Individual Forecasts of Exchange Rates

Magnus Dahlquist Paul Söderlind*

First draft: February 2022
This version: August 2023

Abstract

We study the expectations of individual forecasters in the foreign exchange market. The expectations depart from rationality such that survey risk premia are acyclical in contrast to countercyclical rational risk premia. We find that forecasters learn from their own forecast errors (rather than from consensus forecast errors) and that they tend to overreact when forming expectations (as indicated by their forecast revisions). However, we find little support for the sticky and noisy information models that motivate the typical overreaction specification. Finally, while forecasters have worse forecasting performance relative to simple benchmarks, the forecasters who emphasize the real exchange rate and do not overreact have better out-of-sample forecasting performance. Overall, our results highlight the information contained in individual (rather than consensus) exchange rate forecasts.

Keywords: Beliefs, currencies, expectations, foreign exchange rates, predictability.
JEL codes: D84, F31, F37, G12, G15.

*We greatly benefited from discussions with Nina Karnaukh at an early stage of the project. We thank Julien Pénasse, Joshua Traut, and Sonya Zhu for comments and suggestions. Dahlquist gratefully acknowledges support from the Jan Wallander and Tom Hedelius Foundation. Dahlquist: Stockholm School of Economics and CEPR; e-mail: magnus.dahlquist@hhs.se. Söderlind: University of St. Gallen; e-mail: paul.soderlind@unisg.ch.

1 Introduction

Expectations are a core concern in asset pricing and for understanding exchange rates. An early literature using surveys on exchange rates found that consensus forecasts are biased and contain little information about future exchange rates (see, e.g., Dominguez, 1986; Frankel and Froot, 1987; Froot and Frankel, 1989). That and a more recent literature characterizes survey risk premia based on consensus forecasts, relates them to rational risk premia obtained from predictability regressions, and attempts to understand the formation of exchange rate expectations (see, e.g., Bacchetta, Mertens, and van Wincoop, 2009; Nagel and Xu, 2023).

In this paper, we highlight the information contained in individual rather than consensus exchange rate forecasts. We study a monthly unbalanced panel containing 21 forecasters' expectations of the USD versus six currencies (AUD, CAD, CHF, EUR, GBP, and JPY) over 3- and 12-month horizons during the 1992–2019 period. We consider what information is contained in survey expectations of future exchange rates, the formation process of these expectations, and what these expectations convey about risk premia. The individual forecasts allow us to study what is learned from individual forecast errors, overreaction in individual forecast revisions, and individual forecasting ability—information otherwise hidden in consensus forecasts. We also characterize individual survey-expected risk premia.

We begin by confirming that survey expectations are biased and contain little information about future exchange rates. Panel regressions of the realized depreciation rate on the survey-expected depreciation rate show no significant relationship and the hypothesis of unbiasedness can easily be rejected. Regressions of the realized depreciation rate on the forward premium (corresponding to the risk-neutral expected depreciation rate) reveal a negative relationship. The result that the forward premium negatively predicts the depreciation rate is consistent with an extensive literature on deviations from uncovered interest rate parity

(UIP) and time-varying currency risk premia (Fama, 1984). An alternative view is that this result reflects a violation of rational expectations (Froot and Frankel, 1989). Regardless, while the survey-expected depreciation rate is positively related to the forward premium, the results suggest significant differences in risk-neutral and survey expectations.

Under full information and rational expectations, forecast errors should be unpredictable. That is, the difference between the realized and survey-expected risk premium should be unpredictable. Equivalently, we compare rational risk premia with survey risk premia. We predict future currency excess returns from the current forward premium, the current real exchange rate (capturing long-term reversals), and the recent USD depreciation rate (capturing short-term reversals). These predictors form the basis of several trading strategies, including carry, value, and momentum/reversal strategies (see, e.g., Asness, Moskowitz, and Pedersen, 2013; Burnside, Eichenbaum, and Rebelo, 2011; Daniel, Hodrick, and Lu, 2017; Koijen, Moskowitz, Pedersen, and Vrugt, 2018; Lustig, Roussanov, and Verdelhan, 2011; Menkhoff, Sarno, Schmeling, and Schrimpf, 2012, 2017).

Consistent with the literature, we find that both the forward premium and the real exchange rate predict future currency excess returns (see, e.g., Chernov and Creal, 2021; Dahlquist and Pénasse, 2022). The fitted values in the predictability regressions are a measure of rational risk premia and they are countercyclical. However, the survey risk premium is less countercyclical than is the rational risk premium, primarily because it does not relate to the forward premium; moreover, the survey risk premium instead relates to short-term reversals. Our results hold for the consensus forecast as well as individual forecasts. That the survey risk premium is not related to the forward premium echoes the findings of Bacchetta, Mertens, and van Wincoop (2009) and Nagel and Xu (2023), who use consensus survey data.

.

We finally explore the individual forecasts in more detail, obtaining three main findings. First, panel regressions suggest that survey expectations relate negatively to recent individual, but not consensus, forecast errors; hence, forecasters seem to learn from their own, but not others', forecast errors. This learning subsumes the relationship between survey expectations and short-term reversals. Second, individual forecast errors relate negatively to revisions in individual forecasts. According to the framework of Coibion and Gorodnichenko (2015), this indicates "expectations stickiness"—individual forecasters tend to overreact to news and revise their forecasts too much. However, in further analysis we find little support for the sticky and noisy information models used to motivate the overreaction specification. Third, forecasters underperform simple benchmarks (average and panel regressions) in out-of-sample forecasting. However, forecasters who rely more on the real exchange rate and overreact to a lesser extent in forecast revisions tend to perform better. With these findings, we contribute to the literature by highlighting the information regarding individual expectations and the formation of these expectations.

This paper relates to a vast literature that finds that foreign exchange forecasts are biased and depart from rationality; see Dominguez (1986), Frankel and Froot (1987), Froot and Frankel (1989), Ito (1990), and Chinn and Frankel (1994) for early contributions, and Ince and Molodtsova (2017) for a later analysis. These scholars predominantly consider consensus forecasts. An exception is Ito (1990), who analyzes individual forecasters of the USD–JPY exchange rate over two years. Interestingly, he finds that there are time-invariant individual effects. In our panel regressions, we allow for such heterogeneity via currency and forecaster fixed effects. We also consider individual forecasting ability.

This paper also relates to the literature on survey expectations and risk premia; see, for example, Buraschi, Piatti, and Whelan (2021), Cieslak (2018), Giacoletti, Laursen, and Singleton (2021), Piazzesi, Salomao, and Schneider (2015), and Singleton (2021) on bonds;

Adam, Matveev, and Nagel (2021), Dahlquist and Ibert (2023), and Greenwood and Shleifer (2014) on stocks; and Bacchetta, Mertens, and van Wincoop (2009) and Nagel and Xu (2023) on several asset classes including currencies. In contemporaneous work to us, Nagel and Xu (2023) find that the consensus survey risk premium on a portfolio of developed market currencies is acyclical, in sharp contrast to the rational risk premium obtained from predictive regressions using the currencies' average forward premium (following Lustig, Roussanov, and Verdelhan, 2014) or various business cycle variables as predictors. Similarly, we also find that the survey risk premium is close to acyclical and not countercyclical like the rational risk premium, primarily because it is not related to the forward premium. We also emphasize the great dispersion in individual survey risk premia. Relatedly, Candian and De Leo (2023), Gourinchas and Tornell (2004), Molavi, Tahbaz-Salehi, and Vedolin (2023), and Valente, Vasudevan, and Wu (2022) consider how interest rate expectations matter for currency risk premia and puzzling features of exchange rates; Stavrakeva and Tang (2023) find that macroeconomic surprises, measured from survey forecasts, capture exchange rate movements.

Finally, this paper relates to the literature on expectation formation in finance and macro. We apply the approach of Coibion and Gorodnichenko (2015) to measure expectation stickiness and deviations from rational expectations through the link between forecast errors and past forecast revisions. Coibion and Gorodnichenko (2015) study consensus forecasts of inflation and other macroeconomic variables and find evidence of underreaction to information relative to rational expectations (consistent with information rigidities). Bordalo, Gennaioli, Ma, and Shleifer (2020) study the rationality of individual and consensus forecasts of similar variables, but not foreign exchange rates. They find that individual forecasters typically overreact to news, while consensus forecasts underreact to news relative to rational expectations. Importantly, Stavrakeva and Tang (2020) use both individual and consensus forecasts on exchange rates to characterize deviations from rational expectations and find that the deviations can help resolve exchange rate puzzles. While we also document deviations from rational expectations, we find little evidence in individual exchange rate forecasts in favor of the noisy and sticky information models used in Coibion and Gorodnichenko (2015) to motivate the typical overreaction specification.

The paper proceeds as follows. Section 2 presents and summarizes the data. Section 3 considers common tests of rational expectations and individual forecasting performance. Section 4 studies individual forecasters' learning and overreaction as underlying mechanisms for the deviation from a rational benchmark. Section 5 discusses implications for risk premia and considers common trading strategies in the foreign exchange market. Finally, Section 6 concludes.

2 Data

2.1 Data description

Our main dataset consists of monthly 3-, 6-, and 12-month exchange rate survey forecasts from Blue Chip Financial Forecasts (BCFF) from January 1993 to December 2018 for the USD versus six currencies: AUD, CAD, CHF, EUR (DEM before 1999), GBP, and JPY. The survey contains exchange rate forecasts from 21 large financial institutions and rating agencies.[1] Similar to Bordalo, Gennaioli, Ma, and Shleifer (2020), who use the BCFF survey forecasts of interest rates, we digitize individual forecasts from PDF publications.[2] When

[1]ScotiaBank, The Industrial Bank of Japan, ING Financial Markets, Deutsche Bank, Merrill Lynch & Co., Moody's Analytics, Moody's Capital Markets, Nomura Securities, BMO Capital Markets, WestLB, Barclays Capital, Oxford Economics, Wells Fargo, HSBC Markets, BNP Paribas Americas, IHSMarkit, UBS, Moody's Economy.com, AIG, TS Lombard, and Société Générale.

[2]Alternative data sources for survey forecasts of exchange rates are FX4casts and Consensus Economics. They offer subscriptions of consensus (but not individual) forecasts for major currencies since 1986 and 1989, respectively. To our knowledge, only (Stavrakeva and Tang, 2020, 2023), via the Federal Reserve Bank of

needed, we transform the survey data so that the exchange rates are defined as the number of USD per one foreign currency unit—that is, an increase in the spot exchange rate corresponds to a USD depreciation relative to the foreign currency.

The forecasts are collected over the last days of the month and the survey is published on the first day of the next month. We assume that the forecasts are based on information available on the last day of the month. This assumption implies that the first forecasts in the sample were made at the end of December 1992 and refer to the spot exchange rates at the end of March, July, and December 1993 (3-, 6-, and 12-month forecasts). Similarly, the last forecasts in the sample were made at the end of November 2018 and refer to spot exchange rates at the end of February, June, and November 2019. (In a robustness analysis, we find that the results are little affected by assuming that the forecasts were made earlier than the end of a month.) We also collect individual forecasts of future interest rates.

Figure 1 shows the number of active forecasters over time. The total number is 21, but many forecasters contribute to the survey for only parts of the sample period. In the 1990s and 2000s, 3–5 forecasts are available for each currency; in the 2010s, the number increases to 10–12. In our analyses, each observation has the same weight, which implies that a month with fewer forecasters in the survey has a lower effective weight than does a month with many forecasters. (In a robustness analysis, we find that the results are little affected by giving every month equal weight.)

We retrieve daily spot and forward exchange rates for the January 1976–December 2019 period from *Datastream*, and consumer price indices (CPIs) and industrial production from the Organisation for Economic Co-operation and Development (OECD); see Chernov, Dahlquist,

Boston, have had access to a subset of more than 30 institutions' individual forecasts that contribute to the consensus forecasts of Consensus Economics. Kalemli-Ozcan and Varela (2023) study consensus forecasts but also collect five institutions' individual forecasts from printed Consensus Economics reports from 2001 to 2018.

and Lochstoer (2023) for more details. We let s_t and $f_{k,t}$ denote the log spot exchange rate and k-period forward exchange rate at date t, respectively. Note that the forward exchange rate is the risk-neutral expectation at t of the future exchange rate at $t+k$. We refer to $f_{k,t} - s_t$ as the forward premium. We define the real exchange rate as the spot exchange rate times the foreign CPI level divided by the US CPI level. We let q_t denote the log real exchange rate at t. To keep the notation simple, we suppress a specific notation for a currency i.

Two key variables in the analysis are the depreciation rate of the USD relative to the foreign currency, $s_{t+k} - s_t$, and the currency excess return, $r_{t,t+k} = s_{t+k} - f_{k,t}$. The excess return is a scaled payoff from entering a k-period forward contract at t and selling the currency at $t+k$. If covered interest rate parity (CIP) holds, this is the same as borrowing at a USD interest rate at t, then buying the foreign currency and lending at the foreign interest rate, and finally selling the foreign currency at $t+k$.

From the survey we calculate the expected depreciation rate by an individual forecaster j, $\mathrm{F}_t^j s_{t+k} - s_t$. Under full-information rationality, the forecast would correspond to the conditional expectation, $\mathrm{E}_t s_{t+k} - s_t$. We use forecasts and survey expectations interchangeably but distinguish them from rational expectations. We also calculate the corresponding survey-expected excess return, that is, the survey risk premium, $\mathrm{F}_t^j r_{t,t+k} = \mathrm{F}_t^j s_{t+k} - f_{k,t}$. These definitions assume that current spot and forward exchange rates are known by the forecasters. We assume that the log of the expected spot rate approximates the expected log spot rate (i.e., we disregard a Jensen's inequality term), but none of our results is sensitive to this assumption.

Two other key variables in the analyses are the forecast error, $\mathrm{FE}_{t,t+k}^j = s_{t+k} - \mathrm{F}_t^j s_{t+k}$, and the forecast revision, $\mathrm{FR}_{t-k,t}^j = \mathrm{F}_t^j s_{t+k} - \mathrm{F}_{t-k}^j s_{t+k}$, of an individual forecaster j. The forecast error is simply the difference between the log spot exchange rate and a previous survey forecast, and the forecast revision is the update of the forecast of the same future exchange rate. Finally, note that as the forward exchange rate is known at t, the forecast error of the exchange rate is the same as that of the excess return; that is, we also have that $\mathrm{FE}^j_{t,t+k} = r_{t,t+k} - \mathrm{F}^j_t r_{t,t+k}$. The forecast revision of the excess return requires an expectation of the future forward rate that we can infer from the forecasters' expectations of future spot exchange rates and interest rates under the assumption of CIP.

For some regressions we consider consensus forecasts/errors, which we calculate as the cross-sectional averages of the available individual forecasts/errors.

Finally, we complement the exchange rate, CPI, and industrial production data with J.P. Morgan's global FX implied volatility index (VXY), retrieved from *Bloomberg*.

2.2 Summary statistics

Table 1 reports summary statistics of some of the variables. Panel A is for the 3-month horizon ($k = 3$) and Panel B is for the 12-month horizon ($k = 12$). All statistics are cross-sectional averages; for example, the minimum survey risk premium is calculated by first finding the minimum for each forecaster and currency, and then averaging across forecasters and currencies. The samples are matched so that the realized depreciations refer to only those dates for which we have survey expectations. For instance, the last depreciation is at the end of February 2019 for the 3-month horizon, but at the end of November 2019 for the 12-month horizon.

For the 3-month horizon, the means of the realized depreciation rates and excess returns are close to zero, whereas the means of the expected depreciation rates are significantly negative.

This carries over to the survey risk premium and forecast error. For the 12-month horizon, the means are further from zero, which is reasonable since the statistics are not annualized, but none is statistically significant.

The statistics that describe the volatility and shape of the distribution (skewness, kurtosis, minimum, and maximum) all suggest that the survey expectations are more stable than the realized data—as would be the case of rational expectations (according to which the realized variable equal the expected variable plus noise, with these two parts being uncorrelated). However, there is considerable variation in expectations, both over time and across forecasters. The autocorrelations display a similar pattern, but with much stronger persistence in the survey expectations.

2.3 Forward premia and survey-expected depreciations

Figure 2 shows spot exchange rates (black solid lines) along with 3-month, 6-month, and 12-month forward exchange rates (red lines with squares) and 3-month, 6-month, and 12-month consensus forecasts (blue lines with circles). While forward exchange rates and survey expectations in relation to the spot exchange rate often point in the same direction, they are sometimes completely disconnected. Hence, the forward premium and the survey-expected depreciation rate are distinct.

To gain further insights, we run a simple panel regression of the survey-expected depreciation rate on the forward premium with currency and forecaster fixed effects (untabulated results). The slope coefficient estimate for the 3-month horizon is 1.6, which is significantly different from zero at usual significance levels and significantly different from one at the 7% significance level. The slope coefficient estimate for the 12-month horizon is 1.4, which is significantly different from one at the 11% level. The R^2 values in the regressions are low: 2.5%

and 10% for the 3- and 12-month horizons, respectively. The figure reports corresponding slope coefficients (in the range 0.88–1.72) and R^2 values (in the range 4–12%) for individual currencies. The visual impression from Figure 2 is confirmed: survey expectations typically move in the same direction as does the forward premium, but that most movements in the survey expectations are unrelated to the forward premium.

Figure 3 shows the 12-month forward premium and a shaded band indicating the interquartile range of individual 12-month survey-expected depreciation rate. The figure highlights that the survey-expected depreciation rate tends to move in the same direction as does the forward premium, but that most of the movements in the survey expectations are unrelated to the forward premium. The figure also highlights considerable cross-sectional dispersion of the survey-expected depreciation rate.

This preliminary evidence shows that survey expectations are very different from risk neutral expectations (forward premia), and that there are large differences between forecasters. In the next section, we relate this to important currency market variables (e.g., the real exchange rate and recent exchange rate movements). We also investigate how differences across forecasters affect their forecasting performance.

3 Rational expectations tests

We consider two common tests of rationality: (1) unbiasedness—the survey-expected deprecation rate should provide an unbiased predictor of the future realized depreciation rate; (2) forecast-error unpredictability—the difference between the realized currency excess return and the survey risk premium should be uncorrelated with all information available at the time of forecast (including the past forecast error and the forecast revision). We complement the survey unbiasedness test with a classic forward unbiasedness test.

3.1 Unbiasedness

We begin by evaluating whether forward exchange rates (measuring risk neutral expectations) and survey expectations are unbiased forecasts of future change in the exchange rate.

We run the following panel regressions of the USD depreciation rate on either the forward premium or the survey-expected depreciation rate:

$$s_{t+k} - s_t = a + b\,(f_{k,t} - s_t) + \varepsilon_{t+k}, \tag{1}$$

$$s_{t+k} - s_t = a^j + b\,(\mathrm{F}_t^j s_{t+k} - s_t) + \epsilon_{t+k}^j. \tag{2}$$

While we suppress currency i in the notation, regression (1) is a two-dimensional (time and currency) panel with currency fixed effects, and regression (2) is a three-dimensional (time, currency, and forecaster) panel with currency and forecaster fixed effects. There are six currencies and 21 forecasters.

Regression (1) is a so-called Fama (1984) regression and regression (2) is a so-called Mincer and Zarnowitz (1969) regression used to test for forecast rationality. In both regressions, unbiasedness implies that the slope coefficient is equal to one (i.e., $b = 1$). Unbiasedness thus means that the depreciation rate moves one to one with either the forward premium or the survey-expected depreciation rate. We focus on the slope coefficients and not on the fixed effects as we in the end are mainly interested in the dynamics of rational and survey risk premia.

The data are overlapping (monthly observations of depreciations over three and twelve months), which creates serial correlation; there are also likely cross-sectional correlations.

We therefore use standard errors that are robust to both serial and cross-sectional correlation (Driscoll and Kraay, 1998, with Hansen and Hodrick, 1980, weights up to k lags). Taking into account serial and cross-sectional correlation is of considerable importance for making inference in our regressions throughout the paper.

Table 2 reports the results of the regressions. Specification I shows that the estimates of the slope coefficients for the forward premium are negative and significantly different from one. That is, the forward exchange rate is not an unbiased predictor of the future exchange rate. This result indicates deviations from the UIP and, under rational expectations, it translates into a currency risk premium, $\mathrm{E}_t r_{t,t+k} = \mathrm{E}_t s_{t+k} - f_t$. To interpret this further, recall that the forward premium equals the interest rate differential (US minus foreign) if CIP holds. The negative slope coefficient then suggests that currencies with high interest rates tend to appreciate, which is the basis of the forward premium puzzle as well as the carry trade (see, e.g., Fama, 1984; Bansal and Dahlquist, 2000; Hassan and Mano, 2019). Specification II shows the results for the survey expectations. The estimates of the slope coefficients for the survey-expected depreciation rate are close to zero and significantly different from one (i.e., the survey expectations are not unbiased, either). Moreover, the survey expectations have lower predictive power than does the forward premium.

The results of the unbiasedness regressions suggest significant differences between risk-neutral expectations (driving forward premia) and survey expectations (driving the expected depreciation rate), both on average and for individual forecasters. The next subsection studies this in more detail.

3.2 Forecast error and forecast revision

Rational expectations suggest that forecast errors are serially uncorrelated. To evaluate this, we run the following panel regression of the forecast error on its lagged forecast error:

$$\mathrm{FE}^j_{t,t+k} = a^j + b\,\mathrm{FE}^j_{t-k,t} + \varepsilon^j_{t+k}. \tag{3}$$

Importantly, the forecast errors on the left- and right-hand sides are not overlapping. Under rational expectations, the forecast error should not be predictable by the forecast error known at t. The null hypothesis is therefore that there is no serial correlation (i.e., $b = 0$). Table 3 reports the results. Specification I shows that the forecast error exhibits significant positive serial correlation for the 3-month horizon; the serial correlation is also positive, but statistically insignificant, on the 12-month horizon.

Rational expectations also suggest that forecast errors are unpredictable by information variables at the time of the forecast. We estimate the following panel regressions of the forecast error:

$$\mathrm{FE}^j_{t,t+k} = a^j + b' x_t + \varepsilon^j_{t+k}, \tag{4}$$

where x_t contains the forward premium, the real exchange rate, and the recent USD depreciation rate as information variables (and evaluate whether $b = 0$). Specification II in Table 3 reports the results. The forward premium and recent depreciation rate predict forecast errors on the 3-month horizon, and the real exchange also predicts on the 12-month horizon, which we interpret as deviations from rational expectations. The R^2 values are 5.9% and 14.0%. Specification III combines the previous two specifications and shows similar results (though the lagged forecast error drives out the recent depreciation rate at the 3-month horizon).

We next use the framework of Coibion and Gorodnichenko (2015) and study the predictability of forecast errors from forecast revisions. An advantage with this framework is that we do not need to specify the variables in a forecaster's information set and what prompts a revision. We run panel regressions of the future individual forecast error on the individual forecast revision:

$$\mathrm{FE}^{j}_{t,t+k} = a^{j} + b\,\mathrm{FR}^{j}_{t-k,t} + c'x_t + \varepsilon^{j}_{t+k}, \tag{5}$$

where we again allow for the inclusion of control variables in x_t. The framework is motivated by information rigidities (Mankiw and Reis, 2002): If the forecast revision predicts the forecast error with a positive coefficient, this implies underreaction in the expectations; if the forecast revision predicts the forecast error with a negative coefficient, this implies overreaction in the expectations. To calculate the revision over the 3-month horizon, we compare the current 3-month forecast with the 6-month forecast made 3 months before. In the absence of 24-month forecasts, we proxy the revision of the 12-month forecast by the revision of the 6-month forecast (i.e., the current 6-month forecast compared with the 12-month forecast made six months before).

Table 3 presents the results. In specification IV, the forecast revision negatively predicts the forecast error. For the 3-month horizon, the estimated coefficient is –0.062 and statistically significant at the 5% level; for the 12-month horizon, the estimated coefficient is –0.041 but imprecisely estimated. The table also present results when the forward premium and the real exchange rate are included as control variables. (We exclude the depreciation rate as it is strongly correlated with the revision and this multicollinearity makes it otherwise difficult to interpret coefficient magnitudes.) The coefficient estimates on revisions are slightly more negative and statistically significant when including the control variables in specification V. The negative coefficient estimates suggest that there is overreaction in the individual forecasts: An increase in an individual forecaster's revision predicts a lower forecast error, suggesting that individuals revise their forecasts too much and that this contributes to the predictability of the individual forecasting error. However, the results seem weaker than those reported in the literature on other variables than exchange rates (see, e.g., Bordalo, Gennaioli, Ma, and Shleifer, 2020; Li, Nieuwerburgh, and Renxuan, 2023); see also Stavrakeva and Tang (2020) for mixed evidence on the sign of the revision coefficient estimate for individual and consensus exchange rate forecasts of different currencies. Moreover, when we apply a two-stage instrument variable approach with lagged revision, forecast error, forward premium, real exchange rate, and depreciation rate as instruments, the coefficient estimates are positive and insignificant. Importantly, the significance of other variables also suggests misspecification of the basic overreaction specification. To understand the belief formation and the the underlying mechanism, we dig deeper on this in Section 4.

3.3 Individual beliefs and forecasting performance

In this subsection we evaluate the out-of-sample forecasting performance by comparing mean-squared forecast errors (MSEs) across models and forecasters. Table 4 reports the results.

Panel A compares the performance of the consensus forecast with out-of-sample forecasts from a panel regression of the currency excess return, $r_{t,t+k}$, on information variables, x_t, recursively estimated (see Appendix A.1). For the panel regression, we use the sample period January 1976 to December 1992 as the first training sample and make forecasts for the 3- and 12-month horizons. Then, we add January 1993 to the training sample and make new out-of-sample forecasts, and so forth. The panel reports the ratio of the MSEs aggregated across all currencies, with a ratio above one indicating that the survey forecast performs worse than the forecast from the recursive panel regression. We use a Diebold and Mariano (1995) test, based on a robust variance-covariance estimator, to evaluate whether the MSEs are the same. The main result is that the survey forecasts are significantly worse than forecasts from the recursive panel regression, especially at the 3-month horizon.[3] We have also compared the survey forecasts with simpler forecast benchmarks, such as the historical average and the forecast from a regression with only the forward premium. We found that the survey forecasts are also worse than these benchmarks (untabulated results).

How do individual forecasters fare versus consensus and the panel regression? To characterize the forecasters, we run forecaster-specific regressions and collect forecaster-specific coefficient estimates. Figure 4 shows scatters of estimates for individual forecasters at the 3-month horizon. The first scatter (in brown) is from a regression of the realized depreciation rate on the survey-expected deprecation rate, similar to the unbiasedness regression (2). The estimates are gathered around zero (close to the panel estimates in Table 2). The next scatter (in red) is from regressions of the survey risk premium on the forward premium and shows a great dispersion in estimates (with most weight in the range 0.0–3.0, aligned with the previously reported panel estimate of 1.6). The next three scatters (in blue, green, and magenta) are from regressions of the survey risk premium on the real exchange rate, recent depreciation rate, and recent forecast error, respectively. The last scatter (in black) is from a regression of the forecast error on the forecast revision. The average negative estimates is consistent with the negative estimate in the panel regression (5) but now the great dispersion is also conveyed. The 12-month estimates show a similar pattern.

Panel B compares the performance of different forecasters. We classify the forecasters as "high" and "low" according to the coefficient estimates above. The panel shows the MSEs of forecasters with high (above median) estimates divided by the MSEs of those with low (below median) estimates. Several results stand out for the 3-month horizon. First, forecasters with expectations that correlate positively with the forward premium make worse forecasts than do others; in other words, expectations following UIP perform worse. Second, forecasters

[3] Recently, Kremens, Martin, and Varela (2023) find that consensus expectations are more successful at forecasting on the longer 24-month horizon.

who do not believe in long-term reversal (a negative coefficient of the real exchange rate) also make worse forecasts. Third, forecasters who overreact (more negative coefficient of the revision term) make worse forecasts. We also notice that the degree of learning from mistakes (based on the lagged forecast error) makes little difference in the out-of-sample forecasting performance. The results on the 12-month horizon tend to be qualitative the same but with little statistical significance.

Figure 5 summarizes our forecasting performance results using the survey forecasts versus the panel regressions. It shows the cumulated MSE on a 3-month horizon (for all exchange rates over the three-month horizon) minus that of the panel regressions. The blue solid line represents the consensus forecasts and shows that the panel regression forecast beats the consensus forecast. The shaded dark blue band represents the interquartile range across forecasters and the shaded light blue band the min-max across forecasters. They show that the consensus forecast beats the majority of the individual survey forecasts. It underscores the value of model averaging in forecasting.

4 Understanding belief formation

In this section we characterize the expectations formation of future exchange rates. We use the individual forecasts to analyze the underlying mechanism of information rigidities and learning, linking back to the deviations from rational expectations in the previous section. We also consider the time-series variation in the cross-sectional dispersion of individual forecasts.

4.1 Overreaction

In this subsection, we evaluate whether sticky or noisy information models can explain the predictability of forecast errors. More specifically, we use individual forecasts to analyze the dynamics of forecast errors, the stickiness of forecasts, and noisy forecasts.

The analysis in Subsection 3.2 showed that the link between revisions and subsequent forecast errors is weakly negative, indicating overreaction. We now consider the framework of Kučinskas and Peters (2022) to estimate composite bias coefficients. We apply local projection panel regressions of the impulse response function of 3-month forecast errors. (In shorthand notation, the forecast error e_{t+l} is regressed on (e_t, e_{t-3}, e_{t-6}) for $l = 3, 6, 9, 12$ and the impulse response function is the series of coefficients on e_t). The results show that the first non-overlapping lag, the 3-month lag, has a coefficient of 0.276 and it is strongly significant. However, all further non-overlapping lags $(6, 9, 12)$ give insignificant estimates close to zero. This suggests that the forecasting errors have much more limited autocorrelations than models of occasional (sticky information) or partial updating (noisy information) imply. That is, the forecast errors look more like an MA(1) than an AR(1) process.

The key assumption in the sticky information model is that forecasters only update their forecasts now and then. In addition, it is often assumed that when they do update, they are close to being rational. The consensus forecast then looks like an AR(1) transformation of a rational forecast. Using data on individual forecasters, we study both these assumptions.

The fraction of stale forecasts (a revision of zero, that is, $\mathrm{F}_t^j r_{t,t+k} - \mathrm{F}_{t-k}^j r_{t,t+k} = 0$) is 11.7% on the 3-month horizon and 6.3% on the 6-month horizon. This suggests that stickiness is limited, or, as Farmer, Nakamura, and Steinsson (2023) argue: "professional forecasters pay attention constantly and have precise knowledge of the data in question."

It is still of interest to study whether the forecasters update to something close to rational, when they do update. Specification I in Table 5 reports regressions of the survey risk premium on the rational risk premium (the recursively estimated panel regressions, previously used in Subsection 3.3 and described in Appendix A.1). This is clearly an approximation of a truly rational forecast, but appears to be a reasonable starting point, especially given its strong forecasting performance. For the 3-month horizon, specification I in Table 5 suggests that the survey and rational forecasts are almost unrelated. In contrast, for the 12-month horizon the coefficient is 0.214 (and strongly significant), although far from one. The results are similar if we instead consider, say, the largest 75% of the revisions (in absolute terms), that is, observations when the revisions are indeed non-trivial.

Overall, the findings about the sticky information model are that we find little stickiness and that updating does not bring the forecasts much closer to a rational forecast.

The key assumption in a noisy information model is that today's forecast is formed as a (Bayesian) average of new information, y_t, and the previous forecast (see, e.g., Bouchaud, Krüger, Landier, and Thesmar, 2019; Coibion and Gorodnichenko, 2015):

$$\mathrm{F}_t^j r_{t,t+k} = \lambda y_t + (1-\lambda)\mathrm{F}_{t-k}^j r_{t,t+k}. \tag{6}$$

Note that the two forecasts are for the excess return, $r_{t,t+k}$, but made at different points in time. Using data on individual forecasters allows us to study this. As before, we approximate the new information with the rational recursive forecast. For the 3-month horizon, the results in specification II of Table 5 are completely at odds with the noisy information model since the estimate on the rational forecast is negative. For the 12-month horizon, there is some more support, but the estimates are far from summing to one.

Taken together, as in Farmer, Nakamura, and Steinsson (2023), we find that sticky or noisy information models unlikely explain the rejections of rational expectations.

4.2 Learning

Table 6 presents results of panel regressions of the survey risk premium on the recent individual forecast error, the recent consensus forecast error (denoted $\overline{\mathrm{FE}}$), and the control variables in x_t:

$$\mathrm{F}_t^j r_{t,t+k} = a^j + b\,\mathrm{FE}^j_{t-k,t} + c\,\overline{\mathrm{FE}}_{t-k,t} + d' x_t + \varepsilon^j_{t+k}. \tag{7}$$

Recall that we consider the panel regression for a currency i and an individual forecaster j. To sharpen the interpretation of these regressions, the consensus forecast excludes forecaster j, which means that the consensus forecast error is different for each forecaster. In other words, the first term on the right-hand side of the regression is the recent forecast error of forecaster j, while the second term captures the recent consensus (excluding forecaster j) forecast error. The three information variables from before are used as control variables in x_t.

Table 6 presents the results. We introduce the variables sequentially and then add the three control variables. We note two specific results on the 3-month horizon. First, forecasters seem to adjust their expectations based on their own forecast error and not on the consensus forecast error. Second, the previous result concerning the depreciation rate seems to be driven by the forecast error. The results on the 12-month horizon are similar, but the multicollinearity between the individual and consensus forecast error clouds the results. Nevertheless, we conclude that part of the expectation formation is that the forecasters learn from their own mistakes.

4.3 Dispersion of beliefs

The previous analysis has shown a distinct dispersion of beliefs (or disagreement) and that this matters for, for instance, forecasting performance. This subsection develops this further by regressing a measure of cross-sectional dispersion (CS-Dispersion) on potential drivers.

The main challenge in measuring dispersion is that the size of the cross-section varies over time. However, a traditional unbiased estimate of the cross-sectional standard deviation delivers stable results; that is, results that are similar to other methods like a bias-adjusted mean absolute deviation or range (where the latter requires numerical integration).

It is natural that dispersion is related to uncertainty, so we choose to control for the implied volatility index VXY. We thus run panel regressions of the type:

$$\text{CS-Dispersion}_t = a + b\,x_t + c\,\text{VXY}_t + \varepsilon_t, \tag{8}$$

using currency fixed effects. Each regression uses a different x_t variable. (The coefficient estimates on VXY are not reported, but are around 0.15 and 0.30 for the 3- and 12-month horizons respectively, and statistically significant with t-stats around 5.) The within R^2 values are in the ranges 8–10% and 10–16% for the two horizons, respectively.

Table 7 presents the results. Two results stand out. First, when the real exchange rate is high (an "expensive" currency), the dispersion is larger. Second, when the survey consensus forecasts a depreciation, the dispersion is larger. Both results suggests that currencies expected to depreciate also are currencies with more dispersion (disagreement) across the forecasters.

5 Implications for risk premia

In this section we consider implications for risk premia. We compare rational and individual survey risk premia, and discuss their cyclicality. We also compare actual and survey-expected returns on common trading strategies in the foreign exchange rate market.

We begin by comparing rational and survey risk premia. As a measure of rational risk premia, we use the recursively fitted values of the realized excess return in the panel regression in Appendix A.1. As a measure of survey risk premia, we use the directly observed individual survey risk premia. Hence, both rational and survey risk premia are out-of-sample forecasts.

Figure 6 illustrates the cyclical movements in risk premia by currency for the 12-month horizon. The rational risk premium (green solid line) is countercyclical, increasing during the global financial crisis and NBER recessions (gray shaded areas), and negatively correlated with the annual growth in US industrial production and with annual inflation in US consumer prices. These negative correlations hold also for for the full sample estimated rational risk premium. In the figure, the survey risk premium is represented by the interquartile range of individual survey risk premia (dark blue areas). It shows a low correlation with the rational risk premium. On average, the survey risk premium is weakly correlated (sometimes negatively, sometimes positively) with the US inflation rate and the US industrial production growth.

Taken together, the results suggest that the rational risk premium is countercyclical (due to its relationship with the forward premium and the real exchange rate, and it increases in recessions) but that the survey risk premium seems acyclical.

We close by comparing actual and survey-expected returns on common trading strategies. More specifically, we consider carry, value, and momentum strategies in the cross-section

(CS) as well as in the time series (TS); see Chernov, Dahlquist, and Lochstoer (2023) for the construction of these strategies. These strategies can be seen as dynamic versions of the rational and survey risk premia.

Figure 7 illustrates the cumulated actual and survey-expected excess returns. The actual carry strategies had a stellar performance until the global financial crisis, underperformed during the crisis, and have had a mediocre performance thereafter. The actual value strategies have had a reasonable overall performance, with overperformance during the crisis and mediocre performance after the crisis. The actual momentum strategies have had a flat performance throughout the sample period. The annualized full sample Sharpe ratios are 0.45 and 0.80 for CS and TS carry, 0.34 and 0.50 for CS and TS value, and 0.05 and 0.14 for CS and TS momentum. Note that the strategies are based on only six currencies versus the dollar and do not benefit from the opportunities in other developed or emerging market currencies. Interestingly, relative the actual performances, the survey-expected performance lines ups with the carry strategies, is more positive on the value strategies, and is more negative on the momentum strategies. This is consistent with the view that the forecasters align themselves with the long-term reversals (via the real exchange rate) and short-term reversals (i.e., the opposite of momentum, via the recent depreciation rate).

6 Conclusion

We study a panel of exchange rate forecasters (large financial institutions and rating agencies) for the 1992–2019 period. We find that the expectations of the forecasters differ from both forward premia and rational (regression-based) predictions. While being influenced by the long-term reversal of the real exchange rate, the survey expectations are more in line with UIP and relate negatively to recent depreciations. The expectations also relate negatively to the forecasters' own recent forecast errors, and this learning tends to result in a short-term reversal of the forecasts. These patterns imply that forecast errors exhibit positive serial correlation and that they relate negatively to forecasters' recent revisions of their own forecasts (i.e., the forecasters overreact). However, we find little support for the sticky and noisy information models that motivate the basic overreaction specification. The out-of-sample forecasting performance is inferior to recursive regressions and other simple benchmarks. However, we notice that some forecasters perform better than others. In particular, better forecasting performance correlates with expectations that put less emphasis on UIP, more emphasis on long-term reversal, and involve less overreaction.

Overall, our results highlight the information contained in individual (rather than consensus) exchange rate forecasts.

A Risk premia and simulation experiments

In this appendix we report results on the predictability and risk premia. We also report results from a simulation study.

A.1 Predictability and risk premia

To study rational and survey risk premia, we estimate the following panel regressions of the realized currency excess return and survey risk premium:

$$r_{t,t+k} = a + b'x_t + \varepsilon_{t+k}, \quad (A1)$$

$$\mathrm{F}_t^j r_{t,t+k} = c^j + d'x_t + \epsilon_t^j. \quad (A2)$$

Note that the difference between regressions (A1) and (A2) form a regression of the expectations error, $r_{t,t+k} - \mathrm{F}_t^j r_{t,t+k}$, on x_t.

We let x_t capture the forward premium, the real exchange rate, and the recent USD depreciation rate. These variables are common predictors in the literature and form the basis of so-called carry, value, and momentum/reversal strategies. Notably, the three variables are only weakly correlated. For example, on the 3-month horizon, the forward premium has an average correlation of –0.17 with the real exchange rate (where the average is across the six currencies) and –0.11 with the recent depreciation rate. The real exchange rate and the recent depreciation rate have an average correlation of 0.18.

Table A1 presents the results. Three results for specifications I and II stand out. First, the realized excess return is negatively related to the forward premium, while the survey risk premium is not (the forecasters seem to have expectations more in line with UIP). Second, the realized excess return and the survey risk premium both relate to the real exchange rate, which captures long-term reversals in exchange rates. Third, there is little persistence in the realized excess return (as measured by the coefficient of the recent depreciation rate), while the survey risk premium indicates strong expectations of short-term reversals. The differences in rational and survey risk premia are statistically significant.

We use the fitted values as a rational benchmark of currency risk premia. While the panel regression above uses data for the entire sample period, we also use the panel regression in a recursively fashion to have true out-of-sample predictions.

It is well known that predictability regressions raise econometric concerns in small samples. First, high persistence in regressors and correlations between regressor innovations and excess return innovations causes biases in point estimates (Stambaugh, 1999), typically towards finding predictability. Second, standard errors from commonly used overlapping data estimators are often too small (see Boudoukh, Israel, and Richardson, 2022, for a recent analysis). Taken together, there is a risk of exaggerating the magnitude of predictability. We therefore investigate by simulations the estimator properties in our panel setting.

The simulations use the following data-generating processes for each currency. The predictors in x_t follow a vector autoregression. The realized excess return, the survey risk premium, and the expectations error are linear in the predictors. We use the estimated coefficients in a system including the consensus forecasts (not the individual forecasts) for 3- as well as 12-month horizons. Residuals of the system are block bootstrapped to preserve the correlation structure across (a) variables, (b) currencies, and (c) time. We use 10,000 simulations. Appendix A.2 provides more details.

We indeed find biases in point estimates and standard errors in our simulations, but they do not overturn any of the reported results. For the excess return (the standard predictability regression), the coefficients of the forward premium are less than 15% closer to zero and the coefficients of the real exchange rate are less than 30% closer to zero. (The larger bias for the real exchange rate is expected given that its persistence is greater than that of the forward premium.) The standard errors for the coefficients are 1–5% larger on the 3-month horizon and around 15% larger on the 12-month horizon. The well-known small sample bias in the predictability regression does not carry over to the regression of the survey risk premium. The explanation is that while the regressors are serially correlated, the regressor innovations and the survey risk premium innovations are only weakly correlated. Hence, we confirm a negative relationship with the real exchange rate and a negative relationship with the recent depreciation rate. However, the standard errors are up to 40% larger in the simulations. Despite that, the relationships remain statistically significant. Taken together, we conclude that the difference between the rational risk premium and the survey risk premium is that the survey risk premium is not related to the forward premium but instead relates negatively to the recent depreciation rate.

In the main text, we further explore the negative relationship with the recent trend in exchange rates (the negative coefficient of the recent depreciation rate), which leads us to analyses of overeaction and learning.

A.2 Simulation experiments

This appendix contains more details of the simulation experiments to assess the properties of the panel estimators used in Sections 3 and 4.

The data-generating process (DGP) is:

$$x_{t+k} = Ax_t + \eta_{t+k}, \tag{A3}$$

$$r_{t,t+k} = b'x_t + \varepsilon_{t+k}, \tag{A4}$$

$$\mathrm{F}_t r_{t,t+k} = c'x_t + \epsilon_t, \tag{A5}$$

where, for notational simplicity, we have suppressed currency fixed effects. In (A3) the vector x_t contains three variables (the forward premium, the real exchange rate, and the recent USD depreciation rate), and we model it as a VAR system in which x_{t+k} is driven by x_t; in (A4) the realized currency excess return, $r_{t,t+k}$, is regressed on the x_t vector; and in (A5) the current survey risk premium, $\mathrm{F}_t r_{t,t+k}$, is regressed on the current x_t vector. The DGP is similar to those of Amihud, Hurvich, and Wang (2009) and Boudoukh, Israel, and Richardson (2022), except that it involves multi-period predictions with a vector of predictors.

The parameter values in the simulations are from estimations based on our monthly data during the 1992–2019 period for horizons of three and twelve months ($k = 3$ and $k = 12$): the A matrices are estimated independently for each currency, while the b and d vectors are from panel regressions (with currency fixed effects) and are thus the same across currencies.

In the simulations, we bootstrap the errors $(\eta_t, \varepsilon_t, \epsilon_t)$ from the residuals, drawing blocks of $2k$ consecutive months. To preserve the correlation across the variables and currencies, the blocks are the same for all variables and currencies. For example, if U is a $T \times 5 \times 6$ array with T periods, five variables (related to η_t, ε_t, and ϵ_t), and six currencies, then we effectively draw $U[s, :, :]$, where s selects a block of $2k$ months (e.g., s may be equal to $[11, 12, 13, 14, 15, 16]$).

We run 10,000 simulations of 312 months (using an extra 50 "burn-in" months) and re-estimate the panel regressions on each generated sample ($\hat{b}_l$ and $\hat{c}_l$ for an l that runs from one to 10,000). We assess the bias by comparing the average simulation results with estimates based on a simulated long sample of 100,000 months. We assess the standard errors by comparing the output from the panel regressions on data with the standard deviations across the $10,000$ simulations.

Table A1: Predictability and risk premia

	Realized excess return I	Survey risk premium II
Panel A: 3-month horizon		
Forward premium	-2.298^{***}	0.093
	(0.772)	(0.298)
Real exchange rate	-0.060^{***}	-0.030^{***}
	(0.020)	(0.007)
Recent depreciation	0.060	-0.214^{***}
	(0.050)	(0.026)
Fixed effects	Currency	Currency & Forecaster
R^2	0.036	0.120
N	1872	10664
Panel B: 12-month horizon		
Forward premium	-2.563^{***}	0.140
	(0.656)	(0.209)
Real exchange rate	-0.236^{***}	-0.108^{***}
	(0.060)	(0.024)
Recent depreciation	0.030	-0.209^{***}
	(0.120)	(0.037)
Fixed effects	Currency	Currency & Forecaster
R^2	0.151	0.194
N	1872	10270

The table shows results of regressions of the realized excess return and survey risk premium on the forward premium, the real exchange rate, and the recent depreciation rate for a monthly unbalanced panel with 21 individual forecasters' expectations of the dollar versus six currencies during the 1992–2019 period:

$$r_{t,t+k} = a + b' x_t + \varepsilon_{t+k}, \qquad \text{(I)}$$
$$\mathrm{F}_t^j r_{t,t+k} = c^j + d' x_t + \epsilon_t^j, \qquad \text{(II)}$$

where x_t contains $f_{k,t} - s_t$, q_t, and $s_t - s_{t-k}$. Specification I uses currency fixed effects; specifications II and III use currency and forecaster fixed effects. Panels A and B report results for horizons of three and twelve months ($k = 3$ and $k = 12$), respectively. R^2 is a within R-squared value. N is the number of currency-forecaster-month observations. Standard errors (within parentheses) are calculated using the spatial estimator of Driscoll and Kraay (1998), which allows for both cross-sectional and serial correlations up to k lags in the errors, as in Hansen and Hodrick (1980), as well as for heteroskedasticity in the errors. *, **, and *** denote 10%, 5%, and 1% significance levels, respectively, for the null hypothesis of a coefficient equal to zero.

References

Adam, Klaus, Dmitry Matveev, and Stefan Nagel, 2021, Do survey expectations of stock returns reflect risk adjustments?, *Journal of Monetary Economics* 117, 723–740.

Amihud, Yakov, Clifford M. Hurvich, and Yi Wang, 2009, Multiple-predictor regressions: Hypothesis testing, *Review of Financial Studies* 22, 413–434.

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

Bacchetta, Philippe, Elmar Mertens, and Eric van Wincoop, 2009, Predictability in financial markets: What do survey expectations tell us?, *Journal of International Money and Finance* 28, 406–426.

Bansal, Ravi, and Magnus Dahlquist, 2000, The forward premium puzzle: Different tales from developed and emerging economies, *Journal of International Economics* 51, 115–144.

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer, 2020, Overreaction in macroeconomic expectations, *American Economic Review* 110, 2748–2782.

Bouchaud, Jean-Philippe, Philipp Krüger, Augustin Landier, and David Thesmar, 2019, Sticky expectations and the profitability anomaly, *Journal of Finance* 74, 639–674.

Boudoukh, Jacob, Ronen Israel, and Matthew Richardson, 2022, Biases in long-horizon predictive regressions, *Journal of Financial Economics* 145, 937–969.

Buraschi, Andrea, Ilaria Piatti, and Paul Whelan, 2021, Subjective bond returns and belief aggregation, *Review of Financial Studies* 35, 3710–3741.

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

Candian, Giacomo, and Pierre De Leo, 2023, Imperfect exchange rate expectations, forthcoming in *Review of Economics and Statistics*.

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

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

Chinn, Menzie D., and Jeffrey A. Frankel, 1994, Patterns in exchange rate forecasts for twenty-five currencies, *Journal of Money, Credit and Banking* 26, 759–770.

Cieslak, Anna, 2018, Short-rate expectations and unexpected returns in Treasury bonds, *Review of Financial Studies* 31, 3265–3306.

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

Dahlquist, Magnus, and Markus Ibert, 2023, Equity return expectations and portfolios: Evidence from large asset managers, Working Paper.

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

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

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

Dominguez, Kathryn M., 1986, Are foreign exchange forecasts rational?, *Economic Letters* 21, 277–281.

Driscoll, John, and Aart Kraay, 1998, Consistent covariance matrix estimation with spatially dependent panel data, *Review of Economics and Statistics* 80, 549–560.

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

Farmer, Leland, Emi Nakamura, and Jon Steinsson, 2023, Learning about the long run, forthcoming in *Journal of Political Economy.*

Frankel, Jeffrey A., and Kenneth A. Froot, 1987, Using survey data to test standard propositions regarding exchange rate expectations, *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.

Giacoletti, Marco, Kristoffer T. Laursen, and Kenneth J. Singleton, 2021, Learning from disagreement in the U.S. Treasury bond market, *Journal of Finance* 76, 395–441.

Gourinchas, Pierre-Olivier, and Aaron Tornell, 2004, Exchange rate puzzles and distorted beliefs, *Journal of International Economics* 64, 303–333.

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

Hansen, Lars Peter, 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 K., and Rui C. Mano, 2019, Forward and spot exchange rates in a multi-currency world, *Quarterly Journal of Economics* 134, 397–450.

Ince, Onur, and Tanya Molodtsova, 2017, Rationality and forecasting accuracy of exchange rate expectations: Evidence from survey-based forecasts, *Journal of International Financial Markets, Institutions & Money* 47, 131–151.

Ito, Takatoshi, 1990, Foreign exchange rate expectations: Micro survey data, *American Economic Review* 80, 434–449.

Kalemli-Ozcan, Sebnem, and Liliana Varela, 2023, Five facts about the UIP premium, Working Paper.

Koijen, Ralph S.J., Tobias J. Moskowitz, Lasse Heje Pedersen, and Evert B. Vrugt, 2018, Carry, *Journal of Financial Economics* 127, 197–225.

Kremens, Lukas, Ian Martin, and Liliana Varela, 2023, Long-horizon exchange rate expectations, Working Paper.

Kučinskas, Simas, and Florian S. Peters, 2022, Measuring under- and overreaction in expectation formation, forthcoming in *Review of Economics and Statistics*.

Li, Zigang, Stijn Van Nieuwerburgh, and Wang Renxuan, 2023, Understanding rationality and disagreement in house price expectations, Working Paper.

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

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

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

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

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2017, Currency value, *Review of Financial Studies* 30, 416–441.

Mincer, Jacob A., and Victor Zarnowitz, 1969, The evaluation of economic forecasts, in Jacob A. Mincer, eds.: *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance* (NBER, Cambridge).

Molavi, Pooya, Alireza Tahbaz-Salehi, and Andrea Vedolin, 2023, Model complexity, expectations, and asset prices, forthcoming in *Review of Economic Studies*.

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

Piazzesi, Monika, Juliana Salomao, and Martin Schneider, 2015, Trend and cycle in bond premia, Working Paper.

Singleton, Kenneth J., 2021, Presidential address: How much "rationality" is there in bond-market risk premiums?, *Journal of Finance* 76, 1611–1654.

Stambaugh, Robert F., 1999, Predictive regressions, *Journal of Financial Economics* 54, 375–421.

Stavrakeva, Vania, and Jenny Tang, 2020, Deviations from FIRE and exchange rates: A GE theory of supply and demand, Working Paper.

Stavrakeva, Vania, and Jenny Tang, 2023, A fundamental connection: Exchange rates and macroeconomic expectations, Working Paper.

Valente, Joao Paulo, Kaushik Vasudevan, and Tianhao Wu, 2022, The role of beliefs in asset prices: Evidence from exchange rates, Working Paper.

Table 1: Summary statistics of realized and survey-expected variables

	Realized		Survey		
	Depreciation	Excess return	Expected depreciation	Risk premium	Expectations error
Panel A: 3-month horizon					
Mean	0.052	−0.042	−0.335	−0.430	0.388
Standard deviation	4.949	4.993	2.843	2.781	5.851
Skewness	−0.287	−0.258	−0.055	−0.053	−0.289
Excess kurtosis	2.227	1.988	0.940	1.068	2.249
Minimum	−20.442	−20.533	−10.449	−10.480	−23.731
Maximum	16.633	16.467	8.506	8.634	20.224
AR(3)	0.041	0.060	0.255	0.229	0.256
Panel B: 12-month horizon					
Mean	0.125	−0.323	0.051	−0.397	0.073
Standard deviation	9.857	10.178	4.848	4.448	10.693
Skewness	−0.234	−0.177	0.035	−0.013	−0.050
Excess kurtosis	0.324	0.100	0.962	1.108	0.024
Minimum	−28.522	−28.492	−15.471	−15.161	−28.271
Maximum	25.105	24.850	13.393	11.712	28.511
AR(12)	−0.019	0.045	0.454	0.435	0.083

The table shows summary statistics (in %) for averages of the following variables: currency depreciation rate, $s_{t+k} - s_t$; currency excess return, $r_{t,t+k}$; survey-expected depreciation rate, $\mathrm{F}_t^j s_{t+k} - s_t$; survey risk premium, $\mathrm{F}_t^j r_{t,t+k}$; and survey expectations error, $s_{t+k} - \mathrm{F}_t^j s_{t+k}$. The averages of the variables are taken over currencies, and over forecasters for survey expectations, for a monthly unbalanced panel with 21 forecasters' expectations of the USD versus six currencies during the 1992–2019 period. Panels A and B report results for horizons of three and twelve months ($k = 3$ and $k = 12$), respectively. The numbers of time-series observations in Panels A and B are 310 and 300, respectively.

Table 2: Unbiasedness in risk-neutral and survey expectations

	Realized depreciation	
	Risk-neutral expectations I	Survey expectations II
Panel A: 3-month horizon		
$\hat{b}$	-0.998^{**}	-0.020^{***}
	(0.824)	(0.039)
Fixed effects	Currency	Currency & Forecaster
Within R^2	0.006	0.000
N	1872	10664
Panel B: 12-month horizon		
$\hat{b}$	-1.242^{***}	0.044^{***}
	(0.784)	(0.088)
Fixed effects	Currency	Currency & Forecaster
Within R^2	0.034	0.001
N	1872	10270

The table shows results of regressions of the future depreciation rate on either the current forward premium or the current survey-expected depreciation rate:

$$s_{t+k} - s_t = a + b\,(f_{k,t} - s_t) + \varepsilon_{t+k}, \quad \text{(I)}$$
$$s_{t+k} - s_t = a^j + b\,(\mathrm{F}_t^j s_{t+k} - s_t) + \epsilon_{t+k}^j, \quad \text{(II)}$$

for either a monthly panel with exchange rates of the USD versus six currencies during the 1992–2019 period or a monthly unbalanced panel with 21 forecasters' expectations of the USD versus the six currencies during the 1992–2019 period. Specification I uses the forward exchange rate (the risk-neutral expectations) and allows for currency fixed effects; specification II uses individual forecasters' expectations and allows for currency and forecaster fixed effects. Panels A and B report results for horizons of three and twelve months ($k = 3$ and $k = 12$), respectively. N is the number of currency-forecaster-month observations. Standard errors (within parentheses) are calculated using the spatial estimator of Driscoll and Kraay (1998), which allows for both cross-sectional and serial correlations up to k lags in the errors, as in Hansen and Hodrick (1980), as well as for heteroskedasticity in the errors. *, **, and *** denote 10%, 5%, and 1% significance levels, respectively, for the null hypothesis of a coefficient equal to one.

Table 3: Forecast error and forecast revision

	Forecast error				
	I	II	III	IV	V
Panel A: 3-month horizon					
Lagged forecast error	0.250***		0.257***		
	(0.035)		(0.045)		
Forecast revision				−0.062**	−0.089***
				(0.032)	(0.034)
Forward premium		−2.319***	−1.902***		−2.752***
		(0.790)	(0.827)		(0.921)
Real exchange rate		−0.025	−0.037*		−0.003
		(0.021)	(0.021)		(0.022)
Recent depreciation		0.249***	−0.008		
		(0.056)	(0.067)		
Within R^2	0.063	0.059	0.081	0.004	0.029
N	9909	10664	9909	9926	9926
Panel B: 12-month horizon					
Lagged forecast error	0.112		0.057		
	(0.076)		(0.058)		
Forecast revision				−0.041	−0.121*
				(0.093)	(0.064)
Forward premium		−2.607***	−2.466***		−3.006***
		(0.550)	(0.554)		(0.611)
Real exchange rate		−0.114**	−0.125**		−0.075
		(0.057)	(0.058)		(0.062)
Recent depreciation		0.285***	0.219**		
		(0.101)	(0.102)		
Within R^2	0.040	0.140	0.136	0.001	0.143
N	8534	10270	8534	9053	9053

The table shows results of regressions of the forecast error on the lagged forecast error or the forecast revision, and on the forward premium, the real exchange rate, and the recent depreciation rate for a monthly unbalanced panel with 21 forecasters' expectations of the USD versus six currencies during the 1992–2019 period:

$$\mathrm{FE}^j_{t,t+k} = a^j + b\,\mathrm{FE}^j_{t-k,t} + c\,\mathrm{FR}^j_{t-k,t} + d'x_t + \varepsilon^j_{t+k},$$

where $\mathrm{FE}^j_{t,t+k} = r_{t,t+k} - \mathrm{F}^j_t r_{t,t+k}$ is the forecast error, $\mathrm{FR}^j_{t-k,t} = \mathrm{F}^j_t r_{t,t+k} - \mathrm{F}^j_{t-k} r_{t,t+k}$ is the forecast revision, and x_t contains $f_{k,t} - s_t$, q_t, and $s_t - s_{t-k}$. All specifications use currency and forecaster fixed effects. Panels A and B report results for horizons of three and twelve months ($k = 3$ and $k = 12$), respectively. The 12-month forecast revision is approximated by the 6-month forecast revision. N is the number of currency-forecaster-month observations. Standard errors (within parentheses) are calculated using the spatial estimator of Driscoll and Kraay (1998), which allows for both cross-sectional and serial correlations up to k lags in the errors, as in Hansen and Hodrick (1980), as well as for heteroskedasticity in the errors. *, **, and *** denote 10%, 5%, and 1% significance levels, respectively, for the null hypothesis of a coefficient equal to zero.

Table 4: Out-of-sample forecast comparisons

	3-month	12-month
Panel A. Panel regression and consensus survey forecasts		
Full regression model	1.415***	1.154
Panel B. Panel regression and individual survey forecasts		
Forward premium	1.198***	1.167
Real exchange rate	1.245***	1.165**
Recent depreciation	0.887*	1.043
Lagged forecast error	1.032	1.049
Revision of forecast	0.804***	0.827*

This table shows results of out-of-sample forecast comparisons of mean squared forecasting errors (MSEs).
Panel A shows the ratio:

$$\frac{\text{MSE(consensus)}}{\text{MSE(panel regression)}},$$

where MSE(consensus) and MSE(panel regression) refer to the MSE of survey forecasts and panel regression forecasts, respectively. The panel regression is with currency fixed effects during the 1992–2019 period and estimated using an expanding window so that all forecasts are out of sample.
Panel B shows the ratio:

$$\frac{\text{MSE(high coefficient)}}{\text{MSE(low coefficient)}},$$

where MSE(high coefficient) and MSE(low coefficient) refer to the MSEs of individual survey forecasts conditioning on high and low panel regression coefficients of a variable (the forward premium, real exchange rate, recent depreciation rate, lagged forecast error, and revision of forecast), respectively. These ratios are constructed using the following three steps. First, run a panel regression for each individual forecaster with currency fixed effects. Second, form the forecast average of those forecasters with high coefficients of a variable and another average of those with low coefficients. Third, compare the predictive performance of forecasts for the high and low coefficients. The number of time-series observations is 310 for the 3-month horizon and 300 for the 12-month horizon. We use a Diebold and Mariano (1995) test, using a robust variance-covariance estimator that allows for serial correlations up to k lags, as in Hansen and Hodrick (1980), as well as for heteroskedasticity, to test the null hypothesis of equal MSEs between survey and panel regression forecasts. *, **, and *** denote 10%, 5%, and 1% significance levels, respectively.

Table 5: Sticky and noisy information models

	Survey risk premium	
	Sticky information I	Noisy information II
Panel A: 3-month horizon		
Rational risk premium	−0.036	−0.099
	(0.115)	(0.102)
Previous survey risk premium		0.420***
		(0.036)
Within R^2	0.000	0.071
N	9504	9871
Wald	[0.000]	[0.000]
Panel B: 12-month horizon		
Rational risk premium	0.214***	0.164**
	(0.055)	(0.067)
Previous survey risk premium		0.471**
		(0.044)
Within R^2	0.037	0.110
N	9648	9172
Wald	[0.000]	[0.000]

The table shows results of regressions of the survey risk premium on the rational risk premium and the previous own survey risk premium for a monthly unbalanced panel with 21 forecasters' expectations of the dollar versus six currencies during the 1992–2019 period:

$$\mathrm{F}_t^j r_{t,t+k} = a^j + b\,\hat{\mathrm{E}}_t r_{t,t+k} + c\,\mathrm{F}_{t-k}^j r_{t,t+k} + \varepsilon_t^j,$$

where $\hat{\mathrm{E}}_t r_{t,t+k}$ is the rational risk premium, approximated by a recursively estimated forecasting model, and $\mathrm{F}_{t-k}^j r_{t,t+k}$ is the previous survey risk premium. Panels A and B report results for horizons of three and twelve months ($k = 3$ and $k = 12$), respectively. Specification I uses only those observations for which the revision is non-zero. The Wald statistic reports p-values from a test of the null hypothesis of a slope coefficient equal to one ($b = 1$). Specification II uses all observations. The previous 12-month survey risk premium is approximated by the 6-month survey risk premium. The Wald statistic reports p-values from a joint test of the null hypothesis that the slope coefficients sum to one ($b + c = 1$). N is the number of currency-forecaster-month observations. Standard errors (within parentheses) are calculated using the spatial estimator of Driscoll and Kraay (1998), which allows for both cross-sectional and serial correlations up to k lags in the errors, as in Hansen and Hodrick (1980), as well as for heteroskedasticity in the errors. *, **, and *** denote 10%, 5%, and 1% significance levels, respectively, for the null hypothesis of a coefficient equal to zero.

Table 6: Learning

	Survey risk premium			
	I	II	III	IV
Panel A: 3-month horizon				
Recent forecast error (individual)	−0.230***		−0.253***	−0.245***
	(0.022)		(0.021)	(0.024)
Recent forecast error (consensus)		−0.211***	0.029	0.064**
		(0.024)	(0.026)	(0.029)
Forward premium				−0.140
				(0.292)
Real exchange rate				−0.023***
				(0.007)
Recent depreciation				−0.038
				(0.039)
Within R^2	0.146	0.102	0.147	0.154
N	9909	10581	9909	9909
Panel B: 12-month horizon				
Recent forecast error (individual)	−0.162***		−0.277***	−0.249***
	(0.038)		(0.034)	(0.032)
Recent forecast error (consensus)		−0.114***	0.137***	0.231***
		(0.042)	(0.048)	(0.047)
Forward premium				0.173
				(0.216)
Real exchange rate				−0.116***
				(0.020)
Recent depreciation				−0.184***
				(0.043)
Within R^2	0.092	0.041	0.104	0.176
N	8534	9918	8534	8534

The table shows results of regressions of the survey risk premium on the recent individual forecast error, the recent consensus forecast error, the forward premium, the real exchange rate, and the recent depreciation rate for a monthly unbalanced panel with 21 forecasters' expectations of the dollar versus six currencies during the 1992–2019 period:

$$\mathrm{F}_t^j r_{t,t+k} = a^j + b\,\mathrm{FE}_{t-k,t}^j + c\,\overline{\mathrm{FE}}_{t-k,t} + d' x_t + \varepsilon_{t+k}^j,$$

where $\mathrm{FE}_{t-k,t}^j = r_{t-k,t} - \mathrm{F}_{t-k}^j r_{t-k,t}$ is the recent individual forecast error, $\overline{\mathrm{FE}}_{t-k,t}$ is the corresponding recent consensus forecast error, and x_t contains $f_{k,t} - s_t$, q_t, and $s_t - s_{t-k}$. The consensus forecast for a specific forecaster is an average of all forecasters but the specific forecaster. All specifications use currency and forecaster fixed effects. Panels A and B report results for horizons of three and twelve months ($k = 3$ and $k = 12$), respectively. N is the number of currency-forecaster-month observations. Standard errors (within parentheses) are calculated using the spatial estimator of Driscoll and Kraay (1998), which allows for both cross-sectional and serial correlations up to k lags in the errors, as in Hansen and Hodrick (1980), as well as for heteroskedasticity in the errors. *, **, and *** denote 10%, 5%, and 1% significance levels, respectively, for the null hypothesis of a coefficient equal to zero.

Table 7: Dispersion

	Cross-sectional dispersion	
	3-month horizon I	12-month horizon II
Forward premium	-0.265^{*}	0.203^{*}
	(0.148)	(0.104)
Real exchange rate	0.008^{**}	0.037^{***}
	(0.003)	(0.007)
Recent depreciation	-0.006	-0.019
	(0.009)	(0.022)
US industrial production	-0.016	-0.002
	(0.017)	(0.054)
US inflation	-0.147^{***}	-0.068
	(0.049)	(0.164)
Consensus forecast	-0.050^{***}	-0.051^{*}
	(0.016)	(0.028)

The table shows results of regressions the cross-sectional dispersion of survey forecasts (for a given exchange rate) on the VXY and one more variable, x_t:

$$\text{CS-Dispersion}_t = a + b\,x_t + c\,\text{VXY}_t + \varepsilon_t,$$

for a monthly panel with six currencies during the 1992–2019 period, allowing for currency fixed effects. The left and right panels report results for the estimates of b for horizons of three and twelve months ($k = 3$ and $k = 12$), respectively. The number number of currency-month observations is 1872 for all regressions. Standard errors (within parentheses) are calculated using the spatial estimator of Driscoll and Kraay (1998), which allows for both cross-sectional and serial correlations up to k lags in the errors, as in Hansen and Hodrick (1980), as well as for heteroskedasticity in the errors. *, **, and *** denote 10%, 5%, and 1% significance levels, respectively, for the null hypothesis of a coefficient equal to zero.

Figure 1: Number of forecasters

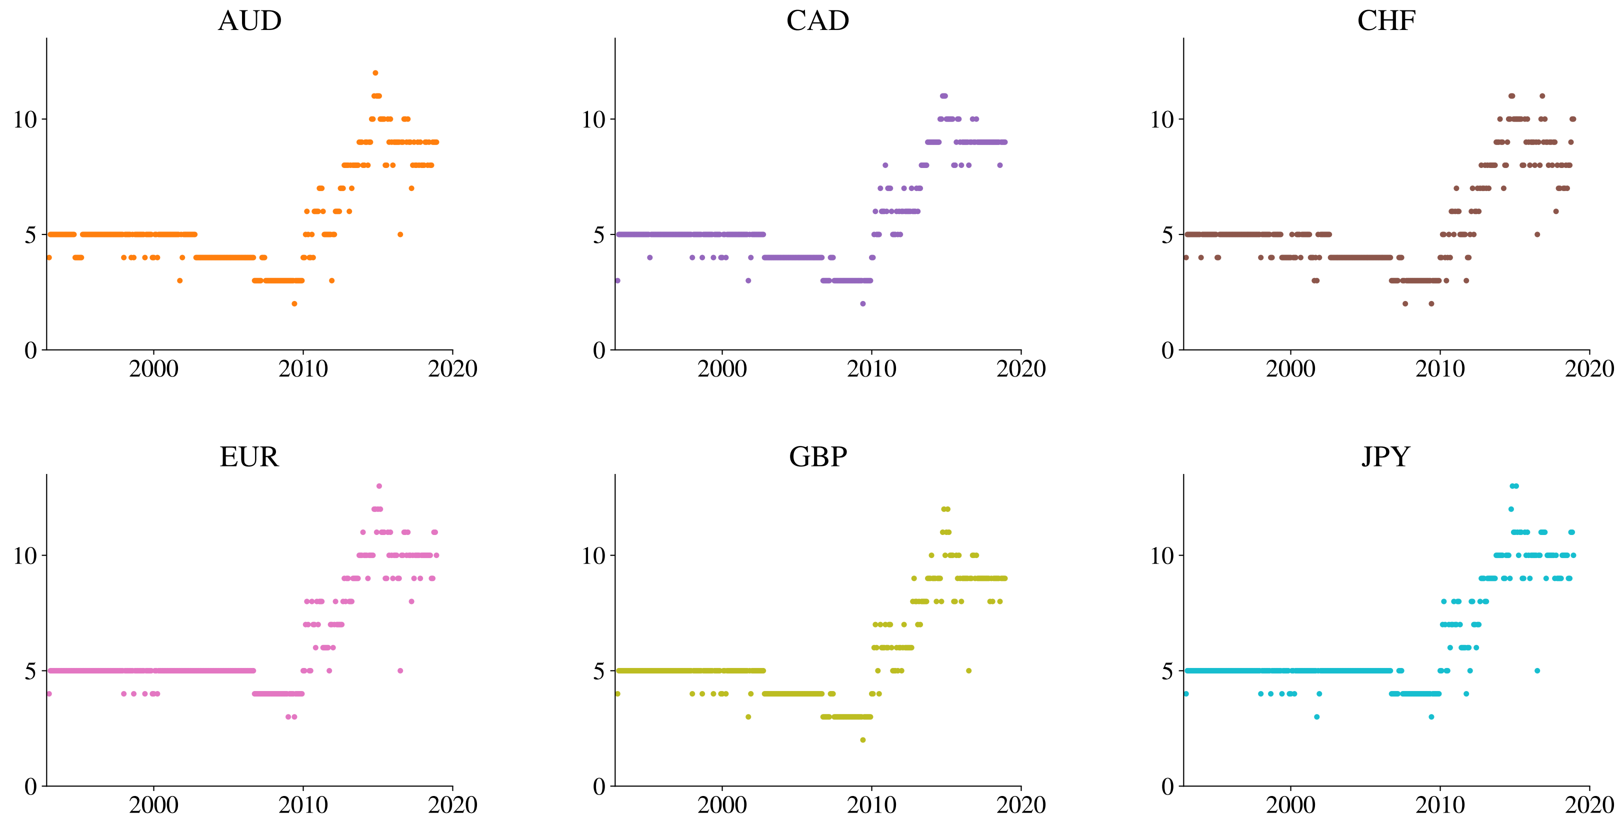

The figure shows the number of forecasters of future USD spot exchange rates (versus AUD, CAD, CHF, EUR, GBP, and JPY) over time.

Figure 2: Forward exchange rates and survey forecasts

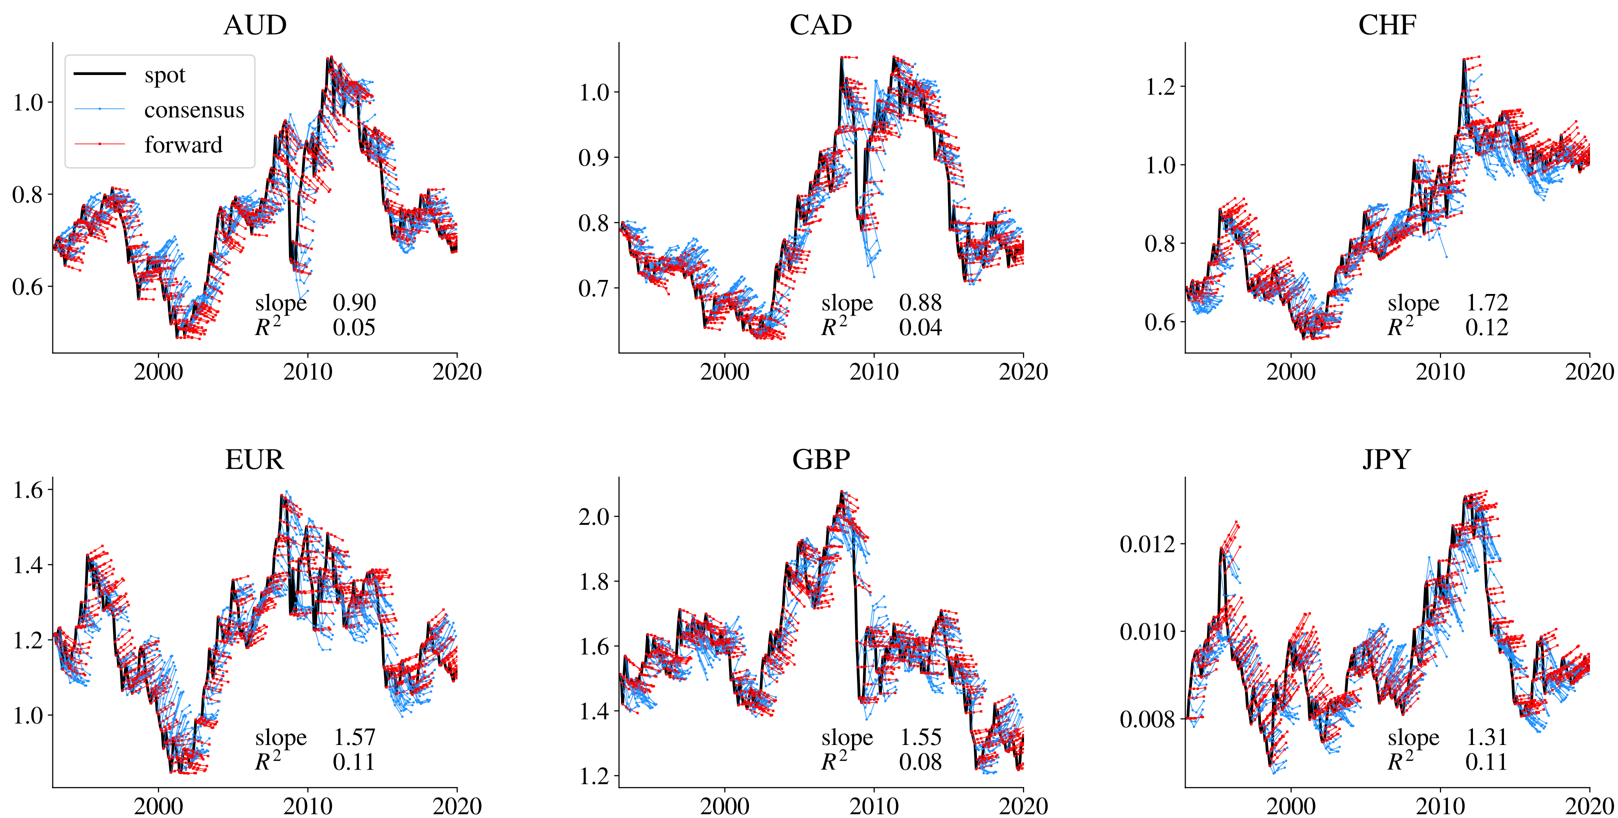

The figure shows USD spot exchange rates (versus the AUD, CAD, CHF, EUR, GBP, and JPY) over time (black solid line). The figure also shows 3-month, 6-month, and 12-month forward exchange rates (red solid line with squares) and 3-month, 6-month, and 12-month consensus expectations (blue solid line with circles) at each date. Each subfigure reports estimated slope coefficients and R^2s from panel regressions of the survey-expected depreciation rate on the forward premium, allowing for forecaster fixed effects.

Figure 3: Forward premia and survey-expected depreciation rates

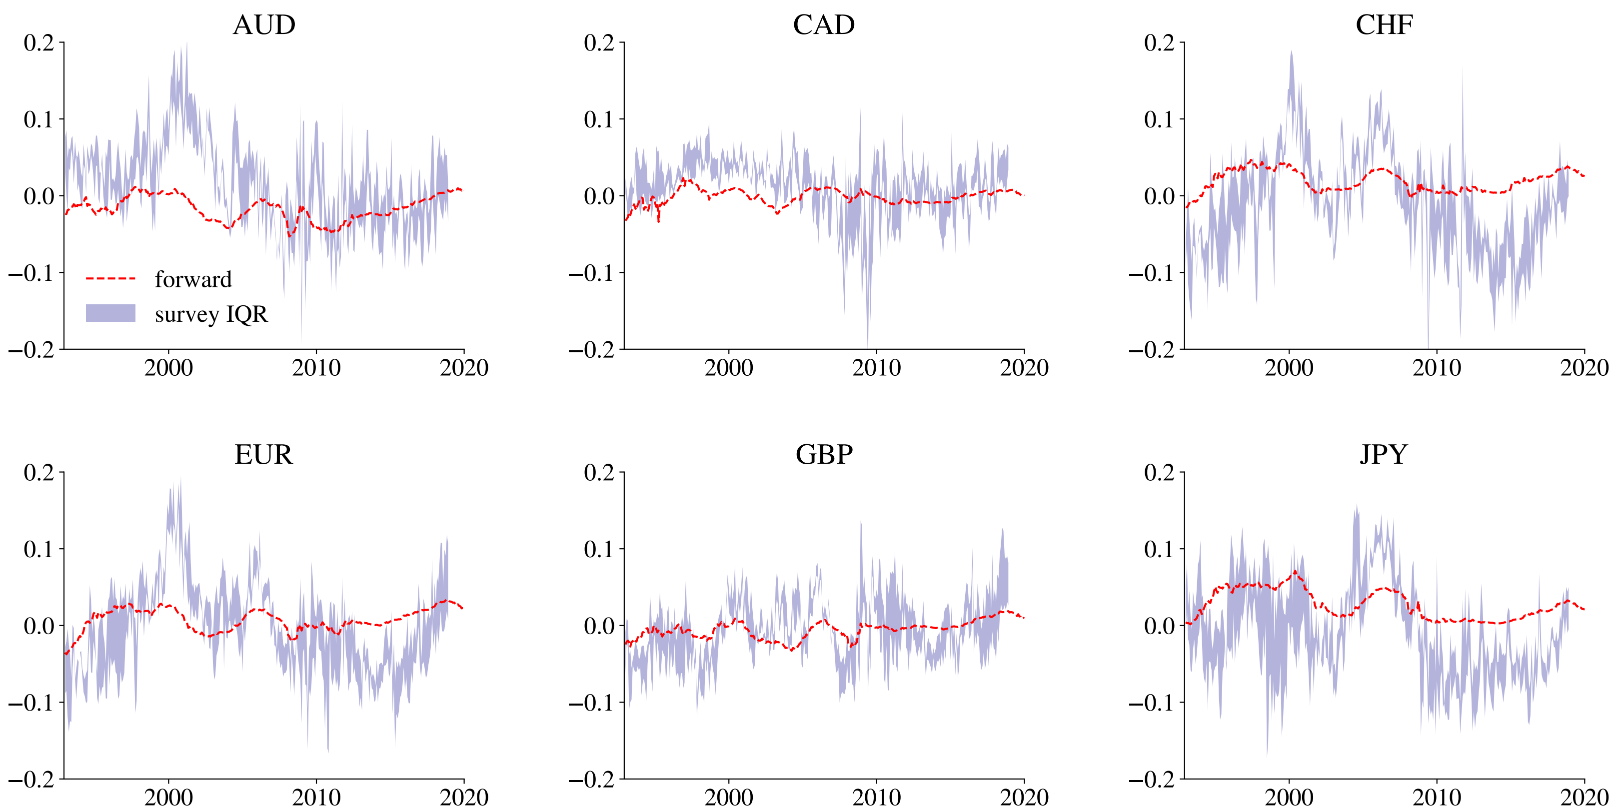

The figure shows US 12-month forward premia (versus AUD, CAD, CHF, EUR, GBP, and JPY) over time (dashed red line). The figure also shows a shaded band indicating the interquartile range of individual 12-month survey-expected depreciation rate (dark blue areas).

Figure 4: Distribution of individual coefficients

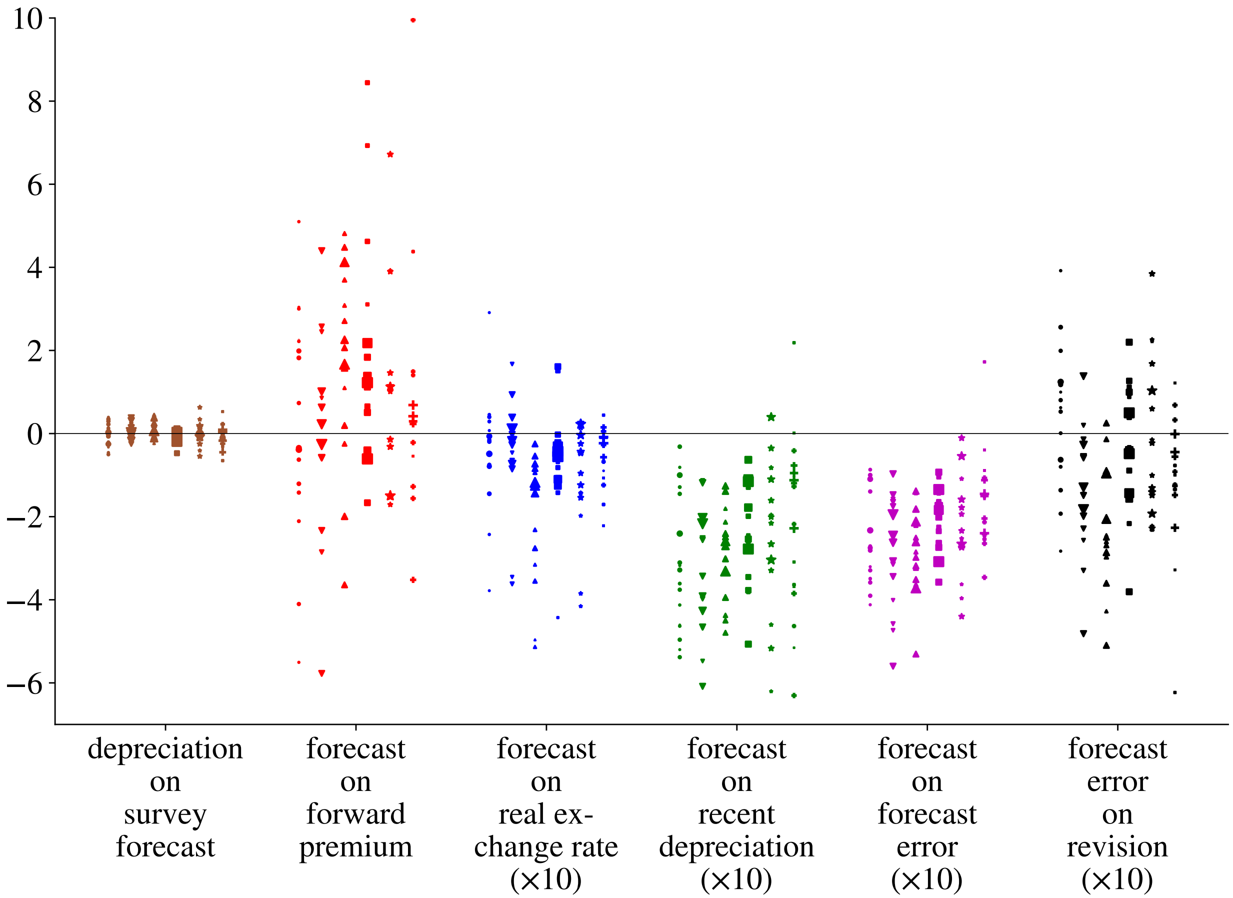

The figure shows regression coefficients from simple regressions with a 3-month horizon. The first scatter of coefficients (in brown) is from a regression of the realized depreciation rate on the survey-expected deprecation rate, similar to regression (2); the next four scatters are from regressions of the survey risk premium on the forward premium (in red), real exchange rate (in blue), recent depreciation rate (green), and recent forecast error (in magenta); the last scatter (in black) is from a regression of the forecast error on the forecast revision, similar to regression (5). Each dot in the scatters shows one forecaster-currency specific regression coefficient estimate, with the size of the dot indicating the relative number of observations. For each regression, a vertical segment illustrates the cross-sectional dispersion of estimation results for a given currenc (with the order AUD, CAD, CHF, EUR, GBP, and JPY). The horizontal axis indicates the regressions and scaling.

Figure 5: Cumulated MSE for survey versus panel regression forecasts

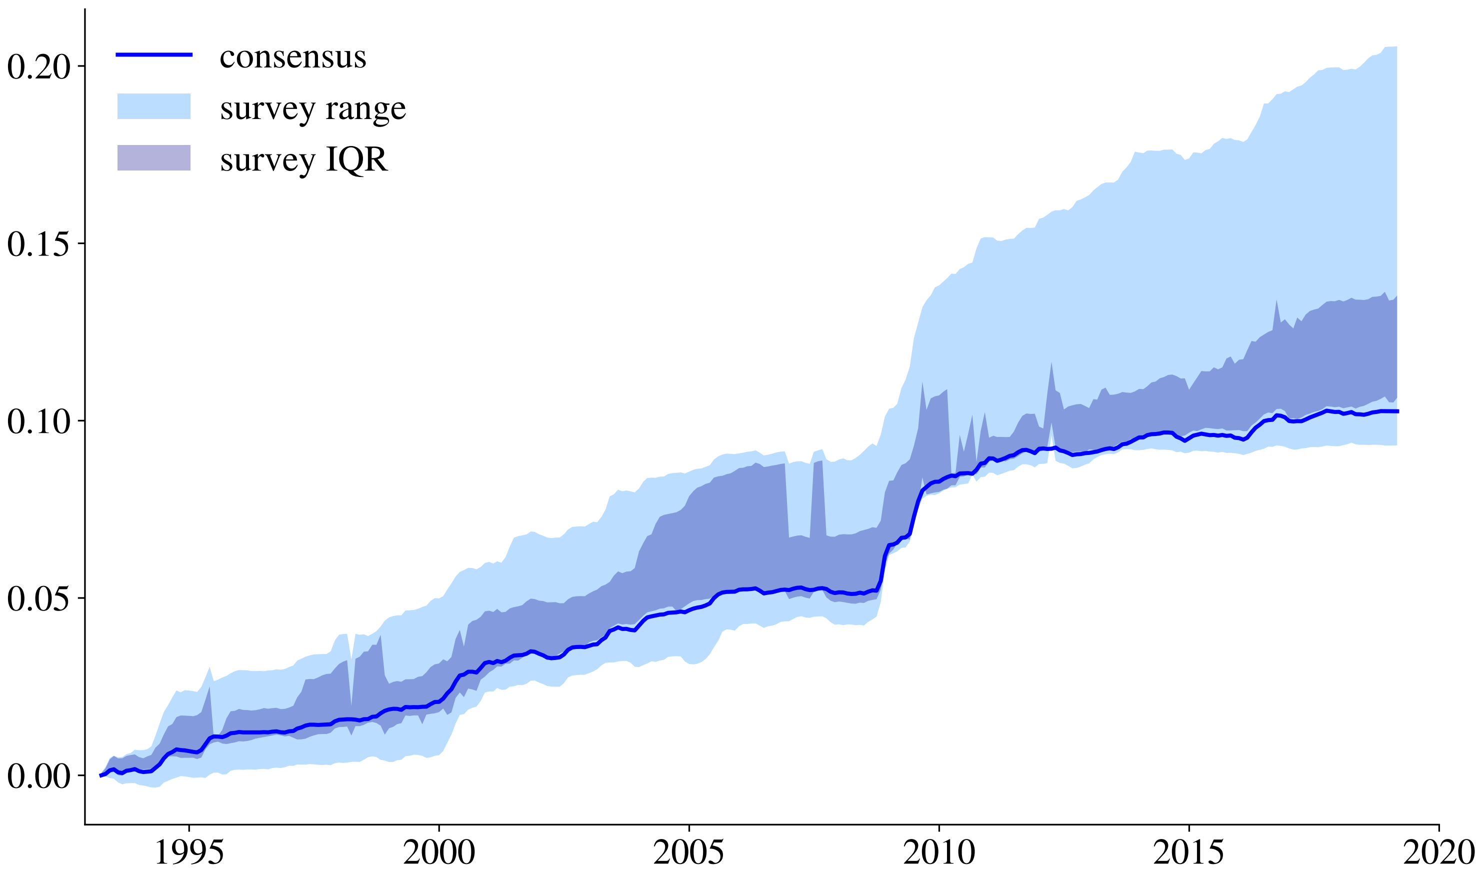

The figure shows the cumulated mean squared forecasting error (MSE) on a 3-month horizon (for all exchange rates over the three-month horizon) minus that of the panel regressions. The blue solid line represents the consensus forecasts, the shaded dark blue band the interquartile range across forecasters, and the shaded light blue band the min-max across forecasters.

Figure 6: Rational and survey risk premia

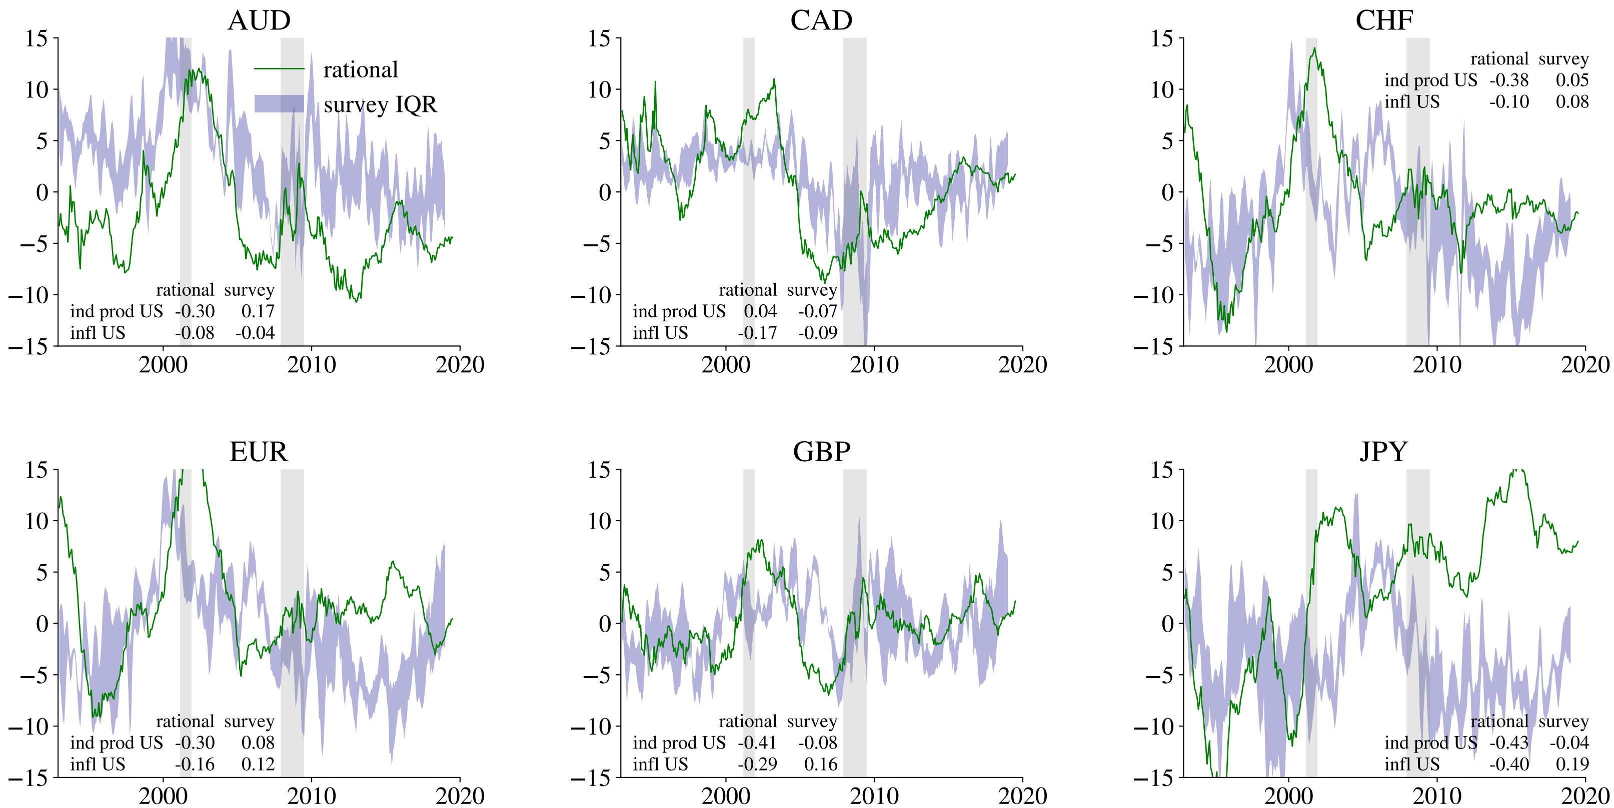

The figure shows 12-month rational risk premia (green solid lines) and the interquartile range of individual survey risk premia (dark blue areas) per currency (AUD, CAD, CHF, EUR, GBP, and JPY) expressed in %. The rational risk premia are the recursively fitted values from regressions on the forward premium, the real exchange rate, and the recent depreciation rate as discussed in Appendix A.1. The gray shaded areas indicate US contractions (peak to trough) as dated by the National Bureau of Economic Research (NBER). Each subfigure reports the correlation between the risk premia and the annual US industrial production growth rate and the annual US inflation rate.

Figure 7: Actual and survey-expected performance of trading strategies

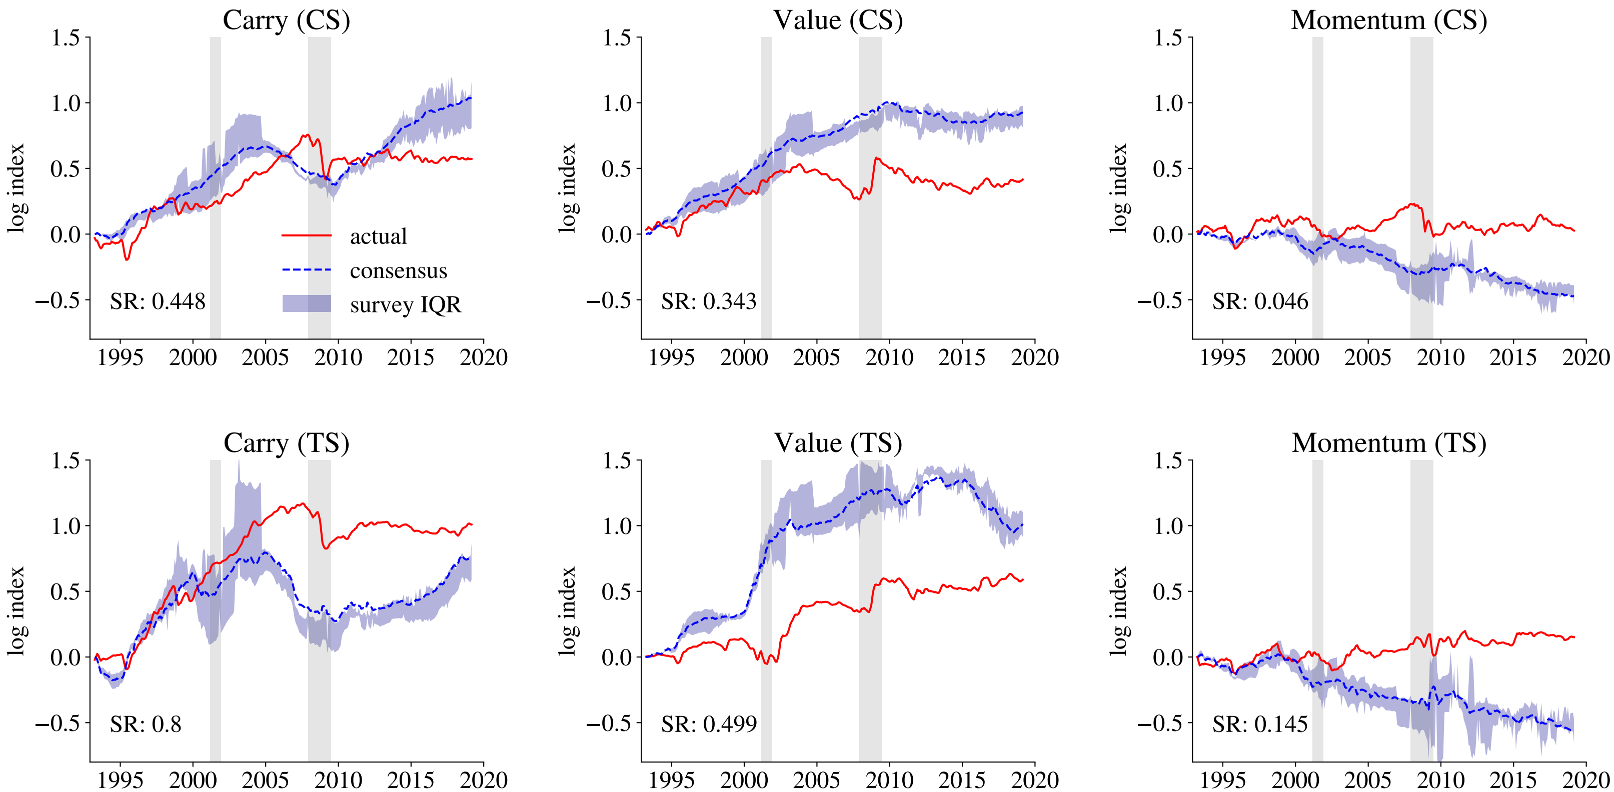

The figure shows cumulated actual and survey-expected excess returns for carry (sorting on the forward premium), value (sorting on the demeaned real exchange rate), and momentum (sorting on the recent excess return) trading strategies. The first row shows cross-sectional (CS) strategies; the second row shows time series (TS) strategies. The holding period is three months and one new portfolio is formed each month; hence, a strategy holds three subportfolios, each with a weight of 1/3, simultaneously. For the CS strategies, three currencies are in the long and short legs, with rank-based weights. TS strategies go long (short) in each currency where the sorting variable is positive (negative). The absolute size of the TS strategies is scaled ex post to have the same volatility as the corresponding CS strategy. The figure also shows the interquartile range of individual survey-expected excess return for those forecasters that submitted a forecast for all currencies in a given month; for the purpose of cumulating, missing interquartile range data are filled with the consensus forecast (for a given month, currency). Each subfigure reports the annualized Sharpe ratio (SR) of the actual excess returns.