF}^L$, estimated under different choices of maximum leverage level constraints ranging from 2 to 100, and the unconstrained empirical SDF, $\widehat{SDF}^U$. A priori, we expect the unconstrained empirical SDF to perform relatively worse than the constrained ones in pricing a given cross section of assets, particularly in larger cross sections, due to over-fitting of $\widehat{SDF}^U$ in relatively short estimation samples. However, the degree of the potential pricing performance differences is an open question. Hence, we evaluate both types of SDFs. It is important to note that all the empirical SDFs we consider are out-of-sample measures as their parameters are estimated using 15-year rolling windows with data only up to time t. We thus conduct our cross-sectional pricing tests exclusively on the out-of-sample evaluation period.

We compare the pricing performance of our empirical SDFs with factor models shown to be relevant for the cross section of currency returns in the extant literature. We begin by considering the models proposed by Lustig et al. (2011) and Menkhoff et al. (2012a). Both models are comprised of two factors: the dollar factor (i.e., the average return of a U.S. investor who buys all available foreign currencies forward) as the first factor, and either Carry or global FX volatility innovations as the second factor. We measure global FX volatility innovations as monthly increments in the option-implied global currency volatility index (ΔVXY).[23]

The literature also finds that currency expected returns are related to FX liquidity (e.g., Mancini et al. (2013), Karnaukh et al. (2015)), global financial market volatility (e.g., Lustig et al. (2011)), and funding liquidity (e.g., Brunnermeier et al. (2009)). We

[23] Considering the realized FX volatility measure constructed as in Menkhoff et al. (2012a), instead of the option-implied FX volatility, produces similar, but slightly weaker, results.

thus consider alternative specifications which substitute Carry (or the global FX volatility factor) with (i) the aggregate FX liquidity factor (ΔFX liq.), as in Karnaukh et al. (2015); (ii) changes in the VIX index as a proxy for global market volatility; or (iii) changes in the TED spread as a proxy for funding liquidity. It is important to note that the dollar factor serves as the level factor and adds little to the cross-sectional pricing performance. Specifically, dollar factor betas from the first stage regressions are around one for all the currency portfolios we consider (this empirical fact is well-documented in prior literature and is confirmed in our setting). Given that our main specification includes a constant in the second-stage regression, we omit dollar factor betas from the second-stage regression, but retain the dollar factor in the first-stage regressions to ensure consistency (we present the specifications with no intercept in the Internet Appendix). Lastly, we also evaluate a simple market model (using the value-weighted U.S. market excess return as the market factor, MKT) and the He et al. (2017) financial intermediary two-factor model (using MKT and the shocks to the primary dealer capital ratio, HKM, as the two factors).

5.3.1 Pricing currency carry portfolios ("proof of method")

In this subsection, we examine the pricing performance of the empirical out-of-sample SDFs (constructed solely from the five carry portfolios) on the cross-section of the five carry portfolios. Table 2 reports the risk price estimates and several measures of model fit. Columns (1) and (2) present the results for the $\widehat{\text{SDF}}^{L_{50}}_{Carry}$ and $\widehat{SDF}^{U}_{Carry}$, respectively (we add a subscript to differentiate these empirical SDFs from our benchmark empirical SDFs constructed using the full cross-section of currency portfolios). Columns (3)–(9) present the results for the models from the existing literature (namely, Carry, ΔVXY, ΔFX liq., ΔVIX, ΔTED, MKT, and HKM models).

We begin by determining whether there is enough spread in the betas to obtain reliable risk price estimates. For all the models, we reject the null hypothesis that first-stage betas are jointly zero. Moreover, the difference between the smallest and the largest beta is positive for all the models, and statistically significant for all but ΔTED. Therefore, we can draw reasonable inference from these asset pricing results.

We find that the risk price of $\widehat{\text{SDF}}^{L_{50}}_{Carry}$ is negative, -0.179, and strongly statistically significant (t-stat $= 3.34$). Moreover, the $\widehat{\text{SDF}}^{L_{50}}_{Carry}$ model delivers the tightest cross-sectional fit ($R^2 = 0.96$) and the lowest MAPE (0.042%) out of all the candidate models, which is surprising because existing models already explain the cross-section of carry portfolios well. The price of risk and model fit are relatively insensitive to the choice of maximum leverage level, particularly for levels above 45. Figure 5 plots risk price estimates and measures of model fit for empirical SDFs with different leverage levels. The price of risk estimates range between -0.057 (t-stat $= 3.29$) for $L = 10$ and -0.246 (t-stat $= 3.10$) for $L = 100$ (despite the differences in point estimates, the 95% confidence intervals of these estimates intersect with $\widehat{\mathrm{SDF}}_{Carry}^{L_{50}}$ risk price estimate's 95% confidence interval). We find that both the risk price of $\widehat{SDF}_{Carry}^{U}$ (-0.257, t-stat $= 3.0$) and its cross-sectional fit ($R^2 = 0.94$, MAPE $= 0.044\%$) are similar to those of the constrained empirical SDFs. This finding is consistent with our intuition: overfitting is expected to be less severe in small cross-sections, hence leverage constraints should play a lesser role for empirical SDFs' out-of-sample performance. However, as we see in the next sub-section, the unconstrained empirical SDF's pricing performance deteriorates substantially once we expand the cross-section of assets.

Our results for other candidate models are in line with the existing literature. The risk price of Carry, 0.006 (t-stat $= 3.50$), is similar to the original estimates of Lustig et al. (2011).[24] Consistent with the findings of Menkhoff et al. (2012a), we find that the ΔVXY risk price estimate is negative and statistically significant, and that the FX volatility model delivers a good cross-sectional fit ($R^2 = 0.91$). However, for these two models, the null that the pricing errors are jointly zero is rejected (although only marginally). Examining the specifications using ΔFX liq., ΔVIX, or ΔTED as the main pricing factor, we find that these factors' risk prices are of the expected negative sign and statistically significant, but their cross-sectional fits are inferior to the Carry and FX volatility models (particularly in the case of the ΔTED model that only delivers an R^2 of 0.57). Cross-sectional tests without a constant, but with the dollar factor, yield very similar results (the Internet Appendix presents those results). Examining the market and the financial-intermediary models, we find that, in line with the findings of He et al. (2017), the HKM model performs well: the HKM risk price estimate is positive and statistically significant (0.044, t-stat $= 2.32$), and the R^2 is relatively high (0.89). Surprisingly, we find that the MKT risk price estimate is positive and statistically significant (0.032, t-stat $= 2.91$), and that the market model's cross-sectional fit compares favorably to the currency-based factor models, contrasting the earlier findings that carry returns have little exposure to the stock market (e.g., Burnside (2012)). These results should not, however, be taken at face value as MKT is a traded factor and, thus, its risk price estimate should be close to its average return (around 0.007 per month).[25]

[24] Carry is a traded factor and its average excess return during our out-of-sample evaluation period is around 0.006 per month, which is the same is its price of risk estimate, thus the estimates make sense economically.

[25] In untabulated results, we repeat the MKT model asset pricing exercise with the additional restriction that MKT should price itself. The price of the risk point estimate is reduced to 0.008 and remains

The results of this subsection demonstrate that our out-of-sample empirical SDFs have good pricing power and offer a marginal improvement even when explaining the expected returns in a small yet widely studied cross-section. Having ascertained that the proposed method performs as expected, we proceed to our main asset pricing tests.

5.3.2 Pricing currency carry, momentum, and value portfolios

In this subsection, we examine the pricing performance of the empirical out-of-sample SDFs (constructed from the five carry, momentum, and value portfolios) on the cross-section of carry, momentum, and value portfolios. Table 3 reports the two-pass regression estimates. Similar to Table 2, columns (1) and (2) present the results for the $\widehat{\mathrm{SDF}}^{L_{50}}$ and $\widehat{SDF}^{U}$, respectively, and columns (3)–(9) present the results for the models from the existing literature. For all the models, we reject the null hypothesis that the first-stage betas are jointly zero, thus we can draw reasonable inference from the second-stage results.

We find that the risk price of $\widehat{\mathrm{SDF}}^{L_{50}}$ is negative, -0.207, and highly statistically significant (t-stat $= 4.34$). The model fit is excellent: the cross-sectional R^2 is 0.91 and MAPE is 0.054%. In contrast, the pricing performance of $\widehat{SDF}^{U}$ is much weaker. Although its price of risk is statistically significant, it appears economically unreasonable, -2.882 (i.e., the expected return on an asset with unit exposure to this factor, if such an asset exists, would be around -288% per month). Moreover, the cross-sectional fit of $\widehat{SDF}^{U}$ is inferior to $\widehat{\mathrm{SDF}}^{L_{50}}$: the R^2 is only 0.61 and MAPE is 0.116% (double $\widehat{\mathrm{SDF}}^{L_{50}}$'s MAPE). Hence, in line with our expectations, $\widehat{SDF}^{U}$ appears to suffer from in-sample overfitting, which significantly hampers its out-of-sample performance. These results highlight the importance of the introduction of leverage constraints to the exponential-tilting empirical SDF estimation problem, particularly in richer cross-sections.

Importantly, the out-of-sample asset pricing performance is robust to the leverage constraint choice, provided the maximum leverage levels are economically meaningful. Figure 6 plots risk price estimates and some measures of model fit for empirical SDFs with different leverage levels. The price of risk estimates range between -0.082 (t-stat $= 3.96$) for $L = 10$ and -0.482 (t-stat $= 4.17$) for $L = 100$ (the 95% confidence intervals of these estimates intersect with the 95% confidence of the $\widehat{\mathrm{SDF}}^{L_{50}}$ risk price estimate). It is worth noting that, for empirical SDFs with leverage constraints ranging from around 1:25 to 1:100, the pricing performance is very similar (interestingly, this range contains the most common maximum leverage limits observed in practice). However, the pricing performance dete-

statistically significant, but the R^2 decreases to 0.69, confirming our intuition that the unexpectedly good pricing performance of the market model comes at the expense of an unrealistic risk price estimate.

riorates for leverage limits of greater than 1:100, with R^2 decreasing to around 0.60 and MAPE increasing to 0.123% for leverage levels 1:250 and above (the Internet Appendix plots risk price estimates and model fit measures for empirical SDFs with leverage limits ranging from 1:1 to 1:1000, clearly illustrating the differences in performance across leverage limit levels). The pricing performance appears worse for the empirical SDFs with the strictest leverage limits. For example, for SDF with 1:2 leverage constraint, the R^2 is only 0.66 and MAPE is 0.092. This pattern is primarily driven by over-penalization and, we believe, has little economic significance as the pricing performance improves rapidly as the leverage constraints are loosened (the lowest MAPE is obtained by an empirical SDF with a 1:40 leverage constraint).[26]

Turning our attention to the factors from the existing literature, we find that none of the factor models are able to maintain their good pricing performance of the cross-section of carry portfolios when confronted with a richer cross-section of currency portfolio returns. The best-performing factor model is the Carry model: Carry factor's risk price is positive, 0.007, and statistically significant (t-stat = 3.5), its R^2, however, is lower than it was in the cross-section of carry portfolios, 0.61, and its MAPE is double $\widehat{\text{SDF}}^{L_{50}}$'s MAPE (0.116). The second best-performing model is the HKM model. The HKM factor risk price estimate is similar to its risk price estimated on the cross-section of carry portfolios, 0.047, and is marginally statistically significant (t-stat = 1.9). The R^2, however, is lower (0.45) and the MAPE is 0.131, which is around 2.4 times the $\widehat{\text{SDF}}^{L_{50}}$'s MAPE. The two worst-performing factor models are the ΔFX liq. and ΔTED models. The risk price of both ΔFX liq. and ΔTED are not statistically significant, and the two models attain a R^2 of only 0.01 and 0.05, respectively. In the case of ΔFX liq., we can reject the null hypothesis that all the pricing errors are jointly zero (we are unable to reject the null of all pricing errors being jointly zero for the ΔTED model due to the lack of variation in the first-stage betas). Surprisingly, the FX volatility model performs poorly: the ΔVXY risk price is negative, -0.002, and marginally statistically significant (t-stat = 1.68), but the R^2 is only 0.07 and we can reject the null hypothesis that all the pricing errors are jointly zero (the results are similar using a specification without a constant, but with the dollar factor, and are reported in the Internet Appendix). Also, surprisingly, we find that the equity-market factors, ΔVIX and MKT, perform relatively well. Their risk price estimates are -0.022 (t-stat = 2.7) and 0.017 (t-stat = 2.1), respectively (again, the MKT risk price is, however, too high given the average return on the market portfolio is only 0.007 per month), and the

[26] Regularization via the leverage constraint reduces the variability of the empirical SDF. However, at the very low (i.e., the strictest) levels of L, the empirical SDFs variability reduces to the extent that pricing performance is severely hindered.

R^2 are 0.21 and 0.15, respectively. These results are in line with Lilley, Maggiori, Neiman, and Schreger (2020) who find that proxies for global risk appetite like the VIX explain a larger share of currency returns after the Global Financial Crisis.

In sum, the results of this subsection amount to two significant findings: first, imposing economic constraints on data-driven pricing factors is necessary to ensure good out-of-sample pricing performance; second, the family of leverage-constrained empirical SDFs manages to explain the expected returns on the cross-section of FX portfolios better than extant factors, while the empirical SDFs' properties remain economically plausible.

5.3.3 Pricing individual currencies and FX hedge funds

In the previous subsection, we establish that our out-of-sample empirical SDF is able to price the cross-section of currency carry, momentum, and value portfolios. Although the empirical SDF estimation and asset pricing tests are conducted using different sample periods, there may still be a concern that pricing performance is mechanically driven by the fact that our empirical SDF is constructed from the same set of assets that they subsequently price. To alleviate this concern, in this subsection, we conduct asset pricing tests on individual currency returns and FX hedge fund returns.

Our approach follows that of the previous subsections. However, as we are no longer dealing with characteristic-sorted portfolios of currencies, we use rolling betas because risk loadings are expected to vary over time for individual currencies and hedge funds. For example, as the Norwegian Krone moves from being a high-interest-rate currency to being a low-interest-rate one, we expect its risk loading to change. Therefore, we use the Fama and MacBeth (1973) regression with rolling betas to estimate the risk prices. We use a rolling window of 48 months to balance stability and timeliness, but our results are similar for alternative rolling windows. We report the results solely for $\widehat{\mathrm{SDF}}^{L_{50}}$, which is constructed from the five carry, momentum, and value portfolios (the results for out-of-sample empirical SDFs with different leverage limits follow the same pattern as in the asset pricing tests discussed earlier in the paper and are thus not tabulated).

Individual currencies We begin by examining 39 individual currency returns.[27] Column 1 of Table 4 reports the results. We find that the risk price of $\widehat{\mathrm{SDF}}^{L_{50}}$ is negative, -0.098, and statistically significant (t-stat = 1.96). Although this risk price point estimate is lower than our main $\widehat{\mathrm{SDF}}^{L_{50}}$ risk price estimated on the cross-section of currency

[27] We lose five currencies in this analysis, relative to our full sample, due to their short time series (namely, Cyprus, Finland, Greece, Slovakia, and Slovenia).

portfolios (Table 3), the 95% confidence intervals of the two estimates overlap, hence one cannot establish that they are statistically significantly different. As is customary in the literature, we also separately examine a sub-sample of 15 developed countries.[28] Column 2 of Table 4 reports the results. We find that the risk price of $\widehat{\text{SDF}}^{L_{50}}$ is negative, -0.066, and marginally statistically significant (t-stat = 1.90). We note, however, that after the introduction of the euro in 1999, this sub-sample of developed countries is reduced to only nine currencies. This relatively small cross-section renders our estimates sensitive to country-specific idiosyncrasies. In sum, our empirical SDF is priced in the large cross-section of individual currencies, and even in the smaller cross-section of developed country currencies.

FX hedge funds FX hedge funds are essentially dynamically managed currency portfolios. However, as researchers, we have no knowledge of the exact strategies pursued by these hedge funds beyond that they trade currencies. This feature, we believe, renders FX hedge funds an ideal out-of-sample laboratory for evaluating currency risk factors. We posit that, irrespective of the exact strategies pursued by the FX hedge funds, to earn currency risk premiums they would need to be exposed to their source, which, based on the results presented earlier, ought to be captured by our empirical SDF. Hence, our empirical SDF should also be priced in the cross-section of FX hedge fund returns. We evaluate this conjecture using a sample of FX hedge funds.

The main data is sourced from the BarclayHedge Currency Traders Database, a database tracking hedge funds whose principal investment strategy involves directional currency trading (in either the spot or the derivative markets).[29] Due to its specialized focus, this database contains many tiny funds with extreme returns. To alleviate this feature of the data, we restrict our analysis to the cross-section of 100 largest FX hedge funds (based on the average assets under management, AUM).[30] We also winsorize hedge fund returns at the 1^{st} and 99^{th} percentiles to minimize the effects of outliers, a common practice in the literature (see, e.g., Lim, Sensoy, and Weisbach (2016)). After implementing these filters, the basic characteristics of our sample of hedge funds are comparable to samples used in the existing studies and are similar to the characteristics of our secondary hedge fund sample discussed below. For instance, the average (median) size and age of the FX hedge funds in our sample, is 467 (173) million and 6.5 (6.1) years, respectively (the Internet Appendix reports the descriptive statistics). We also use a supplementary sample of 287 hedge funds

[28] We follow the definition of Lustig et al. (2011) and consider Australia, Belgium, Canada, Denmark, euro area, France, Germany, Italy, Japan, Netherlands, New Zealand, Norway, Sweden, Switzerland, and the United Kingdom.

[29] These data are available from January 1994 to December 2019 and contain both dead and alive funds.

[30] Restricting our analysis to, for example, the 50 or 75 largest hedge funds yields similar results.

that follow a global macro investment style sourced from Eurekahedge.[31]

Column 3 of Table 4 reports our baseline one-factor specification. We find that, in the sample of FX hedge funds, the risk price of $\widehat{\mathrm{SDF}}^{L_{50}}$ is negative, -0.064, and marginally statistically significant (t-stat = 1.89). However, it is well-established in the literature that certain hedge fund characteristics are important determinants of the cross-section of hedge fund returns (see Agarwal, Mullally, and Naik (2015) for a literature survey), and it is thus imperative to control for known hedge fund characteristics in cross-sectional studies (see, e.g., Bali, Brown, and Caglayan (2011) for a typical application). Hence, we include a number of standard controls, namely fund age, AUM, management fee, incentive fee, high watermark (a dummy variable that equals one if fund i has a high watermark provision and zero otherwise), and mandated redemption notice period. Column 4 of Table 4 reports the specification with all the controls. We find that the risk price of $\widehat{\mathrm{SDF}}^{L_{50}}$ is negative, -0.112, statistically significant (t-stat = 2.60), and, interestingly, similar to the risk price estimated in the cross-section of individual currencies. In untabulated results, we find that introducing controls one by one does not alter the general pattern: the risk price estimate of our empirical SDF remains statistically significant in all specifications.

Importantly, the risk price point estimate lies within the 95% confidence interval of our main $\widehat{\mathrm{SDF}}^{L_{50}}$ risk price estimate (Table 3), implying that the risk price is consistent across a wide spectrum of FX assets. The coefficients on the controls are largely in line with the existing literature, suggesting that our sample of hedge funds is representative.[32] We also repeat the analysis using FX market factors from the existing literature. We find that, in a specification with hedge fund controls, none of the alternative currency-risk factors display a statistically significant price of risk in the cross-section of FX hedge funds (the Internet Appendix reports these results).

Lastly, to mitigate potential concerns that our results are database specific, we repeat our analysis on a sample of 287 global macro hedge funds. Although global macro funds are not explicitly FX hedge funds, they are typically considered exposed to currency risks (e.g., Della Corte et al. (2016)). Column 5 of Table 4 reports the results. We find that, in the cross-section of global macro hedge funds, the risk price estimate is similar to our other

[31] To filter these data, we follow Dahlquist, Sokolovski, and Sverdrup (2021), who use the same database. Specifically, we consider one share class per fund, and retain in our sample only the funds with average AUMs above 20 million USD, that report at least 48 months of returns, and have complete characteristics information. We then keep only the hedge funds that report global macro as their main investment style, leaving us with 287 unique hedge funds for the sample period running from January 2001 to October 2020.

[32] For example, the coefficient on a fund's AUM is negative and statistically significant in line with Aggarwal and Jorion (2010), while the coefficient on the redemption notice period is positive and statistically significant, which is in line with Aragon (2007).

estimates (-0.071), is statistically significant (t-stat = 1.97), and that its 95% confidence interval overlaps with that of our benchmark risk price estimate, corroborating our results using the FX funds.

Taken together, these findings speak to the robustness of the pricing performance of our empirical SDF and suggest that it does indeed capture important drivers of currency risk premiums.

6 Investment performance

A portfolio that perfectly inversely tracks the true SDF would attain the maximal possible Sharpe ratio. However, our out-of-sample empirical SDFs are potentially just approximations of the true SDF, hence no such clear predictions can be made. Nevertheless, it is reasonable to expect that a portfolio mimicking the inverse of an out-of-sample factor with superior pricing performance would display favorable investment properties, as such a factor is likely capturing the underlying risk premiums. Whether this is the case is an empirical question that we examine in this section.

For a given a leverage constraint, our empirical SDF corresponds to the solution of a portfolio optimization problem (see Section 3.1 for details). In this section, we investigate the investment performance of a normalized version this optimal portfolio, i.e., normalized $\widehat{\mathrm{SDF}}^{L_{50}}$ in our benchmark case (SDF portfolio, hereafter). We note that normalization is required to facilitate a comparison with alternative strategies. We compare the SDF portfolio's out-of-sample performance with the investment performance of long-short currency carry, momentum, and value strategies, Barclay Currency Traders Index (a proxy for the average FX hedge fund performance), and a buy-and-hold investment in the value-weighted portfolios of U.S. public firms (U.S. stock market portfolio). Our evaluation sample period runs from March 1991 to October 2020.

We define the SDF portfolio as the following normalization of the portfolio that solves the constrained minimum SDF distance problem in Definition 1. The unnormalized portfolio takes the form

$$R_{t+1}^{C} = \theta_0^{t,C} R_f + \sum_{j=1}^{N} \theta_j^{t,C} (R_{jt+1} - R_f)\,, \tag{26}$$

with the sum the risky assets' weights satisfying the constraint $\sum_{j=1}^{N} \left|\theta_j^{t,C}\right| \leq 50$ in each period. In turn, the investments in basic FX strategies, like Carry, are comprised of long-short positions in the high- and low-interest-rate portfolios corresponding to $\theta_h R_{hi} - \theta_{lo} R_{lo}$,

$\theta_{hi} = \theta_{lo} = 1$, $|\theta_{hi}| + |\theta_{lo}| = 2$. Therefore, the normalized weights of the SDF portfolio are constructed in two steps, first by taking

$$\widetilde{\theta}_j^{t,N} \equiv \frac{\theta_j^{t,C}}{\sum_{k=1}^{N} \left|\theta_k^{t,C}\right|},$$

which produces a portfolio with risky investments limited to the leverage of 1:2. Second, an additional normalization ensures that neither the short nor the long legs of the portfolio require investing more than \$1. This is achieved by calculating the sums S_+^θ and S_-^θ of the absolute values of negative and positive portfolio weights, and then further normalizing

$$\theta_j^{t,N} \equiv \frac{\widetilde{\theta}_j^{t,N}}{\max\left\{S_+^\theta, S_-^\theta\right\}}.$$

To ensure all capital is productive, the remainder, $2 - \sum_{j=1}^{N} \left|\theta_j^{t,N}\right|$, is invested at the risk-free rate.

Table 5 presents the results. We report average returns, their standard deviations, the annualized Sharpe ratio, skewness, kurtosis, the maximum drawdown, and the correlation with the U.S. stock market. Examining first the performance without transaction costs (Panel A), we find that the SDF portfolio achieves a Sharpe ratio of 1.01, which is higher than the Sharpe ratios of currency carry, momentum, and value strategies (0.87, 0.62, and 0.73, respectively). Importantly, the SDF portfolio's returns are positively skewed (0.53), contrasting the characteristic negative skewness of the carry strategy. Momentum and value strategy returns, however, are also positively skewed (0.24 and 1.00, respectively). Given that the SDF portfolio encompasses all three basic currency strategies, our finding that its returns are positively skewed is consistent with Berge et al. (2010), Jordà and Taylor (2012), and Barroso and Santa-Clara (2015) who document, using different methodologies, that augmenting currency carry trades with signals on economic fundamentals and FX momentum enhances investment performance and eliminates negative return skewness.[33] The SDF portfolio also displays relatively low correlation (0.19) with the U.S. stock market, thus potentially offering good diversification benefits for investors.

Next, we turn to the performance adjusted for transaction costs (Panel B). For the currency strategies, we compute excess returns, as in equation (23), using bid and ask quotes from Datastream for the relevant forward and spot rates. We compute returns

[33] In a similar spirit, Kroencke, Schindler, and Schrimpf (2014) show that adding FX strategies to an international stock or bond portfolio increases its Sharpe ratio without adversely impacting the skewness of portfolio returns.

using 50% of the quoted bid-ask spreads for our benchmark analysis, as the bid-ask spreads provided by Datastream are known to be unrealistically high (e.g., Lyons (2001)).[34] The relative performance, however, is not affected but these choices regarding transaction costs. Taking transaction costs into account, the SDF portfolio achieves a Sharpe ratio of 0.77, which is higher than that of carry, momentum, and value strategies, and is also higher than the Sharpe ratios of Barclay Currency Traders Index (0.22) and the U.S. stock market (0.56). Positive skewness of SDF portfolio's returns is also preserved after transaction costs. Moreover, The SDF portfolio offers the greatest improvement relative to the long-short currency strategies, average FX hedge funds, and the U.S. stock market when it comes to periods of prolonged losses. The SDF portfolio produces the lowest maximum drawdowns of only -11.82%, compared to -28.52, -39.93, and -21.57 for the carry, momentum, and value strategies, respectively, and a staggering -54.36% for the U.S. stock market.

To gauge the performance across time, in Figure 7 we graphically present the cumulative investment performance of the strategies. To facilitate the comparison of relative performance across different strategies, we first rescale the return series so that all the strategies attain the same annual volatility as the U.S. stock market portfolio (14.53%). The most salient observations are the following: (i) the periods of largest currency strategy drawdowns typically do not coincide with recessions, (ii) the volatility of the SDF portfolio does not vary substantially over time (similarly to the value and momentum strategies, and in contrast to the carry strategy), (iii) the SDF portfolio benefits from a diversification effect, and unlike the carry strategy, does not, for example, crash in the end of the 1990s, (iv) the volatility-adjusted performance of the SDF portfolio relative to the U.S. stock market is exceptionally good between 2000 and 2007, however it is weaker after the Great Financial Crisis (which could be, in part, due to the widespread compression of short-term interest rates that has taken place since 2008).

In sum, the SDF portfolio performs very well out-of-sample: it offers a stable quantity of risk; attractive return properties, such as positive skewness; and its performance does not seem to deteriorate when stock markets are in distress.

[34] We also avoid double counting transaction costs when calculating SDF portfolio returns by netting positions at the individual currency level. For example, if the SDF portfolio is long one momentum portfolio 5, short one value portfolio 4, and long carry portfolio 3, and CAD is found in all three of these portfolios, the SDF portfolio would net only a single position in CAD and transaction costs would only be incurred once.

7 Robustness

In this section, we conduct a number of robustness tests for our main result – the cross-sectional pricing of carry, momentum, and value portfolios. We focus our discussion on $\widehat{\text{SDF}}^{L_{50}}$ (the robustness results for the out-of-sample empirical SDFs with different leverage limits follow the same pattern as in the main empirical tests and are, thus, not tabulated). Table 6 presents the results. For ease of exposition, column 1 reports the results for our benchmark specification $\widehat{\text{SDF}}^{L_{50}}$ (from Table 3) that is estimated using a 15-year rolling window and is rebalanced monthly.

First, we consider different rebalancing frequencies. In particular, instead of updating portfolio weights every month, we keep the same portfolio weights for 3, 6, 12, or 24 months before updating. Columns 2−5 present these results. We see no noticeable differences in the risk price estimates. Relative to our benchmark specification, we observe a very slight reduction in the cross-sectional R^2s, which range between 0.87 to 0.89, and a slight increase in the MAPEs, which range between 0.058% and 0.060%. These results suggest that the choice of rebalancing frequency has little effect on the asset pricing performance of our empirical SDF.

Second, we consider different rolling windows. We begin by examining an out-of-sample empirical SDF estimated on 20-year rolling window, and rebalanced monthly. Column 6 presents the results. The results are similar to our benchmark specification. We observe a slight change in the risk price estimate (−0.171), and a slight improvement in both the cross-sectional R^2 (0.93) and MAPE (0.043). However, it is important to note that a longer rolling window shortens the out-of-sample evaluation period, hence these slight differences are likely just driven by the different sample periods. Next, we consider an expanding estimation window. Specifically, we estimate our empirical SDF recursively, with an initial estimation period of 15 years, and rebalance monthly. Column 7 presents the results. We observe only a slight change in the risk price estimate (−0.190), suggesting that our risk price estimate is consistent. We do, however, observe a deterioration in the cross-sectional R^2 (0.81) and MAPE (0.077), suggesting that putting more weight on recent time periods when estimating empirical SDFs is beneficial for pricing performance. We do stress, however, that the differences are slight: the risk price is similar to the benchmark specification and the pricing performance of the recursively-estimated empirical SDF is still substantially superior to the existing factors, and to the unconstrained empirical SDF. In sum, our analysis shows that the choice of estimation window has little impact on our results.

Third, we consider an alternative specification for the constraint in Definition 1. Specifically, we replace the leverage constraint, $\sum_{j=1}^{N} |\theta_j| \leq L$, by the following constraint

$$\sum_{j=1}^{N} \theta_j^2 \leq R\,, \tag{27}$$

which is closely related to the concept of ridge regression. In contrast to our baseline leverage constraint (which, as previously mentioned, is closely related to the concept of LASSO), the ridge constraint, when binding, does not set some coefficients to zero. Instead, it impacts all the coefficients in a more uniform manner by shrinking them toward zero. The most important difference between the leverage and ridge constraints lies in their economic interpretations. As a quadratic function of portfolio allocations, the ridge constraint favors allocations of similar magnitude across assets. It can thus be interpreted as a concentration constraint.[35]

We estimate a set of out-of-sample empirical SDFs for different levels of R, ranging from one to 40000, and repeat our benchmark asset pricing tests using these empirical SDFs. The Internet Appendix plots risk prices estimates and fit statistics for the different levels of the square root of R (ranging from one to 100), similar to Figure 6. We observe a clear pattern, where for a certain range of R the risk price estimates are most economically reasonable, and the cross-sectional fit is the best. Importantly, this optimal range coincides with maximum leverage levels of 1:2 to 1:100 (by examining the sum of the squares of the portfolio loadings of our leverage-constrained SDF, it is possible to loosely map the leverage constrain to the concentration constraint). In column 8 of Table 6 we present the empirical out-of-sample SDF with $R = 484$, as it is, on average, approximately the sum of the squares of the portfolio loadings of our benchmark 1:50 leveraged empirical SDF.[36] The estimated price of risk of the concentration-constrained SDF reported in column 8 of Table 6 is almost identical to our other estimates, and its pricing performance is also similar, offering a slight improvement in R^2 and MAPE over the baseline empirical SDF, but yielding a slightly lower p-value of the asset pricing test statistic.

These results demonstrate the robustness of our procedure, as we converge on similar empirical SDF estimates even when starting from a different optimization problem. However, these results also highlight the importance of having a clear economic interpretation to guide our choice of nuisance parameters (such as R) so as to better extract meaningful

[35]In the discussion that follows, we refer to constraint in equation 27 as a concentration constraint.

[36]With $R = 484$, square root of R is equal to 22, the total leverage taken by the investor is then around 1:44.

estimates. For instance, in our analysis above, we needed to rely on the leverage limits to isolate a set of meaningful estimates from the procedure using concentration constraints. Although position limits are an economically intuitive concept, we have little knowledge of the actual limits used in the practice; position limits are not ubiquitous, vary substantially, and are often not implemented in a consistent manner (see, e.g., Acworth (2009)). We thus believe our benchmark formulation is the most relevant one.

8 Concluding remarks

The foreign exchange market is the largest and most liquid market in the world. Yet, existing models are unable to fully explain the cross-section of currency strategies' returns.

We show that it is possible to construct out-of-sample pricing factors from FX portfolio returns that jointly explain the cross-section of carry, momentum, and value portfolio returns. Our pricing factors are also priced in the cross-section of individual currencies and, notably, in a cross-section of FX hedge funds, which constitutes a true out-of-sample test. This thus suggests that our factors capture essential currency risk premiums. To do so, we develop a new methodology that combines an existing minimum-dispersion SDF framework with economically-motivated parameter restrictions (that can be interpreted as leverage constraints), which renders the resulting factors robust in out-of-sample pricing and investment performance. Our results are largely insensitive to the choice of leverage constraints, provided that the leverage limits are broadly consistent with institutional settings. This highlights the importance for empirical methods to have a clear economic interpretation and demonstrates the utility of our approach.

Our study sets out several points of reference for future empirical international asset pricing models. First, when constructing empirical SDFs from large cross-sections, institutionally mandated constraints on investor leverage can serve as an economically-motivated guideline for the degree of regularization, thus enabling a parsimonious (single-factor) representation of the SDF with good out-of-sample pricing performance. Second, the peaks of our empirical SDFs coincide with currency-specific crises, and not adverse business cycle conditions, suggesting the existence of a FX-specific risk. Third, long-short portfolios typically used in the literature do not effectively capture the pricing information present in the cross-section of currency returns.

Our empirical SDFs, however, are data-driven constructs and, hence, the question remains open as to which economic forces are behind them. Nevertheless, our empirical SDFs could serve as parsimonious and robust benchmarks for international asset pricing models.

References

Acworth, Will, 2009, CFTC examines position limits, *The Magazine of the Futures Industry* 19, 24–6.

Agarwal, Vikas, Kevin Mullally, and Narayan Y. Naik, 2015, The economics and finance of hedge funds: A survey of the academic literature, *Foundations and Trends in Finance, Forthcoming* .

Aggarwal, Rajesh K., and Philippe Jorion, 2010, The performance of emerging hedge funds and managers, *Journal of Financial Economics* 96, 238–256.

Akram, Q. Farooq, Dagfinn Rime, and Lucio Sarno, 2008, Arbitrage in the foreign exchange market: Turning on the microscope, *Journal of International Economics* 76, 237–253.

Almeida, Caio, and René Garcia, 2012, Assessing misspecified asset pricing models with empirical likelihood estimators, *Journal of Econometrics* 170, 519–537.

Almeida, Caio, and René Garcia, 2017, Economic implications of nonlinear pricing kernels, *Management Science* 63, 3361–3380.

Ang, Andrew, Sergiy Gorovyy, and Gregory B. Van Inwegen, 2011, Hedge fund leverage, *Journal of Financial Economics* 102, 102–126.

Aragon, George O., 2007, Share restrictions and asset pricing: Evidence from the hedge fund industry, *Journal of Financial Economics* 83, 33–58.

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

Bali, Turan G., Stephen J. Brown, and Mustafa Onur Caglayan, 2011, Do hedge funds' exposures to risk factors predict their future returns?, *Journal of Financial Economics* 101, 36–68.

Barroso, Pedro, and Pedro Santa-Clara, 2015, Beyond the carry trade: Optimal currency portfolios, *Journal of Financial and Quantitative Analysis* 50, 1037–1056.

Barth, Daniel, Laurel Hammond, and Phillip Monin, 2020, Leverage and risk in hedge funds, Working paper, Office of Financial Research.

Bekaert, Geert, and George Panayotov, 2020, Good carry, bad carry, *Journal of Financial and Quantitative Analysis* 55, 1063–1094.

Berge, Travis J., Òscar Jordà, and Alan M. Taylor, 2010, Currency carry trades, Working Paper 16491, National Bureau of Economic Research.

Borwein, J. M., and A. S. Lewis, 1991, Duality relationships for entropy-like minimization problems, *SIAM Journal on Control and Optimization* 29, 325–338.

Boyarchenko, Nina, Thomas M. Eisenbach, Pooja Gupta, Or Shachar, and Peter Van Tassel, 2020, Bank-intermediated arbitrage, Working paper, Federal Reserve Bank of New York.

Branger, Nicole, 2004, Pricing derivative securities using cross-entropy: an economic analysis, *International Journal of Theoretical and Applied Finance* 7, 63–81.

Brunnermeier, Markus K., Stefan Nagel, and Lasse H. Pedersen, 2009, Carry trades and currency crashes, *NBER Macroeconomics Annual* 23, 313–347.

Bryzgalova, Svetlana, Markus Pelger, and Jason Zhu, 2020, Forest through the trees: Building cross-sections of stock returns, Working paper, London Business School.

Burnside, Craig, 2012, Carry trades and risk, in Jessica James, Ian Marsh, and Lucio Sarno, eds., *Handbook of Exchange Rates*, chapter 10, 283–312 (John Wiley & Sons, New Jersey).

Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo, 2011a, Do peso problems explain the returns to the carry trade?, *Review of Financial Studies* 24, 853–891.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo, 2007, The returns to currency speculation in emerging markets, *American Economic Review* 97, 333–338.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo, 2011b, Carry trade and momentum in currency markets, *Annual Review Financial Economics* 3, 511–535.

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

Chen, Luyang, Markus Pelger, and Jason Zhu, 2020, Deep learning in asset pricing, Working paper, Stanford University.

Chernov, Mikhail, Magnus Dahlquist, and Lars A. Lochstoer, 2020, Pricing currency risks, Working Paper 28260, National Bureau of Economic Research.

Chernov, Mikhail, Jeremy Graveline, and Irina Zviadadze, 2018, Crash risk in currency returns, *Journal of Financial and Quantitative Analysis* 53, 137–170.

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

Dahlquist, Magnus, and Henrik Hasseltoft, 2020, Economic momentum and currency returns, *Journal of Financial Economics* 136, 152–167.

Dahlquist, Magnus, Valeri Sokolovski, and Erik Sverdrup, 2021, Hedge funds and financial intermediaries, Working Paper 19-8, Swedish House of Finance.

Daniel, Kent, Robert J. Hodrick, and Zhongjin Lu, 2017, The carry trade: Risks and drawdowns, *Critical Finance Review* 6, 211 – 262.

Della Corte, Pasquale, Roman Kozhan, and Anthony Neuberger, 2021, The cross-section of currency volatility premia, *Journal of Financial Economics* 139, 950–970.

Della Corte, Pasquale, Tarun Ramadorai, and Lucio Sarno, 2016, Volatility risk premia and exchange rate predictability, *Journal of Financial Economics* 120, 21–40.

DeMiguel, Victor, Lorenzo Garlappi, Francisco J. Nogales, and Raman Uppal, 2009, A Generalized Approach to Portfolio Optimization: Improving Performance by Constraining Portfolio Norms, *Management Science* 55, 798–812.

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

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

Farhi, Emmanuel, Samuel Fraiberger, Xavier Gabaix, Romain Ranciere, and Adrien Verdelhan, 2015, Crash risk in currency markets, Working Paper 20948, Harvard University.

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

Gabaix, Xavier, and Matteo Maggiori, 2015, International liquidity and exchange rate dynamics, *Quarterly Journal of Economics* 130, 1369–1420.

Gârleanu, Nicolae, and Lasse Heje Pedersen, 2013, Dynamic trading with predictable returns and transaction costs, *The Journal of Finance* 68, 2309–2340.

Ghojogh, Benyamin, and Mark Crowley, 2019, The theory behind overfitting, cross validation, regularization, bagging, and boosting: Tutorial.

Ghosh, Anisha, Christian Julliard, and Alex P. Taylor, 2019, An information-theoretic asset pricing model, Working paper, London School of Economics.

Giannone, Domenico, Lucrezia Reichlin, and David Small, 2008, Nowcasting: The real-time informational content of macroeconomic data, *Journal of Monetary Economics* 55, 665–676.

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

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

Hansen, L. P., and R. Jagannathan, 1991, Implications of security market data for models of dynamic economies, *Journal of Political Economy* 99, 225–262.

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

Hassan, Tarek Alexander, and Tony Zhang, 2020, The economics of currency risk, Working Paper w27847, National Bureau of Economic Research.

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

Heimer, Rawley, and Alp Simsek, 2019, Should retail investors' leverage be limited?, *Journal of Financial Economics* 132, 1–21.

Itskhoki, Oleg, and Dmitry Mukhin, 2021, Exchange rate disconnect in general equilibrium, *Journal of Political Economy* 129, 2183–2232.

Jagannathan, Ravi, and Tongshu Ma, 2003, Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps, *Journal of Finance* 58, 1651–1683.

Jeanneret, Alexandre, and Valeri Sokolovski, 2021, Commodity prices and currencies, Working paper, HEC Montreal.

Jordà, Òscar, and Alan M. Taylor, 2012, The carry trade and fundamentals: Nothing to fear but feer itself, *Journal of International Economics* 88, 74–90.

Jurek, Jakub W., 2014, Crash-neutral currency carry trades, *Journal of Financial Economics* 113, 325–347.

Jylhä, Petri, 2018, Margin requirements and the security market line, *Journal of Finance* 73, 1281–1321.

Karnaukh, Nina, Angelo Ranaldo, and Paul Söderlind, 2015, Understanding FX liquidity, *Review of Financial Studies* 28, 3073–3108.

Korsaye, Sofonias, Alberto Quaini, and Fabio Trojani, 2019, Smart SDFs, Working paper, Université de Génève.

Korsaye, Sofonias, Fabio Trojani, and Andrea Vedolin, 2020, The global factor structure of exchange rates, Working paper, National Bureau of Economic Research.

Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh, 2020, Shrinking the cross-section, *Journal of Financial Economics* 135, 271–292.

Kroencke, Tim A., Felix Schindler, and Andreas Schrimpf, 2014, International diversification benefits with foreign exchange investment styles, *Review of Finance* 18, 1847–1883.

Lettau, Martin, Matteo Maggiori, and Michael Weber, 2014, Conditional risk premia in currency markets and other asset classes, *Journal of Financial Economics* 114, 197–225.

Lilley, Andrew, Matteo Maggiori, Brent Neiman, and Jesse Schreger, 2020, Exchange rate reconnect, *Review of Economics and Statistics* 1–28.

Lim, Jongha, Berk A. Sensoy, and Michael S. Weisbach, 2016, Indirect incentives of hedge fund managers, *Journal of Finance* 71, 871–918.

Lu, Zhongjin, and Zhongling Qin, 2021, Leveraged funds and the shadow cost of leverage constraints, *Journal of Finance* 76, 1295–1338.

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan, 2011, Common risk factors in currency markets, *Review of Financial Studies* 24, 3731–3777.

Lustig, Hanno, and Adrien Verdelhan, 2007, The cross section of foreign currency risk premia and consumption growth risk, *American Economic Review* 97, 89–117.

Lyons, Richard K., 2001, *The microstructure approach to exchange rates*, volume 333 (MIT press, Cambridge, MA).

Mancini, Loriano, Angelo Ranaldo, and Jan Wrampelmeyer, 2013, Liquidity in the foreign exchange market: Measurement, commonality, and risk premiums, *Journal of Finance* 68, 1805–1841.

Martin, Ian, and Stefan Nagel, 2019, Market efficiency in the age of big data, Working paper, National Bureau of Economic Research.

Maurer, Thomas A, Thuy-Duong Tô, and Ngoc-Khanh Tran, 2019, Pricing risks across currency denominations, *Management Science* 65, 5308–5336.

Maurer, Thomas Andreas, Thuy Duong To, and Ngoc-Khanh Tran, 2020, Market timing and predictability in FX markets, Working paper, University of Hong Kong.

Meese, Richard A., and Kenneth Rogoff, 1983, Empirical exchange rate models of the seventies: Do they fit out of sample?, *Journal of International Economics* 14, 3–24.

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2012a, Carry trades and global foreign exchange volatility, *Journal of Finance* 67, 681–718.

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2012b, Currency momentum strategies, *Journal of Financial Economics* 106, 660–684.

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2016, Currency value, *Review of Financial Studies* 30, 416–441.

Nagel, Stefan, 2021, *Machine Learning in Asset Pricing*, volume 1 (Princeton University Press).

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–08.

Okunev, John, and Derek White, 2003, Do momentum-based strategies still work in foreign currency markets?, *Journal of Financial and Quantitative Analysis* 38, 425–447.

Opie, Wei, and Steven J. Riddiough, 2020, Global currency hedging with common risk factors, *Journal of Financial Economics* 136, 780–805.

Pojarliev, Momtchil, and Richard M. Levich, 2008, Do professional currency managers beat the benchmark?, *Financial Analysts Journal* 64, 18–32.

Sandulescu, Mirela, Fabio Trojani, and Andrea Vedolin, 2021, Model-free international stochastic discount factors, *Journal of Finance* 76, 935–976.

Stutzer, Michael, 1996, A simple nonparametric approach to derivative security valuation, *Journal of Finance* 51 (5), 1633–1652.

Table 1: Descriptive statistics

Panel A of this table presents summary statistics (mean, standard deviation, skewness, and kurtosis) of returns of the carry, momentum and value currency portfolios. All returns are monthly excess returns in USD. Carry, momentum and value portfolios are constructed by sorting currencies on interest rates, past one-month returns, and five-year real exchange rate changes, respectively (the real exchange rates are defined such that a higher real exchange rate indicates a depreciated foreign currency). For each strategy, portfolios are ranked from low to high characteristic such as interest rates; portfolio 1 contains the currencies with the lowest characteristic, and portfolio 5 contains the currencies with the highest. The average return and the standard deviation are annualized and reported in %. Newey and West (1987) standard errors are reported in parentheses. The Sharpe ratio is the ratio of the annualized mean and standard deviation. H/L denotes a long-short portfolio that is long in portfolio 5 and short in portfolio 1. Panel B presents the correlation matrix. Carry, Mom, and Val denote the H/L carry, momentum, and value long-short portfolios, respectively. MKT, HML, and WML denote excess return on the U.S. stock market, Fama and French (1993) value factor (return on a long-short portfolio of high book-to-market stocks and low book-to-market stock), and Carhart (1997) momentum factor (return on a long-short stock portfolio of the winners and losers of the past year), respectively. The sample period is from 1976-02 to 2020-10.

Panel A: Summary statistics

	P1	P2	P3	P4	P5	H/L
Carry portfolios						
Mean	−1.939	1.061	1.219	3.461	6.035	7.975
	(1.459)	(1.576)	(1.671)	(1.475)	(2.210)	(1.676)
Standard deviation	0.085	0.088	0.088	0.087	0.106	0.086
Sharpe ratio	−0.228	0.120	0.139	0.396	0.572	0.925
Skewness	−0.090	−0.158	−0.263	−0.100	−0.406	−0.420
Kurtosis	4.481	4.222	4.471	4.618	4.660	4.862
Momentum portfolios						
Mean	−0.817	1.347	1.692	2.907	4.248	5.065
	(1.530)	(1.687)	(1.601)	(1.531)	(1.906)	(1.368)
Standard deviation	0.098	0.090	0.088	0.089	0.090	0.086
Sharpe ratio	−0.083	0.149	0.191	0.326	0.473	0.591
Skewness	−0.250	−0.405	−0.103	0.002	0.079	0.217
Kurtosis	4.858	4.961	3.958	4.760	4.162	5.444
Value portfolios						
Mean	1.249	1.479	0.746	2.732	5.946	4.698
	(1.719)	(1.694)	(1.781)	(1.740)	(2.234)	(1.805)
Standard deviation	0.093	0.095	0.099	0.097	0.102	0.091
Sharpe ratio	0.135	0.155	0.075	0.281	0.585	0.518
Skewness	−0.304	−0.240	−0.119	−0.196	0.648	1.029
Kurtosis	4.374	4.725	3.729	4.010	5.512	7.460

Panel B: Correlation matrix

	Carry	Mom	Val	MKT	HML	WML
Carry	1.000					
Mom	−0.042	1.000				
Val	0.267	−0.010	1.000			
MKT	0.236	−0.112	0.077	1.000		
HML	0.033	0.002	−0.041	−0.230	1.000	
WML	−0.120	0.051	−0.079	−0.125	−0.232	−0.232

Table 2: Cross-sectional asset pricing results: carry trade portfolios

The table presents factor premiums for factor models estimated by a two-pass regression. Factor risk exposures (betas) are estimated in a first-stage time-series regression. Risk prices are the mean slopes of period-by-period cross-sectional regressions of portfolio excess returns on betas. The test assets are excess returns to five carry portfolios. The factors are the out-of-sample empirical SDF with a 1:50 leverage constraint ($\widehat{\text{SDF}}^{L_{50}}_{Carry}$), the unconstrained out-of-sample empirical SDF ($\widehat{SDF}^{U}_{Carry}$), the difference between the excess return on the high-interest rate currency portfolio and the low-interest-rate portfolio (Carry) as in Lustig et al. (2011), changes to the global option-implied FX volatility (ΔVXY), the aggregate FX liquidity factor (ΔFX liq.) as in Karnaukh et al. (2015), excess return on the U.S. stock market (MKT), changes to the CBOE Volatility Index (ΔVIX), changes to the TED spread (ΔTED), and the shocks to the primary dealer capital ratio (HKM) as in He et al. (2017). The empirical SDFs are constructed from the five carry portfolios. Standard errors, in parentheses, are computed using GMM with the Newey and West (1987) kernel and a twelve-month bandwidth. R^2 refers to the adjusted R-squared from the cross-sectional regression of average excess returns on factor betas, and MAPE is mean absolute pricing error of monthly returns (in %). The $\chi^2(\alpha = 0)$ test statistic, $\alpha' [V_\alpha]^{-1} \alpha$, tests the null hypothesis that all the pricing errors, α, are jointly zero (V_α is based on the GMM variance-covariance matrix as described above). The $\chi^2(\beta = 0)$ test statistic, $\beta' [V_\beta]^{-1} \beta$, tests the null hypothesis that all the betas, β, estimated in the first-stage time-series regressions are jointly zero (V_β is based on the asymptotic variance-covariance matrix). $\beta max - \beta min$ refers to the difference in estimated factor betas of the portfolios with the largest and the smallest beta estimates (Newey and West (1987) standard errors are in parentheses). In column (9) the difference in betas test refers to the HKM factor. The full sample period is from 1976-02 to 2020-10. The empirical SDFs are estimated using a rolling estimation window of 180 months. The cross-sectional tests are conducted on the out-of-sample evaluation period which runs from 1991-03 to 2020-10. *** indicates significance at 1 percent, ** – at 5 percent, * – at 10 percent.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$\widehat{\text{SDF}}^{L_{50}}_{Carry}$	$-0.179^{\star\star\star}$								
	(0.054)								
$\widehat{SDF}^{U}_{Carry}$		$-0.257^{\star\star\star}$							
		(0.086)							
Carry			$0.006^{\star\star\star}$						
			(0.002)						
ΔVXY				$-0.007^{\star\star\star}$					
				(0.002)					
ΔFX liq.					$-0.004^{\star\star}$				
					(0.002)				
ΔVIX						$-0.043^{\star\star\star}$			
						(0.014)			
ΔTED							$-0.004^{\star}$		
							(0.002)		
MKT								$0.032^{\star\star\star}$	$0.026^{\star\star\star}$
								(0.011)	(0.008)
HKM									$0.044^{\star\star}$
									(0.019)
Intercept	−0.002	−0.002	0.002	0.002	0.002	0.002	0.002	$-0.005^{\star\star}$	$-0.004^{\star\star}$
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)
R^2	0.96	0.94	0.91	0.91	0.89	0.85	0.57	0.89	0.92
MAPE (%)	0.042	0.046	0.055	0.062	0.063	0.082	0.134	0.057	0.055
$p(\chi^2(\alpha = 0))$	0.74	0.70	0.09	0.09	0.37	0.12	0.12	0.37	0.07
$p(\chi^2(\beta = 0))$	0.00	0.00	0.00	0.00	0.00	0.00	0.05	0.00	0.00
$\beta_{\max} - \beta_{\min}$	$0.036^{\star\star\star}$	$0.025^{\star\star\star}$	$1.000^{\star\star\star}$	$1.055^{\star\star\star}$	$2.111^{\star\star\star}$	$0.187^{\star\star\star}$	1.291	$0.195^{\star\star\star}$	$0.129^{\star\star\star}$
	(0.01)	(0.00)	(0.07)	(0.23)	(0.78)	(0.05)	(0.88)	(0.06)	(0.04)

Table 3: Cross-sectional asset pricing results carry, momentum and value

This table present factor premiums for factor models estimated by a two-pass regression. Factor risk exposures (betas) are estimated in a first-stage time-series regression. Risk prices are the mean slopes of period-by-period cross-sectional regressions of portfolio excess returns on betas. The test assets are excess returns to five carry, momentum, and value portfolios. The factors are the out-of-sample empirical SDF with a 1:50 leverage constraint ($\widehat{\text{SDF}}^{L_{50}}$), the unconstrained out-of-sample empirical SDF ($\widehat{\text{SDF}}^{U}$), the difference between the excess return on the high-interest rate currency portfolio and the low-interest-rate portfolio (Carry) as in Lustig et al. (2011), changes to the global option-implied FX volatility (ΔVXY), the aggregate FX liquidity factor (ΔFX liq.) as in Karnaukh et al. (2015), excess return on the U.S. stock market (MKT), changes to the CBOE Volatility Index (ΔVIX), changes to the TED spread (ΔTED), and the shocks to the primary dealer capital ratio (HKM) as in He et al. (2017). The empirical SDFs are constructed from the five carry, momentum, and value portfolios. Standard errors, in parentheses, are computed using GMM with the Newey and West (1987) kernel and a twelve-month bandwidth. R^2 refers to the adjusted R-squared from the cross-sectional regression of average excess returns on factor betas, and MAPE is mean absolute pricing error of monthly returns (in %). The $\chi^2(\alpha = 0)$ test statistic, $\alpha' \left[V_\alpha\right]^{-1} \alpha$, tests the null hypothesis that all the pricing errors, α, are jointly zero (V_α is based on the GMM variance-covariance matrix as described above). The $\chi^2(\beta = 0)$ test statistic, $\beta' \left[V_\beta\right]^{-1} \beta$, tests the null hypothesis that all the betas, β, estimated in the first-stage time-series regressions are jointly zero (V_β is based on the asymptotic variance-covariance matrix). $\beta max - \beta min$ refers to the difference in estimated factor betas of the portfolios with the largest and the smallest beta estimates (Newey and West (1987) standard errors are in parentheses). In column (9) the difference in betas test refers to the HKM factor. The full sample period is from 1976-02 to 2020-10. The empirical SDFs are estimated using a rolling estimation window of 180 months. The cross-sectional tests are conducted on the out-of-sample evaluation period which runs from 1991-03 to 2020-10. *** indicates significance at 1 percent, ** – at 5 percent, * – at 10 percent.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$\widehat{\text{SDF}}^{L_{50}}$	−0.207***								
	(0.048)								
$\widehat{\text{SDF}}^{U}$		−2.882***							
		(0.916)							
Carry			0.007***						
			(0.002)						
ΔVXY				−0.002*					
				(0.001)					
ΔFX liq.					−0.001				
					(0.001)				
ΔVIX						−0.022***			
						(0.008)			
ΔTED							−0.001		
							(0.001)		
MKT								0.017**	0.002
								(0.008)	(0.012)
HKM									0.047*
									(0.025)
Intercept	−0.002	−0.001	0.002	0.002	0.002	0.002	0.002	−0.002	0.001
	(0.002)	(0.003)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.003)
R^2	0.91	0.61	0.55	0.07	0.01	0.21	0.05	0.15	0.45
MAPE (%)	0.054	0.116	0.107	0.154	0.160	0.142	0.157	0.145	0.131
$p(\chi^2(\alpha = 0))$	0.60	0.31	0.34	0.00	0.00	0.23	0.22	0.29	0.61
$p(\chi^2(\beta = 0))$	0.00	0.00	0.00	0.00	0.00	0.00	0.09	0.00	0.00
$\beta_{\max} - \beta_{\min}$	0.042***	0.002	1.000***	1.055***	2.111***	0.186***	1.538*	0.195***	0.129***
	(0.01)	(0.00)	(0.07)	(0.23)	(0.78)	(0.05)	(0.93)	(0.06)	(0.04)

Table 4: Cross-sectional asset pricing results: individual currencies and hedge funds

This table presents factor premiums for factor models estimated by Fama and MacBeth (1973) regression. Factor risk exposures (betas) are estimated using 48-month rolling first-stage time-series regressions. Risk prices are the mean slopes of period-by-period cross-sectional regressions of portfolio excess returns on betas. The test assets are as follows. Column 1, 39 individual currencies (All Ccy); column 2, 15 developed country currencies (Dev. Ccy); columns 3 and 4, the 100 largest FX hedge funds in BarclayHedge Currency Traders database (FX HFs); column 5, 287 hedge funds following global macro investment style in Eurekahedge database (GM HFs). The factor is the out-of-sample empirical SDF with a 1:50 leverage constraint ($\widehat{\text{SDF}}^{L_{50}}$). The empirical SDF is constructed from the five carry, momentum, and value portfolios. Regression specifications in columns 4 and 5 include the following hedge fund characteristics: logarithm of each hedge fund's AUM (in USD million), fund age (scaled as tenth of a decade for legibility), management fee (in percent), performance fee (in percent), a dummy indicating if the fund has a high watermark (HWM), and the redemption notice period (in years). Newey and West (1987) standard errors with 12 lags are reported in parenthesis. R^2_{FMB} refers to the time-series average of the adjusted R-squareds of period-by-period cross-sectional regressions. The full sample period is from 1976-02 to 2020-10. The empirical SDF is estimated using a rolling estimation window of 180 months. The cross-sectional tests are conducted on the out-of-sample evaluation period which runs from 1991-03 to 2020-10 for currencies, from 1994-01 to 2017-12 for the FX hedge funds, and from 2001-01 to 2020-10 for the global macro hedge funds. *** indicates significance at 1 percent, ** – at 5 percent, * – at 10 percent.

	(1)	(2)	(3)	(4)	(5)
$\widehat{\text{SDF}}^{L_{50}}$	$-0.098^{\star\star}$	$-0.066^{\star}$	$-0.064^{\star}$	$-0.112^{\star\star\star}$	$-0.071^{\star\star}$
	(0.050)	(0.035)	(0.034)	(0.043)	(0.035)
Age				−0.001	$-0.002^{\star}$
				(0.002)	(0.001)
AUM/10				$-0.007^{\star\star\star}$	$-0.005^{\star}$
				(0.003)	(0.003)
Mgmt. fee				$0.002^{\star\star}$	0.000
				(0.001)	(0.001)
Inct. fee				−0.000	0.000
				(0.000)	(0.000)
HWM				−0.000	0.000
				(0.001)	(0.002)
Redemption				$0.059^{\star\star\star}$	$0.007^{\star}$
				(0.018)	(0.004)
Intercept	−0.001	−0.001	$0.005^{\star\star\star}$	$0.007^{\star\star}$	$0.005^{\star\star\star}$
	(0.001)	(0.001)	(0.001)	(0.003)	(0.001)
R^2_{FMB}	0.17	0.23	0.14	0.43	0.20
N_{Assets}	39	15	100	100	287
Test assets	All Ccy	Dev. Ccy	FX HFs	FX HFs	GM HFs

Table 5: Investment performance

This table presents the summary statistics (mean, standard deviation, Sharpe ratio, skewness, kurtosis, and maximum drawdown, defined as the maximum observed percentage loss from a peak to a trough of a portfolio, before a new peak is attained, and correlation with the U.S. stock market, ρ_{MKT}) for the excess returns on FX factor strategies. For comparison, summary statistics for the Barclay FX hedge fund index and the U.S. stock market are also presented. The average return and the standard deviation are annualized and are expressed in %. The Sharpe ratio is the ratio of the annualized mean and standard deviation. Carry, Momentum and Value are the excess returns of the high-minus-low portfolios for the carry, momentum and value sorted portfolios. The SDF Portfolio excess return is constructed by rescaling the out-of-sample optimal Lagrange multipliers obtained in the constrained SDF problem (1) so that (a) the absolute values of the weights sum to two (i.e., the strategy goes long and short at most $1 worth of foreign currency similarly to e.g., Carry), and (b) the strategy never requires borrowing at the risk-free rate (the sum of the negative weights is capped at minus one and is always greater in absolute value than the sum of positive weights, the remainder is invested in U.S. T-Bills). The after-transaction-costs returns on the FX strategies are calculated using bid and ask quotes from Datastream for the relevant forward and spot rates, with an adjustment to the bid-ask spread that reduces it by 50%. The full sample period is from 1976-02 to 2020-10. The weights of the SDF Portfolio are estimated using a rolling window of 180 months, and they are taken one period ahead out-of-sample. The performance statistics are calculated on the out-of-sample evaluation period which runs from 1991-03 to 2020-10.

Panel A: Returns without Transaction Costs							
Strategy	Mean	SD	Sharpe ratio	Skewness	Kurtosis	Max. draw.	ρ_{MKT}
SDF Portfolio	5.51	5.45	1.01	0.53	5.25	−8.87	0.19
Carry	8.21	9.41	0.87	−0.39	4.54	−26.12	0.31
Momentum	5.47	8.78	0.62	0.24	5.60	−33.09	−0.19
Value	7.35	10.09	0.73	1.00	6.69	−17.66	0.07

Panel B: Returns with Transaction Costs							
Strategy	Mean	SD	Sharpe ratio	Skewness	Kurtosis	Max. draw.	ρ_{MKT}
SDF Portfolio	4.13	5.37	0.77	0.45	5.15	−11.82	0.19
Carry	6.10	9.33	0.65	−0.46	4.56	−28.52	0.31
Momentum	3.47	8.72	0.40	0.17	5.79	−39.93	−0.19
Value	5.28	9.95	0.53	0.92	6.46	−21.57	0.07
Barclays FX Index	1.62	6.99	0.23	1.67	12.32	−19.83	−0.02
US Stock Market	8.39	14.92	0.56	−0.66	4.42	−54.36	

Table 6: Robustness: cross-sectional asset pricing results

Columns 1 to 8 of this table present factor premiums for one-factor models estimated by a two-pass regression. Factor risk exposures (betas) are estimated in a first-stage time-series regression. Risk prices are the mean slopes of period-by-period cross-sectional regressions of portfolio excess returns on betas. The test assets are excess returns to five carry, momentum, and value portfolios. In columns 1 to 7, the factors are the out-of-sample empirical SDFs ($\widehat{\text{SDF}}$) with 1:50 leverage constraint constructed from the five carry, momentum, and value portfolios. The factors are estimated using different rolling estimation windows (15 or 20 years). In column 6 the factor is estimated using an expanding (Exp.) window with initial size of 180 months. The estimated weights are used out-of-sample (OSS) to construct empirical SDF for horizons of one, three, six, 12, or 24 months before being updated. In column 8, the factor is the out-of-sample empirical SDF constructed from the five carry, momentum, and value portfolios estimated under a constraint on position concentration. To ease comparison, the concentration constraint parameter, R, is chosen such that the average value of leverage for the concentration-constrained portfolio is approximately 1:50. Standard errors, in parentheses, are computed using GMM with the Newey and West (1987) kernel and a twelve-month bandwidth. R^2 refers to the adjusted R-squared from the cross-sectional regression of average excess returns on factor betas, and MAPE is mean absolute pricing error of monthly returns (in %). The $\chi^2(\alpha = 0)$ test statistic, $\alpha' \left[V_\alpha\right]^{-1} \alpha$, tests the null hypothesis that all the pricing errors, α, are jointly zero (V_α is based on the GMM variance-covariance matrix as described above). The full sample period is from 1976-02 to 2020-10. The cross-sectional tests are conducted on the out-of-sample evaluation period which runs from 1991-03 to 2020-10. *** indicates significance at 1 percent, ** – at 5 percent, * – at 10 percent.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\widehat{SDF}$	−0.207***	−0.214***	−0.215***	−0.206***	−0.211***	−0.171***	−0.190***	−0.214***
	(0.048)	(0.050)	(0.051)	(0.048)	(0.055)	(0.041)	(0.046)	(0.043)
Intercept	−0.002	−0.002	−0.003	−0.003	−0.003	−0.002	−0.003*	−0.003
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.002)	(0.002)
R^2	0.91	0.89	0.88	0.87	0.89	0.93	0.81	0.93
MAPE (%)	0.054	0.058	0.058	0.060	0.059	0.043	0.077	0.044
$p(\chi^2(\alpha = 0))$	0.60	0.58	0.57	0.61	0.61	0.77	0.76	0.59
OOS horizon	1m	3m	6m	12m	24m	1m	1m	1m
Est. window	15y	15y	15y	15y	15y	20y	Exp.	15y
Constr. type	Lev.	Lev.	Lev.	Lev.	Lev.	Lev.	Lev.	Con.

Figure 1: The out-of-sample empirical SDF pricing errors

The figure shows pricing errors using a one-factor out-of-sample empirical SDF model on the cross-section of monthly currency excess returns of the carry, momentum and value portfolios. The pricing factor is the out-of-sample empirical SDF with a 1:50 leverage constraint. Carry, momentum and value portfolios are constructed by sorting currencies on interest rates, past one-month returns, and five-year real exchange rate changes, respectively (the real exchange rates are defined such that a higher real exchange rate indicates a depreciated foreign currency). For each strategy, portfolios are ranked from low to high characteristic such as interest rates; portfolio 1 contains the currencies with the lowest characteristic, and portfolio 5 contains the currencies with the highest. The full sample period is from 1976-02 to 2020-10. The empirical SDF is constructed from the five carry, momentum, and value portfolios. The empirical SDF is estimated using a rolling window of 180 months, and the out-of-sample step size is one month. The cross-sectional test is conducted on the out-of-sample evaluation period which runs from 1991-03 to 2020-10. The diagonal dashed line denotes the 45-degree line.

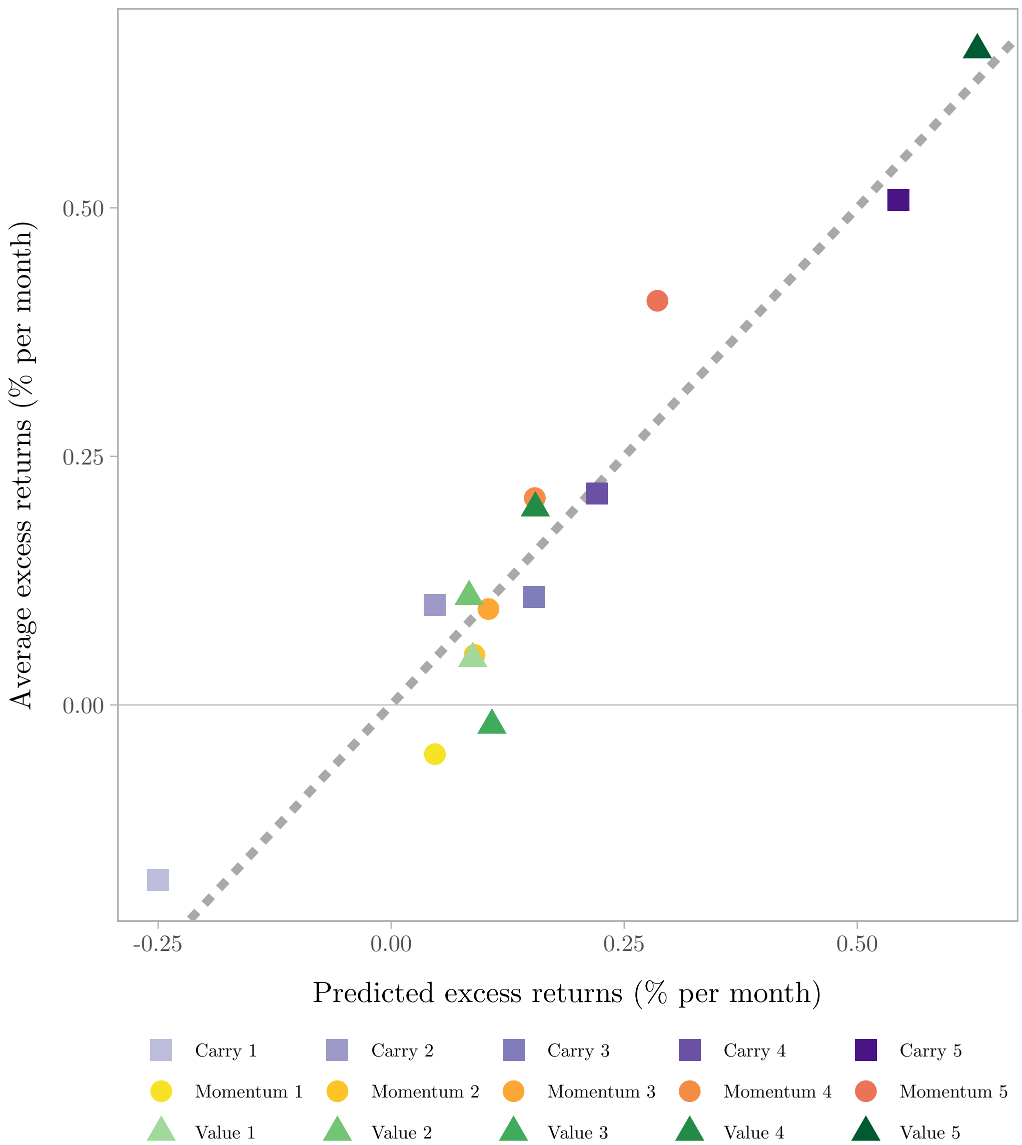

Figure 2: In-sample pricing performance of leverage-constrained SDFs

The figures show the Kullback-Leibler (KL) distance of the leverage-constrained SDFs in the complete sample (1976-01 to 2020-10) and the corresponding p-value of the test of the null hypothesis that the distance is zero, and hence the leverage-constrained SDF is indistinguishable from the SDF that prices the assets in-sample. The test statistic is given in equation (21) in Section 3.3. In Panel (a) the leverage-constrained SDFs are constructed from the five carry portfolios. In Panel (b) the leverage-constrained SDFs are constructed from the five carry, momentum, and value portfolios. Carry, momentum and value portfolios are constructed by sorting currencies on interest rates, past one-month returns, and five-year real exchange rate changes, respectively (the real exchange rates are defined such that a higher real exchange rate indicates a depreciated foreign currency). Panel (a) plots the measures for the empirical SDFs constructed from only the five carry portfolios, and pricing the cross-section of the five carry portfolios. Panel (b) plots the measures for the empirical SDFs constructed from five carry, momentum, and value portfolios, and pricing the full cross-section of the 15 currency portfolios. The data are monthly and the sample period runs from 1976-2 to 2020-10.

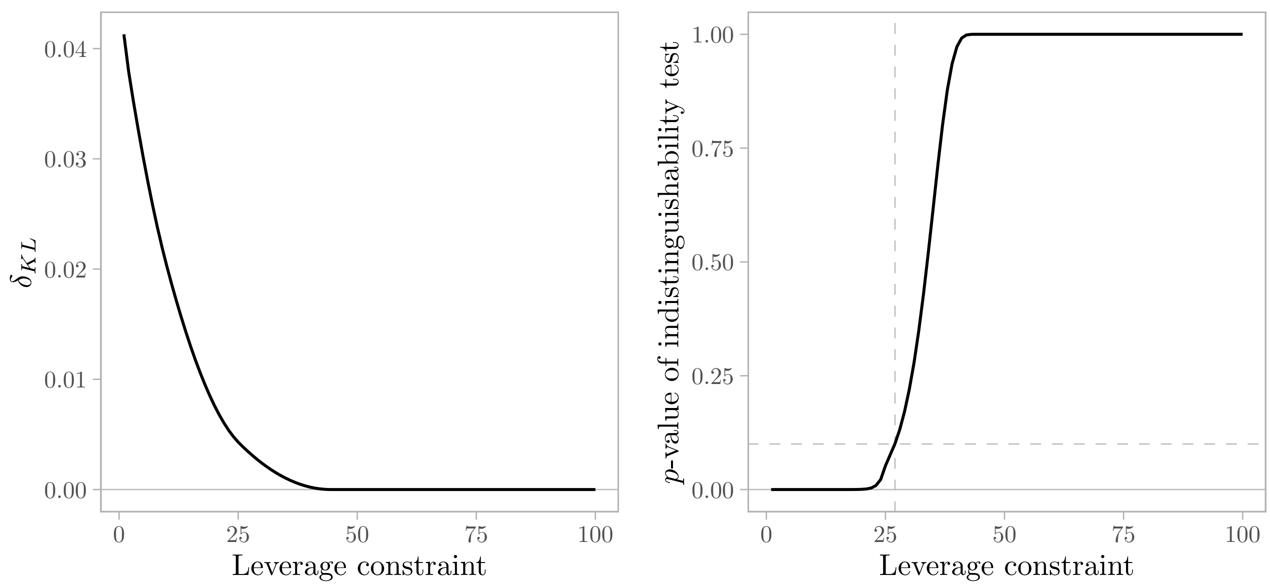

(a) KL distances for SDFs based on only the carry portfolios

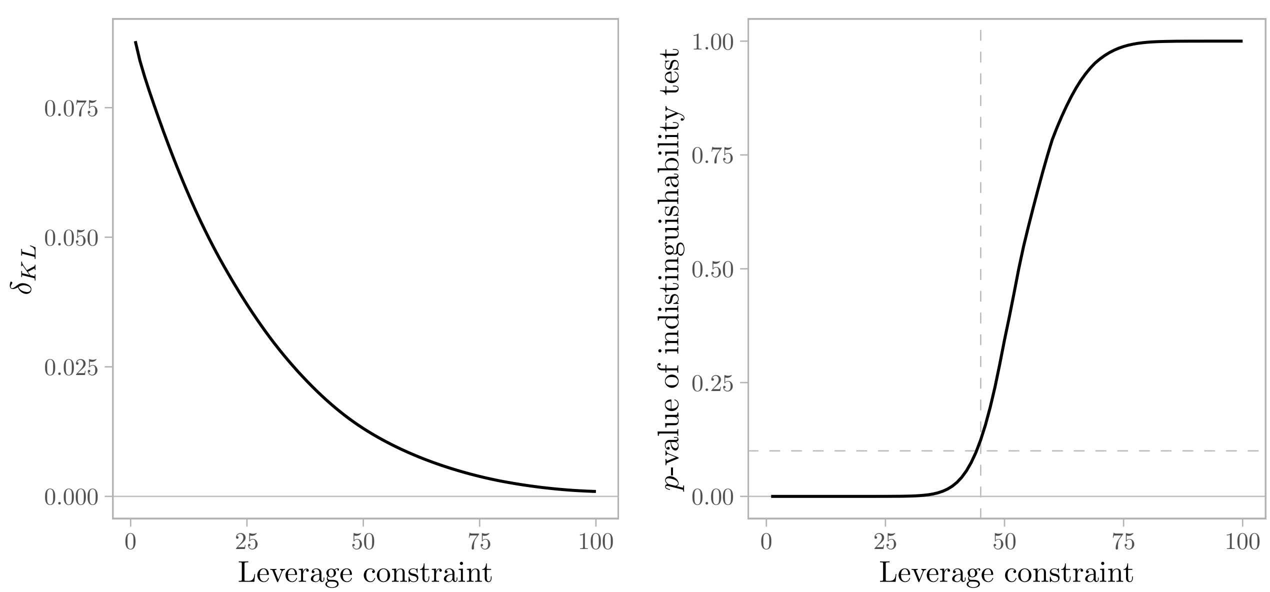

(b) KL distance for SDFs based on the carry, momentum, and value portfolios

Figure 3: Out-of-sample empirical SDF estimated on carry portfolios

The top panel of the figure shows the monthly time series of the out-of-sample empirical SDF with a 1:50 leverage constraint ($\widehat{\text{SDF}}^{L_{50}}_{Carry}$), estimated on five currency carry portfolios. The bottom panel shows the SDF portfolio weights. The empirical SDF is estimated using a rolling window of 180 months, and the out-of-sample step size is one month. Carry portfolios are constructed by sorting currencies on interest rates. The data are monthly and the full sample period runs from 1976-2 to 2020-10, while the out-of-sample period starts in 1991-03. NBER recessions and FX market crises periods are shaded in grey and blue, respectively.

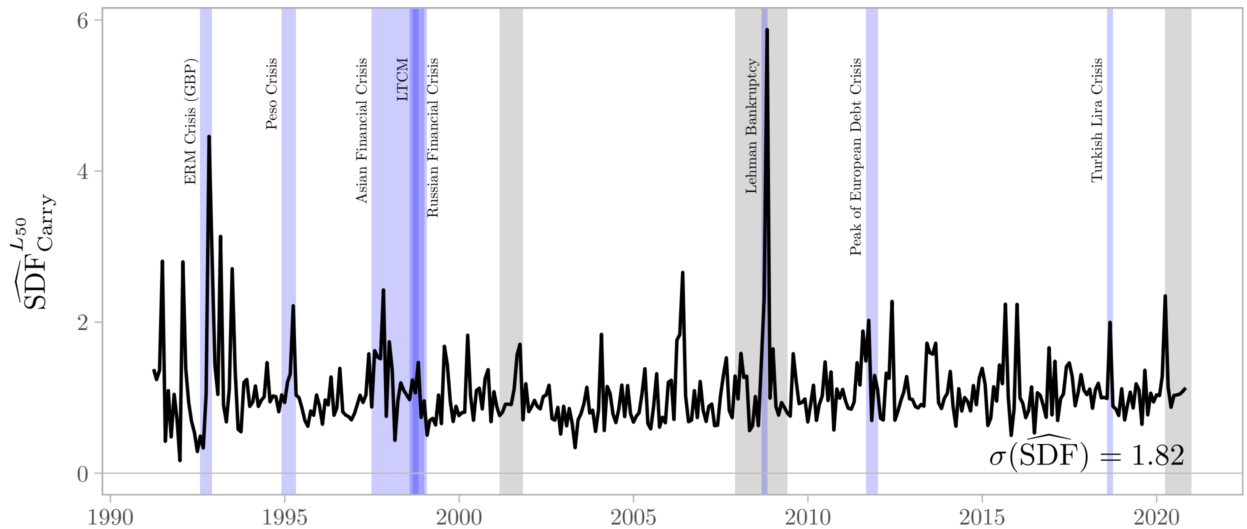

(a) Empirical out-of-sample SDF

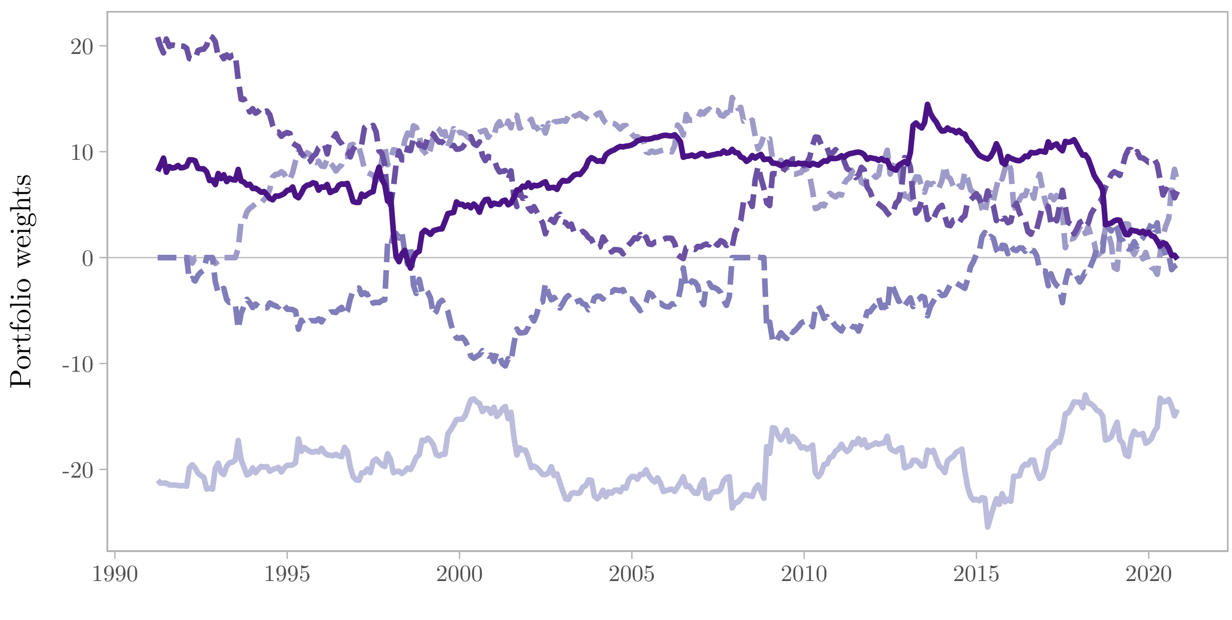

(b) Unscaled SDF portfolio weights

Figure 4: Out-of-sample SDF estimated on carry, momentum, and value portfolios

The top panel of the figure shows the monthly time series of the out-of-sample empirical SDF with a 1:50 leverage constraint ($\widehat{\text{SDF}}^{L_{50}}$), estimated on five currency carry, momentum, and value portfolios. The bottom panel shows the SDF portfolio weights. The empirical SDF is estimated using a rolling window of 180 months, and the out-of-sample step size is one month. Carry, momentum and value portfolios are constructed by sorting currencies on interest rates, past one-month returns, and five-year real exchange rate changes, respectively (the real exchange rates are defined such that a higher real exchange rate indicates a depreciated foreign currency). The data are monthly and the full sample period runs from 1976-2 to 2020-10, while the out-of-sample evaluation period runs from 1991-03 to 2020-10. NBER recessions and FX market crises periods are shaded in grey and blue, respectively.

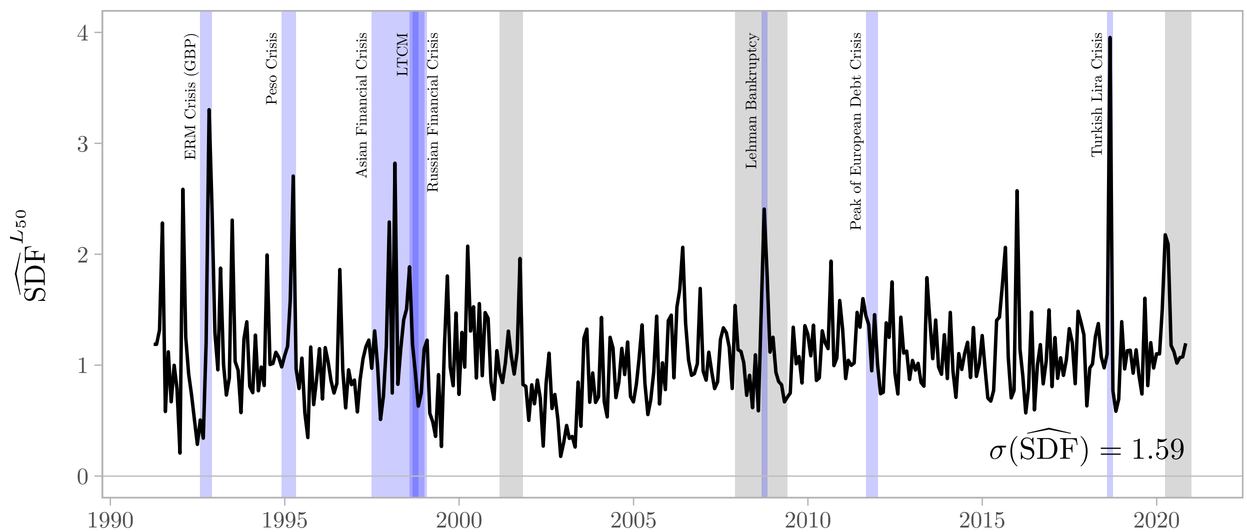

(a) Empirical out-of-sample SDF

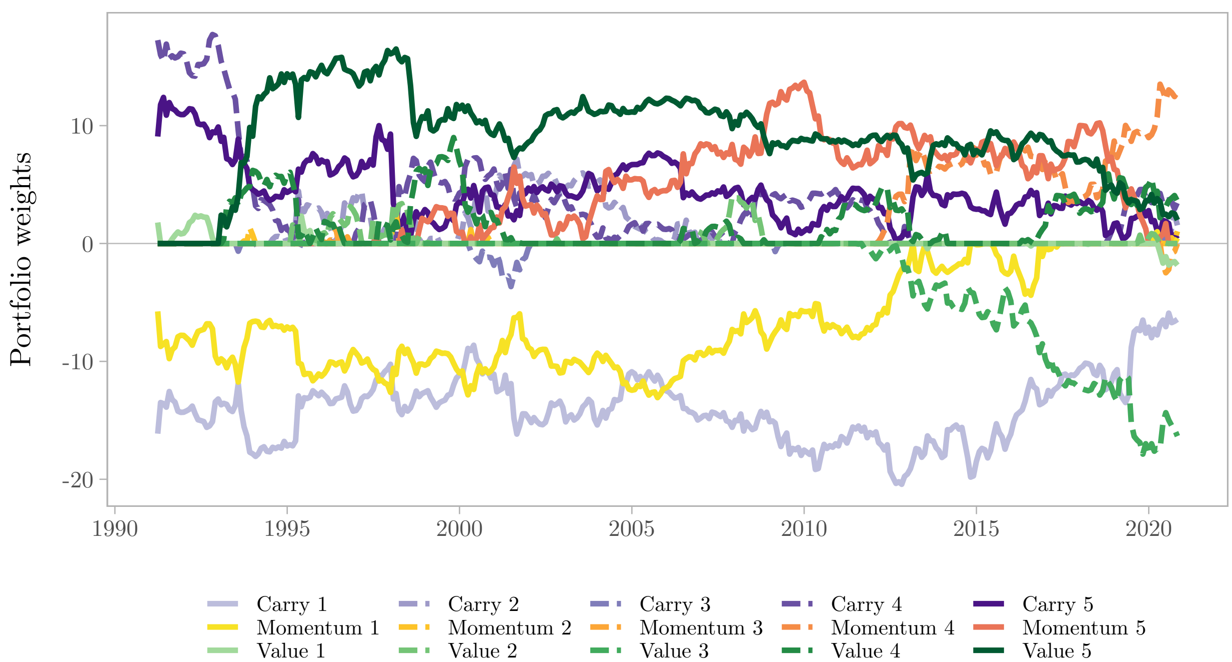

(b) Unscaled SDF portfolio weights

Figure 5: Cross-sectional asset pricing results for carry trade portfolios across leverage constraints

The figure shows risk price estimates and several measures of model fit for out-of-sample empirical SDFs estimated with different levels of maximum leverage constraint. The empirical SDFs are constructed from the five carry portfolios. The test assets are five carry portfolios' excess returns. Factor risk exposures (betas) are estimated in a first-stage time-series regression. Risk prices, shown in panel (a), are the mean slopes of period-by-period cross-sectional regressions of portfolio excess returns on betas. The 95% confidence interval is displayed in red along each risk price estimate. Confidence interval are constructed using standard errors that are computed using GMM with the Newey and West (1987) kernel and a 12-month bandwidth. Panel (b) shows MAPEs, mean absolute pricing error of monthly returns (in %). Panel (c) shows adjusted R^2s from the cross-sectional regression of average excess returns on factor beta. Panel (d) shows the p-values of the $\chi^2(\alpha = 0)$ test statistic, $\alpha' \left[V_\alpha\right]^{-1} \alpha$, which tests the null hypothesis that all the pricing errors, α, are jointly zero (V_α is based on the GMM variance-covariance matrix as described above). The full sample period is from 1976-02 to 2020-10. The empirical SDFs are estimated using a rolling estimation window of 180 months. The cross-sectional tests are conducted on the out-of-sample evaluation period which runs from 1991-03 to 2020-10.

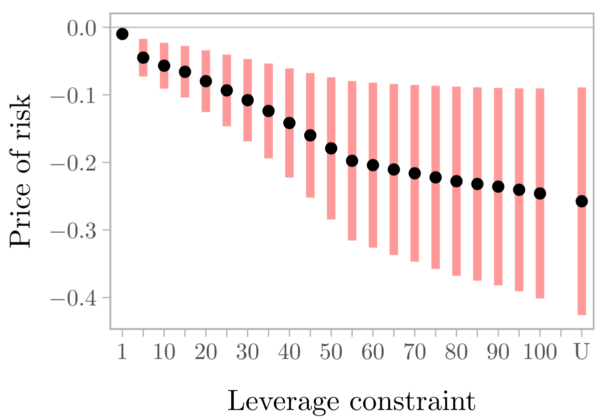

(a) Estimated price of risk

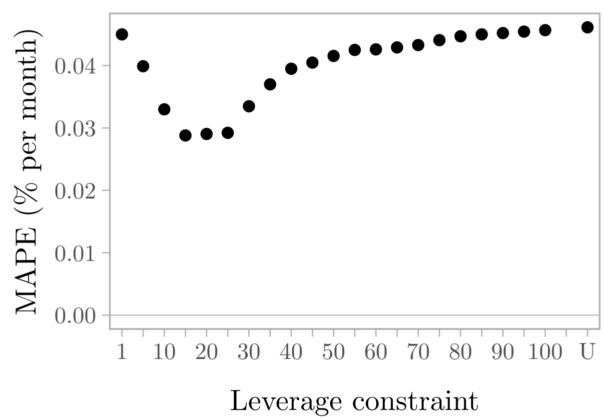

(b) Mean absolute pricing error

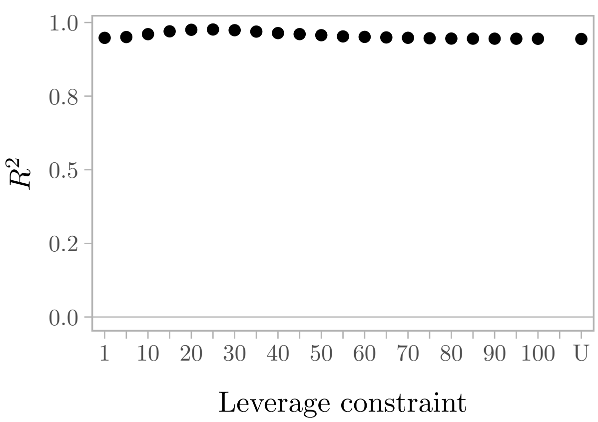

(c) Adjusted R^2

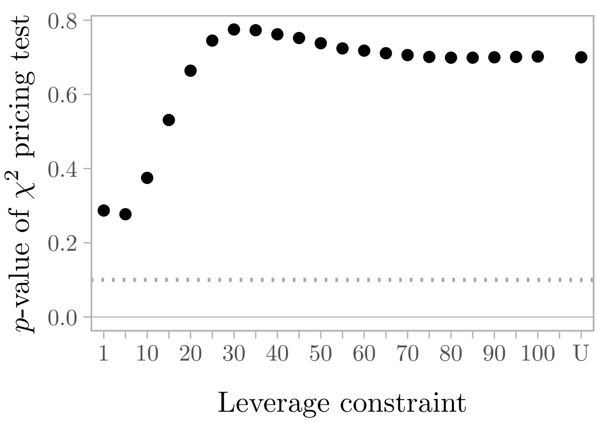

(d) Asset pricing test p-values

Figure 6: Cross-sectional asset pricing results carry, momentum and value

The figure shows risk price estimates and several measures of model fit for out-of-sample empirical SDFs estimated with different levels of maximum leverage constraint. The empirical SDFs are constructed from the five carry, momentum and value portfolios. The test assets are excess returns on five carry, momentum and value portfolios. Factor risk exposures (betas) are estimated in a first-stage time-series regression. Risk prices, shown in panel (a), are the mean slopes of period-by-period cross-sectional regressions of portfolio excess returns on betas. The 95% confidence interval is displayed in red along each risk price estimate. Confidence interval are constructed using standard errors that are computed using GMM with the Newey and West (1987) kernel and a 12-month bandwidth. Panel (b) shows MAPEs, mean absolute pricing error of monthly returns (in %). Panel (c) shows adjusted R^2s from the cross-sectional regression of average excess returns on factor beta. Panel (d) shows the p-values of the $\chi^2(\alpha = 0)$ test statistic, $\alpha' \left[V_\alpha\right]^{-1} \alpha$, which tests the null hypothesis that all the pricing errors, α, are jointly zero (V_α is based on the GMM variance-covariance matrix as described above). The full sample period is from 1976-02 to 2020-10. The empirical SDFs are estimated using a rolling estimation window of 180 months. The cross-sectional tests are conducted on the out-of-sample evaluation period which runs from 1991-03 to 2020-10.

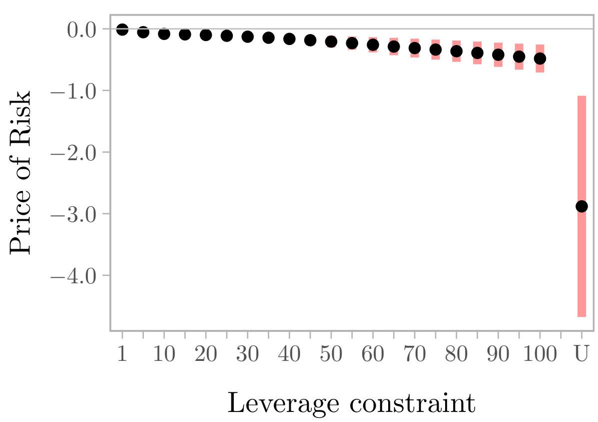

(a) Estimated price of risk

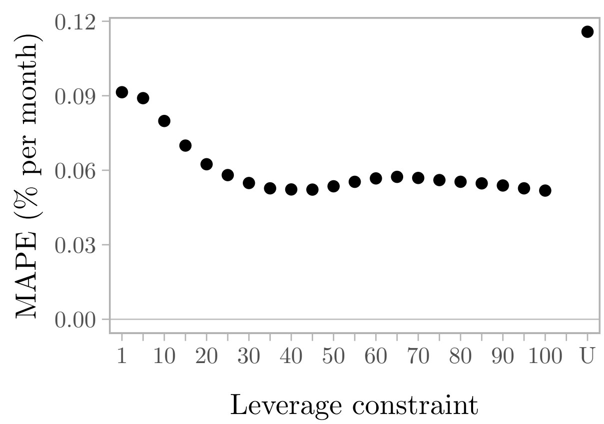

(b) Mean absolute pricing error

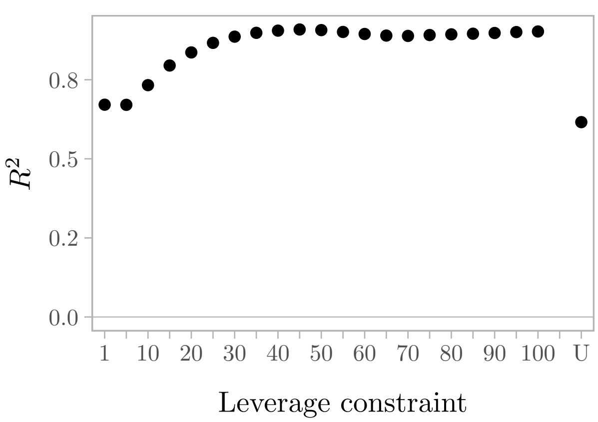

(c) Adjusted R^2

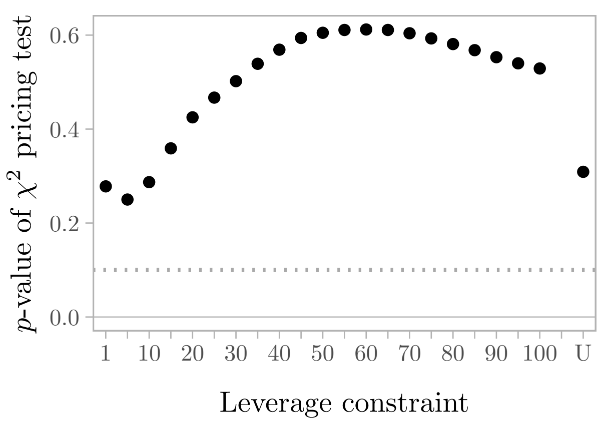

(d) Asset pricing test p-values

Figure 7: Performance of factor portfolios

This figure presents the volatility-adjusted cumulative return after transaction costs (on a log scale) on the out-of-sample SDF Portfolio, carry, momentum, value, U.S. Stock Market and U.S. T-Bills. The returns on the FX strategies are calculated using bid and ask quotes from Datastream for the relevant forward and spot rates, with an adjustment to the bid-ask spread that reduces it by 50%. The FX strategy return series are adjusted to have equal volatility. The volatility adjustment scales the returns on the FX strategy i by $\sigma_{\mathrm{MKT}}/\sigma_i$, where $\sigma_{\mathrm{MKT}} = 14.53\%$ is the volatility of the U.S. Stock Market return. All FX strategies are constructed so that they employ 1:2 leverage before scaling. Carry, Momentum and Value are the returns of the H/L portfolios for the currency carry, momentum and value sorted portfolios, which are long the high-characteristic portfolio 5 and short the low-characteristic portfolio 1. The SDF Portfolio return is constructed by rescaling the out-of-sample optimal Lagrange multipliers obtained in the constrained SDF problem in Definition 1 so that (a) the strategy leverage is scaled to 1:2 (i.e., the strategy goes long and short at most $1 worth of foreign currency similarly to e.g., the H/L carry strategy), and (b) the strategy never requires borrowing at the risk-free rate. The details of the normalization are given in Section 6. NBER recessions (peaks/troughs) are shaded. The data are monthly, the full sample period runs from 1976-02 to 2020-10, the SDF Portfolio is estimated on a rolling window of 180 observations, and updated every month. The out-of-sample evaluation period runs from 1991-03 to 2020-10.

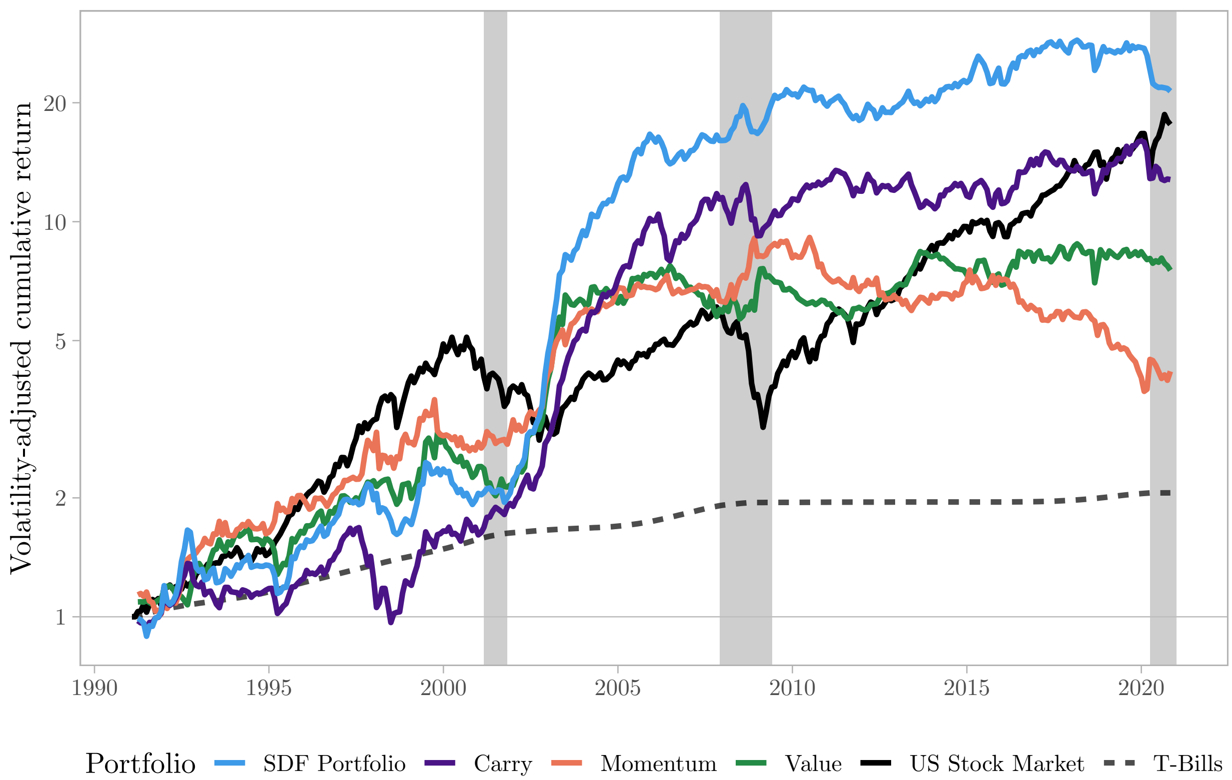

Internet Appendix

Benchmark Currency Stochastic Discount Factors *

November 15, 2021

*Orłowski: HEC Montréal, e-mail: piotr.orlowski@hec.ca; Sokolovski: HEC Montréal, email: valeri.sokolovski@hec.ca; Sverdrup: Stanford University, email: erikcs@stanford.edu.

List of tables

List of figures

Table IA.1: Cross-sectional asset pricing results: carry trade portfolios

The table presents factor premiums for two-factor models estimated by a two-pass regression with no intercept. Factor risk exposures (betas) are estimated in a first-stage time-series regression. Risk prices are the mean slopes of period-by-period cross-sectional regressions of portfolio excess returns on betas. The test assets are excess returns to five carry portfolios. The factors are the dollar factor as in Lustig, Roussanov, and Verdelhan (2011), difference between the excess return on the high-interest rate currency portfolio and the low-interest-rate portfolio (Carry) as in Lustig et al. (2011), changes to the global option-implied FX volatility (ΔVXY), the aggregate FX liquidity factor (ΔFX liq.) as in Karnaukh, Ranaldo, and Söderlind (2015), changes to the CBOE Volatility Index (ΔVIX), and changes to the TED spread (ΔTED). Standard errors, in parentheses, are computed using GMM with the Newey and West (1987) kernel and a twelve-month bandwidth. R^2 refers to the adjusted R-squared from the cross-sectional regression of average excess returns on factor betas, and MAPE is mean absolute pricing error of monthly returns (in %). The $\chi^2(\alpha = 0)$ test statistic, $\alpha' [V_\alpha]^{-1} \alpha$, tests the null hypothesis that all the pricing errors, α, are jointly zero (V_α is based on the GMM variance-covariance matrix as described above). The $\chi^2(\beta = 0)$ test statistic, $\beta' [V_\beta]^{-1} \beta$, tests the null hypothesis that all the betas, β, estimated in the first-stage time-series regressions are jointly zero (V_β is based on the asymptotic variance-covariance matrix). $\beta max - \beta min$ refers to the difference in estimated factor betas of the portfolios with the largest and the smallest beta estimates (Newey and West (1987) standard errors are in parentheses). The difference in betas test refers to the factor other than the DOL. The full sample period is from 1976-02 to 2020-10. The empirical SDFs are estimated using a rolling estimation window of 180 months. The cross-sectional tests are conducted on the out-of-sample evaluation period which runs from 1991-03 to 2020-10. *** indicates significance at 1 percent, ** – at 5 percent, * – at 10 percent.

	(1)	(2)	(3)	(4)	(5)
Carry	0.006***				
	(0.002)				
ΔVXY		$-0.006^{\star\star}$			
		(0.003)			
ΔFX liq.			-0.004		
			(0.008)		
ΔVIX				$-0.040^{\star\star\star}$	
				(0.012)	
ΔTED					-0.004
					(0.017)
DOL	0.002	0.002	0.002	0.002	0.002
	(0.001)	(0.001)	(0.002)	(0.001)	(0.005)
R^2	0.93	0.91	0.90	0.87	0.65
MAPE (%)	0.050	0.059	0.060	0.077	0.125
$p(\chi^2(\alpha = 0))$	0.18	0.25	0.98	0.46	1.00
$p(\chi^2(\beta = 0))$	0.00	0.00	0.00	0.00	0.05
$\beta_{\max} - \beta_{\min}$	1.000***	1.055***	2.111***	0.187***	1.291
	(0.07)	(0.23)	(0.78)	(0.05)	(0.88)

Table IA.2: Cross-sectional asset pricing results carry, momentum and value

This table present factor premiums for two-factor models estimated by a two-pass regression, with no intercept. Factor risk exposures (betas) are estimated in a first-stage time-series regression. Risk prices are the mean slopes of period-by-period cross-sectional regressions of portfolio excess returns on betas. The test assets are excess returns to five carry, momentum, and value portfolios. The factors are the dollar factor as in Lustig et al. (2011), difference between the excess return on the high-interest rate currency portfolio and the low-interest-rate portfolio (Carry) as in Lustig et al. (2011), changes to the global option-implied FX volatility (ΔVXY), the aggregate FX liquidity factor (ΔFX liq.) as in Karnaukh et al. (2015), changes to the CBOE Volatility Index (ΔVIX), and changes to the TED spread (ΔTED). Standard errors, in parentheses, are computed using GMM with the Newey and West (1987) kernel and a twelve-month bandwidth. R^2 refers to the adjusted R-squared from the cross-sectional regression of average excess returns on factor betas, and MAPE is mean absolute pricing error of monthly returns (in %). The $\chi^2(\alpha = 0)$ test statistic, $\alpha' \left[V_\alpha\right]^{-1} \alpha$, tests the null hypothesis that all the pricing errors, α, are jointly zero (V_α is based on the GMM variance-covariance matrix as described above). The $\chi^2(\beta = 0)$ test statistic, $\beta' \left[V_\beta\right]^{-1} \beta$, tests the null hypothesis that all the betas, β, estimated in the first-stage time-series regressions are jointly zero (V_β is based on the asymptotic variance-covariance matrix). $\beta max - \beta min$ refers to the difference in estimated factor betas of the portfolios with the largest and the smallest beta estimates (Newey and West (1987) standard errors are in parentheses). The difference in betas test refers to the factor other than the DOL. The full sample period is from 1976-02 to 2020-10. The empirical SDFs are estimated using a rolling estimation window of 180 months. The cross-sectional tests are conducted on the out-of-sample evaluation period which runs from 1991-03 to 2020-10. *** indicates significance at 1 percent, ** – at 5 percent, * – at 10 percent.

	(1)	(2)	(3)	(4)	(5)
Carry	0.007***				
	(0.002)				
ΔVXY		-0.002			
		(0.002)			
ΔFX liq.			-0.000		
			(0.001)		
ΔVIX				-0.021^{**}	
				(0.009)	
ΔTED					-0.001
					(0.004)
DOL	0.002	0.002	0.002	0.002	0.002
	(0.001)	(0.001)	(0.002)	(0.001)	(0.002)
R^2	0.53	0.11	0.05	0.21	0.09
MAPE (%)	0.108	0.152	0.158	0.141	0.154
$p(\chi^2(\alpha = 0))$	0.29	0.00	0.00	0.17	0.39
$p(\chi^2(\beta = 0))$	0.00	0.00	0.00	0.00	0.09
$\beta_{\max} - \beta_{\min}$	1.000***	1.055***	2.111***	0.186***	1.538*
	(0.07)	(0.23)	(0.78)	(0.05)	(0.93)

Table IA.3: Summary statistics of FX hedge fund characteristics

The table presents summary statistics of FX hedge fund characteristics for the sample of hedge funds in the BarclayHedge Currency Traders Database. The sample period runs from January 1994 to December 2019. N is the number of unique hedge funds for which data on a particular characteristic are available. The statistics are for hedge fund AUM (in USD million), fund age (in years), management fee (in %), incentive fee (in %), a dummy indicating if the fund has a high watermark (HWM), and the redemption notice (in days). The statistics are time-series averages of monthly cross-sectional statistics and represent a typical distribution of hedge fund characteristics available in a given month in the sample.

Panel A: FX hedge funds returns

N	Mean	SD	Min	Max	SD^{CS}
100	0.53	1.39	−2.66	5.35	0.38

Panel B: Global macro hedge funds returns

N	Mean	SD	Min	Max	SD^{CS}
287	0.76	1.75	−3.46	6.97	0.53

Panel C: FX hedge funds characteristics

	N	Mean	Median	25 percentile	75 percentile	Min	Max
AUM	100	466.92	173.86	83.32	477.99	5.06	2771.30
Age	100	6.53	6.10	3.60	8.56	0.61	17.65
Mgmt. fee	100	1.40	1.55	0.91	2.00	0.00	2.39
Inct. fee	100	18.63	20.00	19.97	20.00	2.13	28.89
HWM	100	0.29	0.00	0.00	0.57	0.00	1.00
Redemption	100	3.39	0.00	0.00	0.58	0.00	33.40

Panel D: Global macro hedge funds characteristics

	N	Mean	Median	25 percentile	75 percentile	Min	Max
AUM	287	291.81	98.07	28.89	316.56	1.66	3560.45
Age	287	5.29	4.36	2.53	7.35	0.50	18.01
Mgmt. fee	287	1.55	1.66	1.13	2.00	0.19	3.13
Inct. fee	287	17.58	20.00	18.51	20.00	0.00	30.82
HWM	287	0.85	1.00	0.99	1.00	0.00	1.00
Redemption	287	0.07	0.06	0.02	0.09	0.01	0.29

Table IA.4: Cross-sectional asset pricing FX hedge funds: existing models

The tables present factor premiums for factor models estimated by Fama and MacBeth (1973) regression. Factor risk exposures (betas) are estimated using 48-month rolling first-stage time-series regressions. Risk prices are the mean slopes of period-by-period cross-sectional regressions of portfolio excess returns on betas. The test assets are the 100 largest FX hedge funds in BarclayHedge Currency Traders database. The factors are the difference between the excess return on the high-interest rate currency portfolio and the low-interest-rate portfolio (Carry) as in Lustig et al. (2011), changes to the global option-implied FX volatility (ΔVXY), the aggregate FX liquidity factor (ΔFX liq.) as in Karnaukh et al. (2015), excess return on the U.S. stock market (MKT), changes to the CBOE Volatility Index (ΔVIX), changes to the TED spread (ΔTED), and the shocks to the primary dealer capital ratio (HKM) as in He, Kelly, and Manela (2017). Panel B reports regression specifications that include the following hedge fund characteristics: logarithm of each hedge fund's AUM, fund age, management fee (in percent), performance fee, a dummy indicating if the fund has a high watermark, and the redemption notice period. Newey and West (1987) standard errors with 12 lags are reported in parenthesis. R^2_{FMB} refers to the time-series average of the adjusted R-squareds of period-by-period cross-sectional regressions. The out-of-sample evaluation from 1994-01 to 2017-12.

Panel A: Specification without hedge fund controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Carry	0.004*						
	(0.002)						
ΔVXY		−0.001					
		(0.001)					
ΔFX liq.			−0.001*				
			(0.000)				
ΔVIX				0.001			
				(0.005)			
ΔTED					−0.000		
					(0.000)		
MKT						−0.002	−0.002
						(0.004)	(0.004)
HKM							0.005
							(0.007)
Intercept	0.005***	0.006***	0.006***	0.005***	0.005***	0.005***	0.005***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Controls	No	No	No	No	No	No	No
R^2_{FMB}	0.11	0.14	0.12	0.13	0.09	0.14	0.24
N_{Assets}	100	100	100	100	100	100	100

Panel B: Specification with hedge fund controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Carry	0.003						
	(0.003)						
ΔVXY		−0.001					
		(0.001)					
ΔFX liq.			−0.000				
			(0.000)				
ΔVIX				−0.003			
				(0.005)			
ΔTED					−0.000		
					(0.000)		
MKT						−0.001	−0.002
						(0.004)	(0.004)
HKM							0.006
							(0.007)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2_{FMB}	0.39	0.42	0.40	0.40	0.38	0.41	0.50
N_{Assets}	100	100	100	100	100	100	100

Table IA.5: Correlation of $\widehat{SDF}^{L_{50}}$ with the existing factors

The table presents the factor correlation matrix. The factors are the out-of-sample empirical SDF with a 1:50 leverage constraint ($\widehat{\text{SDF}}^{L_{50}}$), the difference between the excess return on the high-interest rate currency portfolio and the low-interest-rate portfolio (Carry) as in Lustig et al. (2011), changes to the global option-implied FX volatility (ΔVXY), the aggregate FX liquidity factor (ΔFX liq.) as in Karnaukh et al. (2015), excess return on the U.S. stock market (MKT), changes to the CBOE Volatility Index (ΔVIX), changes to the TED spread (ΔTED), and the shocks to the primary dealer capital ratio (HKM) as in He et al. (2017). The sample period runs from 1991-03 to 2020-10.

	Carry	ΔVXY	ΔFX liq.	ΔVIX	ΔTED	MKT	HKM
$\widehat{SDF}^{L_{50}}$	−0.65	0.29	0.17	0.21	0.09	−0.18	−0.16
Carry		−0.42	−0.27	−0.31	−0.10	0.31	0.31
ΔVXY			0.55	0.55	0.23	−0.43	−0.37
ΔFX liq.				0.25	0.18	−0.21	−0.15
ΔVIX					0.26	−0.73	−0.53
ΔTED						−0.20	−0.11
MKT							0.72

Figure IA.1: Financial sector leverage

The figure shows the quarterly time series of investment bank leverage. Leverage is defined as (Book Debt + Market Equity)/Market Equity. The time series are constructed from a cross-section of 29 individual banks that are designated as New York Federal Reserve Primary Dealers i.e., key financial institutions. The plotted series are cross-sectional statistics. The tabulated series are the time-series averages of the cross-sectional statistics. NBER recessions are shaded in grey. The sample period runs from 1999-Q1 to 2020-Q4.

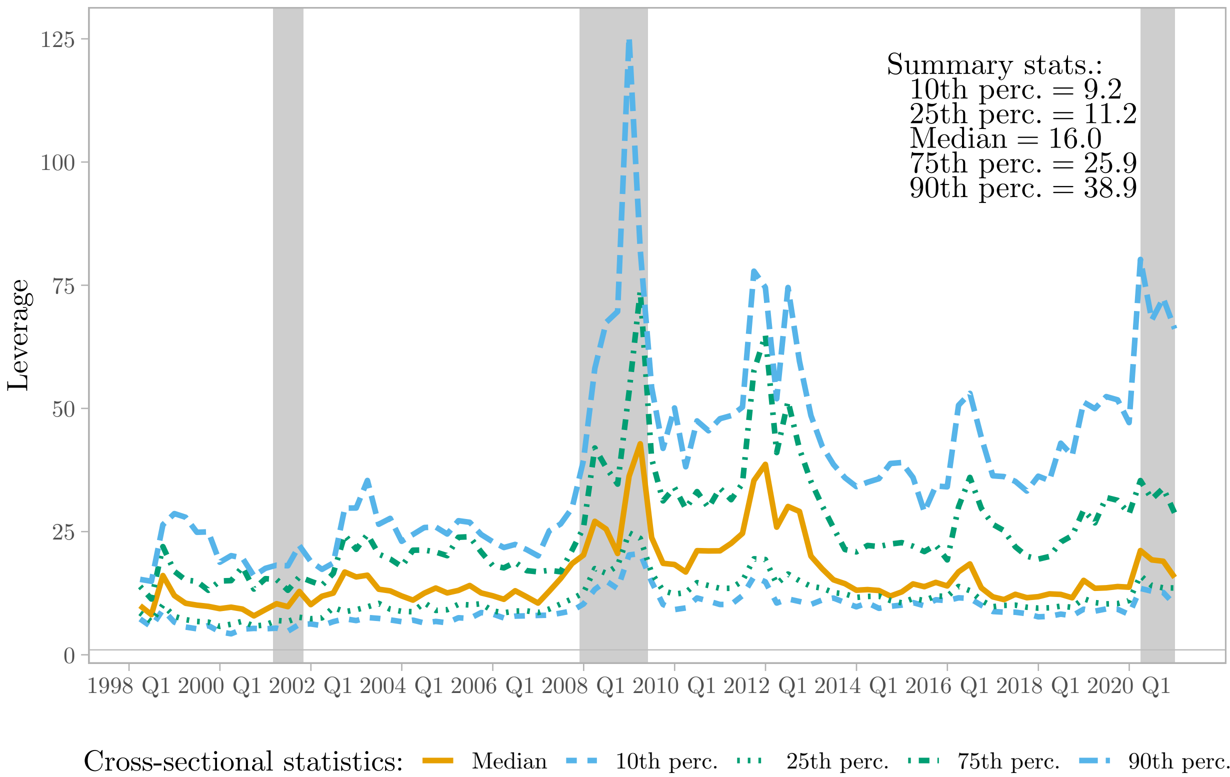

Figure IA.2: Time series of the unconstrained out-of-sample $\widehat{\mathrm{SDF}}^{U}$

The figure shows the monthly time series of the out-of-sample empirical SDF without a leverage constraint ($\widehat{\mathrm{SDF}}^{U}$), estimated on five currency carry, momentum, and value portfolios. The empirical SDF is estimated using a rolling window of 180 months, and the out-of-sample step size is one month. The data are monthly and the full sample period runs from 1976-2 to 2020-10, while the out-of-sample evaluation period runs from 1991-03 to 2020-10. NBER recessions and FX market crises periods are shaded in grey and blue, respectively.

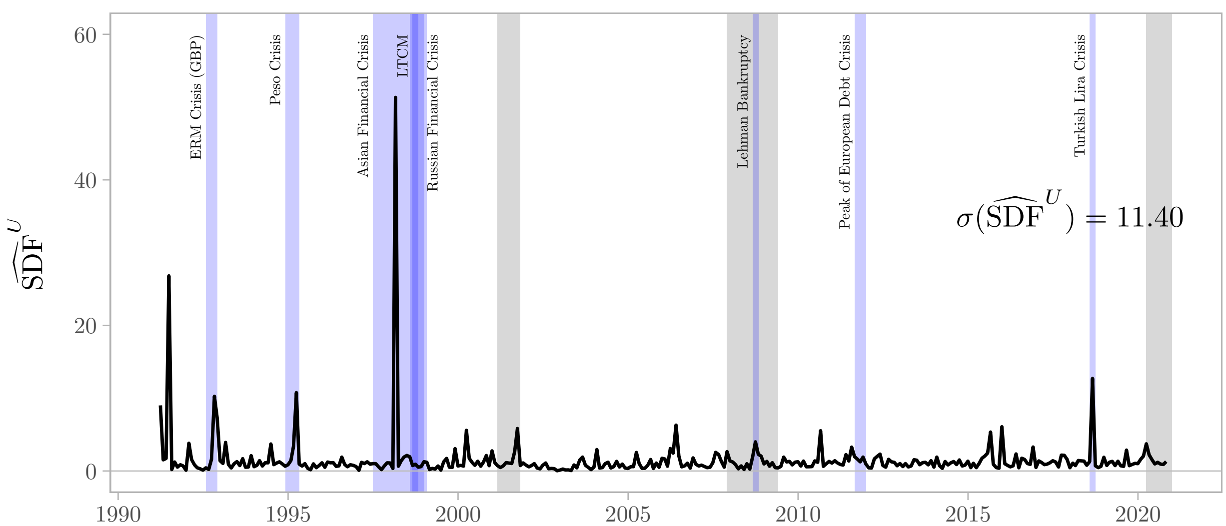

Figure IA.3: Time series of the leverage-constrained SDF and global FX implied volatility

The figure shows the standardized monthly time series of the out-of-sample empirical SDF with a 1:50 leverage constraint ($\widehat{\text{SDF}}^{L_{50}}$), estimated on five currency carry, momentum, and value portfolios, and the standardized option-implied aggregate FX volatility (VXY). The empirical SDF is estimated using a rolling window of 180 months, and the out-of-sample step size is one month. The data are monthly and the full sample period runs from 1976-2 to 2020-10, while the out-of-sample evaluation period runs from 1991-03 to 2020-10. NBER recessions and FX market crises periods are shaded in grey and blue, respectively.

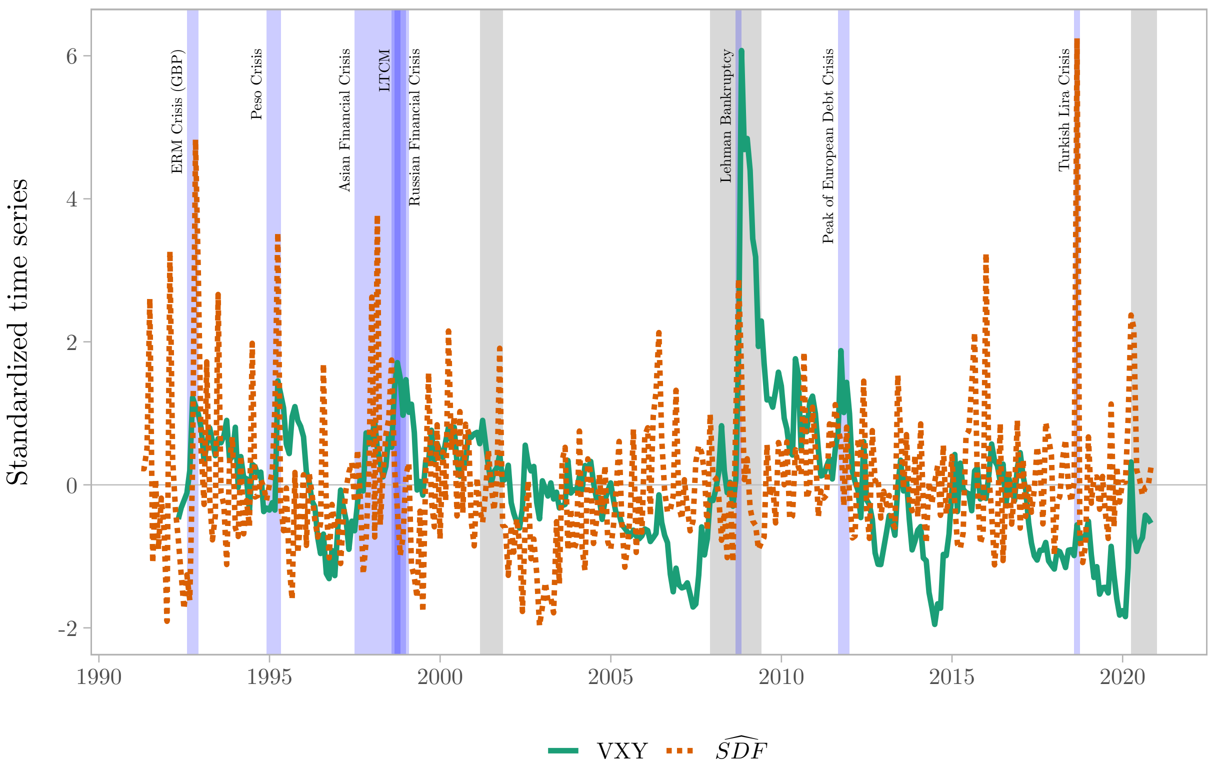

Figure IA.4: Cross-sectional asset pricing results carry, momentum and value

The figure shows risk price estimates and several measures of model fit for out-of-sample empirical SDFs (estimated with different levels of maximum leverage constraint). The empirical SDFs are constructed from the five carry, momentum, and value portfolios. The test assets are excess returns to five carry, momentum, and value portfolios. Risk prices, shown in panel (a), are the mean slopes of period-by-period cross-sectional regressions of portfolio excess returns on betas. The 95% confidence interval is displayed in red along each risk price estimate. Confidence interval are constructed using standard errors that are computed using GMM with the Newey and West (1987) kernel and a twelve-month bandwidth. Panel (b) shows MAPEs, mean absolute pricing error of monthly returns (in %). Panel (c) shows adjusted R^2s from the cross-sectional regression of average excess returns on factor beta. Panel (d) shows the p-values of the $\chi^2(\alpha = 0)$ test statistic, $\alpha' \left[V_\alpha\right]^{-1} \alpha$, which tests the null hypothesis that all the pricing errors, α, are jointly zero (V_α is based on the GMM variance-covariance matrix as described above). The full sample period is from 1976-02 to 2020-10. The empirical SDFs are estimated using a rolling estimation window of 180 months. The cross-sectional tests are conducted on the out-of-sample evaluation period which runs from 1991-03 to 2020-10.

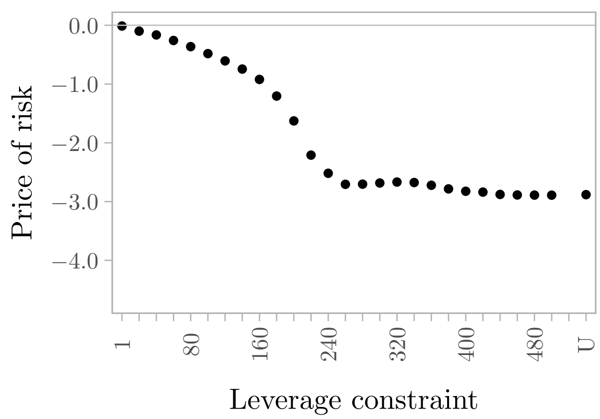

(a) Estimated price of risk

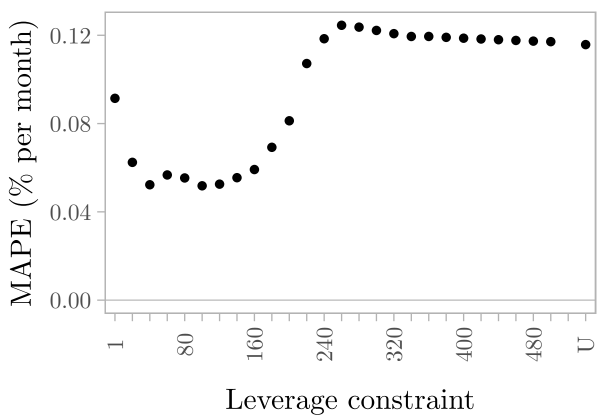

(b) Mean absolute pricing error

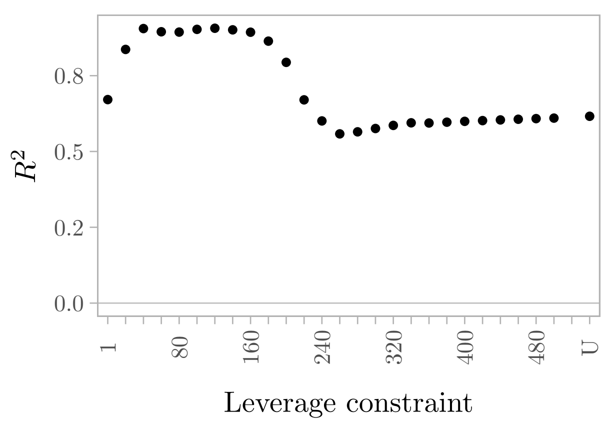

(c) Adjusted R^2

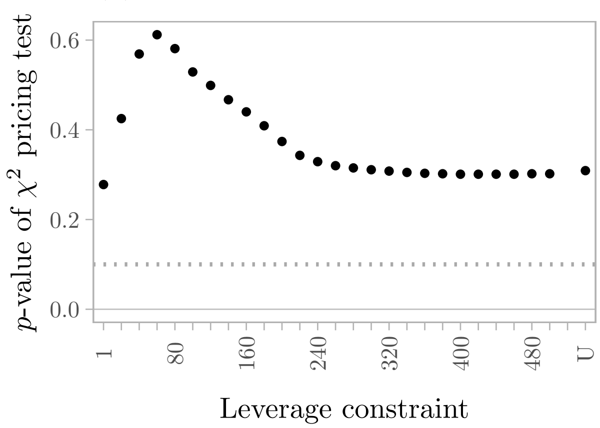

(d) Asset pricing test p-values

Figure IA.5: Cross-sectional asset pricing: empirical SDF with concentration constraint

The figure shows risk price estimates and several measures of model fit for out-of-sample empirical SDFs (estimated with different levels of ridge/concentration constraint R). The empirical SDFs are constructed from the five carry, momentum, and value portfolios. The test assets are excess returns to five carry, momentum, and value portfolios. Risk prices, shown in panel (a), are the mean slopes of period-by-period cross-sectional regressions of portfolio excess returns on betas. The 95% confidence interval is displayed in red along each risk price estimate. Confidence interval are constructed using standard errors that are computed using GMM with the Newey and West (1987) kernel and a twelve-month bandwidth. Panel (b) shows MAPEs, mean absolute pricing error of monthly returns (in %). Panel (c) shows adjusted R^2s from the cross-sectional regression of average excess returns on factor beta. Panel (d) shows the p-values of the $\chi^2(\alpha = 0)$ test statistic, $\alpha' \left[V_\alpha\right]^{-1} \alpha$, which tests the null hypothesis that all the pricing errors, α, are jointly zero (V_α is based on the GMM variance-covariance matrix as described above). The full sample period is from 1976-02 to 2020-10. The empirical SDFs are estimated using a rolling estimation window of 180 months. The cross-sectional tests are conducted on the out-of-sample evaluation period which runs from 1991-03 to 2020-10.

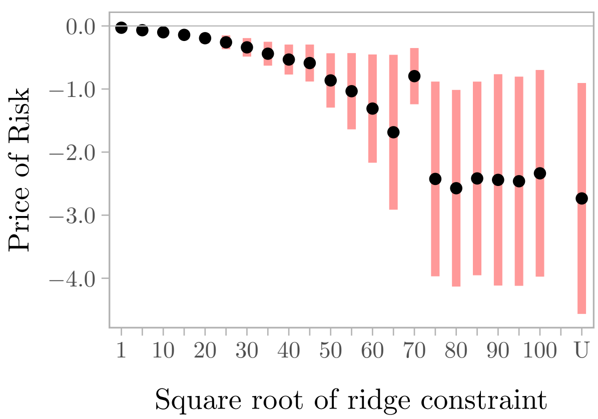

(a) Estimated price of risk

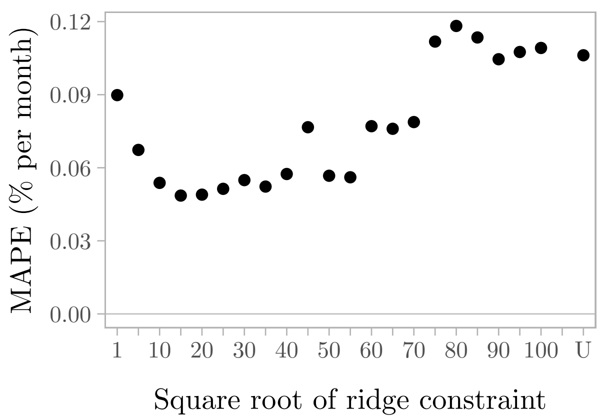

(b) Mean absolute pricing error

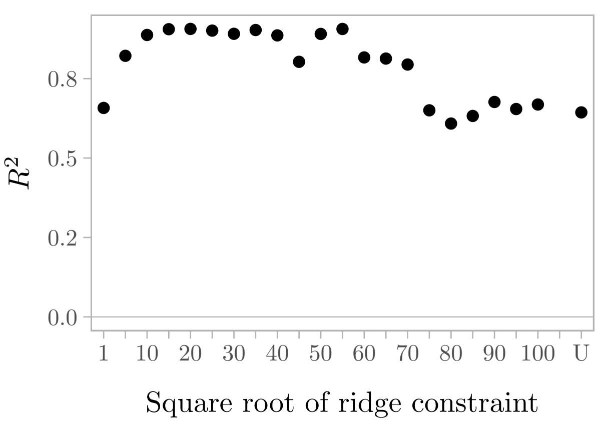

(c) Adjusted R^2

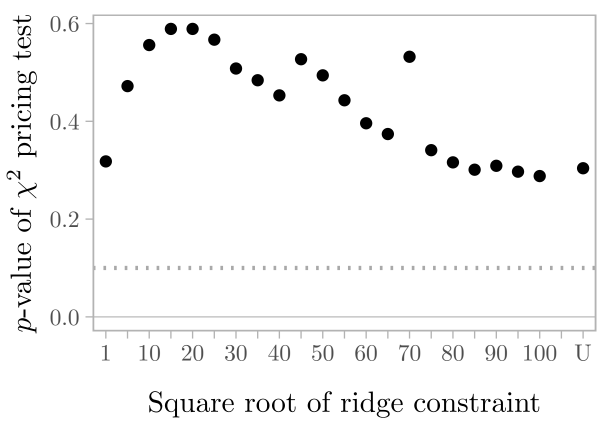

(d) Asset pricing test p-values

References

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

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

Karnaukh, Nina, Angelo Ranaldo, and Paul Söderlind, 2015, Understanding FX liquidity, *Review of Financial Studies* 28, 3073–3108.

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan, 2011, Common risk factors in currency markets, *Review of Financial Studies* 24, 3731–3777.

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–08.