Sticking to Their Guns: Short-Horizon Exchange Rate Expectations

Lukas Kremens Liliana Varela*

March 27, 2026

Abstract

Short-horizon exchange rate forecasts systematically (and incorrectly) predict a partial reversal of their previous errors, both in consensus measures and forecaster-by-forecaster. This simple pattern spans almost two-thirds (two-fifths) of the variation in consensus (individual) forecasts and explains their poor predictive performance at short horizons. We decompose short-term forecasts into three distinct components: (i) a component aligned with long-term forecasts, (ii) a component loading on the past error, and (iii) a residual component. The first two components account for three-quarters of forecast variance and predict exchange rate realizations. However, the second component predicts with a negative sign, offsetting the strong performance of the first component in the overall forecast.

*Lukas Kremens: University of Washington, lkremens@uw.edu. Liliana Varela: London School of Economics and CEPR, l.v.varela@lse.ac.uk. We thank Ric Colacito, Eben Lazarus, Ian Martin, Fabricius Somogyi, and participants at the Junior Valuation Workshop 2025 for helpful comments and suggestions. We are grateful to Francisco Legaspe and Jialiang Lou for excellent research assistance. This work was funded by UK Research and Innovation (UKRI) under the UK government's Horizon Europe funding guarantee [grant number EP/X020916/1].

1 Introduction

The foreign exchange (FX) market is the largest financial market in the world. Exchange rates are a key price for trillions of dollars in cross-border trade and investment flows. Consequently, investors, firms, and financial intermediaries devote considerable resources to research and analysis. Yet, exchange rate predictability is notoriously limited, especially at short horizons (e.g., Meese and Rogoff, 1983), where most market activity takes place.[1] In this paper, we study how professional forecasters form short-horizon exchange rate expectations. We document three striking properties that explain their predictive failure.

First, variation in 1-month forecasts is largely captured by a simple univariate relationship: Forecasts predict a partial reversal of their previous 1-month forecast error over the next month. This pattern, which we term the 'stick-to-their-guns' result ("STTG"), captures over 60% of the variation in consensus forecasts, and almost two-fifths of total variation across individual forecasters, time, and currencies.

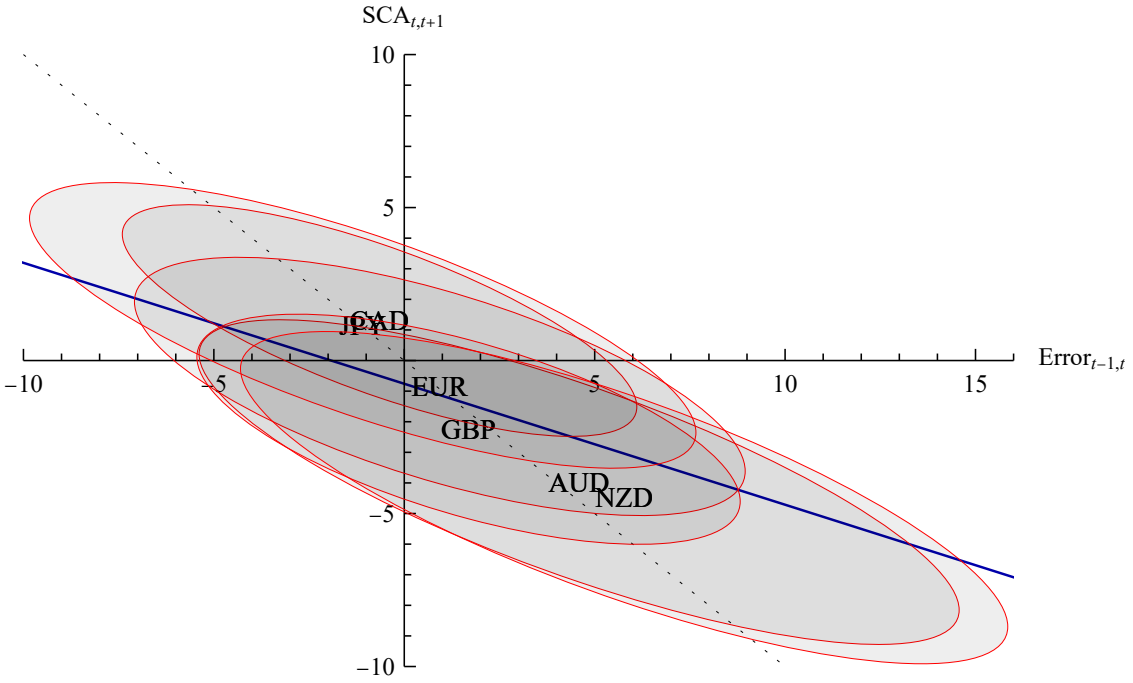
Figure 1 illustrates the basic finding, plotting forecasts (SCA; on the vertical axis) against realized errors (on the horizontal axis) at the 1-month horizon. Rather than making forecasts independently of past errors (data points aligned roughly along the horizontal axis), forecasters tend to be pessimistic for the next month if they were too pessimistic the previous month (a positive error). This holds both across currencies, as indicated by the relative positions of the ellipses, and within currency, as indicated by the orientation of individual ellipses. Indeed, the average relationship illustrated by the blue regression line is almost as far from the benchmark in which forecasts are unpredictable by past errors as it is from another extreme in which forecasts predict a full reversal of past errors (the dotted line representing $y = -x$). Under this benchmark, forecasters echo Michael Burry (alias Christian Bale) in *The Big Short*, insisting *"I may have been early, but I'm not wrong."*[2] On average, forecasts

[1] The foreign exchange (FX) market's trading volume exceeding that of global equity market by more than a factor of fifteen. FX trading volume is overwhelmingly concentrated in over-the-counter derivatives—accounting for roughly 70% of total volume—and at very short horizons, with more than 80% of contracts maturing within three months.

[2] See here.

are roughly half-way between these two extremes. This pattern is pervasive across currencies and time. Using individual forecasts, we find that it is also pervasive across forecasters, rather than arising in consensus forecasts from aggregation over distinct subpopulations of rational agents and would-be Michael Burrys.

Figure 1: Forecasts and forecast errors



Note: For each currency, the figure plots mean 1-month survey forecasts of currency appreciation (SCA$_{t,t+1}$ against lagged 1-month forecast errors (Error$_{t-1,t}$), surrounded by a confidence ellipse whose orientation reflects the time-series correlation between the two variables, and whose size reflects their volatilities (scaled to contain 2% of the observations under joint normality). The horizontal axis also reflects a benchmark under which forecasts are unpredictable from past errors. The solid blue line represents a univariate panel regression with a slope coefficients of -0.40 , while the dotted line is the 45° line on which forecasts predict an exact reversal of past errors.

The apparent stubbornness indicated by this striking pattern is particularly surprising in light of our second result. Forecasters predict several horizons at the same time and an individual forecaster's long-term forecast systematically beats the short-term forecast *at short-term forecasting*. This within-forecaster result holds both in and out of sample.[3] That

[3]Kremens, Martin, and Varela (2025) show that a related result holds in-sample for *consensus* forecasts. Using

is, forecasters have information at their disposal that allows them to make better forecasts than the ones they choose to make.

Why do forecasters seemingly ignore a successful forecasting model in favor of an unsuccessful one? They don't. There is substantial overlap between short- and long-term forecasts. Specifically, the 24-month forecast spans 39% of the total variation in 1-month forecasts. This long-term component in short-term forecasts and a second component that loads only on the past error, together account for three-quarters of 1-month forecast variation.

In our third finding, the above decomposition unearths substantial exchange rate predictability from the ostensibly useless 1-month forecasts. This predictability comes from two distinct—and offsetting—sources. The component shared between long- and short-term forecasts strongly predicts exchange rates at 1-month, both in and out of sample. That is, short-term forecasts load on the 'right' variation of long-term forecasts, the variation that does well in the short-term.

To explain why short-term forecasts do poorly despite the success of the long-term component, it would be sufficient for the second component, reflecting the STTG pattern, to add sufficient noise. It does more damage than that. Unlike Michael Burry in his bet against the U.S. housing market, forecasters are not (on average) rewarded for sticking to their guns—on the contrary, this component systematically pushes forecasts in the *wrong* direction. We view this result as suggestive evidence that the STTG distortion matters for *market prices*. Our findings are consistent with a view of currency markets on which the STTG bias prevents forecasters from adequately reacting to news, which spills over into a delayed market reaction. Through this lens, the STTG pattern in forecasts creates a short-term momentum pattern in exchange rates, and we show that the STTG component in short-term forecasts better captures this pattern than a simple, unconditional autocorrelation in currency movements.

Unchained from one another, the two predictive components forge a formidable short-term forecasting model that beats the random walk benchmark in an out-of-sample test.

a sample of individual forecasts, we show that this holds within forecaster rather than arising from compositional differences in the forecaster pools across horizons. We further extend it to out-of-sample tests.

Tied together, however, the predictive coefficient on the STTG component offsets close to half of the correct forecast from the long-term component. To make matters worse, the residual component, which accounts for around one-quarter of forecast variation, contains no predictive information but adds noise. The overall forecast is left statistically uncorrelated with subsequent outcomes.

Related literature. We relate to a large literature studying expectation formation in asset pricing and international finance. In asset pricing, a recent literature has examined return expectations of professional forecasters (e.g., Greenwood and Shleifer, 2014; Nagel and Xu, 2023; Gandhi, Gormsen, and Lazarus, 2025), price and return patterns associated with biased forecasts (e.g., Bordalo, Gennaioli, La Porta, and Shleifer, 2019; Bouchaud, Krüger, Landier, and Thesmar, 2019; De La O and Myers, 2021), alignment between beliefs and portfolios (e.g., Giglio, Maggiori, Stroebel, and Utkus, 2021; Dahlquist and Ibert, 2024; Andonov, Couts, Goncalves, Loudis, and Rossi, 2025), and mental models of financial forecasters (Bastianello, Decaire, and Guenzel, 2025). We contribute to this literature by providing evidence on expectation formation and its interaction with price formation and predictability in currency markets. Specifically, one of our predictability results points to an explanation for currency momentum that is rooted in distorted expectation formation, thus linking the literature on currency momentum (e.g., Menkhoff, Sarno, Schmeling, and Schrimpf, 2012) with the literature on beliefs-based momentum (e.g. Barberis, Shleifer, and Vishny, 1998; Hong and Stein, 1999). The notion that 1-month forecasts affect market equilibrium is consistent with Bräuer (2026) showing that they strongly predict currency exposure choices of ETF investors.

In international finance, a large literature uses surveys on exchange rate expectations. Frankel and Froot (1987) and Froot and Frankel (1989) use survey-based exchange rate expectations and show that, in advanced economies, survey expectations are inconsistent with Full Information Rational Expectations (FIRE).[4] A recent literature uses the surveys of Consensus Economics on exchange rate expectations to reassess first-order questions in

[4]These findings also hold in other advanced and emerging market economies. See for example Bacchetta, Mertens, and van Wincoop (2009), Chinn and Meredith (2005) and Chinn and Frankel (1999).

international finance, such as uncovered interest rate parity violations, delayed overshooting, international portfolio holdings (De Marco, Macchiavelli, and Valchev 2021; Kalemli-Özcan and Varela 2024; Lloyd and Marin 2020; Candian and De Leo 2025; Candian, De Leo, and Gemmi 2025; Ballensiefen, Somogyi, and Winterberg 2025, among others). Filippou, Li, Liu, and Taylor (2024) show that 3- and 12-month forecasts contain a component that loads on value, momentum, and carry signals. They find that portfolios sorted on this forecast component are profitable. Instead, our aim is to describe how short-horizon forecasts are formed, especially given the sharp contrast between long- and short horizons even within forecaster, in order to explain why they predict so poorly. Two simple components provide the answer: one—overlapping with long-term forecasts and spanning two-fifths of 1-month forecast variation—is a highly successful predictor; the other—the STTG component predicting error reversal and spanning 35% of the variation—largely offsets this success.

Another recent strand in international finance focuses on intermediation of cross-currency flows and intermediary constraints (Gabaix and Maggiori, 2015; Itskhoki and Mukhin, 2025; Gourinchas, Ray, and Vayanos, 2025) in determining exchange rate dynamics. We contribute to this literature by studying exchange rate forecasts produced by these intermediaries. Focusing on advanced economies, Stavrakeva and Tang (2020) map medium-term forecasts to intermediary positions. They also find that 3- and, to a lesser extent, 12-month expectations predict a partial reversal of realized appreciation. Instead, we find that the explanatory power of this pattern more than doubles in short-term forecasts when expressed as a reversal of realized *forecast errors*. This kind of error reversal prediction arises—as, for instance, in Gourinchas and Tornell (2004) for interest rates rather than exchange rates—from overattribution to transitory shocks (see also Augenblick, Lazarus, and Thaler, 2024). Gemmi and Valchev (2026) argue that strategic incentives lead survey forecasts to systematically underreact to common news and overreact to idiosyncratic signals.

Structure of the paper: Section 2 describes the data. Section 3 documents the STTG pattern and the overlap between short- and long-term expectations. Section 4 tests exchange rate predictability using different components of short-term forecasts. Section 5 concludes.

2 Data and definitions

We conduct our analysis using two different datasets on exchange rate expectations at 1-, 3-, 12- and 24-month horizons, both produced by Consensus Economics. The first dataset uses consensus–i.e., the mean across forecasters–exchange rate expectations for 25 currencies from high-income and emerging market economies against the U.S. dollar (Argentina, Australia, Brazil, Canada, China, Chile, Colombia, Euro Area, Hungary, India, Israel, Japan, Korea, Mexico, Malaysia, New Zealand, Peru, Philippines, Poland, Romania, Russia, Singapore, South Africa, Turkey, and the United Kingdom) between June 1997 and February 2025. This dataset brings vast cross-sectional variation across currencies. We complement our analysis with the second dataset that provides expectations at the forecaster level for a group of 6 high-income economies (Australia, Canada, Euro Area, Japan, New Zealand, and the United Kingdom). The small cross-section is set by the data provider but the individual forecasts allows us to zoom in on forecaster behavior and ensure that our results are not affected by compositional differences across consensus measures. To maintain consistency across horizons, we restrict the analysis of the second dataset to forecaster-currency pairs for which forecasts are available at all horizons. Throughout, we refer to the first sample as the "consensus sample" and the second as the "individual sample".

The forecasters surveyed are primarily global banks and institutional investors that actively participate in currency markets. Global banks are key players in the FX market, accounting for around 46% of total trading volume in 2025 (Bank for International Settlements, 2025). The majority of FX trading (69%) takes place in derivatives markets, with a significant share of these instruments traded over the counter by financial intermediaries with their corporate and buy-side clients. These intermediaries maintain research divisions that produce analysis and forecasts to support their sales and trading activities. Consensus Economics collects these forecasts, which constitute our primary source of exchange rate expectations data.

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 and therefore do not violate covered interest parity (CIP). 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 loadings on it, $\beta^{\$}$) following Lustig, Roussanov, and Verdelhan (2014) and the high-minus-low factor (HML) following Lustig, Roussanov, and Verdelhan (2011) (loadings β^{HML}). We use the current account balance scale by GDP as a proxy for capital flows, both obtained from International Financial Statistics (IFS) of the International Monetary Fund (IMF). We obtain the real exchange rate (RER) from the Bank for International Settlements to proxy for deviations from purchasing power parity.

We denote by $e_{j,t}$ the exchange rate in US dollars per unit of currency j in month t, as of the exact date on which Consensus Economics conducts its survey of forecasters. We define the key variables measuring currency appreciation (as expected by survey participants, SCA, and realized, RCA) as follows:

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

$$SCA_{t,t+h} = \frac{1}{\sum_i 1} \sum_i \mathbb{E}_t^i \frac{e_{j,t,t+h}}{e_{j,t}} - 1 \tag{2}$$

$$RCA_{i,t,t+h} = \frac{e_{j,t,t+h}}{e_{j,t}} - 1. \tag{3}$$

The first definition is for forecaster-level expectations, meaning the expectations operator $\mathbb{E}_t^i$ is the subjective expectation of forecaster i, articulated at time t. Throughout the paper, we use the terms expectation and forecast interchangeably, but acknowledge that we cannot know whether an individual forecaster's articulated forecast coincides with her expectation in the sense of a genuinely held belief. The second definition drops the i subscript and refers to the average across forecasters as provided by Consensus Economics in the *consensus sample*, and as constructed by us in the *individual sample*. The third definition is for realized currency appreciation, expressed as a percentage change in the exchange rate between months

t and $t+h$. Given our definition of the exchange rate, $RCA > 0$ implies appreciation of the respective currency against the dollar. In the interest of brevity, we omit a subscripts indicating the currency in all variables but note that we run panel regressions with multiple currencies throughout. We define forecast errors as the difference between the realized and expected exchange rate:

$$\text{Error}_{i,t,t+h} = RCA_{t,t+h} - SCA_{i,t,t+h}. \tag{4}$$

We denote h-month, risk-free interest rate differentials versus the dollar by

$$IRD_{t,t+h} \equiv R^{\$}_{f,t,t+h}/R^{j}_{f,t,t+h} - 1, \tag{5}$$

meaning this variable is negative for currencies in which the interest rate exceeds that for the dollar. Following KMV, we use IRD along with the VIX, the real exchange rate (RER), the current account balance over GDP (CA/GDP), and the factors loadings on "Dollar" and "Carry" ($\beta^{\$}$) and ($\beta^{HML}$) as macro-finance predictors of the exchange rate and potential determinants of survey forecasts. We refer to these variables as "fundamentals."

3 What determines short-horizon expectations?

In this section, we ask which observable variables statistically describe short-horizon exchange rate expectations. In Section 3.1, we perform the analysis at the consensus (i.e., averaged across forecasters) level, and in Section 3.2 we further extend our analysis by exploiting currency expectations at the individual-forecaster level. Lastly, in Section 3.3, we ask how the forecasters' models underlying short-term expectations relates to their model used to form long-term expectations.

3.1 Consensus (Average) Expectations

We first consider 1-month consensus (i.e., average across forecasters) expectations. KMV show that *long*-horizon expectations of currency *excess returns* correlate strongly with a small set of macro-finance predictors but find no such association for short-horizon forecasts. We conduct a similar exercise for *short*-horizon expectations of currency appreciation and, given these prior results, also consider lagged expectations, realizations, and forecast errors, in addition to macro-finance covariates. In a currency-month panel, we estimate

$$SCA_{t,t+1} = \gamma_0 + \gamma_1 RCA_{t-1,t} + \gamma_2 SCA_{t-1,t} + \gamma_3 \text{Error}_{t-1,t} + \gamma_4 X_t + \varepsilon_{t,t+1}, \quad (6)$$

where X_t is a vector of macro-finance variables: IRD, VIX, RER, CA/GDP, $\beta^{\$}$ and β^{HML}.

Column 1 in Table 1 includes only the macro-finance covariates and shows that these do not meaningfully explain consensus forecasts. Only the real exchange rate and the dollar beta carry individually significant coefficients—e.g., a higher real exchange rate is associated with lower expected appreciation, directionally consistent with expected reversal to relative PPP—but the share of total variation explained by these variables, R^2, is only 5.5%.

Column 2 adds lagged 1-month forecasts and lagged 1-month realizations to the right-hand side. R^2 jumps to 64.2% and both regressors enter with strong statistical significance: a higher previous forecast is associated with a higher forecast today, and higher realized appreciation is associated with lower expected appreciation going forward. Noticeably, dropping the macro-finance fundamentals only reduces the R^2 by one percentage point and does not meaningfully change the coefficients on lagged forecasts and realizations (Column 3). Columns 4 and 5 consider these two variables individually, finding that each contributes roughly half of R^2, and that the inclusion of one does not meaningfully change the coefficient of the other.

In column 6, we combine the lagged 1-month forecasts and their realizations into the realized 1-month forecast error, given that their point estimates take similar magnitudes but opposite signs. Forecast error enters with a point estimate negative of -0.496, which

Table 1: Consensus forecasts and errors

	1 month Consensus SCA								Error
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$RCA_{t-1,t}$		-0.455	-0.452	-0.418					
		(0.038)	(0.037)	(0.075)					
$SCA_{t-1,t}$		0.562	0.577		0.531				
		(0.106)	(0.118)		(0.077)				
$\text{Error}_{t-1,t}$						-0.496	-0.485	-0.525	0.495
						(0.035)	(0.027)	(0.043)	(0.025)
$IRD_{t,t+1}$	0.112	0.047							
	(0.104)	(0.064)							
RER_t	-0.202	-0.055							
	(0.029)	(0.018)							
VIX_t	0.003	-0.002							
	(0.002)	(0.001)							
CA/GDP_t	-0.158	0.101							
	(0.179)	(0.056)							
$\beta_t^{\$}$	-0.030	-0.017							
	(0.006)	(0.003)							
β_t^{HML}	0.018	0.005							
	(0.014)	(0.009)							
Within R^2	0.055	0.642	0.632	0.301	0.282	0.622	0.650	0.603	0.246
N	5956	5925	5925	5925	5925	5949	5949	5949	5949
FE	None	None	None	None	None	None	FX	Time	None

Note: This table reports regressions of 1-month Consensus (mean) expectations of exchange rate appreciations on a constant, lagged 1-month realized currency appreciation ($RCA_{t-1,t}$), lagged 1-month survey expectations ($SCA_{t-1,t}$), 1-month interest rate differentials (IRD), real exchange rates (RER), VIX, current account over GDP (CA/GDP), and dollar and carry betas of Lustig, Roussanov, and Verdelhan (2011, 2014). We annualize RCA, SCA, Error, and IRD. The sample runs from 06/1997 to 2/2025 for 25 currencies.

indicates that, on average, forecasters expect a reversal of half of their previous error. That is, forecasters 'stick to their guns' as they appear to (partially) roll-over unsuccessful past forecasts to the next month. Strikingly, this simple STTG pattern captures 62.2% of the variation in 1-month consensus forecasts. We also that the R^2 more than doubles when this pattern is expressed as predicted reversal of realized *errors*, relative to reversal of realized *appreciation* (column 4). Columns 7 and 8 show that this result arises roughly equally within and across currencies.[5]

Lastly, column 9 regresses current forecast errors on lagged forecast errors and shows that these are highly autocorrelated. The R^2 is close to 25%, meaning the STTG pattern in forecasts explains around one-quarter of the variation in 1-month forecast errors. Except for the sign flip, the point estimate is practically identical to that in column 6.

Together, these results suggest that the STTG pattern is detrimental to forecast performance, and that forecasters could reduce the variance of their errors by around one quarter if they were to adopt a random-walk forecast. We return to both of these issues in the subsequent sections.

3.2 Individual Forecaster Expectations

To further investigate these results and ensure that they are not driven by aggregation across forecasters, we estimate analogs of the regressions in Table 1 on forecaster-level data (i.e., in the *individual sample*).[6] We exclude all observations for which $\text{Error}_{i,t-1,t} = -SCA_{i,t,t+1}$. This ensures that we exclude analyst-currency-month observations, for which the target exchange rate is unchanged for two consecutive months. While these may be genuine data points, they could also arise from stale data and would lead to a mechanical STTG result. We report the results in Table 2.

[5]To a lesser extent, the pattern of predicted reversal of appreciation also shows up in 3- and 12-month forecasts (Stavrakeva and Tang, 2020). Table A1 in the appendix reports the 1-month results with standarized variables for easier interpretation of economic magnitudes.

[6]This sample spans only high-income economies so, for comparability, Table A2 in the appendix shows that the consensus results in Table 1 hold similarly in a subsample of high-income economies, that is, in a similar

Table 2: Individual forecasts and forecast errors

	1-month individual forecasts, $SCA_{i,t,t+1}$							Error$_{i,t,t+1}$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$SCA_{i,t-1,t}$	0.487	0.489					0.492		
	(0.019)	(0.020)					(0.018)		
$SCA_{t-1,t}$	-0.118	-0.114	-0.095		-0.106				
	(0.032)	(0.029)	(0.034)		(0.034)				
$RCA_{t-1,t}$	-0.461	-0.455							
	(0.031)	(0.032)							
Error$_{i,t-1,t}$			-0.469	-0.451	-0.460	-0.441		0.483	0.471
			(0.021)	(0.019)	(0.021)	(0.019)		(0.028)	(0.028)
$IRD_{t,t+1}$	0.022								
	(0.030)								
RER_t	0.003								
	(0.005)								
VIX_t	-0.000								
	(0.000)								
CA/GDP_t	0.015								
	(0.027)								
$\beta_t^{\$}$	-0.000								
	(0.002)								
β_t^{HML}	-0.002								
	(0.002)								
Within R^2	0.393	0.389	0.388	0.385	0.376	0.373	0.239	0.230	0.218
N	30042	31197	31197	31197	31190	31190	31197	31197	31190
Forecaster-FX FE	No	No	No	No	Yes	Yes	No	No	Yes
FX-time FE	No	No	No	No	No	No	Yes	No	No

Note: This table reports regressions of 1-month forecaster-level expectations of exchange rate appreciations on a constant, lagged 1-month realized currency appreciation ($RCA_{t-1,t}$), lagged 1-month survey expectations ($SCA_{i,t-1,t}$), lagged 1-month consensus expectations ($SCA_{t-1,t}$), interest rate differentials (IRD), real exchange rates (RER), VIX, current account over GDP (CA/GDP), and dollar and carry betas of Lustig, Roussanov, and Verdelhan (2011, 2014). We annualize RCA, SCA, Error, and IRD. The sample runs from 06/1997 to 2/2025 for 25 currencies.

Column 1 shows that fundamentals, lagged forecasts (individual and consensus), and lagged realizations explain almost 40% of the total variation in 1-month forecasts across currencies, forecasters, and time. Dropping the fundamentals leaves the R^2 unchanged to the percentage point (Column 2), confirming our previous results that fundamentals do not

cross-section of currencies as in the individual sample

meaningfully explain the variation in short-horizons survey expectations either at the consensus or individual forecaster levels. Interestingly, the lagged consensus forecast enters with a negative and statistically significant coefficient, suggesting that conditional on sticking to their own guns—as evidenced by the strong positive coefficient on previous *individual* forecasts—forecasters move against the previous consensus.

We then combine lagged realizations and individual forecasts into the lagged individual error. As for consensus forecast, the individual forecast error is statistically significant and explains a large fraction of the variation of individual short-horizon expectations. In particular, the individual error alone statistically explains 38.5% of the total variation in individual forecasts (column 4), and 37.3% of the time-series variation in each forecaster-currency pair (column 6). The lagged consensus forecast remains statistically significant, but contributes only marginally (around 0.3 percentage points) to the bivariate R^2 (columns 3 and 5); forecasts load much more strongly on individual prior errors than on prior consensus forecasts.

Column 7 estimates the STTG result with respect to deviations from consensus (by adding currency-time fixed effects). The univariate pattern—loading on prior deviations from consensus—explains close to one-quarter of the variation across forecasters. Throughout these different specifications, the point estimate is similar to that in consensus data, close to but just under one-half.[7]

Analogous to the last column of Table 1, columns 8 and 9 show that individual errors are strongly autocorrelated and explain 23.0% and 21.8% of total and, respectively, time-series variation in individual forecast errors.

Importantly, these average results arise from a pervasive pattern across currencies, forecasters and time. We estimate the univariate specifications in columns 4 and 6 of Table 2 (a) for 24-month rolling windows, (b) currency by currency, and (c) forecaster by forecaster. We plot each coefficient on the realized error in Figure 2, along with its 95% confidence interval. All but one point estimates are negative and statistically significant. The only exception is

[7] Table A3 in the appendix reproduces this analysis using standarized variables and shows that a larger association of previous forecasts errors to forecast with point estimate of -0.609 when forecaster-FX FE are included.

a single forecaster in panels C and D, whose coefficient is statistically zero with or without currency fixed effects.

3.3 Short- vs Long-Horizon Expectations

The weak loading of short-term expectations on macroeconomic and financial fundamentals, together with their strong loading on past forecast errors, stands in sharp contrast to the patterns observed in long-horizon expectations (c.f. Kremens, Martin, and Varela, 2025, hereafter "KMV25"). Moreover, KMV find that long-horizon consensus expectations predict exchange rate realizations at both long and short horizons. In consensus forecasts, this could arise from a sharp separation of models across horizons, or from differences in the respective forecaster pools underlying the consensus. We therefore first assess whether a similar result holds within-forecaster.

We use our individual sample and construct consensus estimates among only forecasters that report forecasts for all horizons. We then compare the performance of 1-month SCA with that of 24-month SCA in forecasts of 1-month RCA. To better interpret the magnitudes, we annualize all three objects and estimate

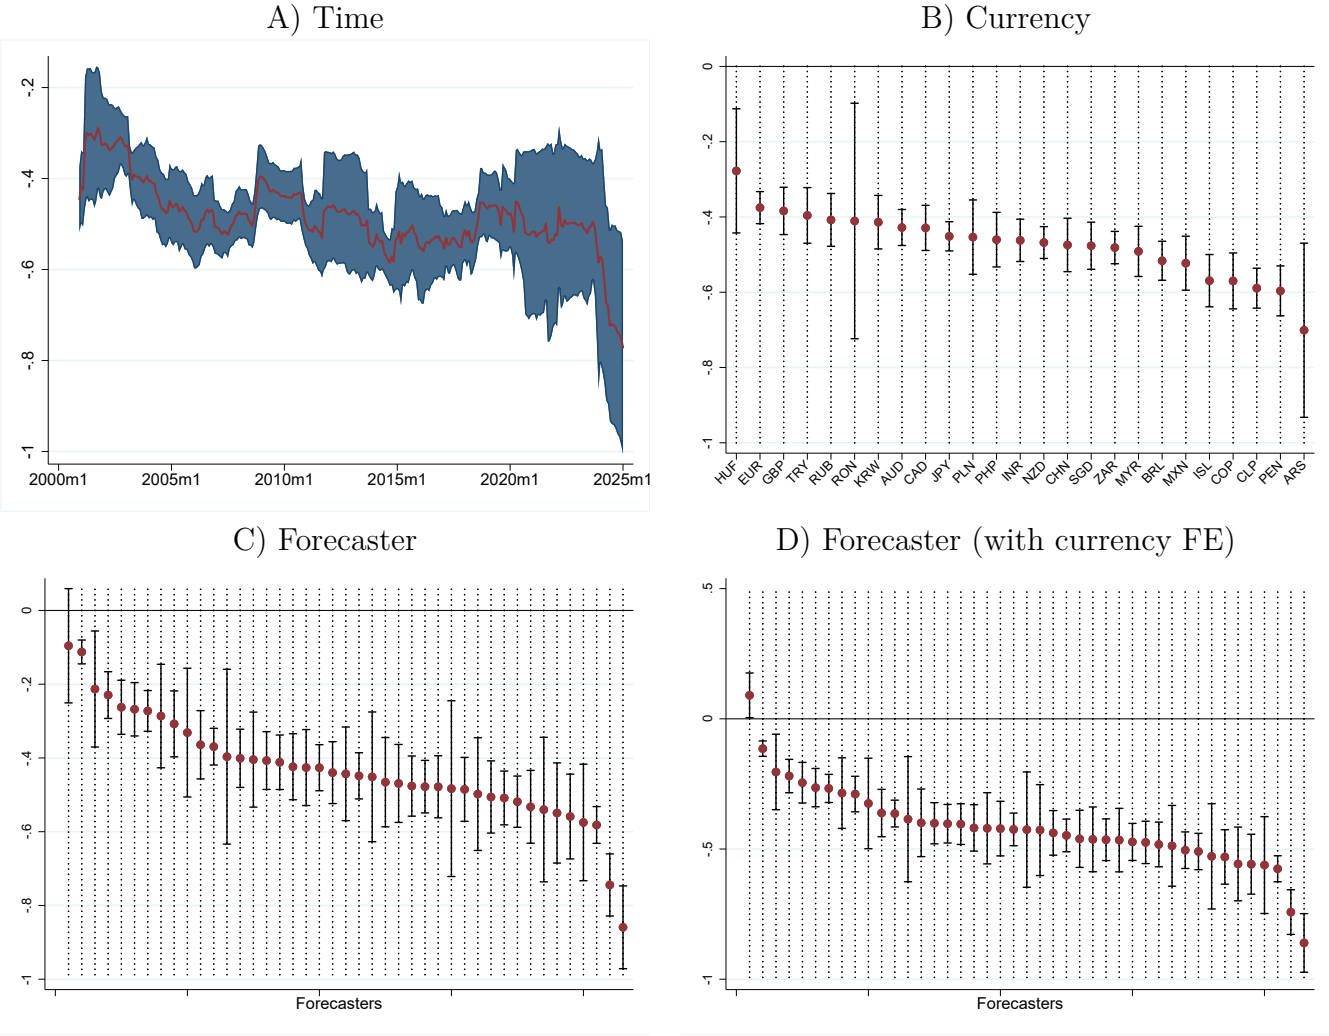
$$RCA_{t,t+1} = \gamma_0 + \gamma_1 SCA_{t,t+1} + \gamma_2 SCA_{t,t+24} + \varepsilon_{t,t+1}, \tag{7}$$

For an out-of-sample test, we compute out-of-sample R^2 for the annualized 1-month prediction following Goyal and Welch (2008):

$$R^2_{OS} = 1 - \frac{\sum_i \sum_t (\epsilon^M_{i,t,t+h})^2}{\sum_i \sum_t (\epsilon^C_{i,t,t+h})^2}, \tag{8}$$

where M is the annualized 1-month forecast from $SCA_{t,t+1}$ evaluated against two different competitor models C: the random walk, that is, a forecast of zero currency appreciation, and the annualized 24-month prediction from $SCA_{t,t+24}$. A positive R^2_{OS} indicates a smaller mean-squared error of the model relative to the competitor. We compute p-values using a

Figure 2: Stick-to-their-guns across time, currencies, and forecasters



Note: Each graph plots the coefficient of $SCA_{t,t+1}$ on $RCA_{t-1,t} - SCA_{t-1,t}$. In panel A, this is a set of regressions of consensus forecasts pooled across currencies for rolling 24-month windows. In panel B, this is a set of time-series regressions of consensus forecasts by currency. In panels C and D, this is a set of regressions of individual forecasts pooled across currencies. Relative to panel C, panel D includes a currency fixed effect. Each coefficient is accompanied by a 95% confidence interval.

non-parametric bootstrap procedure as in KMV.[8] Table 3 reports the results.

Table 3: ST vs LT forecasts in short-term forecasting

	Dependent Variable: $RCA_{t,t+1}$				
	In-sample			Out of sample	
	(1)	(2)	(3)	(4)	(5)
$SCA_{t,t+h}$	-0.019		-1.063		
	(0.060)		(0.843)		
$SCA_{t,t+24}$		0.862	1.113		
		(0.419)	(0.476)		
$SCA_{t,t+1}$ vs RW				-0.463	
				[0.000]	
$SCA_{t,t+1}$ vs $SCA_{t,t+24}$					-0.471
					[0.000]
R^2	0.000	0.007	0.013		
N	1775	1775	1775	1775	1775

Note: Columns 1–3 report regressions of annualized 1-month exchange rate appreciation on a constant and annualized consensus forecasts made for 1-and 24-month horizons. Standard errors in parentheses are clustered by currency and date. Columns 4 and 5 report out-of-sample R^2 measures following Goyal and Welch (2008) for annualized 1-month SCA against, respectively, the random walk (column 4) and against annualized 24-month SCA (column 5). In brackets, we report bootstrapped p-values computed following Footnote 8. The sample constructs consensus measures from individual forecasts, restricting to forecasters who report all horizons, for 6 high-income economies. The sample period runs from 06/1997 to 2/2025.

Column 1 confirms that short-horizon forecasts do not predict 1-month realizations, but column 2 shows that long-horizon forecasts do—with a point estimate indistinguishable from one, but statistically different from zero. Column 3 runs a horse race confirming that long-

[8] We draw, with replacement, blocks with a cross-sectional width uniformly distributed between 2 and the number of currencies, and time-series length equal to the forecasting horizon. (We extend the time-series block length to 6 and 12 months to account for potential persistence in predictors but note that this does not meaningfully change the p-values.) 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. The bootstrapped p-value is the fraction of bootstrap draws for which the model performance is reversed and the model with the higher full-sample MSE beats the one with the lower full-sample MSE. Conceptually, this is similar to a Diebold–Mariano test (Diebold and Mariano, 1995) of the null hypothesis that the forecasts perform equally well for all currencies. Unlike the Diebold-Mariano procedure, which is designed for time-series data, this bootstrap accounts for cross-sectional correlation.

horizon forecasts clearly outperform. Columns 4 and 5 take this exercise out of sample and the negative out-of-sample R^2 confirm that 1-month forecasts perform poorly. They are resoundingly beaten by the random walk, and by 24-month forecasts.

Since these forecasts are made at the same time by the same pool of institutions, the information underlying the 24-month forecasts is clearly available to the forecasters making the 1-month forecasts. This raises the question whether forecasters incorporate some of the predictive information from 24-month expectations in their short-horizon forecasts, and whether any such overlap generates predictive success.

We assess the overlap between 1- and 24-month expectations by decomposing 1-month consensus expectations into three orthogonal pieces: (i) a component spanned by 24-month expectations, (ii) a component loading on the realized 1-month error but orthogonal to 24-month expectations, and (iii) a residual component. Our analysis proceeds in the following steps. We first regress (annualized) 1-month expectations on (annualized) 24-month expectations,

$$SCA_{t,t+1} = \theta_0^{LT} + \theta_1^{LT} SCA_{t,t+24} + \theta_t^{LT} \tag{9}$$

and define the LT component as the fitted value:

$$LT_t \equiv \hat{\theta}_0^{LT} + \hat{\theta}_1^{LT} SCA_{t,t+24}. \tag{10}$$

We then orthogonalize the lagged error with respect to annualized long-term forecasts

$$\text{Error}_{t-1,t} = \eta_0 + \eta_1 SCA_{t,t+24} + \eta_t. \tag{11}$$

We then regress the residuals from regression equation (9) on the orthogonalized 1-month errors, that is, the residuals from (11)

$$\theta_t^{LT} = \theta_1^E \hat{\eta}_t + \theta_t^E. \tag{12}$$

We then define the "STTG component" as the fitted value from (12).

$$STTG_t \equiv \hat{\theta}_1^E \times \hat{\eta}_t. \tag{13}$$

By construction, this is orthogonal to LT_t. Lastly, we define the "residual component" as

$$Res_t \equiv SCA_{t,t+1} - LT_t - STTG_t. \tag{14}$$

We find that the LT component captures 39% of the variance in 1-month expectations. The "STTG component" captures another 35% of the variation, with the residual accounting for the remaining 26%. Table A4 reports these variance decompositions, and (9), (11), and (12).

The STTG result in 1-month expectations is striking in isolation: A pattern that can be condensed into a single variable as simple as the realized forecast error explains close to two-thirds of the variation in consensus expectations, two-fifths of individual forecasts, one-quarter of individual deviations from consensus, and close to one quarter of variation in *forecast errors* (both for the consensus and deviations from it). But short-term forecasts also contain substantial overlap with long-term forecasts that have been shown to predict well even at short horizons. Given this overlap, we now turn to assess exchange rate predictability at the 1-month horizon.

4 Exchange rate predictability

Using the decomposition developed in the previous section, we revisit whether exchange rates are predictable in the short term, in sample (Section 4.1) and out of sample (Section 4.2). Throughout this section, we account for sampling uncertainty in the predictors generated via the above decomposition by bootstrapping standard errors in currency clusters.

4.1 In-sample forecasting

We regress monthly realized currency appreciation first on its consensus forecast and then on the orthogonal components of consensus forecasts constructed in the previous section: LT, $STTG$, and Res.

$$RCA_{t,t+1} = \gamma_0 + \gamma_1 SCA_{t,t+1} + \varepsilon_{t,t+1}, \tag{15}$$

$$RCA_{t,t+1} = \gamma_0 + \gamma_1 LT_{t,t+1} + \gamma_2 STTG_{t,t+1} + \gamma_3 Res_{t,t+1} + \varepsilon_{t,t+1}, \tag{16}$$

Column 1 of Table 4 shows the result, familiar from the literature, that short-term expectations largely fail to predict short-term exchange rate appreciation. The point estimate is economically close to and statistically indistinguishable from zero.

The specification in column 2 uses only the LT component in short-term expectations and show that this predicts exchange rates with a strongly significant coefficient of 0.382—albeit still far from unity—rises almost fourfold relative to the raw forecast in column 1. The R^2 jumps to 3.3%, rising more than five-fold compared to column 1. This result indicates that survey expectations contain predictive information about exchange rates even at the horizons as short as one month.

Column 3 uses only the STTG component. While this, too, has a strongly significant coefficient, it is negative, indicating that the tendency of forecasters to stick to their guns as documented in Section 3 offsets the informative component in their forecasts. In terms of magnitude, the coefficient is close to half of that in column 2. Given that the two orthogonal components have similar variance, and the overall forecast coefficient in column 1 is the variance weighted average of the coefficients in columns 2–4, this indicates that the significant negative predictability from the STTG component is largely to blame for the economically (and statistically) zero correlation between the overall forecast and subsequent realizations.

Column 4 uses only the residual component. This has close to zero explanatory power and point estimate of economically and statistically zero. This adds noise to the overall forecast, but unlike, the STTG component, does not offset the economic magnitude of the predictive

Table 4: Currency forecasting

	$RCA_{t,t+1}$							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$SCA_{t,t+1}$	0.104							
	(0.102)							
LT_t		0.382			0.382	0.391	0.382	
		(0.083)			(0.083)	(0.075)	(0.109)	
$STTG_t$			-0.173		-0.173		-0.182	
			(0.039)		(0.039)		(0.100)	
Res_t				0.058	0.058	0.067	0.058	
				(0.127)	(0.127)	(0.123)	(0.135)	
$RCA_{t-1,t}$						0.064	-0.004	
						(0.017)	(0.047)	
$SCA_{t,t+24}$								1.840
								(0.188)
$\text{Error}_{t-1,t}$								0.070
								(0.024)
R^2	0.006	0.033	0.006	0.001	0.039	0.037	0.039	0.039
Adjusted R^2	0.006	0.033	0.006	0.000	0.039	0.037	0.039	0.039
N	5949	5949	5949	5949	5949	5925	5925	5949

Note: This table reports regressions of 1-month realized exchange rate appreciations on survey currency appreciations, a long-term component (LT component), a stick to your guns component (STTG component), a residual and lagged realized currency appreciations. Standard errors in parentheses are clustered by currency and time in columns 1 and 8, and bootstrapped by currency clusters in columns 2–7 to account for sampling uncertainty in the generated regressors. The sample runs from 06/1997 to 2/2025. All variables are annualized. The sample spans 25 currencies.

information contained in the LT component.

Column 5 combines all three pieces. These add up to the raw forecasts in column 1, but now take three freely estimated coefficients. Given the three components are orthogonal, their point estimates are identical to those in their respective univariate specifications in columns 2–4, and the R^2 of 3.9% is the sum of across those three specifications.

The majority of this R^2 comes from the LT component (3.3%). While the predictable share is lower for the STTG component, the monthly R^2 of 0.6% is comparable to those of frequently used equity market return predictors like the earnings-to-price ratio or the term spread (Campbell and Thompson, 2008).

Overall, the R^2 in short-horizon predictive regression warrants thoughtful benchmarking. Explanatory power in return prediction is naturally bounded in competitive markets since even small R^2 imply large economic gains. In the words of Campbell and Thompson (2008) (p. 1526): "Regressions with large R^2 statistics would be too profitable to believe. The saying 'If you're so smart, why aren't you rich?' applies with great force here, and should lead investors to suspect that highly successful predictive regressions are spurious. Since the squared Sharpe ratio and average real interest rate increase in proportion with the investment horizon, however, larger R^2 statistics are believable at longer horizons." In this light, the 3.3% for the LT component is especially striking. For comparison, this exceeds the in-sample R^2 of even the most successful equity market return predictors like the consumption-wealth ratio of Lettau and Ludvigson (2001), which reaches 2.60% in Campbell and Thompson (2008).

Taking stock of these short-term predictability results in Table 4, we note that

1. Short-horizon expectations *do* contain substantial predictive information.
2. The STTG bias documented in Section 3 is the key distortion that needs to be stripped out of 1-month forecasts to render them successful.
3. This STTG bias not only contributes noise crowding out predictive information. It individually predicts currency appreciation (in the wrong direction).

The first result is striking given the widespread finding—both in previous literature and in column 1—that survey forecasts generally do poorly (Nagel and Xu, 2023; Dahlquist and Söderlind, 2023), and that exchange rates are difficult to predict at short-horizons, even for sophisticated econometric models (Meese and Rogoff, 1983; Rossi, 2013; Engel and Wu, 2023). Instead, forecasters appear to *have* predictive information about short-horizon currency appreciation, but this is drowned out by other distortions in their forecasts.

The second result identifies the predominant distortion and echoes a large literature in behavioral macroeconomics documenting 'underreaction' patterns in forecasts of macro variables (e.g., Coibion and Gorodnichenko, 2015; Angeletos, Huo, and Sastry, 2020). The third result warrants further inspection with respect to exchange rate determination. It suggests

that the bias in forecast formation may also affect price formation in currency markets. The result says that currencies tend to appreciate when they have previously appreciated by more than forecasters had expected, and when the same forecasters predict this unexpected appreciation to partially revert over the coming month. That is, the result is consistent with a model of currency markets in which professional forecasters tend to treat unexpected appreciation as transitory rather than persistent, and in which other market participants either independently share this view or are influenced by it. As a result, the market exchange rate initially only partially reflects new information and fully incorporates it only with delay, creating short-term momentum.

Columns 6 and 7 of Table 4 test whether this result simply repackages a more generic momentum pattern in 1-month currency appreciation. Column 6 replaces the STTG component with lagged currency appreciation. The coefficient is positive and significant, indicating that momentum generally exists. But the R^2 is smaller than with the STTG component in column 5, suggesting that the momentum-like pattern in the data is better described by the STTG distortion in short-term forecasts than by autocorrelation in exchange rate movements. Lastly, column 7 includes both predictors at the same time: the coefficient on lagged RCA is driven to zero, while that on the STTG component is largely unchanged compared to column 5, as is the R^2, reinforcing the conclusion that the predictive power from the STTG forecast distortion is more subtle than simple autocorrelation in price changes.

Column 8 confirms the central result of the decomposition exercise more directly. Rather than projecting short-term forecasts onto long-term forecasts and short-term errors, we directly use these variables as predictors. Both are significant positive predictors. For the short-term error, this means unexpected appreciation predicts further appreciation—consistent with the results for the STTG component in columns 3 and 5.

In Table A5, we further show that the results continue to hold with time-fixed effects. This cross-sectional manifestation implies that the baseline results are not driven by a few global shocks or time periods associated with particular dollar movements.

4.2 Out-of-sample forecasting

To further substantiate the first and third result, we turn to out-of-sample forecasting. Using the first ten years of the sample, until (excluding) June 2007, we run regressions (9), (11), and (12) to construct subsample analogs of the LT and $STTG$ components following, respectively, definitions (10) and (13). Over the same sample, we regress $RCA_{t,t+1}$ on the constructed LT and $STTG$ components. This bivariate regression has an R^2 of 1.6% in this training sample of 1514 observations, with a positive coefficient of 0.268 on the LT component and a negative coefficient of -0.206 on the STTG component. We then compute fitted values $\widehat{RCA}^{ST}_{t,t+1}$ for the remaining sample. This contains 4396 non-missing observations across 22 currencies.

Mirroring the in-sample specification in column (7) of Table 4, we proceed analogously for a competitor model based on short-term momentum. Here, the forecasts are obtained from a regression of 1-month RCA on its 1-month lag over the same training sample. The fitted values are denoted by $\widehat{RCA}^{1M}_{t,t+1}$. We add four more competitor models: a carry model using 1-month IRD to predict ($\widehat{RCA}^{IRD}_{t,t+1}$), a 2-factor model using IRD and 1-month momentum ($\widehat{RCA}^{2F}_{t,t+1}$), a 3-factor model further adding RER ($\widehat{RCA}^{3F}_{t,t+1}$), and a UIP forecast using IRD to predict with a fixed coefficient of 1 ($\widehat{RCA}^{UIP}_{t,t+1}$.

We then compute the out-of sample 1-month currency appreciation forecast for the different models. The random walk (RW) has a constant forecast of zero, and forecast error equal to $\epsilon^{RW}_{t,t+1} = RCA_{t,t+1}$. The other models have forecast errors equal to $\epsilon^{M}_{t,t+1} = RCA_{t,t+1} - \widehat{RCA}^{M}_{t,t+1}$ for $M \in \{ST, 1M, IRD, 2F, 3F, UIP\}$. We then evaluate out-of-sample R^2 following (8) for model $M \in \{ST\}$ and competitor model $C \in \{RW, 1M, IRD, 2F, 3F, UIP\}$. As shown in Table 5, the disaggregated ST model beats all models with out-of-sample R^2 around 3% and p-values < 0.05. The only exception is the UIP model which generates smaller but very volatile errors, such that its performance is not statistically distinguishable from that of the disaggregated ST model.

Short-horizon expectations fail to forecast exchange rates. This section shows that this is predominantly due to a component that predicts reversal of previous errors (the STTG

Table 5: Out-of-sample short-run forecasting

Competitor model (→)	RW	1M	IRD	2F	3F	UIP
ST	0.026	0.030	0.033	0.031	0.029	-0.065
	[0.024]	[0.003]	[0.009]	[0.007]	[0.002]	[0.231]

Note: This table reports out-of-sample R^2 measures as defined in (8), for the disaggregated short-term model (ST) against six competitor models: random walk (RW), 1-month momentum (1M), carry (IRD), 2-factor (2F), 3-factor (3F), and UIP, all as described in the text. Forecasts are for 1-month currency appreciation between 6/2007 and 2/2025. Models are estimated over a training sample spanning 6/1997 to 5/2007. The sample is for the 25 currencies of the consensus sample. In brackets, we report bootstrapped p-values computed following Footnote 8.

result). But expectation formation is more subtle: forecasters *do* have access to and incorporate information that strongly predicts exchange rates even at short horizons. This component overlaps with long-term forecasts but is offset by the STTG distortion, which predicts with the wrong sign.

This last result, negative predictability from the STTG component suggests that this type of distortion affects market prices. Either because forecasters influence other market participants, or because other market participants suffer from similar biases, the exchange rate appears to underreact to news, leading to short-term momentum. The STTG pattern captures this particular momentum pattern better than simple autocorrelation in exchange rate movements, indicating that the predictability result from STTG is not due to spurious correlation with an unrelated momentum pattern.

5 Conclusion

Our findings shed light on the expectation formation process in currency markets. The STTG pattern—the striking tendency to predict a partial reversal of realized errors—captures a large share of the variation in short-horizon forecasts and, simultaneously, explains why these forecasts predict so poorly. Stripping out the associated component and focusing instead on a component shared with long-term expectations reveals strong exchange rate predictability

even at the 1-month horizon.

Moreover, we find that forecast errors themselves positively predict exchange rates. This result is consistent with a delayed full incorporation of news into prices arising from a distortion in short-term forecasts that systematically underreacts to news that are at odds with previous forecasts. That said, we do not identify the source of the STTG distortion. A pattern like this may arise from behavioral biases or strategic incentives in a signaling or reputation game among analysts (cf. Gemmi and Valchev, 2026). Our data do not let us conclusively pin down the precise motivation that produces this strikingly pervasive and distortive pattern so we leave this to future work.

References

Andonov, A., S. Couts, A. S. Goncalves, J. Loudis, and A. Rossi (2025): "Subjective Beliefs and the Portfolio Allocations of Institutional Investors," *Working Paper*.

Angeletos, G.-M., Z. Huo, and K. A. Sastry (2020): "Imperfect Macroeconomic Expectations: Evidence and Theory," Working Paper 27308, National Bureau of Economic Research.

Augenblick, N., E. Lazarus, and M. Thaler (2024): "Overinference from Weak Signals and Underinference from Strong Signals*," *The Quarterly Journal of Economics*, 140(1), 335–401.

Bacchetta, P., E. Mertens, and E. van Wincoop (2009): "Predictability in Financial Markets: What Do Survey Expectations Tell Us?," *Journal of International Money and Finance*, 28(3), 406–426.

Ballensiefen, B., F. Somogyi, and H. Winterberg (2025): "Demand for Dollars: Evidence from Survey Expectations," *Working Paper*.

Bank for International Settlements (2025): "OTC foreign exchange turnover in April 2025," *Monetary and Economic Department*.

Barberis, N., A. Shleifer, and R. Vishny (1998): "A Model of Investor Sentiment," *Journal of Financial Economics*, 49(3), 307–343.

Bastianello, F., P. Decaire, and M. Guenzel (2025): "Mental Models and Financial Forecasts," *Working Paper*.

Bordalo, P., N. Gennaioli, R. La Porta, and A. Shleifer (2019): "Diagnostic Expectations and Stock Returns," *The Journal of Finance*, 74(6), 2839–2874.

Bouchaud, J.-P., P. Krüger, A. Landier, and D. Thesmar (2019): "Sticky Expectations and the Profitability Anomaly," *The Journal of Finance*, 74(2), 639–674.

Bräuer, L. (2026): "Exchange Rate Expectations and Currency Demand," *Working Paper*.

Campbell, J. Y., and S. B. Thompson (2008): "Predicting Excess Stock Returns Out of Sample: Can Anything Beat The Historical Average?," *Review of Financial Studies*, 21, 1509–1531.

Candian, G., and P. De Leo (2025): "Imperfect Exchange Rate Expectations," *The Review of Economics and Statistics*, 107(5), 1406–1423.

Candian, G., P. De Leo, and L. Gemmi (2025): "Managing Expectations with Exchange Rate Policy," *Working Paper*.

Chinn, M., and J. Frankel (1999): "More survey data on exchange rate expectations: More currencies, more horizons, more tests," International Finance 9508003, Monetary Policy, Capital Flows and Financial Market Developments in the Era of Financial Globalisation: Essays in Honour of Max Fry, edited by Bill Allen and David Dickinson.v.

Chinn, M. D., and G. Meredith (2005): "Testing Uncovered Interest Parity at Short and Long Horizons during the Post-Bretton Woods Era," Working Paper 11077, National Bureau of Economic Research.

Coibion, O., and Y. Gorodnichenko (2015): "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review*, 105(8), 2644–78.

Dahlquist, M., and M. Ibert (2024): "Equity Return Expectations and Portfolios: Evidence from Large Asset Managers," *The Review of Financial Studies*, 37(6), 1887–1928.

Dahlquist, M., and P. Söderlind (2023): "Individual Forecasts of Exchange Rates," *Working Paper*.

De La O, R., and S. Myers (2021): "Subjective Cash Flow and Discount Rate Expectations," *Journal of Finance*, 76(3), 1339–1387.

De Marco, F., M. Macchiavelli, and R. Valchev (2021): "Beyond Home Bias: International Portfolio Holdings and Information Heterogeneity," *The Review of Financial Studies*, 35(9), 4387–4422.

Diebold, F. X., and R. S. Mariano (1995): "Comparing Predictive Accuracy," *Journal of Business and Economic Statistics*, 13(3), 253–263.

Engel, C., and S. P. Y. Wu (2023): "Forecasting the U.S. Dollar in the 21st Century," *Journal of International Economics*, 141, 103715.

Filippou, I., J. Li, X. Liu, and M. P. Taylor (2024): "Short-Horizon Currency Expectations," *Working Paper*.

Frankel, J. A., and K. A. Froot (1987): "Using Survey Data to Test Standard Propositions Regarding Exchange Rate Expectations," *The American Economic Review*, 77(1), 133–153.

Froot, K. A., and J. A. Frankel (1989): "Forward Discount Bias: Is it an Exchange Risk Premium," *Quarterly Journal of Economics*, 104(1), 139–161.

Gabaix, X., and M. Maggiori (2015): "International Liquidity and Exchange Rate Dynamics," *Quarterly Journal of Economics*, 130(3), 1369–1420.

Gandhi, M., N. J. Gormsen, and E. Lazarus (2025): "Forward Return Expectations," *Working Paper*.

Gemmi, L., and R. Valchev (2026): "Biased surveys," *Journal of Monetary Economics*, 157, 103868.

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus (2021): "Five Facts about Beliefs and Portfolios," *American Economic Review*, 111(5), 1481,Äì1522.

Gourinchas, P., W. D. Ray, and D. Vayanos (2025): "A Preferred,ÄêHabitat Model of Term Premia, Exchange Rates, and Monetary Policy Spillovers," *American Economic Review*, 115(11), 3788–3824.

Gourinchas, P.-O., and A. Tornell (2004): "Exchange rate puzzles and distorted beliefs," *Journal of International Economics*, 64(2), 303–333.

Goyal, A., and I. Welch (2008): "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction," *Review of Financial Studies*, 21(4), 1455–1508.

Greenwood, R., and A. Shleifer (2014): "Expectations of Returns and Expected Returns," *Review of Financial Studies*, 27(3), 714–746.

Hong, H., and J. C. Stein (1999): "A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets," *Journal of Finance*, 54(6), 2143–2184.

Itskhoki, O., and D. Mukhin (2025): "Mussa Puzzle Redux," *Econometrica*, 93(1), 1–39.

Kalemli-Özcan, S., and L. Varela (2024): "Five Facts About the UIP Premium," *NBER Working Paper 28923*.

Kremens, L., I. W. R. Martin, and L. Varela (2025): "Long-Horizon Exchange Rate Expectations," *Journal of Finance*, 80(6), 3695–3724.

Lettau, M., and S. Ludvigson (2001): "Consumption, Aggregate Wealth, and Expected Stock Returns," *The Journal of Finance*, 56(3), 815–849.

Lloyd, S., and E. Marin (2020): "Exchange rate risk and business cycles," Bank of England working papers 872, Bank of England.

Lustig, H. N., N. L. Roussanov, and A. Verdelhan (2011): "Common Risk Factors in Currency Markets," *Review of Financial Studies*, 24(11).

——— (2014): "Countercyclical Currency Risk Premia," *Journal of Financial Economics*, 111(3), 527–553.

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

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2012): "Currency Momentum Strategies," *Journal of Financial Economics*, 106(3), 660–684.

Nagel, S., and Z. Xu (2023): "Dynamics of Subjective Risk Premia," *Journal of Financial Economics*, 150(2), 103713.

Rossi, B. (2013): "Exchange Rate Predictability," *Journal of Economic Literature*, 51(4), 1063–1119.

Stavrakeva, V., and J. Tang (2020): "Deviations from FIRE and Exchange Rates: a GE Theory of Supply and Demand," *Working Paper*.

A Supplementary Tables

Table A1: Consensus Forecasts and Errors (Standarized Variables)

	1-month consensus forecasts, $SCA_{t,t+1}$					
	(1)	(2)	(3)	(4)	(5)	(6)
$RCA_{t-1,t}$		-0.597	-0.594	-0.549		
		(0.049)	(0.049)	(0.099)		
$SCA_{t-1,t}$		0.562	0.577		0.531	
		(0.106)	(0.118)		(0.077)	
$\text{Error}_{t-1,t}$						-0.788
						(0.055)
$IRD_{t,t+1}$	0.169	0.070				
	(0.156)	(0.096)				
RER_t	-0.147	-0.040				
	(0.021)	(0.013)				
VIX_t	0.066	-0.042				
	(0.039)	(0.021)				
CA/GDP_t	-0.029	0.018				
	(0.032)	(0.010)				
$\beta_t^{\$}$	-0.089	-0.050				
	(0.018)	(0.010)				
β_t^{HML}	0.049	0.013				
	(0.037)	(0.024)				
R^2	0.055	0.642	0.632	0.301	0.282	0.622
N	5956	5925	5925	5925	5925	5949

Note: This table reports regressions of annualized 1-month Consensus (mean) expectations of exchange rate appreciations on a constant, lagged 1-month realized currency appreciation ($RCA_{t-1,t}$), lagged 1-month survey expectations ($SCA_{t-1,t}$), interest rate differentials (IRD), real exchange rates (RER), VIX, current account over GDP (CA/GDP), and dollar and carry betas of Lustig, Roussanov, and Verdelhan (2011, 2014). Variables are standarized. The sample runs from 06/1997 to 2/2025 for 25 currencies.

Table A2: Consensus forecasts and errors (High-income economies)

	1 month consensus forecasts, $SCA_{t,t+1}$						$\text{Error}_{t,t}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$RCA_{t-1,t}$		-0.457	-0.455	-0.453			
		(0.021)	(0.022)	(0.022)			
$SCA_{t-1,t}$		0.392	0.397		0.392		
		(0.029)	(0.030)		(0.038)		
$\text{Error}_{t-1,t}$						-0.438	0.465
						(0.020)	(0.035)
$IRD_{t,t+1}$	1.001	0.154					
	(0.439)	(0.271)					
RER_t	-0.189	-0.010					
	(0.090)	(0.058)					
VIX_t	0.001	-0.001					
	(0.002)	(0.001)					
CA/GDP_t	-0.029	0.224					
	(0.363)	(0.186)					
$\beta_t^{\$}$	-0.012	0.009					
	(0.025)	(0.016)					
β_t^{HML}	-0.027	-0.015					
	(0.034)	(0.018)					
R^2	0.035	0.666	0.660	0.500	0.156	0.658	0.218
N	2207	2193	2193	2194	2193	2200	2192

Note: This table reports regressions of annualized 1-month Consensus (mean) expectations of exchange rate appreciations on a constant, lagged 1-month realized currency appreciation ($RCA_{t-1,t}$), lagged 1-month survey expectations ($SCA_{t-1,t}$), interest rate differentials (IRD), real exchange rates (RER), VIX, current account over GDP (CA/GDP), and dollar and carry betas of Lustig, Roussanov, and Verdelhan (2011, 2014). The sample runs from 06/1997 to 2/2025 for high-income economies currencies.

Table A3: Individual forecasts and forecast errors (Standarized Variables)

	1-month individual forecast, $SCA_{i,t,t+1}$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$SCA_{i,t-1,t}$	0.487	0.489					0.492
	(0.019)	(0.020)					(0.018)
$SCA_{t-1,t}$	-0.070	-0.068	-0.057		-0.064		
	(0.019)	(0.017)	(0.020)		(0.020)		
$RCA_{t-1,t}$	-0.425	-0.419					
	(0.028)	(0.029)					
$\text{Error}_{i,t-1,t}$			-0.648	-0.623	-0.636	-0.609	
			(0.029)	(0.026)	(0.029)	(0.026)	
$IRD_{t,t+1}$	0.019						
	(0.026)						
RER_t	0.013						
	(0.022)						
VIX_t	-0.032						
	(0.026)						
CA/GDP_t	0.017						
	(0.030)						
$\beta_t^{\$}$	-0.002						
	(0.020)						
β_t^{HML}	-0.028						
	(0.034)						
Within R^2	0.393	0.389	0.388	0.385	0.376	0.373	0.240
N	30044	31199	31199	31199	31192	31192	31199
Forecaster-FX FE	No	No	No	No	Yes	Yes	No
FX-time FE	No	No	No	No	No	No	Yes

Note: This table reports regressions of 1-month forecaster-level expectations of exchange rate appreciations on a constant, lagged 1-month realized currency appreciation ($RCA_{t-1,t}$), lagged 1-month survey expectations ($SCA_{t-1,t}$), interest rate differentials (IRD), real exchange rates (RER), VIX, current account over GDP (CA/GDP), and dollar and carry betas of Lustig, Roussanov, and Verdelhan (2011, 2014). Variables are standarized. The sample runs from 06/1997 to 2/2025 for 25 currencies.

Table A4: Decomposing 1-month consensus forecasts

	$SCA_{t,t+1}$	$\text{Error}_{t-1,t}$	θ_t^{LT}	LT_t	$STTG_t$	Res_t
	(1)	(2)	(3)	(4)	(5)	(6)
$SCA_{t,t+24}$	4.061	-4.121				
	(0.328)	(0.973)				
η_t			-0.405			
			(0.018)			
$SCA_{t,t+1}$				0.388	0.348	0.265
				(0.076)	(0.100)	(0.030)
R^2	0.388	0.158	0.567	0.388	0.348	0.265
N	5956	5949	5949	5956	5949	5949

Note: This table reports the regressions decomposing 1-month consensus forecasts into orthogonal LT, $STTG$, and Res components (columns 1–3). Columns 4–6 report regressions to estimate the variance share of each component as a univariate regression coefficient on the full 1-month forecast. Standard errors in parentheses are clustered by currency and time. The sample runs from 06/1997 to 2/2025 for 25 currencies.

Table A5: Cross-sectional currency forecasting

	$RCA_{t,t+1}$							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$SCA_{t,t+1}$	0.162							
	(0.093)							
LT_t		0.377			0.368	0.375	0.370	
		(0.078)			(0.078)	(0.076)	(0.097)	
$STTG_t$			-0.137		-0.103		-0.079	
			(0.048)		(0.046)		(0.094)	
Res_t				0.110	0.100	0.104	0.102	
				(0.136)	(0.137)	(0.135)	(0.142)	
$RCA_{t-1,t}$						0.043	0.014	
						(0.022)	(0.044)	
$SCA_{t,t+24}$								1.673
								(0.183)
$\text{Error}_{t-1,t}$								0.041
								(0.021)
Within R^2	0.017	0.044	0.004	0.002	0.048	0.047	0.047	0.046
N	5949	5949	5949	5949	5949	5925	5925	5949

Note: This table reports regressions analogous to those in Table 4, adding time fixed effects. Standard errors in parentheses are clustered by currency and time in columns 1 and 8, and bootstrapped by currency clusters in columns 2–7 to account for sampling uncertainty in the generated regressors. The sample runs from 06/1997 to 2/2025. All variables are annualized. The sample spans 25 currencies.