Exchange Rate Expectations and Currency Demand

Leonie Bräuer[*][†]

November 19, 2025

Access the latest version here

Abstract

This paper develops a new method to extract exchange rate expectations from investment positions. I use relative allocations between otherwise identical exchange-traded funds (ETFs) offered with and without a currency hedge to measure investors' pure currency demand and infer a distribution of currency return expectations. These portfolio-implied expectations predict future exchange rates more accurately than survey-based expectations or expectations derived from macroeconomic models or currency pricing factors. Dispersion in portfolio-implied expectations accounts for 27% of exchange rate volatility, consistent with models of heterogeneous beliefs.

JEL Codes: E44, F31, F32, G11, G15, G23

Keywords: Foreign Exchange, Exchange Rate Expectations, Exchange Traded Funds, Global Currency Hedging

[*]I would like to thank Harald Hau, Karen Lewis, and Fabio Trojani for their invaluable guidance. I thank Tony Berrada, Ines Chaieb, Wenxin Du, Itamar Drechsler, Robin Greenwood, Travis Johnson, Sean Myers, Nick Roussanov, Clemens Sialm, Yao Zeng and seminar participants at Wharton, GFRI and SFI Research Days for valuable comments and suggestions.

[†]GFRI, University of Geneva and Swiss Finance Institute. Email: leonie.braeuer@unige.ch.

1 Introduction

Economists increasingly rely on survey expectations to explain asset prices and constrain macroeconomic models.[1] Yet, how well such survey expectations relate to investors' portfolio decisions remains an unsettled question.[2] A survey sample could diverge from the relevant group of investors that decides on financial allocation and from the marginal group that influences prices. Moreover, surveys may fail to elicit investors' true beliefs and expectations inferred from equilibrium models or risk factors may be more informative.

An alternative promising path toward more accurate return expectations is to recover them directly from portfolio allocations. Such holding-inferred expectations are tied to actual invested capital, giving investors skin in the game and thus an incentive to reveal their true beliefs. But inferring exchange rate expectations from currency holdings is particularly challenging, as currency exposure is often incidental to foreign asset investment, which makes it difficult to isolate pure currency demand from asset demand. This paper addresses this challenge by exploiting the observable asset allocation between share classes of exchange-traded funds (ETFs). Specifically, for a given foreign index, many funds can be purchased currency-hedged or unhedged, allowing a clean separation of currency demand from asset demand. The analysis here combines two elements.

First, I construct a novel dataset covering the universe of international ETFs for 2014–2025. I match ETFs domiciled in six currency areas outside the US that track the same underlying US benchmark index but differ in their currency hedging strategy.[3] My final dataset comprises 350 currency-hedged ETFs with their 350 paired unhedged twins, with aggregate holdings accounting for roughly 5% and 1% of total foreign holdings in USD-denominated equities and bonds, respectively.[4]

Second, I build on the mean-variance framework for optimal currency hedging pioneered by Anderson and Danthine (1981). This theory predicts that the relative size of two otherwise identical

[1]For example, survey expectations are used to identify the Uncovered Interest Parity (UIP) premium and explain the forward premium anomaly (Kalemli-Özcan and Varela, 2021; Dao et al., 2025), isolate expected inflation from inflation risk premia (Chernov and Mueller, 2012), or quantify information rigidity to discipline expectations formation (Coibion and Gorodnichenko, 2015).

[2]Cochrane (2011) notes a disconnect between investors' portfolio choice and their survey responses, while recent research establishes a close connection between beliefs and holdings in equity markets (Giglio et al., 2021).

[3]International ETFs represent a substantial portion of global ETF investment. For example, in December 2024, ETFs domiciled in the euro area that invest in international equity had assets under management totaling USD 482 billion, while those investing in European equity managed just under half of that, with USD 188 billion in assets.

[4]Du and Huber (2024) find that in June 2021 foreigners held USD 12 trillion USD-denominated equities and USD 20 trillion in USD-denominated bonds. Their sample includes all non-US based investors, while I focus on holdings by six currency areas.

ETFs, differing only in their currency exposure, reflects two separable components: a hedging motive, given by the covariance between currency and asset returns scaled by exchange rate volatility, and a speculative motive given by the risk-adjusted expected currency excess return. Based on this framework, I can identify pure currency demand and recover a distribution of currency return expectations from investors' relative allocations between unhedged and hedged ETFs.

The combination of ETF holding data and financial theory that maps relative holdings to exchange rate expectations allows me to address the following three research questions:

1. How well can survey-based exchange rate expectations account for the relative portfolio allocation of ETF investors compared to expectations derived from macroeconomic theory or asset pricing factors?

2. If we use relative ETF allocations as a new inference method to deduce exchange rate expectations, how accurate can such portfolio-implied expectations predict future returns relative to expectations based on surveys, macro models or asset pricing factor models?

3. Is a higher dispersion in portfolio-implied exchange rate expectations (across ETF pairs) related to higher foreign exchange (FX) volatility as predicted by heterogeneous belief models (Basak, 2005; Buraschi and Jiltsov, 2006)?

The answers to these three questions can be summarized as follows:

First, survey expectations outperform expectations derived from macro or asset pricing models when it comes to explaining the observed portfolio allocations between hedged and unhedged ETF pairs. I focus on two existing model-implied estimates of conditional currency return expectations for benchmarking: macro-based expectations that combine a mean reversion in the real exchange rate implied by purchasing power parity with a trend signal from past returns as in Chernov et al. (2023), and factor-based expectations derived from the two most prominent currency pricing factors, the dollar and carry factor (Lustig et al., 2011; Verdelhan, 2018) as in Opie and Riddiough (2020). ETF investors' USD holdings show no significant sensitivity to either estimate of the latter two model-implied expectations, once controlling for aggregate time trends. In contrast, survey expectations sourced from Consensus Economics are highly significant in explaining the cross-section of currency holdings: Investors in currency areas that are expected to depreciate more against the USD according to surveys, allocate a larger share of their portfolios to the unhedged ETF, increasing their exposure to USD fluctuations. This shows that FX survey expectations are indeed a suitable inference method for rationalizing international portfolio allocations.

Second, I show that the accuracy of FX survey expectations is still surpassed by expectations deduced from relative ETF allocations. A horse race between portfolio-implied exchange rate expectations, survey expectations and those embedded in macro and factor models shows that only the first have predictive power for future exchange rates, over the 11-year period under consideration. I highlight that the predictive power of portfolio-implied expectations stems primarily from active flows into the unhedged ETF relative to the hedged ETF, however, the theory-guided portfolio-implied expectations contain incremental and stronger predictive information beyond flows.

Third, I find that the cross-sectional dispersion of portfolio-implied expectations increases significantly during volatile FX market periods, consistent with models of heterogeneous beliefs. Specifically, the mean absolute deviation of portfolio-implied expectations can account for 27% of FX option-implied volatility and is positively correlated with it at 71% for the aggregate USD basket. Quantitatively, the explanatory power is comparable to that of order flow or macroeconomic variables.[5]

My paper provides an indicator with immediate practical use: By recovering expectations from daily ETF trading, I construct a high-frequency gauge of investors' sentiment towards the USD. As an illustrative example, following the announcement of US tariffs on April 2, 2025, portfolio-implied expectations as well as flows into unhedged relative to hedged ETFs fell sharply (see Figure A.1), consistent with expectations of a USD depreciation. This pattern aligns with Jiang et al. (2025) who argue, based on price movements, that investors started questioning the dominant role of the USD in the international financial system, following the day of the announcement.[6] The example underscores that investors' traded quantities can deliver a near-real-time indicator of foreign investors' USD perceptions.

2 Related Literature

My study contributes to the literature on subjective beliefs derived from survey expectations (Greenwood and Shleifer, 2014; Adam and Nagel, 2023), particularly, on the link between be-

[5] Order flows explain 40% and 67% of variation in Evans and Lyons (2002), and macroeconomic variables account for roughly 26% of exchange rate fluctuations in Koijen and Yogo (2020).

[6] Similarly, Shin et al. (2025) document increased hedging by non-US institutions in USD index futures during this period, using the Commodity Futures Trading Commission's (CFTC) weekly data.

liefs and portfolio decisions. Previous literature documents a significant link between individual investor holdings and their survey responses (Amromin and Sharpe, 2014; Ameriks et al., 2020; Giglio et al., 2021). Equally important is establishing this link for publicly available average (consensus) forecasts, on which many studies rely on. For example, in the context of currency survey expectations, sourced from Consensus Economics, studies include Kremens et al. (2025); Nagel and Xu (2023); Kalemli-Özcan and Varela (2021); Della Corte et al. (2023); Candian and De Leo (2025); De Marco et al. (2022); Pesch et al. (2024); Bartram et al. (2025); Dao et al. (2025).[7] Stavrakeva and Tang (2020) show that FX survey expectations align closely with dealer banks' aggregate futures positions. In contrast to their paper and more generally to approaches that use only FX derivative data, I observe each ETF's exact underlying assets and can thus control for hedging demand to cleanly isolate the belief-driven currency demand. I provide new cross-sectional validation of average (consensus) currency survey expectations. More specifically, I show a strong cross-sectional relationship between the USD share and expected USD returns from surveys, with a sensitivity quantitatively comparable to investor-level estimates in Giglio et al. (2021).[8]

My paper relates to the literature on recovering beliefs Ross (2015); Ghosh and Roussellet (2023), in particular the strand that infers beliefs from holdings data. A recent example includes Crescini et al. (2025) who recover market return expectations from option holdings data without relying on a specific asset pricing model. I instead adopt a mean–variance framework whose virtue is its tractable and parsimonious structure, akin to French and Poterba (1991). The latter infer international investors' expected equity returns from aggregate country-level portfolio holdings. By contrast, I exploit a cross-section of 350 ETF pairs and show that this level of granularity is consequential: unlike my advocated portfolio-implied expectations, expectations derived from a single aggregate equity and bond index do not significantly predict future returns. A study closely related to my paper is Egan et al. (2022). They also recover beliefs from ETFs and infer market return expectations from demand for leveraged ETFs. In their paper, identification comes from investors' choice between leveraged and non-leveraged ETFs, a margin likely used by a selected, more risk-tolerant clientele. In my study, I exploit the choice between unhedged and hedged ETFs,

[7] In the context of inflation expectations, Nagel and Yan (2022) show a strong link between investors' aggregate financial decisions and their expectations, using average inflation expectations and retail flows into TIPS ETFs.

[8] Specifically, a one standard deviation increase in expected spot returns is associated with a 0.21 standard deviation increase in the USD share, comparable to the 0.16 sensitivity reported by Giglio et al. (2021) between individual investors' equity shares and their expected equity returns from surveys.

a ubiquitous decision for most non-US investors. I find that the average expected return of investors forecasts future returns, which contrasts with Egan et al. (2022). However, consistent with existing literature, I document substantial and time-varying heterogeneity in beliefs among investors.

I also contribute to the literature on heterogeneous beliefs (e.g., Scheinkman and Xiong (2003); Bacchetta and Van Wincoop (2006); Jouini and Napp (2007); Dumas et al. (2009, 2017)) by providing new empirical evidence of a close link between differences in investor expectations and volatility. In Buraschi and Jiltsov (2006) and Buraschi et al. (2014) volatility and belief dispersion move together as optimistic investors insure pessimists in high volatility states, earning a premium that rises with disagreement. Beber et al. (2010) confirm this theory for three currency pairs over the period 1993-2006 using dispersion in survey forecasts, and find that differences in beliefs explain on average 45% of option-implied FX volatility. I complement this evidence by showing that, for a cross-section of six currency pairs, in a much more recent time period, the correlation between differences in beliefs derived from portfolio holdings and implied volatility reaches up to 71% for the aggregate USD basket.

This study also relates to the growing empirical literature on currency hedging which builds on theoretical work on optimal currency hedging by Anderson and Danthine (1980, 1981), or more recently by Campbell et al. (2010). Du and Huber (2024) is the first comprehensive study to examine global currency hedging at an aggregate level, comparing hedging quantities to mean-variance optimal benchmarks. With the growing availability of data, recent advances shift toward comparing optimal hedging strategies at the fund level. For example, Sialm and Zhu (2024), Cheema-Fox and Greenwood (2024), Opie and Riddiough (2024) and Chen and Zhou (2025) analyze fund-level hedging behavior by US investment funds, while Bräuer and Hau (2024) and Kubitza et al. (2024) focus on hedging by European investment funds. This study differs from the existing literature in one key aspect: rather than examining the hedging behavior of funds, I focus on investors of ETFs that are either fully hedged or completely unhedged. As a result, the optimal currency hedging problem transitions from the fund level to the investor level, placing the emphasis on investor decisions regarding hedging strategies rather than on the funds' hedging strategies. Furthermore, this study highlights the speculative component of foreign currency demand, the expectations about exchange rates implied by currency portfolios.

My study also contributes to the expanding field of international asset pricing, where the number of currency pricing factors continues to grow beyond the well-known carry and dollar factors (Lustig et al., 2011; Verdelhan, 2018), with new additions such as momentum (Burnside et al., 2011; Menkhoff et al., 2012b), value (Menkhoff et al., 2017), volatility (Menkhoff et al., 2012a), illiquidity (Mancini et al., 2013), downside beta risk (Lettau et al., 2014), and economic momentum (Dahlquist and Hasseltoft, 2020). In contrast to this existing literature on FX factors, this study asks whether investors rely on the predictive signals implied by asset pricing models when deciding on global currency allocations. This question is largely overlooked in the literature, even though asset pricing models should not only replicate asset price behavior but also align with observable investor expectations. I show that aggregate investors do not systematically tilt their portfolio toward conditional expectations derived from the two most prominent factors (dollar and carry) to generate excess returns. Meanwhile, I find that trading strategies built on information from portfolio-implied expectations deliver higher Sharpe ratios than traditional factors.

Lastly, this paper adds new evidence to the literature on ETF investment and its effect on prices (e.g., Ben-David et al. (2017, 2018); Da and Shive (2018); Glosten et al. (2021); Pan and Zeng (2017); Koont et al. (2025)). While the asset allocation is most often passive, when investing internationally, investors face an active decision regarding currency exposure. This paper shows that this decision is rational and that it forecasts currency returns, contrasting with Brown et al. (2021) who attribute ETF flows to non-fundamental demand. Thus, while some ETF trading might be driven by unsophisticated individuals, in the aggregate the relative trading between unhedged and hedged ETFs contains information about future currency returns.

3 From FX Hedging Choices to Exchange Rate Expectations

This section explains how hedging choices map into exchange rate expectations. I first derive optimal foreign currency demand under a predetermined foreign portfolio, then introduce common estimates of exchange rate expectations, and finally describe the estimation of currency and asset risk.

3.1 Optimal Currency Demand

Model Setup. Suppose there exist two ETFs that invest in the same foreign asset index. One does not hedge its currency risk and the other fully hedges its currency risk. Suppose further that there exists one risk-averse domestic investor i in a specific currency area, who has decided to invest in the foreign asset index, i, and can choose between the two products. Her total wealth equals $A_{i,t} = A_{i,t}^{unhedged} + A_{i,t}^{hedged}$ of which she can allocate $w_{i,t}$ to the unhedged ETF so that

$$w_{i,t} = \frac{A_{i,t}^{unhedged}}{A_{i,t}^{unhedged} + A_{i,t}^{hedged}} \tag{1}$$

where assets under management, $A_{i,t}$, are measured as the product of net asset value (NAV) and outstanding shares, and $1 - w_{i,t}$ is the weight in the hedged ETF. The conditional one-month expected returns on the hedged and unhedged ETFs are given by

$$\begin{aligned} \mathbb{E}_t[r_{i,t+1}^{unhedged}] &= \mathbb{E}_t[r_{i,t+1}] + \mathbb{E}_t[\Delta s_{t+1}] \\ \mathbb{E}_t[r_{i,t+1}^{hedged}] &= \underbrace{\mathbb{E}_t[r_{i,t+1}] + \mathbb{E}_t[\Delta s_{t+1}]}_{\text{Asset Return in Home Currency}} + \underbrace{f_t - \mathbb{E}_t[s_{t+1}]}_{\text{Forward Contract Return}} = \mathbb{E}_t[r_{i,t+1}] + f_t - s_t \end{aligned}$$

where $\mathbb{E}_t[r_{i,t+1}]$ represents the expected return of the foreign equity or bond index, s_t represents the log spot rate, and f_t denotes the log one-month forward rate. All exchange rates are denoted such that the base currency corresponds to the foreign currency and the quote currency corresponds to the domestic currency. I use the one-month forward rate, consistent with industry practice of hedging with short-dated forward contracts (Bräuer and Hau, 2024).

The difference in return between the unhedged and hedged ETF equals the expected log currency excess return from buying the dollar in the forward market and selling it in the spot market after one month against the foreign currency, defined as

$$\mathbb{E}_t[rx_{t+1}] = \mathbb{E}_t[s_{t+1}] - f_t = \underbrace{\mathbb{E}_t[\Delta s_{t+1}]}_{\text{Expected Spot Return}} - \underbrace{(f_t - s_t).}_{\text{Forward Premium}} \tag{2}$$

The expected currency excess return comprises the expected USD appreciation, $\mathbb{E}_t[\Delta s_{t+1}]$, and the forward premium, $f_t - s_t$. The forward premium represents the time-t known (deterministic) component of the expected excess return on the unhedged ETF and, under covered interest parity, equals the interest rate differential between the domestic and foreign (US) rate, $i^{dom} - i^* = f_t - s_t$. For investors hedging foreign-denominated assets, the negative of the forward premium $(s_t - f_t)$ is often referred to as hedging cost, since empirically the domestic currency does not appreciate sufficiently to offset the negative interest differential between the domestic and foreign (US) interest rate.

A portfolio consisting of an unhedged and a currency-hedged ETF with return $r^p_{i,t}$ has a conditional mean and variance equal to

$$\mathbb{E}_t[r^p_{i,t+1}] = \mathbb{E}_t[r_{i,t+1}] + w_{i,t}\,\mathbb{E}_t[\Delta s_{t+1}] + (1 - w_{i,t})(f_t - s_t)$$
$$\mathbb{V}\mathrm{ar}_t[r^p_{i,t+1}] = \mathbb{V}\mathrm{ar}_t[r_{i,t+1}] + w^2_{i,t}\,\mathbb{V}\mathrm{ar}_t[\Delta s_{t+1}] + 2w_{i,t}\,\mathbb{C}\mathrm{ov}_t[r_{i,t+1}, \Delta s_{t+1}].$$

Sequential Optimization. The investor's optimal allocation to the unhedged index, $w_{i,t}$, versus the hedged index can be characterized by maximizing the following objective function

$$\max_{w_{i,t}} \quad \mathbb{E}_{i,t}[r^p_{i,t+1}] - \frac{\gamma}{2}\,\mathbb{V}\mathrm{ar}_{i,t}[r^p_{i,t+1}] \tag{3}$$

where γ is the level of risk aversion and $\mathbb{E}_{i,t}[\cdot]$ is the conditional subjective expectation of investor i. Crucially, and in contrast to the original model in Anderson and Danthine (1980, 1981), I allow investor-specific subjective beliefs about the first moment. Otherwise, the model follows the original setup and assumes homogeneous beliefs about second moments and a common risk aversion parameter across investors.[9]

The first-order condition of the maximization problem in Eq. (3) for the optimal currency

[9] In heterogeneous-beliefs models, it is standard to assume that agents can deduce volatility but must estimate the conditional mean (e.g., Basak (2005)) as mean returns are much harder to estimate than volatility (Merton, 1980), such that disagreement centers on the first moment. The baseline specification of Egan et al. (2022) likewise imposes common risk aversion and perceptions about risk. They also show that extending the model by allowing for heterogeneous risk aversion, or equivalently, for heterogeneous perceptions about second moments, yields qualitatively similar results.

demand follows as

$$w_{i,t}^{*} = \underbrace{\frac{\mathbb{E}_{i,t}[\Delta s_{t+1}] - (f_t - s_t)}{\gamma\,\mathbb{V}\mathrm{ar}_t[\Delta s_{t+1}]}}_{\text{Speculative term}} - \underbrace{\frac{\mathbb{C}\mathrm{ov}_t[r_{i,t+1}, \Delta s_{t+1}]}{\mathbb{V}\mathrm{ar}_t[\Delta s_{t+1}]}}_{\text{Benchmark hedge}} \tag{4}$$

where $w_{i,t}^{*}$ is between zero and one. With the speculative term equal to zero, full hedging ($w_{i,t}^{*} = 0$) is optimal when the correlation between exchange rates and local currency asset returns is zero, i.e., $\mathbb{C}\mathrm{ov}_t[r_{i,t+1}, \Delta s_{t+1}] = 0$, while zero hedging ($w_{i,t}^{*} = 1$) is optimal when exchange rates and local currency asset returns are perfectly negatively correlated and have the same volatility, i.e., $\mathbb{C}\mathrm{ov}_t[r_{i,t+1}, \Delta s_{t+1}] = -\,\mathbb{V}\mathrm{ar}[\Delta s_{t+1}]$.

Portfolio-implied Expectations. One of the paper's key contributions is to recover investors' exchange rate expectations from ETF holdings by inverting Eq. (4), using the observed portfolio weights, $w_{i,t}^{observed}$, as

$$\mathbb{E}_{i,t}^{PF}[rx_{t+1}] = \gamma(w_{i,t}^{observed}\,\mathbb{V}\mathrm{ar}_t[\Delta s_{t+1}] + \mathbb{C}\mathrm{ov}_t[r_{i,t+1}, \Delta s_{t+1}]). \tag{5}$$

The average portfolio-implied expectations are then equal to

$$\mathbb{E}_{t}^{PF}[rx_{t+1}] = \frac{1}{N_t}\sum_{i}^{N_t}\mathbb{E}_{i,t}^{PF}[rx_{t+1}] \tag{6}$$

and their dispersion, measured by the mean absolute deviation (MAD), is equal to

$$Disp_t^{PF} = \frac{1}{N_t}\sum_{i=1}^{N_t}\left|\mathbb{E}_{i,t}^{PF}[rx_{t+1}] - \mathbb{E}_{t}^{PF}[rx_{t+1}]\right| \tag{7}$$

where N_t equals the number of funds in each period.[10]

As a special case, under homogeneous beliefs, we have

$$\mathbb{E}_{t}^{PF}[rx_{t+1}] = \mathbb{E}_{i,t}^{PF}[rx_{t+1}]\;\forall\, i \quad \text{and} \quad Disp_t^{PF} = 0$$

where expectations are taken with respect to either the objective or a common subjective probability

[10] I follow Beber et al. (2010) and define the dispersion using the *mean* absolute deviation rather than the *median* absolute deviation. The results in section 7 are robust to the alternative definition.

measure. Controlling for exchange rate volatility and the benchmark hedging term, the first-order condition in Eq. (4) then maps portfolio weights to a common expectation about exchange rates. Three such candidates for exchange rate expectations are given in the next section.

3.2 Common Estimates of Exchange Rate Expectations

This section introduces the three most prominent estimates for conditional currency excess return expectations from the literature. Estimates are given per currency area, c.

Survey Expectations. Building on the growing literature on subjective beliefs, I construct expected currency returns and excess returns from survey data, denoted as $\mathbb{E}_t^{survey}[rx_{c,t+1}]$. I compute the return using the estimate of the expected future (one-month ahead) spot rate, $Spot_{c,t+1}^{survey}$. Formally,

$$\mathbb{E}_t^{survey}[rx_{c,t+1}] = log(Spot_{c,t+1}^{survey}) - s_{c,t} - (f_{c,t} - s_{c,t}). \tag{8}$$

Macro Expectations. I construct conditional expected currency returns as in Chernov et al. (2023). They draw on prominent economic theories to derive conditional expected currency excess returns, which are used in the construction of the unconditional mean-variance efficient portfolio. I refer to these expected currency returns as macro expectations, $\mathbb{E}_t^{macro}[rx_{c,t+1}]$. The macro expectations consist of three components. The first component is anchored in the random-walk hypothesis (Meese and Rogoff, 1983) that implies that expected currency excess returns are predictable using the forward discount: $s_{c,t} - f_{c,t}$, akin to the carry factor. The second component is based on the purchasing power parity, which implies that the real exchange rate (RER) is mean reverting in the long run and can thus be used to forecast movements in the nominal exchange rate (Chernov and Creal, 2021; Dahlquist and Pénasse, 2022), akin to the value factor in the asset pricing literature. The mean-reversion component, z_{c,Q_t}, is built as the relative level of the current real exchange rate compared to its centered moving average from five years ago. The last component, the trend signal, builds on evidence from the forecasting literature that past returns predict future returns, echoing the momentum factor. Combining all three components, I compute macro expectations as

$$\mathbb{E}_t^{macro}[rx_{c,t+1}] = \gamma_{c,t} \underbrace{(s_{c,t} - f_{c,t})}_{\text{Forward Discount}} + \delta_{c,t} \underbrace{z_{c,Q_t}}_{\text{RER}} + \phi_{c,t} \underbrace{(s_{c,t} - s_{c,t-12})}_{\text{Trend}}. \tag{9}$$

The regression is run month-by-month in an expanding window fashion. Following Chernov et al. (2023) I set $\gamma_{c,t} = 1$ to align with the random walk baseline, noting that this choice has little impact on the results compared to leaving $\gamma_{c,t}$ unrestricted. Details of the estimation technique are provided in Appendix B.1, with coefficient estimates reported in Figure B.1.

Factor Expectations. I construct conditional expected currency returns based on the two most popular currency pricing factors, the dollar and carry factors (Lustig et al., 2011) and compute expected excess returns as in Opie and Riddiough (2020). The latter show that a hedging strategy that applies the mean-variance optimal hedging framework of the previous section with expected currency returns based on the forecastable component of global factor returns can generate superior investment performance. The key idea is that expected currency returns are determined using the linear beta pricing model

$$\mathbb{E}_t^{factor}[rx_{c,t+1}] = \lambda_t^{carry} \beta_{c,t}^{carry} + \lambda_t^{dollar} \beta_{c,t}^{dollar} \tag{10}$$

where λ_t^{carry} and λ_t^{dollar} are replaced with their expected components, $\mathbb{E}_t[\lambda_{t+1}^{carry}]$ and $\mathbb{E}_t[\lambda_{t+1}^{dollar}]$. These expected factor returns are the forecastable components of the risk premia, predicted using other variables that have been shown to successfully forecast factor returns. More precisely, in a first step, the factor loadings, $\beta_{c,t}$, are estimated by running a rolling OLS regression of excess returns on the dollar and carry factors using past data up to t. The dollar factor involves taking a long position in foreign currencies and a short position in the dollar, while the carry factor takes a long position in currencies with relatively high forward discounts and a short position in currencies with relatively low forward discounts. In a second step, the average forward discount, the foreign exchange volatility and the commodity index return are used to predict one-month ahead expected factor returns $\mathbb{E}_t[\lambda_{t+1}^{carry}]$ and $\mathbb{E}_t[\lambda_{t+1}^{dollar}]$. Combining the factor loadings and the expected factor returns yields what I refer to as factor expectations. Details of the estimation technique are provided in Appendix B.2, with coefficient estimates reported in Figure B.2.

For all estimates, expected exchange rate returns, $\mathbb{E}_t[\Delta s_{c,t+1}]$, are then derived from expected excess returns by adding the forward premium (see Eq. (2)).

3.3 Estimation of Currency and Asset Risk

The mean-variance framework for optimal currency demand requires estimates for conditional variances and covariances between currency and asset returns. I estimate conditional covariance matrices, Ω_t, from daily exchange rate changes and asset returns. For each fund i and currency c, I estimate the covariance matrix

$$\Sigma_{i,c,t} = \begin{bmatrix} \mathbb{V}\text{ar}_t[\Delta s_{c,t+1}] & \mathbb{C}\text{ov}_t[\Delta s_{c,t+1}, r_{i,c,t+1}] \\ \mathbb{C}\text{ov}_t[\Delta s_{c,t+1}, r_{i,c,t+1}] & \mathbb{V}\text{ar}_t[r_{i,c,t+1}] \end{bmatrix} \tag{11}$$

where the asset return $r_{i,c,t+1}$ corresponds to daily realized changes in net asset values of the unhedged ETF i, domiciled in currency area c, and measured in local currency. The covariance is estimated over the preceding 25 trading days to arrive at monthly covariances for each day. I follow Chernov et al. (2023) and compute the exponentially weighted average of the matrix using a decay factor of $\lambda = 0.94$. The conditional covariance matrix for month $t+1$ is then equal to

$$\Omega_{i,c,t} = (1-\lambda)\Sigma_{i,c,t} + \lambda\Omega_{i,c,t-1}. \tag{12}$$

I check whether the conditional variance is plausible and, following Chernov et al. (2023), regress the variance of realized currency excess returns on the estimated conditional variance. To get from the estimated conditional variance of spot depreciation rates, $\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]$, to the conditional variance of excess returns, $\mathbb{V}\text{ar}_t[rx_{c,t+1}]$, I add the squared forward discount. The variance of realized currency excess returns is constructed using the three different estimates of exchange rate expectations introduced in the previous section. Under the null hypothesis that the correct conditional means and variances are identified, the coefficient should equal unity. Results are reported in Table A.1. For the model-implied expectations (macro and factor) the coefficient is relatively close to one, with macro expectations yielding a coefficient of 1.3. In contrast, the coefficient for survey expectations is ten times larger, indicating that its estimated variance substantially underestimates the realized squared errors.

4 Data

4.1 Exchange-Traded Funds

I retrieve the universe of currently active and inactive ETFs from LSEG (formerly Refinitiv) Workspace's Advanced Search. I restrict the sample to ETFs that invest in US equity or debt and are domiciled in one of six currency areas: Australia, Canada, the euro area, Japan, Switzerland, the United Kingdom. These ETFs make up the largest share of international currency-hedged ETFs, i.e., ETFs investing abroad for which there exists a hedged counterpart to the unhedged ETFs. I keep only funds that offer both a currency-hedged and an unhedged share class. For example, the sample includes a Canada-domiciled ETF investing in the United States, such as the *BMO S&P 500 Hedged to CAD Index ETF*, along with its corresponding unhedged counterpart, the *BMO S&P 500 Index ETF*. For funds investing in a global index, I allocate assets under management (AUM) across currencies proportionally to the fund's stated currency allocation and retain only observations with a USD share above 50%; smaller USD allocations are dropped. I exclude observations with AUM below USD 1 million.

My final dataset includes 350 currency-hedged ETFs and their 350 unhedged counterparts over 11 years, from January 2014 to June 2025. While ETFs have existed since the early 2000s, currency-hedged funds reached meaningful AUM only from 2014 onward. I rely on Bloomberg for the underlying data of the ETFs, since the data are more comprehensive and complete. Panel A of Figure 1 plots the total AUMs in billions of USD for hedged and unhedged ETFs in blue and red, respectively, from 2014 to 2025. By the end of 2025, unhedged funds hold nearly five times the assets of currency-hedged funds, with totals reaching USD 600 billion for unhedged funds and USD 120 billion for currency-hedged funds. These aggregate holdings accounting for roughly 5% and 1% of total foreign holdings in USD-denominated equities and bonds, respectively (Du and Huber, 2024).

For each fund-pair i domiciled in currency area c at time t, I calculate its weight in the unhedged index as

$$w_{i,c,t} = \frac{A_{i,c,t}^{unhedged}}{A_{i,c,t}^{unhedged} + A_{i,c,t}^{hedged}}. \tag{13}$$

where $A_{i,c,t}^{unhedged}$ and $A_{i,c,t}^{hedged}$ are the assets under management for the unhedged and hedged ETF in domestic currency, respectively.

Table 1, Panel A, reports summary statistics for the sample of funds, including assets under management in USD, the weight in the unhedged index and the return of the hedged and unhedged ETF in domestic currency. Note that the unhedged ETF earns on average a higher return but entails higher risk compared to the hedged ETF.

Additionally, I follow Curcuru et al. (2011) and define active flows into the unhedged ETF, $flow_{i,c,t}$, as the active change in the weight in USD as

$$flow_{i,c,t} = w_{i,c,t} - w_{i,c,t-1} \frac{1 + r_{i,c,t}^{unhedged}}{1 + r_{i,c,t}^{p}} \tag{14}$$

where $r_{i,c,t}^{unhedged}$ is the return of the unhedged ETF per currency and $r_{i,c,t}^{p}$ is the portfolio return of the unhedged and hedged ETF using weights $w_{i,c,t-1}$. As is standard practice in mutual fund flow analysis, I winsorize the flow distribution at the 1st and 99th percentiles. Furthermore, I define aggregate flows per currency area as

$$flow_{c,t} = w_{c,t}^{agg} - w_{c,t-1}^{agg} \frac{1 + r_{c,t}^{unhedged}}{1 + r_{c,t}^{p}} \tag{15}$$

where returns are (AUM-weighted) currency averages across funds and aggregate weights, $w_{c,t}^{agg}$, are computed using currency-level aggregates of assets under management. Panel B of Figure 1 plots the average aggregate active flows over time.

Lastly, I note a potential bias from cross-border ownership in my currency-demand measure, which assumes that ETF shares are held only by investors in the fund's domicile. In principle, investors outside the ETF's domicile could also invest in the ETF, and if their home currency matches the fund's base currency, they may have little incentive to hedge. However, in practice, this bias should be limited, as retail investors are generally legally constrained from buying ETFs outside their home market (e.g., euro area investors cannot access US-listed ETFs) and institutional investors, who are generally subject to fewer such restrictions, hold only about one-quarter of the shares in my sample of funds. Any resulting bias should be modest.

4.1.1 Institutional Details of ETFs

Since this study analyzes holdings from ETFs, I provide a brief, self-contained overview of key ETF mechanics for context. I summarize the main concepts here and refer to Hill et al. (2015) for details.

An exchange-traded fund is a pooled investment vehicle with shares that are traded throughout the day at stock exchanges in the secondary market at a price, P. This price stays close to the market value of the underlying securities, the net asset value, NAV, due to arbitrage by authorized participants (AP). APs are typically large financial institutions that have agreements with the ETF issuer. These agreements allow them to create and redeem ETF shares at the NAV at the end of each day in the primary market to arbitrage any difference between the NAV and the price of an ETF. The arbitrage process, when the ETF trades at a premium ($P > NAV$), is depicted in Figure A.2. The basic mechanism works as follows: (1) an AP buys the underlying securities of the ETF, (2) sells it to the ETF sponsor at the NAV in the primary market, receives the ETF shares, and (3) sells the ETF shares to the ETF investors in the secondary market at price P.

More precisely, the ETF sponsor publishes a daily creation basket, which details the specific securities and quantities that reflect the ETF's target portfolio. The AP delivers this basket in-kind to the sponsor in order to receive a "creation unit" of ETF shares. Typically, one creation unit consists of a number of ETF shares that range between 25,000 to 200,000 shares. When the underlying securities are difficult and illiquid, the ETF may allow the AP, instead of in-kind delivery, to substitute cash or derivatives (total return swaps) for some or all of the assets in the creation basket. The value of the creation basket and any cash adjustment equal the value of the creation unit based on the ETF's NAV at the end of the day on which the transaction was initiated. In practice, the ETF sponsor also charges a fee to the AP, which is higher when in-kind delivery does not occur, to compensate for the transaction costs incurred when the sponsor has to purchase the underlying securities directly. Moreover, the AP does not need to sell all the shares, it can also hold some on its balance sheet. Per share sold in the secondary market, the AP makes a profit of $P - NAV$. The increased supply of ETF shares puts downward pressure on the price of the ETF, so that the ETF trades again at a price close to the NAV. If the ETF trades at a discount to its NAV, the reverse arbitrage mechanism functions: the AP redeems ETF shares in the primary market in exchange for the underlying securities, which can then be sold in the asset market.

Importantly, the ETF fund flows represent newly created or redeemed shares in the primary market. Unlike mutual funds, these flows do not represent end investors trading with the asset market. Instead, they represent APs transacting with the ETF sponsor to create or redeem ETF shares. However, as highlighted by Converse et al. (2023), when the ETF arbitrage mechanism functions effectively, the creation and redemption of shares directly mirror the excess demand of end investors and can be interpreted as equivalent to investor dollar flows into or out of ETFs. If ETF arbitrage works well, the difference between NAV and the price (the mispricing) is small.

Mispricing is generally higher for my sample of international ETFs compared to a sample of domestic ETFs. This is due to the different trading hours between the ETF and the underlying assets. For example, an ETF trading on the Frankfurt Stock Exchange, that invests in US equity, trades at Frankfurt trading hours, while its underyling assets trade at US trading hours. The NAV of the ETF is thus stale for around 6.5 hours during Frankfurt trading hours, while the underlying assets of the ETF are still traded for 4.5 more hours when the Frankfurt exchange is already closed. This makes the arbitrage mechanism by APs more complicated and less efficient. In Figure A.3 I plot the average relative percentage mispricing over time for hedged (red) and unhedged (blue) ETFs, respectively. The relative mispricing is defined as the difference between the last trading price of the ETF and the NAV divided by the NAV, i.e., $\epsilon = \frac{P_t - NAV_t}{NAV_t} \times 100$. The average mispricing is particularly pronounced during periods of market stress. For instance, at the onset of the COVID-19 crisis in March 2020, mispricing spikes to 3%. On average, hedged ETFs exhibit slightly higher mispricing than unhedged ETFs, with mean values of 0.036% and 0.019%, respectively. These figures are similar to those reported by Petajisto (2017), who document an average mispricing of 0.06% for all US-listed ETFs and 1 to 2% for the sample of US-listed international ETFs over the period 2007-2014.

In the context of my study, it is important to verify that differences in flows between hedged and unhedged ETFs coming from AP activity are driven by investor demand rather than by supply-side constraints of APs. In other words, the willingness of APs to arbitrage the hedged ETF versus the unhedged ETF must be the same. The only difference between the constituents of the creation basket of a hedged and unhedged ETF is the forward contract, which the AP delivers in cash. Thus, the liquidity of the hedged and unhedged creation basket is essentially the same. Liquidity constraints of APs typically arise in highly volatile times. Thus, especially for these times, it is necessary to ensure that the mispricing of hedged and unhedged ETFs does not differ systematically. To check how mispricing between hedged and unhedged ETFs differs, in Table A.2 I report the results of regressing the hedged ETF premium/discount against the premium/discount of the unhedged ETF, where a premium is defined as $\epsilon > 0$ and a discount as $\epsilon < 0$. Column (1) and (4) show a clear positive relationship between the mispricing of the two ETFs. Importantly, this correlation increases when the VIX increases, as documented in Column (3) and (6) where I add the VIX (as a proxy for liquidity constraints as in Pan and Zeng (2017)) and the interaction term. Thus, when constraints of APs are more likely to bind, the mispricing of hedged and unhedged ETFs comove even more strongly. In less constrained times, the comovement is weaker and potentially driven by other demand-side factors. This provides evidence that investors' demand is transmitted to an AP's activity similarly for both types of ETFs, suggesting that the currency hedge of an ETF does not matter for an AP's willingness to arbitrage mispricing and that flows can be interpreted as investor demand.

4.2 Exchange Rate Prices and Expectations

Exchange rates for six currencies EUR, CAD, GBP, CHF, AUD, JPY against the USD are sourced from Bloomberg. I use spot and one-month forward rates and compute excess returns as defined in Eq. (2). I also obtain one-month option-implied FX volatility. For the monthly sample, I use end-of-month observations.

Survey expectations for one-month future spot rates are obtained from Consensus Economics. The surveys are typically collected on the second Wednesday of each month. Accordingly, I align the monthly sample to each survey's collection date. Data are available through August 2024.

Portfolio-implied expectations are computed according to Eq. (5), using conditional variance-covariance estimates and setting the risk aversion parameter to three as in other studies (e.g., Opie and Riddiough (2020)). Additionally, it is important to account for the fact that changes in expectations can arise solely from valuation effects. Specifically, portfolio-implied expectations based on the weight may shift over time due to exchange rate fluctuations, even in the absence of any inflows or outflows. To address this, I compute a flow-only version of the weight using active flows into the unhedged ETF share class:

$$w_{i,c,t}^{flow} = w_{i,c,t-1} + flow_{i,c,t} \tag{16}$$

which I then use in the computation of $\mathbb{E}_{i,t}^{PF}[rx_{c,t+1}]$.[11]

In Table 1, Panel D, I report the average expected currency excess return for all four estimates. Notably, survey and factor expectations are much larger in scale. For example, portfolio-implied and macro expectations average +3 bp and −3 bp, respectively, with an s.d. of roughly 12 bp, while survey expectations average 19 bp with an s.d. of 162 bp. However, note that portfolio-implied expectations scale with the risk aversion parameter, and a higher value (e.g., $\gamma = 30$) would bring their magnitude closer to survey estimates. Figure 2 plots the four different estimates of expected currency excess returns over time averaged across all currency pairs in bps, dividing survey and factor expectations by 10 for comparability. The four estimates of currency expectations feature some similar trends. For example, all exchange rate expectations drop in March 2020 during the onset of COVID-19. Quantitatively, in March 2020, the average ETF investor expected an annualized currency risk premium on the USD basket of about −4% (= 12 × 33 bp).

5 Observed and Optimal Mean-Variance Currency Demand

While expectations derived from surveys or constructed from macro or factor models are widely used in the literature, it is not clear whether these benchmarks correspond to investors' revealed currency demand. The theory of optimal currency demand provides a framework to test this link, as it specifies how expected returns and second moments map into observed portfolio holdings.

Table 2 presents results of regressing monthly USD holdings, $w_{i,c,t}$, on the optimal currency demand components: the expected exchange rate change, $\mathbb{E}_t[\Delta s_{c,t+1}]$, the forward premium, $f_{c,t} - s_{c,t}$, and the covariance of real asset returns with exchange rates, $\mathbb{C}\mathrm{ov}_t[r_{i,c,t+1}, \Delta s_{c,t+1}]$. All terms are scaled by the exchange rate volatility, $\mathbb{V}\mathrm{ar}_t[\Delta s_{c,t+1}]$, with corresponding summary statistics reported in Panel B of Table 1. For readability, the speculative components are divided by 100. I estimate the model separately for the three expectation measures (survey, macro, and factor

[11] Theoretically, the difference between w_t^{flow} and w_t is important, however, empirically, this difference is very small and amounts, on average, to only 1.14 basis points (bps). Also, note that the value of the risk aversion parameter is merely a scaling factor in Eq. (5) and does not change the results qualitatively.

expectations) and run the regressions separately for equity funds (Panel A) and debt funds (panel B), as hedging practices are documented to differ considerably (Bräuer and Hau, 2024; Cheema-Fox and Greenwood, 2024). Columns (1)–(3) present the baseline specification without fixed effects.[12]

The baseline model matches the data extremely well for equity and debt funds with an R^2 ranging between 51% and 56%. All components of optimal currency hedging yield the sign predicted by theory and all estimates of expected currency returns correlate highly significantly with currency demand in the baseline model. Specifically, a one standard deviation increase in (variance-adjusted) expected spot returns implied by surveys is associated with a 0.21 ($= 0.18 \times \frac{0.33}{0.28}$) standard deviation, or 6.5 percentage points, increase in investors' dollar share. The response to macro and factor expectations is slightly stronger and amounts to a 0.38 ($= 5.25 \times \frac{0.02}{0.28}$) and 0.28 ($= 0.25 \times \frac{0.31}{0.28}$) standard deviation increase in investors' dollar share, respectively. The estimated sensitivity of portfolio shares to beliefs is comparable to investor-level evidence in the equity market: Giglio et al. (2021) find that a one standard deviation increase in expected one-year stock returns increases the equity share by 0.16 standard deviations.[13]

Consistent with theory, the forward premium enters the regression with a significantly negative coefficient for both asset classes. The size of the effect is also substantial in economic terms. Specifically, a one standard deviation increase in the forward premium comes with a 1.3 ($= 10.6 \times \frac{0.034}{0.28}$) standard deviation decrease in investors' dollar share. Higher hedging costs (the negative of the forward premium) thus decrease USD-hedging and increase USD holdings substantially.

According to theory, the benchmark coefficient should be equal to -1. Empirically, this holds more closely for debt funds than for equity funds: the estimated coefficient ranges from -0.31 for equity funds to -0.46 for debt funds. This is in line with Bräuer and Hau (2024) who also find a smaller coefficient for the benchmark hedge for equity funds compared to fixed-income funds. With zero speculation on future currency returns, the benchmark coefficient between -0.31 and -0.46 implies an average weight in the unhedged index between 15% ($= -0.47 \times -0.31$) and 22%

[12] Survey data are only available through August 2024. However, the results for the macro and factor expectations presented in Table 2 remain robust when the samples are truncated at the same endpoint.

[13] Compared to Giglio et al. (2021), my main specification regresses the USD share on *variance-adjusted* expected exchange rate changes, which limits one-to-one comparability. In Appendix A.3 I replicate their specification in Eq. (1) and regress the USD share on expected spot rate changes only. Among the high statistically significant coefficients, the estimated sensitivity of USD allocations to beliefs shows that a one standard deviation increase in expected returns is associated with a 0.08 to 0.27 standard deviation increase in the USD share. Thus, using their regression specifications, the estimated sensitivity also closely matches the baseline specification of Table 2 and falls within the range reported by Giglio et al. (2021).

($= -0.47 \times -0.46$). Since the average observed weight is much higher, at around 68%, most of the unhedged exposure cannot be explained by the hedging component alone.

Columns (4)–(6) add time fixed effects, yielding a key result: only survey-based expectations remain statistically significant for both investor groups; macro- and factor-based measures do not. For equity investors, the forward premium remains highly significant. The result also shows that across funds the covariance term is irrelevant to both investor groups.

Columns (7) to (9) examine the relationship between observed and optimal currency demand within individual funds by adding fund fixed effects. The results show that equity investors respond only to changes in fund-specific covariances and do not speculate on expected currency returns. In contrast to theoretical predictions, the forward premium enters with a positive coefficient, but the effect is weak and only marginally significant at the 10% level. For debt fund investors, none of the components of optimal currency exposure matters to their allocation between the unhedged and hedged share classes within a given fund.

In the Appendix, I show results for two alternative specifications of the regression. First, I combine the expected spot change with the forward premium into the excess return and re-estimate the regression in order to obtain an estimate for the risk aversion parameter. Specifically, in the model of optimal currency demand, the estimated coefficient of the expected currency excess return equals the inverse of the risk aversion parameter. Recall also that the speculative components are divided by 100. Table A.4 shows that macro expectations imply a risk-aversion parameter of ≈ 13.8 for equity investors and ≈ 15.6 for debt fund investors. These estimates are relatively close to estimates of the risk aversion parameter documented in the literature, which typically range from 3 to 10. In contrast, survey expectations imply a much higher risk aversion parameter of more than 300. This finding aligns with Giglio et al. (2021) who also find a higher risk aversion parameter using survey data of equity investors (of 36), compared to usual estimates. Excess returns based on factor expectations yield a similarly high risk aversion parameter, but are not significant in explaining USD allocations at conventional significance levels or negative in some specifications.

Second, in Table A.5, I run the regression at the currency level by aggregating the weights (AUMs) per currency. That is, instead of applying the theory at the fund-level, the analysis assumes one representative agent per currency area that solves for the optimal currency demand. The results are very similar to the fund-level regression for the baseline model, with some interesting differences. In the baseline model of the currency-level regression, only survey expectations yield a significant coefficient, while the other two expectation estimates are only weakly significant. However, macro expectations turn significant for equity funds once controlling for average time effects in Column (5). At the currency level, aggregate investors demand more dollars in currency areas with a positive (less negative) covariance in Columns (4) to (6), which is in contrast to the fund-level regression and the theory. A possible explanation is measurement error in the covariance matrix at the aggregate level. An aggregate investor potentially holds the full portfolio of currencies. Consequently, portfolio risk depends on the entire FX covariance matrix, including cross-currency correlations, rather than the covariance of a single currency with the asset return. Consistent with this, Bräuer and Hau (2024) demonstrate that optimal hedging based on the full FX covariance structure yields substantially lower risk than hedging strategies that consider each currency in isolation.

In Figure 3, I visualize the cross-sectional relationship between USD holdings and expectations at the currency level. It plots the average weight in unhedged ETFs on the vertical axis against the three estimates of the (variance-adjusted) expected excess returns (Panel A-C) per currency for equity funds (left column) and debt funds (right column), respectively. All variables are scaled by the inverse of the exchange rate variance and the benchmark hedge is subtracted as in Eq. (4). Panel A confirms earlier regression evidence that survey-based exchange rate expectations are strongly related to the weight in unhedged ETFs. The correlation between the two variables across funds and time is positive and significant and amounts to 97% for equity and 98% bond funds. For example, while Swiss investors held their investment in the US almost completely unhedged, Japanese investors hedged more than half of their US investment, which aligns with higher survey expectations for currency excess returns on the USD against the CHF compared to the JPY. The correlation between return expectations and currency demand is lower for macro expectations with 60% or 76% for equity and debt funds, respectively, while factor expectations are only weakly correlated and negative, around -8%.[14]

I highlight that, absent speculation, the average observed weights per currency area cannot be reconciled with the theory, which predicts that those should match the benchmark hedging

[14] Note that the figure is virtually identical when using only expected exchange rate changes on the horizontal axis, i.e., without variance scaling and without subtracting the forward premium and the benchmark hedging term.

term. This discrepancy is documented in Table A.6, which reports the *theory-implied* optimal benchmark weights, w_f^*, and hedge ratios, HR^*, computed from the estimated covariance and variance matrices, alongside the *observed* weights and hedge ratios. On average, the absolute difference between observed and theory-implied benchmark weights is about 60%. The largest gap arises for Japanese equity ETF investors, who hedge roughly have of their currency exposure, whereas benchmark hedging implies a hedge ratio greater than one.

Taken together, the analysis highlights the following findings: Expected currency returns are only relevant in the cross-section of funds and currency areas, with survey expectations best capturing currency (USD) holdings. Time-variation in the estimates of currency expectations obtained from surveys or proposed by the literature does not determine investors' USD demand. Hedging costs, captured by the forward premium, play an especially important role for equity investors. Equity investors also tilt their currency positions with the FX-asset covariance, however, the direction of its effect is ambiguous and depends on the regression specification (fund versus currency-level). Abstracting from heterogeneity across funds, currencies, and time, the model of optimal currency demand explains in the baseline model the overall patterns in observed currency demand remarkably well.

5.1 Heterogeneity among ETFs

This section explores whether the currency demand of investors depends on specific ETF characteristics. Specifically, I perform the previous baseline regression for each fund separately and sort the coefficients on the expected currency return, the forward premium, and the benchmark hedge in quartiles of fund characteristics. In this section, I focus exclusively on survey expectations, $\mathbb{E}_t^{survey}[\Delta s_{c,t+1}]$, since the previous section found this estimate to best explain currency holdings.

I sort the ETFs based on four characteristics: size, relative expense ratio, institutional investor share, and hedging strategy. Summary statistics for the four characteristics are reported in Panel C of Table 1. The first characteristic is the log size, which is the log of the sum of the hedged and unhedged assets under management. The second characteristic is the relative or net expense ratio that captures potential cost considerations when investors switch between the two ETFs. It is the difference between the expense ratio of the unhedged ETF and the hedged ETF, so that a lower ratio indicates a relatively costlier hedged product. On average, the hedged share class is 2.6 bps more expensive than the unhedged share class, reflecting hedging costs associated with forward trading. The upper quartile of funds has a zero net expense ratio, with the hedged share class as expensive as the unhedged share class. The third characteristic is the institutional investor share, which captures the average (across time and share class) proportion of shares outstanding held by institutional investors for a given ETF-pair. Institutional investors typically own about one-quarter of the shares outstanding in most funds, with the average institutional share being 23% and the median at 17%. The last characteristic is the hedging strategy, which measures how much a fund's hedging strategy deviates from a naive hedging strategy that rolls over one-month forward contracts at the end of each month. Specifically, I take for each fund the one-month return difference between the unhedged and hedged share class and subtract the one-month realized excess return,

$$\Delta_{i,c,t}^{strategy} = \underbrace{r_{i,c,t}^{unhedged} - r_{i,c,t}^{hedged}}_{\text{return difference}} - \underbrace{(s_{c,t} - f_{c,t-1})}_{\text{realized excess return}} \quad . \tag{17}$$

If a fund can trade the forward and spot rate at the mid rate and rolls over its forward contract at the end of the month, the deviation is equal to zero, $\Delta_{i,c,t}^{strategy} = 0$. Conversely, the difference is non-zero if the fund incurs transaction costs or actively optimizes the hedge, i.e., deviates from a naive end-of-month roll-over strategy, and strategically times its forward transaction. A positive (negative) difference indicates that a fund outperforms (underperforms) a naive hedging strategy with zero transaction costs. Table 1 shows that the average deviation is positive, $\Delta_{i,c,t}^{strategy} > 0$, and amounts to 4.4 bps. Thus, on average funds' strategic timing more than compensates for transaction costs, resulting in a net gain relative to a naive hedge. For reference, the average quoted bid-ask spread on a one-month outright EUR/USD forward contract is about 3.1 bps over the sample period 2014-2024.

Figure 4 presents results of the fund-level regression with the average of the estimated coefficient for the expected currency return in Column (1), the forward premium in Column (2) and the benchmark hedge in Column (3) per quartile. Column (1) shows that survey expectations are positively correlated with currency demand for larger funds (quartile 3 and 4), while they are negatively correlated for very small funds (quartile 1). Thus, survey expectations are most clearly reflected in the currency holdings of large funds, with smaller funds betting against the signal. Funds with a relatively cheap hedged product (with a net expense ratio > 0) or a relatively expensive hedged product (with a net expense ratio < 2 bp) take positions most consistent with survey expectations. Conversely, funds with an average expensive hedged product do not significantly load on survey expectations. Furthermore, speculation based on survey expectations tends to be concentrated in funds with a relatively low share of institutional investors (below 34%) and with medium hedging strategy profits (quartiles 2 and 3 with $-1.8 < \Delta_{i,c,t}^{strategy} < 5.5$).

Column (2) shows that the forward premium coefficient is stable and negative across characteristics: estimates are statistically indistinguishable across quartiles, with no evidence of heterogeneity. Thus, hedging costs play an important role for all types of investors and funds.

Column (3) presents the coefficient on the benchmark hedge. A coefficient that deviates further from -1 indicates that investors do not fully align their currency allocation with an optimal mean-variance strategy. Across characteristics, the coefficient averages about -0.4 and is slightly less negative (further away from the optimal hedge) for small funds, for funds with a low institutional ownership share, and funds with high hedging-strategy profits.

The heterogeneity analysis shows that benchmark hedging and forward-premium sensitivity are broadly similar across funds, whereas speculative behavior varies significantly across funds and drives the largest cross-sectional differences. This finding is also important for section 7 on dispersion in portfolio-implied expectations, as such dispersions are likely to capture differences in speculative activity rather than in benchmark hedging or cost sensitivity.

6 Predictability of Exchange Rate Expectations

The finding that agents' aggregate portfolios correlate with common expectation estimates only cross-sectionally raises the question of what their actual time-varying expectations are and what properties these portfolio-implied expectations exhibit. In particular, in this section, I examine whether portfolio-implied expectations are more accurate than the other expectation estimates in predicting future returns.

To that end, I conduct a horse race across survey, macro, factor, and portfolio-implied expectations to assess their predictive content for future exchange rate movements. Formally, I test the forecasting power of each estimate with the following panel regression

$$rx_{c,t+1} = \alpha_t + \beta_1 \mathbb{E}_t[rx_{c,t+1}] + \beta_2(f_{c,t} - s_{c,t}) + \epsilon_{c,t} \tag{18}$$

where $rx_{c,t+1}$ are one-month realized currency excess returns, $\mathbb{E}_t[rx_{c,t+1}]$ is one of the four expectation estimates, $f_{c,t} - s_{c,t}$ is the one-month forward premium, and α_t are time fixed effects. I remove currency-specific levels of expected excess returns by subtracting a within-currency trailing mean that is based on data up to t to avoid look-ahead bias.

According to the uncovered interest rate parity, the coefficient for the forward premium, β_2, should equal zero. I also run regressions using the one-month realized exchange rate returns Δs_{t+1} as the dependent variable, in which β_2 should equal one.

Table 3 reports the results. In Column (1) and (5) I report results for the advocated portfolio-implied return expectation, $\mathbb{E}_t^{PF}[rx_{c,t+1}]$, in Column (2) and (6) I use monthly survey expectations, $\mathbb{E}_t^{survey}[rx_{c,t+1}]$, in Column (3) and (7) macro expectations and Column (4) and (8) factor expectations. Portfolio-implied expectations predict future monthly dollar excess returns with a highly significant coefficient of 1.12.[15] In contrast, neither survey expectations, nor macro expectations, nor factor expectations significantly predict future currency returns. In line with existing literature, the forward premium yields a non-zero coefficient in Column (1)-(4) and a coefficient that is not equal to one in Column (5)-(8), however, neither is statistically significant.

It is important to note that the predictive power of survey, macro and factor expectations increases when evaluated over a longer time period or at a different time horizon. In particular, Chernov et al. (2023) show that their return expectation forecasts returns in the period from 1985-2020. Similarly, Kremens et al. (2025) find that survey-based expectations of currency returns can successfully predict realized currency appreciations at the two-year horizon. Generally, a shared predictive power of survey-based and portfolio-implied expectations is consistent with a common priced risk channel: Kremens et al. (2025) show that surveys are to a large extent explained by the covariance between the equity market and the (real) exchange rate, which is also a key ingredient for the construction of portfolio-implied expectations of equity ETFs.

As a robustness check for the advocated portfolio-implied expectations, I alternatively construct

[15] In unreported regressions, I find that daily portfolio-implied expectations also forecast future one-day ahead currency appreciations with a t-statistics of 2.081.

expectations based on representative equity and bond index returns. More precisely, in Eq. (5) the ETF-specific covariance matrices are replaced by a covariance matrix from returns on a portfolio of the S&P aggregate sovereign bond index and the MSCI equity index, with an equity weight of 0.43 reflecting the fraction of equity ETF in the sample relative to debt ETFs. The results, not reported, show no statistically significant predictive power (t-statistic of 1.39) for these index-based expectations.

6.1 What Drives the Predictive Power of Portfolio-implied Expectations?

Portfolio-implied expectations consist of three components: the weight in the unhedged ETF, the variance term, and the covariance term. The weight in the unhedged ETF can further be decomposed into the active flows into the unhedged ETF (defined in Eq. (15)) and the prior period's weight. Formally, recall that

$$\mathbb{E}_t^{PF}[rx_{c,t+1}] = \gamma \underbrace{w_{c,t}^{flow}}_{= w_{c,t-1} + flow_{c,t}} \mathbb{V}\mathrm{ar}_t[\Delta s_{c,t+1}] + \gamma\, \mathbb{C}\mathrm{ov}_t[r_{c,t+1}, \Delta s_{c,t+1}].$$

This section tests the predictive content of each component of the portfolio-implied expectation to determine whether predictability is driven by a single component or by their combination.

Columns (1)-(3) of Table 4 show the results of regressing realized one-month currency excess returns on each of the components of portfolio-implied expectations separately. None of the components are significant on their own. Notably, however, once the weight is decomposed into its two components, Column (4) shows that monthly flows can significantly predict subsequent exchange rate excess returns. Specifically, a one standard deviation increase in active USD flows predicts a positive USD excess return of 13 bps ($= 0.09 \times 0.014$). This estimate is similar in magnitude to the findings of Engel and Wu (2023). They show that a one standard deviation decrease in the US term spread, the only significant predictor in the in-sample monthly regressions, predicts an 18.3 bps appreciation of the USD. Furthermore, a back-of-the-envelope calculation suggests that a one standard deviation of active flows corresponds to USD 38 million in 2024.[16] In contrast to flows, past weights have no predictive power for future returns. However, in the monthly sample, flows

[16] Average assets under management in 2024 of hedged and unhedged ETFs correspond to USD 2.7 billion, multiplied by one standard deviation of active flows this yields USD 38 million.

do not outperform the forecasting power of portfolio-implied expectations as shown in Column (5) of Table 4. When entered jointly, both variables are significant, but flows are only significant at the 10% level.

I further compare the predictive power of flows and portfolio-implied expectations across different forecasting horizons, using realized excess returns, $rx^k_{c,t+1}$, over $k \in \{1, 3, 12\}$ months. For example, for $k = 3$, the realized excess return is calculated as the difference between the (log) spot rate three months ahead and the (log) three-month forward rate observed at time t. To expand the sample size, I use daily data, constructing flow measures as one-month rolling averages of past flows. Standard errors are obtained via a moving-block bootstrap with block length set to 22, 66, and 264 for $k = 1, 3, 12$, respectively, to account for overlapping returns. Significance is based on bootstrap p-values. Table 5 presents the regression results for one-month returns in Panel A, three-month returns in Panel B, and twelve-month returns in Panel C.

For the one-month horizon (Panel A), the evidence echoes the finding of the monthly sample. Portfolio-implied expectations possess stronger predictive information than flows, but when both predictors are included, both remain statistically significant. At the three-month horizon (Panel B), overall predictability weakens, with portfolio-implied expectations being significant at the 10% level. In contrast, flows and portfolio-implied expectations both exhibit strong predictive power for twelve-month excess returns (Panel C). This pattern mirrors survey-based exchange rate expectations that also feature improved predictive power at longer horizons (Kremens et al., 2025).

In sum, flows provide a close proxy for portfolio-implied expectations. Nevertheless, the theory-guided portfolio-implied expectations contain incremental and stronger predictive information beyond flows.

In Appendix C I additionally compare several monthly trading strategies, including three built on the signal obtained from ETF portfolio-implied expectations or ETF active flows and traditional strategies such as the carry, dollar, or momentum strategy. Over the period 2014-2025, the strategies based on ETF signals deliver the highest Sharpe ratios (up to 0.58), whereas the traditional strategies reach at most 0.24. This pattern suggests that signals based on ETF allocations (hedged vs. unhedged) contain incremental information about currency risk premia beyond traditional predictors.

7 Dispersion in Exchange Rate Expectations

The unique feature of portfolio-implied expectations compared to most other estimates of exchange rate expectations is that they provide a cross-section of individual investor beliefs for each currency. This allows me to extract information about the dispersion of investor expectations, which may arise from differences in currency demand, $w_{i,c,t}$, and/or covariances, $\mathbb{C}\mathrm{ov}_{i,c,t}[r_{t+1}, \Delta s_{c,t+1}]$, (see Eq. (5), where both terms are fund-specific). Figure 5 visualizes the distribution of individual expectations, $\mathbb{E}_{i,t}^{PF}[rx_{c,t+1}]$ by currency, with each dot representing a fund-day observation. The figure documents substantial heterogeneity of exchange rate expectations across investors, which is in line with prior studies by MacDonald and Marsh (1996); Chionis and MacDonald (2002). The dispersion of expectations increases notably during periods of heightened market volatility, most prominently in March 2020 at the onset of the COVID-19 crisis and again in April 2025 following the announcement of US tariffs. Currency-specific spikes in dispersion are also evident around major shocks such as the third phase of the euro area sovereign debt crisis in July 2015, the Brexit referendum in June 2016, or the Swiss National Bank's decision to unpeg the Swiss franc from the euro in January 2015.

These patterns suggest a strong link between dispersion in beliefs and exchange rate volatility. Such a link can be rationalized by a simple consumption-based asset pricing model under heterogeneous beliefs. Specifically, in Appendix D I show that with heterogeneous beliefs, the standard stochastic discount factor (SDF) is augmented by a belief component that is a wealth-weighted average of how each investor tilts probabilities relative to a reference probability, e.g., the equally-weighted average belief.[17] Holding the reference probability measure fixed, greater dispersion among beliefs thus increases the volatility of the SDF.[18]

State-contingent equilibrium consumption shares can further amplify the volatility of the SDF as derived in Buraschi and Jiltsov (2006) and Buraschi et al. (2014). They show that when investors can trade state contingent risk-sharing contracts (options), the optimistic investor insures the pessimistic investor in bad states of the world, such that the pessimist's relative consumption

[17] The form of the SDF is similar to the one of Dumas et al. (2009, 2017) who develop a general equilibrium model with two groups of investors that hold different beliefs about the state of the economy. In these models heterogeneity arises because one group of investors overreacts to a public signal or because one group misinterprets the public signals of others. I abstract from the underlying reasons for differences in investors' beliefs.

[18] Empirically, I hold the reference probably fixed by computing the mean absolute deviation which centers individual expectations around the average belief per time t.

share is higher in those states. The optimist earns a premium on these contracts that increases with the degree of heterogeneity in beliefs. As option-implied volatility is derived from the market price of these contracts, it also increases with the degree of belief dispersion. Thus, the model also predicts a strong correlation between differences in beliefs and exchange rate volatility.

The link between exchange rate volatility and disagreement is particularly salient in Figure 6. The figure plots the mean absolute deviation of the portfolio-implied expectations, defined in Eq. (7), (solid blue line) against the option-implied one-month FX volatility (red dashed line). The correlation between these two series is remarkably high, ranging from 47% (GBP) to 63% (JPY), suggesting that higher dispersion in expectations is systematically associated with higher exchange rate volatility. In Appendix Figure A.4, I present the daily dispersion of portfolio-implied expectations, calculated across all funds and currency areas. The disagreement measure exhibits an even stronger correlation with average FX volatility amounting to 71%.

To formally evaluate the link, I run monthly panel regressions of option-implied FX volatility on the mean absolute deviation of portfolio-implied expectations. Formally,

$$IVol_{c,t} = \alpha_t + \alpha_c + \beta_1 Disp^{PF}_{c,t} + \beta_2 IVol_{c,t-1} + \epsilon_{c,t} \tag{19}$$

where $IVol_{c,t}$ is the one-month option-implied FX volatility and $Disp^{PF}_{c,t}$ denotes the mean absolute deviation of portfolio-implied expectations, and α_t and α_c represent time and currency fixed effects, respectively. I also add the one-month lag of implied volatility $IVol_{c,t-1}$ to account for its persistence. I use a log–log specification to stabilize variance and scale across currencies.[19] Table 6 reports the results. In Columns (1) and (2) I estimate the regression in levels, in Columns (3) and (4), I run the regression using log differences, in Columns (5) and (6) I run a predictive regression with lagged dispersion. In all specifications, I include both currency and time fixed effects.[20]

The results mirror the visual evidence from the previous figure: dispersion in portfolio-implied

[19] Since only one fund is available for AUD until 2015, the sample for this currency begins in 2015. For all other observations, the variables $Disp^{PF}_{c,t}$ and $IVol_{c,t}$ are strictly positive and taking logs is feasible. The results remain qualitatively unchanged when the regression is estimated in levels rather than logs. Note also that the results are robust when restricting the sample to observations with an above-average number of funds per currency–month, i.e., when $N_{c,t} > 27$.

[20] Note that the dynamic panel bias (Nickell, 1981) from including the lagged volatility with fixed effects that affects $\hat{\beta}_2$ is negligible. Using $T = 132$ implies a bias of about -0.013 on the persistence coefficient, raising $\hat{\beta}_2$ in Column (2) from 0.663 to 0.675 after correction. Alternatively, when I exclude time fixed effects in the first-difference specification (Columns 3–4), the results are qualitatively unchanged: coefficient magnitudes shift only modestly, signs are preserved, and statistical significance is unchanged.

expectations is strongly and positively associated with expected future exchange rate volatility, even after controlling for lagged volatility and time and currency specific variations. The point estimate in Column (1) implies that a one standard deviation increase of dispersion raises the exchange rate volatility by roughly 0.122 points ($= 0.135 \times 0.9$), which is about half its standard deviation, indicating an economically sizable effect. Approximately 27% of the exchange rate volatility is explained by the dispersion of expectations alone.[21] Column (2) shows that the dispersion coefficient remains positive and significant at the 1% level when including lagged volatility. Estimating the specification in first differences (Columns (3) and (4)) still yields a statistically significant coefficient for dispersion, though the t-statistics decline to 2.75 and 3.49, respectively.

The explanatory power of differences in beliefs from portfolio-implied expectations for option-implied FX volatility is comparable to that of survey-based measures. In particular, Beber et al. (2010) use mean absolute deviations in survey forecasts to measure differences in beliefs from May 1993 to December 2006 for the currency pairs JPY, EUR, and GBP against USD. In their most parsimonious regression (Table 5), differences in survey expectations explain on average 45% of the variance of the one-month option-implied volatility. In my sample, time-series regressions for the three currency pairs yield similar R^2 values, averaging 40%, with the GBP showing the lowest explanatory power as in Beber et al. (2010).

Columns (5)–(6) show that dispersion also predicts future implied volatility. On its own, dispersion strongly predicts future volatility (t-statistic of 4.3) and explains 19% of future volatility. Adding lagged volatility preserves a positive, significant coefficient, though the t-statistic falls to 2.42.

Overall, heterogeneity in beliefs inferred from portfolio holdings explains a significant share of exchange rate volatility. The explanatory power of 27% is comparable to that of order flow variables or macroeconomic variables: For order flows, Evans and Lyons (2002) document an R^2 between 40 and 67%, while for macro variables Stavrakeva and Tang (2025) document an R^2 between 23% and 37%, and Koijen and Yogo (2020) attribute roughly 26% of exchange rate variations to macro variables.

[21] Following Beber et al. (2010), the main regression employs implied volatility. Replacing FX option-implied volatility with realized monthly volatility increases the R^2 in Column (1) to 37%. Note that implied volatility correlates with realized volatility at 80%.

8 Conclusion

Survey-based exchange rate expectations are informative about aggregate investors' portfolio decisions and remain a defensible input for model calibration and asset pricing tests. Using investors' relative allocations to unhedged and hedged ETFs, I find that investors' USD demand is strongly linked to survey-based expectations in the cross-section of currencies, unlike expectations derived from macroeconomic models or asset pricing factors. However, time-series variation in currency demand shows little sensitivity to common estimates of exchange rate expectations, leaving open what investors' time-varying expectations actually are.

I propose a simple and tractable method to recover time-varying exchange rate expectations directly from investors' currency positions. The approach yields a daily indicator of average expectations and disagreement about future exchange rates. Empirically, these portfolio-implied expectations are more accurate in predicting returns than other expectations estimates. They also capture cross-sectional dispersion in expectations, which comoves strongly with FX volatility, suggesting that disagreement across agents contributes to exchange rate fluctuations.

Applications of portfolio-implied expectations are straightforward. For policymakers, portfolio-implied expectations provide a daily gauge of investors' sentiment about future currency returns. Specifically, they help assess how closely expectations align with planned policy actions, how quickly policy is transmitted to investors, and the extent to which a policy decision reduces or amplifies investor disagreement. Thus, expectations derived from holdings offer a promising alternative to survey-, macro-, or factor-based measures for calibration, monitoring, and policy evaluation.

References

Adam, Klaus and Stefan Nagel (2023) "Expectations data in asset pricing," in *Handbook of economic expectations*, 477–506: Elsevier.

Ameriks, John, Gábor Kézdi, Minjoon Lee, and Matthew D Shapiro (2020) "Heterogeneity in expectations, risk tolerance, and household stock shares: The attenuation puzzle," *Journal of Business & Economic Statistics*, 38 (3), 633–646.

Amromin, Gene and Steven A Sharpe (2014) "From the horse's mouth: Economic conditions and investor expectations of risk and return," *Management Science*, 60 (4), 845–866.

Anderson, Ronald W and Jean-Pierre Danthine (1980) "Hedging and joint production: Theory and illustrations," *Journal of Finance*, 35 (2), 487–498.

——— (1981) "Cross hedging," *Journal of Political Economy*, 89 (6), 1182–1196.

Bacchetta, Philippe and Eric Van Wincoop (2006) "Can information heterogeneity explain the exchange rate determination puzzle?" *American Economic Review*, 96 (3), 552–576.

Back, Kerry (2010) *Asset pricing and portfolio choice theory*: Oxford University Press.

Backus, David K, Silverio Foresi, and Chris I Telmer (2001) "Affine term structure models and the forward premium anomaly," *Journal of Finance*, 56 (1), 279–304.

Backus, David K and Gregor W Smith (1993) "Consumption and real exchange rates in dynamic economies with non-traded goods," *Journal of International Economics*, 35 (3-4), 297–316.

Bartram, Söhnke M, Leslie Djuranovik, Anthony Garratt, and Yan Xu (2025) "Mispricing and risk premia in currency markets," *Journal of Financial and Quantitative Analysis*, 60 (2), 695–733.

Basak, Suleyman (2005) "Asset pricing with heterogeneous beliefs," *Journal of Banking & Finance*, 29 (11), 2849–2881.

Beber, Alessandro, Francis Breedon, and Andrea Buraschi (2010) "Differences in beliefs and currency risk premiums," *Journal of Financial Economics*, 98 (3), 415–438.

Ben-David, Itzhak, Francesco A Franzoni, and Rabih Moussawi (2017) "Exchange traded funds (ETFs)," *Annual Review of Financial Economics*, 9, 2016–22.

Ben-David, Itzhak, Francesco Franzoni, and Rabih Moussawi (2018) "Do ETFs increase volatility?" *Journal of Finance*, 73 (6), 2471–2535.

Bräuer, Leonie and Harald Hau (2024) "Do Funds Engage in Optimal FX Hedging?" *Swiss Finance Institute Research Paper No. 24-103*.

Brown, David C, Shaun William Davies, and Matthew C Ringgenberg (2021) "ETF arbitrage, non-fundamental demand, and return predictability," *Review of Finance*, 25 (4), 937–972.

Buraschi, Andrea and Alexei Jiltsov (2006) "Model uncertainty and option markets with heterogeneous beliefs," *Journal of Finance*, 61 (6), 2841–2897.

Buraschi, Andrea, Fabio Trojani, and Andrea Vedolin (2014) "When uncertainty blows in the orchard: Comovement and equilibrium volatility risk premia," *Journal of Finance*, 69 (1), 101–137.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo (2011) "Carry trade and momentum in currency markets," *Annu. Rev. Financ. Econ.*, 3 (1), 511–535.

Campbell, John Y, Karine Serfaty-De Medeiros, and Luis M Viceira (2010) "Global currency hedging," *Journal of Finance*, 65 (1), 87–121.

Candian, Giacomo and Pierre De Leo (2025) "Imperfect exchange rate expectations," *Review of Economics and Statistics*, 107 (5), 1406–1423.

Cheema-Fox, Alexander and Robin Greenwood (2024) "How do Global Portfolio Investors Hedge Currency Risk?" *Available at SSRN 5036404*.

Chen, Nanyu and Haonan Zhou (2025) "Managing Emerging Market Currency Risk," *Available at SSRN 5119836*.

Chernov, Mikhail and Drew Creal (2021) "The PPP view of multihorizon currency risk premiums," *Review of Financial Studies*, 34 (6), 2728–2772.

Chernov, Mikhail, Magnus Dahlquist, and Lars Lochstoer (2023) "Pricing currency risks," *Journal of Finance*, 78 (2), 693–730.

Chernov, Mikhail and Philippe Mueller (2012) "The term structure of inflation expectations," *Journal of Financial Economics*, 106 (2), 367–394.

Chionis, Dionysios and Ronald MacDonald (2002) "Aggregate and disaggregate measures of the foreign exchange risk premium," *International Review of Economics & Finance*, 11 (1), 57–84.

Cochrane, John H (2011) "Presidential address: Discount rates," *Journal of finance*, 66 (4), 1047–1108.

Coibion, Olivier and Yuriy Gorodnichenko (2015) "Information rigidity and the expectations formation process: A simple framework and new facts," *American Economic Review*, 105 (8), 2644–2678.

Converse, Nathan, Eduardo Levy-Yeyati, and Tomas Williams (2023) "How ETFs amplify the global financial cycle in emerging markets," *Review of Financial Studies*, 36 (9), 3423–3462.

Crescini, Alessandro, Fabio Trojani, and Andrea Vedolin (2025) "Demand-based Expected Returns."

Curcuru, Stephanie E, Charles P Thomas, Francis E Warnock, and Jon Wongswan (2011) "US international equity investment and past and prospective returns," *American Economic Review*, 101 (7), 3440–3455.

Da, Zhi and Sophie Shive (2018) "Exchange traded funds and asset return correlations," *European Financial Management*, 24 (1), 136–168.

Dahlquist, Magnus and Henrik Hasseltoft (2020) "Economic momentum and currency returns," *Journal of Financial Economics*, 136 (1), 152–167.

Dahlquist, Magnus and Julien Pénasse (2022) "The missing risk premium in exchange rates," *Journal of*

Financial Economics, 143 (2), 697–715.

Dao, Mai Chi, Pierre-Olivier Gourinchas, and Oleg Itskhoki (2025) "Breaking Parity: Equilibrium Exchange Rates and Currency Premia," *Working Paper forthcoming, IMF*.

De Marco, Filippo, Marco Macchiavelli, and Rosen Valchev (2022) "Beyond home bias: International portfolio holdings and information heterogeneity," *Review of Financial Studies*, 35 (9), 4387–4422.

Della Corte, Pasquale, Can Gao, and Alexandre Jeanneret (2023) "Survey expectations meet option prices: New insights from the FX markets," *Working paper*.

Du, Wenxin and Amy Huber (2024) "Dollar Asset Holding and Hedging Around the Globe," *Available at SSRN 4478513*.

Dumas, Bernard, Alexander Kurshev, and Raman Uppal (2009) "Equilibrium portfolio strategies in the presence of sentiment risk and excess volatility," *Journal of Finance*, 64 (2), 579–629.

Dumas, Bernard, Karen K Lewis, and Emilio Osambela (2017) "Differences of opinion and international equity markets," *Review of Financial Studies*, 30 (3), 750–800.

Egan, Mark, Alexander MacKay, and Hanbin Yang (2022) "Recovering investor expectations from demand for index funds," *The Review of Economic Studies*, 89 (5), 2559–2599.

Engel, Charles and Steve Pak Yeung Wu (2023) "Forecasting the US Dollar in the 21st Century," *Journal of International Economics*, 141, 103715.

Evans, Martin DD and Richard K Lyons (2002) "Order flow and exchange rate dynamics," *Journal of Political Economy*, 110 (1), 170–180.

French, Kenneth R. and James M. Poterba (1991) "Investor Diversification and International Equity Markets," *American Economic Review*, 81 (2), 222–226.

Ghosh, Anisha and Guillaume Roussellet (2023) "Identifying beliefs from asset prices," in *Proceedings of Paris December 2019 Finance Meeting EUROFIDAI-ESSEC*.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus (2021) "Five facts about beliefs and portfolios," *American Economic Review*, 111 (5), 1481–1522.

Glosten, Lawrence, Suresh Nallareddy, and Yuan Zou (2021) "ETF activity and informational efficiency of underlying securities," *Management Science*, 67 (1), 22–47.

Greenwood, Robin and Andrei Shleifer (2014) "Expectations of returns and expected returns," *Review of Financial Studies*, 27 (3), 714–746.

Hill, Joanne M, Dave Nadig, and Matt Hougan (2015) *A comprehensive guide to exchange-traded funds (ETFs)*: CFA Institute Research Foundation.

Jiang, Zhengyang, Arvind Krishnamurthy, Hanno N Lustig, Robert Richmond, and Chenzi Xu (2025) "Dollar upheaval: This time is different," *Available at SSRN*.

Jouini, Elyes and Clotilde Napp (2007) "Consensus consumer and intertemporal asset pricing with heterogeneous beliefs," *The Review of Economic Studies*, 74 (4), 1149–1174.

Kalemli-Özcan, Sebnem and Liliana Varela (2021) "Five facts about the UIP Premium," *National Bureau of Economic Research*.

Koijen, Ralph SJ and Motohiro Yogo (2020) "Exchange rates and asset prices in a global demand system," *National Bureau of Economic Research*.

Koont, Naz, Yiming Ma, Lubos Pastor, and Yao Zeng (2025) "Steering a ship in illiquid waters: Active management of passive funds," *The Review of Financial Studies*, 38 (10), 2887–2935.

Kremens, Lukas, Ian Martin, and Liliana Varela (2025) "Long-horizon exchange rate expectations," *Journal of Finance*.

Kubitza, Christian, Jean-David Sigaux, and Quentin Vandeweyer (2024) "The implications of CIP deviations for international capital flows," *Fama-Miller Working Paper, Chicago Booth Research Paper*, 24–18.

Lettau, Martin, Matteo Maggiori, and Michael Weber (2014) "Conditional risk premia in currency markets and other asset classes," *Journal of Financial Economics*, 114 (2), 197–225.

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan (2011) "Common risk factors in currency markets," *Review of Financial Studies*, 24 (11), 3731–3777.

MacDonald, Ronald and Ian W Marsh (1996) "Currency forecasters are heterogeneous: confirmation and consequences," *Journal of International Money and Finance*, 15 (5), 665–685.

Mancini, Loriano, Angelo Ranaldo, and Jan Wrampelmeyer (2013) "Liquidity in the foreign exchange market: Measurement, commonality, and risk premiums," *Journal of Finance*, 68 (5), 1805–1841.

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 (1-2), 3–24.

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf (2012a) "Carry trades and global foreign exchange volatility," *Journal of Finance*, 67 (2), 681–718.

——— (2012b) "Currency momentum strategies," *Journal of Financial Economics*, 106 (3), 660–684.

——— (2017) "Currency value," *Review of Financial Studies*, 30 (2), 416–441.

Merton, Robert C (1980) "On estimating the expected return on the market: An exploratory investigation," *Journal of Financial Economics*, 8 (4), 323–361.

Nagel, Stefan and Zhengyang Xu (2023) "Dynamics of subjective risk premia," *Journal of Financial Economics*, 150 (2), 103713.

Nagel, Stefan and Zhen Yan (2022) "Inflation Hedging on Main Street? Evidence from Retail TIPS Fund Flows," *National Bureau of Economic Research*.

Nickell, Stephen (1981) "Biases in dynamic models with fixed effects," *Econometrica*, 1417–1426.

Opie, Wei and Steven Riddiough (2024) "On the use of currency forwards: Evidence from international equity mutual funds," *Available at SSRN 4432796*.

Opie, Wei and Steven J Riddiough (2020) "Global currency hedging with common risk factors," *Journal of Financial Economics*, 136 (3), 780–805.

Pan, Kevin and Yao Zeng (2017) "ETF arbitrage under liquidity mismatch," *Available at SSRN 3723406*.

Pesch, Daniel, Ilaria Piatti, and Paul Whelan (2024) "Subjective risk premia in bond and fx markets," in *AFA in New Orleans meetings paper*.

Petajisto, Antti (2017) "Inefficiencies in the pricing of exchange-traded funds," *Financial Analysts Journal*, 73 (1), 24–54.

Ross, Steve (2015) "The recovery theorem," *Journal of Finance*, 70 (2), 615–648.

Scheinkman, Jose A and Wei Xiong (2003) "Overconfidence and speculative bubbles," *Journal of political Economy*, 111 (6), 1183–1220.

Shin, Hyun Song, Philip Wooldridge, and Dora Xia (2025) "US dollar's slide in April 2025: the role of FX hedging," *Bank for International Settlements, Working Paper*.

Sialm, Clemens and Qifei Zhu (2024) "Currency Management by International Fixed-Income Mutual Funds," *Journal of Finance*, 79 (6), 4037–4081.

Stavrakeva, Vania and Jenny Tang (2020) "Individual Beliefs, Demand for Currency and Exchange Rate Dynamics," *Working Paper*.

——— (2025) "Explaining the great moderation exchange rate volatility puzzle," *IMF Economic Review*, 73 (1), 196–238.

Verdelhan, Adrien (2018) "The share of systematic variation in bilateral exchange rates," *Journal of Finance*, 73 (1), 375–418.

Table 1: Summary Statistics

Panel A reports assets under management for hedged and unhedged ETFs, the weight in the unhedged index, and the index returns of hedged and unhedged ETFs. Panel B reports three different expected currency returns, along with the forward premium and the covariance between asset and currency returns, all scaled by exchange rate volatility. Panel C reports the fund characteristics: the average log of assets under management, the difference between the expense ratio of the unhedged versus the hedged ETF, the share of shares outstanding held by institutional investors, and the hedging strategy. Panel D (next page) reports aggregates per currency area for portfolio-implied expectations, survey expectations, macro expectations, factor expectations, currency excess returns, exchange rate changes, the one-month forward premium, active flows, dispersion in expectations, and one-month option-implied FX volatility. For readability, all variables in Panel D are multiplied by 100. The sample period spans from January 2014 to June 2025.

	Obs.	Mean	Std	Min	25%	50%	75%	Max
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Fund Sample								
$A_{i,t}^{hedged}$ (mil USD)	350	197	516	1	12	53	185	6,198
$A_{i,t}^{unhedged}$ (mil USD)	350	1,474	4,472	1	24	170	899	39,646
$w_{i,t}$	350	0.678	0.284	0.008	0.494	0.761	0.914	0.998
$r_{i,t}^{hedged}$ (%)	350	0.303	0.566	−1.746	−0.119	0.309	0.724	2.276
$r_{i,t}^{unhedged}$ (%)	350	0.414	0.671	−2.323	−0.075	0.394	0.891	4.195
Panel B: Explanatory Components								
$\frac{\mathbb{E}_t^{survey}[\Delta s_{c,t+1}]}{\mathrm{Var}_t[\Delta s_{t+1}]}$	15,622	0.020	0.326	−1.360	−0.189	0.043	0.192	1.321
$\frac{\mathbb{E}_t^{macro}[\Delta s_{c,t+1}]}{\mathrm{Var}_t[\Delta s_{t+1}]}$	20,429	0.014	0.020	−0.137	0.005	0.012	0.024	0.124
$\frac{\mathbb{E}_t^{factor}[\Delta s_{c,t+1}]}{\mathrm{Var}_t[\Delta s_{t+1}]}$	20,429	0.035	0.313	−2.735	−0.082	0.061	0.211	1.142
$\frac{f_{c,t}-s_{c,t}}{\mathrm{Var}_t[\Delta s_{t+1}]}$	20,429	−0.026	0.034	−0.385	−0.034	−0.017	−0.003	0.084
$\frac{\mathrm{Cov}_t[r_{ic,,t+1},\Delta s_{t+1}]}{\mathrm{Var}_t[\Delta s_{t+1}]}$	20,429	−0.465	0.669	−6.307	−0.790	−0.359	−0.049	7.304
Panel C: Fund Characteristics								
Log Average Size	350	5.664	2.209	0.874	4.091	5.760	7.153	12.122
Net Expense Ratio (%)	315	−0.026	0.065	−0.390	−0.050	−0.020	0.000	0.610
Institutional Share	290	0.225	0.202	0.000	0.071	0.168	0.339	0.935
$\Delta^{strategy}$ (in bp)	350	4.405	21.807	−58.321	−1.808	2.803	5.547	243.934

Table 1 continued.

	Obs.	Mean	Std	Min	25%	50%	75%	Max
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel D: Aggregates								
$\mathbb{E}_t^{PF}[rx_{c,t+1}]$	763	−0.034	0.140	−1.338	−0.100	−0.035	0.018	1.285
$\mathbb{E}_t^{survey}[rx_{c,t+1}]$	638	−0.192	1.622	−5.753	−1.179	−0.227	0.657	6.225
$\mathbb{E}_t^{macro}[rx_{c,t+1}]$	763	0.037	0.109	−0.384	−0.008	0.044	0.115	0.292
$\mathbb{E}_t^{factor}[rx_{c,t+1}]$	763	0.090	1.099	−4.957	−0.692	−0.004	0.847	3.870
$rx_{c,t+1}$	763	0.230	2.390	−7.126	−1.415	0.228	1.845	9.065
$\Delta s_{c,t+1}$	763	0.132	2.399	−7.424	−1.496	0.139	1.775	9.130
$f_{c,t} - s_{c,t}$	763	−0.098	0.134	−0.543	−0.162	−0.069	−0.015	0.228
$flow_{c,t}$	763	0.019	1.399	−11.102	−0.174	0.013	0.247	14.881
$Disp_{c,t}^{PF}$	748	−7.853	0.895	−12.860	−8.311	−7.846	−7.313	−5.116
$IVol_{c,t}$	748	2.053	0.274	1.342	1.858	2.029	2.241	2.963

Table 2: Observed and Optimal Mean-Variance Hedging

The table reports monthly panel regressions of the observed currency weight $w_{i,c,t}$ on the optimal mean–variance components: three currency return expectations $\mathbb{E}_t[\Delta s_{c,t+1}]$, the forward premium $f_{c,t} - s_{c,t}$, and the covariance with exchange rates $\mathbb{C}\text{ov}_t[r_{i,c,t+1}, \Delta s_{c,t+1}]$, all scaled by $\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]$. Panel A shows results for equity ETFs and Panel B for debt ETFs. The sample covers the period from January 2014 to May 2025 (survey expectations to August 2024). I cluster standard errors at the time- and fund-level and mark statistical significance at the 10%, 5%, and 1% level by *, **, and ***, respectively.

Dep. Variable:	$w_{i,c,t}$								
	Baseline Model			Across Funds			Within Funds		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Equity ETFs									
$\frac{\mathbb{E}_t^{survey}[\Delta s_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$	0.18**			0.18***			0.01		
	(0.08)			(0.04)			(0.01)		
$\frac{\mathbb{E}_t^{macro}[\Delta s_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$		5.25***			0.34			0.16	
		(1.07)			(0.71)			(0.22)	
$\frac{\mathbb{E}_t^{factor}[\Delta s_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$			0.25***			−0.13			−0.01
			(0.09)			(0.08)			(0.02)
$\frac{f_{c,t}-s_{c,t}}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$	−10.62***	−8.15***	−9.97***	−2.35***	−1.38***	−1.04**	0.33*	0.25*	0.26
	(0.80)	(0.98)	(0.93)	(0.56)	(0.45)	(0.49)	(0.20)	(0.14)	(0.17)
$\frac{\mathbb{C}\text{ov}_t[r_{i,c,t+1},\Delta s_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$	−0.36***	−0.31***	−0.33***	0.00	0.02	0.02	−0.01***	−0.01***	−0.01***
	(0.03)	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)	(0.00)	(0.00)	(0.00)
Time FE	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Fund FE	No	No	No	No	No	No	Yes	Yes	Yes
R^2 (%)	56.21	55.03	54.23	5.50	1.88	2.28	0.80	0.52	0.50
No. Obs.	9,866	12,713	12,713	9,866	12,713	12,713	9,866	12,713	12,713
Panel B: Debt ETFs									
$\frac{\mathbb{E}_t^{survey}[\Delta s_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$	0.24**			0.20***			−0.01		
	(0.10)			(0.06)			(0.01)		
$\frac{\mathbb{E}_t^{macro}[\Delta s_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$		7.90***			0.43			−0.51	
		(1.25)			(0.98)			(0.45)	
$\frac{\mathbb{E}_t^{factor}[\Delta s_{t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{t+1}]}$			0.69***			−0.08			0.00
			(0.12)			(0.13)			(0.02)
$\frac{f_{c,t}-s_{c,t}}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$	−8.77***	−5.68***	−9.51***	−1.43*	−0.83	−0.54	0.19	0.01	0.06
	(0.87)	(0.99)	(0.88)	(0.75)	(0.60)	(0.69)	(0.21)	(0.15)	(0.16)
$\frac{\mathbb{C}\text{ov}_t[r_{i,c,t+1},\Delta s_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$	−0.46***	−0.44***	−0.41***	0.01	0.04	0.04	0.01	0.00	0.00
	(0.06)	(0.07)	(0.06)	(0.04)	(0.04)	(0.05)	(0.01)	(0.01)	(0.01)
Time FE	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Fund FE	No	No	No	No	No	No	Yes	Yes	Yes
R^2 (%)	51.17	51.37	51.87	2.77	0.75	0.79	0.11	0.34	0.02
No. Obs.	5,756	7,716	7,716	5,756	7,716	7,716	5,756	7,716	7,716

Table 3: Forecast Performance of Exchange Rate Expectations

The table presents results of monthly regressions where the realized currency excess return $rx_{c,t+1}$ (Column 1-4) or currency appreciation $\Delta s_{c,t+1}$ (Column 5-8) is regressed on estimates for average expected currency excess returns, $\mathbb{E}_t[rx_{c,t+1}]$, and the one-month forward premium, $f_{c,t} - s_{c,t}$. $\mathbb{E}_t^{PF}[rx_{c,t+1}]$ are the portfolio-implied expectations, $\mathbb{E}_t^{survey}[rx_{c,t+1}]$ are expectations obtained from surveys, $\mathbb{E}_t^{macro}[rx_{c,t+1}]$ are expectations derived from prominent predictors of exchange rates as in Chernov et al. (2023) and $\mathbb{E}_t^{factor}[rx_{c,t+1}]$ are expectations constructed from two currency factors as in Opie and Riddiough (2020). All regressions include time fixed effects. Expected currency returns are demeaned at the currency level using backward-looking averages. The sample covers the period from January 2014 to June 2025 (for Column (4) only until August 2024). I cluster standard errors at the time-level and mark statistical significance at the 10%, 5%, and 1% level by *, **, and ***, respectively.

Dep. Variable:	$rx_{c,t+1}$				$\Delta s_{c,t+1}$			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\mathbb{E}_t^{PF}[rx_{c,t+1}]$	1.12**				1.12**			
	(0.497)				(0.497)			
$\mathbb{E}_t^{survey}[rx_{c,t+1}]$		0.097				0.097		
		(0.078)				(0.078)		
$\mathbb{E}_t^{macro}[rx_{c,t+1}]$			−0.370				−0.370	
			(0.828)				(0.828)	
$\mathbb{E}_t^{factor}[rx_{c,t+1}]$				0.112				0.112
				(0.151)				(0.151)
$f_t - s_t$	−0.640	−0.600	−0.690	−0.776	0.360	0.400	0.310	0.224
	(0.699)	(0.865)	(0.719)	(0.711)	(0.699)	(0.865)	(0.719)	(0.711)
R^2 (%)	1.19	0.00	0.08	0.21	1.00	0.00	0.11	0.19
No. Obs.	763	701	763	763	763	701	763	763

Table 4: Forecast Performance of Components of PF-implied Expectations

The table presents results of monthly regressions where the dependent variable is the realized one-month currency excess return, $rx_{c,t+1}$. In Column (1)-(3) the independent variable is one of the components of portfolio-implied expectations: weight in the unhedged ETF, $w_{c,t}^{flow}$, variance of spot returns, $\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]$, covariance of underlying average asset returns and spot returns, $\mathbb{C}\text{ov}_t[r_{c,t+1}, \Delta s_{c,t+1}]$. In Column (4) and (5) the weight in the unhedged ETF is further decomposed into the active flows into the unhedged ETF, $flow_{c,t}$, and previous month's weight, $w_{c,t-1}$. All regressions control for the one-month forward premium, $f_{c,t} - s_{c,t}$, and include time fixed effect. The sample covers the period from January 2014 to June 2025. I cluster standard errors at the time-level and mark statistical significance at the 10%, 5%, and 1% level by *, **, and ***, respectively.

Dep. Variable:	$rx_{c,t+1}$				
	(1)	(2)	(3)	(4)	(5)
$w_{c,t}^{flow}$	-0.005				
	(0.003)				
$\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]$		-0.541			
		(0.955)			
$\mathbb{C}\text{ov}_t[r_{c,t+1}, \Delta s_{c,t+1}]$			2.306		
			(1.760)		
$flow_{c,t}$				0.090^{**}	0.088^{*}
				(0.045)	(0.045)
$w_{c,t-1}$				-0.005	
				(0.003)	
$\mathbb{E}_t^{PF}[rx_{c,t+1}]$					1.079^{**}
					(0.482)
$f_{c,t} - s_{c,t}$	-0.691	-0.660	-0.441	-0.735	-0.704
	(0.707)	(0.705)	(0.728)	(0.707)	(0.700)
R^2 (%)	0.62	0.29	0.81	1.38	1.56
No. Obs.	763	763	763	763	763

Table 5: Forecast Performance at Different Horizons

The table presents results of daily regressions where the dependent variable is the realized k-months currency excess return, $rx^k_{c,t+1}$, and the independent variables are active flows, $flow_{c,t}$, into the unhedged versus hedged ETF and portfolio-implied expectations, $\mathbb{E}^{PF}_t[rx_{c,t+1}]$. All regressions include the k-month forward premium, $f^k_{c,t} - s_{c,t}$. Each panel computes realized returns and forward premia for a different horizon. Standard errors are computed using a moving-block bootstrap (10,000 replications) to account for autocorrelation from overlapping observations and significance is based on bootstrap p-values. I denote by *, **, and ***, the significance levels at the 10%, 5%, and 1%, respectively.

Dep. Variable:		$rx^k_{c,t+1}$	
	(1)	(2)	(3)
Panel A: 1-month horizon ($k=1$)			
$flow_{c,t}$	0.214*		0.188*
	(0.110)		(0.109)
$\mathbb{E}^{PF}_t[rx_{c,t+1}]$		1.006***	0.937***
		(0.275)	(0.272)
$f^k_{c,t} - s_{c,t}$	−0.787	−1.422**	−1.296*
	(0.618)	(0.675)	(0.657)
R^2 (%)	0.97	1.38	1.89
Panel B: 3-months horizon ($k=3$)			
$flow_{c,t}$	0.394		0.350
	(0.255)		(0.254)
$\mathbb{E}^{PF}_t[rx_{c,t+1}]$		1.783*	1.655*
		(0.838)	(0.830)
$f^k_{c,t} - s_{c,t}$	−0.605	−0.986	−0.911
	(0.566)	(0.660)	(0.645)
R^2 (%)	1.48	1.91	2.59
Panel C: 12-months horizon ($k=12$)			
$flow_{c,t}$	1.173***		1.051***
	(0.361)		(0.300)
$\mathbb{E}^{PF}_t[rx_{c,t+1}]$		4.929**	4.538**
		(2.410)	(2.309)
$f^k_{c,t} - s_{c,t}$	−0.670	−0.934*	−0.879*
	(0.412)	(0.490)	(0.465)
R^2 (%)	5.45	6.05	7.75
No. Obs.	14,389	14,389	14,389

Table 6: Dispersion in Portfolio-implied Expectations and FX volatility

The table presents results of monthly regressions where the dependent variable is the log of the one-month option-implied FX volatility, $IVol_{c,t}$, and the independent variables are the log mean absolute deviation of portfolio-implied expectations, $Disp^{PF}_{c,t}$, and the lagged one-month option-implied volatility, $IVol_{c,t-1}$. In Column (1) and (2) the regression is performed in levels, in Column (3) and (4) the regression is performed in changes. The sample covers the period from January 2014 to June 2025 (for AUD, I start in 2015, as there are not enough observations to compute the mean absolute deviation). In all regressions, I add currency and time fixed effects. I cluster standard errors by time and currency and mark statistical significance at the 10%, 5%, and 1% level by *, **, and ***, respectively.

Dep. Variable:	$IVol_{c,t}$		$\Delta IVol_{c,t}$		$IVol_{c,t}$	
	(1)	(2)	(3)	(4)	(5)	(6)
$Disp^{PF}_{c,t}$	0.134***	0.055***				
	(0.025)	(0.009)				
$Disp^{PF}_{c,t-1}$					0.112***	0.019*
					(0.026)	(0.008)
$IVol_{c,t-1}$		0.662***				0.721***
		(0.045)				(0.038)
$\Delta Disp^{PF}_{c,t}$			0.027**	0.027**		
			(0.009)	(0.007)		
$\Delta IVol_{c,t-1}$				−0.269***		
				(0.038)		
Overall R^2	0.81	0.90	0.68	0.70	0.78	0.89
Within R^2	0.27	0.61	0.03	0.10	0.19	0.58
No. Obs.	749	749	749	749	742	742

Figure 1: Assets Under Management and Flows for FX-hedged and unhedged ETFs

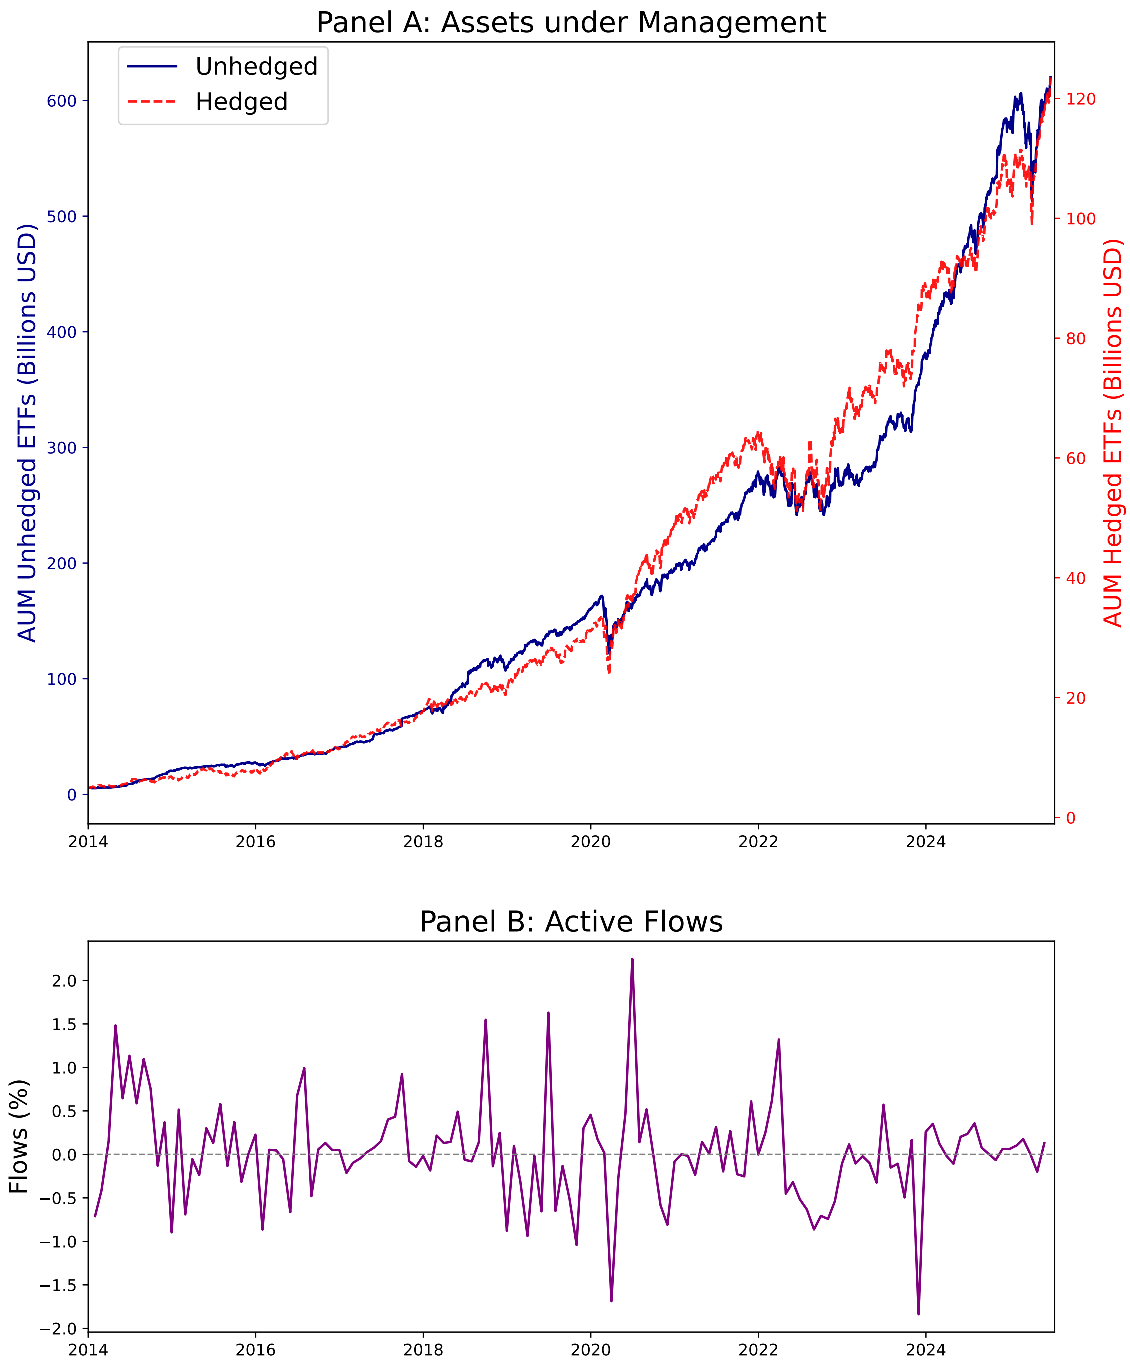

Notes: The figure depicts the assets under management in USD billions in Panel A and active flows in Panel B of currency-hedged ETFs (red, right axis) and their unhedged counterparts (blue, left axis) over the period 2014-2025.

Figure 2: Expected Currency Excess Returns

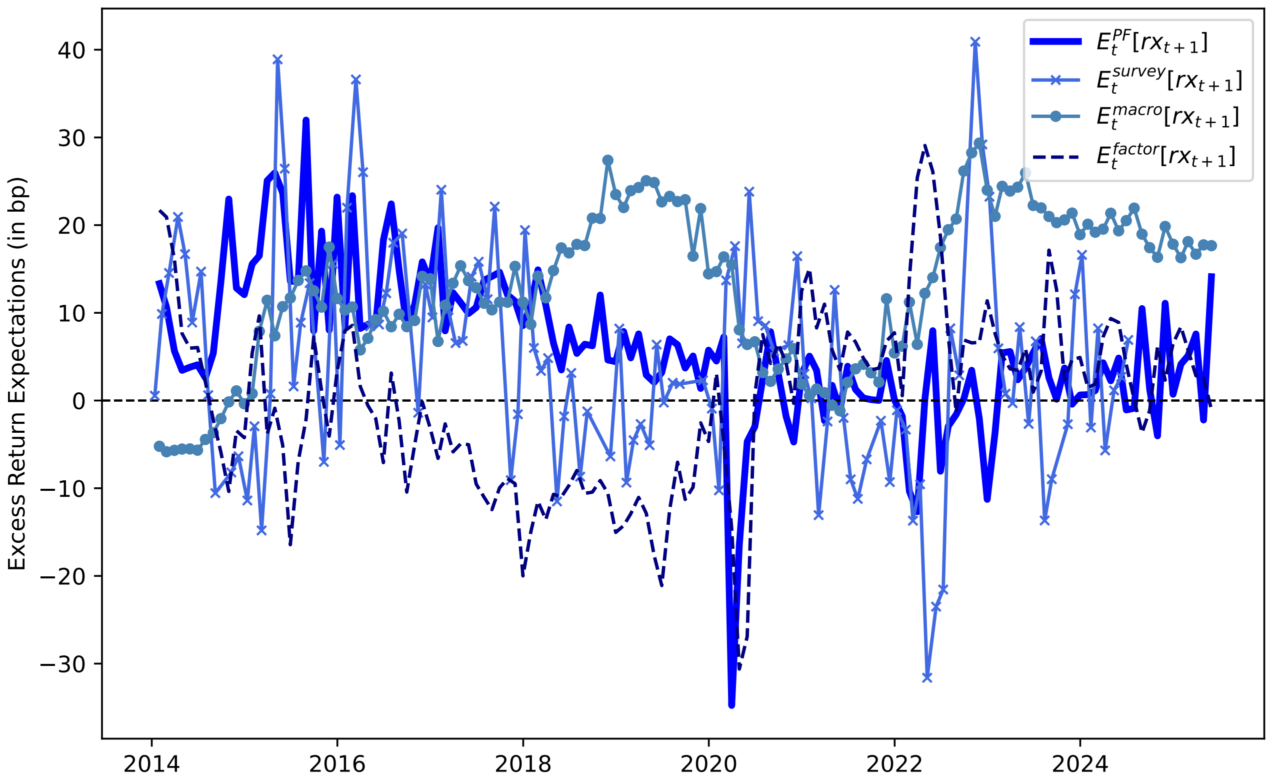

Notes: The figure plots four different estimates of monthly expected currency excess returns averaged across currency areas over the period 2014-2025. The estimates include portfolio-implied expectation (solid line), survey expectations (crosses on line), macro expectations (dots on line) and factor expectations (dashed line). The survey and factor expectations are divided by 10 to match the scale of the other expectation estimates.

Figure 3: Currency Demand and Optimal Hedging Components

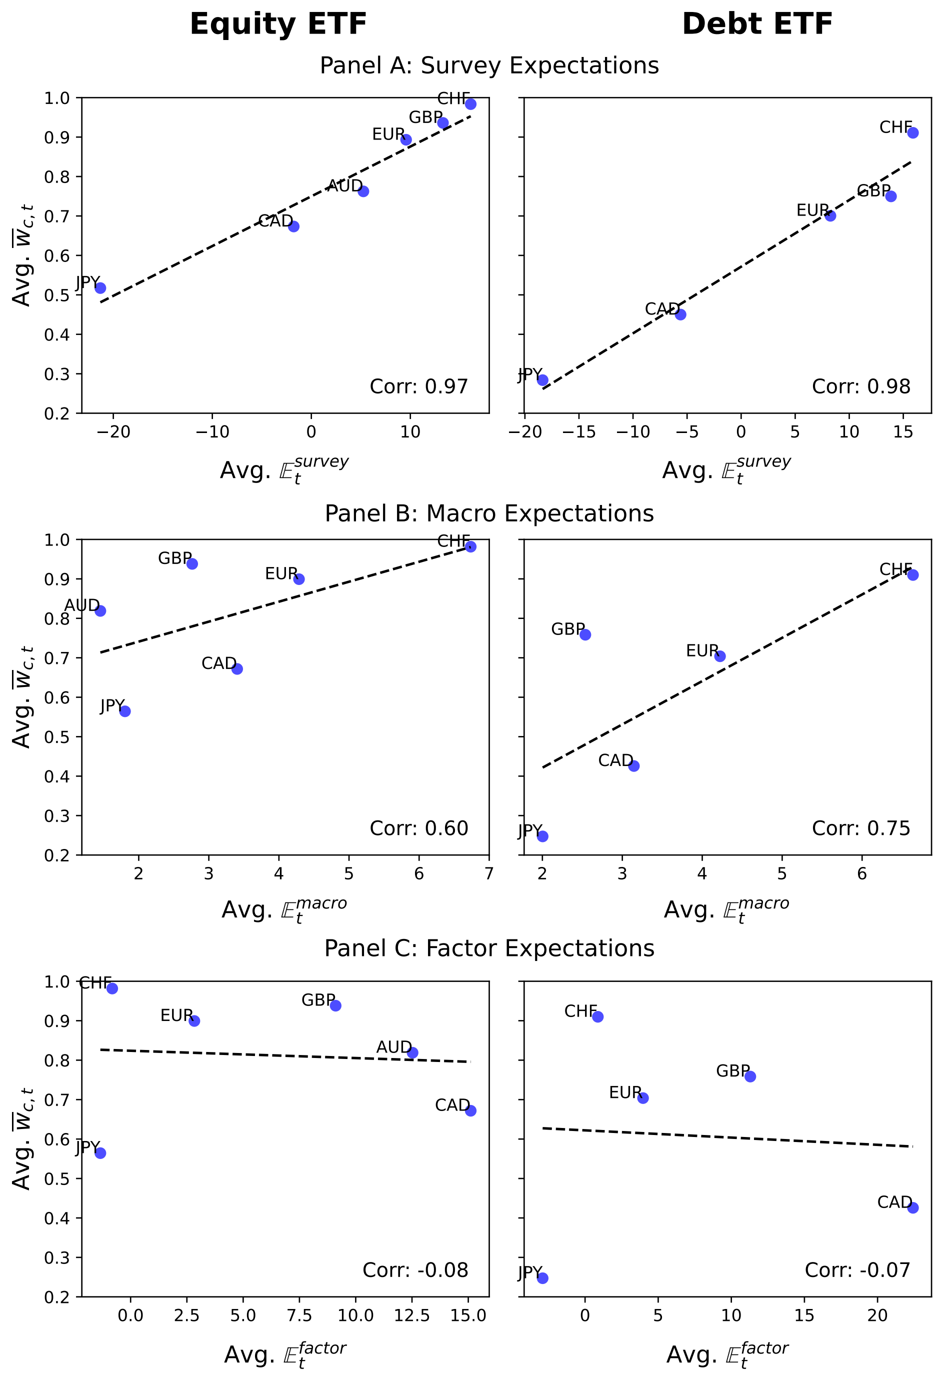

Notes: The figure shows the time-averaged weight in unhedged ETFs on the vertical axis against three estimates of the expected excess returns (Panel A-C) per currency for equity funds (left column) and debt funds (right column). All variables are scaled by the inverse of the exchange rate variance and the benchmark hedge is subtracted as in Eq. (4). Panel A uses survey-based, Panel B macro-based and Panel C factor-based expectations of currency excess returns. The dashed line is the trendline.

Figure 4: Heterogeneity in Currency Speculation and Hedging

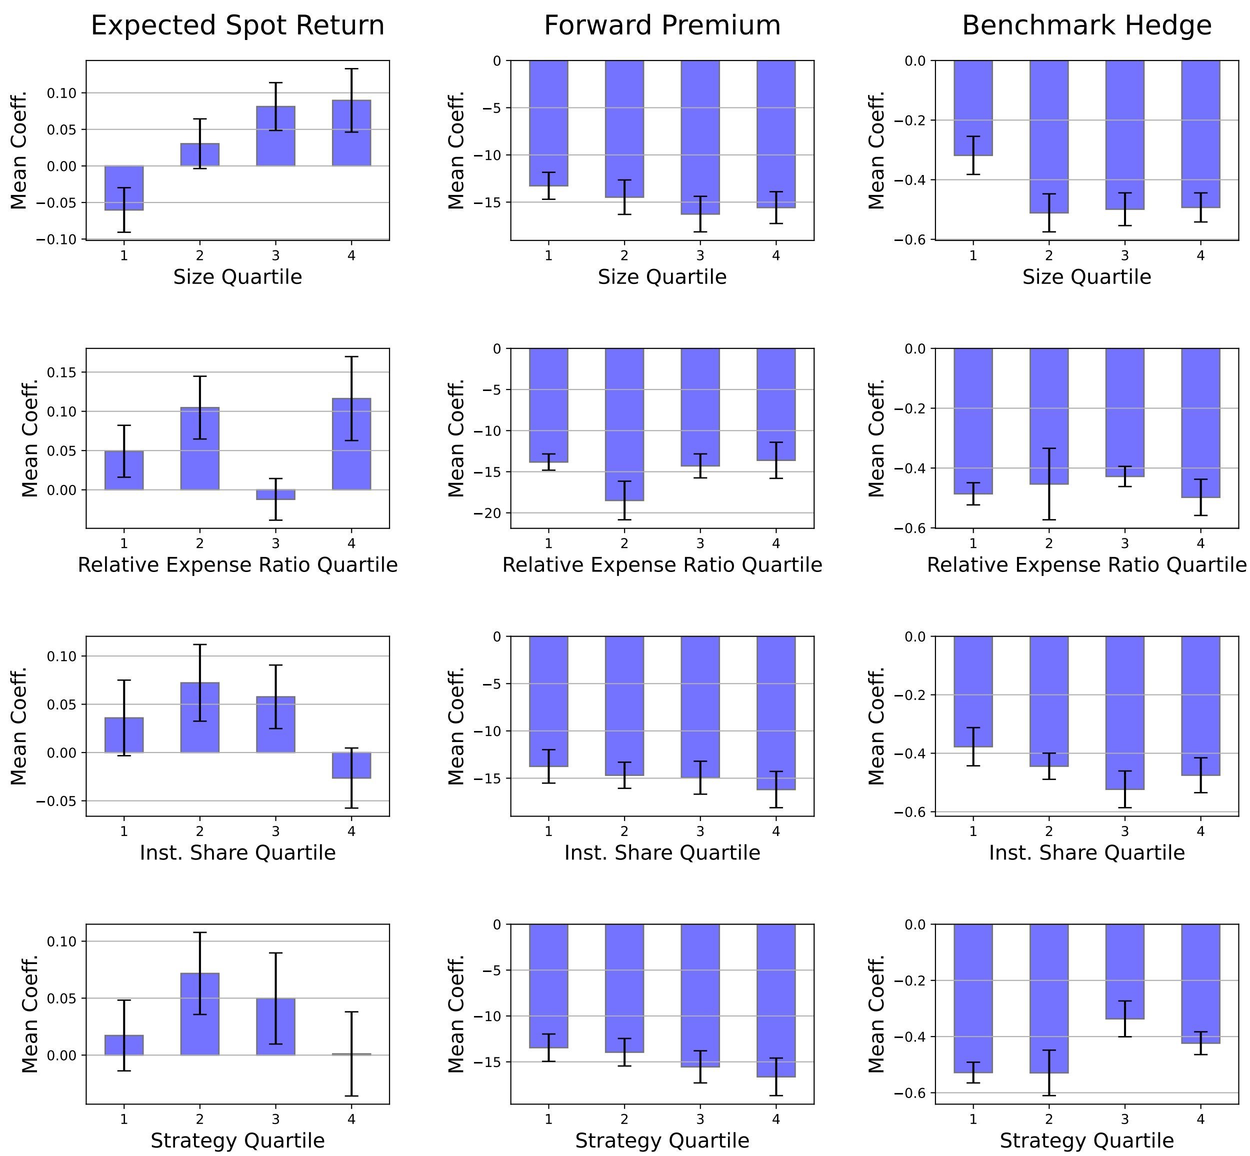

Notes: The figure presents the mean coefficient per quartile of the expected currency return in Column (1), the forward premium in Column (2) and the benchmark hedge in Column (3), categorized by fund characteristics, when running the baseline regression in Table 2 for each fund separately. The black bar indicates the standard errors of the mean, i.e., the standard deviation divided by the square root of the number of observations. Funds are sorted into quartiles by log total assets, net expense ratio, average institutional share, and average hedging strategy of the hedged and unhedged ETFs.

Figure 5: Portfolio-implied Expectations

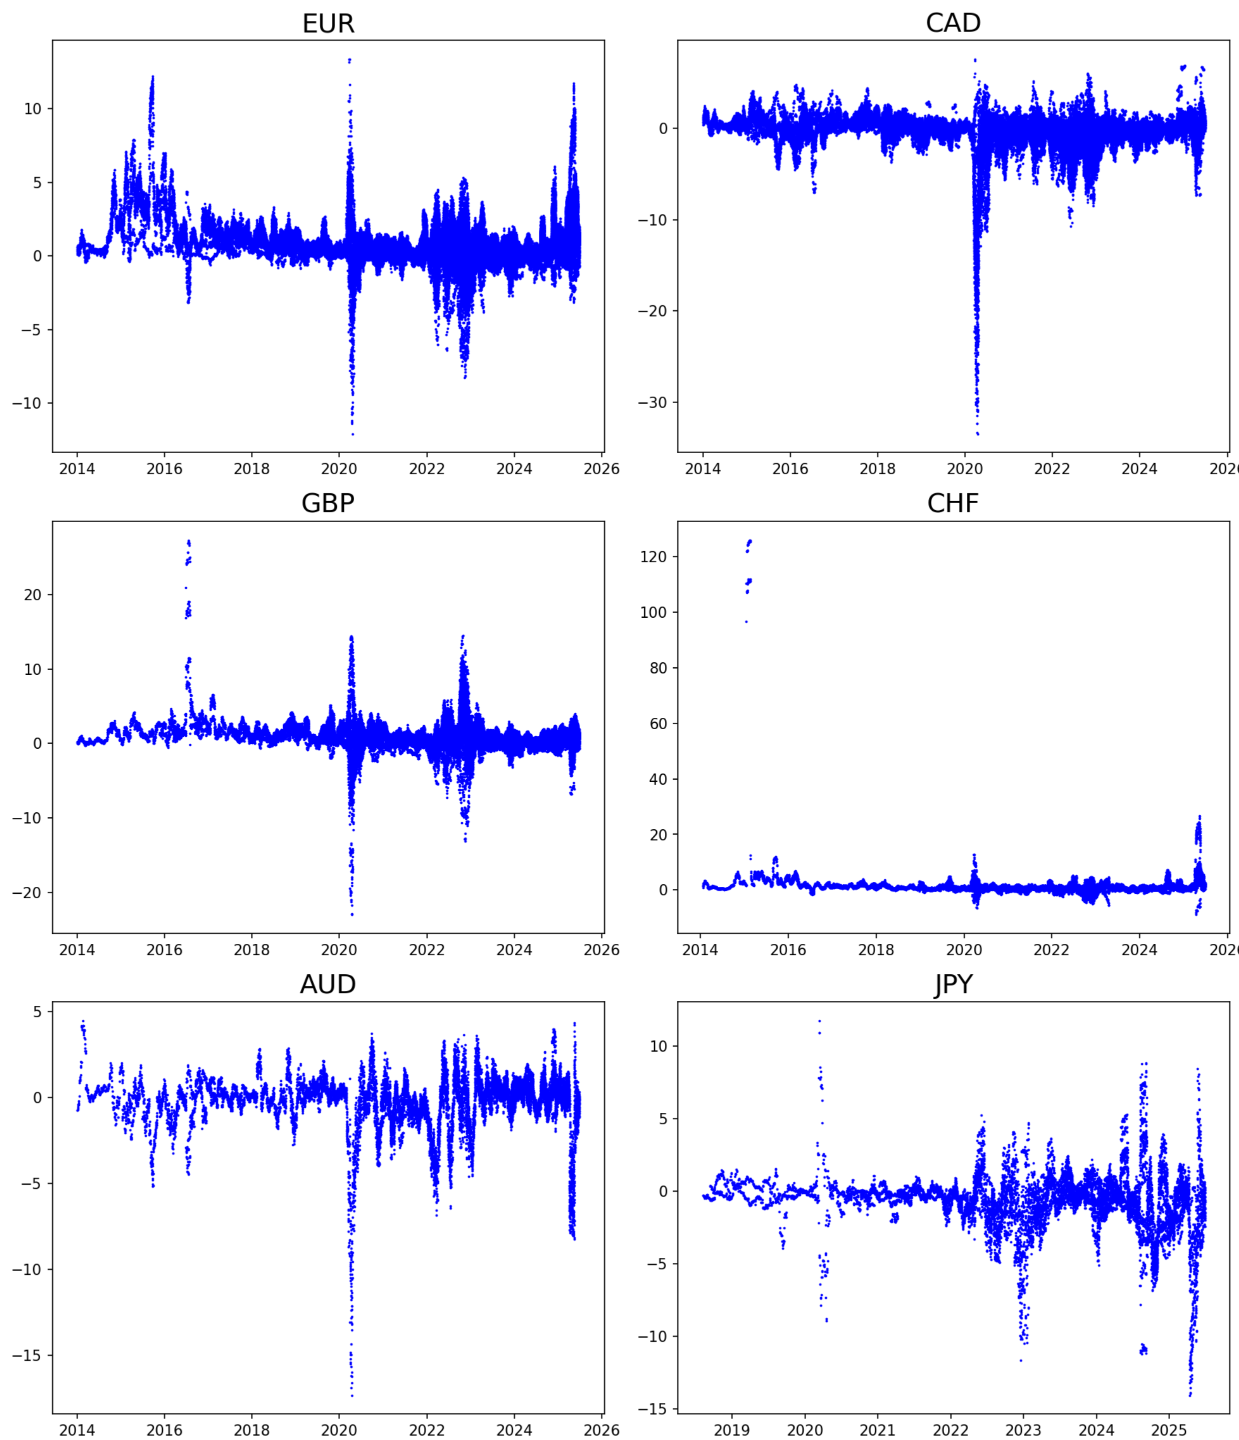

Notes: The figure shows the portfolio-implied expectations per currency. Each dot represents a fund-day observation.

Figure 6: Dispersion in Portfolio-implied Expectations and FX volatility

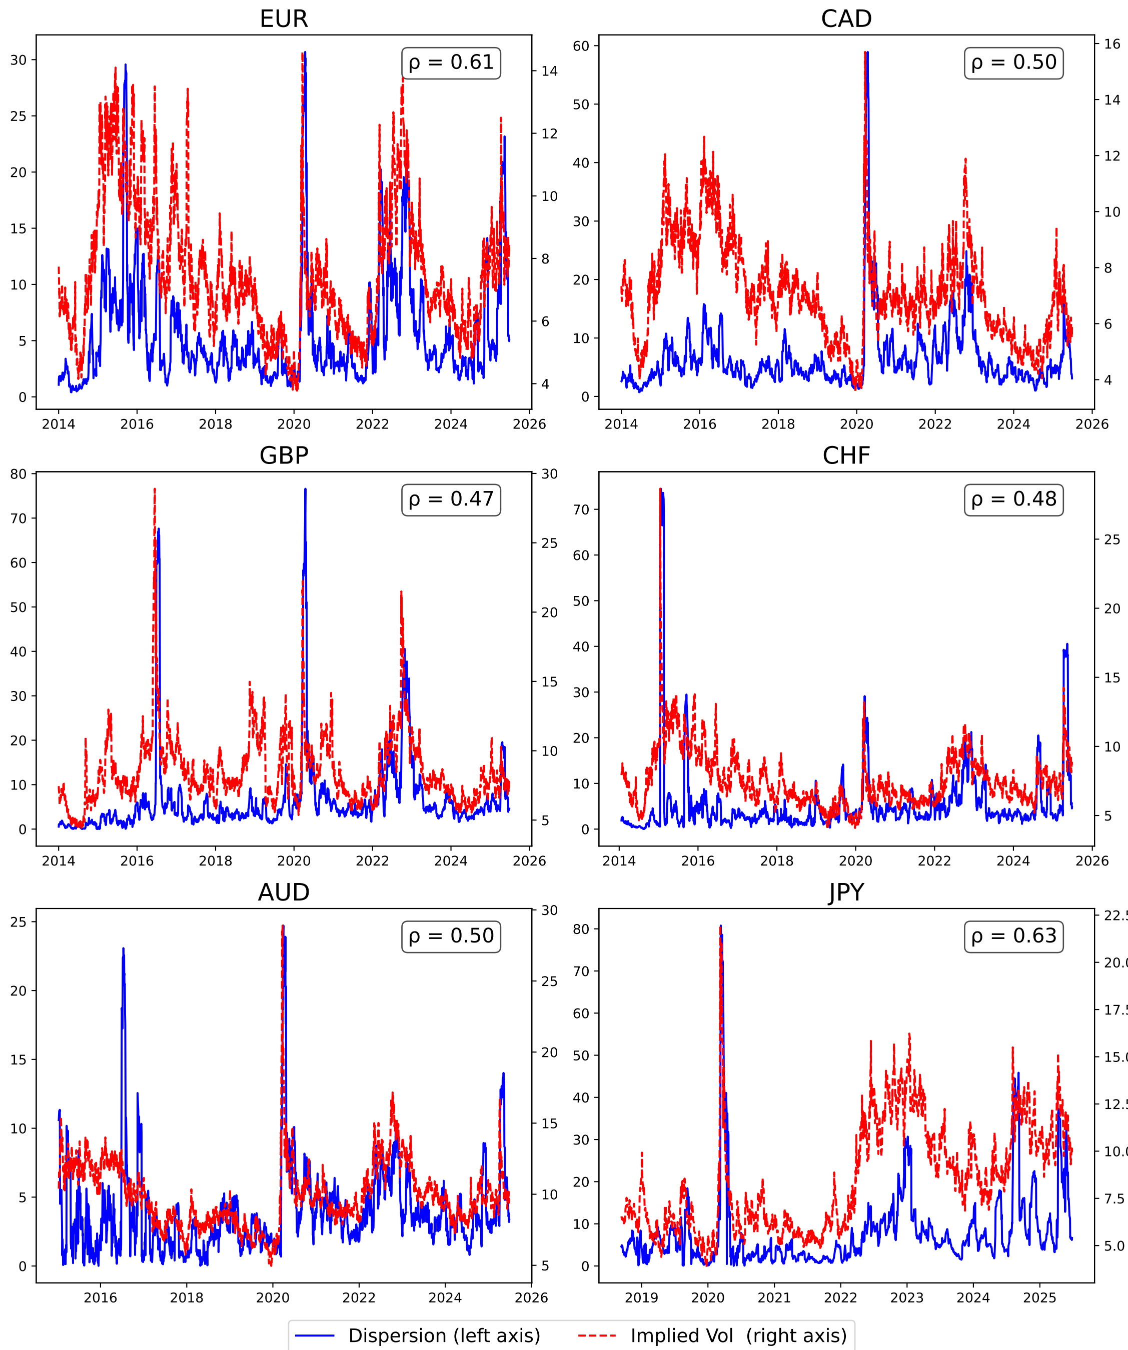

Notes: The figure shows the mean absolute deviation of portfolio-implied expectations across funds defined in Eq. (7) in blue (left axis) against the option-implied FX volatility in red (right axis) for each currency area.

Appendix

Exchange Rate Expectations and Currency Demand

A Additional Tables and Figures

Table A.1: Predictive Ability of Conditional Variance

The table presents results of monthly regressions where the variance of realized currency excess returns is regressed on the estimated conditional variance $Var_t(rx_{c,t+1})$. In Column (1) the realized variance is computed using surveys, in Column (2) using macro expectations, and in Column (3) using factor expectations. All regressions include currency and time fixed effects. I double-cluster standard errors by time and currency and mark statistical significance at the 10%, 5%, and 1% level by *, **, and ***, respectively.

Dep. Variable:	$(rx_{c,t+1} - \mathbb{E}_t^{survey}[rx_{c,t+1}])^2$	$(rx_{c,t+1} - \mathbb{E}_t^{macro}[rx_{c,t+1}])^2$	$(rx_{c,t+1} - \mathbb{E}_t^{factor}[rx_{c,t+1}])^2$
	(1)	(2)	(3)
$Var_t(rx_{c,t+1})$	11.690*	1.175	1.407*
	(6.688)	(1.142)	(0.806)
R^2	0.032	0.001	0.002
No. Obs.	701	763	763

Table A.2: Mispricing between Hedged and Unhedged ETFs

The table presents results of daily panel regressions where the premia and discounts of the hedged ETF is regressed on the premia and discounts of its unhedged counterpart, the VIX and the interaction term. The premia and discounts of an ETF are calculated using the mispricing of an ETF $\epsilon = \frac{P_t - NAV_t}{NAV_t} \times 100$, where a premium corresponds to a positive value of ϵ and a discount to a negative value of ϵ. The VIX is normalized to have mean zero. All regressions include fund fixed effects. I double-cluster standard errors by time and currency and mark statistical significance at the 10%, 5%, and 1% level by *, **, and ***, respectively.

Dep. Variable:	$Premium_{i,t}^{hedged}$			$Discount_{i,t}^{hedged}$		
	(1)	(2)	(3)	(4)	(5)	(6)
$Premium_{i,t}^{unhedged}$	0.313***		0.222***	0.255***		0.182***
	(0.048)		0.034***	(0.044)		0.035
VIX_t		0.023***	0.015***		−0.019***	−0.012***
		(0.003)	(0.001)		(−8.737)	0.002***
$Premium_{i,t}^{unhedged} \times VIX_t$			1.077***			0.722***
			(0.350)			(0.128)
R^2	0.09	0.05	0.14	0.06	0.04	0.09
No. Obs.	194, 223	194, 223	194, 223	159, 097	159, 097	159, 097

Table A.3: Expected Currency Returns and Portfolio USD Shares

The table reports monthly panel regressions of the observed currency weight $w_{i,c,t}$ on three currency excess return expectations $\mathbb{E}_t[\Delta s_{c,t+1}]$. Panel A shows results for equity ETFs and Panel B for debt ETFs. The row *Sensitivity* reports the change in the USD share in terms of standard deviations to a one standard deviation change in expected currency returns, computed as $\hat{\beta} \times (\text{std}(\mathbb{E}_t[\Delta s_{c,t+1}])/\text{std}(w_{i,c,t}))$. The sample covers the period from January 2014 to May 2025 (survey expectations to August 2024). I cluster standard errors at the time- and fund-level and mark statistical significance at the 10%, 5%, and 1% level by *, **, and ***, respectively.

Dep. Variable:	$w_{i,c,t}$								
	Baseline Model			Across Funds			Within Funds		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Equity ETFs									
$\mathbb{E}_t^{survey}[\Delta s_{c,t+1}]$	1.457***			4.097***			0.379*		
	(0.455)			(0.887)			(0.220)		
$\mathbb{E}_t^{macro}[\Delta s_{c,t+1}]$		26.842**			33.846**			7.452	
		(12.993)			(16.255)			(6.415)	
$\mathbb{E}_t^{factor}[\Delta s_{c,t+1}]$			−0.263			−3.236			−0.496
			(0.674)			(2.712)			(0.866)
Sensitivity	0.08	0.08	−0.01	0.22	0.10	−0.13	0.02	0.02	−0.02
Time FE	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Fund FE	No	No	No	No	No	No	Yes	Yes	Yes
R^2 (%)	0.620	0.580	0.010	1.930	0.710	0.270	0.110	0.120	0.030
No. Obs.	10,006	12,713	12,713	10,006	12,713	12,713	10,006	12,713	12,713
Panel B: Debt ETFs									
$\mathbb{E}_t^{survey}[\Delta s_{c,t+1}]$	0.971*			5.046***			0.092		
	(0.534)			(1.389)			(0.261)		
$\mathbb{E}_t^{macro}[\Delta s_{c,t+1}]$		17.979			45.004**			−10.309	
		(21.917)			(20.783)			(14.600)	
$\mathbb{E}_t^{factor}[\Delta s_{c,t+1}]$			−3.859***			−5.252			−1.410
			(1.172)			(4.723)			(1.851)
Sensitivity	0.05	0.05	−0.13	0.27	0.13	−0.18	0.00	−0.03	−0.05
Time FE	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Fund FE	No	No	No	No	No	No	Yes	Yes	Yes
R^2 (%)	0.280	0.260	1.760	2.120	1.300	0.540	0.000	0.190	0.090
No. Obs.	5,756	7,716	7,716	5,756	7,716	7,716	5,756	7,716	7,716

Table A.4: Observed and Optimal Mean-Variance Hedging Using Excess Returns

The table reports monthly panel regressions of the observed currency weight $w_{i,c,t}$ on the optimal mean–variance components: three currency excess return expectations $\mathbb{E}_t[rx_{c,t+1}]$, the forward premium $f_{c,t} - s_{c,t}$, and the covariance with exchange rates $\mathbb{C}\text{ov}_t[r_{i,c,t+1}, \Delta s_{c,t+1}]$, all scaled by $\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]$. Panel A shows results for equity ETFs and Panel B for debt ETFs. The sample covers the period from January 2014 to May 2025 (survey expectations to August 2024). I cluster standard errors at the time- and fund-level and mark statistical significance at the 10%, 5%, and 1% level by *, **, and ***, respectively.

Dep. Variable:	$w_{i,c,t}$								
	Baseline Model			Across Funds			Within Funds		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Equity ETFs									
$\frac{\mathbb{E}_t^{survey}[rx_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$	0.31***			0.19***			0.01		
	(0.11)			(0.04)			(0.01)		
$\frac{\mathbb{E}_t^{macro}[rx_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$		7.24***			1.06***			−0.11	
		(0.73)			(0.33)			(0.14)	
$\frac{\mathbb{E}_t^{factor}[rx_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$			0.10			−0.17**			−0.01
			(0.13)			(0.07)			(0.02)
$\frac{\mathbb{C}\text{ov}_t[r_{i,c,t+1},\Delta s_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$	−0.44***	−0.30***	−0.43***	0.02	0.02	0.03	−0.01***	−0.01***	−0.01***
	(0.03)	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)	(0.00)	(0.00)	(0.00)
Time FE	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Fund FE	No	No	No	No	No	No	Yes	Yes	Yes
R^2 (%)	35.12	54.72	30.59	2.69	1.61	1.31	0.54	0.34	0.28
No. Obs.	9,866	12,713	12,713	9,866	12,713	12,713	9,866	12,713	12,713
Panel B: Debt ETFs									
$\frac{\mathbb{E}_t^{survey}[rx_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$	0.33***			0.19***			−0.00		
	(0.13)			(0.05)			(0.01)		
$\frac{\mathbb{E}_t^{macro}[rx_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$		6.35***			0.73			−0.18	
		(0.82)			(0.50)			(0.16)	
$\frac{\mathbb{E}_t^{factor}[rx_{t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{t+1}]}$			0.23			−0.15			0.01
			(0.19)			(0.12)			(0.02)
$\frac{\mathbb{C}\text{ov}_t[r_{i,c,t+1},\Delta s_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$	−0.79***	−0.45***	−0.84***	0.00	0.03	0.04	0.01	0.00	0.01
	(0.07)	(0.07)	(0.07)	(0.04)	(0.05)	(0.05)	(0.01)	(0.01)	(0.01)
Time FE	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Fund FE	No	No	No	No	No	No	Yes	Yes	Yes
R^2 (%)	30.32	51.13	27.30	1.77	0.71	0.55	0.04	0.16	0.01
No. Obs.	5,756	7,716	7,716	5,756	7,716	7,716	5,756	7,716	7,716

Table A.5: Observed and Optimal Mean-Variance Hedging at the Currency Level

The table reports monthly panel regressions of the observed aggregate currency weight $w_{c,t}$ on the optimal mean–variance components: three currency return expectations $\mathbb{E}_t[\Delta s_{c,t+1}]$, the forward premium $f_{c,t} - s_{c,t}$, and the covariance with exchange rates $\mathbb{C}\text{ov}_t[r_{c,t+1}, \Delta s_{c,t+1}]$, all scaled by $\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]$. Panel A shows results for equity ETFs and Panel B for debt ETFs. The sample covers the period from January 2014 to May 2025 (survey expectations to August 2024). I cluster standard errors at the time- and fund-level and mark statistical significance at the 10%, 5%, and 1% level by *, **, and ***, respectively.

Dep. Variable:	$w_{c,t}$								
	Baseline Model			Across Funds			Within Funds		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Equity ETFs									
$\frac{\mathbb{E}_t^{survey}[\Delta s_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$	0.47***			0.19**			0.01		
	(0.10)			(0.08)			(0.02)		
$\frac{\mathbb{E}_t^{macro}[\Delta s_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$		4.96*			1.90***			0.73	
		(2.94)			(0.40)			(0.22)	
$\frac{\mathbb{E}_t^{factor}[\Delta s_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$			0.32			0.19*			0.03
			(0.23)			(0.11)			(0.03)
$\frac{f_{c,t}-s_{c,t}}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$	−9.95***	−7.52***	−9.59***	−1.11	−0.26	−0.98	−0.40	−0.10	−0.28
	(1.50)	(1.12)	(1.99)	(0.95)	(0.82)	(0.71)	(0.60)	(0.36)	(0.31)
$\frac{\mathbb{C}\text{ov}_t[r_{c,t+1},\Delta s_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$	−0.52***	−0.48***	−0.47***	0.07***	0.12***	0.11***	0.00	0.01	0.01
	(0.10)	(0.10)	(0.12)	(0.02)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)
Time FE	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Currency FE	No	No	No	No	No	No	Yes	Yes	Yes
R^2 (%)	54.01	50.72	49.37	17.60	19.12	14.16	1.67	8.28	1.02
No. Obs.	636	763	763	636	763	763	636	763	763
Panel B: Debt ETFs									
$\frac{\mathbb{E}_t^{survey}[\Delta s_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$	0.36**			0.22***			0.00		
	(0.16)			(0.06)			(0.02)		
$\frac{\mathbb{E}_t^{macro}[\Delta s_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$		2.26			0.21			−0.90	
		(2.18)			(0.98)			(0.73)	
$\frac{\mathbb{E}_t^{factor}[\Delta s_{t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{t+1}]}$			0.47*			0.17			−0.03
			(0.24)			(0.14)			(0.05)
$\frac{f_{c,t}-s_{c,t}}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$	−9.31***	−6.38***	−8.74***	−1.64	−0.95	−1.49	−0.31	−0.18	−0.01
	(0.92)	(1.36)	(1.54)	(1.36)	(1.43)	(1.22)	(0.57)	(0.47)	(0.27)
$\frac{\mathbb{C}\text{ov}_t[r_{c,t+1},\Delta s_{c,t+1}]}{\mathbb{V}\text{ar}_t[\Delta s_{c,t+1}]}$	−0.47**	−0.38*	−0.29	0.24	0.38**	0.37**	−0.02	−0.01	0.01
	(0.21)	(0.20)	(0.18)	(0.16)	(0.17)	(0.17)	(0.06)	(0.04)	(0.04)
Time FE	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Currency FE	No	No	No	No	No	No	Yes	Yes	Yes
R^2 (%)	44.65	36.50	38.62	19.25	21.20	22.35	0.38	3.54	0.22
No. Obs.	451	579	579	451	579	579	451	579	579

Table A.6: Model-implied optimal USD Demand without Speculation

The table presents time-averaged variances and covariances in Columns (1) and (2), model-implied and observed USD Demand in Columns (3) and (4), and model-implied and observed hedge ratios (HR) in Columns (5) and (6) per currency area, for equity ETFs in Panel A and debt ETFs in Panel B. The optimal weight in the unhedged index, w_f^*, is computed as $-\mathbb{C}\text{ov}[\Delta s_{t+1}, r_{t+1}]/\mathbb{V}\text{ar}[\Delta s_{t+1}]$, while the hedge ratio is equal to $1 - w_f$.

	$\mathbb{V}\text{ar}[\Delta s_{t+1}]$	$\mathbb{C}\text{ov}[\Delta s_{t+1}, r_{t+1}]$	w_f^*	$w_f^{observed}$	HR^*	$HR^{observed}$
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Equity ETFs						
AUD	1.48	-1.03	0.70	0.78	0.30	0.22
CAD	0.70	-0.70	1.00	0.60	0.00	0.40
CHF	1.08	0.14	-0.13	0.92	1.13	0.08
EUR	0.84	-0.20	0.24	0.83	0.76	0.17
GBP	0.84	-0.35	0.42	0.85	0.58	0.15
JPY	0.94	0.52	-0.55	0.56	1.55	0.44
Panel B: Debt ETFs						
CAD	1.48	0.02	-0.01	0.42	1.01	0.58
CHF	0.70	0.03	-0.05	0.88	1.05	0.12
EUR	1.08	-0.10	0.09	0.64	0.91	0.36
GBP	0.84	-0.05	0.05	0.67	0.95	0.33
JPY	0.84	-0.03	0.03	0.30	0.97	0.70

Figure A.1: USD Expectations around US Tariff Shock 2025

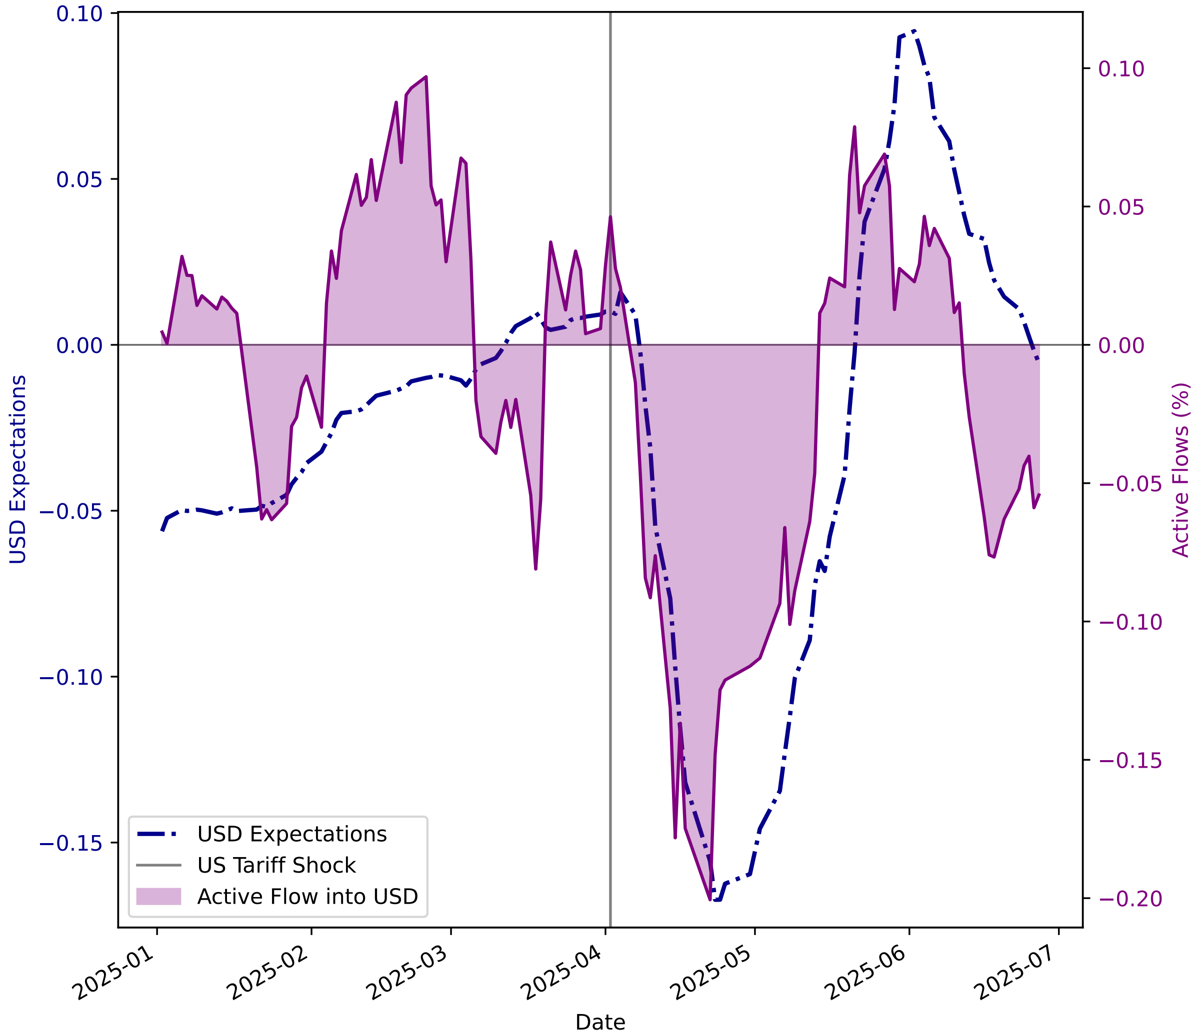

Notes: The figure shows portfolio-implied expectations extracted from ETF trading (blue dashed line, left axis) and active flows into unhedged ETFs relative to their hedged counterparts expressed in percent (purple shaded area, right axis). Expectations are plotted as a two-week rolling mean, flows are plotted as a two-week rolling sum. The gray vertical line marks the event date of April 2, 2025. The sample period spans from January to June 2025.

Figure A.2: ETF Arbitrage Mechanism

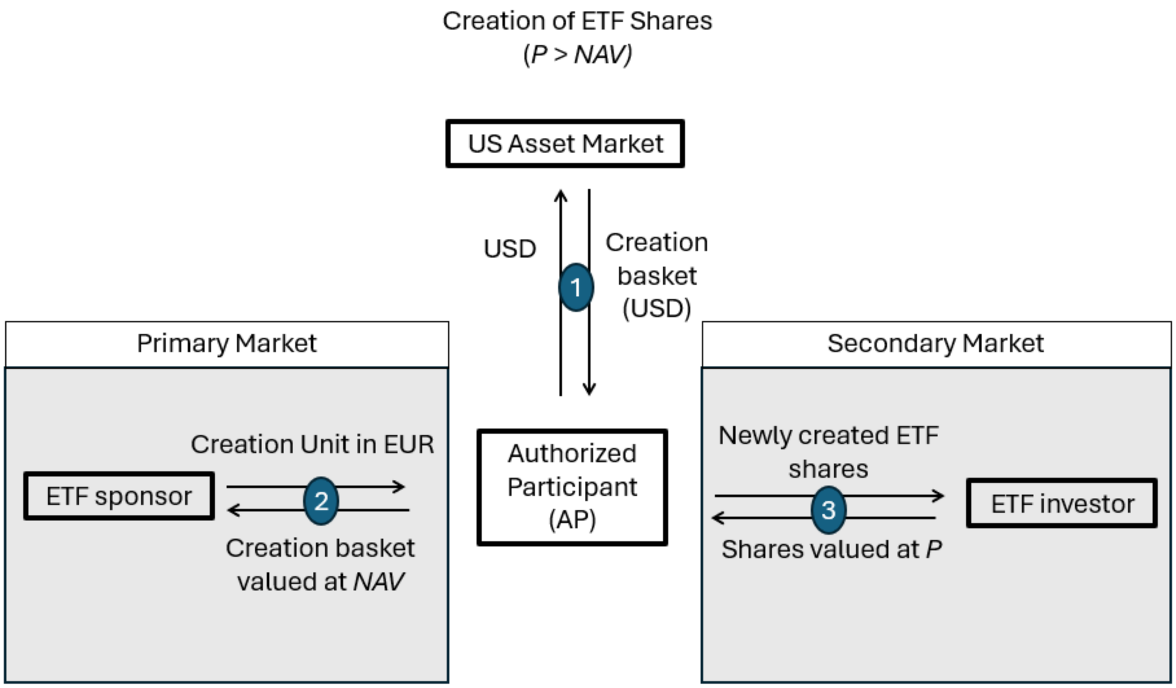

Notes: The figure shows the arbitrage mechanism for ETFs when the ETF trades at a premium and the Authorized Participant (AP) creates new shares.

Figure A.3: Time-series of relative Mispricing of ETFs

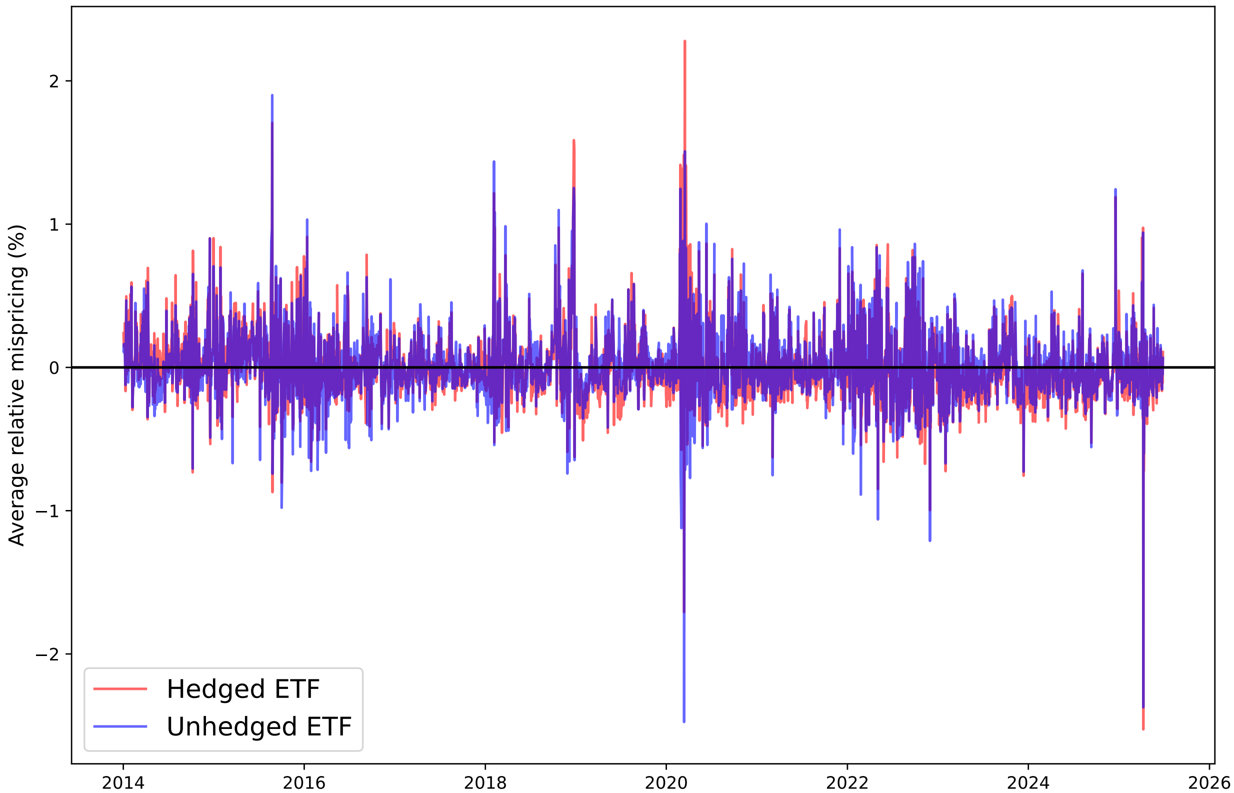

Notes: The figure plots the relative mispricing averaged across hedged ETFs (red) and unhedged ETFs (blue). Relative mispricing is defined as the difference between the last trading price of the ETF and the net asset value (NAV) divided by the NAV, i.e., $\epsilon = \frac{P_t - NAV_t}{NAV_t} \times 100$.

Figure A.4: Dispersion of PF-implied Expectations and FX volatility USD basket

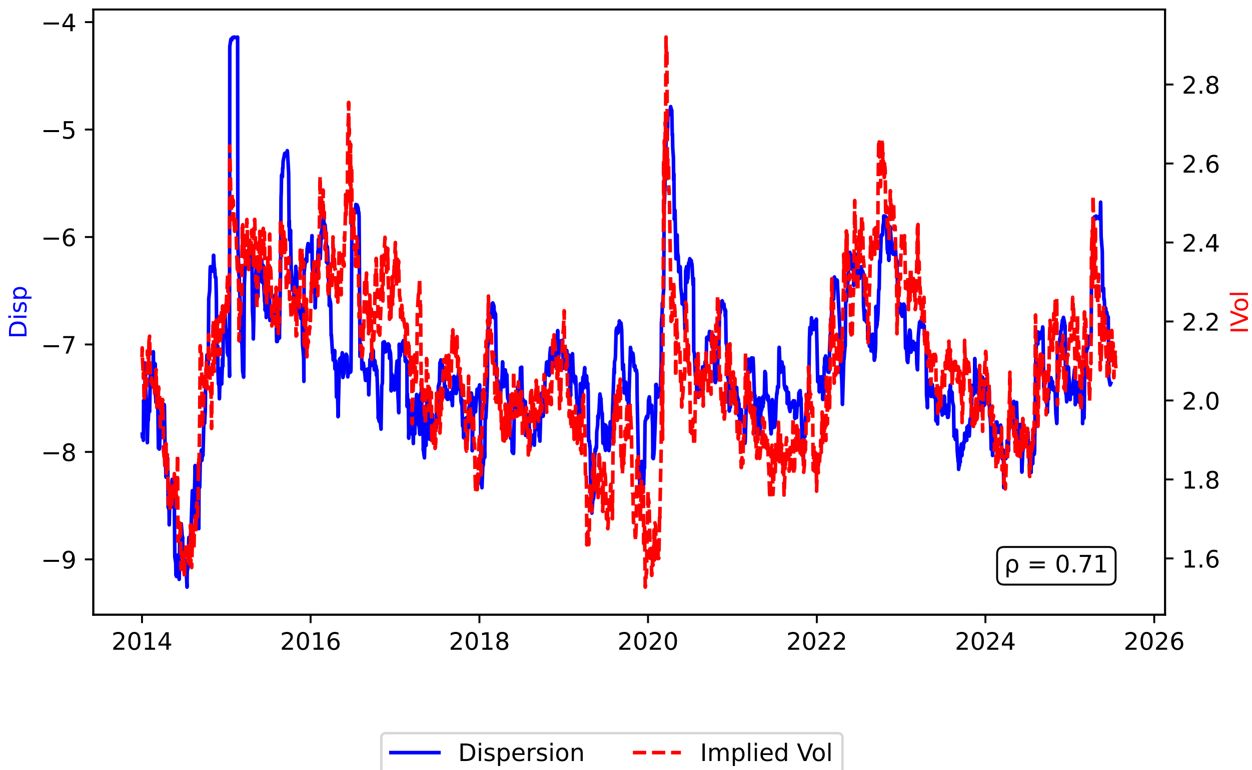

Notes: The figure shows the mean absolute deviation of portfolio-implied expectations defined in Eq. (7) for each day across all funds in all currency areas in blue (left axis) against the log of the average option-implied FX volatility in red (right axis).

B Estimation Details for Expected Returns

This section details the estimation techniques and results for the construction of macro and factor expectations.

B.1 Macro Expectations

The macro expectations, $\mathbb{E}_t^{macro}[rx_{c,t+1}]$, are estimated by following Chernov et al. (2023). I present the key equations for estimating expected currency excess returns and refer to the original paper for further details, in particular the underlying motivation behind each step. For comparability with Chernov et al. (2023), I define exchange rates as USD per unit of foreign currency and estimate expected excess currency returns in percentage changes in this section. The expected returns used in the main paper correspond to the negative of these estimates. Furthermore, the sample covers all G10 currencies: AUD, CAD, EUR (or DEM before 1999), JPY, NZD, NOK, SEK, CHF, GBP and USD and starts in 1990. I source consumer price indices from Datastream and, for AUD and NZD, forward-fill them to obtain monthly observations.

First, I compute an estimate of the mean-reversion component, z_{c,Q_t}, using the real exchange rate (RER)

$$Q_{c,t} = S_{c,t}P_{c,t}/P_t$$

where $P_{c,t}$ and P_t are the foreign and US consumer price index (CPI), respectively. In a next step, the RER is divided by its five-year smoothed lag (the average RER from 4.5 and 5.5 years ago)

$$\tilde{Q}_{c,t} = Q_{c,t}\left(\frac{1}{13}\sum_{j=-6}^{6} Q_{c,t-60+j}\right)^{-1}$$

and then cross-sectionally demeaned

$$z_{c,Q_t} = \tilde{Q}_{c,t} - \frac{1}{N}\sum^{i=1}\tilde{Q}_{c,t}.$$

Second, I run monthly panel regressions of realized spot exchange rate changes on the real exchange rate signal and the trend signal:

$$S_{c,t+1}/S_{c,t} - 1 = \overline{\delta_t} z_{c,Q_t} + \overline{\phi_t}(S_{c,t}/S_{c,t-12} - 1) + \epsilon_{c,t+1} \tag{1}$$

in an expanding window setup month by month. To get to expected excess returns, I multiply the coefficients by the forward discount

$$\begin{aligned}
\mathbb{E}_t^{macro}[rx_{c,t+1}] &= (S_{c,t}/F_{c,t})\, \mathbb{E}_t^{macro}[S_{c,t+1}/S_{c,t}] - 1 \\
\mathbb{E}_t^{macro}[rx_{c,t+1}] &= (S_{c,t}/F_{c,t} - 1) + (S_{c,t}/F_{c,t})\overline{\delta_t} z_{c,Q_t} + (S_{c,t}/F_{c,t})\overline{\phi_t}(S_{c,t}/S_{c,t-12} - 1)
\end{aligned}$$

where $(S_{c,t}/F_{c,t})\overline{\delta_t}$ and $(S_{c,t}/F_{c,t})\overline{\phi_t}$ correspond to $\delta_{c,t}$ and $\phi_{c,t}$ in Eq. (9).

Figure B.1 replicates Figure 1 of Chernov et al. (2023) and plots the standardized values of cross-sectional averages of the annualized estimated coefficients of the RER, $\delta_{c,t}$, and trend component, $\phi_{c,t}$, together with the standardized unit coefficient of the forward discount, $\gamma_{c,t}$. The figure shows that the estimated coefficients are very similar to those reported in Chernov et al. (2023); small differences arise at the beginning of the sample (which is not used in the main paper) since they start their estimation already from 1973. The coefficients stabilize by the time the sample period of the main paper begins in 2014. By then, the average annualized coefficients across the seven currency pairs used in the main paper are approximately -0.173 for the RER coefficient, 0.047 for the trend coefficient, and 0.026 for the forward discount.

Figure B.1: Loadings on signals in expected excess currency returns.

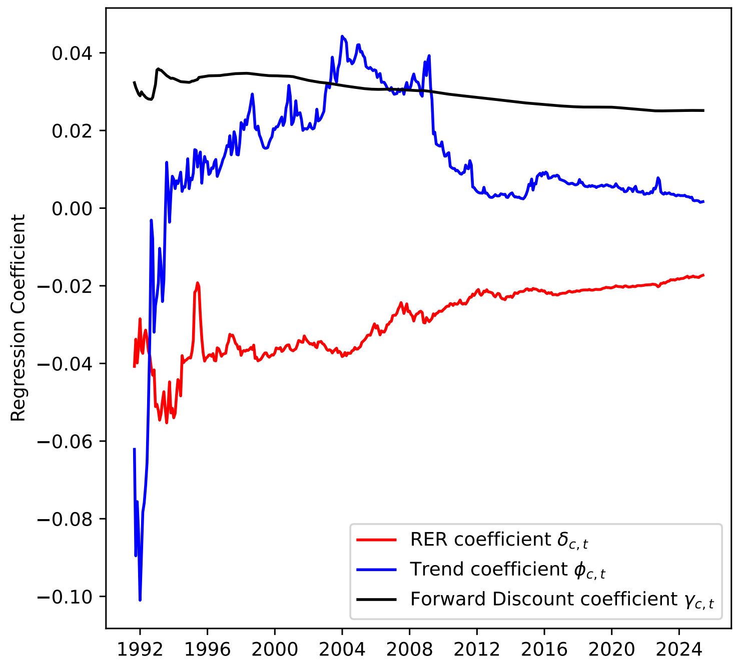

Notes: The figure shows the standardized values of cross-sectional averages of the annualized estimated coefficients of the RER, $\delta_{c,t}$, and trend component, $\phi_{c,t}$, together with the standardized unit coefficient of the forward discount, $\gamma_{c,t}$. The sample period spans from 1990 to 2025.

B.2 Factor Expectations

For the construction of factor expectations, $\mathbb{E}_t^{factor}[rx_{c,t+1}]$, I proceed in two steps following Opie and Riddiough (2020). I present the key equations for estimating expected currency excess returns and refer to the original paper for further details, in particular the underlying motivation behind each step. As with macro expectation, I define exchange rates as USD per unit of foreign currency. The expected returns used in the main paper correspond to the negative of these estimates. The sample covers all G10 currencies: AUD, CAD, EUR (or DEM before 1999), JPY, NZD, NOK, SEK, CHF, GBP and USD and starts in 1990.

First, I estimate rolling betas of currency excess returns with respect to the carry and dollar factors. The carry factor corresponds to a strategy that goes long high-interest-rate currencies and short low-interest-rate currencies. More precisely, a rank weight is given to each currency as

$$\omega_{c,t} = \text{rank}(s_{c,t} - f_{c,t}) - N^{-1} \sum_{c=1}^{N} \text{rank}(s_{c,t} - f_{c,t}) \tag{2}$$

where $s_{c,t} - f_{c,t}$ is the forward discount and $N = 9$. With nine currencies versus the USD, the possible weight values are $+0.4, +0.3, +0.2, +0.1, 0, -0.1, -0.2, -0.3$ and -0.4. The carry factor then equals $R^{pf}_{t+1} = \sum_{c=1}^{N} \omega_{c,t} R_{c,t+1}$. The dollar factor is defined as the average return of a basket of currencies against the USD. To obtain time-varying loadings, $\hat{\beta}_{c,t}$, I regress excess returns, $rx_{c,t} = s_t - f_{t-1}$, on the two factors using a 60-month rolling window. The estimated carry and dollar beta for the currencies used in the main paper are plotted in Panel A and B in Figure B.2, respectively.

Second, I forecast the factor prices of risk, λ_t, using predictor variables, X_{t-1}, in a rolling window regression. The predictors include the average forward discount, foreign exchange volatility, $\Delta\sigma_t^{FX}$, and commodity returns from the Commodity Research Bureau's raw industrials index, constructed as $\Delta CRB_t = \frac{1}{3} log \left(\frac{CRB_t}{CRB_{t-3}} \right)$, together with a constant. This yields expected factor prices of risk, $\mathbb{E}_t[\lambda^{carry}_{t+1}]$ and $\mathbb{E}_t[\lambda^{dollar}_{t+1}]$, that are shown in Panel C and Panel D of Figure B.2, respectively.

To summarize, first, run

$$rx_{c,t} = \beta^{carry}_{c,t} \lambda^{carry}_t + \beta^{dollar}_{c,t} \lambda^{dollar}_t + \epsilon_{c,t}$$

and get $\hat{\beta}^{carry}_{c,t}$ and $\hat{\beta}^{dollar}_{c,t}$. Then run

$$\lambda^{carry}_t = \alpha^{carry}_t + \phi^{carry}_t X_{t-1} + \epsilon^{carry}_t$$
$$\lambda^{dollar}_t = \alpha^{dollar}_t + \phi^{dollar}_t X_{t-1} + \epsilon^{dollar}_t$$

and get predicted values $\hat{\lambda}^{carry}_{t+1}$ and $\hat{\lambda}^{dollar}_{t+1}$ as estimates of the conditional expectations, $\mathbb{E}_t[\lambda^{carry}_{t+1}]$ and $\mathbb{E}_t[\lambda^{dollar}_{t+1}]$, respectively.

Combining the estimated betas with the expected factor prices of risk, factor expectations of currency excess returns are given by

$$\mathbb{E}^{factor}_t [rx_{c,t+1}] = \hat{\beta}^{carry}_{c,t} \hat{\lambda}^{carry}_{t+1} + \hat{\beta}^{dollar}_{c,t} \hat{\lambda}^{dollar}_{t+1}$$

as in Eq. (10).

Figure B.2: Factor Loadings and Expected Prices of Risk for Dollar and Carry Factors

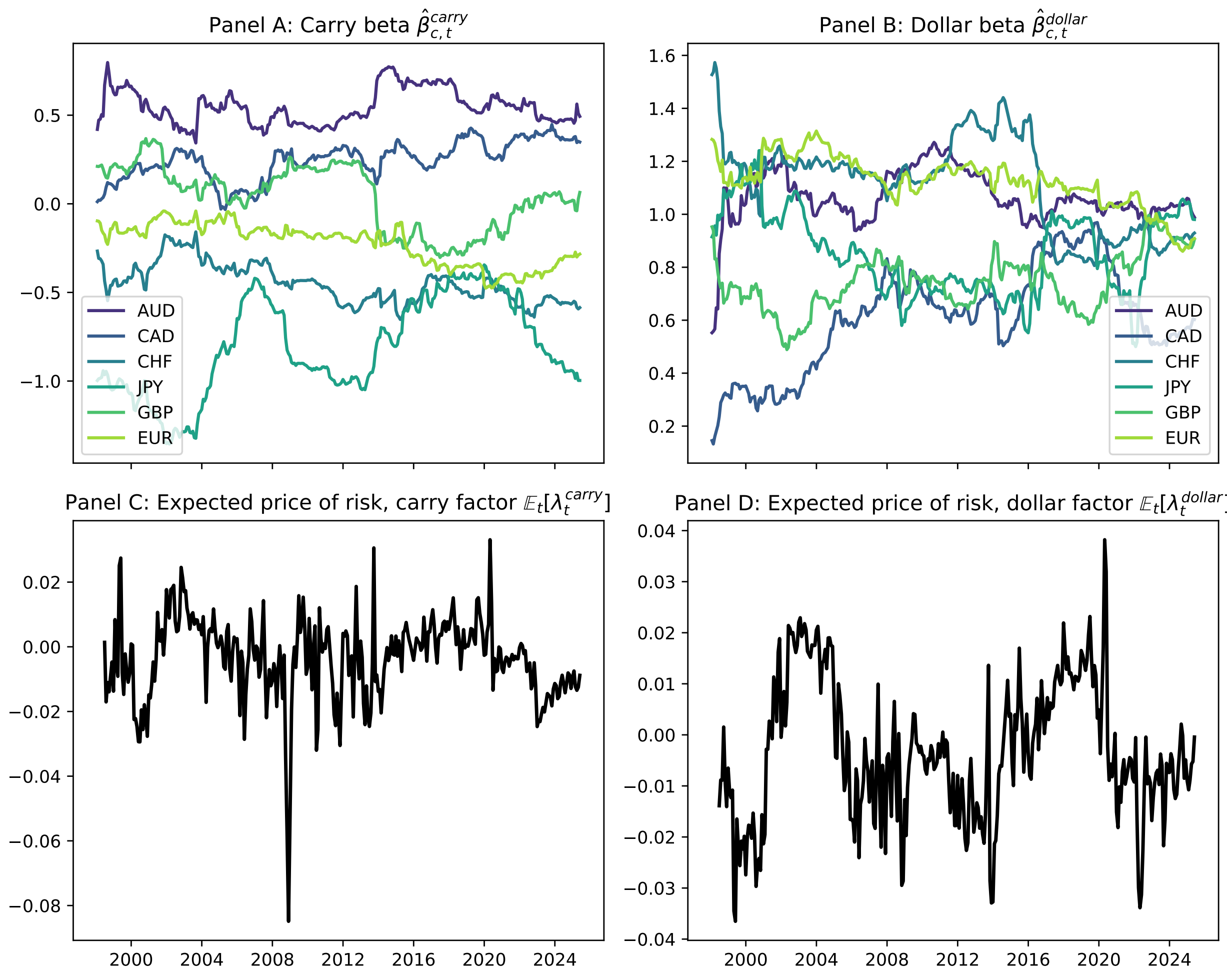

Notes: The figure reports rolling estimates of currency betas with respect to the carry and dollar factors for each currency used in the main paper (Panels A and B) and the expected factor prices of risk (Panels C and D). Betas are estimated from 60-month rolling regressions of currency excess returns on the carry and dollar factors. Expected factor prices of risk are obtained from predictive regressions using forward discounts, FX volatility, and commodity returns as predictors. The sample period spans 1990–2025.

C Trading Strategies Built on ETF Signals

This section reports statistics for several monthly cross-sectional and time-series strategy returns. I replicate all major factor trading strategies outlined in Chernov et al. (2023) and additionally add strategies based on portfolio-implied expectations and active flows into the USD obtained from trading in international ETFs. The portfolio excess return of a trading strategy is given by $R^{pf}_{t+1} = \sum_{c=1}^{N} \omega_{c,t} R_{c,t+1}$ where $\omega_{c,t}$ is the weight in currency c at time t and N are the number of currencies.

I construct cross-sectional (CS) strategies with rank weights

$$\omega_{c,t} = \text{rank}(z_{c,t}) - N^{-1} \sum_{c=1}^{N} \text{rank}(z_{c,t}) \qquad \text{(C1)}$$

where $z_{c,t}$ is the signal. The net exposure to the USD in cross-sectional strategies is thus zero. Weights for time-series (TS) strategies are equal to $+1$ or -1, depending on the sign of the signal. The net exposure to the USD in time-series strategies can be positive or negative. Returns and rebalancing is monthly. Table C.1 defines all strategies with the corresponding signals. The strategies based on ETF portfolio-implied expectations (denoted *ETF-Exp*) and active ETF flows (denoted *ETF-Flow*) are novel and derived from the mechanism analyzed in the main paper. All other strategies are standard.

Table C.2 presents average annualized returns in percentage, t-statistics (of non-zero average returns), annualized Sharpe ratios (SR) and annualized Sharpe ratios including transaction costs τ from all trading strategies. Transaction costs are equal to the total rebalancing times the average round-trip cost per trade which I assume to be constant and equal to 0.7 bps (see Bräuer and Hau (2024)). In Columns (1)-(4) the main paper's time period, 2014–2025, is used, in Columns (5)-(8) the period starts in 1990, for which no data exist for the ETF trading strategies.

I note three main observations. First, for a given portfolio construction (cross-sectional, time-series, dollar) ETF strategies based on portfolio-implied expectations and on flows yield comparable performances, with the largest difference for the ETF-dollar strategy. Second, ETF strategies achieve an average excess return of 3.25%, substantially exceeding that of traditional strategies. Among traditional strategies, the CS-Carry strategy also yields high average returns of 1.64%.

However, only CS-ETF strategies have an average excess return statistically different from zero. Third, the Sharpe ratios are highest for the cross-sectional ETF trading strategies, and reach up to 0.58 (or 0.56 including transaction costs) for the strategy based on portfolio-implied ETF expectations. The highest Sharpe Ratio in the full sample is obtained by the TS-Carry strategy (of 0.46) as in Chernov et al. (2023).

Table C.1: Currency Trading Strategy Definitions

CS-ETF-Exp	Goes long (short) currencies with relatively low (high) portfolio-implied expectations, i.e., it buys (sells) the USD whenever there are relatively high (low) portfolio-implied expectations in the previous month.
TS-ETF-Exp	Goes long (short) currencies with negative (positive) portfolio-implied expectations, i.e., it buys (sells) the USD whenever there are positive (negative) portfolio-implied expectations in the previous month.
Dollar-ETF-Exp	Goes (short) long all currencies versus the USD, whenever average portfolio-implied expectations are negative (positive).
CS-ETF-Flow	Goes long (short) currencies with relatively low (high) active flows into the USD, i.e., it buys (sells) the USD whenever there are relatively high (low) active flows into the USD (unhedged vs. hedged share class) in the previous month.
TS-ETF-Flow	Goes long (short) currencies with negative (positive) active flows, i.e., it buys (sells) the USD whenever there are positive (negative) active flows into the USD (unhedged vs. hedged share class) in the previous month.
Dollar-ETF-Flow	Goes (short) long all currencies versus the USD, whenever average active flows are negative (positive).
Dollar	Goes long all currencies versus the USD.
Dollar-Carry	Goes (short) long all currencies versus the USD, whenever the average forward discount is positive (negative).
CS-Carry	Goes long (short) currencies with a relatively high (low) forward discount.
TS-Carry	Goes long (short) currencies with a positive (negative) forward discount.
CS-Mom 1	Goes long (short) currencies that have performed relatively well in the past month.
CS-Mom 12	Goes long (short) currencies that have performed relatively well in the past year.
TS-Mom 1	Goes long (short) currencies that had positive (negative) returns in the past months.
TS-Mom 12	Goes long (short) currencies that had positive (negative) returns in the past 12 months.

Table C.2: Currency Trading Strategy Statistics

The table presents average annualized returns (%), t-statistics (of non-zero average returns), annualized Sharpe ratios (SR) and annualized Sharpe ratios including transaction costs τ from trading strategies outlined in Table C.1. Columns (1)-(4) report results for the sample from January 2014 to May 2025, Columns (5)-(8) report results for the sample from January 1990 to May 2025. For the latter sample no results are reported for the ETF strategies as data start only in 2014. t-statistics test whether the mean monthly return equals zero, using Newey–West (lag 3) heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. The sample includes all currencies used in the main paper: AUD, GBP, EUR, CHF, JPY, CAD against the USD.

	Period 2014-2025				Period 1990-2025			
	Avg. Return (%)	t-statistics	SR	SR (incl. τ)	Avg. Return (%)	t-statistics	SR	SR (incl. τ)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
CS-ETF-Exp	3.25	2.01	0.58	0.56	–	–	–	–
TS-ETF-Exp	0.89	0.73	0.20	0.17	–	–	–	–
Dollar-ETF-Exp	0.12	0.07	0.02	−0.00	–	–	–	–
CS-ETF-Flow	3.03	1.90	0.53	0.50	–	–	–	–
TS-ETF-Flow	1.13	0.95	0.27	0.22	–	–	–	–
Dollar-ETF-Flow	1.51	0.75	0.23	0.20	–	–	–	–
Dollar	−2.34	−1.22	−0.36	−0.36	0.82	0.60	0.10	0.10
Dollar-Carry	−0.30	−0.15	−0.05	−0.05	2.64	2.02	0.33	0.33
CS-Carry	1.64	1.03	0.24	0.24	2.96	2.15	0.34	0.34
TS-Carry	0.75	0.51	0.14	0.14	2.03	2.94	0.46	0.46
CS-Mom 1	−2.10	−1.07	−0.33	−0.37	0.48	0.35	0.06	0.04
CS-Mom 12	−0.34	−0.20	−0.05	−0.06	0.37	0.32	0.05	0.04
TS-Mom 1	−1.07	−0.78	−0.21	−0.26	1.36	1.72	0.27	0.23
TS-Mom 12	−0.69	−0.44	−0.14	−0.15	0.60	0.78	0.13	0.12

D Heterogeneous Beliefs in Asset Pricing

This section briefly outlines the main mechanism by which heterogeneous beliefs in complete markets can affect exchange rate volatility, built on Back (2010).

General Setup. A model with heterogeneous beliefs can be re-expressed to one with homogeneous beliefs and state-dependent utility. This can be done by defining some reference probability, for example, the average belief as

$$\mathbb{P}(\omega) = \frac{1}{N}\sum_{i=1}^{N}\mathbb{P}_i(\omega) \tag{C1}$$

over N investors for a finite number of states ω where $\mathbb{P}_i$ and $\mathbb{P}$ are two σ-finite measures defined on a measurable space $(X, \mathcal{F}_t)$. The Radon-Nikodym theorem states that if $\mathbb{P}_i$ is absolutely continuous with respect to $\mathbb{P}_i$ ($\mathbb{P}_i \ll \mathbb{P}$), there exists a nonnegative random variable $\Lambda_{i,t}$, such that for any measurable set $A \in \mathcal{F}_t$

$$\mathbb{P}_i(A) = \int_A \Lambda_{i,t} d\mathbb{P} \text{ over all } A \in \mathcal{F}_t \tag{C2}$$

where $\Lambda_{i,t}$ is a strictly positive martingale process and is called the Radon-Nikodym derivative of $\mathbb{P}_i$ with respect to $\mathbb{P}$. It can be written as

$$\Lambda_{i,t} = \frac{d\mathbb{P}_i}{d\mathbb{P}}\bigg|\mathcal{F}_t. \tag{C3}$$

The change of measure, $\Lambda_{i,t}$, transforms the belief of any investor i about a random variable x as

$$\mathbb{E}_t^{\mathbb{P}_i}[x_{t+1}] = \mathbb{E}_t^{\mathbb{P}}\left[\frac{\Lambda_{i,t+1}}{\Lambda_{i,t}}x_{t+1}\right]. \tag{C4}$$

with the one-step likelihood ratio defined as $\xi_{i,t+1} = \frac{\Lambda_{i,t+1}}{\Lambda_{i,t}}$.

In a consumption-based asset pricing model, heterogeneous beliefs thus transform the stochastic discount factor (SDF) to

$$m_{i,t+1} = \kappa_t\beta\frac{u_i'(C_{i,t+1})}{u_i'(C_{i,t})}\frac{\Lambda_{i,t+1}}{\Lambda_{i,t}} \tag{C5}$$

where κ_t is a scaling factor, β is the discount factor, $u_i'(C_{i,t})$ is investor i's marginal utility of consumption $C_{i,t}$.

Example with Log Utility. With log-utility and aggregate consumption $\sum_i C_{i,t} = C_t$, the first-order condition of the social planner's problem yields equilibrium consumption shares that satisfy

$$C_{i,t} = C_t \left(\frac{\lambda_i \Lambda_{i,t}}{\sum_i \lambda_i \Lambda_{i,t}} \right) \tag{C6}$$

where λ_i denotes pareto weights assigned by the planner. The SDF then equals

$$m_{t+1} = \beta \frac{C_t \left(\frac{\lambda_i \Lambda_{i,t}}{\sum_i \lambda_i \Lambda_{i,t}} \right)}{C_{t+1} \left(\frac{\lambda_i \Lambda_{i,t+1}}{\sum_i \lambda_i \Lambda_{i,t+1}} \right)} \frac{\Lambda_{i,t+1}}{\Lambda_{i,t}} = \beta \frac{C_t}{C_{t+1}} \frac{BI_{t+1}}{BI_t} \tag{C7}$$

where $BI_t = \sum_i \lambda_i \Lambda_{i,t}$ is the wealth-weighted average belief index (BI). Thus, compared to a model of homogenous beliefs, the SDF also depends on the time variation in wealth-weighted average beliefs $\frac{BI_{t+1}}{BI_t}$.

Dispersion in Beliefs. When investors disagree more about probabilities of future states, the wealth-weighted belief index, $\frac{BI_{t+1}}{BI_t}$, features greater variation. To see this, define

$$\frac{BI_{t+1}}{BI_t} = \sum_i \alpha_{i,t} \xi_{i,t+1} \quad \text{with } \alpha_{i,t} = \frac{\lambda_i \Lambda_{i,t}}{\sum_k \lambda_k \Lambda_{k,t}} \quad \text{and } \xi_{i,t+1} = \frac{\Lambda_{i,t+1}}{\Lambda_{i,t}} \tag{C8}$$

then, assuming a discrete state space, its conditional variance under $\mathbb{P}$ is given by

$$\mathbb{V}\text{ar}_t^{\mathbb{P}} \left[\frac{BI_{t+1}}{BI_t} \right] = \mathbb{V}\text{ar}_t^{\mathbb{P}} \left[\sum_i \alpha_{i,t} \xi_{i,t+1} \right] \tag{C9}$$

$$= \sum_s p(s) \left(\frac{\sum_i \alpha_{i,t} p_i(s)}{p(s)} - 1 \right)^2 \tag{C10}$$

$$= \sum_s \frac{\left(\sum_i \alpha_{i,t} p_i(s) - p(s) \right)^2}{p(s)} \tag{C11}$$

where $p(s)$ and $p_i(s)$ are the time-t conditional probabilities under $\mathbb{P}$ and $\mathbb{P}_i$, respectively, such that $\xi_{i,t+1}(s) = \frac{p_i(s)}{p(s)}$. The second line follows from the martingale property of the Radon-Nikodym derivative under $\mathbb{P}$, i.e., $\mathbb{E}_t^{\mathbb{P}}[\sum_i \alpha_{i,t} \xi_{i,t+1}] = \mathbb{E}_t^{\mathbb{P}} \left[\frac{\sum_i \alpha_{i,t} p_i}{p} \right] = 1$ since $\sum_i \alpha_{i,t} = 1$. Thus, the larger the difference between the wealth-weighted mean beliefs (probabilities), $\sum_i \alpha_{i,t} p_i(s)$, from the reference probability, $p(s)$, the higher the variance of $\mathbb{V}\text{ar}_t \left[BI_{t+1}/BI_t \right]$.

As an example, consider a two-agent two-state model: an optimist and a pessimist who hold different beliefs about the parameters of a random variable that follows a Bernoulli distribution.

Specifically, the optimist assigns a probability of $p^{opti}(u) \in [0.1, 1]$ to the up state while the pessimist always assigns a probability of $p^{pessi}(u) = 0.1$ to the up state. The equilibrium share of the optimist takes a value of $\lambda^{opti} \in \{0.2, 0.5, 1\}$ with the pessimist taking $\lambda^{pessi} = 1 - \lambda^{opti}$. The reference probability is the equal-weighted average belief as defined in Eq. (C1). Figure D.1 plots the difference in beliefs (probabilities) between the agents on the horizontal axis against the variance in the belief index on the vertical axis, computed as[22]

$$\mathbb{V}\mathrm{ar}\left[\frac{BI_{t+1}}{BI_t}\right] = \sum_{s\in\{u,d\}} p(s)\left(\lambda^{opti}\frac{p^{opti}(s)}{p(s)} + \lambda^{pessi}\frac{p^{pessi}(s)}{p(s)} - 1\right)^2. \tag{C12}$$

The figure shows that as the belief gap, $|p^{opti} - p^{pessi}|$, widens, the variance of the wealth-weighted belief update increases whenever $\lambda^{opti} \neq \lambda^{pessi}$ (lines marked with diamonds and circles). If agents hold equal shares $\lambda^{opti} = \lambda^{pessi} = 0.5$, the wealth-weighted average belief equals the reference probability and thus any difference in beliefs does not affect the variance of the belief update.

Dispersion in Exchange Rate Expectations. The log change in the nominal exchange rate between the home country and foreign country (*) is equal to the difference in the stochastic discount factors (see, e.g., Backus et al. (2001))

$$\Delta s_{t+1} = log(m_{t+1}) - log(m^*_{t+1}) \tag{C13}$$

where I define an increase in Δs_{t+1} as a USD appreciation and refer to the home country as US. Note that if the SDF is in real terms, one needs to additionally subtract the inflation differential between countries in Eq. (C13).

In the log-utility example before, the exchange rate change then equals

$$\Delta s_{t+1} = [\Delta c^*_{t+1} - \Delta c_{t+1}] + [\Delta bi_{t+1} - \Delta bi^*_{t+1}] \tag{C14}$$

where small letters denote logs and I assume no difference in foreign and home discount factors, i.e., $\beta^* = \beta$. Excluding the belief index (the second term), Eq. (C14) equals the standard Backus-Smith condition (Eq. (4.8) in Backus and Smith (1993)).

[22]Note that in this example $\alpha_{i,t} = \lambda_i$ as I normalize $\Lambda_{i,t} = 1$ at the conditioning date t, such that $\alpha_{i,t} = \frac{\lambda_i \Lambda_{i,t}}{\sum_k \lambda_k \Lambda_{k,t}} = \lambda_i$ as $\sum_k \lambda_k = 1$.

Eq. (C14) shows that under heterogeneous beliefs, the exchange rate depends on the random variables Δbi^*_{t+1} and Δbi_{t+1}, the average belief indices of the home and foreign country. In the main paper, I argue that dispersion in beliefs among foreign investors (here, e.g., euro area ETF investors) varies with the volatility of the USD exchange rate. By Eq. (C11), an increase in dispersion among foreign investors raises $\mathbb{V}\mathrm{ar}_t\left[BI^*_{t+1}/BI^*_t\right]$. Under a first-order log approximation, this also raises $\mathbb{V}\mathrm{ar}_t\left[\Delta bi^*_{t+1}\right]$, and thus increases $\mathbb{V}\mathrm{ar}_t[\Delta s_{t+1}]$ for fixed consumption and a fixed US belief index.

Figure D.1: Difference in Beliefs and Average Belief Index

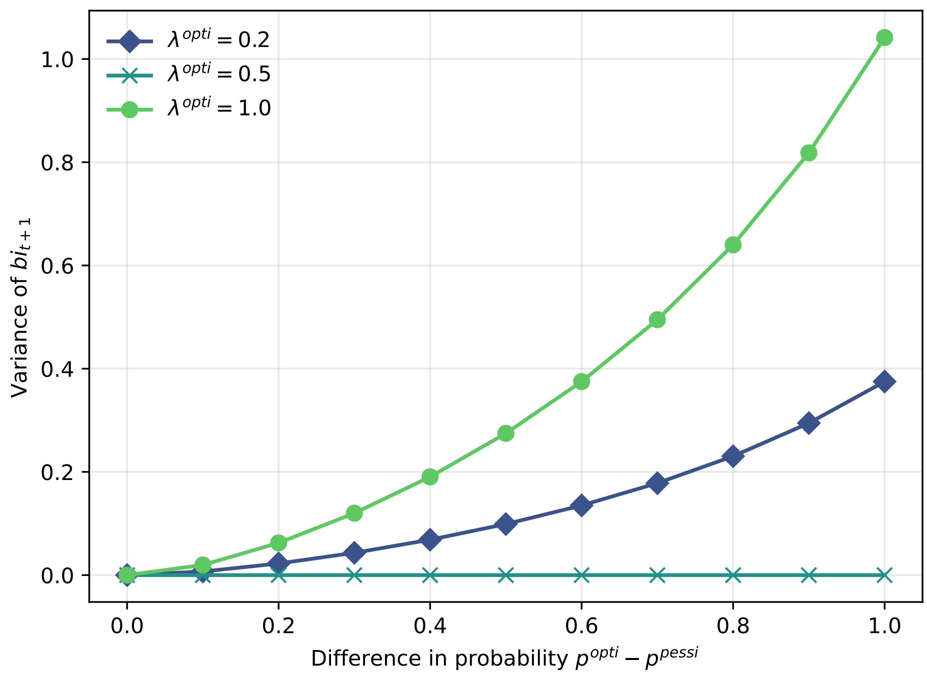

Notes: The figure depicts the difference in beliefs (probabilities) between two agents, an optimist and a pessimist, on the horizontal axis against the variance in the belief index on the vertical axis. The variance of the belief index is computed as in Eq. (C12). I use three different values for the equilibrium share of the optimist $\lambda^{opti} \in \{0.2, 0.5, 1\}$.