Short-Horizon Currency Expectations

Abstract

Previous research finds that analysts' survey-based exchange rate forecasts are at most weakly informative at short horizons. We show, using a microstructural framework, that only the *systematic* component of professional survey forecasts is priced in the cross-section of currency returns, while the idiosyncratic component diversifies away. Using Consensus Economics data, we extract this systematic element by projecting forecasts onto forward discounts, value, and momentum. Portfolios sorted on the resulting risk-based expectation deliver annualized Sharpe ratios between 0.3 and 0.7, remain significant after transaction costs, and are robust across G10 currencies, forecast horizons of 3, 12, and 24 months, and alternative specifications including machine-learning estimators. The systematic component earns a distinct risk premium, with compensation rising when dealer balance-sheet capacity is constrained, and its predictive power strengthens at longer horizons as fundamentals accumulate. Unlike raw forecasts, the risk-based measure also filters out behavioral biases such as overreaction and does not exhibit return reversals. Taken together, these results re-interpret survey expectations as an economically meaningful channel through which global risks are transmitted into exchange rates.

JEL Classification: F31, G12, G15

Keywords: currency expectations, systematic risk, foreign exchange

1 Introduction

Expectations collected from surveys of professional financial analysts occupy a central place in asset pricing. They offer a direct measure of how investors perceive future returns, yet their role in pricing has remained controversial. A large literature documents that survey-based expectations differ systematically from rational benchmark models. For equities, Greenwood and Shleifer (2014) show that investor forecasts track past stock market movements and indexes, but move negatively with model-implied risk premia. In foreign exchange, early work found that consensus survey forecasts were biased and contained little information about near-term currency changes (e.g., Frankel and Froot, 1987; Froot and Frankel, 1989). These findings contributed to the view that surveys are unreliable and largely irrelevant for understanding risk premia.

This paper revisits that conclusion and shows that survey expectations, far from being uninformative, contain a powerful priced component once decomposed into systematic and idiosyncratic elements. We demonstrate that the *systematic component* of analyst forecasts—the part aligned with priced global risk exposures—is strongly predictive of currency excess returns at both short and long horizons. By contrast, the *idiosyncratic component* diversifies away and carries no premium. This decomposition reconciles two longstanding puzzles: (i) why raw survey forecasts have appeared unreliable at short horizons, and (ii) why they nonetheless contain information that some traders and policymakers rely upon.

Our contribution is threefold.

First, we develop a new microstructural framework (Section 2) that connects analyst expectations, customer order flow, and constrained dealer balance sheets. The central mechanism is simple: systematic analyst forecasts generate correlated speculative flows that dealers must warehouse, and because dealer capacity is limited, these positions require compensation. The model yields sharp predictions: only systematic expectations are priced; their associated premium rises when intermediation capacity is tight; and

predictability strengthens at longer horizons as fundamentals accumulate.

Second, we test these predictions empirically. Section 3 describes our dataset, which combines Consensus Economics survey forecasts with spot and forward exchange rates from Barclays/Reuters and auxiliary macro-financial variables. Section 4 presents the main results. Portfolios sorted on the systematic component of expectations earn annualized excess returns of 3–5%, with Sharpe ratios between 0.3 and 0.7, comparable to established strategies such as carry and momentum. By contrast, portfolios sorted on residual forecasts earn no premia. Cross-sectional Fama–MacBeth regressions (Fama and MacBeth, 1973) confirm that the systematic component remains priced after controlling for momentum, value, and forward discounts. Stochastic discount factor tests show that a factor-mimicking portfolio based on systematic expectations earns a significant and distinct price of risk, whereas the dollar factor does not (e.g., Lustig et al., 2011). These results provide direct validation of our theoretical propositions.

Third, we explore mechanisms, horizon dependence, and robustness. We show that systematic expectations line up with contemporaneous order imbalances and trading activity in CLS data, consistent with the model's transmission channel through dealer balance sheets (e.g., Cespa et al., 2022; Hasbrouck and Levich, 2021). We further demonstrate that predictability strengthens at longer horizons, consistent with systematic signals compounding over time, while residual forecasts display behavioral biases such as overreaction and reversals (Coibion and Gorodnichenko, 2015; Bordalo et al., 2020). Robustness checks confirm that the premium survives when we extend to 24-month forecast horizons, use lagged forecasts to address timing, apply alternative weighting schemes, or expand the predictor set with machine learning methods. Across specifications, only the systematic component of survey expectations carries priced information.

Our paper relates to and extends two strands of the literature. The first is the work on survey expectations in asset pricing. While early studies highlighted biases and limited short-run power of FX surveys (e.g., Frankel and Froot, 1987; Froot and Frankel,

1989), more recent contributions emphasize their role at longer horizons. Kremens et al. (2023) show that two-year-ahead Consensus forecasts predict realized appreciations and are largely explained by a handful of macro-financial state variables. Other work links survey expectations to fundamentals, portfolios, and risk premia: Stavrakeva and Tang (2020) highlight alignment with macroeconomic expectations; De Marco et al. (2022) connect information heterogeneity to international portfolio holdings; and Kalemli-Ozcan and Varela (2021) document stylized facts about the UIP premium. Candian and De Leo (2023) model under- and overreaction consistent with horizon-dependent reversals, while Della Corte et al. (2023) and Beckmann and Reitz (2023) show that survey forecasts comove strongly with option-implied quanto premia and intermediary conditions. Related studies examine mispricing and risk premia in currency markets (Bartram et al., 2023) and document the cyclical dynamics of subjective premia and individual forecasts (Nagel and Xu, 2023; Dahlquist and Söderlind, 2023). Our contribution is to move beyond reduced-form predictability by providing a structural mechanism: survey expectations matter because their systematic component drives order flow and is priced through dealer balance-sheet constraints, whereas idiosyncratic noise diversifies away.[1]

The second strand is intermediary asset pricing and FX microstructure. Our framework synthesizes insights from intermediary-based models (e.g., Adrian et al., 2014; He et al., 2017) with the order-flow literature (e.g., Lyons, 1997; Evans and Lyons, 2002). The result is a unified account in which analyst-driven order flow transmits global risks into currency premia through constrained dealer balance sheets.

Our results reposition survey expectations from being dismissed as a behavioral curiosity to being recognized as an economically meaningful channel through which global risk exposures are priced in FX markets. By filtering out idiosyncratic noise, we un-

[1]Relatedly, Adam et al. (2025) show with macro survey data that forecasters overreact to private information and over-anchor to priors, consistent with overconfidence in private signals. While they do not study pricing, their evidence is consistent with our interpretation—see Section 2.8—that overconfidence-driven dispersion in survey responses can mute short-horizon predictability when respondent panels are small.

cover a signal that forecasts returns, earns a distinct premium, and remains robust across horizons, states, and methods.

The rest of the paper is organized as follows. Section 2 develops the theoretical framework and empirical hypotheses. Section 3 describes the data and portfolio construction. Section 4 presents the main empirical results, including portfolio sorts, cross-sectional regressions, factor pricing, machine learning extensions, and horizon dependence. Section 5 provides robustness checks and alternative specifications. Section 6 concludes.

2 Theoretical Framework

We develop a microstructural framework that links survey-based expectations to currency risk premia through the transmission of order flow across constrained dealer balance sheets. The key distinction is between the *systematic component* of analyst forecasts, which aggregates and is priced, and the *idiosyncratic component*, which diversifies and vanishes in equilibrium. The framework is parsimonious: it isolates the institutional mechanisms through which expectations map into market-clearing prices while remaining consistent with the modern theory of intermediary asset pricing.

Building on the general order–flow framework of Evans and Lyons (2002), we develop a stochastic discount factor (SDF) model of foreign exchange (Lustig and Verdelhan, 2012) that incorporates intermediary constraints (e.g., He and Krishnamurthy, 2013; Adrian et al., 2014; He et al., 2017). In this setting, analysts' expectations generate speculative customer flow, dealers absorb that flow subject to balance–sheet limits, and only the systematic component of expectations survives aggregation to earn a risk premium.

The model delivers three central insights. Systematic components of forecasts generate correlated order flow that shifts equilibrium prices, while idiosyncratic noise cancels in aggregation. The premia implied by this mechanism are *state dependent*, rising when dealer capacity is tight and compressing when balance sheets are unconstrained. Finally,

the framework yields sharp, testable predictions: survey-based premia should covary with measures of dealer balance-sheet tightness and strengthen at longer horizons as fundamentals accumulate.

For expositional clarity, the subsections that follow develop the model at a largely intuitive level; a complete formalization—with assumptions, aggregation results, and proofs of all propositions—is given in Internet Appendix Section A.

2.1 Model Overview

The model links analyst forecasts to currency risk premia through a three–stage mechanism. First, analysts issue short–horizon forecasts of currency appreciation against the U.S. dollar. Second, end–users (corporates and investors) translate these forecasts into order flow, generating both a systematic component—correlated with an observed vector of state variables—and idiosyncratic noise. Third, FX dealers absorb the resulting net order flow subject to balance–sheet constraints, thereby determining equilibrium prices.

In large analyst panels, the systematic component of forecasts dominates, while residual noise diversifies away. Dealer balance–sheet capacity governs the price of risk: tighter capacity amplifies premia, while slack capacity compresses them. The resulting equilibrium delivers sharp testable predictions: only systematic expectations are priced, premia are countercyclical in dealer capacity, and predictability strengthens with horizon and in tight intermediation states.

2.2 Notation

Let $i \in \{1, \ldots, \mathrm{J}\}$ index currencies against the U.S. dollar. We adopt the *USD-per-foreign-currency* quotation convention: the spot exchange rate $\mathrm{S}_{i,t}$ denotes the U.S.-dollar price of one unit of foreign currency, so that an increase in $\mathrm{S}_{i,t}$ corresponds to a dollar depreciation. Let $\mathrm{F}_{i,t}$ denote the forward exchange rate contracted at t for delivery at $t{+}1$, with $s_{i,t} \equiv \log \mathrm{S}_{i,t}$ and $f_{i,t} \equiv \log \mathrm{F}_{i,t}$ the corresponding log rates. The U.S. risk-

free interest rate is r_t, and the foreign risk-free rate for currency i is $r_{i,t}$. We impose covered interest parity (in logs),

$$f_{i,t} = s_{i,t} + r_t - r_{i,t},$$

and take the forward to settle at the end-of-(t+1) spot, denoted $s_{i,t+1}$. The (log) excess return on a long foreign-currency position is

$$rx_{i,t+1} \equiv s_{i,t+1} - f_{i,t}, \quad \text{so that} \quad \mathbb{E}_t[rx_{i,t+1}] = \mathbb{E}_t[s_{i,t+1} - s_{i,t}] - (r_t - r_{i,t}).$$

Under this quotation convention, a positive forecast $\mathrm{AE}^j_{i,t}$ corresponds to an expected *appreciation of the foreign currency* (a USD depreciation) between t and t+1.

2.3 Environment

We adopt a CARA–Normal environment: currency payoffs are conditionally Normal and dealers have constant absolute risk aversion (CARA). Three types of agents interact:

- **Analysts.** N professional forecasters who issue predictions of currency appreciation against the U.S. dollar based on a vector of state variables observed at time t. Their forecasts serve as signals but they do not trade.

- **End-users.** Either corporates who trade for hedging or liquidity, generating largely idiosyncratic order flows, or speculators who trade by following analysts' forecasts. The number of end-users is assumed large but finite and is shared among analysts. We thus adopt a finite-clientele normalization: the client flow associated with analyst j—both the speculative demand of j's followers and a proportionate share of corporate idiosyncratic flow—is scaled by 1/N. The exact number of end-users is immaterial for equilibrium.

- **Dealers.** A competitive sector of market-making intermediaries that absorbs net customer flow and warehouses inventory. Their effective risk tolerance varies with balance-sheet capacity. In equilibrium, they are the marginal investors who set prices.

2.4 Signals and Net Inventory

At the start of period t, each analyst j observes the vector of state variables $\mathbf{Z}_t$ and issues a forecast $\mathrm{AE}^j_{i,t}$ of currency i's appreciation against the USD, $(s_{i,t+1} - s_{i,t})$. The systematic component of forecasts, common across analysts, is defined as

$$\widehat{\mathrm{AE}}_{i,t} \;\equiv\; \mathbb{E}[\,\mathrm{AE}^j_{i,t} \mid \mathbf{Z}_t\,],$$

and we write

$$\mathrm{AE}^j_{i,t} \;=\; \widehat{\mathrm{AE}}_{i,t} \;+\; \mathrm{AE}^{r,j}_{i,t}, \qquad \mathbb{E}[\mathrm{AE}^{r,j}_{i,t} \mid \mathbf{Z}_t] = 0. \tag{1}$$

End-user flow associated with analyst j is

$$c^j_{i,t} \;=\; \frac{1}{\mathrm{N}}\Big(\beta\,\mathrm{AE}^j_{i,t} \;+\; \tilde{c}^j_{i,t}\Big), \tag{2}$$

where $\beta > 0$ is a demand sensitivity and $\tilde{c}^j_{i,t}$ captures idiosyncratic (non-informational) trading. The factor $1/\mathrm{N}$ imposes a finite-clientele normalization: as N grows, analysts share a fixed pool of end-users so aggregate volume does not scale mechanically with N.[2]

[2]One could allow a nonlinear demand function $c^j_{i,t} \propto \beta(\mathrm{AE}^j_{i,t})$ with $\beta(\cdot)$ monotone increasing. We adopt the linear form for tractability; the main results are unaffected since idiosyncratic flows still diversify and only the systematic component survives aggregation. Similarly, allowing β to vary across currencies would complicate notation but not alter the central results, because the law of large numbers ensures that heterogeneity in scaling factors washes out in aggregation, leaving only the systematic component priced.

Aggregating across analysts yields

$$q_{i,t} \equiv \sum_{j=1}^{\mathrm{N}} c_{i,t}^{j} = \underbrace{\beta\,\overline{\mathrm{AE}}_{i,t}}_{\text{signal}} + \underbrace{\xi_{i,t}}_{\text{noise}}, \qquad \overline{\mathrm{AE}}_{i,t} \equiv \tfrac{1}{\mathrm{N}} \sum_{j=1}^{\mathrm{N}} \mathrm{AE}_{i,t}^{j}, \tag{3}$$

with residual $\xi_{i,t} \equiv \frac{1}{\mathrm{N}} \sum_{j=1}^{\mathrm{N}} \tilde{c}_{i,t}^{j}$.[3] Stacking across currencies gives $q_t = \beta\,\overline{\mathrm{AE}}_t + \xi_t \in \mathbb{R}^{\mathrm{J}}$.

This decomposition highlights that order flow consists of a systematic *signal* component and an idiosyncratic *noise* component. Under weak cross-analyst dependence and bounded variance,[4] $\operatorname{plim}_{\mathrm{N}\to\infty} \overline{\mathrm{AE}}_{i,t} = \widehat{\mathrm{AE}}_{i,t}$, so the systematic component is $\mathrm{O}_p(1)$, while the noise term satisfies $\xi_{i,t} = \mathrm{O}_p(\mathrm{N}^{-1/2})$ and vanishes in probability. Consequently, the signal-to-noise ratio increases with analyst panel size (is $\mathrm{O}_p(\sqrt{\mathrm{N}})$), and

$$\operatorname*{plim}_{\mathrm{N}\to\infty} q_t = \beta\,\widehat{\mathrm{AE}}_t, \tag{4}$$

where $\widehat{\mathrm{AE}}_t$ stacks the systematic components $\{\widehat{\mathrm{AE}}_{i,t}\}$.[5]

2.5 Pricing Kernel

Equation (4) shows that only the systematic component of analysts' forecasts drives net order–flow imbalance in large-N panels. Dealers accommodate this imbalance one-for-one, so their *inventory* is the opposite of customer flow. Define the dealer's portfolio and its normalized composition as

$$w_t \equiv -\,q_t, \qquad \pi_t \equiv \frac{w_t}{\|w_t\|_1}, \qquad \|w_t\|_1 \equiv \sum_{i=1}^{\mathrm{J}} |w_{i,t}|, \tag{5}$$

[3] For tractability, we assume dealers begin each period with zero net inventory. In FX markets, inventories are rapidly rebalanced across the interdealer network and via prime brokerage, so at the start of each information period only the new net customer imbalance must be warehoused. Allowing carry-over inventories would add persistence—premia would load on current and lagged systematic expectations—but would not alter the core implication that idiosyncratic forecast components diversify away while only the systematic component is priced. We therefore adopt the zero-initial inventory benchmark for clarity.

[4] See Assumption 1(i) in the Appendix for a formal statement.

[5] See Lemma 1 in the Appendix.

so that $\sum_{i=1}^{J} |\pi_{i,t}| = 1$. Since w_t is measured in USD notionals, $\|w_t\|_1$ is total gross USD exposure and π_t captures only portfolio composition. The dealer's excess return per one dollar of gross notional is

$$rx_{t+1}^{\pi} \;=\; \pi_t^{\top} rx_{t+1},$$

where rx_{t+1} is the vector of currency excess returns.

As shown in Internet Appendix Section A.3, with CARA preferences and Normal payoffs, when the representative dealer chooses portfolio positions to maximize expected utility subject to convex balance-sheet constraints (e.g., variance budgets, leverage limits, or capital requirements), the stochastic discount factor (SDF) takes the affine form

$$m_{t+1} = a_t - b_t\, rx_{t+1}^{\pi}, \qquad b_t \equiv \frac{1}{\varphi_t} > 0, \tag{6}$$

where φ_t is a reduced-form measure of aggregate dealer risk tolerance that reflects both baseline risk aversion and the shadow costs of binding balance-sheet constraints. When constraints are slack, φ_t equals the inverse of the coefficient of absolute risk aversion; as constraints tighten, φ_t declines. Lower φ_t (indicating higher effective risk aversion and/or tighter constraints) implies a higher price of risk.[6] This embeds the Evans–Lyons order-flow channel into a linear SDF framework, comparable to intermediary models (He and Krishnamurthy, 2013; Adrian et al., 2014).

Imposing $\mathbb{E}_t[m_{t+1}] = 1$ pins down a_t, and the Euler condition

$$\mathbb{E}_t[m_{t+1} rx_{i,t+1}] = 0$$

implies

$$\mathbb{E}_t[rx_{i,t+1}] = b_t\, \mathrm{Cov}_t(rx_{i,t+1}, rx_{t+1}^{\pi}). \tag{7}$$

[6] See Adrian and Shin (2011) on balance-sheet capacity. A formal derivation of (6) and the reduced-form expression for φ_t is provided in Internet Appendix Section A.3.

Vectorizing,

$$\mathbb{E}_t[rx_{t+1}] = b_t \, \Sigma_t \pi_t, \tag{8}$$

where $\Sigma_t \equiv \mathrm{Cov}_t(rx_{t+1})$.[7]

2.6 Vector Pricing Relation

Equation (8) implies linearity in *dealer* exposures π_t. Combining the large–N result from (4) with $w_t = -q_t$ and the definition in (5) yields

$$\operatorname*{plim}_{\mathrm{N}\to\infty} \pi_t \;=\; -\frac{\widehat{\mathrm{AE}}_t}{\|\widehat{\mathrm{AE}}_t\|_1},$$

so that π_t reflects the (opposite of the) composition of systematic expectations that dealers warehouse. Substituting into (8) yields, in large analyst panels,[8]

$$\mathbb{E}_t[rx_{t+1}] \;=\; -\, b_t \, \Sigma_t \left(\frac{\widehat{\mathrm{AE}}_t}{\|\widehat{\mathrm{AE}}_t\|_1}\right). \tag{9}$$

Equation (9) is the Central Pricing Equation (CPE): expected returns reflect the covariance structure Σ_t scaled by the price of dealer risk b_t and applied to the normalized vector of systematic forecasts, with the minus sign reflecting market clearing ($w_t = -q_t$).[9] Idiosyncratic forecast noise vanishes in aggregation, so only the systematic component $\widehat{\mathrm{AE}}_t$ is priced.[10] Intuitively, expected premia compensate dealers according to how strongly each currency's expected payoff co-moves with the portfolio of systematic expectations they are forced to warehouse, with the magnitude of compensation rising when dealer capacity is tight (low φ_t, high b_t) (He and Krishnamurthy, 2013; Adrian

[7] A residual-explicit version appears in Internet Appendix Lemma 2.

[8] Componentwise, $\mathbb{E}_t[rx_{i,t+1}] \;=\; b_t \, e_i^\top \Sigma_t \pi_t$, where e_i denotes the ith basis vector. In the special case in which Σ_t is close to diagonal, this simplifies to an "own-signal" form $\mathbb{E}_t[rx_{i,t+1}] \approx \tilde{c}_{i,t} \, \widehat{\mathrm{AE}}_{i,t}$, with $\tilde{c}_{i,t}$ capturing the effective risk loading on currency i and its sign determined by the s-quote and inventory convention. See Internet Appendix Proposition 3.

[9] See Proposition 2 of the Internet Appendix for a formal statement and proof.

[10] See Corollary 1 of the Internet Appendix for a formal treatment.

et al., 2014). [11]

This framework embeds the classic Evans–Lyons order-flow channel within a modern asset-pricing structure: systematic analyst expectations generate order flow, which in turn shapes dealer portfolios; dealer capacity determines the price of risk; and idiosyncratic flows diversify away.

2.7 Hypotheses

The theoretical framework delivers four central, complementary hypotheses—each tied to a distinct element of the model—which we set out in this section. All hypotheses are developed from formal propositions and corollaries proved in Internet Appendix Section A. Hypothesis 1 follows directly from the aggregation result in Section 2.4 and its formal counterpart, Corollary 1, which shows that only systematic components of analyst forecasts survive diversification and are priced. Hypotheses 2 and 3 derive from the equilibrium pricing relation in Section 2.6, which builds on the CARA–Normal pricing kernel in Proposition 1 and is formalized in Proposition 2; cross-sectional predictability corresponds to Proposition 3, and the positive premium on a –AE-*mimicking* portfolio corresponds to Proposition 4. Finally, Hypothesis 4 builds on the pricing kernel in Section 2.5 and is formalized in Proposition 5, which shows how balance-sheet tightness and persistence of signals jointly shape the state and horizon dependence of predictability.

Hypothesis 1 (Systematic vs. Non-Systematic Expectations)**.** *Only the systematic component* $\widehat{\mathrm{AE}}_{i,t}$ *is priced in currency* i *excess returns. The residual component* $\mathrm{AE}^{r,j}_{i,t} = \mathrm{AE}^{j}_{i,t} - \widehat{\mathrm{AE}}_{i,t}$ *represents forecast noise that is unpriced and diversifiable. Formal counterpart: Corollary 1 in Internet Appendix Section A.*

[11] Note also that analyst participation sharpens the law of large numbers but does not mechanically scale up aggregate order flow. Because each analyst's clientele is scaled by 1/N in (2), the net volume of systematic flow remains $\mathrm{O}_p(1)$ as N grows. Thus, increasing the number of analysts reduces noise through diversification but does not alter the overall magnitude of premia, which depend instead on the composition of $\widehat{\mathrm{AE}}_t$ and the price of risk b_t.

Hypothesis 2 (Cross-Sectional Predictability). *Systematic expectations* $\widehat{\mathrm{AE}}_{i,t}$ *predict the cross-section of currency excess returns. Under the* USD-per-foreign-currency *convention and the dealer-centric pricing relation in* (9)*, currencies with higher* $\widehat{\mathrm{AE}}_{i,t}$ *(stronger expected* foreign *appreciation, i.e., USD depreciation) tend to have* lower *expected excess returns on the foreign-currency leg; equivalently, sorting by* $-\widehat{\mathrm{AE}}_{i,t}$ *yields a ranking that aligns in large panels with the ranking of* $\mathbb{E}_t[rx_{i,t+1}]$*. Formal counterpart: Proposition 3 in Internet Appendix Section A.*

Hypothesis 3 (Portfolio Performance). *A factor-mimicking portfolio that is long currencies with high* $-\widehat{\mathrm{AE}}_{i,t}$ *and short those with low* $-\widehat{\mathrm{AE}}_{i,t}$ *earns a positive premium, and this premium increases when dealer capacity* φ_t *is constrained (equivalently, when the price of risk* $b_t = 1/\varphi_t$ *is high). Formal counterpart: Proposition 4 in Internet Appendix Section A.*

Hypothesis 4 (Horizon and State Dependence). *The predictive power of systematic expectations* $\widehat{\mathrm{AE}}_{i,t}$ *is both state- and horizon-dependent. At short horizons, predictability is strongest when dealer capacity* φ_t *is tight (i.e., when the price of risk* $b_t = 1/\varphi_t$ *is high). At longer horizons, persistence of* $\widehat{\mathrm{AE}}_t$ *implies that predictive power persists and can strengthen as systematic signals accumulate. Formal counterpart: Proposition 5 in Internet Appendix Section A.*

These hypotheses map directly to the empirical analysis in Section 4, where we test whether systematic expectations forecast returns, whether the $-\widehat{\mathrm{AE}}$ mimicking factor earns a risk premium that varies with intermediation tightness, and whether predictability depends on horizon and market conditions.

2.8 Implications for Survey Data and Predictive Regressions

Because empirical work relies on survey data, the theoretical aggregation result has direct implications for how survey forecasts map into both predictive regressions and

portfolio strategies. Section 2.4 shows that, under weak cross-analyst dependence, the cross-sectional average forecast converges in probability to the systematic component, $\text{plim}_{N\to\infty} \overline{AE}_{i,t} = \widehat{AE}_{i,t}$, and aggregate net order flow satisfies

$$\operatorname*{plim}_{N\to\infty} q_t \;=\; \beta\, \widehat{AE}_t \quad \text{(Eq. (4))}.$$

Because $\widehat{AE}_{i,t}$ drives the systematic part of q_t, the dealer-centric pricing relation in Section 2.6 implies

$$\mathbb{E}_t[rx_{t+1}] \;=\; -\, b_t\, \Sigma_t \left(\frac{\widehat{AE}_t}{\|\widehat{AE}_t\|_1} \right) \quad \text{(Eq. (9))},$$

so $\widehat{AE}_{i,t}$ is directly priced in equilibrium (up to the market-clearing sign) and idiosyncratic forecast noise diversifies away. Consequently, in the population predictive regression

$$rx_{i,t+1} \;=\; \alpha_i \;+\; \beta_i^{(rx)}\, \widehat{AE}_{i,t} \;+\; u_{i,t+1}, \tag{10}$$

the slope $\beta_i^{(rx)}$ is pinned down by Equation (9): it reflects how currency i's payoff covaries (via Σ_t) with the normalized portfolio implied by $\widehat{AE}_t$, with the overall *sign* determined by the quoting and inventory conventions.[12] Importantly, $\beta_i^{(rx)}$ is state-dependent: both Σ_t and the price of risk $b_t \equiv 1/\varphi_t$ may vary over time, so regression slopes need not be constant. Even if the covariance structure is assumed time-invariant ($\Sigma_t = \Sigma$), the model highlights that slopes should tighten when dealer balance-sheet capacity is constrained (high b_t). Even when capacity is relatively stable, however, the framework can be used to clarify how survey-based expectations map into predictive regressions.

While our theoretical framework is predicated on forecasts from the full population

[12] Under the *USD-per-foreign-currency* convention with dealer-centric pricing, higher $\widehat{AE}_{i,t}$ (stronger expected *foreign* appreciation / USD depreciation) typically corresponds to *lower* expected excess returns on the foreign-currency leg, so regressions using $\widehat{AE}_{i,t}$ tend to have negative slopes, whereas using $-\widehat{AE}_{i,t}$ flips the sign.

of N analysts, in practice survey data services such as *Consensus Economics* report the mean across a relatively small subset of $\mathrm{M} \ll \mathrm{N}$ respondents, yielding a noisy proxy

$$\widetilde{\mathrm{AE}}_{i,t} \;=\; \overline{\mathrm{AE}}_{i,t} + u_{i,t},$$

where $u_{i,t}$ captures sampling error and residual idiosyncratic noise. Thus, when researchers attempt to estimate $\beta_i^{(rx)}$ by using the survey-mean forecast in a time-series regression (sample size T) of the form[13]

$$rx_{i,t+1} \;=\; \alpha_i \;+\; \beta_i\, \widetilde{\mathrm{AE}}_{i,t} \;+\; \varepsilon_{i,t+1}, \tag{11}$$

this induces classic errors-in-variables attenuation. Under standard stationarity/ergodicity and orthogonality conditions, the estimated slope satisfies

$$\operatorname*{plim}_{\mathrm{T}\to\infty} \hat{\beta}_i \;=\; \beta_i^{(rx)} \cdot \lambda_{\mathrm{M}}, \qquad \lambda_{\mathrm{M}} \;=\; \frac{1}{1 + \dfrac{\mathrm{Var}(u_{i,t})}{\mathrm{Var}(\widehat{\mathrm{AE}}_{i,t})}} \in (0,1). \tag{12}$$

Projecting survey means on observable state variables Z_t to recover $\widehat{\mathrm{AE}}_{i,t} = \mathrm{B}'\mathrm{Z}_t$ is equivalent to instrumenting $\widetilde{\mathrm{AE}}_{i,t}$ with Z_t in two-stage least squares. This instrumental-variables (IV) solution removes measurement error, recovers the systematic component, and restores the structural predictive relation implied by Equation (9); see Internet Appendix Section A.6, Lemma 4 for a formal treatment. [14]

[13] An analogous errors-in-variables attenuation applies to alternative predictive setups used in the literature. In Frankel and Froot (1987) and Froot and Frankel (1989), the dependent variable is typically the log spot change $\Delta s_{i,t+h}$ (often with the forward discount or interest differential as an additional regressor); in Kremens, Martin, and Varela (2023) long-horizon regressions are run for both realized appreciation and excess returns using survey expectations. In each case, replacing the systematic component $\widehat{\mathrm{AE}}_{i,t}$ with the finite-M survey mean $\widetilde{\mathrm{AE}}_{i,t}$ induces the same attenuation factor $\lambda_{\mathrm{M},h} = 1\big/\big(1 + \frac{\mathrm{Var}(u_{i,t|t+h})}{\mathrm{Var}(\widehat{\mathrm{AE}}_{i,t|t+h})}\big)$ (up to the definition of h and any controls such as the interest differential). We focus on regressions of $rx_{i,t+1}$ on average analyst expectations—cf. Equations (10)–(11)—because this aligns most directly with the Central Pricing Equation (9) and avoids confounding the exposition with interest differentials).

[14] Relatedly, Adam et al. (2025) analyze macroeconomic forecasts from the Federal Reserve Bank of Philadelphia's *Survey of Professional Forecasters*. They show that—with a panel of roughly forty

Attenuation also explains why survey-based predictors perform better at longer horizons, as documented by Kremens, Martin, and Varela (2023). The attenuation factor at horizon h is

$$\lambda_{\mathrm{M},h} \;=\; \frac{1}{1+\dfrac{\mathrm{Var}(u_{i,t|t+h})}{\mathrm{Var}(\widehat{\mathrm{AE}}_{i,t|t+h})}} \;\in (0,1). \tag{13}$$

Here $u_{i,t|t+h}$ is the sampling error in the h-horizon survey mean formed at time t; its variance is essentially horizon-invariant (reflecting finite-M sampling noise), so for brevity we write $\mathrm{Var}(u_{i,t})$ elsewhere. By contrast, the variance of the systematic component $\mathrm{Var}(\widehat{\mathrm{AE}}_{i,t|t+h})$ rises with h: as the horizon lengthens, transitory shocks wash out and forecasts increasingly reflect persistent fundamentals shared across analysts. Consequently, $\lambda_{\mathrm{M},h}$ is increasing in h, so predictive slopes estimated on survey means converge toward their true values as the horizon grows.[15] Hence the signal-to-noise ratio rises with horizon, reducing measurement-error attenuation. In this sense, given our theoretical framework that maps analysts' expectations into currency premia, the superior long-horizon performance of survey-based predictors is not a paradox but a natural implication of persistent signals, weakly dependent noise, and finite-M sampling.

The same attenuation logic applies to portfolio construction, which we turn to next.

2.9 Implications for Portfolio Construction Using Survey Data

Portfolio strategies provide another lens on the same pricing relation developed in Section 2.6. Just as predictive regressions are attenuated when survey means are used in place of systematic expectations, so too are portfolio premia and Sharpe ratios. In large

respondents—professional forecasters tend to overreact to private (idiosyncratic) signals and over-anchor on prior beliefs, consistent with overconfidence in private information generating systematic forecast biases. In our setting, overconfidence in (i.e., overweighting) private signals appears as greater cross-analyst dispersion in the residual component $\mathrm{AE}^{r,j}_{i,t}$, which increases $\mathrm{Var}(u_{i,t})$ and amplifies bias in mean forecasts by strengthening finite-M attenuation (i.e., reducing λ_{M} in (12) and (13)).

[15] Formally, this aligns with Proposition 5 in the Internet Appendix (the main theoretical result) and Proposition 7 (its survey-data counterpart): systematic signals accumulate roughly linearly in h, while idiosyncratic noise grows at most at rate $\sqrt{h}$.

analyst populations, sorting currencies by the average forecast $\overline{\mathrm{AE}}_{i,t}$ is equivalent to sorting on the systematic component $\widehat{\mathrm{AE}}_{i,t}$, since $\overline{\mathrm{AE}}_{i,t} \xrightarrow{p} \widehat{\mathrm{AE}}_{i,t}$ as $\mathrm{N} \to \infty$. Under the dealer-centric Central Pricing Equation (9), the *positively priced* mimicking portfolio loads on the *opposite* of systematic expectations. Hence, factor-mimicking portfolios based on $-\overline{\mathrm{AE}}_{i,t}$ deliver the premia predicted by (9) (cf. Internet Appendix Proposition 4).

In practice, however, researchers and portfolio managers observe only the finite-M survey mean $\widetilde{\mathrm{AE}}_{i,t}$, which contains sampling error. Constructing a portfolio strategy by sorting on $\widetilde{\mathrm{AE}}_{i,t}$ therefore attenuates both expected premia and Sharpe ratios, just as regression slopes are attenuated in Section 2.8. Formally, let r^{AE}_{t+1} denote the excess return on a factor-mimicking portfolio formed by sorting currencies on *negative* expected appreciation and normalizing weights to unit ℓ_1 norm, $\sum_i |w^{\mathrm{AE}}_{i,t}| = 1$. If the systematic component were observed, the population expectation would satisfy

$$\mathbb{E}_t[r^{\mathrm{AE}}_{t+1}] \;=\; \theta^\top \Big(\frac{-\widehat{\mathrm{AE}}_t}{\|\widehat{\mathrm{AE}}_t\|_1} \Big), \tag{14}$$

for some vector of risk loadings θ pinned down by Σ_t and b_t. Replacing $\widehat{\mathrm{AE}}_{i,t}$ with $\widetilde{\mathrm{AE}}_{i,t}$ yields

$$\mathbb{E}_t[r^{\mathrm{AE}}_{t+1} \mid \widetilde{\mathrm{AE}}_t] \;=\; \lambda_{\mathrm{M}} \cdot \theta^\top \Big(\frac{-\widehat{\mathrm{AE}}_t}{\|\widehat{\mathrm{AE}}_t\|_1} \Big), \tag{15}$$

with the same attenuation factor $\lambda_{\mathrm{M}} \in (0,1)$ defined in Eq. (12). Appendix Section A.7, Proposition 9 formally establishes this portfolio-level attenuation result. Thus, portfolio premia constructed directly from survey consensus forecasts are biased toward zero when $\mathrm{M} \ll \mathrm{N}$ or when dispersion across analysts is high.[16]

The same attenuation logic applies to Sharpe ratios. Intuitively, using survey means does not materially change the riskiness of the factor portfolio— the volatility of FX returns is driven by market covariances, not by whether expectations are measured from

[16] If one instead sorts on $+\overline{\mathrm{AE}}_{i,t}$ (or $+\widetilde{\mathrm{AE}}_{i,t}$), the factor's sign flips under our convention, so its expected premium is the negative of (14)–(15).

20 respondents or 2000. What survey noise does is blur the signal used to rank currencies, so that expected premia are diluted while portfolio variance is essentially unaffected, and the combined effect shrinks the Sharpe ratio. Formally, let σ_r^2 denote the variance of the factor-mimicking portfolio return. Replacing $\widehat{\mathrm{AE}}_t$ with $\widetilde{\mathrm{AE}}_t$ perturbs expected returns at first order (via the attenuation factor λ_{M}), but alters the variance only at second order (O(1/M)) because the added sampling noise averages out across currencies (cf. Internet Appendix Lemma 3). This variance effect is negligible relative to the mean effect, so we may treat σ_r^2 as common across both cases to a close approximation. Hence

$$\mathrm{SR}^* \;=\; \frac{\theta^\top\big(-\,\widehat{\mathrm{AE}}_t/\|\widehat{\mathrm{AE}}_t\|_1\big)}{\sigma_r}, \quad \mathrm{SR}^{\mathrm{survey}} \;\approx\; \lambda_{\mathrm{M}}\cdot \mathrm{SR}^* \;\le\; \mathrm{SR}^*, \tag{16}$$

where the approximation reflects the fact that portfolio variance differs only at second order O(1/M). Internet Appendix Corollary 3 formalizes this result.

Instrumenting survey means with observable state variables Z_t again restores the systematic component. Sorting on the fitted values $\widehat{\mathrm{AE}}_{i,t} = \mathrm{B}'\mathrm{Z}_t$ eliminates measurement error, so factor premia and Sharpe ratios estimated from IV-based portfolios align with the theoretical benchmark implied by the Central Pricing Equation (9). Appendix Lemma 5 shows that IV sorting recovers $\widehat{\mathrm{AE}}_{i,t}$ and attains SR^* up to $o_p(1)$.

In short, the model predicts that portfolio strategies based on raw consensus forecasts will appear weak when the number of survey respondents, M, is small, but their true risk premia and Sharpe ratios are revealed once the systematic component is recovered. This offers a direct explanation for why Consensus-based portfolio returns often understate the role of expectations, and it provides a clear econometric solution for empirical implementation.

3 Data and Portfolio Construction

This section describes the data sources and the construction of currency excess returns and portfolios used in the empirical analysis.

Spot and Forward Exchange Rates. We obtain daily spot and 1-, 3-, 6-, and 12-month forward exchange rates against the U.S. dollar from Barclays and Reuters via Datastream/Refinitiv, spanning October 1990 to October 2016. The main analysis uses mid quotes, defined as the mean of bid and ask quotes. Transaction costs are incorporated in robustness checks. Following Burnside, Eichenbaum, Kleshchelski, and Rebelo (2011), we construct end-of-month series using the rate on the last trading day of each month, which replicates an investor's real-time implementation.

The baseline sample covers 34 currencies: Australian Dollar, Brazilian Real, Bulgarian Lev, Canadian Dollar, Croatian Kuna, Cypriot Pound, Czech Koruna, Danish Krone, Egyptian Pound, Euro, British Pound, Hungarian Forint, Indian Rupee, Indonesian Rupiah, Israeli Shekel, Japanese Yen, Malaysian Ringgit, Mexican Peso, New Zealand Dollar, Norwegian Krone, Philippine Peso, Polish Złoty, Russian Ruble, Saudi Riyal, Singapore Dollar, Slovak Koruna, South African Rand, South Korean Won, Swedish Krona, Swiss Franc, Thai Baht, and Ukrainian Hryvnia.[17]

To address tradability concerns, we also analyze a smaller liquid subset, namely the G10 group: Australia, Canada, Denmark, the Euro area, Japan, New Zealand, Norway, Sweden, Switzerland, and the United Kingdom. Currencies pegged (partially or fully) to the U.S. dollar remain in the sample since forward contracts were available to investors. Euro area currencies exit the sample upon euro adoption in January 1999 or later entry dates for subsequent members.

To account for implementation costs, we drop currency–time observations with large

[17] Appendix Table [THIS TABLE NEEDS TO BE ADDED] provides a complete list of currencies and coverage by horizon. The coverage is dictated by the availability of analysts' forecasts.

deviations from covered interest parity (CIP), following Lustig, Roussanov, and Verdelhan (2011).[18] These screens ensure empirical returns are consistent with the CIP environment assumed in the theoretical model.

Currency Excess Returns. We define s_t and $f_{t|t+h}$ the logarithm of the time t spot and forward exchange rates. Each currency is quoted against the U.S. dollar such that an appreciation of the U.S. dollar is associated with an increase in s_t. The log excess return (rx_{t+h}) is defined as the payoff of a strategy that buys a foreign currency in the forward market at time t and then sells it in the spot market at maturity (i.e. at time $t+h$). We consider horizons of one ($h=1$), three ($h=3$), six ($h=6$), and twelve ($h=12$) months. The excess return at time $t+h$ can be computed as

$$rx_{t+h} = f_{t|t+h} - s_{t+h}, \tag{17}$$

Thus, the excess return can also be expressed as the sum of the forward discount and the exchange rate return:

$$rx_{t+h} = f_{t|t+h} - s_t - \Delta s_{t+h}, \tag{18}$$

The CIP condition states that the forward discount should be equal to the interest rate differentials, i.e. $f_t - s_t \simeq \hat{i}_t - i_t$, where $\hat{i}_t$ and i_t denote the foreign and domestic riskless interest rates, respectively. Therefore, the excess return could also be written as $rx_{t+h} \simeq (\hat{i}_t - i_t) - \Delta s_{t+h}$. Thus, the currency excess returns can be approximated by the interest rate differential subtracted by the rate of depreciation of the exchange rate.

Analysts' Forecasts. We obtain professional forecasts from Consensus Economics, following Kalemli-Ozcan and Varela (e.g., 2021); Stavrakeva and Tang (e.g., 2020). Each month, Consensus Economics collects forecasts and reports the mean consensus at 3-, 12-,

[18] For example, South Africa from December 2001–May 2004, Malaysia from August 1998–June 2005, and Indonesia from December 2000–May 2007. The precise exclusion rule follows Lustig, Roussanov, and Verdelhan (2011) and is detailed in Appendix C.

and 24-month horizons, released mid-month. For portfolio formation at the end of month t, we use the most recent release available by that date (results are robust to lagging one month for conservatism).

In the theoretical framework (Section 2), we assume that analysts condition their forecasts on contemporaneous state variables $\mathbf{Z}_t$. In the empirical implementation, however, we project forecasts onto lagged predictors $\mathbf{Z}_{i,t-1}$ (see Section 4.1). This adjustment reflects the timing of actual data releases and guards against look-ahead bias, ensuring that the predictors used to extract $\widehat{\mathrm{AE}}_{i,t}$ are information available at the portfolio formation date.

Foreign Exchange Market Volume. We obtain foreign exchange trading volume from CLS, the world's largest settlement institution. CLS settlement members provide transaction details for authentication and matching. Following CLS classification, we treat corporates, funds, and non-bank financials as price takers, and banks as price makers. For each currency, CLS records buy and sell orders separately by participant type.[19] These measures allow us to map the model's systematic flow channel into observables.

Other Predictors. We also collect macro-financial predictors from Global Financial Data (GFD): dividend–price ratio (DP), price-to-earnings ratio (PE), inflation (INF), unemployment (UN), and market capitalization (MKTCAP).[20] The sample spans October 1990 to October 2016.

Currency Systematic-Expectation Portfolios. At the end of each month t, we sort currencies into quintiles based on the systematic component of analyst expectations, $-\widehat{\mathrm{AE}}_{i,t}$. The lowest quintile contains currencies with low values of $-\widehat{\mathrm{AE}}_{i,t}$, and the highest quintile contains those with high values of $-\widehat{\mathrm{AE}}_{i,t}$. Excess returns are equally

[19] See Hasbrouck and Levich (2019, 2021); Cespa, Gargano, Riddiough, and Sarno (2022) for related use of CLS data.

[20] All macro variables are lagged to account for publication delays. The unemployment gap is estimated via a Hodrick–Prescott filter.

weighted within quintiles. The high-minus-low (HML) portfolio buys high-$(-\widehat{\mathrm{AE}})$ currencies and sells low-$(-\widehat{\mathrm{AE}})$ currencies. Portfolios are rebalanced monthly using all available currencies, are USD-neutral by construction, and skip missing observations. Sorting on $-\widehat{\mathrm{AE}}_{i,t}$ is consistent with the theoretical sign convention and, under diagonal or weakly off-diagonal Σ_t, is asymptotically rank-equivalent to sorting on expected returns (Proposition **??**). This portfolio implements Hypotheses 1–3.

Currency Carry Trade Portfolios. At the end of each month t, we sort currencies into quintiles by the forward discount,

$$\mathrm{FD}_{i,t|t+h} \equiv f_{i,t|t+h} - s_{i,t}.$$

The lowest quintile contains low–interest-rate currencies, and the highest quintile contains high–interest-rate currencies. Excess returns are equally weighted within quintiles. The zero-cost carry portfolio (CAR) goes long high–interest-rate currencies and short low–interest-rate currencies. Portfolios are rebalanced monthly with equal weights and USD neutrality as above.

Currency Momentum Portfolios. At the end of each month t, we sort currencies into quintiles based on past excess returns (12-month cumulative return skipping the most recent month, i.e., 12–1). The lowest quintile contains past losers, the highest quintile contains past winners. Excess returns are equally weighted within quintiles. The zero-cost momentum portfolio (MOM) goes long winners and short losers.

Currency Value Portfolios. At the end of each month t, we sort currencies into quintiles based on real exchange rate valuations (log CPI-based real exchange rate relative to its long-run mean). The lowest quintile contains overvalued currencies, the highest quintile contains undervalued currencies. Excess returns are equally weighted within

quintiles. The zero-cost value portfolio (VAL) goes long undervalued and short overvalued currencies.

4 Empirical Results

In this section, we construct the systematic (risk-based) currency expectation measure and present empirical evidence on its asset-pricing implications. The goal is to test the empirical hypotheses derived from the theoretical framework in Section 2: namely, that only the systematic component of analyst expectations is priced, while the idiosyncratic residual is not (Hypotheses 1–4).

4.1 Construction of Systematic Appreciation Expectations

Appreciation expectation. In the data, we only observe the *average* forecast across analysts from Consensus Economics, which corresponds to $\overline{\mathrm{AE}}_{i,t}$ in the theoretical framework of Section 2. Consistent with that framework, we measure analysts' expected *change in the log spot rate* for currency i at horizon $h = 3$ months as

$$\mathrm{AE}_{i,t|t+3} \;=\; \tilde{s}_{i,t|t+3} - s_{i,t}, \tag{19}$$

where $\tilde{s}_{i,t|t+3}$ is the Consensus Economics mean forecast (formed at t) of $\log \mathrm{S}_{i,t+3}$ and $s_{i,t} \equiv \log \mathrm{S}_{i,t}$ is the current log spot. By our quotation convention ($\mathrm{S}_{i,t}$ is the foreign-currency price of one U.S. dollar), a positive $\mathrm{AE}_{i,t|t+3}$ implies expected *dollar* appreciation (i.e., expected depreciation of currency i).[21]

Systematic vs. residual expectations. In line with the decomposition $\mathrm{AE}_{i,t} = \widehat{\mathrm{AE}}_{i,t} + \mathrm{AE}^{r}_{i,t}$ in (1), we extract the *systematic* component of 3-month expectations by

[21] Since all exchange rates are quoted as foreign currency per U.S. dollar, a decrease in $s_{i,t}$ corresponds to an appreciation of the foreign currency. For brevity, we refer to $\tilde{s}_{i,t|t+3} - s_{i,t}$ as the "appreciation expectation," with sign as described.

projecting the cross-analyst mean forecast $\overline{\text{AE}}_{i,t|t+3}$ onto predictors that are known to span cross-sectional variation in currency premia,

$$\mathbf{Z}_{i,t-1} = (\text{Mom}_{i,t-1}, \text{Value}_{i,t-1}, \text{FD}_{i,t-1})'.$$

We use three canonical predictors—momentum (*Mom*), value (*Value*), and the forward discount (*FD*)—shown to price currency returns (e.g., Lustig et al., 2011; Menkhoff et al., 2012b; Chernov et al., 2023). Following Menkhoff et al. (2012b), *Mom* is the cumulative excess return over the past month (alternative formation windows yield similar results). The forward discount is defined as $FD_{i,t-1} \equiv f_{i,t-1|t} - s_{i,t-1}$, and *Value* is a cross-sectional valuation measure constructed as in Chernov et al. (2023) (details in Appendix B).

At each month t, we run the cross-sectional regression across currencies,

$$\overline{\text{AE}}_{i,t|t+3} = \beta_{0,t} + \beta_{1,t}\,\text{Mom}_{i,t-1} + \beta_{2,t}\,\text{Value}_{i,t-1} + \beta_{3,t}\,\text{FD}_{i,t-1} + \varepsilon_{i,t}, \tag{20}$$

and define the fitted values as the *systematic* (risk-based) expectation,

$$\widehat{\text{AE}}_{i,t|t+3} \equiv \widehat{\beta}_{0,t} + \widehat{\beta}_{1,t}\,\text{Mom}_{i,t-1} + \widehat{\beta}_{2,t}\,\text{Value}_{i,t-1} + \widehat{\beta}_{3,t}\,\text{FD}_{i,t-1}, \tag{21}$$

$$\text{AE}^{r}_{i,t|t+3} \equiv \overline{\text{AE}}_{i,t|t+3} - \widehat{\text{AE}}_{i,t|t+3}. \tag{22}$$

Equations (21)–(22) implement the theoretical decomposition in (1): $\widehat{\text{AE}}_{i,t|t+3}$ is the component of forecasts spanned by priced predictors $\mathbf{Z}_{i,t-1}$, while $\text{AE}^{r}_{i,t|t+3}$ captures the idiosyncratic residual. Under the bounded-moment and weak-dependence assumptions from Section 2, the cross-analyst mean $\overline{\text{AE}}_{i,t|t+3}$ provides a consistent proxy for the systematic component $\widehat{\text{AE}}_{i,t|t+3}$, so that the fitted values in (21) inherit the theoretical interpretation in (1).

By Hypothesis 1, only $\widehat{\text{AE}}_{i,t|t+3}$ should be priced, while $\text{AE}^{r}_{i,t|t+3}$ is diversifiable and

unpriced. By Hypothesis 3, portfolios sorted on $-\widehat{\text{AE}}_{i,t|t+3}$ should earn a positive premium, amplified when dealer capacity is tight. This construction ensures that predictors are predetermined, Consensus forecasts are observed in real time, and excess returns are measured out-of-sample.

4.2 Univariate Sorts

Risk-based Appreciation Expectation. We begin by testing the core implications of the theory. Hypothesis 1 states that only the systematic component $\widehat{\text{AE}}_{i,t}$ aligns with priced risk exposures, while residuals are unpriced. Hypothesis 3 further predicts that portfolios formed on $-\widehat{\text{AE}}$ earn positive risk premia.

Accordingly, we examine the cross-sectional pricing ability of (i) raw analyst forecasts ($-$AE), (ii) the systematic component ($-\widehat{\text{AE}}$), and (iii) the residual ($-\text{AE}^r$). Panel A of Table 1 reports average excess returns for quintile portfolios sorted on $-\widehat{\text{AE}}$, for holding periods $h \in \{1, 3, 6, 12\}$ months and forecast horizons $f \in \{3, 12\}$ months.[22] For each configuration, the table displays annualized currency excess returns (left panel) and exchange-rate changes (right panel).

Example. To fix ideas, consider a June formation using f=3 months and a one-month holding period (h=1): we use June forecasts for t+3 (September), form $-\widehat{\text{AE}}$ at end-June, and hold until end-July. With h=3, positions are held until end-September. Panel A of Figure 1 illustrates f=3, h=1; Panel B shows f=3, h=3. Each month, we regress forecasts on lagged predictors, extract $-\widehat{\text{AE}}$, sort into quintiles, and hold for h months.

Long–short portfolios formed on $-\widehat{\text{AE}}$ deliver annualized excess returns between 2.88 and 5.46 for f=3 and between 3.02 and 5.45 for f=12, with stronger results at shorter holding periods. This directly confirms Hypotheses 1 and 3: because $\widehat{\text{AE}}_{i,t}$ maps into the priced dealer exposure (formal counterpart: Proposition 2; cross-sectional link: Propo-

[22] Results for the 24-month forecast horizon appear in Internet Appendix Table I.A. 1.

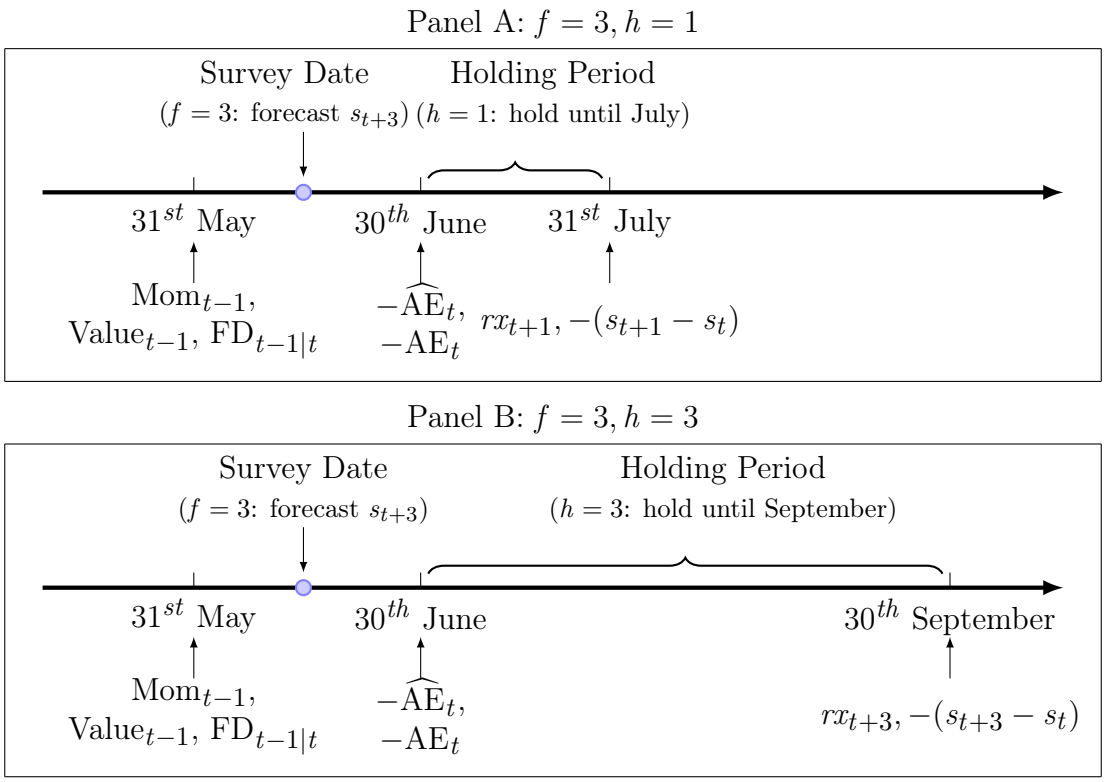


Figure 1 Main Analysis Timeline with Months

The figure illustrates the timeline of our main analysis. In Panel A, we illustrate the 3-month forecast horizon and 1-month holding period. In Panel B, we extend the holding period to 3 months based on the same forecast. The arrows below the main timeline point to the date when the variables are used. The variables included in this figure are momentum (Mom), value (Value), forward discount (FD), risk-based appreciation expectation ($-\widehat{\text{AE}}$), analyst forecast ($-\text{AE}$), excess returns (rx), and change in spot rate ($-(s_{t+3} - s_t)$). The blue dot at the timeline indicates the survey date, when the analysts make the forecasts.

sition 3), it forecasts returns in the cross-section and earns a premium. The right-hand panels show that predictability reflects realized spot movements rather than only forward-discount variation. Predictive power is strongest at short horizons for f=3 and remains robust across holding periods for f=12, consistent with Hypothesis 4 (Proposition 5) that systematic signals compound with horizon. This pattern echoes the classic exchange-rate disconnect (Meese and Rogoff, 1983): short-run forecasts are noisy, while longer horizons reveal fundamentals.[23]

[23]See also Jordà and Taylor (2012) and Filippou, Rapach, Taylor, and Zhou (2023) for complementary evidence on fundamentals and short-run predictability.

Residuals-based Sorts. Panel B of Table 1 reports portfolios sorted on the residual $-\mathrm{AE}^{r}_{i,t}$. Consistent with Hypothesis 1 (formal counterpart: Corollary 1), returns are statistically indistinguishable from zero across horizons and holding periods. Idiosyncratic forecast noise diversifies and does not load on the priced dealer portfolio.

Analyst Forecast Portfolios. Panel C of Table 1 shows results for raw forecasts $-\mathrm{AE}$. In line with prior evidence, analysts do not reliably forecast short-horizon exchange-rate changes: with f=3, spreads are insignificant for h=1 and h=3, becoming significant only for h=6 and h=12. With f=12, the same pattern holds, with modest predictability at longer horizons. This is consistent with Hypothesis 4: raw forecasts mix systematic and residual components, so predictability emerges predominantly when the systematic component dominates at longer horizons.

Cumulative Returns. Figure 2 plots cumulative returns to the $-\widehat{\mathrm{AE}}$ strategy across horizons and rebalancing frequencies. In line with Hypotheses 3 and 4, strategies with less frequent rebalancing capture more persistent fundamental information and generate higher long-run profits. The stronger performance for f=12 relative to f=3 highlights the compounding of systematic expectations. Sharp drawdowns, such as around the Lehman collapse, are consistent with Proposition 2: premia compress when dealer capacity is scarce and rebound as capacity normalizes.

Behavioral Robustness: Over/Undereaction. A potential concern is that analyst forecasts embed systematic biases (over/underreaction) (e.g., Coibion and Gorodnichenko, 2015; Bordalo et al., 2020; Bartram et al., 2023). We therefore study revisions

Table 1 Single Portfolio Sorting on Analyst Forecasts: Excess Return and Spot Rate Change

The table presents the single-sorting quintile annualized excess returns and spot rate changes on the risk-based analyst forecast ($-\widehat{\text{AE}}$), expectation residual ($-\text{AE}^r$), and analyst expectation ($-\text{AE}$). The rows show analyst forecast horizons (f), whereas the columns indicate holding periods (h). The left part of the table shows currency excess returns ($rx_{t+h} = \text{FD}_{t|t+h} - (s_{t+h} - s_t)$), whereas the right part shows spot rate changes. In Panel A (Panel B, Panel C), we report the sorting results based on $-(\widehat{\text{AE}})$ ($-\text{AE}^r, -\text{AE}$). At the end of each month t, we sort currencies based on the Panel-specific signal into quintiles and calculate the average excess return over the holding periods, respectively. We report the time-series average annualized returns of the high-minus-low quintile. Newey and West (1987) adjusted t-statistics are reported in parentheses. Our sample ranges from February 1990 to October 2016.

Forecast horizon f	Holding period h							
	$rx_{t+h} = \text{FD}_{t\|t+h} - (s_{t+h} - s_t)$				-$(s_{t+h} - s_t)$			
	1	3	6	12	1	3	6	12
	Panel A: -$\widehat{\text{AE}}$							
3	5.46	3.36	3.88	2.88	3.79	1.82	2.28	1.45
	(3.17)	(2.65)	(3.85)	(3.68)	(2.19)	(1.44)	(2.25)	(1.71)
12	5.45	4.01	3.99	3.02	4.51	3.08	2.93	2.08
	(3.48)	(3.25)	(3.94)	(3.97)	(2.78)	(2.40)	(2.86)	(2.51)
	Panel B: $-\text{AE}^r$							
3	-1.28	-1.05	-1.49	-0.68	-0.87	-0.75	-1.17	-0.65
	(-0.88)	(-0.86)	(-1.58)	(-0.93)	(-0.61)	(-0.63)	(-1.21)	(-0.84)
12	-0.55	-0.46	-0.80	-0.25	-0.09	-0.14	-0.46	-0.20
	(-0.39)	(-0.40)	(-0.87)	(-0.35)	(-0.07)	(-0.12)	(-0.49)	(-0.26)
	Panel C: -AE							
3	-0.07	0.90	1.84	1.32	-0.85	0.15	1.15	0.48
	(-0.05)	(0.73)	(1.82)	(1.72)	(-0.56)	(0.13)	(1.13)	(0.61)
12	1.70	1.37	1.66	1.44	2.14	1.72	1.97	1.50
	(1.05)	(1.04)	(1.61)	(1.84)	(1.33)	(1.31)	(1.89)	(1.85)

and subsequent forecast errors. For currency i at time t, define for the 3-month horizon:

$$\text{FR}_{i,t,3} = \tilde{s}_{i,t|t+3} - \tilde{s}_{i,t-9|t+3}, \tag{23}$$

$$\text{FE}_{i,t,3} = s_{i,t+3} - \tilde{s}_{i,t|t+3}, \tag{24}$$

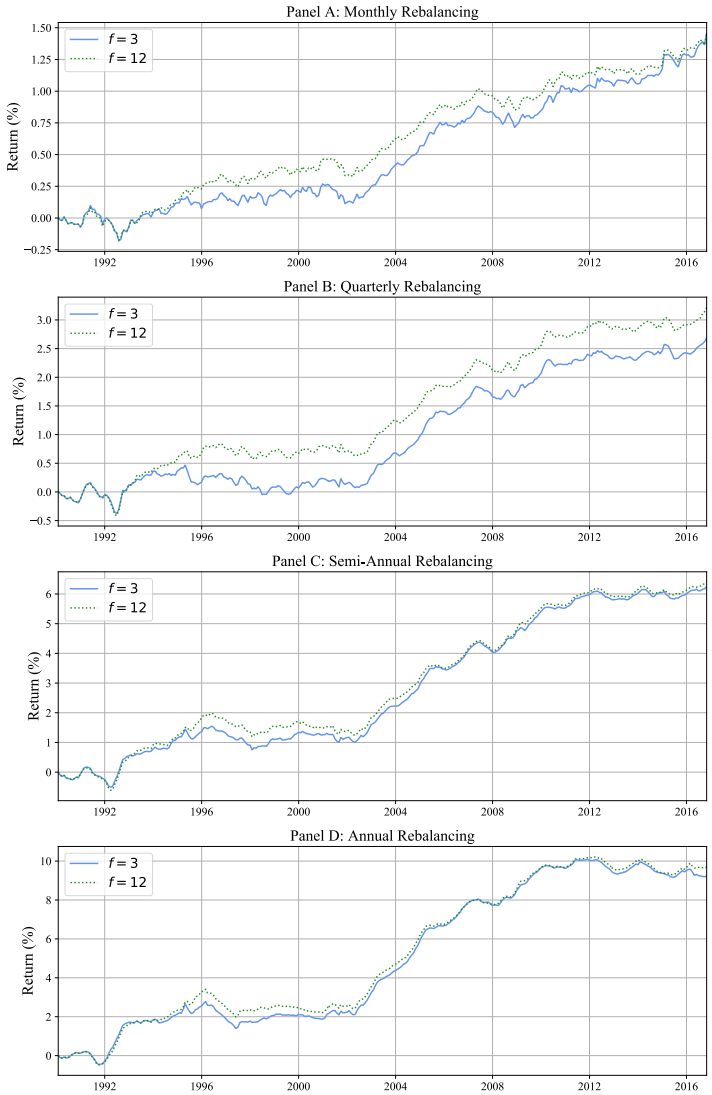


Figure 2 Cumulative Returns

This figure displays cumulative returns of a high-minus-low risk-based appreciation expectation strategy. We report monthly, quarterly, semi-annual, and annual rebalancing; forecast horizons are 3 and 12 months. Sample: February 1990–October 2016.

where $\tilde{s}_{i,t|t+3}$ is the forecast at t for t+3, and $s_{i,t+3}$ is realized log spot. Overreaction implies a negative relation between FE and FR, underreaction implies a positive one.

Table 2 shows that raw expectations overreact: FE loads negatively on FR. Once we isolate $\widehat{\text{AE}}_{i,t}$, this relation disappears. This is consistent with Hypothesis 1: the risk-based measure filters behavioral noise and retains only priced exposures.

Table 2 Error-in-Revision Regression Results

This table reports coefficients from Coibion and Gorodnichenko (2015)'s forecast error on forecast revision regressions. Firstly for currency i at month t, we define forecast revisions as $\text{FR}_{t,3} = \tilde{s}_{i,t|t+3} - \tilde{s}_{i,t-9|t+3}$ and forecast errors $\text{FE}_{t,3} = s_{i,t+3} - \tilde{s}_{i,t|t+3}$. Second, we project $\text{FR}_{t,3}$ onto 1-month lagged Momentum, Value, and Forward Discount and take the predictions (residuals) as systematic (idiosyncratic) forecast revisions. Lastly, we run Fama and MacBeth (1973)'s regressions of forecast error on FR, systematic FR ($\widehat{\text{FR}}$), and idiosyncratic FR ($\widehat{\text{FR}}^r$) in columns 1-3, respectively. Newey and West (1987) adjusted t-statistics are reported in parentheses.

	Panel A: Full Sample			Panel B: G10 Sample		
Spec.	(1) FR	(2) $\widehat{\text{FR}}$	(3) $\widehat{\text{FR}}^r$	(1) FR	(2) $\widehat{\text{FR}}$	(3) $\widehat{\text{FR}}^r$
Intercept	-0.03	-0.03	-0.01	-0.08	-0.07	-0.03
	(-1.65)	(-1.55)	(-0.88)	(-3.52)	(-3.08)	(-1.39)
FR	-0.72	-0.19	-0.98	-0.83	-0.65	-1.07
	(-3.10)	(-0.68)	(-2.99)	(-2.16)	(-1.26)	(-2.10)
Adj-R^2	0.10	0.01	0.13	0.11	0.05	0.09

Behavioral Robustness: Reversals. If residual forecasts overreact, residual-based strategies should exhibit medium-horizon return reversals. Table 3 confirms significant reversals for AE^r portfolios over t+6:t+12 and t+12:t+24, whereas $-\widehat{\text{AE}}$ portfolios do not reverse. This reinforces Hypotheses 1 and 3: profitability is not a byproduct of misreaction, but compensation for risk exposures.

Mechanism: Trading Activity and Order Flow. The framework implies that systematic expectations shape speculative order flow transmitted through constrained dealer

Table 3 Single Portfolio Sorting on Analyst Forecasts: Reversal

This table reports the single-sorting quintile annualized excess returns and spot rate changes on risk-based analyst forecasts ($-\widehat{\text{AE}}$). We consider currency excess returns for the periods of $t+6:t+12$ and $t+12:t+24$ and forecast horizons of three (Panel A) and 12 months (Panel B). Newey and West (1987) adjusted t-statistics are reported in parentheses.

Spec.	Excess Return				Changes in Spot Rate			
	P1(Lo)	P3	P5 (Hi)	Hi-Lo	P1(Lo)	P3	P5 (Hi)	Hi-Lo
	Panel A: f =3							
$t+6:t+12$	-0.61	-0.97	1.73	2.34	-2.00	-2.01	-1.40	0.60
	(-0.47)	(-0.92)	(1.29)	(2.08)	(-1.49)	(-1.91)	(-1.05)	(0.50)
$t+12:t+24$	-0.67	0.31	0.97	1.65	-1.93	-0.68	-2.00	-0.07
	(-0.82)	(0.37)	(0.93)	(2.01)	(-2.16)	(-0.84)	(-1.99)	(-0.08)
	Panel B: f =12							
$t+6:t+12$	-0.51	-0.48	1.88	2.40	-2.25	-1.59	-1.09	1.16
	(-0.39)	(-0.47)	(1.41)	(2.19)	(-1.67)	(-1.57)	(-0.83)	(1.00)
$t+12:t+24$	-0.58	0.26	0.84	1.41	-2.15	-0.69	-2.01	0.14
	(-0.69)	(0.31)	(0.80)	(1.71)	(-2.36)	(-0.87)	(-2.02)	(0.17)

balance sheets (Propositions 2 and 4). We examine contemporaneous participant imbalances using CLS classifications. Define order imbalance

$$\text{OIB}_{i,t} \equiv \frac{\text{Buy}_{i,t} - \text{Sell}_{i,t}}{\text{Buy}_{i,t} + \text{Sell}_{i,t}},$$

where *Buy* (*Sell*) is the dollar value of buyer- (seller-) initiated trades; positive OIB indicates net USD buying.

Univariate sorts. Table 4 sorts OIB on $-\widehat{\text{AE}}$. Consistent with Hypothesis 1, the buy vs. sell split and corporates display systematic co-movement with $-\widehat{\text{AE}}$: corporates are net buyers in high-($-\widehat{\text{AE}}$) portfolios, while funds and non-bank financials do not drive the effect.

Fama–MacBeth regressions. Table 5 regresses OIB on $-\widehat{\text{AE}}$ controlling for *Mom*, *Value*, and *FD*. $-\widehat{\text{AE}}$ significantly predicts buy/sell imbalances and corporates' imbalances, matching the mechanism in which systematic expectations translate into aggregate

Table 4 Single Portfolio Sorting on Risk-based Analyst Forecast: Order Imbalance

This table presents the single-sorting quintile order imbalance on risk-based analyst forecast ($-\widehat{\text{AE}}$). Following Filippou, Taylor, and Wang (2023), we define currency order imbalance ($\text{OIB}_{i,t}$) for currency pair i at t as

$$\text{OIB}_{i,t} = \frac{\text{B}uy_{i,t} - \text{S}ell_{i,t}}{\text{B}uy_{i,t} + \text{S}ell_{i,t}},$$

where $\text{B}uy_{i,t}$ ($\text{S}ell_{i,t}$) represents the dollar volume of buyer (seller)-initiated transactions of the currency pair i at time t. Based on the classification by the CLS dataset, in Panel A, the transactions of currency pair are executed between B*uy side* and S*ell side*, and in Panel B (Panel C, Panel D) transactions between C*orporate* (F*und*, N*on* − B*ank*F*inancial*) as price taker and B*ank* as market maker. Newey and West (1987) adjusted t-statistics are reported in parentheses. Limited to the trade activity data, our sample ranges from September 2012 to October 2016.

Forecast Horizon f	Portfolio					
	1(Lo)	2	3	4	5(Hi)	Hi-Lo
	Panel A: Buy side vs Sell side					
3	-0.54	-0.18	0.09	0.28	0.15	0.68
	(-3.45)	(-1.55)	(1.04)	(3.32)	(2.86)	(4.25)
12	-0.49	-0.22	0.09	0.29	0.14	0.63
	(-4.00)	(-1.67)	(0.98)	(3.55)	(2.70)	(4.66)
	Panel B: Corp. vs Bank					
3	-0.73	-0.20	0.37	2.79	2.97	3.69
	(-1.24)	(-0.29)	(0.50)	(3.74)	(3.78)	(5.10)
12	-0.56	-0.14	0.43	2.29	3.18	3.73
	(-0.93)	(-0.25)	(0.59)	(2.96)	(4.14)	(4.98)
	Panel C: Fund vs Bank					
3	-0.53	-0.11	-0.60	0.19	-0.15	0.38
	(-1.76)	(-0.45)	(-2.32)	(0.88)	(-0.49)	(0.75)
12	-0.56	-0.04	-0.31	0.05	-0.16	0.40
	(-1.93)	(-0.14)	(-1.27)	(0.24)	(-0.56)	(0.78)
	Panel D: Non-Bank Financials vs Bank					
3	-0.84	-0.40	-0.56	-0.34	-0.29	0.55
	(-1.88)	(-1.95)	(-1.99)	(-1.64)	(-2.54)	(1.21)
12	-0.88	-0.35	-0.63	-0.47	-0.21	0.67
	(-1.83)	(-1.50)	(-2.27)	(-1.88)	(-3.11)	(1.39)

order flow at the end of Trading Round 2 (Hypothesis 3).

Mechanism: Trading Volume. We also study unsigned activity, $\text{Volume}_{i,t} \equiv \text{Buy}_{i,t} + \text{Sell}_{i,t}$ (USD billions). Table 6 shows that $-\widehat{\text{AE}}$ negatively predicts cross-sectional trading

Table 5 Fama-MacBeth Regression of Order Imbalance on Risk-based Analyst Forecast

This table presents the Fama-MacBeth regressions of order imbalance on risk-based analyst forecast ($-\widehat{\text{AE}}$). Based on the classification by the CLS dataset, in Panel A, the transactions of currency pair are executed between B*uy side* and S*ell side*, and in Panel B (Panel C, Panel D) transactions between C*orporate* (F*und*, N*on*$-$B*ank*F*inancial*) as price taker and B*ank* as market maker. Currency momentum (*Mom*), value (*Value*) and forward discount (*FD*) serve as controls in this table. Newey and West (1987) adjusted t-statistics are reported in parentheses. Limited to the trade activity data, our sample ranges from September 2012 to October 2016.

Forecast Horizon f	Coef. (t-stat)					
	Intercept	$-\widehat{\text{AE}}$	Mom	Value	FD	adj-R^2
	Panel A: Buy side vs Sell side					
3	0.03	0.25	-0.09	-0.01	5.25	-0.07
	(1.14)	(1.59)	(-0.33)	(-0.23)	(0.74)	
12	0.03	0.23	-0.08	-0.00	5.62	-0.07
	(1.08)	(1.54)	(-0.32)	(-0.07)	(0.74)	
	Panel B: Corp. vs Bank					
3	0.58	3.06	0.62	0.52	-128.71	-0.04
	(1.66)	(1.66)	(0.33)	(0.83)	(-1.49)	
12	0.62	3.20	0.85	0.64	-137.96	-0.04
	(1.59)	(1.57)	(0.44)	(0.88)	(-1.50)	
	Panel C: Fund vs Bank					
3	-0.03	-0.45	1.16	-0.04	42.25	-0.01
	(-0.61)	(-1.15)	(1.57)	(-0.28)	(2.20)	
12	0.00	-0.36	1.18	0.02	34.35	-0.01
	(0.03)	(-0.85)	(1.57)	(0.15)	(1.65)	
	Panel D: Non-Bank Financials vs Bank					
3	-0.06	-0.85	1.23	-0.01	34.41	-0.12
	(-0.76)	(-1.20)	(1.90)	(-0.09)	(1.45)	
12	-0.06	-0.73	1.32	-0.02	34.27	-0.12
	(-1.04)	(-1.27)	(2.03)	(-0.17)	(1.77)	

volume for corporates and funds, with no strong relation for the buy/sell aggregate or non-bank financials.

Table 7 confirms these relations in Fama–MacBeth regressions with controls. Overall, the evidence indicates that the systematic component of expectations helps explain contemporaneous order flow and volume for key participant groups, supporting the model's

Table 6 Single Portfolio Sorting on Risk-based Analyst Forecast: Trading Volume

This table presents the single-sorting quintile trading volume on risk-based analyst forecast ($-\widehat{\text{AE}}$). Following Cespa, Gargano, Riddiough, and Sarno (2022), we define currency trading volume (*Volume*) based on the buy and sell volume of the currency pair i at t as

$$\text{Volume}_{i,t} = \text{Buy}_{i,t} + \text{Sell}_{i,t},$$

where *Buy* (*Sell*) represents the dollar volume of buyer (seller)-initiated transactions of the currency pair i at time t. *Volume* is in the unit of billion. Based on the classification by the CLS dataset, in Panel A, the transactions of currency pair are executed between B*uy side* and S*ell side*, and in Panel B (Panel C, Panel D) transactions between C*orporate* (F*und*, N*on* − B*ank*F*inancial*) as price taker and B*ank* as market maker. Newey and West (1987) adjusted t-statistics are reported in parentheses. Limited to the span of trade activity data, our sample ranges from September 2012 to October 2016.

Forecast Horizon f	Portfolio					
	1(Lo)	2	3	4	5(Hi)	Hi-Lo
	Panel A: Buy side vs Sell side					
3	3137.32	2555.56	1929.07	1778.49	2453.56	-683.76
	(9.03)	(12.21)	(8.09)	(12.47)	(12.46)	(-1.71)
12	2883.41	2755.96	1805.80	1828.42	2490.16	-393.25
	(8.36)	(14.78)	(9.08)	(13.40)	(12.04)	(-0.97)
	Panel B: Corp. vs Bank					
3	17.96	17.39	7.47	3.87	2.73	-15.23
	(5.97)	(7.99)	(3.75)	(3.04)	(6.30)	(-4.89)
12	16.20	18.84	7.18	3.91	2.78	-13.42
	(6.49)	(9.98)	(3.69)	(3.11)	(6.27)	(-5.14)
	Panel C: Fund vs Bank					
3	253.59	226.49	136.14	101.46	125.87	-127.72
	(10.36)	(9.38)	(4.95)	(4.69)	(6.72)	(-4.04)
12	237.56	238.45	132.31	99.92	130.69	-106.88
	(8.83)	(11.59)	(5.04)	(4.98)	(6.48)	(-3.25)
	Panel D: Non-Bank Financials vs Bank					
3	121.34	92.48	67.73	61.55	93.90	-27.44
	(7.11)	(7.98)	(7.33)	(9.47)	(8.72)	(-2.05)
12	108.58	100.83	64.02	64.61	94.82	-13.76
	(7.21)	(8.66)	(7.80)	(9.61)	(8.95)	(-1.08)

transmission channel (Hypothesis 3).

Table 7 Fama-MacBeth Regression of Volume on Expectation Prediction

This table presents the Fama-MacBeth regressions of trading volume (in log) on risk-based analyst forecast ($-\widehat{\text{AE}}$). Based on the classification by the CLS dataset, in Panel A, the transactions of currency pair are executed between B*uy side* and S*ell side*, and in Panel B (Panel C, Panel D) transactions between C*orporate* (F*und*, N*on* − B*ank*F*inancial*) as price taker and B*ank* as market maker. Currency momentum (*Mom*), value (*Value*) and forward discount (*FD*) serve as controls in this table. Newey and West (1987) adjusted t-statistics are reported in parenthesis. Limited to the trade activity data, our sample ranges from September 2012 to October 2016.

Forecast Horizon f	Coef. (t-stat)					
	Intercept	$-\widehat{\text{AE}}$	Mom	Value	FD	adj-R^2
	Panel A: Buy side vs Sell side					
3	5.61	5.43	-6.44	1.99	-348.76	-0.19
	(9.38)	(1.42)	(-1.25)	(1.89)	(-2.13)	
12	5.54	5.41	-6.03	1.99	-329.65	-0.19
	(9.54)	(1.47)	(-1.15)	(1.84)	(-2.08)	
	Panel B: Corp. vs Bank					
3	0.57	2.63	-1.88	5.84	-1218.19	0.03
	(0.63)	(0.58)	(-0.27)	(3.49)	(-4.90)	
12	0.61	3.44	-1.30	5.99	-1230.34	0.03
	(0.62)	(0.69)	(-0.19)	(3.14)	(-4.63)	
	Panel C: Fund vs Bank					
3	3.23	4.62	-6.08	3.29	-542.13	-0.13
	(7.67)	(1.64)	(-1.10)	(3.49)	(-3.11)	
12	3.09	4.00	-5.64	3.10	-505.32	-0.13
	(8.32)	(1.44)	(-1.01)	(3.57)	(-3.07)	
	Panel D: Non-Bank Financials vs Bank					
3	2.60	12.19	-8.21	2.95	-487.95	-0.20
	(3.32)	(2.84)	(-1.27)	(2.25)	(-2.39)	
12	2.45	11.70	-7.64	2.87	-448.71	-0.20
	(3.26)	(2.79)	(-1.16)	(2.15)	(-2.29)	

Transaction Costs. Table 8 reports results adjusted for effective bid–ask spreads. Following Neely and Weller (2013), we scale WM/R quoted spreads by 25% and 50%. Consistent with Hypothesis 3, predictability remains statistically and economically significant after costs, especially for shorter holding periods where turnover is higher.

Table 8 Single Portfolio Sorting on the Risk-Based Analyst Forecast: Transaction Costs

The table presents the single-sorting quintile annualized excess returns on the risk-based analyst forecasts ($-\widehat{\text{AE}}$) after adjusting for the transaction costs. The rows show analyst forecast horizons (f) whereas the columns indicate holding periods (h). According to Menkhoff, Sarno, Schmeling, and Schrimpf (2012b), we proxy the transaction cost by quote spread, which in Panel A (Panel B) we take 25% (50%). At the end of each month t, we sort currencies based on the Panel-specific signal into the quintiles and calculate the average excess return over the holding periods, respectively. We report the time-series average annualized returns of the high-minus-low quintile after adjusting for the quote spread. Newey and West (1987) adjusted t-statistics are reported in parentheses. Our sample ranges from February 1990 to October 2016.

Forecast Horizon f	Holding Period h			
	1	3	6	12
	Panel A: 25% Quote Spread			
3	4.51	2.97	3.64	2.71
	(2.62)	(2.35)	(3.62)	(3.47)
12	4.49	3.62	3.74	2.85
	(2.87)	(2.94)	(3.70)	(3.75)
	Panel B: 50% Quote Spread			
3	3.56	2.59	3.41	2.55
	(2.07)	(2.05)	(3.38)	(3.26)
12	3.52	3.23	3.50	2.69
	(2.25)	(2.62)	(3.46)	(3.53)

G10 Currencies. Table 9 focuses on G10. Excess returns remain in the 3–5% range and are robust to costs, corroborating Hypotheses 1 and 3: systematic expectations are priced even in the most liquid FX segments.

Media Sentiment. Finally, we test whether profitability varies with currency-specific media tone. Following Filippou et al. (2023), define $\text{Sent}_{i,t} = (\text{Positive}_{i,t} - \text{Negative}_{i,t})/n_{i,t}$ using the Loughran and McDonald (2011) dictionary. Double sorts in Table 10 show that the $-\widehat{\text{AE}}$ premium concentrates among *low-sentiment* currencies, consistent with higher required compensation when perceived risk is elevated (Hypothesis 3).

Table 9 Single Portfolio Sorting on the Risk-Based Analyst Forecast: G10 Subsample

The table presents the single-sorting quintile annualized excess returns on the risk-based analyst expectation ($-\widehat{\text{AE}}$) in the G10 subsample. The rows show analyst forecast horizons (f), whereas the columns indicate holding periods (h). At the end of each month t, we sort currencies based on the Panel-specific signal into the quintiles and calculate the average excess return over the holding periods, respectively. We report the time-series average annualized returns of the high-minus-low quintile. Newey and West (1987) adjusted t-statistics are reported in parentheses. Our sample ranges from February 1990 to October 2016.

Forecast Horizon f	Holding Period h			
	1	3	6	12
	Panel A: Without Transaction Costs			
3	5.20	3.93	3.16	3.05
	(2.89)	(2.55)	(2.44)	(3.06)
12	5.14	4.06	3.42	3.10
	(2.83)	(2.67)	(2.62)	(3.22)
	Panel B: 25% Effective Spread			
3	4.60	3.72	3.04	2.96
	(2.56)	(2.42)	(2.35)	(2.97)
12	4.54	3.85	3.29	3.01
	(2.50)	(2.53)	(2.53)	(3.13)
	Panel C: 50% Effective Spread			
3	4.02	3.51	2.92	2.87
	(2.23)	(2.28)	(2.25)	(2.88)
12	3.96	3.64	3.17	2.93
	(2.18)	(2.39)	(2.44)	(3.04)

Sharpe Ratios. Panel A of Table 11 and Figure 3 show that the $-\widehat{\text{AE}}$ strategy delivers annual Sharpe ratios between 0.32 and 0.67, with peaks above 2.0 before the Global Financial Crisis. Performance is comparable to established currency factors (Chernov et al., 2023). Time-variation in Sharpe ratios aligns with Hypothesis 3: premia compress when intermediation capacity expands and spike when dealer balance sheets are tight. While Sharpe ratios are somewhat below carry or momentum, they are economically meaningful given the purely survey-based signal. Positive, significant alphas relative to carry, momentum, value, and dollar factors indicate distinct predictive content.

Table 10 Double-Sorted Risk-Based Analyst Forecast Portfolio: conditional on Sentiment

The table presents annualized average conditional excess returns for double-sorted portfolios. We first allocate all currencies into 2 portfolios on Filippou, Taylor, and Wang (2023)'s sentiment measure with 1-month (12-month) formation period in Panel A (Panel B). Then, within each of the 2 portfolios, currencies are sorted into 3 portfolios based on the risk-based analyst forecasts ($-\widehat{\text{AE}}$). The rows show analyst forecast horizons (f), whereas the columns indicate holding periods (h). We report the time-series average annualized returns of the high-minus-low portfolios. Newey and West (1987) adjusted t-statistics are reported in parentheses. Our sample ranges from February 1990 to October 2016.

Port.	Forecast Horizon f	Holding Period h			
		1	3	6	12
	Panel A: 1-month Sentiment				
Low Sentiment	3	2.35	2.62	3.30	1.73
		(1.39)	(2.13)	(3.21)	(2.20)
	12	2.11	2.27	3.05	1.44
		(1.26)	(1.82)	(2.93)	(1.81)
High Sentiment	3	1.03	0.10	1.34	1.00
		(0.61)	(0.09)	(1.50)	(1.45)
	12	0.75	-0.08	1.31	1.03
		(0.46)	(-0.07)	(1.54)	(1.58)
	Panel B: 12 month Sentiment				
Low Sentiment	3	2.11	1.66	2.93	2.27
		(1.08)	(1.32)	(2.65)	(2.61)
	12	1.67	1.66	2.56	2.04
		(0.80)	(1.26)	(2.22)	(2.31)
High Sentiment	3	2.83	1.03	1.12	0.72
		(1.50)	(0.79)	(1.17)	(1.02)
	12	3.18	1.97	1.74	0.96
		(1.74)	(1.56)	(1.78)	(1.42)

Risk-Adjusted Currency Excess Returns. Panel B of Table 11 shows positive and significant alphas for the $-\widehat{\text{AE}}$ strategy relative to carry, value, momentum, and dollar factors. This confirms Hypothesis 3: the premium associated with systematic expectations is distinct from existing sources and is stronger when intermediation is constrained.

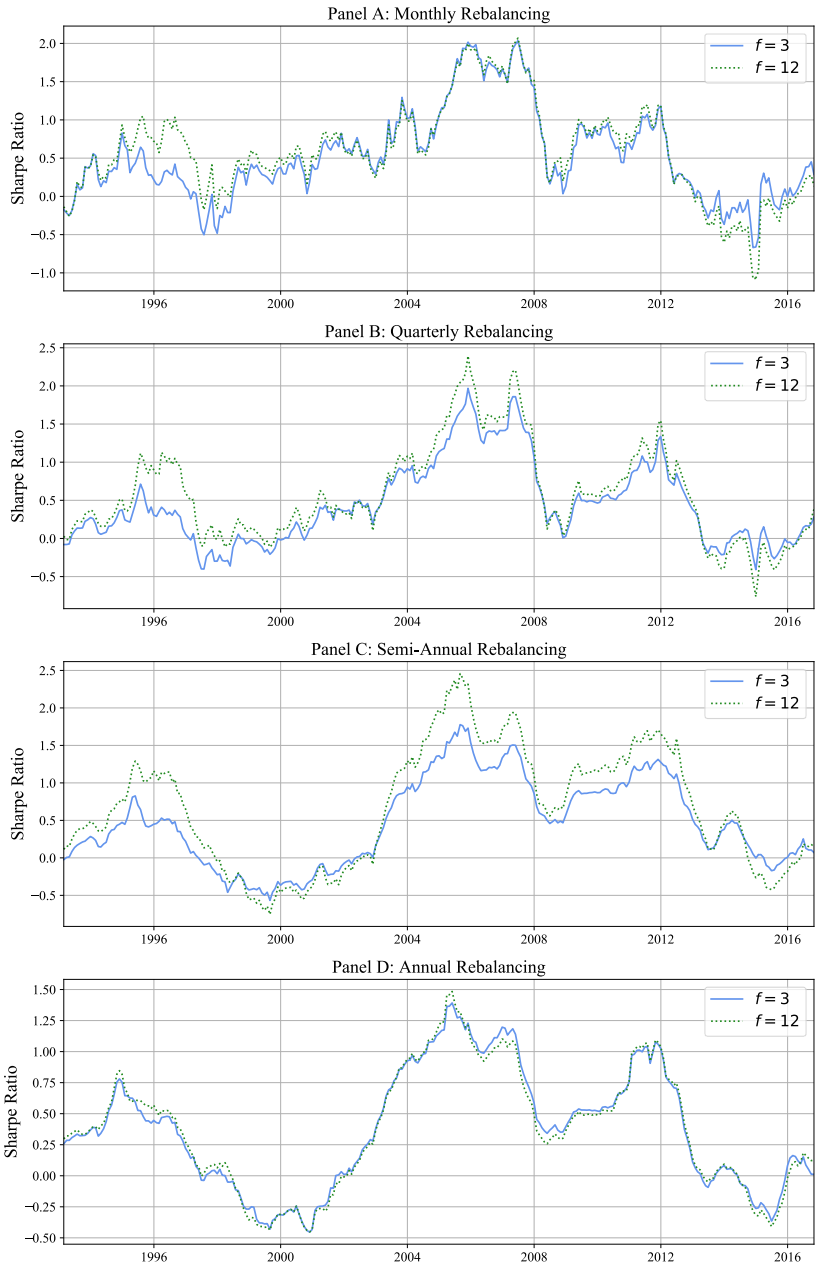


Figure 3 Rolling Sharpe Ratios

Rolling annualized Sharpe ratios (36-month window) for the high-minus-low risk-based appreciation expectation strategy, across rebalancing frequencies and forecast horizons (3 and 12 months). Sample: February 1993–October 2016.

Table 11 Single Portfolio Sorting on the Risk-Based Analyst Forecast: Sharpe Ratio and Alphas

This table reports the Sharpe ratio and alphas of risk-based analyst forecast strategy. The rows show analyst forecast horizons (f) whereas the columns indicate holding periods (h). In Panel A, we report the Sharpe ratio for ($-\widehat{\text{AE}}$) strategy. In Panel B1 (Panel B2, Panel B3), we consider two factors consisting of Dollar and Carry (Value, Momentum), and we include 4 factors in Panel B4. Dollar is defined as the cross-sectional means of currency returns (Lustig, Roussanov, and Verdelhan, 2011). The Carry, Value, and Momentum factors are constructed as the high-minus-low quintile returns sorted on Forward Discount, Value, and Momentum, respectively. We run simultaneous regressions of high-minus-low ($-\widehat{\text{AE}}$) quintiles on these factors and take the constants as the reported alphas. Newey and West (1987) adjusted t-statistics are reported in parentheses. Our sample ranges from February 1990 to October 2016.

Forecast Horizon f	Holding Period h							
	1	3	6	12				
	Panel A: Sharpe Ratio							
3	0.64	0.40	0.45	0.32				
12	0.67	0.47	0.46	0.35				
	Panel B1: Dollar + Carry				Panel B2: Dollar + Value			
3	5.88	3.64	4.16	2.83	5.16	3.00	3.48	2.11
	(3.21)	(2.81)	(4.01)	(3.42)	(2.77)	(2.47)	(3.53)	(2.70)
12	6.00	4.36	4.26	2.97	5.28	3.72	3.56	2.29
	(3.53)	(3.38)	(4.03)	(3.70)	(3.05)	(3.03)	(3.49)	(3.02)
	Panel B3: Dollar + Momentun				Panel B4: Dollar + Carry + Value +Momentum			
3	5.88	3.64	4.00	2.83	5.88	0.83	3.70	2.20
	(3.56)	(2.83)	(3.93)	(3.64)	(3.23)	(2.71)	(3.86)	(2.83)
12	6.00	4.32	4.10	2.98	6.12	4.08	3.78	2.39
	(4.06)	(3.45)	(4.05)	(3.94)	(3.63)	(3.29)	(3.85)	(3.18)

4.3 Cross-sectional Regressions

The portfolio sorts above provide a nonparametric assessment of predictability, but they do not jointly control for other priced predictors. We therefore estimate Fama–MacBeth (Fama and MacBeth, 1973) cross-sectional regressions to test whether the systematic component $-\widehat{\text{AE}}_{i,t}$ remains priced conditional on known currency characteristics. This maps most directly to Proposition 3 (systematic expectations align with priced exposures spanned by $\mathbf{Z}_{i,t}$) and Corollary 1 (idiosyncratic components are unpriced), and underpins Hypothesis 1.

Specification. For each month t, and for each return horizon $h \in \{1, 3, 6, 12\}$, we run the cross-sectional regression

$$rx_{i,t+h} \;=\; a_{t,h} + b_{t,h}\,(-\widehat{\mathrm{AE}}_{i,t}) + \boldsymbol{\gamma}_{t,h}^{\top}\mathbf{X}_{i,t-1} + u_{i,t+h}, \qquad \mathbf{X}_{i,t-1} \equiv (Mom_{i,t-1},\ Value_{i,t-1},\ FD_{i,t-1}) \tag{25}$$

where *Mom*, *Value*, and *FD* are the standard currency predictors (Section 4.1) that span $\mathbf{Z}_{i,t}$ in the theory. We also consider an augmented regression that includes the residual forecast component:

$$rx_{i,t+h} \;=\; a_{t,h} \;+\; b_{t,h}\,(-\widehat{\mathrm{AE}}_{i,t}) \;+\; c_{t,h}\,(-\mathrm{AE}^{r}_{i,t}) \;+\; \boldsymbol{\gamma}_{t,h}^{\top}\mathbf{X}_{i,t-1} \;+\; u_{i,t+h}. \tag{26}$$

We then time-average the cross-sectional slopes $(b_{t,h}, c_{t,h}, \boldsymbol{\gamma}_{t,h})$ to obtain $\bar{b}_h$, $\bar{c}_h$, and $\bar{\boldsymbol{\gamma}}_h$, and compute t-statistics using the time series of monthly estimates with Newey–West corrections (lag $h-1$ to account for overlap). For completeness, we repeat (25)–(26) with $\Delta s_{i,t+h}$ on the left-hand side to show that effects come from realized spot movements as well as forward discounts.

Results. Panel A of Table 12 reports average coefficients for $f=3$ (left: excess returns $rx_{i,t+h}$; right: exchange-rate changes $\Delta s_{i,t+h}$). Columns (1)–(4) correspond to $h \in \{1, 3, 6, 12\}$. The estimated price of the systematic expectation, $\bar{b}_h$, is positive and statistically significant across horizons, and the regressions explain roughly 30–33% of the cross-sectional variation in returns. Panel B presents the corresponding results for $f=12$ with qualitatively similar patterns. In the augmented specification (26), the residual coefficient $\bar{c}_h$ is indistinguishable from zero once $-\widehat{\mathrm{AE}}_{i,t}$ and controls are included.

These findings provide regression-based confirmation of Proposition 3 and Hypothesis 1: the systematic component of survey expectations predicts currency excess returns even conditional on *Mom*, *Value*, and *FD*. They also corroborate Corollary 1: idiosyncratic residuals do not survive conditional tests and are diversifiable in the cross-section.

Table 12 Fama-MacBeth Regressions of Future Excess Returns on the Risk-Based Analyst Forecast

This table reports the results from the Fama and MacBeth (1973) regressions of future excess returns and pure spot rate changes on the risk-based analyst forecast ($-\widehat{\text{AE}}$). The left part of the table shows currency excess returns ($rx_{t+h} = \text{FD}_{t|t+h} - (s_{t+h} - s_t)$) whereas the right part shows spot rate changes. For each month t, we run the cross-sectional regressions as follows:

$$rx_{i,t+h} = \beta_0 + \beta_1(-\widehat{\text{AE}})_{i,t} + \beta_2 \text{Mom}_{i,t} + \beta_3 \text{Value}_{i,t} + \beta_4 \text{FD}_{i,t} + \varepsilon_{t+h}$$

where rx is the currency excess returns, ($-\widehat{\text{AE}}$) is the risk-based appreciation expectation, *Mom* is the currency momentum, *Value* is the currency value, and *FD* is the currency forward discount. In Panel A, the forecast horizon f for analysts' expectations is equal to 3 months, and in Panel B, f is equal to 12. Within each Panel, the h for excess returns are 1, 3, 6, and 12 from Column (1) to (4), respectively. Newey and West (1987) adjusted t-statistics are reported in parentheses. Our sample ranges from February 1990 to October 2016.

	y = rx_{t+h}				y = -$(s_{t+h} - s_t)$			
Spec.	(1) 1 month	(2) 3 month	(3) 6 month	(4) 12 month	(1) 1 month	(2) 3 month	(3) 6 month	(4) 12 month
Panel A: $f = 3$								
Intercept	0.15	0.49	0.66	1.20	0.14	0.30	0.41	0.32
	(0.95)	(1.50)	(1.13)	(1.44)	(0.84)	(0.95)	(0.64)	(0.33)
$-\widehat{\text{AE}}$	3.97	7.70	14.21	20.80	4.03	7.01	13.52	17.85
	(3.02)	(2.60)	(3.05)	(3.45)	(2.87)	(2.43)	(2.81)	(2.78)
Mom	7.15	14.31	6.79	15.48	6.95	14.60	5.45	11.36
	(2.51)	(2.64)	(0.91)	(1.25)	(2.57)	(2.73)	(0.77)	(0.92)
Value	0.19	0.50	0.34	-0.81	0.20	0.08	-0.25	-2.85
	(0.72)	(0.82)	(0.39)	(-0.66)	(0.82)	(0.15)	(-0.27)	(-2.11)
FD	54.49	85.00	160.32	279.45	-48.55	-145.86	-328.60	-536.99
	(1.70)	(0.95)	(1.34)	(1.64)	(-1.49)	(-1.82)	(-2.47)	(-2.66)
adj-R^2	0.31	0.33	0.32	0.30	0.28	0.30	0.30	0.31
Panel B: $f = 12$								
Intercept	0.05	0.34	0.02	0.11	0.06	0.16	-0.22	-0.70
	(0.31)	(1.12)	(0.03)	(0.15)	(0.41)	(0.54)	(-0.38)	(-0.85)
$-\widehat{\text{AE}}$	2.15	6.20	8.16	10.40	2.44	5.56	7.57	8.00
	(2.22)	(2.68)	(1.67)	(1.75)	(2.41)	(2.52)	(1.47)	(1.33)
Mom	7.31	14.09	7.56	16.58	7.03	14.28	6.09	12.27
	(2.55)	(2.58)	(1.03)	(1.35)	(2.58)	(2.63)	(0.88)	(1.01)
Value	0.02	0.32	-0.27	-1.27	0.05	-0.11	-0.83	-3.26
	(0.06)	(0.52)	(-0.29)	(-0.95)	(0.21)	(-0.20)	(-0.81)	(-2.20)
FD	62.61	96.02	201.28	307.81	-42.57	-134.65	-292.57	-519.85
	(1.89)	(1.04)	(1.60)	(1.79)	(-1.26)	(-1.63)	(-2.07)	(-2.52)
adj-R^2	0.31	0.33	0.32	0.30	0.28	0.30	0.30	0.31

4.3.1 Risk-based Factor Pricing.

As a complementary test, we embed the risk-based analyst expectations directly into a stochastic discount factor (SDF) framework. In Section **??**, we showed that the representative dealer's SDF prices a single factor—the normalized dealer excess return rx^{π}_{t+1}—with price of risk $b_t = 1/\varphi_t$, which varies inversely with balance-sheet capacity. Empirically, we proxy this factor by constructing the portfolio-mimicking factor $\mathrm{F}_{-\widehat{\mathrm{AE}}}$, which translates the composition of systematic analyst expectations into excess returns.[24] For comparison with the standard FX asset-pricing literature, we also include the dollar factor DOL, which goes long all foreign currencies against the U.S. dollar.[25]

Following Cochrane (2005), expected excess returns should satisfy the Euler condition

$$\mathbb{E}_t\big[m_{t+1}\, rx_{i,t+1}\big] = 0,$$

for any test asset i, where m_{t+1} is the U.S. investor's SDF and $rx_{i,t+1}$ denotes the excess return. We estimate a two-factor linear SDF of the form

$$m_{t+1} = 1 - b_{\mathrm{DOL}}\big(\mathrm{DOL}_{t+1} - \mu_{\mathrm{DOL}}\big) - b_{-\widehat{\mathrm{AE}}}\big(\mathrm{F}_{-\widehat{\mathrm{AE}},t+1} - \mu_{-\widehat{\mathrm{AE}}}\big),$$

using as test assets six excess-return portfolios: the five quintile portfolios sorted on $-\widehat{\mathrm{AE}}$ and the corresponding high-minus-low spread.

Table 12 reports the second-pass Fama–MacBeth estimates for 3-month and 12-month forecast horizons. The estimated price of risk for $-\widehat{\mathrm{AE}}$, $\lambda_{-\widehat{\mathrm{AE}}}$, is positive and statistically significant across holding periods, consistent with Proposition 4 and Hypothesis 3: compensation for systematic expectations risk is embedded in the cross-section of currency excess returns. By contrast, the dollar factor carries no significant price of risk in

[24] This mapping follows Proposition 4: the $-\widehat{\mathrm{AE}}$ factor mimics the normalized dealer portfolio and therefore captures the priced exposure to systematic expectations.

[25] The inclusion of DOL is not required by the theoretical model but provides a benchmark against existing work such as Lustig et al. (2011) and related studies.

this setting, highlighting that our expectation-based factor captures distinct variation in premia. The magnitudes of $\lambda_{-\widehat{\mathrm{AE}}}$ line up closely with the annualized returns reported in Table 1, reinforcing the interpretation of $-\widehat{\mathrm{AE}}$ as a priced risk factor. Formally, these estimates should be interpreted as *unconditional* average prices of risk; in Section **??** (Proposition 2), the price of risk is state dependent, rising when dealer capacity φ_t is tight and compressing when balance sheets are unconstrained.

This exercise therefore complements the nonparametric portfolio sorts and regression evidence, providing a full asset-pricing validation of Hypothesis 3.

4.4 Estimation of Risk-based Expectations via Machine Learning

Our baseline analysis extracts the systematic component of expectations using only forward discounts, momentum, and value predictors. This parsimonious approach emphasizes economic interpretability and aligns directly with the risk exposures highlighted in Propositions 2 and 3, and with Hypothesis 1. As a robustness check, we now consider a richer predictor set and employ machine learning methods to guard against overfitting. This extension also speaks to Proposition 2 and Hypothesis 3, as it tests whether the priced component is robust to broader information environments and whether its influence varies with intermediation tightness.

Elastic Net Prediction. The small cross-section of currency returns makes Elastic Net (ENet) particularly suitable, since it combines shrinkage and variable selection. Each month, we allow ENet to select among a broad set of macroeconomic and financial variables, producing a fitted systematic expectation $-\widehat{\mathrm{AE}}^{\mathrm{ENet}}$. Details of the variable set and estimation are provided in Appendix C.

Panel A of Table 13 reports annualized average returns of spread portfolios sorted on $-\widehat{\mathrm{AE}}^{\mathrm{ENet}}$. Results closely mirror our main specification: strategies yield positive and statistically significant returns at most horizons, with stronger significance at longer

holding periods. The only exception is the $h = 1$ strategy, which is weaker due to the mismatch between forecast horizon and holding period. These findings reinforce the theoretical prediction in Proposition 3 and Hypothesis 1: systematic components drive premia regardless of the method used to extract them, while residual elements remain unpriced.

Risk-based Currency Momentum. We next examine whether our results are driven by the risk-based currency momentum documented by Filippou, Li, Yuan, and Zhou (2023). Using the same ENet approach, we extract the systematic component of past exchange rate changes, denoted $-\widehat{\Delta s}_{i,t-f}$, and form quintile portfolios. Panel B of Table 13 shows that momentum-based strategies deliver strong short-run premia. However, as predicted by Proposition 5 and Hypothesis 4, our risk-based expectation measure dominates at longer horizons, reflecting the accumulation of fundamental information in analyst forecasts. Panel C formally tests this by regressing the $-\widehat{\text{AE}}$ spread portfolio on risk-based momentum. The positive and significant alpha confirms that risk-based expectations contain incremental information beyond momentum.

Variable Importance. Finally, we examine the variable-importance rankings from the ENet estimation using the SHAP method of Lundberg and Lee (2017). Figure 4 plots the average contribution of each characteristic. Consistent with our parsimonious baseline, the forward discount, momentum, and value emerge as the most influential predictors of systematic expectations. This provides an empirical microfoundation for Proposition 2 and Hypothesis 1: the systematic component of analyst forecasts is spanned by the same priced exposures $\mathbf{Z}_{i,t}$ that dealers require compensation for. Other predictors such as volatility, illiquidity, and macro variables play supporting roles, but the core priced risk drivers are precisely those identified in the theoretical framework.

Table 13 Single Portfolio Sorting on Elastic-Net-Based Analyst Forecasts: Excess Return

The table presents the single-sorting quintile annualized excess returns on the Elastic-Net-based analyst forecast ($-\widehat{\text{AE}}^{\text{ENet}}$) and Elastic-Net-based exchange rate changes. The rows show analyst forecast horizons (f) whereas the columns indicate holding periods (h). For the Elastic-Net-based analyst forecast in Panel A, we obtain the prediction from the predictive cross-sectional Elastic-Net (ENet) regression as detailed in Internet Appendix C. In Panel A, we report the zero-cost portfolio returns sorted on $-\widehat{\text{AE}}^{\text{ENet}}$. In Panel B, we forecast the change in the exchange rate during the past f month ($-\Delta s_{t-f} = s_{t-f} - s_t$ in logs) via ENET regressions and combine the ENet-estimated betas and predictors at the end of month $t-f$ as Elastic-Net-based exchange rate changes. We report the annualized returns of zero-cost portfolios sorted on $-\widehat{\Delta s_{t-f}}^{\text{ENet}}$ in Panel B. In Panel C, we regress the high-minus-low returns in Panel A on the dollar factor and high-minus-low quintile returns in Panel B, and report the constants as αs. Newey and West (1987) adjusted t-statistics are reported in parentheses. Limited to the estimation window requirement, our sample ranges from March 1992 to October 2016.

Forecast Horizon f	Holding Period h			
	1	3	6	12
	Panel A: $-\widehat{\text{AE}}^{\text{ENet}}$			
3	2.26	2.33	2.20	1.87
	(1.55)	(2.66)	(3.75)	(4.54)
12	1.66	2.43	2.06	2.10
	(1.11)	(2.67)	(3.40)	(5.25)
	Panel B: $-\widehat{\Delta s_{t-f}}^{\text{ENet}}$			
1	4.66	1.69	-0.03	0.47
	(2.67)	(1.69)	(-0.04)	(0.94)
3	3.14	0.55	-0.57	0.31
	(1.88)	(0.50)	(-0.72)	(0.59)
12	0.93	0.55	0.40	-0.51
	(0.52)	(0.53)	(0.54)	(-1.05)
	Panel C: αs			
3	0.19	0.61	1.07	1.87
	(1.55)	(1.99)	(2.15)	(2.61)
12	0.16	0.62	1.04	2.02
	(1.15)	(1.90)	(2.02)	(2.99)

4.5 Horizon and State Dependence of Systematic Expectations

This section provides a direct empirical test of Proposition 5 and Hypothesis 4, which predict that the systematic component of analyst expectations should matter differently

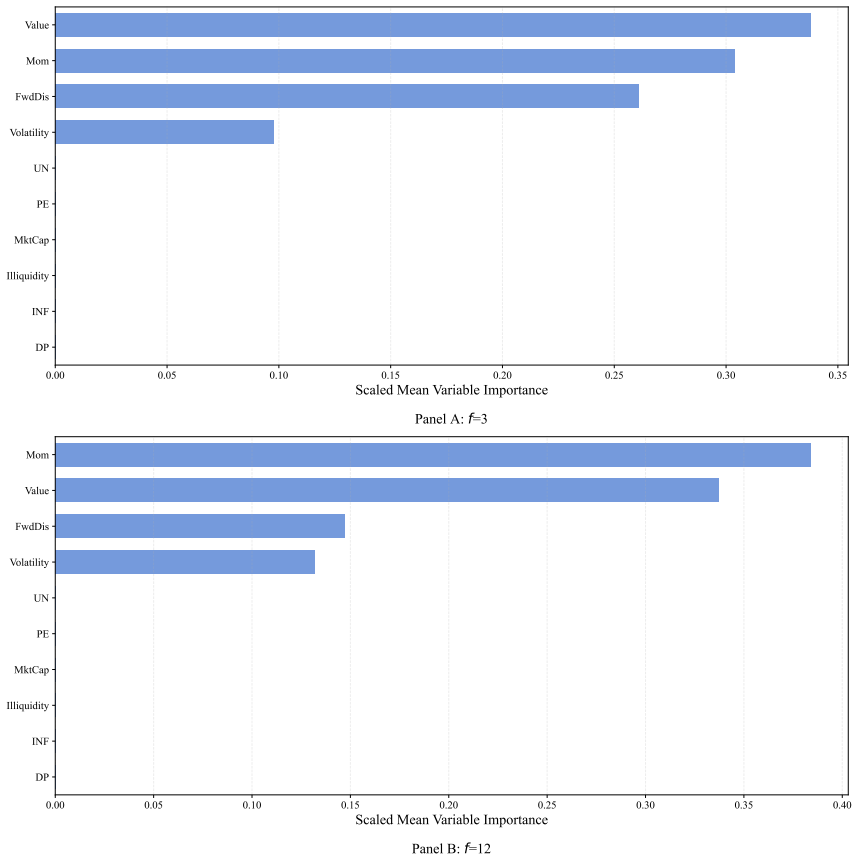


Figure 4 Variable Importance

This figure displays the variable importance of the ENet regression of the appreciation expectation on a broad predictor set. Variable importance is measured using the average absolute SHAP value scaled by the sum across all variables. In Panel A (Panel B), the forecast horizon is 3 months (12 months). The sample ranges from February 1990 to October 2016.

across horizons and states: most influential at short horizons when markets are calm, but also persistent—and potentially stronger—at long horizons as fundamentals accumulate, and amplified when dealer balance-sheet capacity is tight.

The Term Structure of Expectations. In the previous section, we identified forward discount, momentum, and value as the key priced exposures underlying systematic

expectations. We now examine whether these exposures maintain predictive power across horizons. To distinguish between short- and long-term horizons, for each currency and month we construct the difference between the 24-month and the 12-month appreciation expectation:

$$\mathrm{AE}_{t+12,t+24} \;=\; \mathrm{AE}_{t,t+24} - \mathrm{AE}_{t,t+12}.$$

This measure captures the slope of the expectation term structure. Prior work in equities shows that forward return expectations are excessively volatile and countercyclical (Gandhi, Gormsen, and Lazarus, 2023), while in FX the term structure of forward risk premia helps forecast future spot rates (Clarida et al., 2003). Both literatures highlight that horizon fundamentally shapes how expectations map into risk premia.

Table 14 presents estimated coefficients from cross-sectional regressions of appreciation expectations on past forward discounts, momentum, and value. Panel A uses a momentum strategy with a one-month formation and holding period, while Panel B uses a 12-month formation period. In both cases, value and forward discount significantly influence analyst forecasts, while past excess returns contribute little explanatory power. Consistent with Proposition 5, for $\mathrm{AE}_{t+12,t+24}$ the forward discount loads positively on expectations—opposite in sign from its short-horizon relation to spot expectations. This shift highlights that at longer horizons the systematic component becomes more tightly tied to fundamentals, while transitory signals decay.

Raw versus Systematic Forecasts. Kremens, Martin, and Varela (2023) provide related evidence using raw analyst forecasts, documenting long-run predictability but little short-run power. Building on this insight, we compare raw forecasts and risk-based forecasts side by side.

Table 15 reports cross-sectional regressions of currency excess returns on $-\widehat{\mathrm{AE}}$, raw expectations AE, and standard controls, at forecast horizons of 3, 12, and 24 months (Panels A–C). The left panels show results for a one-month holding period and the right

Table 14 The Term Structure of Currency Appreciation Expectation

The table presents the coefficients from Fama and MacBeth (1973)'s regressions of appreciation expectations on common risk factors. The common risk factors indicate currency-level momentum (Mom), value (Value), and forward discount (FD). For dependent variables, we include appreciation expectations with forecast horizon equal to 3, 12, 24 months as well as the difference between the 24-month and 12-month appreciation expectations ($AE_{t+12,t+24}$) in column (1)-(4), respectively. The appreciation expectations are in line with **??**. Especially in column (4), we define the 12-month forward 12-month appreciation expectations as $AE_{t+12,t+24} = AE_{t,t+24} - AE_{t,t+12}$. In Panel A (Panel B), we consider a momentum factor with a formation period of 1 (12) month. Newey and West (1987) adjusted t-statistics are reported in parentheses. Our sample ranges from December 1994 to October 2016.

	(1) $AE_{t,t+3}$	(2) $AE_{t,t+12}$	(3) $AE_{t,t+24}$	(4) $AE_{t+12,t+24}$
	Panel A: 1-month momentum			
Constant	0.03	0.03	0.03	-0.00
	(1.41)	(1.25)	(1.22)	(-0.52)
Mom1	0.03	0.08	0.08	-0.00
	(0.04)	(0.12)	(0.11)	(-0.19)
Value	0.13	0.14	0.15	0.01
	(3.05)	(3.23)	(3.42)	(3.47)
FD	-20.13	-18.23	-16.31	1.92
	(-3.95)	(-3.57)	(-3.10)	(5.55)
Adj-R^2	0.18	0.18	0.18	0.36
	Panel B: 12-month momentum			
Constant	0.05	0.05	0.05	-0.00
	(1.83)	(1.76)	(1.76)	(-0.58)
Mom12	0.42	0.45	0.46	0.00
	(1.67)	(1.84)	(1.88)	(0.14)
Value	0.10	0.12	0.13	0.01
	(2.12)	(2.45)	(2.71)	(4.24)
FD	-20.89	-19.68	-18.06	1.62
	(-3.36)	(-3.22)	(-2.94)	(4.33)
Adj-R^2	0.17	0.17	0.17	0.36

panels for a one-year holding period.

Three findings stand out. First, raw expectations display only marginal predictive power at short horizons, though they gain significance at longer horizons. Second, risk-based expectations significantly predict returns at both short and long horizons, consistent with Propositions 3 and 5. Third, in horse-race regressions, $-\widehat{AE}$ subsumes the

predictive content of raw expectations, which lose significance once both are included, consistent with Corollary 1.

These patterns confirm that the systematic component, rather than idiosyncratic noise, drives currency premia, and that its influence persists—and even strengthens—at longer horizons, exactly as predicted by Hypothesis 4. Moreover, the contrast with raw forecasts highlights the value of filtering out residual noise and behavioral biases: only the systematic component remains priced once horizons extend.

Taken together, these results validate Proposition 5 and Hypothesis 4: systematic expectations are priced at both short and long horizons, their predictive power strengthens as horizons extend, and their premia amplify when dealer balance-sheet capacity is constrained.

5 Robustness and Other Specification Tests

Longer Forecast Horizons. Our main analysis considers forecast horizons of 3 and 12 months. Here, we extend to a 24-month forecast horizon. Table I.A. 1 of the Internet Appendix shows spread portfolio results based on risk-based appreciation expectations at the 24-month horizon. Panel A shows that returns remain statistically significant before transaction costs. After incorporating transaction costs of 25% (Panel B) or 50% (Panel C) of the quoted spread, results remain economically meaningful. Consistent with Proposition 5 and Hypothesis 4, significance is concentrated at long holding periods ($h = 12$), reflecting that systematic information in expectations accumulates and matters most over extended horizons.

Lagged Expectations. In our baseline analysis, we sort currencies using the most recent month's risk-based appreciation expectation. To guard against publication bias, we also sort portfolios using two-month lagged forecasts. Table I.A. 2 of the Internet Appendix shows that strategies based on lagged $-\widehat{\mathrm{AE}}$ remain statistically and economi-

Table 15 Fama-MacBeth Regressions of Future Excess Returns: Short vs Long Horizons

This table reports the results from the Fama and MacBeth (1973) regressions of future excess returns on the risk-based appreciation forecast ($-\widehat{\text{AE}}$) and appreciation forecast ($-\text{AE}$) across short and long forecast horizons. The left part of the table shows 1-month currency excess returns whereas the right part shows 12-month returns. In Panel A (Panel B, Panel C), the forecast horizon f for analysts' expectations is equal to 3 months (12, 24). Controls indicate currency 1-month momentum, value, and forward discount. Newey and West (1987) adjusted t-statistics are reported in parentheses. Our sample ranges from December 1994 to October 2016.

	h= 1-month				h = 12-month			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
	Panel A: $f = 3$							
Intercept	0.11	0.04	0.11	0.00	0.60	0.07	0.60	0.29
	(0.84)	(0.29)	(0.84)	(0.00)	(0.60)	(0.08)	(0.60)	(0.32)
$-\widehat{\text{AE}}$	1.93		1.91	2.25	13.28		12.36	16.01
	(2.67)		(2.72)	(2.29)	(2.75)		(2.55)	(2.53)
$-\text{AE}$		0.23	0.03	0.01		1.87	0.95	1.09
		(1.77)	(0.26)	(0.11)		(2.13)	(1.05)	(1.23)
Control	No	No	No	Yes	No	No	No	Yes
Adj-R^2	0.07	0.04	0.09	0.32	0.09	0.05	0.13	0.36
	Panel B: $f = 12$							
Intercept	0.10	0.03	0.10	0.07	0.49	0.06	0.49	-0.29
	(0.73)	(0.28)	(0.73)	(0.43)	(0.52)	(0.08)	(0.52)	(-0.35)
$-\widehat{\text{AE}}$	1.72		1.69	2.53	11.63		10.69	12.88
	(2.51)		(2.54)	(2.46)	(3.27)		(3.07)	(1.98)
$-\text{AE}$		0.23	0.03	0.02		1.88	0.96	1.09
		(1.81)	(0.27)	(0.12)		(2.15)	(1.06)	(1.22)
Control	No	No	No	Yes	No	No	No	Yes
Adj-R^2	0.06	0.04	0.09	0.32	0.08	0.05	0.12	0.37
	Panel C: $f = 24$							
Intercept	0.06	0.03	0.06	-0.00	0.41	0.05	0.41	-0.36
	(0.47)	(0.27)	(0.47)	(-0.02)	(0.46)	(0.07)	(0.46)	(-0.42)
$-\widehat{\text{AE}}$	1.52		1.49	1.97	11.29		10.13	12.77
	(2.39)		(2.39)	(2.27)	(3.83)		(3.69)	(1.94)
$-\text{AE}$		0.22	0.04	0.02		1.96	1.14	1.20
		(1.79)	(0.30)	(0.16)		(2.21)	(1.26)	(1.35)
Control	No	No	No	Yes	No	No	No	Yes
Adj-R^2	0.06	0.04	0.08	0.32	0.07	0.05	0.10	0.36

cally significant. By contrast, strategies based directly on raw analyst forecasts weaken substantially. This provides further evidence consistent with Proposition 3 and Hypothesis 1: it is the systematic component, rather than idiosyncratic forecast noise, that drives

premia.

Volatility and Illiquidity States. We next examine whether results vary across volatility and illiquidity states. Following Menkhoff, Sarno, Schmeling, and Schrimpf (2012a); Filippou, Gozluklu, and Taylor (2018), we construct monthly volatility and illiquidity measures.[26] Table I.A. 3 of the Internet Appendix shows that predictability is strongest among low-volatility currencies, though results remain significant at longer horizons even for high-volatility portfolios.

The illiquidity results are strongly consistent with the model's prediction that premia rise when dealer balance sheets are constrained and the price of risk $b_t = 1/\varphi_t$ increases—these results should be emphasized. For illiquidity (Table I.A. 4), results are much stronger in high-illiquidity portfolios. This aligns with Proposition 2 and Hypothesis 3, which predict that premia linked to systematic expectations are more pronounced when intermediation frictions are elevated. They also echo Proposition 5, since the amplifying effect of constrained intermediation is clearest at short horizons. By contrast, while volatility captures transitory noise and uncertainty, it does not map one-for-one into balance-sheet tightness. Higher volatility increases idiosyncratic variance (which diversifies away) more than systematic exposure, which can dampen the measured $-\widehat{\mathrm{AE}}$ premium even if the structural channel remains intact. This distinction highlights that it is intermediation constraints, not raw volatility per se, that drive the pricing of systematic expectations.

Uncertainty States. We condition on currencies' exposures to global uncertainty measures. Each month we estimate exposures to: (i) economic policy uncertainty (Baker et al., 2016), (ii) monetary policy uncertainty (Baker et al., 2016), (iii) geopolitical risk (Caldara and Iacoviello, 2022), (iv) the financial stress indicator of Püttmann (2018), and

[26] Volatility is defined as $\mathrm{VOL}_{i,t} = (1/\mathrm{T}_t)\sum_{d\in\mathrm{T}_t}|\Delta s_{i,d}|$, where $\Delta s_{i,d}$ is the daily log spot change. Illiquidity is $\mathrm{ILLIQ}_{i,t} = (1/\mathrm{T}_t)\sum_{d\in\mathrm{T}_t}\mathrm{BAS}_{i,d}$, where $\mathrm{BAS}_{i,d}$ is the daily bid–ask spread. T_t is the number of trading days in month t.

(v) the firm-level political risk measure of Hassan et al. (2019). Currencies are split into low- and high-beta portfolios with respect to each measure, and within each group sorted on $-\widehat{\text{AE}}$. Table I.A. 4 shows that premia are stronger for currencies with low sensitivity to economic policy uncertainty, but also for currencies highly exposed to monetary policy uncertainty and geopolitical risk. These findings reinforce Proposition 2 and Hypothesis 3: compensation for systematic expectations risk is state dependent and interacts with broader measures of uncertainty. In particular, the findings are strongly consistent with the model's prediction that premia rise when dealer balance sheets are constrained and the price of risk $b_t = 1/\varphi_t$ increases.

Rank-weighted Portfolios. Our baseline portfolios are equally weighted. To check robustness, we implement rank-weighted portfolios following Chernov, Dahlquist, and Lochstoer (2023); Zhang (2022). Specifically,

$$\text{R}_{p,t+1} = \sum_{i=1}^{\text{N}} w_{pt}^i \text{R}_{t+1}^{ei}, \quad w_{pt}^i = \lambda \left(\text{rank}(z_{pt}^i) - \frac{1}{\text{N}} \sum_{j=1}^{\text{N}} \text{rank}(z_{pt}^j) \right),$$

where z_{pt}^i is the signal, N the number of currencies, and λ a scaling constant that ensures USD neutrality. Panel A of Table I.A. 5 of the Internet Appendix shows that returns are even stronger under rank-weighting. Notably, the one-month holding period strategy, which is weaker in our baseline, becomes highly significant. This provides further evidence consistent with Proposition 3 and Hypothesis 1: returns are robust to alternative weighting schemes and reflect the systematic component of expectations. *This robustness also reflects the theoretical emphasis on composition rather than scale in dealer portfolios (Section* **??***).*

Pooled ENet-Based Measure. We further estimate a pooled ENet model that incorporates a wide predictor set using 12-month rolling windows. Table I.A. 6 of the Internet Appendix reports the results. Panel A shows that $-\widehat{\text{AE}}^{\text{ENet}}$ strategies are sta-

tistically significant across horizons and holding periods. Panel B shows that risk-based currency momentum remains robust at short horizons, while Panel C demonstrates that $-\widehat{\text{AE}}^{\text{ENet}}$ delivers significant alpha relative to momentum. These findings are consistent with Propositions 3 and 1, as well as Hypotheses 1 and 3: risk-based expectations contain incremental priced information beyond momentum.

Alternative Specification: Risk-based Currency Return Expectations. Finally, we compute an alternative measure of risk-based currency return expectations (RCRE) by replacing realized exchange rate appreciation in (**??**) with the systematic component of survey expectations:

$$\text{RCRE}_{i,t} = \text{FD}_{i,t} - \widehat{\text{AE}}_{i,t|t+3},$$

where $\text{FD}_{i,t}$ is the 3-month forward discount. Table I.A. 7 of the Internet Appendix shows that RCRE-based portfolios deliver strong returns across horizons and holding periods, robust to transaction costs. This provides further structural evidence consistent with Propositions 3 and 1, as well as Hypothesis 1: it is the systematic component of analyst expectations—rather than residual noise—that drives currency premia. *This specification mirrors the theoretical condition in Proposition 2, where only* $\widehat{\text{AE}}_t$ *enters pricing via its impact on dealer portfolios.*

6 Conclusion

This paper demonstrates that currency survey expectations contain a priced component that maps systematically into risk premia. Decomposing forecasts into systematic and idiosyncratic parts, we show that only the systematic component—proxied by risk-based appreciation expectations—is priced in the cross-section of currency returns. This resolves a long-standing puzzle: while raw survey forecasts often appear uninformative at short

horizons, once separated into systematic and residual elements, the systematic component predicts both short- and long-run returns, whereas the residual component diversifies away and is unpriced.

Our empirical analysis tests the three hypotheses derived from the microstructural framework. First, in line with Hypothesis 1, only systematic expectations are priced, while residual forecast noise is not. Second, consistent with Hypothesis 3, a factor-mimicking portfolio based on systematic expectations earns a positive and significant risk premium, which rises when dealer balance-sheet capacity is constrained. Third, confirming Hypothesis 4, predictability is horizon- and state-dependent: systematic expectations generate strong short-horizon premia in calm states, and their influence strengthens at longer horizons as fundamentals accumulate.

We further show that the premium associated with systematic expectations is robust across forecast horizons, holding periods, transaction costs, and currency samples, including the G10. Risk-based strategies deliver Sharpe ratios comparable to established currency factors and remain significant after controlling for carry, value, momentum, and dollar factors. Importantly, they retain positive alpha even when estimated via machine learning methods, underscoring that the priced component is robust to the information set used for extraction.

Additional evidence supports the mechanism implied by the model: systematic expectations align with order imbalances and trading volume of key market participants, consistent with dealers transmitting correlated speculative flow through constrained balance sheets. At the same time, the risk-based component filters out behavioral distortions present in raw analyst forecasts, such as overreaction and reversals, ensuring that the priced signal reflects compensation for systematic risk rather than forecast bias.

Taken together, our results show that survey-based expectations, long regarded as uninformative, contain a systematic component that is priced in currency markets. By filtering out idiosyncratic forecast noise, we uncover a signal that forecasts excess returns,

earns a distinct risk premium, and remains robust across horizons, market conditions, and specifications. This reinterpretation positions survey expectations not as a behavioral curiosity, but as a fundamental channel through which global risks are transmitted into exchange rates.

More broadly, our analysis bridges intermediary asset pricing and FX microstructure by showing how survey expectations generate systematic order flow that is transmitted through constrained dealer balance sheets into risk premia. In doing so, we provide a unified interpretation of survey data, order-flow dynamics, and dealer capacity within a coherent asset-pricing framework. The evidence thus repositions survey expectations from a peripheral behavioral artifact to a central component of international asset pricing, highlighting their role in connecting global macro risks to exchange rate determination.

References

Adam, K., Kuang, P., Xie, S., 2025. Overconfidence in Private Information Explains Biases in Professional Forecasts. Journal of Monetary Economics Available online.

Adrian, T., Etula, E., Muir, T., 2014. Financial intermediaries and the cross-section of asset returns. Journal of Finance 69, 2557–2596.

Adrian, T., Shin, H. S., 2011. Financial intermediary balance sheet management. FRB of New York Staff Report .

Baker, S. R., Bloom, N., Davis, S. J., 2016. Measuring economic policy uncertainty. Quarterly Journal of Economics 131, 1593–1636.

Bartram, S. M., Djuranovic, L., Garratt, A., Xu, Y., 2023. Mispricing and risk premia in currency markets. Journal of Financial and Quantitative Analysis Forthcoming.

Beckmann, J., Reitz, S., 2023. Dealer risk premiums in fx forecasts. Available at SSRN 4513275 .

Bordalo, P., Gennaioli, N., Ma, Y., Shleifer, A., 2020. Overreaction in macroeconomic expectations. American Economic Review 110, 2748–2782.

Burnside, C., Eichenbaum, M., Kleshchelski, I., Rebelo, S., 2011. Do peso problems explain the returns to the carry trade? The Review of Financial Studies 24, 853–891.

Caldara, D., Iacoviello, M., 2022. Measuring geopolitical risk. American Economic Review 112, 1194–1225.

Candian, G., De Leo, P., 2023. Imperfect exchange rate expectations. Review of Economics and Statistics Forthcoming.

Cespa, G., Gargano, A., Riddiough, S. J., Sarno, L., 2022. Foreign exchange volume. The Review of Financial Studies 35, 2386–2427.

Chernov, M., Dahlquist, M., Lochstoer, L., 2023. Pricing currency risks. The Journal of Finance 78, 693–730.

Clarida, R. H., Sarno, L., Taylor, M. P., Valente, G., 2003. The out-of-sample success of term structure models as exchange rate predictors: a step beyond. Journal of International Economics 60, 61–83.

Cochrane, J. H., 2005. Asset pricing: Revised edition. Princeton university press.

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

Dahlquist, M., Söderlind, P., 2023. Individual forecasts of exchange rates. Swedish House of Finance Research Paper .

De Marco, F., Macchiavelli, M., Valchev, R., 2022. Beyond home bias: International portfolio holdings and information heterogeneity. The Review of Financial Studies 35, 4387–4422.

Della Corte, P., Gao, C., Jeanneret, A., 2023. Survey expectations meet option prices: New insights from the fx market .

Evans, M. D. D., Lyons, R. K., 2002. Order flow and exchange rate dynamics. Journal of Political Economy 110, 170–180.

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

Filippou, I., Gozluklu, A. E., Taylor, M. P., 2018. Global political risk and currency momentum. Journal of Financial and Quantitative Analysis 53, 2227–2259.

Filippou, I., Li, S. Z., Yuan, P., Zhou, G., 2023. Risk momentum in currencies .

Filippou, I., Rapach, D., Taylor, M. P., Zhou, G., 2023. Out-of-sample exchange rate prediction: A machine learning perspective. Available at SSRN 3455713 .

Filippou, I., Taylor, M. P., Wang, Z., 2023. Media sentiment and currency reversals. Journal of Financial and Quantitative Analysis pp. 1–47.

Frankel, J. A., Froot, K. A., 1987. Using survey data to test some standard propositions regarding exchange rate expectations. American Economic Review 77, 133–153.

Froot, K. A., Frankel, J. A., 1989. Forward discount bias: Is it an exchange risk premium? Quarterly Journal of Economics 104, 139–161.

Gandhi, M., Gormsen, N. J., Lazarus, E., 2023. Forward return expectations. Tech. rep., National Bureau of Economic Research.

Greenwood, R., Shleifer, A., 2014. Expectations of returns and expected returns. The Review of Financial Studies 27, 714–746.

Hasbrouck, J., Levich, R. M., 2019. Fx liquidity and market metrics: New results using cls bank settlement data. Available at SSRN 2912976 .

Hasbrouck, J., Levich, R. M., 2021. Network structure and pricing in the fx market. Journal of Financial Economics 141, 705–729.

Hassan, T. A., Hollander, S., Van Lent, L., Tahoun, A., 2019. Firm-level political risk: Measurement and effects. The Quarterly Journal of Economics 134, 2135–2202.

He, Z., Kelly, B., Manela, A., 2017. Intermediary asset pricing: New evidence from many asset classes. Journal of Financial Economics 126, 1–35.

He, Z., Krishnamurthy, A., 2013. Intermediary asset pricing. American Economic Review 103, 732–770.

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

Kalemli-Ozcan, S., Varela, L., 2021. Five facts about the uip premium. Working Paper .

Kremens, L., Martin, I., Varela, L., 2023. Long-horizon exchange rate expectations. Available at SSRN 4545603 .

Loughran, T., McDonald, B., 2011. When is a liability not a liability? textual analysis, dictionaries, and 10-ks. The Journal of Finance 66, 35–65.

Lundberg, S. M., Lee, S.-I., 2017. A unified approach to interpreting model predictions. Advances in neural information processing systems 30.

Lustig, H., Roussanov, N., Verdelhan, A., 2011. Common risk factors in currency markets. The Review of Financial Studies 24, 3731–3777.

Lustig, H., Verdelhan, A., 2012. Exchange rates in a stochastic discount factor framework. In: James, J., Marsh, I. W., Sarno, L. (eds.), *Handbook of Exchange Rates*, John Wiley & Sons, chap. 14, pp. 391–420.

Lyons, R. K., 1997. A simultaneous trade model of the foreign exchange hot potato. Journal of international Economics 42, 275–298.

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

Menkhoff, L., Sarno, L., Schmeling, M., Schrimpf, A., 2012a. Carry trades and global foreign exchange volatility. The Journal of Finance 67, 681–718.

Menkhoff, L., Sarno, L., Schmeling, M., Schrimpf, A., 2012b. Currency momentum strategies. Journal of Financial Economics 106, 660–684.

Nagel, S., Xu, Z., 2023. Dynamics of subjective risk premia. Journal of Financial Economics 150, 103713.

Neely, C. J., Weller, P. A., 2013. Lessons from the evolution of foreign exchange trading strategies. Journal of Banking & Finance 37, 3783–3798.

Newey, W. K., West, K. D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703–708.

Püttmann, L., 2018. Patterns of panic: Financial crisis language in historical newspapers. Available at SSRN 3156287 .

Stavrakeva, V., Tang, J., 2020. A fundamental connection: Exchange rates and macroeconomic expectations. Working Paper .

Zhang, S., 2022. Dissecting currency momentum. Journal of Financial Economics 144, 154–173.

Internet Appendix for Short-Horizon Currency Expectations

A Formal Framework, Assumptions, and Proofs

This section provides the formal underpinnings for the theoretical results in Section 2 of the main text. We (i) restate notation and standing conventions and specify the decomposition of analyst forecasts (§A.1); (ii) impose residual aggregation and idiosyncratic flow assumptions that deliver a cross-sectional Law of Large Numbers (LLN)/Central Limit Theorem (CLT) (Assumption 1); (iii) establish large–N aggregation and portfolio normalization results (§A.2); (iv) derive the Central Pricing Equation (CPE), linking expected returns to the ℓ_1–normalized systematic expectations (§A.4); and (v) present the main propositions and their proofs, including cross-sectional predictability, the AE-mimicking factor premium, and horizon/state dependence (§A.5). Notation, conventions, and asymptotics are set out in §A.1.

A.1 Setup, Assumptions, and Notation

Standing conventions. All random variables are defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with filtration $\{\mathcal{F}_t\}$; $\mathbb{E}_t[\cdot] \equiv \mathbb{E}[\cdot \mid \mathcal{F}_t]$, and $\mathbb{E}[\cdot \mid \mathbf{Z}_t]$ denotes conditioning on the state vector.

Convergences: $\mathrm{O}_p(\cdot)$, $o_p(\cdot)$, $\xrightarrow{p}$, and $\xrightarrow{\mathrm{L}^2}$ are with respect to $\mathrm{N} \to \infty$ for fixed (i, t). We retain the *USD-per-foreign-currency* quotation convention from the main text (so $\Delta s > 0$ denotes USD *depreciation*).

Norms: $\|\cdot\|_1$ denotes the ℓ_1 norm on $\mathbb{R}^{\mathrm{J}}$; $\|\cdot\|_2$ the Euclidean norm; and for a matrix A, $\|\mathrm{A}\|_{\mathrm{op}} \equiv \sup_{\|x\|_2=1} \|\mathrm{A}x\|_2 = \sigma_{\max}(\mathrm{A})$ (for symmetric positive semidefinite (p.s.d.) A, this equals $\lambda_{\max}(\mathrm{A})$).

Risk and payoffs: $\Sigma_t \equiv \mathrm{Cov}_t(rx_{t+1})$ is symmetric p.s.d.; the dealer price of risk is $b_t \equiv 1/\varphi_t > 0$. Net customer flow aggregates as $q_t = \beta\, \overline{\mathrm{AE}}_t + \xi_t$, where $\overline{\mathrm{AE}}_t \in \mathbb{R}^{\mathrm{J}}$ stacks cross-analyst means across currencies and $\xi_t \in \mathbb{R}^{\mathrm{J}}$ is the stacked idiosyncratic flow residual (defined below). Where unambiguous, we suppress (i, t) subscripts for readability.

For each currency i and time t, analyst forecasts decompose as

$$\mathrm{AE}^{j}_{i,t} \;=\; \widehat{\mathrm{AE}}_{i,t} \;+\; \mathrm{AE}^{r,j}_{i,t}, \qquad \widehat{\mathrm{AE}}_{i,t} \;=\; \boldsymbol{\theta}^{\top}_{i,t}\mathbf{Z}_t,$$

where $\widehat{\mathrm{AE}}_{i,t}$ is the systematic component and $\mathrm{AE}_{i,t}^{r,j}$ is an idiosyncratic residual with mean zero conditional on $\mathbf{Z}_t$. Let the cross-analyst mean be

$$\overline{\mathrm{AE}}_{i,t} \;=\; \tfrac{1}{\mathrm{N}}\sum_{j=1}^{\mathrm{N}} \mathrm{AE}_{i,t}^{j} \;=\; \widehat{\mathrm{AE}}_{i,t} \;+\; \overline{\mathrm{AE}}_{i,t}^{r}, \qquad \overline{\mathrm{AE}}_{i,t}^{r} \equiv \tfrac{1}{\mathrm{N}}\sum_{j=1}^{\mathrm{N}} \mathrm{AE}_{i,t}^{r,j}.$$

Stacking across currencies yields $\overline{\mathrm{AE}}_t = \widehat{\mathrm{AE}}_t + \overline{\mathrm{AE}}_t^{r} \in \mathbb{R}^{\mathrm{J}}$.

Assumption 1 (Residual aggregation and idiosyncratic flow)**.**

(i) Residual forecasts. *For each* (i,t), $\{\mathrm{AE}_{i,t}^{r,j}\}_{j=1}^{\mathrm{N}}$ *are mean–zero conditional on* $\mathbf{Z}_t$, *uniformly* L^2*–bounded, and weakly dependent across* j *so that a cross–sectional Law of Large Numbers (LLN) or Central Limit Theorem (CLT) holds:* $\mathbb{E}_t[\overline{\mathrm{AE}}_{i,t}^{r}] = 0$ *and* $\mathrm{Var}_t(\overline{\mathrm{AE}}_{i,t}^{r}) = \mathrm{O}(1/\mathrm{N})$, *hence* $\overline{\mathrm{AE}}_{i,t}^{r} \xrightarrow{p} 0$ *as* $\mathrm{N}\to\infty$. *Moreover,* $\mathbb{E}\Big[\overline{\mathrm{AE}}_{i,t}^{r} \mid \mathbf{Z}_t, \widehat{\mathrm{AE}}_t, \Sigma_t, b_t\Big] = 0$.

(ii) Idiosyncratic customer flow. *For each analyst* j, *let* $\tilde{c}_t^{\,j} \in \mathbb{R}^{\mathrm{J}}$ *collect idiosyncratic (non-informational) flows across currencies, and define*

$$\xi_t \;\equiv\; \tfrac{1}{\mathrm{N}}\sum_{j=1}^{\mathrm{N}} \tilde{c}_t^{\,j} \in \mathbb{R}^{\mathrm{J}}.$$

Then $\mathbb{E}_t[\xi_t] = 0$, $\|\xi_t\|_2 = \mathrm{O}_p(\mathrm{N}^{-1/2})$, *and* $\mathbb{E}\Big[\xi_t \mid \mathbf{Z}_t, \widehat{\mathrm{AE}}_t, \Sigma_t, b_t\Big] = 0$.

(iii) Risk structure and non–degeneracy. Σ_t *is symmetric p.s.d., with* $\|\Sigma_t\|_{\mathrm{op}} = \mathrm{O}_p(1)$, $b_t = \mathrm{O}_p(1)$ *(equivalently,* $\|b_t\Sigma_t\|_{\mathrm{op}} = \mathrm{O}_p(1)$*), and* $\mathbb{P}(\|\widehat{\mathrm{AE}}_t\|_1 > 0) = 1$.

A.2 Aggregation and Normalization

Dealers accommodate customer flow one-for-one; their inventory is $w_t \equiv -\,q_t$. Define the dealer's ℓ_1-normalized portfolio as

$$\pi_t \;\equiv\; \frac{w_t}{\|w_t\|_1} \;=\; -\,\frac{q_t}{\|q_t\|_1}.$$

The next lemma collects the large-N consequences of Assumption 1.

Lemma 1 (Aggregation and portfolio composition)**.** *Under Assumption 1,*

$$\operatorname*{plim}_{\mathrm{N}\to\infty} q_t = \beta\,\widehat{\mathrm{AE}}_t, \qquad \operatorname*{plim}_{\mathrm{N}\to\infty} \pi_t = -\,\frac{\widehat{\mathrm{AE}}_t}{\|\widehat{\mathrm{AE}}_t\|_1}.$$

Proof. Write $q_t = \beta(\widehat{\mathrm{AE}}_t + \overline{\mathrm{AE}}_t^r) + \xi_t$. By Assumption 1(i)–(ii), $\overline{\mathrm{AE}}_t^r \xrightarrow{p} 0$ and $\xi_t = \mathrm{O}_p(\mathrm{N}^{-1/2})$, so $q_t \xrightarrow{p} \beta\,\widehat{\mathrm{AE}}_t$. Since $\|\widehat{\mathrm{AE}}_t\|_1 > 0$ a.s. by Assumption 1(iii), the map $v \mapsto v/\|v\|_1$ is continuous on $\{v : \|v\|_1 > 0\}$; with $w_t = -q_t$, the continuous mapping theorem yields

$$\pi_t = \frac{w_t}{\|w_t\|_1} = -\frac{q_t}{\|q_t\|_1} \xrightarrow{p} -\frac{\widehat{\mathrm{AE}}_t}{\|\widehat{\mathrm{AE}}_t\|_1}.$$

□

A.3 Microfoundation of the Pricing Kernel

Dealers have exponential (CARA) utility and portfolio payoffs are jointly Normal. Let rx_{t+1} denote the vector of excess returns and let π_t be the ℓ_1-normalized *dealer* portfolio from Lemma 1 (i.e., $\pi_t = w_t/\|w_t\|_1$ with $w_t = -q_t$). Define $rx_{t+1}^{\pi} \equiv \pi_t^{\top} rx_{t+1}$, with variance $\sigma_{\pi,t}^2 \equiv \mathrm{Var}_t(rx_{t+1}^{\pi}) > 0$. A representative competitive dealer chooses a scalar position y_t in rx_{t+1}^{π} (measured in units of the ℓ_1-normalized dealer portfolio π_t), so the wealth increment from t to $t{+}1$ is $\Delta \mathrm{W}_{t+1} \equiv y_t\, rx_{t+1}^{\pi}$.

Under CARA–Normal structure, maximizing expected utility

$$\max_{y_t}\ \mathbb{E}_t[-\exp(-\gamma_t \Delta \mathrm{W}_{t+1})]$$

is equivalent to mean–variance optimization with baseline risk tolerance $\varphi_t^b \equiv 1/\gamma_t$:

$$\max_{y_t}\ \mathbb{E}_t[\Delta \mathrm{W}_{t+1}] - \tfrac{1}{2\varphi_t^b}\,\mathrm{Var}_t(\Delta \mathrm{W}_{t+1}) \quad \text{s.t.} \quad g_{k,t}(y_t) \le 0, \quad k = 1, \ldots, \mathrm{K},$$

where each $g_{k,t}$ is a convex balance-sheet constraint (e.g. variance budgets, leverage limits, or capital requirements).

The Lagrangian is

$$\mathcal{L}(y_t, \mu_t) = \mathbb{E}_t[y_t\, rx_{t+1}^{\pi}] - \tfrac{1}{2\varphi_t^b} y_t^2 \sigma_{\pi,t}^2 - \sum_{k=1}^{\mathrm{K}} \mu_{k,t}\, g_{k,t}(y_t),$$

with multipliers $\mu_{k,t} \ge 0$. The Karush–Kuhn–Tucker (KKT) first-order condition is

$$\mathbb{E}_t[rx_{t+1}^{\pi}] - \tfrac{1}{\varphi_t^b} y_t\, \sigma_{\pi,t}^2 - \sum_{k=1}^{\mathrm{K}} \mu_{k,t}\, g_{k,t}'(y_t) = 0,$$

together with $\mu_{k,t}\, g_{k,t}(y_t) = 0$, $\mu_{k,t} \geq 0$, and $g_{k,t}(y_t) \leq 0$.

Imposing market clearing $y_t = 1$ (the ℓ_1-normalized dealer portfolio π_t is defined to have unit net supply, so the representative dealer holds one unit) gives

$$\mathbb{E}_t[rx^{\pi}_{t+1}] \;=\; \Big(\frac{1}{\varphi^b_t} + \sum_{k=1}^{\mathrm{K}} \mu_{k,t} \frac{g'_{k,t}(1)}{\sigma^2_{\pi,t}}\Big)\, \sigma^2_{\pi,t} \;\equiv\; b_t\, \sigma^2_{\pi,t}. \tag{27}$$

Hence the effective price of risk is $b_t = 1/\varphi_t$, where

$$\varphi_t \;\equiv\; \frac{1}{\frac{1}{\varphi^b_t} + \sum_{k=1}^{\mathrm{K}} \mu_{k,t} \frac{g'_{k,t}(1)}{\sigma^2_{\pi,t}}} \tag{28}$$

is the reduced-form measure of aggregate risk tolerance. Here $\mu_{k,t} \geq 0$ are the KKT multipliers (shadow prices) on the balance-sheet constraints $g_{k,t}$: when constraints bind ($\mu_{k,t} > 0$), their shadow costs lower φ_t below the baseline φ^b_t, whereas if all constraints are slack ($\mu_{k,t} = 0$) then $\varphi_t = \varphi^b_t$.

Proposition 1 (Pricing kernel under CARA–Normal)**.** *With CARA preferences and Normal payoffs, the representative dealer's stochastic discount factor takes the linear (affine) form*

$$m_{t+1} = a_t - b_t\, rx^{\pi}_{t+1}, \qquad b_t \equiv \tfrac{1}{\varphi_t} > 0,$$

where φ_t is aggregate risk tolerance. Lower φ_t (tighter balance sheets) implies a higher price of risk.

Proof. Equation (27) implies $\mathbb{E}_t[rx^{\pi}_{t+1}] = b_t\, \sigma^2_{\pi,t}$, where $\sigma^2_{\pi,t} \equiv_t (rx^{\pi}_{t+1})$. By the beta relation, for any i,

$$\mathbb{E}_t[rx_{i,t+1}] = \frac{\mathrm{Cov}_t(rx_{i,t+1}, rx^{\pi}_{t+1})}{\sigma^2_{\pi,t}}\, \mathbb{E}_t[rx^{\pi}_{t+1}] = b_t\, \mathrm{Cov}_t(rx_{i,t+1}, rx^{\pi}_{t+1}).$$

Now consider the affine kernel

$$m_{t+1} = a_t - b_t\, rx^{\pi}_{t+1}.$$

Its conditional mean is

$$\mathbb{E}_t[m_{t+1}] \;=\; a_t - b_t\, \mathbb{E}_t[rx^{\pi}_{t+1}],$$

so choose

$$a_t \;=\; 1 + b_t\,\mathbb{E}_t[rx^{\pi}_{t+1}] \quad\Longrightarrow\quad \mathbb{E}_t[m_{t+1}] = 1.$$

Next, for any asset i,

$$\begin{aligned}\mathbb{E}_t[m_{t+1} rx_{i,t+1}] &= \mathbb{E}_t\big[(a_t - b_t rx^{\pi}_{t+1})\, rx_{i,t+1}\big] \\ &= a_t\,\mathbb{E}_t[rx_{i,t+1}] \;-\; b_t\,\mathbb{E}_t[rx^{\pi}_{t+1} rx_{i,t+1}] \\ &= \big(a_t - b_t\,\mathbb{E}_t[rx^{\pi}_{t+1}]\big)\,\mathbb{E}_t[rx_{i,t+1}] \;-\; b_t\,\mathrm{Cov}_t(rx_{i,t+1}, rx^{\pi}_{t+1}),\end{aligned}$$

using $\mathbb{E}_t[\mathrm{XY}] = \mathbb{E}_t[\mathrm{X}]\mathbb{E}_t[\mathrm{Y}] + \mathrm{Cov}_t(\mathrm{X},\mathrm{Y})$. From (27) and the beta relation,

$$\mathbb{E}_t[rx_{i,t+1}] = \frac{\mathrm{Cov}_t(rx_{i,t+1}, rx^{\pi}_{t+1})}{\sigma^2_{\pi,t}}\,\mathbb{E}_t[rx^{\pi}_{t+1}] = b_t\;\mathrm{Cov}_t(rx_{i,t+1}, rx^{\pi}_{t+1}).$$

Substituting, we get

$$\mathbb{E}_t[m_{t+1} rx_{i,t+1}] = \big(a_t - b_t\,\mathbb{E}_t[rx^{\pi}_{t+1}] - 1\big)\; b_t\;\mathrm{Cov}_t(rx_{i,t+1}, rx^{\pi}_{t+1}) \;=\; 0,$$

since $a_t = 1 + b_t\,\mathbb{E}_t[rx^{\pi}_{t+1}]$. Hence m_{t+1} prices all assets and is normalized. Therefore, under CARA–Normal preferences with convex balance-sheet constraints, the representative dealer's SDF is

$$m_{t+1} \;=\; a_t - b_t\, rx^{\pi}_{t+1}, \qquad b_t \equiv \tfrac{1}{\varphi_t} > 0,$$

where φ_t is the reduced-form measure of risk tolerance defined in Equation (28). □

A.4 Pricing and the Central Pricing Equation

Combining the linear SDF derived in Proposition 1 with the aggregation result of Lemma 1 yields the pricing relation. With CARA–Normal structure $m_{t+1} = a_t - b_t\, rx^{\pi}_{t+1}$ and $rx^{\pi}_{t+1} = \pi_t^{\top} rx_{t+1}$, the Euler condition implies

$$\mathbb{E}_t[rx_{t+1}] \;=\; b_t\,\Sigma_t\,\pi_t.$$

Under the dealer-centric normalization in Lemma 1, $\pi_t \xrightarrow{p} -\widehat{\mathrm{AE}}_t/\|\widehat{\mathrm{AE}}_t\|_1$, so the main-text relation (9)—the *Central Pricing Equation (CPE)*—is

$$\mathbb{E}_t[rx_{t+1}] \;=\; -\,b_t\,\Sigma_t\left(\frac{\widehat{\mathrm{AE}}_t}{\|\widehat{\mathrm{AE}}_t\|_1}\right).$$

For convenience we also record a residual-explicit form useful for proofs.

Lemma 2 (Expected returns under residual aggregation)**.** *Under Assumption 1,*

$$\mathbb{E}_t[rx_{t+1}] \;=\; -\,\frac{b_t}{\|\widehat{\mathrm{AE}}_t\|_1}\,\Sigma_t\,\widehat{\mathrm{AE}}_t \;+\; \varepsilon_t, \qquad \varepsilon_t \equiv -\,\frac{b_t}{\|\widehat{\mathrm{AE}}_t\|_1}\,\Sigma_t\,\overline{\mathrm{AE}}^{r}_t, \tag{29}$$

with $\varepsilon_t \xrightarrow{\mathrm{L}^2} 0$ *as* $\mathrm{N}\to\infty$.

Proof. From $\mathbb{E}_t[rx_{t+1}] = b_t\Sigma_t\pi_t$ and Lemma 1, $\pi_t = -\,\widehat{\mathrm{AE}}_t/\|\widehat{\mathrm{AE}}_t\|_1 + o_p(1)$. Substitute and expand:

$$\mathbb{E}_t[rx_{t+1}] = -\,\frac{b_t}{\|\widehat{\mathrm{AE}}_t\|_1}\,\Sigma_t\widehat{\mathrm{AE}}_t - \,\frac{b_t}{\|\widehat{\mathrm{AE}}_t\|_1}\,\Sigma_t\,\overline{\mathrm{AE}}^{r}_t + r_t^{(\mathrm{N})},$$

where $r_t^{(\mathrm{N})} = b_t\Sigma_t\, o_p(1)$. By Assumption 1(i)–(iii), $\overline{\mathrm{AE}}^{r}_t = \mathrm{O}_p(\mathrm{N}^{-1/2})$ and $\|b_t\Sigma_t\|_{\mathrm{op}} = \mathrm{O}_p(1)$, so the second term and $r_t^{(\mathrm{N})}$ vanish in L^2 as $\mathrm{N}\to\infty$. Define ε_t accordingly to obtain (29). □

A.5 Main Results and Proofs

Proposition 2 (Determinants of currency risk premia (CPE))**.** *In large analyst panels,*

$$\mathbb{E}_t[rx_{t+1}] \;=\; -\,b_t\,\Sigma_t\left(\frac{\widehat{\mathrm{AE}}_t}{\|\widehat{\mathrm{AE}}_t\|_1}\right).$$

Expected excess returns scale with $b_t = 1/\varphi_t$ *and with each currency's exposure* $e_i^\top\Sigma_t\widehat{\mathrm{AE}}_t$*, normalized by* $\|\widehat{\mathrm{AE}}_t\|_1$ *(the overall sign reflects the dealer–customer market-clearing convention).*

Proof. Immediate from Lemma 1 and $\mathbb{E}_t[rx_{t+1}] = b_t\Sigma_t\pi_t$ with $\pi_t \to -\,\widehat{\mathrm{AE}}_t/\|\widehat{\mathrm{AE}}_t\|_1$. □

Corollary 1 (Residual expectations are unpriced)**.** *Under Assumption 1, any zero-cost portfolio formed solely from* $\{\mathrm{AE}^{r,j}_{i,t}\}_j$ *has zero expected premium (up to sampling error).*

Proof. By Assumption 1, $\overline{\mathrm{AE}}^{r}_t \xrightarrow{p} 0$ and $\xi_t = \mathrm{O}_p(\mathrm{N}^{-1/2})$ while $\beta\,\widehat{\mathrm{AE}}_t = \mathrm{O}_p(1)$, hence $\pi_t \xrightarrow{p} -\,\widehat{\mathrm{AE}}_t/\|\widehat{\mathrm{AE}}_t\|_1$. Portfolios measurable w.r.t. $\{\mathrm{AE}^{r,j}_{i,t}\}_j$ are orthogonal to the priced component, so their expected premium is zero by (29). □

Proposition 3 (Cross-sectional predictability from systematic expectations)**.** *At a given* t*, cross-sectional variation in* $\mathbb{E}_t[rx_{i,t+1}]$ *is driven by* $e_i^\top\Sigma_t\widehat{\mathrm{AE}}_t$*. If* Σ_t *is diagonal (or* $\widehat{\mathrm{AE}}_t$

aligns with an eigenvector of Σ_t), then $\mathbb{E}_t[rx_{i,t+1}] \propto -\widehat{\mathrm{AE}}_{i,t}$ with $\sigma^2_{i,t} \equiv e_i^\top \Sigma_t e_i > 0$, so ranking by $-\widehat{\mathrm{AE}}_{i,t}$ coincides with ranking by $\mathbb{E}_t[rx_{i,t+1}]$ under the USD-per-foreign-currency convention.

Proof. From Lemma 2, cross-sectional variation is proportional (up to a common negative factor) to $e_i^\top \Sigma_t \widehat{\mathrm{AE}}_t$. The diagonal/eigenvector cases yield proportionality to $-\widehat{\mathrm{AE}}_{i,t}$. □

Proposition 4 (Factor premium for a –AE-mimicking portfolio)**.** *Let $\overline{\widehat{\mathrm{AE}}}_t \equiv \frac{1}{\mathrm{J}} \sum_{i=1}^{\mathrm{J}} \widehat{\mathrm{AE}}_{i,t}$ and $z_{i,t} \equiv \widehat{\mathrm{AE}}_{i,t} - \overline{\widehat{\mathrm{AE}}}_t$. Define zero-cost, unit-ℓ_1 weights $w_{i,t}^{-\mathrm{AE}} \equiv \frac{-z_{i,t}}{\sum_{j=1}^{\mathrm{J}} |z_{j,t}|}$ and the portfolio return $rx_{t+1}^{-\mathrm{AE}} \equiv (w_t^{-\mathrm{AE}})^\top rx_{t+1}$. Then*

$$\mathbb{E}_t[rx_{t+1}^{-\mathrm{AE}}] = \frac{b_t}{\|\widehat{\mathrm{AE}}_t\|_1} (w_t^{-\mathrm{AE}})^\top \Sigma_t \widehat{\mathrm{AE}}_t \; > \; 0,$$

so this factor earns a positive premium that increases with $b_t = 1/\varphi_t$.

Proof. Take expectations under Lemma 2 and note $w_t^{-\mathrm{AE}}$ aligns with $-\widehat{\mathrm{AE}}_t$; since Σ_t is p.s.d., the stated inner product is positive. □

Proposition 5 (Horizon and state dependence)**.** *If $(\widehat{\mathrm{AE}}_t, \Sigma_t)$ (and/or b_t) are persistent and idiosyncratic components are weakly serially correlated, then for $h \geq 1$,*

$$\mathbb{E}_t\Big[\sum_{\tau=1}^{h} rx_{i,t+\tau}\Big] = -\sum_{\tau=0}^{h-1} \frac{b_{t+\tau}}{\|\widehat{\mathrm{AE}}_{t+\tau}\|_1} e_i^\top \Sigma_{t+\tau} \widehat{\mathrm{AE}}_{t+\tau},$$

so the signal accumulates at least linearly in h while noise grows at most at rate $\sqrt{h}$, implying a weakly increasing h-horizon signal-to-noise ratio. Since $b_t = 1/\varphi_t$, tighter capacity (higher b_t) amplifies short-horizon predictability.

Proof. Iterate expectations and apply standard variance bounds under weak dependence. □

A.6 Measurement Error in Survey Means and Attenuation Bias

This section formalizes the errors-in-variables problem that arises when using survey averages of a small panel of analysts, as discussed informally in Section 2.8 of the main text. We show (i) how survey-reported means embed finite-M sampling noise; (ii) how this generates attenuation bias in predictive regressions; (iii) how the severity of attenuation depends on forecast dispersion and horizon; and (iv) how projection on state variables provides an instrumental-variables correction.

Survey-reported averages. Let $\mathrm{M} \ll \mathrm{N}$ denote the number of analysts surveyed. The reported mean is

$$\widetilde{\mathrm{AE}}_{i,t} \;\equiv\; \frac{1}{\mathrm{M}} \sum_{j=1}^{\mathrm{M}} \mathrm{AE}_{i,t}^{j} \;=\; \overline{\mathrm{AE}}_{i,t} + u_{i,t},$$

where $u_{i,t} \equiv \frac{1}{\mathrm{M}} \sum_{j=1}^{\mathrm{M}} \mathrm{AE}_{i,t}^{r,j}$ captures sampling error and residual idiosyncratic noise. By construction, $\mathbb{E}[u_{i,t} \mid \mathbf{Z}_t] = 0$, $\mathrm{Var}(u_{i,t}) = \mathrm{O}(1/\mathrm{M})$, and $u_{i,t}$ is orthogonal to $\widehat{\mathrm{AE}}_{i,t}$ conditional on $\mathbf{Z}_t$.

Lemma 3 (Variance of survey noise)**.** *Under Assumption 1, the conditional variance of the survey noise satisfies*

$$\mathrm{Var}(u_{i,t} \mid \mathbf{Z}_t) = \frac{1}{\mathrm{M}^2}\, \mathbf{1}^{\top} \mathrm{Cov}\big((\mathrm{AE}_{i,t}^{r,1}, \ldots, \mathrm{AE}_{i,t}^{r,\mathrm{M}})^{\top} \mid \mathbf{Z}_t\big)\, \mathbf{1} = \mathrm{O}\Big(\tfrac{1}{\mathrm{M}}\Big).$$

Proof. Since $u_{i,t} = \frac{1}{\mathrm{M}} \sum_{j=1}^{\mathrm{M}} \mathrm{AE}_{i,t}^{r,j}$ with mean zero conditional on $\mathbf{Z}_t$, weak cross-analyst dependence implies the sum of covariances is $\mathrm{O}(\mathrm{M})$; dividing by M^2 yields $\mathrm{O}(1/\mathrm{M})$ under uniform L^2-boundedness. □

Proposition 6 (Attenuation bias in predictive regressions)**.** *Consider the predictive regression*

$$rx_{i,t+1} = \alpha_i + \beta_i^{(rx)}\, \widehat{\mathrm{AE}}_{i,t} + \varepsilon_{i,t+1},$$

with slope $\beta_i^{(rx)}$ pinned down by the dealer-centric Central Pricing Equation (9). Replacing $\widehat{\mathrm{AE}}_{i,t}$ with the survey proxy $\widetilde{\mathrm{AE}}_{i,t}$ yields an errors-in-variables estimator

$$\mathrm{plim}\, \hat{\beta}_i^{(rx)} = \beta_i^{(rx)} \cdot \lambda_{\mathrm{M}}, \qquad \lambda_{\mathrm{M}} \equiv \frac{1}{1 + \dfrac{\mathrm{Var}(u_{i,t})}{\mathrm{Var}(\widehat{\mathrm{AE}}_{i,t})}} \in (0,1).$$

Proof. Standard errors-in-variables algebra applies since $u_{i,t}$ is mean-zero and orthogonal to $\widehat{\mathrm{AE}}_{i,t}$ (conditional on $\mathbf{Z}_t$). The attenuation factor follows by direct computation. □

Corollary 2 (Effect of dispersion and panel size)**.** *Attenuation is most severe when (i) forecast dispersion is high, inflating* $\mathrm{Var}(u_{i,t})$*, or (ii) the number of respondents* M *is small, so that averaging reduces residual variance only at rate* $1/\mathrm{M}$*. In these cases,* λ_{M} *is close to zero and slopes estimated on survey means are strongly biased toward zero.*

Proof of Corollary 2. From Proposition 6,

$$\lambda_{\mathrm{M}} \;=\; \Big(1 + \frac{\mathrm{Var}(u_{i,t})}{\mathrm{Var}(\widehat{\mathrm{AE}}_{i,t})}\Big)^{-1}.$$

By Lemma 3, $\mathrm{Var}(u_{i,t}) = \mathrm{O}(1/\mathrm{M})$, with the constant of proportionality increasing in forecast dispersion through

$$\mathrm{Var}(u_{i,t} \mid \mathbf{Z}_t) \;=\; \tfrac{1}{\mathrm{M}^2}\,\mathbf{1}^\top \mathrm{Cov}\big((\mathrm{AE}_{i,t}^{r,1}, \ldots, \mathrm{AE}_{i,t}^{r,\mathrm{M}})^\top \mid \mathbf{Z}_t\big)\,\mathbf{1}.$$

Thus higher residual dispersion raises $\mathrm{Var}(u_{i,t})$, lowering λ_{M}, while smaller M slows the $1/\mathrm{M}$ variance reduction, again lowering λ_{M}. In both cases attenuation is more severe, with $\lambda_{\mathrm{M}} \to 0$ as dispersion grows and $\lambda_{\mathrm{M}} \to 1$ only as $\mathrm{M} \to \infty$. □

Proposition 7 (Horizon dependence of attenuation)**.** *At horizon* h*, the attenuation factor is*

$$\lambda_{\mathrm{M},h} \;=\; \frac{1}{1 + \dfrac{\mathrm{Var}(u_{i,t})}{\mathrm{Var}(\widehat{\mathrm{AE}}_{i,t|t+h})}} \;\in (0,1).$$

While $\mathrm{Var}(u_{i,t}) = \mathrm{O}(1/\mathrm{M})$ *is essentially horizon-invariant,* $\mathrm{Var}(\widehat{\mathrm{AE}}_{i,t|t+h})$ *is weakly increasing in* h *under persistence of* $(\widehat{\mathrm{AE}}_t, \Sigma_t)$ *and weak dependence of residuals. Hence* $\lambda_{\mathrm{M},h}$ *is increasing in* h*, so attenuation bias is weaker at longer horizons. This matches Proposition 5, where systematic signals accumulate at least linearly in* h *while noise grows at most at rate* $\sqrt{h}$*.*

Proof. Immediate from Lemma 3 and Proposition 5. □

Lemma 4 (Instrumental variables recovery of systematic expectations)**.** *Let* $\widetilde{\mathrm{AE}}_{i,t} = \widehat{\mathrm{AE}}_{i,t} + u_{i,t}$ *with* $\mathbb{E}[u_{i,t} \mid \mathbf{Z}_t] = 0$*. Since* $\widehat{\mathrm{AE}}_{i,t} = \boldsymbol{\theta}_{i,t}^\top \mathbf{Z}_t$*, the linear projection of* $\widetilde{\mathrm{AE}}_{i,t}$ *onto* $\mathbf{Z}_t$ *recovers* $\widehat{\mathrm{AE}}_{i,t}$*:*

$$\mathrm{Proj}\big(\widetilde{\mathrm{AE}}_{i,t} \mid \mathbf{Z}_t\big) = \widehat{\mathrm{AE}}_{i,t}.$$

Hence instrumenting $\widetilde{\mathrm{AE}}_{i,t}$ *with* $\mathbf{Z}_t$ *in two-stage least squares eliminates measurement error and restores the predictive regression slope implied by the Central Pricing Equation* (9).

Proof. Decompose $\widetilde{\mathrm{AE}}_{i,t} = \widehat{\mathrm{AE}}_{i,t} + u_{i,t}$ with $u_{i,t}$ mean-zero and orthogonal to $\mathbf{Z}_t$. Since $\widehat{\mathrm{AE}}_{i,t}$ is linear in $\mathbf{Z}_t$, the projection operator yields $\widehat{\mathrm{AE}}_{i,t}$. Instrumenting $\widetilde{\mathrm{AE}}_{i,t}$ with $\mathbf{Z}_t$ therefore eliminates bias. □

Proposition 8 (Monotone long-horizon attenuation)**.** *Suppose the state vector follows a stable VAR(1),*

$$\mathbf{Z}_{t+1} = \mathrm{A}\,\mathbf{Z}_t + \varepsilon_{t+1}, \qquad \mathbb{E}[\varepsilon_{t+1} \mid \mathcal{F}_t] = 0, \quad \mathrm{Cov}(\varepsilon_{t+1}) = \Omega, \quad \rho(\mathrm{A}) < 1,$$

and the one–period systematic expectation for currency i *is linear in the state,*

$$\widehat{\mathrm{AE}}_{i,t|t+1} = \mathrm{B}_i^\top \mathbf{Z}_t,$$

so that the h*–horizon systematic expectation is*

$$\widehat{\mathrm{AE}}_{i,t|t+h} = \sum_{\tau=0}^{h-1} \mathrm{B}_i^\top \mathrm{A}^\tau \mathbf{Z}_t \;=\; \Big(\sum_{\tau=0}^{h-1} \mathrm{A}^{\tau\top} \mathrm{B}_i\Big)^\top \mathbf{Z}_t.$$

Let $\Gamma \equiv \mathrm{Var}(\mathbf{Z}_t)$ *(well-defined under* $\rho(\mathrm{A}) < 1$*). Then:*

1. $\mathrm{Var}(\widehat{\mathrm{AE}}_{i,t|t+h})$ *is weakly increasing in* h*; more precisely,*

$$\mathrm{Var}(\widehat{\mathrm{AE}}_{i,t|t+h}) = \Big(\sum_{\tau=0}^{h-1} \mathrm{A}^{\tau\top} \mathrm{B}_i\Big)^\top \Gamma \Big(\sum_{\tau=0}^{h-1} \mathrm{A}^{\tau\top} \mathrm{B}_i\Big),$$

   and the sequence of Gram matrices $\mathrm{K}_h \equiv \sum_{\tau=0}^{h-1} \mathrm{A}^{\tau\top} \mathrm{B}_i \mathrm{B}_i^\top \mathrm{A}^\tau$ *satisfies* $\mathrm{K}_{h+1} = \mathrm{K}_h + \mathrm{A}^{h\top} \mathrm{B}_i \mathrm{B}_i^\top \mathrm{A}^h \succeq \mathrm{K}_h$ *(Loewner order).*

2. *With survey noise* $u_{i,t}$ *satisfying* $\mathrm{Var}(u_{i,t}) = \sigma_u^2 = \mathrm{O}(1/\mathrm{M})$ *and independent of* $\mathbf{Z}_t$*, the attenuation factor*

$$\lambda_{\mathrm{M},h} \;=\; \frac{1}{1 + \dfrac{\sigma_u^2}{\mathrm{Var}(\widehat{\mathrm{AE}}_{i,t|t+h})}}$$

   is weakly increasing in h*.*

Proof. For (i), write $\widehat{\mathrm{AE}}_{i,t|t+h} = \mathrm{C}_h^\top \mathbf{Z}_t$ with $\mathrm{C}_h \equiv \sum_{\tau=0}^{h-1} \mathrm{A}^{\tau\top}\mathrm{B}_i$. Then

$$\mathrm{Var}(\widehat{\mathrm{AE}}_{i,t|t+h}) = \mathrm{C}_h^\top \Gamma \mathrm{C}_h = \mathrm{tr}\big(\Gamma^{1/2}\mathrm{C}_h\mathrm{C}_h^\top\Gamma^{1/2}\big).$$

Since $\mathrm{C}_{h+1}\mathrm{C}_{h+1}^\top = \mathrm{C}_h\mathrm{C}_h^\top + \mathrm{A}^{h\top}\mathrm{B}_i\mathrm{B}_i^\top\mathrm{A}^h \succeq \mathrm{C}_h\mathrm{C}_h^\top$, we have $\mathrm{C}_{h+1}^\top\Gamma\mathrm{C}_{h+1} \geq \mathrm{C}_h^\top\Gamma\mathrm{C}_h$ because $\Gamma \succeq 0$. Thus $h \mapsto \mathrm{Var}(\widehat{\mathrm{AE}}_{i,t|t+h})$ is weakly increasing.

For (ii), define $g(x) \equiv 1/(1+\sigma_u^2/x)$ on $x > 0$. Then $g'(x) = \sigma_u^2/(x+\sigma_u^2)^2 > 0$, so g is strictly increasing. Since $\mathrm{Var}(\widehat{\mathrm{AE}}_{i,t|t+h})$ is weakly increasing in h by (i), it follows that $\lambda_{\mathrm{M},h} = g(\mathrm{Var}(\widehat{\mathrm{AE}}_{i,t|t+h}))$ is weakly increasing in h.

Economically, the increase in $\mathrm{Var}(\widehat{\mathrm{AE}}_{i,t|t+h})$ reflects the accumulation of persistent systematic signals over the horizon, consistent with Proposition 5; the survey noise variance σ_u^2 is horizon–invariant, so the signal–to–noise ratio rises with h, reducing measurement–error attenuation. □

A.7 Survey-Based Portfolio Construction: Attenuation and IV Recovery

We formalize the portfolio counterparts of the regression results in Section 2.8. Under the dealer-centric CPE, the *positively priced* factor loads on the *opposite* of systematic expectations. Let $r_{t+1}^{-\mathrm{AE}}$ denote the excess return on a –AE–mimicking portfolio formed by sorting currencies on *negative* expected appreciation and normalizing to unit ℓ_1 exposure.[27]

Assumption 2 (Asymptotic linearity of the factor signal)**.** *There exists a (random) loading vector θ_t with $\|\theta_t\|_2 = \mathrm{O}_p(1)$ such that*

$$\mathbb{E}_t[r_{t+1}^{-\mathrm{AE}}] \;=\; \theta_t^\top\Big(\frac{-\widehat{\mathrm{AE}}_t}{\|\widehat{\mathrm{AE}}_t\|_1}\Big), \qquad \mathrm{Var}_t(r_{t+1}^{-\mathrm{AE}}) \;=\; \sigma_{r,t}^2 = \mathrm{O}_p(1).$$

Assumption 2 holds for standard long–short portfolios sorted on $-\widehat{\mathrm{AE}}_{i,t}$ under the dealer-centric CPE and ℓ_1 normalization (cf. Proposition 4).

Proposition 9 (Attenuation of survey-based portfolio premia)**.** *Let $\widetilde{\mathrm{AE}}_{i,t} = \widehat{\mathrm{AE}}_{i,t} + u_{i,t}$ as in Section A.6, with $u_{i,t}$ mean-zero and orthogonal to $\mathbf{Z}_t$. Consider the portfolio formed*

[27] For example, take cross-sectional ranks of $-\widehat{\mathrm{AE}}_{i,t}$, form long–short $w_{i,t}^{-\mathrm{AE}}$ with $\sum_i |w_{i,t}^{-\mathrm{AE}}| = 1$, and set $r_{t+1}^{-\mathrm{AE}} \equiv (w_t^{-\mathrm{AE}})^\top r x_{t+1}$. Any construction that is a measurable functional of $-\widehat{\mathrm{AE}}_t$ and is asymptotically linear in the signal (e.g., quantile sorts with a large cross-section) fits the results below.

by sorting on the survey mean with the dealer-consistent sign, yielding return $r_{t+1}^{-\mathrm{AE},\ survey}$ *(i.e., sort on* $-\widetilde{\mathrm{AE}}_{i,t}$*). Under Assumption 2,*

$$\mathbb{E}_t\big[r_{t+1}^{-\mathrm{AE},\ survey}\big] \;=\; \lambda_{\mathrm{M}}\,\theta_t^{\top}\Big(\frac{-\,\widehat{\mathrm{AE}}_t}{\|\widehat{\mathrm{AE}}_t\|_1}\Big), \qquad \lambda_{\mathrm{M}} \;=\; \frac{1}{1+\dfrac{\mathrm{Var}(u_{i,t})}{\mathrm{Var}(\widehat{\mathrm{AE}}_{i,t})}} \in (0,1),$$

so expected premia are attenuated by the same factor λ_{M} *as in Proposition 6.*

Proof. By asymptotic linearity, the sorting functional is equivalent to applying $\theta_t^{\top}(\cdot)$ to the (signed) signal used to sort. Replacing $-\widehat{\mathrm{AE}}_t$ by $-\widetilde{\mathrm{AE}}_t$ introduces classical measurement error, and the conditional mean contracts by λ_{M} as in Proposition 6. □

Corollary 3 (Sharpe ratio attenuation)**.** *Suppose* $\mathrm{Var}_t(r_{t+1}^{-\mathrm{AE},\ survey}) = \sigma_{r,t}^2 + \mathrm{O}_p(1/\mathrm{M})$*, i.e., replacing* $-\widehat{\mathrm{AE}}_t$ *by* $-\widetilde{\mathrm{AE}}_t$ *alters portfolio variance only at second order* $\mathrm{O}(1/\mathrm{M})$ *because sampling noise averages out across currencies (cf. Lemma 3). Then*

$$\mathrm{SR}_t^{survey} \;\equiv\; \frac{\mathbb{E}_t[r_{t+1}^{-\mathrm{AE},\ survey}]}{\sigma_{r,t}} \;\approx\; \lambda_{\mathrm{M}}\cdot\frac{\theta_t^{\top}\big(-\,\widehat{\mathrm{AE}}_t/\|\widehat{\mathrm{AE}}_t\|_1\big)}{\sigma_{r,t}} \;=\; \lambda_{\mathrm{M}}\cdot\mathrm{SR}_t^{*} \;\le\; \mathrm{SR}_t^{*},$$

where $\mathrm{SR}_t^{*} \equiv \theta_t^{\top}\big(-\,\widehat{\mathrm{AE}}_t/\|\widehat{\mathrm{AE}}_t\|_1\big)/\sigma_{r,t}$ *and* $\approx$ *indicates that variance differences are negligible relative to mean effects. Thus the Sharpe ratio is attenuated by approximately the same factor* λ_{M}*.*

Proof. Immediate from Proposition 9 and the variance invariance condition. □

Lemma 5 (IV recovery for portfolio construction)**.** *Let* $\widehat{\mathrm{AE}}_{i,t} = \mathrm{B}_i^{\top}\mathbf{Z}_t$ *and* $\widetilde{\mathrm{AE}}_{i,t} = \widehat{\mathrm{AE}}_{i,t} + u_{i,t}$ *with* $\mathbb{E}[u_{i,t} \mid \mathbf{Z}_t] = 0$*. Define fitted values* $\widehat{\mathrm{AE}}_{i,t}^{\mathrm{IV}} \equiv \mathrm{Proj}(\widetilde{\mathrm{AE}}_{i,t} \mid \mathbf{Z}_t)$*. Then* $\widehat{\mathrm{AE}}_{i,t}^{\mathrm{IV}} = \widehat{\mathrm{AE}}_{i,t}$ *and the portfolio formed by sorting on* $-\widehat{\mathrm{AE}}_{i,t}^{\mathrm{IV}}$ *attains*

$$\mathbb{E}_t\big[r_{t+1}^{-\mathrm{AE},\ IV}\big] \;=\; \theta_t^{\top}\Big(\frac{-\,\widehat{\mathrm{AE}}_t}{\|\widehat{\mathrm{AE}}_t\|_1}\Big), \qquad \mathrm{SR}_t^{IV} \;=\; \mathrm{SR}_t^{*} + o_p(1).$$

Proof. Same projection argument as Lemma 4 recovers the systematic component. Applying Assumption 2 to the signed signal yields the stated mean; variance equality follows from the $\mathrm{O}(1/\mathrm{M})$ perturbation condition in Corollary 3. □

B Definition of Currency Value

We define the variable *value* to reflect the currency' relatively cross-sectional cheapness or expensiveness (Chernov, Dahlquist, and Lochstoer, 2023). The definition of *value* starts with a currency's real exchange rate (RER). We calculate the RER (Q^i) by taking the product of spot rate (S_t) and foreign consumer price index (CPI) divided by U.S. CPI. We include an additional lag so as to account for CPI publication delay, which is given by

$$\mathrm{Q}_t^i = \mathrm{S}_t^i \cdot \mathrm{P}_t^i / \mathrm{P}_t,$$

where S_i, P_t and P_t^i are the spot rate of currency i, U.S. and foreign consumer price index (CPI), respectively.

Next, according to Jordà and Taylor (2012), the long-run mean of RER is not able to be empirically clear defined. Thus, we scale the RER by its five-year smoothed lag (specifically, the average RER from 4.5 to 5.5 years ago) to avoid its long-term reversal trend implied by Purchasing Power Parity,

$$\tilde{\mathrm{Q}}_t^i \equiv \mathrm{Q}_t^i \cdot \left(\frac{1}{13} \sum_{j=-6}^{6} \mathrm{Q}_{t-60+j}^i \right)^{-1}.$$

Lastly, we demean the $\tilde{\mathrm{Q}}_t^i$ cross-sectionally at each time t to reflect each currency's relative "valuation", which is given by

$$\mathrm{V}alue^i = \tilde{\mathrm{Q}}_t^i - \frac{1}{\mathrm{N}} \sum_{i=1}^{\mathrm{N}} \tilde{\mathrm{Q}}_t^i,$$

C Estimation from Elastic-Net Regression

We estimate the ENet-based appreciation expectation ($\text{-}\widehat{\mathrm{AE}}^{\mathrm{EN}et}$) from the Elastic-Net (ENet) regressions and detail our procedure in this subsection. Compared to the common variable selection method (e.g. LASSO, RIDGE), the ENet model allows us to select predictors in a small cross section.

Predictor Candidates. We consider a set of macroeconomic and financial predictors: the momentum (Mom), the value (Value), the forward discount (FD), the illiquidity (ILL), the volatility (Vol), dividend price ratio (DP), the price-to-earnings ratio (PE),

the inflation rate (INF), the unemployment rate (UN), and the market capitalization (MKTCAP).

ENet Estimation. At the end of each month t, we run the predictive ENet regressions as follows. In Table 13, we run a predictive cross-sectional ENet regression, while in Table I.A. 6, we estimate the coefficients in ENet regressions from a past 12-month rolling window. The ENet objective function can be expressed as

$$\arg\min_{\mathbf{w}} \frac{1}{2n}\|\mathrm{X}w - y\|_2^2 + \alpha \mathrm{P}_\delta\left(\mathbf{w}\right), \tag{B.1}$$

where

$$\mathrm{P}_\delta\left(\mathbf{w}\right) = \alpha\left[\frac{(1-\delta)}{2}\|\mathbf{w}\|_2^2 + \delta\|\mathbf{w}\|_1\right]; \tag{B.2}$$

where $\alpha \geq 0$ represents a hyperparameter that captures the degree of shrinkage; n is the number of observations; $\|\cdot\|_1$ and $\|\cdot\|_2$ are the ℓ_1 and ℓ_2 norms, respectively; $0 \leq \delta \leq 1$ is a hyperparameter for combining the ℓ_1 and ℓ_2 components in the penalty term. $\mathbf{X}_{t-1}$ includes 10 predictor candidates, and all macroeconomic variables account for publication bias by adding an extra lag. The ENet objective function reduces to that for OLS for $\lambda = 0$. If $\delta = 1$, then B.1 corresponds to the LASSO objective function, and $\delta = 0$ corresponds to ridge regression. Following the literature, we set δ equal to 0.5. At the end of each month t, we search the optimal α based on the Bayesian Information Criterion (BIC) after setting the number of minimum predictors selected as three. The α ranges from 0 to 3 with step equal to 1 basis point. $||\mathbf{w}||$ and $||\mathbf{w}||^2$ are the first and second-norm penalty terms, respectively. Additionally, we apply equal weights between two penalties.

D Additional Results

Table I.A. 1 Single Portfolio Sorting on Analyst Forecasts: 24-month Forecast Horizon

The table presents the single-sorting quintile annualized excess returns on the risk-based analyst forecast ($-\widehat{\text{AE}}$) and analyst expectation ($-$AE), both of which are based on the forecast horizon of 24 months. The rows show analyst forecast horizons (f), whereas the columns indicate holding periods (h). The left part of the table shows the results of risk-based analysts' expectations, whereas the right part shows ones for analysts' expectations. In Panel A, we report the sorting results without considering the transaction costs and in Panel B (Panel C) we proxy transaction costs for the effective costs of 25% (50%). At the end of each month t, we sort currencies based on (risk-based) analysts' expectations into quintiles and calculate the average excess return over the holding periods. We report the time-series average annualized returns of the high-minus-low quintile. Newey and West (1987) adjusted t-statistics are reported in parentheses. Our sample ranges from December 1994 to October 2016.

Forecast Horizon f	Holding Period h							
	$-\widehat{\text{AE}}$				$-$AE			
	1	3	6	12	1	3	6	12
	Panel A: Without Transaction Costs							
24	4.20	3.22	3.49	2.40	1.75	1.73	1.79	1.80
	(2.56)	(2.70)	(3.63)	(3.15)	(1.09)	(1.27)	(1.59)	(1.97)
	Panel B: 25% Quote Spread							
24	3.18	2.80	3.23	2.22	0.87	1.39	1.58	1.65
	(1.93)	(2.35)	(3.36)	(2.92)	(0.54)	(1.02)	(1.40)	(1.81)
	Panel C: 50% Quote Spread							
24	2.16	2.38	2.97	2.05	0.00	1.05	1.38	1.51
	(1.31)	(2.00)	(3.09)	(2.68)	(0.00)	(0.77)	(1.22)	(1.65)

Table I.A. 2 Single Portfolio Sorting on Analyst Forecasts: Publication Bias

The table presents the single-sorting quintile annualized excess returns on the risk-based analyst forecast ($-\widehat{\text{AE}}$) and analyst expectation ($-$AE), both of which are with one extra lag for the pubilcation bias. The rows show analyst forecast horizons (f), whereas the columns indicate holding periods (h). The left part of the table shows the results of risk-based analysts' expectations, whereas the right part shows ones of analysts' expectations. In Panel A, we report the sorting results without considering the transaction costs and in Panel B (Panel C) we proxy transaction costs for the effective costs of 25% (50%). At the end of each month t, we sort currencies based on (risk-based) analysts' expectations into quintiles and calculate the average excess return over the holding periods. We report the time-series average annualized returns of the high-minus-low quintile. Newey and West (1987) adjusted t-statistics are reported in parentheses. Our sample ranges from March 1990 to October 2016.

Forecast Horizon f	Holding Period h							
	$-\widehat{\text{AE}}$				$-$AE			
	1	3	6	12	1	3	6	12
	Panel A: Without Transaction Costs							
3	2.46	2.86	3.41	2.55	1.79	1.98	2.07	1.67
	(1.48)	(2.21)	(3.04)	(3.17)	(1.12)	(1.62)	(1.99)	(2.12)
12	3.41	3.21	3.61	2.71	1.84	1.64	1.30	1.66
	(2.10)	(2.43)	(3.27)	(3.48)	(1.12)	(1.27)	(1.25)	(2.11)
	Panel B: 25% Quote Spread							
3	1.52	2.48	3.17	2.39	1.05	1.70	1.91	1.55
	(0.92)	(1.91)	(2.83)	(2.97)	(0.66)	(1.39)	(1.83)	(1.97)
12	2.43	2.82	3.36	2.54	1.08	1.35	1.12	1.53
	(1.50)	(2.14)	(3.05)	(3.27)	(0.66)	(1.04)	(1.08)	(1.94)
	Panel C: 50% Quote Spread							
3	0.58	2.10	2.93	2.23	0.33	1.42	1.74	1.43
	(0.35)	(1.62)	(2.62)	(2.77)	(0.21)	(1.16)	(1.66)	(1.81)
12	1.48	2.43	3.12	2.38	0.33	1.06	0.94	1.40
	(0.91)	(1.85)	(2.83)	(3.05)	(0.20)	(0.82)	(0.90)	(1.78)

Table I.A. 3 Double-Sorted Risk-Based Analyst Forecast Portfolio: conditional on Volatility or Illiquidity

The table presents annualized average conditional excess returns for double-sorted portfolios. We first allocate all currencies into 2 portfolios on Volatility (Panel A) or Illiquidity (Panel B). Then, within each of the 2 portfolios, currencies are sorted into 3 portfolios based on the risk-based analyst forecasts ($-\widehat{\text{AE}}$). The rows show analyst forecast horizons (f), whereas the columns indicate holding periods (h). We report the time-series average annualized returns of the high-minus-low portfolios. Newey and West (1987) adjusted t-statistics are reported in parentheses. Our sample ranges from February 1990 to October 2016.

Port.	Forecast Horizon f	Holding Period h			
		1	3	6	12
		Panel A: Volatility			
Low VOL	3	4.10	3.34	3.15	2.59
		(2.72)	(3.41)	(3.95)	(4.34)
	12	3.11	2.86	3.18	2.64
		(2.17)	(3.13)	(4.04)	(4.63)
High VOL	3	0.86	1.59	2.12	1.48
		(0.55)	(1.54)	(2.51)	(2.26)
	12	1.50	1.90	2.26	1.40
		(0.95)	(1.88)	(2.66)	(2.14)
		Panel B: Illiquidity			
Low ILLIQ	3	0.91	1.22	1.94	1.81
		(0.60)	(1.13)	(2.24)	(2.85)
	12	0.85	1.52	2.29	1.85
		(0.57)	(1.42)	(2.57)	(2.89)
High ILLIQ	3	2.92	2.04	2.35	1.84
		(1.96)	(1.88)	(2.72)	(2.94)
	12	3.63	2.81	2.75	1.99
		(2.64)	(2.75)	(3.28)	(3.30)

Table I.A. 4 Double-Sorted Risk-Based Analyst Forecast Portfolio: conditional on Uncertainty Measure

The table presents annualized average conditional excess returns for double-sorted portfolios. We first allocate all currencies into 2 portfolios on the exposure to Baker, Bloom, and Davis (2016)'s Economic Policy Uncertainty, Baker, Bloom, and Davis (2016)'s Monetary Policy Uncertainty, Caldara and Iacoviello (2022)'s Geopolitical Risk, and Püttmann (2018)'s Financial Stress Indicator, aggregated Hassan, Hollander, Van Lent, and Tahoun (2019)'s Firm-level Risk measure in Panel A, Panel B, Panel C, Panel D, and Panel E, respectively. For currency global risk exposure, at each end of month, we run the predictive regressions of currency excess returns on the global risk measure in each panel over the past 60 months (at least 36 required), from which we take the slopes as the global risk exposure. Then within each of the 2 portfolios, currencies are sorted into 3 portfolios based on the risk-based analyst forecasts $(-\widehat{\text{AE}})$. The rows show analyst forecast horizons (f), whereas the columns indicate holding periods (h). We report the time-series average annualized returns of the high-minus-low portfolios. Newey and West (1987) adjusted t-statistics are reported in parentheses. For Panel A, Panel B, Panel C, Panel D, our sample ranges from February 1990 to October 2016 and for Panel D our sample ranges from January 2012 to October 2016.

Forecast Horizon f	Holding Period h							
	Low Betas				High Betas			
	1	3	6	12	1	3	6	12
	Panel A: Economic Policy Uncertainty							
3	3.32	2.05	2.39	1.03	1.53	0.80	1.21	1.25
	(2.12)	(2.03)	(3.49)	(1.70)	(0.98)	(0.66)	(1.18)	(1.79)
12	2.18	1.94	2.32	1.04	2.07	1.35	1.40	1.30
	(1.44)	(2.00)	(3.56)	(1.83)	(1.39)	(1.17)	(1.45)	(1.89)
	Panel B: Monetary Policy Uncertainty							
3	1.95	0.88	0.99	0.37	2.86	2.62	2.57	1.97
	(1.23)	(0.88)	(1.41)	(0.68)	(1.78)	(2.51)	(3.12)	(3.09)
12	1.19	0.78	1.03	0.33	2.35	2.92	2.74	2.15
	(0.76)	(0.78)	(1.48)	(0.62)	(1.48)	(2.97)	(3.56)	(3.47)
	Panel C: Geopolitical Risk							
3	0.18	0.34	1.20	0.50	4.86	3.64	2.83	1.88
	(0.10)	(0.28)	(1.24)	(0.74)	(2.81)	(3.04)	(3.18)	(2.90)
12	-0.13	0.58	1.23	0.52	3.99	3.47	3.01	2.11
	(-0.07)	(0.46)	(1.28)	(0.75)	(2.39)	(3.17)	(3.59)	(3.41)
	Panel D: Financial Stress Indicator							
3	0.01	0.61	1.90	1.31	5.07	3.02	2.70	1.36
	(0.00)	(0.63)	(2.67)	(2.22)	(3.13)	(2.49)	(2.85)	(1.99)
12	-0.20	1.08	1.85	1.20	4.23	2.59	2.44	1.33
	(-0.15)	(1.17)	(2.79)	(2.13)	(2.71)	(2.28)	(2.66)	(1.97)
	Panel E: Aggregate Firm-level Risk							
3	2.97	2.79	3.53	2.73	5.80	3.43	4.17	2.22
	(1.72)	(2.14)	(3.44)	(3.34)	(3.14)	(2.42)	(3.46)	(2.76)
12	2.09	2.37	3.07	2.33	5.50	3.79	4.49	2.58
	(1.21)	(1.85)	(3.10)	(2.98)	(3.05)	(2.61)	(3.85)	(3.62)

Table I.A. 5 Single Value-Weighted Portfolio Sorting on the Risk-Based Analyst Forecast: Alphas

This table reports the alphas of the signal-weighted risk-based appreciation expectation strategy. The rows show analyst forecast horizons (f) whereas the columns indicate holding periods (h). Following Chernov, Dahlquist, and Lochstoer (2023); Zhang (2022), we construct our factor as

$$\mathrm{R}_{p,t+1} = \sum_{i=1}^{\mathrm{N}} w_{pt}^{i} \mathrm{R}_{t+1}^{ei},$$

where w_{pt}^{i} is the portfolio's weight in currency i at time t and N is the number of currencies. The weight of a given portfolio p can be based on a signal z_{pt}^{i} (Forward Discount signal for carry factor, Value signal for value factor, and Mom signal for momentum factor). For the weight,

$$w_{pt}^{i} = \lambda \left(\mathrm{rank}(z_{pt}^{i}) - \frac{\sum_{i=1}^{\mathrm{N}} \mathrm{rank}(z_{pt}^{i})}{\mathrm{N}} \right),$$

where the scaling constant λ makes the portfolio one USD long and one USD short (and hence USD neutral). We run simultaneous regressions of $(-\widehat{\mathrm{AE}})$-based factor on these factors and take the constants as the reported alphas. Newey and West (1987) adjusted t-statistics are reported in parentheses. Our sample ranges from February 1990 to October 2016.

Forecast Horizon f	Holding Period h							
	1	3	6	12				
	Panel A: Portfolio Returns							
3	2.88	2.38	2.81	2.15				
	(2.30)	(2.44)	(3.54)	(3.66)				
12	2.84	2.67	2.91	2.18				
	(2.35)	(2.76)	(3.64)	(3.76)				
	Panel B1: Dollar + Carry				Panel B2: Dollar + Value			
3	3.12	2.72	1.36	2.14	2.64	2.08	1.04	1.80
	(2.41)	(2.47)	(2.47)	(3.00)	(1.98)	(2.08)	(2.08)	(2.78)
12	3.24	3.04	1.52	2.21	2.52	2.36	1.18	1.80
	(2.51)	(2.81)	(2.81)	(3.10)	(2.00)	(2.34)	(2.34)	(2.78)
	Panel B3: Dollar + Momentun				Panel B4: Dollar + Carry + Value +Momentum			
3	3.60	2.68	1.34	2.14	3.48	0.62	1.24	1.93
	(2.95)	(2.62)	(2.62)	(3.66)	(2.66)	(2.35)	(2.35)	(2.87)
12	3.60	2.96	1.48	2.19	3.48	2.84	1.42	1.96
	(3.04)	(2.95)	(2.95)	(3.81)	(2.75)	(2.68)	(2.68)	(3.00)

Table I.A. 6 Single Portfolio Sorting on Pooled Elastic-Net-Based Analyst Forecasts: Excess Return

The table presents the single-sorting quintile annualized excess returns on the Elastic-Net-based analyst forecast ($-\widehat{\text{AE}}^{\text{ENet}}$) and Elastic-Net-based exchange rate changes. The estimation is based on a rolling pooled ENet regression. We consider a 12-month rolling window. The rows show analyst forecast horizons (f) whereas the columns indicate holding periods (h). For the Elastic-Net-based analyst forecast in Panel A, we obtain the prediction from a predictive rolling window Elastic-Net (ENet) regression as detailed in Internet Appendix C. In Panel A, we report the zero-cost portfolio returns sorted on $-\widehat{\text{AE}}^{\text{ENet}}$. In Panel B, we forecast the change in the exchange rate during past f month ($-\Delta s_{t-f} = s_{t-f} - s_t$ in logs) via ENet regressions and combine the ENet-estimated betas and predictors at the end of the month $t-f$ as Elastic-Net-based exchange rate changes. We report the annualized returns of zero-cost portfolios sorted on $-\widehat{\Delta s_{t-f}}^{\text{ENet}}$ in Panel B. In Panel C, we regress the high-minus-low returns in Panel A on the dollar factor and high-minus-low quintile returns in Panel B, and report the constants as αs. Newey and West (1987) adjusted t-statistics are reported in parentheses. Limited to the estimation window requirement, our sample ranges from March 1992 to October 2016.

Forecast Horizon f	Holding Period h			
	1	3	6	12
	Panel A: $-\widehat{\text{AE}}^{\text{ENet}}$			
3	3.47	2.92	2.85	1.97
	(2.08)	(2.17)	(2.52)	(2.30)
12	3.49	3.19	3.28	2.21
	(2.01)	(2.29)	(2.86)	(2.57)
	Panel B: $-\widehat{\Delta s_{t-f}}^{\text{ENet}}$			
1	3.45	1.87	1.06	0.09
	(2.07)	(1.73)	(1.26)	(0.14)
3	0.47	1.54	-0.01	-1.14
	(0.28)	(1.26)	(-0.01)	(-1.67)
12	-0.50	-1.59	-3.17	-3.02
	(-0.28)	(-1.17)	(-2.69)	(-3.91)
	Panel C: αs			
3	0.29	0.78	1.42	1.74
	(2.06)	(2.19)	(2.50)	(2.17)
12	0.28	0.75	1.54	1.98
	(1.94)	(2.09)	(2.50)	(2.20)

Table I.A. 7 Single Portfolio Sorting on the Alternative Risk-Based Appreciation Expectation

The table presents the single-sorting quintile annualized excess returns on the risk-based analyst expectation which is defined as $\text{RCRE}_{i,t} = \text{FD}_{i,t} - \widehat{\text{AE}}_{i,t|t+3}$. The rows show analyst forecast horizons (f) whereas the columns indicate holding periods (h). At the end of each month t, we sort currencies based on the Panel-specific signal into the quintiles and calculate the average excess return over the holding periods, respectively. We report the time-series average annualized returns of high-minus-low quintile. Newey and West (1987) adjusted t-statistics are reported in parentheses. Our sample ranges from February 1990 to October 2016.

Forecast Horizon f	Holding Period h			
	1	3	6	12
	Panel A: RCRE			
3	5.88	3.95	4.24	3.30
	(3.49)	(3.04)	(4.09)	(4.06)
12	5.66	4.25	4.34	3.63
	(3.42)	(3.21)	(4.12)	(4.30)
	Panel A: 25% Quote Spread			
3	4.91	3.56	4.00	3.13
	(2.91)	(2.74)	(3.85)	(3.85)
12	4.73	3.87	4.11	3.46
	(2.86)	(2.93)	(3.91)	(4.11)
	Panel B: 50% Quote Spread			
3	3.94	3.17	3.75	2.97
	(2.33)	(2.44)	(3.62)	(3.65)
12	3.80	3.50	3.88	3.30
	(2.30)	(2.65)	(3.69)	(3.92)

Table I.A. 8 Fama-MacBeth Regression of Order Imbalance on Risk-based Analyst Forecast

This table presents the Fama-MacBeth regressions of order imbalance on risk-based analyst forecast ($-\widehat{\text{AE}}$). Based on the classification by the CLS dataset, in Panel A, the transactions of currency pair are executed between B*uy side* and S*ell side*, and in Panel B (Panel C, Panel D) transactions between C*orporate* (F*und*, N*on* − B*ank*F*inancial*) as price taker and B*ank* as market maker. Currency volatility and illiquidity serve as controls in this table. Newey and West (1987) adjusted t-statistics are reported in parentheses. Limited to the trade activity data, our sample ranges from September 2012 to October 2016.

Forecast Horizon f	Coef. (t-stat)					
	Intercept	-$\widehat{\text{AE}}$	Mom	Value	FD	adj-R^2
	Panel A: Buy side vs Sell side					
3	0.00	-0.03	0.34	-0.05	5.52	0.12
	(0.10)	(-0.11)	(1.49)	(-0.72)	(0.58)	
12	0.01	0.01	0.38	-0.03	4.56	0.11
	(0.30)	(0.03)	(1.64)	(-0.51)	(0.54)	
	Panel B: Corp. vs Bank					
3	0.50	3.87	-1.80	0.24	-88.37	0.26
	(1.16)	(1.36)	(-0.51)	(0.30)	(-0.78)	
12	0.47	3.32	-1.97	0.25	-73.84	0.26
	(1.01)	(1.19)	(-0.56)	(0.28)	(-0.63)	
	Panel C: Fund vs Bank					
3	0.01	-0.04	0.31	0.11	19.13	0.19
	(0.11)	(-0.08)	(0.37)	(0.89)	(0.92)	
12	-0.00	-0.09	0.32	0.09	20.58	0.20
	(-0.02)	(-0.16)	(0.39)	(0.69)	(0.98)	
	Panel D: Non-Bank Financials vs Bank					
3	0.07	0.39	-0.13	0.24	-19.59	0.13
	(0.75)	(0.71)	(-0.17)	(1.22)	(-0.80)	
12	0.07	0.55	-0.04	0.24	-18.67	0.13
	(0.89)	(1.41)	(-0.06)	(1.53)	(-0.92)	

Table I.A. 9 Fama-MacBeth Regression of Volume on Expectation Prediction

This table presents the Fama-MacBeth regressions of trading volume (in log) on risk-based analyst forecast ($-\widehat{\text{AE}}$). Based on the classification by the CLS dataset, in Panel A, the transactions of currency pair are executed between B*uy side* and S*ell side*, and in Panel B (Panel C, Panel D) transactions between C*orporate* (F*und*, N*on* − B*ank*F*inancial*) as price taker and B*ank* as market maker. Currency volatility and illiquidity serve as controls in this table. In parentheses, Newey and West (1987) adjusted t-value is reported. Limited to the trade activity data, our sample ranges from September 2012 to October 2016.

Forecast Horizon f	Coef. (t-stat)					
	Intercept	-$\widehat{\text{AE}}$	Mom	Value	FD	adj-R^2
	Panel A: Buy side vs Sell side					
3	5.58	5.07	-5.46	1.92	-349.97	0.03
	(9.64)	(1.29)	(-1.09)	(1.90)	(-2.23)	
12	5.52	4.94	-5.09	1.95	-333.82	0.03
	(9.72)	(1.29)	(-0.99)	(1.86)	(-2.19)	
	Panel B: Corp. vs Bank					
3	0.05	-0.80	8.15	4.47	-958.21	0.28
	(0.04)	(-0.12)	(1.15)	(2.13)	(-2.87)	
12	-0.12	-0.61	9.50	4.10	-907.00	0.28
	(-0.12)	(-0.11)	(1.27)	(2.24)	(-3.15)	
	Panel C: Fund vs Bank					
3	3.13	3.57	-8.27	3.03	-516.56	0.11
	(7.66)	(1.17)	(-1.47)	(3.34)	(-3.02)	
12	2.98	3.11	-7.80	2.84	-477.63	0.11
	(8.50)	(1.05)	(-1.36)	(3.48)	(-3.01)	
	Panel D: Non-Bank Financials vs Bank					
3	2.62	12.58	-3.21	2.97	-527.97	0.04
	(3.32)	(2.13)	(-0.51)	(2.27)	(-2.39)	
12	2.45	11.62	-2.66	2.88	-480.74	0.04
	(3.22)	(2.13)	(-0.41)	(2.14)	(-2.32)	