The Journal of FINANCE
The Journal of THE AMERICAN FINANCE ASSOCIATION

THE JOURNAL OF FINANCE • VOL. LXXX, NO. 6 • DECEMBER 2025

Long-Horizon Exchange Rate Expectations

LUKAS KREMENS, IAN W. R. MARTIN, and LILIANA VARELA*

ABSTRACT

We study exchange rate expectations in surveys of financial professionals and find that they successfully forecast currency appreciation at the two-year horizon, both in and out of sample. Exchange rate expectations are also interpretable, in the sense that three macro-finance variables—the risk-neutral covariance between the exchange rate and equity market, the real exchange rate, and the current account relative to GDP—explain most of their variation. There is no "secret sauce," however, in expectations: After controlling for the three macro-finance variables, the residual information in survey expectations does not forecast currency appreciation in our sample.

IN A RISK-NEUTRAL WORLD, A currency with a high interest rate would be expected to depreciate against a currency with a low interest rate to equate their expected returns. This is the celebrated prediction of uncovered interest parity (UIP). Empirically, however, a large literature starting with Hansen and

*Lukas Kremens is at the University of Washington. Ian Martin is at the London School of Economics. Liliana Varela is at the London School of Economics and CEPR. We thank Ric Colacito; Charles Engel; Jonathan Lewellen (the Editor); Stefan Nagel; Steve Wu; two anonymous referees; the Associate Editor; and conference participants at the SED Annual Meeting 2023, the NBER Summer Institute 2023, and the Macrofinance Society Workshop 2024 for helpful comments. We are grateful to Oliver Ashtari Tafti for excellent research assistance and to the Systemic Risk Centre at the LSE for their support and for providing access to data. Markit® is a trade name and the property of Markit Group Limited or its affiliate (Markit) and is used by the London School of Economics and Political Science under license. Data provided by Markit®. Nothing in this publication is sponsored, endorsed, sold, or promoted by Markit or its affiliates. Neither Markit nor its affiliates make any representations or warranties, express or implied, to you or any other person regarding the advisability of investing in the financial products described in this report or as to the results obtained from the use of the Markit data. Neither Markit nor any of its affiliates have any obligation or liability in connection with the operation, marketing, trading, or sale of any financial product described in this report or use of the Markit data. Markit and its affiliates shall not be liable (whether in negligence or otherwise) to any person for any error in the Markit data and shall not be under any obligation to advise any person of any error therein. This work was funded by U.K. Research and Innovation (UKRI) under the U.K. government's Horizon Europe funding guarantee (grant number EP/X020916/1). We have read *The Journal of Finance*'s disclosure policy and have no conflicts of interest to disclose.

Correspondence: Lukas Kremens, University of Washington, Seattle, WA; e-mail: lkremens@uw.edu.

DOI: 10.1111/jofi.13504

Hodrick (1980) and Fama (1984) finds that currencies with high interest rates earn higher returns, on average, than currencies with low interest rates.[1]

What explains the failure of UIP, that is, the gap between expected currency appreciation and the interest rate differential? Assuming that frictionless trade in the currencies and interest rates is possible, this gap represents an expected excess return, or risk premium. On the traditional view of international financial markets, this risk premium should reflect the covariation of currency returns with a stochastic discount factor (SDF) whose variation reflects movements in investors' marginal utilities across states.

A recent literature argues that currency markets are profoundly influenced by financial intermediaries who face balance sheet (or other) constraints. On this view, movements in currencies reflect, at least in part, shadow prices on financier constraints, and thus expected currency movements are influenced by variation in these shadow prices and cross-currency flows.

Another branch of the recent literature emphasizes the importance of subjective expectations. In the case of equity markets, for example, Greenwood and Shleifer (2014) argue that investor expectations move in the opposite direction to the forecasts of a rational person, with investors becoming more bullish at times when they should be bearish, and vice versa. In our context, this raises the possibility that realized currency movements do not reflect ex ante expectations. If so, the failure of UIP may simply reflect investor errors. This explanation has a long history: Frankel and Froot (1987) and Froot and Frankel (1989), for example, use survey expectations and find that investors make systematic forecast errors at short horizons.

In this paper, we study expectations drawn from monthly surveys of finance professionals conducted by Consensus Economics and draw two major conclusions. First, survey expectations successfully forecast exchange rate movements over a two-year horizon both in and out of sample. (By contrast, they are considerably less successful in predicting exchange rate appreciation over shorter horizons.) In sample, survey expectations are strongly significant predictors, with an estimated coefficient close to and insignificantly different from one. Figure 1 illustrates the basic finding, plotting realized currency appreciation (RCA; on the vertical axis) against survey expectations (on the horizontal axis) at the 24-month horizon.[2] Realizations are broadly consistent with expectations both across currencies, as indicated by the relative positions of the ellipses, and over time within currency, as indicated by the orientation of individual ellipses. In this sense, survey expectations appear broadly rational. And whereas interest rate differentials alone explain only 3.1% of the variation in realized currency appreciation, interest rate differentials and survey forecasts together explain 16.9% of the variation.

[1] Some papers even find that high-interest currencies *appreciate* on average. In more recent data, Hassan and Mano (2019) find that high-interest currencies depreciate, but not enough to offset interest rate differentials.

[2] Figure 1 does not show standard errors, but these are shown later. Note that Nagel and Xu (2023) instead find that survey forecasts are relatively poor predictors at horizons below one year.

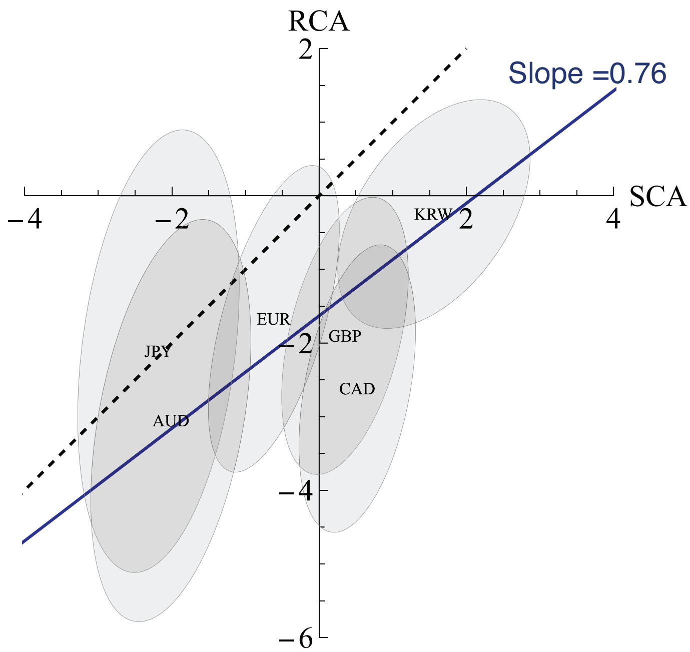

Figure 1. Realized currency appreciation (RCA) versus survey expectations (SCA). For each currency, the figure plots mean RCA against SCA surrounded by a confidence ellipse whose orientation reflects the time-series correlation between RCA and SCA, and whose size reflects their volatilities (scaled to contain 10% of the observations under joint normality). The solid blue line represents a univariate panel regression with a slope coefficient of 0.76, while the dotted line is the 45° line on which realizations equal survey expectations. Six high-income currencies: Australian dollar (AUD), Canadian dollar (CAD), euro (EUR), Great British pound (GBP), Japanese yen (JPY), and Korean won (KRW). (Color figure can be viewed at wileyonlinelibrary.com)

When we compare survey expectations to various predictor variables proposed by the literature—the quanto-implied risk premium (QRP, Kremens and Martin (2019)), which measures the risk-neutral covariance of the exchange rate with the S&P 500 index, the real exchange rate (RER, for example, Dahlquist and Penasse (2022)), the VIX index, the dollar and carry betas of Lustig, Roussanov, and Verdelhan (2011, 2014), interest rate differentials, the current account balance to GDP (CA-to-GDP; see, for example, Gabaix and Maggiori (2015)), capital inflows to GDP, primary balance to GDP, industrial production, and net foreign assets to GDP—we find that survey expectations are the best performing univariate predictor in an R^2 sense.

Out of sample, we find that survey forecasts—which evidently have the advantage of not requiring estimation of free parameters—outperform the random walk benchmark of Meese and Rogoff (1983) both in terms of bilateral exchange rate predictions against the dollar and dollar-neutral relative forecasts of other currencies.

Our second major conclusion is that survey expectations are interpretable, in the sense that they load heavily on a subset of the macro/finance predictor variables above. Specifically, three variables—QRP, RER, and CA-to-GDP—explain more than half of the variation in survey expectations.

It is natural then to wonder whether there is any "secret sauce" in survey expectations. We regress survey expectations onto the three variables and view the residuals—the components of expectations *not* explained by QRP, RER, or CA-to-GDP—as, potentially, the secret sauce. It turns out, however, that these residuals have essentially no predictive power for returns. That is, there is no secret sauce.

Finally, we compare the predictive success of surveys at different horizons, and conclude with a puzzle. The literature shows that survey forecasts are poor predictors of currency appreciation at shorter horizons below a year. We confirm this finding in our sample and show that short-horizon forecasts do not relate to the macro-finance variables associated with predictive success in long-horizon forecasts. This result suggests that forecasters use different models to form expectations at different horizons. Surprisingly, however, we find that *long*-term forecasts do have predictive power for short-term realizations; we find point estimates that are statistically significant in our full sample and close to one when forecasts and realizations are annualized appropriately to make them comparable. The puzzling—bordering on schizophrenic—fact that our forecasters' long-horizon models outperform their own short-horizon models in forecasting short-horizon outcomes appears difficult to explain with information frictions alone (Mankiw and Reis (2002), Coibion and Gorodnichenko (2015)).

Literature review: The Consensus Economics surveys have been used in various other studies in international finance and asset pricing (Stavrakeva and Tang (2023), De Marco, Macchiavelli, and Valchev (2021), Kalemli-Özcan and Varela (2024), Lloyd and Marin (2021), Pesch, Piatti, and Whelan (2024), Bartram et al. (2025)). Candian and De Leo (2025) use these forecasts to estimate a model of under- and overreaction to interest rates that matches the observed reversal of UIP deviations over longer horizons.[3] Della Corte, Gao, and Jeanneret (2023) use the relationship between expectations and quanto-implied risk premia to estimate risk-aversion parameters at different horizons. In contrast, we study both the information that is and is not shared between quantos and surveys, and we assess the predictive power of each component for realized currency returns. In contemporaneous and independent work, Beckmann and Reitz (2023) also find that survey expectations correlate strongly with the quanto-implied risk premium and argue that the quanto-implied risk premium proxies for intermediary capital ratios.

The predictor variables that we use as competitors for survey forecasts come from a number of papers. RER tracks trends in nominal exchange rates as well

[3] See, for example, Froot and Thaler (1990), Bacchetta and van Wincoop (2010, 2021), Engel (2016), Valchev (2020) for evidence and explanations of UIP-reversal and related patterns frequently tied to underreaction and overshooting.

as inflation differentials and has often been linked to currency excess returns (e.g., Asness, Moskowitz, and Pedersen (2013), Koijen et al. (2018), Dahlquist and Penasse (2022)). QRP measures the exposure of currencies to equity-market risk (as also studied by Campbell, Serfaty-De Medeiros, and Viceira (2010), Lettau, Maggiori, and Weber (2014), Cenedese et al. (2016), Kremens (2024)), and the economics broadly resemble arguments rooted in consumption risk (e.g., Lustig and Verdelhan (2007), Verdelhan (2010), Burnside (2011)). The current account balance has been associated with (expected) exchange rate movements (e.g., Kouri (1976), Dornbusch (1976), Gourinchas and Rey (2007)), and cross-border flows with constraints of global financial intermediaries (e.g., Gabaix and Maggiori (2015) and Bianchi, Bigio, and Engel (2023)).

Structure of the paper: Section I describes the data on survey expectations and macro-finance variables. Section II tests the predictive power of long-horizon survey expectations in and out of sample. Section III examines the relationship between survey expectations and various macro-finance variables and interprets our results in terms of two prominent views of excess returns, one based on risk preferences and one based on intermediation constraints. Section IV contrasts the evidence for predictability at short and long horizons. Section V concludes.

I. Data and Definitions

Our sample includes six high-income currencies (Australian dollar, Canadian dollar, euro, Great British pound, Japanese yen, and Korean won) against the U.S. dollar. We observe survey expectations from Consensus Economics, which provides the mean across forecasters of expected exchange rates at one-, three-, 12-, and 24-month horizons from 1994 to 2021. The forecasters who are interviewed principally comprise global banks and investors who actively participate in the foreign exchange (FX) market. We obtain forward discounts from Reuters and use the terms "forward discount" and "interest rate differential" interchangeably. Accordingly, these interest rate differentials are consistent with derivatives prices so do not violate covered interest parity (CIP). We extend the quanto-implied risk premium of Kremens and Martin (2019) until March 2019 using quanto data from Markit. (See Section II.B for more details.) We use the 30-day S&P implied volatility index (VIX) reported by Federal Reserve Economic Data (FRED) to proxy for global risk perception. We construct the dollar carry factor and its loadings, $\beta^{\$}$, following Lustig, Roussanov, and Verdelhan (2014), and the high-minus-low factor (HML) and its loadings, β^{HML}, following Lustig, Roussanov, and Verdelhan (2011). We use various measures of cross-country flows, including the current account balance and capital inflows, both obtained from International Financial Statistics (IFS) of the International Monetary Fund (IMF), which we scale by GDP. Capital inflows are constructed from total debt inflows as the sum of direct investment, portfolio investment, and other investment. We also employ net foreign asset positions over GDP from Lane and Milesi-Ferreti (2018). We obtain RER from the Bank for International Settlements.

As the quanto data from Markit are only reported since December 2009, our baseline specification spans forecasts from December 2009 to March 2019 (with realizations until March 2021). We conduct parallel tests for a longer sample starting in December 1994, wherever the quanto data are not needed. Table IA.I in the Internet Appendix summarizes the data sources.[4]

To set up some notation, let M_{t+h} denote the h-period SDF, which prices payoffs denominated in U.S. dollars, and let $R^{\$}_{f,t,h}$ denote the U.S. riskless rate. The fundamental asset pricing equation states that for any h-period gross dollar return R_{t+h}, we have

$$\mathbb{E}_t \left(M_{t+h}R_{t+h}\right) = 1, \tag{1}$$

or, equivalently,

$$\mathbb{E}_t\, R_{t+h} - R^{\$}_{f,t,h} = R^{\$}_{f,t,h}\, \mathrm{cov}_t \left(-M_{t+h}, R_{t+h}\right). \tag{2}$$

We are interested in the return on a currency trade that converts one U.S. dollar to foreign currency i at time t, invests at the gross h-period riskless rate in currency i, $R^{i}_{f,t,h}$, and then converts back to U.S. dollars at time $t+h$. This is a dollar-denominated trading strategy: Starting from one dollar at time t, it returns $R_{t+h} = R^{i}_{f,t,h}e_{i,t+h}/e_{i,t}$ dollars at time $t+h$, where $e_{i,t}$ is the nominal exchange rate expressed in U.S. dollars per unit of currency i. Substituting this return into equation (2) and rearranging, we have

$$\mathbb{E}_t\, \frac{e_{i,t+h}}{e_{i,t}} - 1 = \underbrace{\frac{R^{\$}_{f,t,h}}{R^{i}_{f,t,h}} - 1}_{\text{UIP}} + \underbrace{R^{\$}_{f,t,h}\, \mathrm{cov}_t \left(-M_{t+h}, \frac{e_{i,t+h}}{e_{i,t}}\right)}_{\text{residual}}. \tag{3}$$

This identity expresses the (net) exchange rate appreciation of currency i in terms of the (net) interest rate differential and a covariance term that captures the risk premium associated with currency i. If the risk premium adjustment is ignored, equation (3) reduces to the traditional prediction of UIP.

Based on identity (3), we define the interest rate differential (IRD) and realized currency appreciation (RCA) at the h-month horizon as

$$IRD_{i,t,h} = \frac{R^{\$}_{f,t,h}}{R^{i}_{f,t,h}} - 1, \tag{4}$$

$$RCA_{i,t,h} = \frac{e_{i,t+h}}{e_{i,t}} - 1. \tag{5}$$

Note that IRD is negative for currencies with high interest rates,[5] for which UIP predicts depreciation. We also define the realized currency excess return

[4] The Internet Appendix is available in the online version of the article on *The Journal of Finance* website.

[5] High relative to the dollar, because we use the dollar as the base currency.

(RXR) as

$$RXR_{i,t,h} = RCA_{i,t,h} - IRD_{i,t,h}. \tag{6}$$

This quantity is an excess return because it has zero price: We can write $RXR_{i,t,h} = \frac{1}{R^i_{f,t,h}}\left(R_{t+h} - R^{\$}_{f,t,h}\right)$ (note that an excess return scaled by a constant is still an excess return).

Analogously, we define survey-based expectations of currency appreciation (SCA) and of currency excess returns (SXR) as

$$SCA_{i,t,h} = \tilde{\mathbb{E}}_t \frac{e_{i,t+h}}{e_{i,t}} - 1, \tag{7}$$

$$SXR_{i,t,h} = SCA_{i,t,h} - IRD_{i,t,h}, \tag{8}$$

where $\tilde{\mathbb{E}}$ denotes the survey consensus expectations operator, computed as a simple average of individual forecasters' reported expectations.

Figure 2 plots the time series of three- and 24-month excess return expectations (SXR) by currency. (For comparison, UIP asserts that every currency should have zero expected excess return. Figure 2 therefore shows that survey expectations deviate from UIP.) Table IA.II reports summary statistics over the post-Global Financial Crisis (GFC) sample.

II. Surveys and Exchange Rate Predictability

Do survey expectations predict exchange rates? We conduct an in-sample test of exchange rate predictability in Section II.A and compare the forecasting power of survey expectations against other exchange rate predictors in Section II.B. Finally, we assess the out-of-sample performance of survey expectations in Section II.C. In this section, we focus exclusively on long-horizon (i.e., two-year) expectations; we will contrast the results for long-horizon forecasts with those for shorter horizon forecasts in Section IV.

A. In-Sample Predictions

We start our analysis by adding survey-based excess return expectations to the UIP regression of currency appreciation on interest rate differentials. That is, we estimate

$$RCA_{i,t,h} = \alpha_h + \gamma_1 SXR_{i,t,h} + \gamma_2 IRD_{i,t,h} + \varepsilon_{i,t,h}. \tag{9}$$

According to the traditional UIP prediction, interest rate differentials explain currency appreciation, so in the event that UIP holds and the deviations from UIP in survey expectations shown in Figure 2 are pure noise, we should find $\gamma_1 = 0$ and $\gamma_2 = 1$. If the estimate of the coefficient γ_1 is positive and significantly different from zero, survey expectations are qualitatively successful

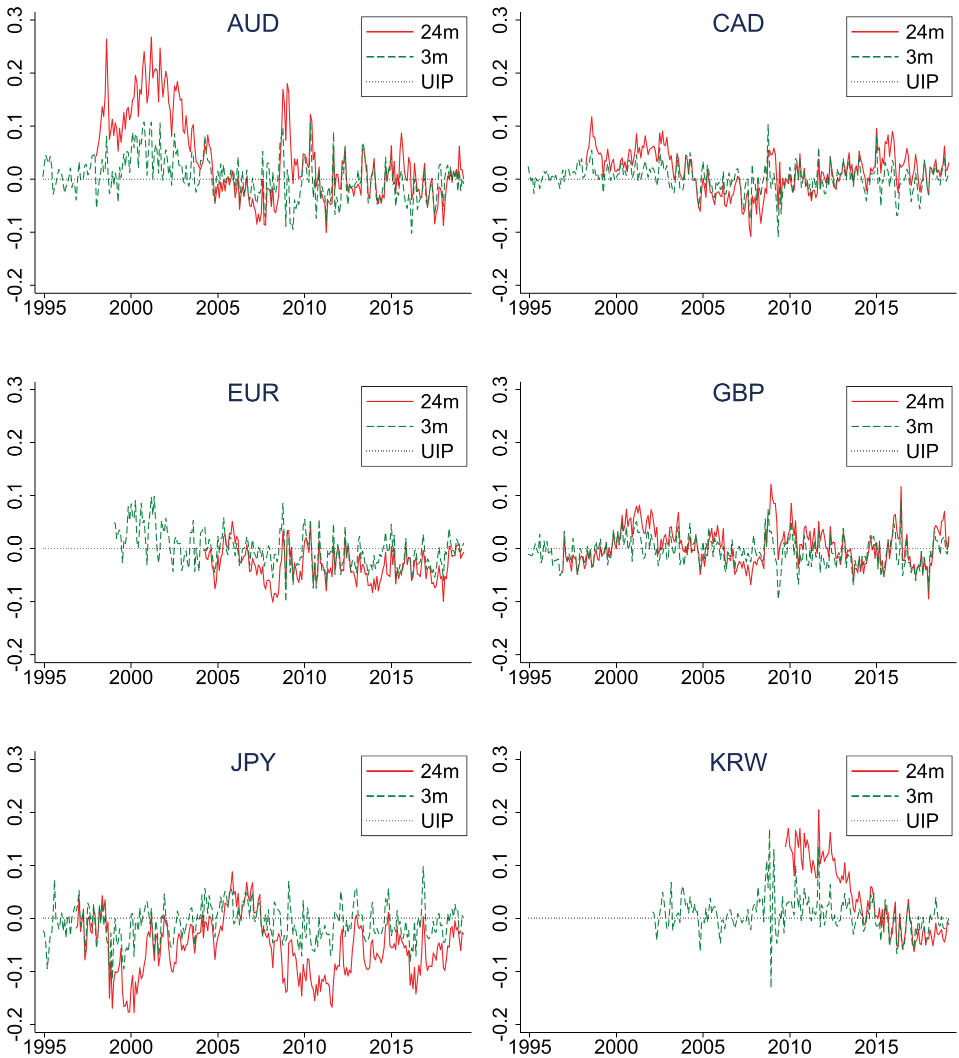

Figure 2. Long-horizon survey expectations over time. This figure plots survey expectations of currency excess returns (*SXR*, not annualized) at three- and 24-month horizons. UIP predicts that expected excess returns are zero. (Color figure can be viewed at wileyonlinelibrary.com)

exchange rate predictors. If both γ_1 and γ_2 are close to one, surveys are also quantitatively successful, in that they predict not just the direction but also the size of currency movements. Throughout the paper, we consider predictor success in terms of these two criteria (coefficient statistically positive and economically close to one) as well as R^2.

We also estimate an alternative specification with realized excess returns on the left-hand side:

$$RXR_{i,t,h} = \alpha_h + \gamma_1 SXR_{i,t,h} + \gamma_2 IRD_{i,t,h} + \varepsilon_{i,t,h}. \tag{10}$$

As excess returns equal currency appreciation minus the interest differential (i.e., $RXR_{i,t,h} = RCA_{i,t,h} - IRD_{i,t,h}$, by equation (6)), the coefficient estimates in (9) and (10) are mechanically related: The estimated γ_1 will be identical in each case, and the estimated γ_2 will differ by exactly one. We include the two specifications to compare the R^2 for both currency appreciation and excess returns.[6]

In both cases, we also estimate specifications with currency and time fixed effects. As our baseline exercise tests long-horizon forecasts over a relatively short sample, we estimate standard errors using a nonparametric block-bootstrap to account for overlapping observations, as in Kremens and Martin (2019).[7]

Columns (1) and (5) of Table I, Panel A, show that, in line with the existing literature, interest rate differentials have limited predictive power for currency movements, with R^2 s in univariate regressions of RCA or RXR on IRD that are close to zero. The coefficient on IRD is imprecisely estimated, however, so we do not statistically reject the prediction of UIP that $\gamma_2 = 1$ in column (1) and $\gamma_2 = 0$ in column (5).

Columns (2) and (6) add survey excess returns as a regressor. The coefficient on SXR is positive, statistically significant, and close to one, indicating that the surveyed forecasters successfully predict the direction and size of currency movements. Moreover, R^2 increases more than fivefold for currency appreciation and nearly tenfold for excess returns, to 16.9% and 15.7%, respectively.[8]

Columns (3) and (7) report similar results with currency fixed effects and indicate that surveys successfully forecast within-currency appreciation. Columns (4) and (8) assess the predictive power across currencies by reporting results with time fixed effects. The coefficients on survey expectations remain significantly different from zero, but are also significantly different from

[6] Similarly, the estimate of γ_1 is identical between using *SXR* or *SCA* as predictors. This choice also mechanically affects the estimate of γ_2, which is reduced by $\hat{\gamma}_1$ when using *SCA*. We prefer to express forecasts in terms of excess returns because doing so cleanly separates the role of interest differentials in excess returns from its mechanical role in currency appreciation. The exception to this preference is Table IX where we construct *forward* expectations, in which case interest rates make it more tedious to address excess returns.

[7] We draw, with replacement, blocks with a time-series length equal to the forecasting horizon and cross-sectional width uniformly distributed between two and six. We permute the cross-section before each draw and randomize the cross-sectional block width to account for cross-sectional correlation. We reconstitute these blocks to form 10,000 bootstrap samples with the same size as our original sample and reestimate the regressions. The bootstrapped standard errors are the standard deviations of the coefficient estimates across bootstrap samples. They are typically more conservative (i.e., larger) than standard errors based on Hansen and Hodrick (1980).

[8] The 10^{th} percentiles of the bootstrapped R^2 distributions are 9.3% and 9.8%, and the respective 90^{th} percentiles are 33.8% and 31.5%.

Table I

In-Sample Forecast Performance

This table reports forecasting regressions (9) and (10) of 24-month realized currency appreciation (RCA) and currency excess returns (RXR) on survey-based expectations of currency excess returns (SXR) and interest rate differentials (IRD). The sample is December 2009 to March 2019 (realizations until March 2021) in Panel A and December 1994 to March 2019 in Panel B, and includes AUD, CAD, EUR, GBP, JPY, and KRW against USD. In brackets, we report standard errors obtained from a nonparametric block-bootstrap to account for overlapping observations in long-horizon forecasts.

	RCA				RXR			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Post-GFC Sample (12/2009 to 3/2019)								
SXR		0.726	0.837	0.523		0.726	0.837	0.523
		[0.212]	[0.251]	[0.213]		[0.212]	[0.251]	[0.213]
IRD	0.577	1.065	1.147	0.693	−0.423	0.065	0.147	−0.307
	[0.599]	[0.601]	[0.674]	[0.548]	[0.599]	[0.601]	[0.674]	[0.548]
Constant (p.a.)	−0.020	−0.017			−0.020	−0.017		
	[0.012]	[0.010]			[0.012]	[0.010]		
Fixed effects	None	None	Currency	Time	None	None	Currency	Time
R^2	0.031	0.169	0.192	0.564	0.017	0.157	0.180	0.558
Within R^2	0.031	0.169	0.165	0.117	0.017	0.157	0.130	0.174
N	672	672	672	672	672	672	672	672
Panel B. Full Sample (12/1994 to 3/2019)								
SXR		0.865	1.066	0.601		0.865	1.066	0.601
		[0.294]	[0.269]	[0.198]		[0.294]	[0.269]	[0.198]
IRD	0.156	0.600	−0.020	0.615	−0.844	−0.400	−1.020	−0.385
	[0.575]	[0.631]	[0.707]	[0.423]	[0.575]	[0.631]	[0.707]	[0.423]
Constant (p.a.)	0.004	0.004			0.004	0.004		
	[0.013]	[0.013]			[0.013]	[0.013]		
Fixed effects	None	None	Currency	Time	None	None	Currency	Time
R^2	0.002	0.145	0.185	0.628	0.058	0.192	0.231	0.649
Within R^2	0.002	0.145	0.173	0.115	0.058	0.192	0.188	0.193
N	1,340	1,340	1,340	1,340	1,340	1,340	1,340	1,340

one. Thus, time-series predictability is an important part of the success of survey forecasts.

In Panel B, we extend our analysis to the period starting in December 1994 and reestimate regressions (9) and (10). The coefficient on SXR remains statistically significant in all specifications and is close to one both in the panel and with currency fixed effects. Here, the R^2 is similar to that of the shorter sample.[9] Table IA.IV in the Internet Appendix reports the full-sample results by currency: The point estimate on SXR is economically and statistically close

[9] Table IA.III in the Internet Appendix reports results for the pre-GFC period (December 1994 to March 2008) separately. The SXR coefficients are again close to one. The statistical difference from zero is more marginal in the panel and with time fixed effects. Here, R^2 is comparable to or higher than in the full sample.

to one for all currencies except EUR and GBP, for which it is statistically larger than one, and is statistically different from zero for all currencies except CAD. Estimating currency-specific coefficients increases the average in-sample R^2 to 25.5% (RCA) and 26.6% (RXR), respectively.[10]

B. Alternative In-Sample Predictors

We now compare the predictive success of surveys with six other predictors of excess returns proposed by the prior literature: the quanto-implied risk premium, the real exchange rate, implied equity market volatility, capital flows, and factor loadings on dollar and carry.

Quanto-implied risk premia (QRP): Kremens and Martin (2019) rewrite equation (3) to show that expected currency appreciation satisfies the model-free identity

$$\mathbb{E}_t \frac{e_{i,t+h}}{e_{i,t}} - \frac{R^{\$}_{f,t,h}}{R^{i}_{f,t,h}} = \underbrace{\frac{1}{R^{\$}_{f,t,h}} \operatorname{cov}^{\mathbb{Q}}_t\left(\frac{e_{i,t+h}}{e_{i,t}}, R_{t+h}\right)}_{\text{QRP}} + \underbrace{\operatorname{cov}_t\left(-M_{t+h}R_{t+h}, \frac{e_{i,t+h}}{e_{i,t}}\right)}_{\text{residual}}, \quad (11)$$

where the superscript $\mathbb{Q}$ indicates a risk-neutral quantity.

The return R_{t+h} that appears in the above identity can be an arbitrary dollar-denominated gross return. For example, setting $R_{t+h} = R^{\$}_{f,t,h}$, we recover equation (3). Kremens and Martin (2019) suggest setting R_{t+h} equal to the gross return on the S&P 500, arguing that this generates a smaller residual covariance term while leaving the first covariance term in (11) directly observable from the prices of so-called quanto forwards.

To see why, recall that the long side of a conventional forward contract on the S&P 500 index pays $P_{t+h} - F_t$ dollars at maturity date $t+h$, where P_{t+h} is the level of the S&P 500 index at time $t+h$ and F_t is the forward price agreed at date t. By contrast, the long side of a quanto forward contract pays $P_{t+h} - Q_{i,t,h}$ units of currency i at maturity date $t+h$ where $Q_{i,t,h}$, the quanto forward price, is agreed at date t. Equivalently, the quanto forward contract pays $e_{i,t+h}(P_{t+h} - Q_{i,t,h})$ dollars at time $t+1$.

The forward and quanto forward prices are chosen to make the initial market value of the contract zero, that is, $F_t = \mathbb{E}^{\mathbb{Q}}_t P_{t+h}$ and $Q_{i,t,h}\, \mathbb{E}^{\mathbb{Q}}_t e_{i,t+h} = \mathbb{E}^{\mathbb{Q}}_t (e_{i,t+h} P_{t+h})$. Consequently, the quanto forward price is sensitive to the risk-neutral covariance between the S&P and the exchange rate. Specifically, we

[10] Table IA.V shows that the results extend to other currencies like the New Israeli shekel and the New Zealand dollar—which we exclude from the baseline sample because we do not observe quanto data for them. Measuring long-horizon IRD for emerging markets is difficult due to sparse and noisy forward data. A univariate regression shows that survey expectations significantly predict currency appreciation in emerging markets, consistent with findings in Kalemli-Özcan and Varela (2024) for the 12-month horizon. Their findings post-date circulation of an earlier draft of the present paper.

have

$$QRP_{i,t,h} = \frac{Q_{i,t,h} - F_t}{R^i_{f,t,h} P_t} = \frac{1}{R^{\$}_{f,t,h}} \mathrm{cov}^{\mathbb{Q}}_t \left(\frac{e_{i,t+h}}{e_{i,t}}, R_{t+h} \right), \tag{12}$$

where $R^i_{f,t,h}$ and $R^{\$}_{f,t,h}$ are risk-free interest rates (see Kremens and Martin (2019, p. 817)). We follow Kremens and Martin (2019) and construct QRP from quotes on 24-month conventional and quanto forwards on the S&P 500 obtained from Markit.

The residual term in identity (3) expresses differences in expected currency appreciation between a risk-neutral benchmark, in which UIP holds, and one in which agents are risk-averse. By contrast, the residual term in identity (11) is zero in a benchmark in which the marginal investor has log utility and is fully invested in the S&P 500, so that $M_{t+h} = 1/R_{t+h}$ and currency risk premia line up perfectly with QRP:[11]

$$\mathbb{E}_t \frac{e_{i,t+h}}{e_{i,t}} - \frac{R^{\$}_{f,t,h}}{R^i_{f,t,h}} = \frac{1}{R^{\$}_{f,t,h}} \mathrm{cov}^{\mathbb{Q}}_t \left(\frac{e_{i,t+h}}{e_{i,t}}, R_{t+h} \right). \tag{13}$$

Equation (13) predicts that a currency should earn a positive excess return if it is risky in the sense of having positive risk-neutral covariance with the market (as measured by the S&P 500 index).[12] Currencies that depreciate when equity markets crash are risky, while those that appreciate are hedges (so-called "safe haven currencies"). QRP reveals whether a currency is one or the other because, unlike measures of FX or equity market volatility, it captures the sign of the correlation between exchange rates and the stock market. Note also—though this point is not made by Kremens and Martin (2019)—that QRP may arise as a predictor of excess returns even in the absence of risk aversion. We discuss examples illustrating why this may be the case for QRP and other predictor variables in Section III.C.

Kremens and Martin (2019) show that QRP predicts 24-month currency excess returns in and out of sample. They also show, however, that other variables capture the empirical counterpart of the residual term in (11)—most notably the real exchange rate.

Real exchange rate (RER): Asness, Moskowitz, and Pedersen (2013), Koijen et al. (2018), and Chernov, Dahlquist, and Lochstoer (2023) show that RER is a persistent predictor of currency excess returns. Dahlquist and Penasse (2022) further argue that RER captures a "missing risk premium" distinct from information in interest rate differentials.

[11] If the investor is more risk averse than log, the residual is increasing in QRP (see Della Corte, Gao, and Jeanneret (2023)) and the slope coefficient of (realized or expected) excess returns on QRP exceeds one. Kremens and Martin (2019) show that this is true for realized returns. Table IV and Della Corte, Gao, and Jeanneret (2023) show that this is true for survey expectations.

[12] For comparison, forward-looking true covariances come out of the theory of the CAPM. Unlike risk-neutral covariances, however, true covariances are not observable, so backward-looking realized covariances must be used as proxies in empirical implementations.

Implied equity market volatility (VIX): Kalemli-Özcan (2019) and Kalemli-Özcan and Varela (2024) show that the VIX correlates with currency excess returns in advanced and emerging market economies. While VIX has no cross-sectional dimension, it is often used as a broad uncertainty proxy that drives risk premia in the time series. Martin (2017) argues that a relative of the VIX ("SVIX", the risk-neutral variance of the S&P 500) represents a lower bound on the equity premium.

Factor loadings on "Dollar" and "Carry" ($\beta^{\$}$, β^{HML}): Lustig, Roussanov, and Verdelhan (2011, 2014) show that the factor structure of exchange rates is well summarized by the returns to two trading strategies, termed Dollar and Carry. The former goes long (short) the dollar against a basket of currencies when dollar interest rates are high (low) relative to the rest of the world; the latter goes long high-interest currencies against low-interest currencies.

Current account balance over GDP (CA/GDP): The international macro-finance literature shows that current account balances are linked to exchange rates (e.g., Kouri (1976), Dornbusch (1976), Obstfeld and Rogoff (2005), Gourinchas and Rey (2007)). A recent literature emphasizes the importance of capital flows in the presence of constraints on global financial intermediaries (e.g., Gabaix and Maggiori (2015), Bianchi, Bigio, and Engel (2023)). Given the connection between trade balances and capital flows, both literatures hypothesize a role for the current account in exchange rate determination. In robustness tests, we employ alternative measures of cross-border financial operations, including the capital inflows-to-GDP ratio and the net foreign asset position-to-GDP ratio.

We estimate univariate regressions of realized excess returns on each of these alternative predictors, the interest rate differential, and survey-based excess returns. Our interest is in comparing the univariate R^2, which we report in Table II. Survey expectations of excess returns have the highest explanatory power with an R^2 of 15.7%, more than one-third higher than the second-best predictor, QRP, at 11.6%. The third-best univariate predictor is RER with an R^2 of 10.4%. Other financial variables have substantially lower explanatory power with R^2 s of 8.5% for the VIX, 7.2% for β^{HML}, 1.7% for IRD, 0.9% for $\beta^{\$}$, and essentially zero for the current account.[13] In the longer sample, survey expectations attain an R^2 of 18.1%, almost twice as high as the runner-up (VIX, with 10.5%; quanto data are unavailable prior to 2009).

C. Out-of-Sample Predictions

Survey expectations predict exchange rates in sample, but the literature has struggled to overturn the result of Meese and Rogoff (1983) that the random walk process is a better out-of-sample predictor of exchange rates than many

[13] The current account is a proxy for net capital flows. For robustness, we also estimate univariate regressions for other macro-finance variables, but they all result in low R^2 s. In particular, the R^2 s are: net foreign asset position-to-GDP 1%, capital inflows-to-GDP 0.2%, industrial production 5.1%, and primary balance 1.2%.

Table II

R^2 of Alternative Predictors

This table reports the univariate R^2 of regressions of 24-month realized currency excess returns (RXR) on the candidate predictors (and a constant). The last column treats the dollar and carry betas of Lustig, Roussanov, and Verdelhan (2011, 2014) as a single model and reports the bivariate R^2. The row labeled p-value reports the fraction of bootstrap draws in which the R^2 for the corresponding variable exceeds that for SXR in the full sample (with the exception of QRP which is calculated using the post-GFC sample due to data availability). The post-GFC sample starts in December 2009 and the full sample in December 1994. Predictor variables run through March 2019, with realizations until March 2021.

	Univariate R^2 of RXR on Each Variable								Bivariate
	SXR	QRP	RER	VIX	β^{HML}	$\beta^{\$}$	IRD	CA / GDP	LRV (β^{HML} & $\beta^{\$}$)
Sample	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Post-GFC	0.157	0.116	0.104	0.085	0.072	0.009	0.017	0.000	0.085
Full	0.181		0.074	0.105	0.012	0.003	0.058	0.003	0.016
p-value		[0.191]	[0.069]	[0.193]	[0.030]	[0.048]	[0.097]	[0.010]	[0.198]

macro models. Survey expectations are well suited for out-of-sample forecasting and a natural competitor of the random walk, because they express ex ante predictions without need to estimate free parameters.

The second-best in-sample univariate predictor, QRP, is also well suited for out-of-sample testing—it describes the ex ante prediction of an unconstrained, rational investor with log utility who holds the stock market. Since QRP beats the random walk in dollar-neutral out-of-sample forecasts (Kremens and Martin (2019)), we add it as a second competitor model.

We define the survey-based forecast error as the difference between the realized appreciation and SCA: $\epsilon^{S}_{i,t,t+h} = RCA_{i,t,h} - SCA_{i,t,h}$. For the random walk, the currency appreciation forecast is zero, so the error is $\epsilon^{RW}_{i,t,t+h} = RCA_{i,t,h}$. For the quanto theory, the forecast error is $\epsilon^{Q}_{i,t,t+h} = RCA_{i,t,h} - (IRD_{i,t,h} + QRP_{i,t,h})$. Focusing again on the 24-month horizon, we compute the out-of sample R^2_{OS} as in Goyal and Welch (2008),

$$R^2_{OS} = 1 - \frac{\sum_i \sum_t (\epsilon^{S}_{i,t,t+h})^2}{\sum_i \sum_t (\epsilon^{C}_{i,t,t+h})^2}, \tag{14}$$

for competitor model $C \in \{RW, Q\}$. A positive R^2_{OS} indicates a smaller mean-squared error of the surveys relative to the competitor model. We refer to this quantity as the "dollar-based" measure, as it computes errors in bilateral exchange rate forecasts against the dollar. Since the dollar has strengthened substantially over the relatively short post-crisis sample, we also calculate a "dollar-neutral" measure that compares different models' performance in forecasting relative appreciation of different currencies (e.g., in forecasting

Table III

Out-of-Sample Forecast Performance

This table reports out-of-sample R^2 s following Goyal and Welch (2008) for surveys against the random walk (RW) and the Quanto Theory (QRP). The different measures for dollar-based and dollar-neutral returns are defined in equations (14) and (15). The last two lines of the table report p-values for a Diebold-Mariano (DM) test as well as bootstrapped p-values for a test of the null hypothesis that survey expectations and the competitor model perform equally well for all currencies.

Sample	Post-GFC				Full Sample
Benchmark	RW		QRP		RW
Dollar-Based/Dollar-Neutral	R^2_{OS}	$\tilde{R}^2_{OS}$	R^2_{OS}	$\tilde{R}^2_{OS}$	R^2_{OS}
	(1)	(2)	(3)	(4)	(5)
All	0.1915	0.1499	0.2095	0.0540	0.1341
AUD	0.3125	0.2257	0.2522	0.1268	0.0944
CAD	−0.0054	−0.0639	0.0274	−0.1421	0.0723
EUR	0.3553	0.0711	0.4511	0.0028	0.3726
GBP	0.0841	0.0102	0.1473	−0.0738	0.1964
JPY	0.2024	0.1444	0.1753	0.0395	0.1400
KRW	0.0098	0.4740	0.1604	0.3775	−0.1923
Diebold-Mariano p-value	0.0809	0.0474	0.0278	0.3468	0.1598
Bootstrapped p-value	0.0881	0.0337	0.0382	0.2446	0.0967

dollar-yen relative to dollar-euro):

$$\tilde{R}^2_{OS} = 1 - \frac{\sum_i \sum_j \sum_t (\epsilon^S_{i,t,t+h} - \epsilon^S_{j,t,t+h})^2}{\sum_i \sum_j \sum_t (\epsilon^C_{i,t,t+h} - \epsilon^C_{j,t,t+h})^2}. \tag{15}$$

The results of this exercise are reported in Table III. Columns (1) and (2) show that surveys outperform the random walk in dollar-based ($R^2_{OS} = 19.15\%$) and dollar-neutral ($\tilde{R}^2_{OS} = 14.99\%$) forecasts. We compute p-values from the bootstrap procedure outlined in Footnote 7 and additionally run Diebold–Mariano tests (Diebold and Mariano (1995)) of the null hypothesis that the forecasts perform equally well for all currencies. In either case, the outperformance relative to the random walk in dollar-based forecasts is at the margins of statistical significance at conventional levels. Perhaps due to unexpected dollar appreciation over the post-GFC sample, outperformance is statistically stronger in cross-sectional (i.e., dollar-neutral) predictions, where survey expectations beat the random walk with a bootstrapped p-value of 3.37%.

To assess whether these results are driven by any particular currency, we additionally estimate individual $R^2_{OS,i}$ and $\tilde{R}^2_{OS,i}$ for each currency, $R^2_{OS,i} = 1 - \frac{\sum_t (\epsilon^S_{i,t,t+h})^2}{\sum_t (\epsilon^C_{i,t,t+h})^2}$ and $\tilde{R}^2_{OS,i} = 1 - \frac{\sum_j \sum_t (\epsilon^S_{i,t,t+h} - \epsilon^S_{j,t,t+h})^2}{\sum_j \sum_t (\epsilon^C_{i,t,t+h} - \epsilon^C_{j,t,t+h})^2}$. Results presented in the second to seventh rows confirm that both the dollar-based and the dollar-neutral measures are positive for all currencies except the Canadian dollar. Survey expectations also beat the quanto-theory forecast, with $R^2_{OS} = 20.95\%$ and

$\tilde{R}^2_{OS} = 5.40\%$, and significantly so for dollar-based predictions with a bootstrapped p-value of 3.82%. The Diebold-Mariano p-values yield similar results.

The results are qualitatively similar over the longer sample (column (5)): Survey expectations beat the random walk ($R^2_{OS} = 13.41\%$), but by a smaller margin and with marginal statistical significance. This comparison suggests that the strong dollar appreciation since the financial crisis does not bias the test against surveys.

We have seen that survey forecasts are successful predictors of exchange rate movements in and out of sample. We next examine what informs these survey expectations.

III. What Informs Expectations?

We run regressions of survey forecasts of excess returns on the interest rate differential and the various candidate covariates described in the previous section,

$$SXR_{i,t,h} = \alpha_h + \gamma_1 X_{i,t} + \gamma_2 IRD_{i,t,h} + \varepsilon_{i,t,h}, \tag{16}$$

where $X_{i,t}$ is a vector containing a subset of the following contemporaneous covariates: QRP, RER, VIX, the CA-to-GDP ratio, $\beta^{\$}$, and β^{HML}. We first assess these covariates individually (or in pairs in the case of $\beta^{\$}$ and β^{HML}) and then jointly. We cluster standard errors by time and currency and standardize the independent variables for ease of comparison.

Table IV reports the results. Columns (1) through (5) report coefficient estimates for univariate regressions of survey excess returns on the various macro-finance predictor variables. QRP and RER are each individually highly significant and have considerable explanatory power, with R^2 around 40%; the other variables are not significant at conventional levels. Coefficients are expressed in percentage points, so column (2) implies that a one-standard-deviation change in QRP corresponds to a 3.737 percentage point increase in (fitted) survey excess returns, and column (3) implies that a one-standard-deviation change in RER corresponds to a 3.090 percentage point decrease in survey excess returns.

Column (6) reports estimates for a bivariate regression of survey excess returns on dollar and HML betas. The HML beta is significant, but the two betas together have limited explanatory power, achieving an R^2 less than half that achieved by QRP or RER on their own.

Column (7) reports estimates for a multivariate regression that includes all of the predictor variables. QRP and RER remain significant—and highly significant in the case of QRP—while CA-to-GDP, which was not significant in a univariate regression, enters significantly. The other variables are not significant, and column (8) shows that they can be dropped entirely at almost no cost

Table IV

What Informs Exchange Rate Expectations?

This table presents regressions of 24-month survey expectations of currency excess returns (SXR, not annualized) onto a constant and various standardized financial and macroeconomic variables: the interest rate differential (IRD), the quanto-implied risk premium (QRP), the real exchange rate (RER), the 30-day S&P implied volatility index (VIX), the current account-to-GDP ratio (CA/GDP), and the 24-month rolling monthly beta of the exchange rate on the dollar and carry factors of Lustig, Roussanov, and Verdelhan (2011, 2014), respectively ($\beta^{\$}$, β^{HML}). Coefficients are expressed in percentage points. Standard errors in parentheses are clustered at the currency and time level. We report asterisks indicating significance at 10%, 5%, and 1%, respectively, for convenience given the large number of columns and regressors.

	Survey Excess Returns (SXR)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
IRD	−2.041	0.851	−3.467 **	−1.895	−1.913	−1.950	0.283	
	(1.142)	(0.612)	(0.905)	(1.192)	(1.460)	(1.232)	(0.861)	
QRP		3.737 ***					3.056 ***	3.052 ***
		(0.491)					(0.239)	(0.515)
RER			−3.090 **				−1.763 **	−1.807 **
			(0.927)				(0.678)	(0.598)
VIX				0.732			0.141	
				(0.798)			(0.304)	
CA / GDP					−0.297		−1.274 **	−1.287 ***
					(1.406)		(0.386)	(0.300)
$\beta^{\$}$						−0.156	−0.308	
						(0.885)	(0.807)	
β^{HML}						1.058 ***	0.347	
						(0.244)	(0.271)	
R^2	0.138	0.402	0.387	0.155	0.140	0.175	0.536	0.528
N	672	639	672	672	672	672	639	639

in terms of explained variation: R^2 drops from 53.6% when all variables are included to 52.8% when only QRP, RER, and CA-to-GDP are included.[14]

Table IA.VI in the Internet Appendix shows that results for the full sample are qualitatively comparable, except that in the absence of QRP, all other covariates jointly explain less than 40% of the variation in survey expectations (compared to more than half in the post-GFC period when QRP is included). IRD and RER covary with excess return expectations. Again, the current account is not individually significant, but complements RER.

Table IA.VII in the Internet Appendix presents analogous short-sample results with currency and time fixed effects. QRP and RER are significant individual covariates with high R^2 both within and across currencies. The current

[14] It is no surprise that the R^2 achieved in regressions that aim to explain *expected* returns, as in Table IV, are much higher than those achieved in regressions that aim to explain *realized* returns, as in Table I. As a hypothetical example, in a CAPM equilibrium betas are known and perfectly explain expected returns so that the R^2 in a regression of expected returns onto betas would have an R^2 of 100%; but, as there may be arbitrary amounts of idiosyncratic risk, regressions of realized returns onto explanatory variables may only achieve low R^2.

account balance is again significant only jointly with other regressors. In the multivariate cross-sectional regression (with time fixed effects), the loadings on dollar and carry also become significantly positively correlated with survey expectations of excess returns. Comparing raw R^2 and within-R^2, we note that time fixed effects explain a larger portion of survey variation than currency fixed effects, indicating that dollar-related elements that are unspanned by these covariates play a larger role in the panel of currency return expectations than currency-specific but time-invariant unspanned components. The observation from Table I that predictability is slightly stronger within than across currencies suggests that such an unspanned dollar component may contribute to surveys' forecasting success.

A. Do Survey Respondents Have a "Secret Sauce"?

The previous section shows that survey forecasts load heavily on QRP, RER, and CA-to-GDP. We now ask whether the survey forecasts include any additional information that improves their performance relative to these variables.

To do so, we compute the fitted values $\widehat{SXR}$ and residuals $\varepsilon(SXR)$ from a regression of survey expected excess returns onto QRP, RER, CA-to-GDP, and a constant (i.e., the trivariate specification of (16) reported in column (8) of Table IV). The fitted values represent the component of survey expectations attributable to movements in QRP, RER, and CA-to-GDP. The residuals represent the component that the macro-finance variables cannot explain (even in sample): If they have predictive power for realized currency movements, then we can think of them as the "secret sauce" in survey expectations.

Column (1) of Table V reproduces our previous finding for a regression of RCA on interest rate differentials and survey excess returns; it is almost but not quite identical to column (2) of Table I, Panel A, because the number of observations decreases slightly due to the lack of quanto data for some currency-time pairs.

Column (2) includes only the fitted value, $\widehat{SXR}$, and the R^2 increases from 17.0% to 25.2%. Column (3) adds the residuals, $\varepsilon(SXR)$. Separately estimated, the coefficient on the residual is economically and statistically close to zero and its inclusion adds little to the R^2, indicating that the residuals do not contain predictive information about excess returns. Survey expectations aggregate useful predictive information from a few predictors that make them the best univariate predictor, but contain little information with predictive power beyond this set of variables. In that sense, there is no secret sauce.

We note, however, that this definition of a secret sauce, based on in-sample residualization, sets a high bar for finding one. The in-sample fit presumes that forecasters knew the right mapping of predictor variables to multivariate forecasts. If, instead, forecasters must learn the true model, researchers studying the data should expect residuals to feature a component that looks like noise ex post (Lewellen and Shanken (2002), Martin and Nagel (2022)). In the next

Table V

Do Survey Respondents Have a Secret Sauce?

This table reports forecasting regressions of 24-month realized currency appreciation (RCA) onto a constant, interest rate differentials (IRD), survey-based excess returns (SXR), and the fitted values ($\widehat{SXR}$) and residuals ($\varepsilon(SXR)$) of SXR. Fitted values and residuals are obtained from the trivariate specification in column (8) of Table IV. The sample runs from December 2009 to March 2019 (realizations until March 2021) and includes AUD, CAD, EUR, GBP, JPY, and KRW against USD. In brackets, we report standard errors obtained from a nonparametric block-bootstrap to account for overlapping observations in long-horizon forecasts.

	RCA		
	(1)	(2)	(3)
IRD	1.137	1.563	1.559
	[0.747]	[0.855]	[0.900]
SXR	0.740		
	[0.246]		
$\widehat{SXR}$		1.415	1.414
		[0.841]	[0.832]
$\varepsilon(SXR)$			0.177
			[0.232]
R^2	0.170	0.252	0.256
N	639	639	639

section, we therefore also benchmark surveys against out-of-sample forecasts that an econometrician could construct from these variables in real time.

B. What Are the Best Predictors of Currency Returns?

Having shown that expectations are individually successful predictors of currency movements and excess returns (Table I), that they are largely explained by QRP, RER, and CA-to-GDP (Table IV), and that they do not contain predictive content beyond those variables (Table V), we now ask which variables are the most successful predictors in multivariate regressions. Of the possible predictor combinations, Table VI reports the univariate, bivariate, and trivariate specifications that produce the highest R^2 in forecasting realized excess returns.

With two predictors, QRP and RER raise R^2 to 26% from 15.7% for the univariate survey-based forecast. The success of this combination partly reflects the fact that the correlation between QRP and RER is low, at 0.111. (Table IA.VIII in the Internet Appendix reports correlations among the macro-finance variables.) The R^2 rises modestly, to 31.4%, when VIX is included, and to 35.9% when all variables are included. Columns (4) and (5) show that survey forecasts contribute little explanatory power when we use the full set of macro-finance predictors.

All predictor variables are standardized to have unit standard deviation, and coefficients are reported in percentage points, as in Table IV, so the coefficients in column (2) of Table VI indicate that one-standard-deviation changes in QRP

Table VI

R^2-Maximizing Predictors

This table reports the R^2-maximizing univariate, bivariate, etc., specifications in regressions of 24-month realized currency excess returns (RXR) onto a constant and combinations of various standardized candidate predictors. The last column reports the specification with all variables except SXR. Coefficients are expressed in percentage points.

	Coefficient Estimates in R^2-Maximizing Specifications				
	Univariate (1)	Bivariate (2)	Trivariate (3)	8-Variate (4)	Excl. SXR (5)
SXR	3.916	.	.	0.639	.
QRP	.	3.705	3.055	0.988	1.343
RER	.	−3.715	−4.183	−4.805	−5.010
VIX	.	.	2.459	2.920	2.936
IRD	.	.	.	−3.018	−2.985
CA/GDP	.	.	.	0.528	0.379
$\beta^{\$}$	.	.	.	1.170	1.135
β^{HML}	.	.	.	1.827	1.867
R^2	0.157	0.260	0.314	0.359	0.357

or RER each move the bivariate regression's forecast of realized excess returns by about 3.7 percentage points (in opposite directions). For comparison, column (8) of Table IV suggests that a one-standard-deviation change in QRP moves expectations by around 3 percentage points while a one-standard-deviation change in RER moves expectations by about 1.8 percentage points.

Together with our finding that a trivariate projection of SXR onto macro-financial covariates captures all of SXR's predictive success (i.e., there is no secret sauce), the above may suggest that surveys are dominated by this small set of macro-finance predictors. We note, however, that this is only true in sample.

Table VII therefore benchmarks surveys against an econometrician's out-of-sample forecast of excess returns based on six macro-finance variables (excluding QRP, which we do not observe prior to December 2009). We construct the econometrician's forecast, $\widehat{RXR}^{OOS}$, from expanding-window regressions, starting with a sample ending in November 2009 (i.e., prior to the start of our post-crisis sample) and adding one month at a time. Column (2) shows that the econometrician's forecast is a statistically significant predictor, but achieves a lower R^2 in predicting excess returns than SXR (11.2% versus 15.7% in column (1)). Column (3) combines the two and shows that the econometrician's forecast raises R^2 only modestly relative to column (1), and is statistically driven out by SXR. This result may indicate that surveys reflect time-variation in the true model of conditional returns that is not captured by the econometrician's model (i.e., a different notion of secret sauce from that in Section III.A). However, it could also arise from the omission of QRP from the econometrician's model.

To disentangle the two effects, we project SXR onto the same six variables underlying the econometrician's forecast. This out-of-sample projection,

Table VII

Surveys versus Out-of-Sample Fitted Forecasts

This table reports forecasting regressions of 24-month realized currency excess returns (RXR) on a constant, survey-based expectations of excess returns (SXR), out-of-sample forecasts of RXR ($\widehat{RXR}^{OOS}$) and SXR ($\widehat{SXR}^{OOS}$), and the residuals of the latter ($\varepsilon(SXR)^{OOS}$). Fitted values and residuals are obtained from expanding-window regressions of RXR and SXR, respectively, on IRD, RER, VIX, CA/GDP, β^{HML}, and $\beta^{\$}$. The estimation window is initialized using dependent variables observed until November 2009 and then expands one month at a time. The sample for the forecasting regressions reported in this table runs from December 2009 to March 2019 (realizations until March 2021). In brackets, we report standard errors obtained from a nonparametric block-bootstrap.

	RXR				
	(1)	(2)	(3)	(4)	(5)
SXR	0.713		0.559		
	[0.200]		[0.218]		
$\widehat{RXR}^{OOS}$		0.745	0.454		
		[0.363]	[0.474]		
$\widehat{SXR}^{OOS}$				1.285	1.407
				[0.696]	[0.763]
$\varepsilon(SXR)^{OOS}$					0.540
					[0.206]
R^2	0.157	0.112	0.191	0.131	0.207
N	672	672	672	672	672

$\widehat{SXR}^{OOS}$, outperforms the econometrician's forecast (albeit narrowly, with an R^2 of 13.1% versus 11.2%). Its advantage is twofold. First, survey forecasts may be less noisy than return realizations, making a projection of the latter more vulnerable to overfitting. Second, the econometrician's forecast requires a time lag to observe 24-month realizations while the projection of surveys can be estimated using the past month's data. Their relative performance indicates that surveys filter out noise in return realizations and/or capture time-variation in the mapping of predictors to returns.

The residual, $\varepsilon(SXR)^{OOS}$, correlates strongly with QRP ($\rho = 0.65$) and thus does not have the secret-sauce interpretation of its in-sample analog in Section III.A. Consistent with our previous findings, SXR outperforms its QRP-deprived projection (R^2 of 15.7% versus 13.1%), and the residual is a strongly significant predictor that substantially raises R^2 (column (5)).[15]

[15] QRP explains around 43% of the variation in $\varepsilon(SXR)^{OOS}$, and accounts for all of the incremental R^2 in column (5). Adding QRP to the specification in column (2) raises the R^2 to 16.9% (above that in column (1)). However, further adding SXR raises R^2 by another 2 percentage points, indicating that a small portion of the outperformance of SXR relative to the econometrician's forecast cannot be explained (even in sample) by QRP.

C. Interpretation

What does the finding that QRP, RER, and CA-to-GDP span much of the variation in expected and realized currency appreciation reveal about the economics of exchange rate determination? Broadly speaking, two views have emerged in the literature: A preference- or risk-based view, according to which risk premia reflect the covariation of currencies with macroeconomic risk factors, and a frictions-based view that emphasizes the importance of constraints, such as those of the financial intermediaries whose expectations we study.

We emphasize that our results are potentially consistent with either view (or both), so our findings do not settle the question of which mechanism is more relevant. Prior literature has given risk-based interpretations for QRP, RER, and CA-to-GDP,[16] but in principle the patterns of expected returns that arise in economies featuring unconstrained, risk-averse investors can equally arise in risk-neutral economies in which marginal investors are subject to constraints, if those constraints are sensitive to risk measures.

Consider, for example, the one-period portfolio choice problem of a financial intermediary that is risk-neutral but subject to a constraint on risk-neutral (or implied) portfolio variance. Write R_{t+h} for the chosen portfolio return, $R^e_{i,t+h}$, and w_i for the excess return on asset i and its portfolio weight, and write κ for the limit on portfolio variance. Such an agent maximizes the expected excess portfolio return subject to the variance constraint, so solves $\max_{w_1,\ldots,w_N} \mathbb{E}_t \sum_i w_i R^e_{i,t+h}$ subject to the constraint that $\mathrm{var}^{\mathbb{Q}}_t(\sum_i w_i R^e_{i,t+h}) \leq \kappa$. The first-order conditions for this problem imply that $\mathbb{E}_t R^e_{i,t+h} = \lambda \,\mathrm{cov}^{\mathbb{Q}}_t(R^e_{i,t+h}, R_{t+h})$ for each i, where λ is the Lagrange multiplier on the variance constraint. Applied to the excess return on currency i, $R^e_{i,t+h} = R^i_{f,t,h} e_{i,t+h}/e_{i,t} - R^{\$}_{f,t,h}$, this implies that

$$\mathbb{E}_t \frac{e_{i,t+h}}{e_{i,t}} - \frac{R^{\$}_{f,t,h}}{R^i_{f,t,h}} = \lambda \,\mathrm{cov}^{\mathbb{Q}}_t\left(\frac{e_{i,t+h}}{e_{i,t}}, R_{t+h}\right). \tag{17}$$

If the intermediary's portfolio return R_{t+h} is (or is perfectly correlated with) the return on the S&P 500, then equation (17) states that the currency expected excess returns are proportional to QRP_i, as in the risk-based view of equation (13).

Essentially the same logic applies if marginal currency investors are risk-neutral but subject to constraints on true variance: In this case, currency excess returns will line up with true covariances between currency appreciation and the investors' portfolio returns. If, for example, variables such as RER and CA-to-GDP are associated with trade or portfolio flows that lead to large, underdiversified currency exposures for marginal financial intermediaries—and

[16] See, for example, Kremens and Martin (2019) in the case of QRP, Asness, Moskowitz, and Pedersen (2013), Koijen et al. (2018), and Chernov, Dahlquist, and Lochstoer (2023) for RER, and Della Corte, Riddiough, and Sarno (2016) and Colacito et al. (2018) for CA-to-GDP.

therefore affect the tightness of variance constraints—then they should be expected to predict currency returns.

IV. Short versus Long Horizons

The predictive success of long-horizon survey expectations is surprising given that previous studies find that short-horizon expectations tend to forecast poorly (e.g., Nagel and Xu (2023), Dahlquist and Söderlind (2022)). In this subsection, we compare the predictive success of forecasts across horizons.

We first reestimate the regression (10) of realized excess returns onto survey excess returns and the interest rate differential for horizons of $h = \{1, 3, 12\}$ months. We annualize variables for ease of comparison across horizons, using asterisks to indicate annualized quantities: For example,

$$SCA^*_{i,t,h} = \frac{12}{h} SCA_{i,t,h}, \tag{18}$$

and similarly for SXR, RXR, IRD, and so on.

Columns (1) to (4) in Table VIII report the results. (We include column (4), which is identical to column (6) of Table I, for convenience.) Consistent with prior literature, the predictive power of surveys is substantially smaller at short horizons: At horizons of one, three, and 12 months, none of the point estimates on survey expectations is statistically distinguishable from zero but all are different from one. We find broadly monotonic patterns on three dimensions: as forecast horizon increases, (i) the coefficient on SXR rises towards one, (ii) the coefficient on IRD shrinks toward zero, and (iii) R^2 rises faster than linearly in horizon.[17]

Given the successful predictive performance of survey forecasts at long horizons, is there something special about long-horizon *realizations*, or is there something special about long-horizon *expectations*? We address this question by comparing the forecasting power of long-horizon forecasts for short-horizon outcomes with the forecasting power of short-horizon forecasts for long-horizon outcomes.

Columns (5) to (7) of Table VIII show that annualized long-horizon forecasts (SXR^*_{24}) successfully predict annualized short-run realizations (RXR^*_h for $h = \{1, 3, 12\}$), with estimated coefficients that are economically and statistically close to one and, in the full sample, significantly different from zero at all horizons. This suggests that the component of excess returns predicted by long-run expectations materializes evenly over the 24-month horizon.

In contrast, columns (8) to (10) of Table VIII show that annualized short-horizon forecasts (SXR^*_h) are broadly unsuccessful at predicting long-horizon realizations (RXR^*_{24}), with point estimates on one- and three-month forecasts that are close to and statistically indistinguishable from zero. The point es-

[17] Campbell and Thompson (2008) show that as a rule of thumb, the economic magnitude of R^2 can be judged against the squared Sharpe ratio, which scales roughly linearly with horizon.

Table VIII

Forecast Performance across Horizons

This table reports forecasting regressions of annualized one-, three-, 12-, and 24-month realized currency excess returns (RXR^*) on a constant and survey-based expectations of annualized excess returns (SXR^*) and interest rate differentials (IRD^*). The sample is an unbalanced panel and includes AUD, CAD, EUR, GBP, JPY, and KRW against USD. The horizon for IRD^* is always equal to that of RXR^*. In brackets, we report standard errors, clustered by currency and time for one-month realizations and obtained from a nonparametric block-bootstrap to account for overlapping observations in three-, 12-, and 24-month forecasts.

	Forecasting RXR^*_h with SXR^*_h				RXR^*_h with SXR^*_{24}			RXR^*_{24} with SXR^*_h		
	1M	3M	12M	24M	1M	3M	12M	1M	3M	12M
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A. Post-GFC Sample (12/2009 to 3/2019)										
SXR^*	0.088	0.093	0.237	0.726	1.548	0.911	0.683	0.007	0.037	0.282
	[0.067]	[0.102]	[0.215]	[0.212]	[0.857]	[0.543]	[0.384]	[0.009]	[0.031]	[0.130]
IRD^*	−1.112	−1.066	−0.689	0.065	−0.083	−0.498	−0.406	−0.425	−0.405	−0.221
	[0.856]	[0.998]	[0.890]	[0.601]	[0.763]	[1.057]	[1.056]	[0.613]	[0.612]	[0.642]
R^2	0.011	0.014	0.038	0.157	0.018	0.028	0.070	0.019	0.025	0.071
N	672	672	672	672	672	672	672	672	672	672
Panel B. Full Sample (12/1994 to 3/2019)										
SXR^*	0.013	−0.010	0.380	0.865	1.007	0.879	0.839	0.012	0.071	0.447
	[0.041]	[0.081]	[0.237]	[0.294]	[0.356]	[0.411]	[0.379]	[0.010])	[0.039]	[0.173]
IRD^*	−1.233	−1.037	−0.835	−0.400	−0.758	−0.610	−0.611	−0.850	−0.816	−0.623
	[0.467]	[0.738]	[0.717]	[0.631]	[0.503]	[0.807]	[0.732]	[0.631]	[0.572]	[0.563]
R^2	0.006	0.011	0.073	0.192	0.013	0.029	0.105	0.061	0.075	0.145
N	1,340	1,340	1,340	1,340	1,340	1,340	1,340	1,340	1,340	1,340

timate on the 12-month forecast is on the border of statistical significance, though far from one.

Table IX examines the relationship between forecasting horizons in a different way. We define the time-t *forward* expectation between horizons h and $H > h$ as

$$sca_{i,t}^{h,H} = \log\left(\frac{1 + SCA_{i,t,H}}{1 + SCA_{i,t,h}}\right) = sca_{i,t,H} - sca_{i,t,h}, \tag{19}$$

where $sca_{i,t,h} = \log(1 + SCA_{i,t,h})$ and $rca_{i,t,h} = \log(1 + RCA_{i,t,h})$. By working in logs, we can decompose the long-horizon expectation as the sum of a short-horizon expectation and forward expectations. For example,

$$sca_{i,t,24} = sca_{i,t,3} + sca_{i,t}^{3,12} + sca_{i,t}^{12,24} \tag{20}$$

expresses expected 24-month currency appreciation as the sum of expected three-month appreciation ($sca_{i,t,3}$) plus expected appreciation from month

Table IX

In-Sample Predictability: Spot and Forward Expectations

This table presents regressions of realized currency appreciation onto a constant and spot and forward survey expectations of currency appreciation. Standard errors in brackets are obtained from a nonparametric block-bootstrap to account for overlapping observations in realizations. Observations range from December 1994 to March 2019.

	RCA_{24}	rca_{24}			rca_3
	(1)	(2)	(3)	(4)	(5)
SCA_{24}	0.812				
	[0.269]				
sca_{24}		0.804			
		[0.266]			
sca_3			0.252	0.246	−0.062
			[0.192]	[0.226]	[0.086]
$sca^{3,12}$			1.070		
			[0.943]		
$sca^{12,24}$			1.136		
			[0.897]		
$sca^{3,24}$				1.102	0.188
				[0.361]	[0.087]
R^2	0.139	0.138	0.162	0.162	0.030
N	1,340	1,340	1,340	1,340	1,340

three to month 12 ($sca_{i,t}^{3,12}$) plus expected appreciation from month 12 to month 24 ($sca_{i,t}^{12,24}$).

The first two columns of Table IX confirm that the predictability of currency appreciation using survey expectations is similar whether we work in logs, regressing $rca_{i,t,h}$ on $sca_{i,t,h}$, or in levels, regressing $RCA_{i,t,h}$ on $SCA_{i,t,h}$.

Columns (3) and (4) predict 24-month log realizations using a spot expectation and the complementary forward expectations. If the predictive information were evenly spread across the different expectation horizons, we would expect all coefficients to be close to one. Instead we find that the three-month spot expectation has a point estimate that is economically and statistically close to zero, while the coefficients on the two forward expectations in column (3) are close to one but (as they are strongly correlated with one another) imprecisely estimated. Column (4) therefore splits the 24-month forecast into a short-term forecast and a complementary long-term forward forecast. The coefficient on the short-term spot forecast is again economically and statistically close to zero. The coefficient on the long-horizon forward expectation is close to one and significantly different from zero.

As an even starker test of this, column (5) predicts three-month realizations using three-month spot expectations and forward expectations from three to 24 months. Again, short-run expectations do not predict successfully. But forward expectations, which reflect what forecasters expect to happen *after* the three-month horizon, reliably predict short-run currency appreciation. Thus,

Table X

What Informs Exchange Rate Expectations at Shorter Horizons?

This table presents regressions of survey expectations of currency excess returns (SXR, not annualized) at one-, three-, 12-, and 24-month horizons on a constant and standardized financial and macroeconomic variables: the horizon-matched interest rate differential (IRD), the 24-month quanto-implied risk premium (QRP), the real exchange rate (RER), the 30-day S&P implied volatility index (VIX), the current account balance relative to GDP (CA/GDP), and the 24-month rolling monthly beta of the exchange rate on the dollar and carry factors of Lustig, Roussanov, and Verdelhan (2011, 2014), respectively ($\beta^{\$}$, β^{HML}). Coefficients are expressed in percentage points. Standard errors in parentheses are clustered at the currency and time level. $*$, $**$, and $***$ indicate significance at 10%, 5%, and 1%, respectively. Observations range from December 2009 to March 2019.

	Survey Excess Returns (SXR)			
	Horizon (Months)			
	1	3	12	24
	(1)	(2)	(3)	(4)
IRD	−0.114	−0.235	0.871	0.283
	(0.274)	(0.420)	(0.728)	(0.861)
QRP	0.094	0.436	2.277 ***	3.056 ***
	(0.218)	(0.231)	(0.173)	(0.239)
RER	−0.098	−0.128	−0.504	−1.763 **
	(0.139)	(0.248)	(0.549)	(0.678)
VIX	0.376	0.423	0.519	0.141
	(0.305)	(0.300)	(0.300)	(0.304)
CA/GDP	−0.049	−0.080	−0.716 **	−1.274 **
	(0.102)	(0.113)	(0.201)	(0.386)
$\beta^{\$}$	0.387*	0.665 **	0.012	−0.308
	(0.191)	(0.225)	(0.535)	(0.807)
β^{HML}	−0.065	−0.066	−0.218	0.347
	(0.269)	(0.253)	(0.222)	(0.271)
R^2	0.030	0.070	0.311	0.536
N	639	639	639	639

consistent with column (6) in Table VIII, long-horizon expectations are informative about short-run realizations, while short-run expectations are largely noise.

We next ask which—if any—of the macro-financial predictor variables help explain short-run expectations. Table X reports results analogous to those in Table IV for different forecast horizons. At one- and three-month horizons, we find that the macro-finance variables explain little of the variation in survey expectations, with only the dollar beta statistically significant. At the 12-month horizon, R^2 is markedly higher, and QRP and the CA-to-GDP ratio are strongly significant, while RER becomes significant only at the 24-month horizon. We standardize the explanatory variables to have unit standard deviation and report coefficients in percentage points, so, for example, a one-standard-deviation change in QRP increases 12-month expectations by about 2.3 percentage points.

Taking stock, the results in this section are consistent with survey forecasters using different models to form expectations at different horizons. Long-horizon forecasts predict well at long and short horizons, and they correlate with observable macro-finance variables. Short-horizon forecasts predict poorly and do not correlate with the macro-finance variables that have performed well over our sample period. Our understanding is that the short-term and long-term forecasts are made not only by the same entities but also by the same teams, and that Consensus Economics reaches out to the same set of forecasters every month.[18] This leaves us with a puzzle—our forecasters' long-horizon models outperform their own short-horizon models in forecasting short-horizon outcomes—that is hard to explain with information frictions.

V. Conclusion

We view our findings as cause for optimism on two fronts. First, the long-horizon expectations of informed market participants about currency movements are broadly rational, in the sense that their forecasts predict realizations, and with an estimated coefficient that is both statistically significant and close to one. Second, their expectations are comprehensible, with three variables that have been studied by macroeconomists and financial economists (QRP, RER, and CA-to-GDP) explaining a substantial fraction of the variation in expectations.

That said, our findings do not identify *how* these variables determine (expected or realized) exchange rates and currency excess returns. Variables such as QRP, RER, and CA-to-GDP may arise as excess return predictors because they reflect loadings on priced risk factors, or because they correlate with the tightness of balance sheet or other constraints faced by the marginal bearers of currency risk (see Section III.C). Thus, while our results are potentially consistent with both views of currency returns, they do not allow us to distinguish between the two—and we note that they are not mutually exclusive.

Similarly, we cannot rule out a role for irrational expectations in exchange rate determination. In particular, we find that *some* individuals' expectations are broadly rational, but these market participants may be correctly anticipating the irrational behavior of others.

Nor are survey expectations consistently rational. The residual component of expectations not correlated with the three macro-finance variables has no predictive power for currency movements. Moreover, although the current account plays a role in shaping long-run exchange rate expectations, it does not predict realizations. Instead, the VIX index (a measure of 30-day implied equity market volatility, and therefore an indicator of short-term market stress) improves forecasts of long-term exchange rate realizations relative to survey forecasts. Even more puzzlingly, while short-horizon expectations fail to forecast short-

[18] While response rates vary across time and currencies, Figure IA.2 shows that the number of forecasters is large relative to its within-currency variation, and thus variation in response rates is unlikely to render consensus estimates noisy within currency and horizon.

run outcomes, long-horizon expectations forecast not only long-run outcomes, but also short-run outcomes. That is, our forecasters' long-horizon models outperform their own short-horizon models in forecasting short-horizon outcomes.

Our data do not let us resolve this puzzle so we leave this for future work and only offer two speculative explanations. The first is based on forecaster composition at different horizons: short-horizon forecasts may be obtained from a systematically noisier forecaster pool. The second, and perhaps more interesting, explanation is that forecasts at different horizons are produced rationally—by the same institution or even individual—for different clienteles, with different objectives, and therefore based on different models.

Initial submission: August 20, 2023; Accepted: March 18, 2025
Editors: Antoinette Schoar, Urban Jermann, Leonid Kogan, Jonathan Lewellen, and Thomas Philippon

REFERENCES

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

Bacchetta, Philippe, and Eric van Wincoop, 2010, Infrequent portfolio decisions: A solution to the forward discount puzzle, *American Economic Review* 100, 870–904.

Bacchetta, Philippe, and Eric van Wincoop, 2021, Puzzling exchange rate dynamics and delayed portfolio adjustment, *Journal of International Economics* 131, 103460.

Bartram, Söhnke, Leslie Djuranovic, Anthony Garratt, and Yan Xu, 2025, Mispricing and risk premia in currency markets, *Journal of Financial and Quantitative Analysis* 60, 695–733.

Beckmann, Joscha, and Stefan Reitz, 2023, Dealer risk premiums in FX forecasts, Working paper.

Bianchi, Javier, Saki Bigio, and Charles Engel, 2023, Scrambling for dollars: Liquidity, banks and exchange rates, NBER Working Paper 29457 .

Burnside, Craig, 2011, The cross section of foreign currency risk premia and consumption growth risk: Comment, *American Economic Review* 101, 3456–3476.

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

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

Candian, Giacomo, and Pierre De Leo, 2025, Imperfect exchange rate expectations, *Review of Economics and Statistics* 107, 513–530.

Cenedese, Gino, Richard Payne, Lucio Sarno, and Giorgio Valente, 2016, What do stock markets tell us about exchange rates?, *Review of Finance* 20, 1045–1080.

Chernov, Mikhail, Magnus Dahlquist, and Lars A. Lochstoer, 2023, Pricing currency risk, *Journal of Finance* 78, 693–729.

Coibion, Olivier, and Yuriy Gorodnichenko, 2015, Information rigidity and the expectations formation process: A simple framework and new facts, *American Economic Review* 105, 2644–78.

Colacito, Riccardo, Mariano M. Croce, Federico Gavazzoni, and Robert Ready, 2018, Currency risk factors in a recursive multicountry economy, *The Journal of Finance* 73, 2719–2756.

Dahlquist, Magnus, and Julien Penasse, 2022, The missing risk premium in exchange rates, *Journal of Financial Economics* 143, 697–715.

Dahlquist, Magnus, and Paul Söderlind, 2022, Individual forecasts of exchange rates, Technical Report 22-06, Swedish House of Finance, Stockholm School of Economics.

De Marco, Filippo, Marco Macchiavelli, and Rosen Valchev, 2021, Beyond home bias: International portfolio holdings and information heterogeneity, *The Review of Financial Studies* 35, 4387–4422.

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

Della Corte, Pasquale, Steven J. Riddiough, and Lucio Sarno, 2016, Currency premia and global imbalances, *Review of Financial Studies* 29, 2161–2193.

Diebold, Francis X., and Roberto S. Mariano, 1995, Comparing predictive accuracy, *Journal of Business and Economic Statistics* 13, 253–263.

Dornbusch, Rüdiger,1976, Expectations and exchange rate dynamics, *Journal of Political Economy* 84, 1161–1176.

Engel, Charles, 2016, Exchange rates, interest rates, and the risk premium, *American Economic Review* 106, 436–474.

Fama, Eugene F., 1984, Forward and spot exchange rates, *Journal of Monetary Economics* 14, 319–338.

Frankel, Jeffrey A., and Kenneth A. Froot, 1987, Using survey data to test standard propositions regarding exchange rate expectations, *The American Economic Review* 77, 133–153.

Froot, Kenneth A., and Jeffrey A. Frankel, 1989, Forward discount bias: Is it an exchange risk premium, *Quarterly Journal of Economics* 104, 139–161.

Froot, Kenneth A., and Richard H. Thaler, 1990, Anomalies: Foreign exchange, *Journal of Economic Perspectives* 4, 179–192.

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

Gourinchas, Pierre-Olivier, and Hélène Rey, 2007, International financial adjustment, *Journal of Political Economy* 115, 665–703.

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

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

Hansen, Lars P., and Robert J. Hodrick, 1980, Forward exchange rates as optimal predictors of future spot rates: An econometric analysis, *Journal of Political Economy* 88, 829–853.

Hassan, Tarek A., and Rui C. Mano, 2019, Forward and spot exchange rates in a multi-currency world, *Quarterly Journal of Economics* 134, 397–450.

Kalemli-Özcan, Şebnem, 2019, U.S. monetary policy and international risk spillovers, Technical report, Proceedings of the Jackson Hole Symposium, 2019.

Kalemli-Özcan, Şebnem, and Liliana Varela, 2024, Five facts about the UIP premium, NBER Working Paper 28923.

Koijen, Ralph S. J., Tobias J. Moskowitz, Lasse H. Pedersen, and Evert B. Vrugt, 2018, Carry, *Journal of Financial Economics* 172, 197–225.

Kouri, Pentti J. K., 1976, The exchange rate and the balance of payments in the short run and in the long run: A monetary approach, *The Scandinavian Journal of Economics* 78, 280–304.

Kremens, Lukas, 2024, Positioning risk, Working paper, University of Washington.

Kremens, Lukas, and Ian W. R. Martin, 2019, The quanto theory of exchange rates, *American Economic Review* 109, 810–843.

Lane, Philip R., and Gian M. Milesi-Ferreti, 2018, The external wealth of nations revisited: International financial integration in the aftermath of the global financial crisis, *IMF Economic Review* 66, 189–222.

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

Lewellen, Jonathan, and Jay Shanken, 2002, Learning, asset-pricing tests, and market efficiency, *Journal of Finance* 57, 1113–1145.

Lloyd, Simon, and Emile Marin, 2021, Exchange rate risk and business cycles, Bank of England working papers 872, Bank of England.

Lustig, Hanno N., Nikolai L. Roussanov, and Adrien Verdelhan, 2011, “Common Risk Factors in Currency Markets,” *Review of Financial Studies* 24: 3731–3777.

Lustig, Hanno N., Nikolai L. Roussanov, and Adrien Verdelhan, 2014, Countercyclical currency risk premia, *Journal of Financial Economics* 111, 527–553.

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

Mankiw, N. Gregory, and Ricardo Reis, 2002, Sticky Information versus sticky prices: A proposal to replace the new Keynesian Phillips curve, *The Quarterly Journal of Economics* 117, 1295–1328.

Martin, Ian W. R., 2017, What is the expected return on the market?, *Quarterly Journal of Economics* 132, 367–433.

Martin, Ian W. R., and Stefan Nagel, 2022, Market efficiency in the age of big data, *Journal of Financial Economics* 145, 154–177.

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

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

Obstfeld, Maurice, and Kenneth S. Rogoff, 2005, Global current account imbalances and exchange rate adjustments, *Brookings Papers on Economic Activity* 36, 67–146.

Pesch, Daniel, Ilaria Piatti, and Paul Whelan, 2024, Subjective risk premia in bond and FX markets, Working paper, Proceedings of the EUROFIDAI-ESSEC Paris December Finance Meeting 2024.

Stavrakeva, Vania, and Jenny Tang, 2023, Individual beliefs, demand for currency and exchange rate dynamics, CEPR Discussion Paper 20321.

Valchev, Rosen, 2020, Bond convenience yields and exchange rate dynamics, *American Economic Journal: Macroeconomics* 12, 124–66.

Verdelhan, Adrien, 2010, A habit-based explanation of the exchange rate risk premium, *Journal of Finance* 65, 123–145.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Appendix S1: Internet Appendix.
Replication Code.