Exchange Rate Disconnect with Disagreeing Arbitrageurs

Rikuto Onishi*

October 30, 2026

Click here for the latest version

Abstract

I build a two-country general equilibrium model in which international arbitrageurs disagree about future exchange rates. Trading gains and losses redistribute wealth among the arbitrageurs, so the set of arbitrageurs whose beliefs determine the equilibrium price changes. This mechanism amplifies the exchange rate response to a standard macro shock and generates an endogenous wedge in uncovered interest parity without the need for a separate financial shock. The model is calibrated to macro, financial, and survey data and can reproduce the high exchange rate volatility and the weak relationship between exchange rates and underlying fundamentals. The calibrated disagreement-wealth multiplier implies that the amplification accounts for about 62% of the exchange rate standard deviation.

1 Introduction

Exchange rates remain difficult to reconcile with standard international macroeconomic models. They are highly volatile and closely resemble random walks. The standard macroeconomic fundamentals explain exchange rate movements very poorly. Interest differentials

*The University of Texas at Austin. Email: rikuto.onishi@utexas.edu. I am grateful for the continuous guidance and support of Olivier Coibion (co-chair), Stefano Eusepi (co-chair), and Christoph Boehm. I am especially grateful to Saroj Bhattarai, Masao Fukui (external committee), and Oliver Pfäuti for their generous time, close attention to this project, and many thoughtful discussions throughout its development. I also thank Michael Bauer, Ryan Chahrour, Jaedo Choi, Andrés Drenik, Xavier Gabaix, Oleg Itskhoki, Rohan Kekre, Husang Kim, Zheng Liu, Thomas Mertens, Nitya Pandalai-Nayar, Thuy Lan Nguyen, Damjan Pfajfar, Timo Reinelt, Sanjay Singh, Mark Spiegel, and Yeji Sung for valuable comments and conversations, as well as seminar and conference participants at numerous institutions and meetings. All remaining errors are my own.

often predict currency movements in the opposite direction that uncovered interest parity (UIP) would predict. Exchange rates co-move far less with relative consumption than the standard complete markets models imply. Recent quantitative work has made significant progress by combining segmented international financial markets with *exogenous* shocks to currency demand or intermediaries' risk-bearing capacity, which lead to movements in the UIP wedge. This paper takes an alternative approach and *endogenizes* the UIP wedge through disagreement among international arbitrageurs and the consequent shifts in their wealth shares.

Consider two groups of arbitrageurs who disagree about the future value of a currency. Relatively optimistic investors take larger long positions, while pessimistic investors take smaller or short positions. When a macro shock causes the currency to appreciate, optimists make the most money and become wealthier relative to pessimists. As arbitrageurs' capacity to bear risk increases with their wealth, optimists' views become more influential in the equilibrium demand for assets. Their greater influence increases demand for the currency and, hence, appreciation. Further appreciation then revalues wealth in favor of investors whose positions did well. This is the wealth revaluation channel through which disagreement can alter the transmission of a macro shock even when average beliefs are rational.

I formalize this mechanism in a two-country general equilibrium model with segmented asset markets and risk-averse international arbitrageurs. The framework builds on Gabaix and Maggiori (2015) and Itskhoki and Mukhin (2021). Home households trade bonds denominated in their own consumption bundle, Foreign households trade bonds denominated in theirs, and international arbitrageurs can hold both assets. Arbitrageurs take currency positions based on their subjective expectations of excess returns, and their risk tolerance scales with their wealth. Disagreement leads to different currency positions, and realized returns alter the distribution of wealth across arbitrageurs.

This environment generates an endogenous UIP wedge. The UIP condition depends on the wealth-weighted average of arbitrageurs' expectations about future exchange rates. When investors who are relatively optimistic about a currency hold a larger share of intermediary wealth, their expectations receive greater weight in equilibrium pricing. The difference between wealth-weighted and average expectations creates a wedge in expected currency returns. As actual exchange rate movements redistribute wealth among arbitrageurs, this wedge evolves endogenously in response to macroeconomic shocks.

A simplified version of the model gives an analytical description of the mechanism. The equilibrium can be reduced to three equations: one for the interest rate differential, one for the exchange rate, and one for the wealth share held by the relatively optimistic arbitrageurs. These equations give a closed-form disagreement-wealth multiplier. A shock of fundamental

news will move the exchange rate. The resulting return will redistribute the arbitrageurs' wealth, and changes in wealth shares will then feed back into equilibrium pricing. Solving for the feedback yields an exchange rate response greater than in an economy without the wealth revaluation channel. The multiplier's strength increases with the level of disagreement and decreases with arbitrageurs' effective risk aversion and the rate at which wealth differences dissipate.

The same mechanism generates several dimensions of exchange-rate disconnect. When macro fundamentals and arbitrageurs' wealth are sufficiently persistent, exchange rates will resemble random walks even as they continue to react to macro shocks. The endogenous UIP wedge can also give rise to a negative Fama coefficient. At the same time, amplifying exchange rate movements relative to consumption can produce a negative Backus-Smith correlation. These results show the conditions under which a common wealth revaluation mechanism can generate exchange rate volatility, UIP deviations, and weak exchange rate comovement with macroeconomic variables.

I quantify the mechanism using macroeconomic, financial, and survey data for the US and advanced foreign economies. Individual forecasts from Blue Chip Financial Forecasts discipline the level and persistence of exchange rate disagreement. US broker-dealer equity relative to GDP disciplines the size and persistence of the intermediary wealth. Macroeconomic variables, yields and exchange rate moments together discipline the remaining parameters.

The baseline model imposes rational expectations on (equal-weighted) average across arbitrageurs and does not have an exogenous UIP shock. Beliefs are still dispersed across individuals, and they can infer persistent fundamentals from the noisy data they observe. The calibrated baseline reproduces the key features of the exchange rate disconnect: high volatility relative to consumption and interest differentials, exchange rates that are essentially a near random walk and negative Fama and Backus-Smith coefficients. With all other parameters held constant, removing disagreement reduces the quarterly standard deviation of the exchange rate by 62% and reverses the signs of both coefficients. Re-estimating the conventional macroeconomic shock processes without disagreement does not recover the same joint exchange-rate behavior.

The augmented model incorporates average belief distortions and an exogenous UIP shock and is separately calibrated. Standard demand and supply shocks account for the vast majority of currency variation in both models (around 80% of the one-quarter exchange rate variance in the augmented model). Direct disagreement innovations contribute very little, and the exogenous UIP shock remains a minor driver when it is disciplined by a risk proxy. Disagreement matters because of the channel through which it operates: fundamental news moves the currency, redistributes wealth among investors with differing beliefs, and changes

how much those investors subsequently influence currency pricing. The augmented model examines how this mechanism works alongside average belief bias and exogenous financial shocks.

I provide further empirical support for two elements of the mechanism using data external to that used to discipline the quantitative model. First, I examine whether survey expectations are reflected in actual currency positions. Using Bloomberg forecast data and CFTC dealer positions for the six major currencies, I find that revisions to the consensus exchange rate forecast predict changes in dealer positions, after controlling for contemporaneous and lagged spot rate movements. The relationship is positive in all six cases. Second, the model predicts that greater disagreement will increase conditional exchange rate volatility, as greater disagreement in expectations will widen position differences and therefore increase wealth redistribution effects following currency movements. This is exactly what I find: unusually high levels of survey disagreement do in fact predict higher subsequent realized currency volatility. A one residual standard deviation increase in disagreement predicts about a 2% increase in realized volatility over the next year. These exercises provide additional support for the links between expectations, currency positions and subsequent currency risk that underlie the model.

The paper contributes to models of exchange-rate determination in segmented financial markets, including Gabaix and Maggiori (2015), Itskhoki and Mukhin (2021, 2025), and Kekre and Lenel (2024). These models show how limited international risk-bearing capacity plays a key role in determining exchange rate dynamics. I add an endogenous source of UIP deviations: macroeconomic shocks redistribute wealth across disagreeing arbitrageurs and change whose beliefs influence prices. This provides a channel through which supply and demand shocks can produce large currency movements even when the corresponding movements in aggregate macroeconomic quantities and interest rates remain modest.[1]

The model also builds on work on subjective and heterogeneous beliefs in exchange-rate determination, including Evans and Lyons (2002), Gourinchas and Tornell (2004), Bacchetta and Van Wincoop (2006), Du et al. (2021), Candian and De Leo (2025), and Molavi et al. (2024). The key channel, distinct from these existing papers, is the interaction between disagreement and investor wealth. Investors' positions determine how shocks redistribute wealth, and these wealth changes subsequently alter the influence of different beliefs on equilibrium prices. The distribution of wealth across investors with different beliefs therefore becomes a state variable linking past currency returns to current exchange-rate determina-

[1]Other explanations for UIP deviations emphasize risk premia and convenience yields (Engel, 2016; Valchev, 2020; Jiang et al., 2024) or portfolio adjustment costs (Chang et al., 2015; Bacchetta and Van Wincoop, 2021).

tion.

Recent empirical evidence finds conditional connections between exchange rates and macroeconomic fundamentals (Hassan et al., 2016; Lilley et al., 2022; Engel and Wu, 2024; Kekre and Lenel, 2024; Chahrour et al., 2024; Stavrakeva and Tang, 2024).[2] In particular, Stavrakeva and Tang (2024) find that macroeconomic news explains much of exchange-rate variation through responses to past announcements. The disagreement-wealth wedge is one source of this history dependence, as prior shocks can alter investors' wealth shares and thus their impact on the current currency level. The average belief distortions in the augmented model also affect how news is transmitted over time, consistent with their argument.

The paper also relates to asset-pricing models in which heterogeneous beliefs and relative wealth affect equilibrium prices (Detemple and Murthy, 1994; Jouini and Napp, 2007; Xiong and Yan, 2010). I bring a similar mechanism into a two-country general equilibrium model and evaluate its implications for the core exchange-rate puzzles. Keeping the permanent disagreement at the approximation point maintains the wealth-belief feedback in the first-order approximation. Caramp and Silva (2026) obtain related tractability with heterogeneous beliefs in a continuous-time, closed-economy model. The discrete-time international framework here links the mechanism to the quantitative segmented-market models of Itskhoki and Mukhin (2021) and Kekre and Lenel (2024).

Section 2 presents the model and information structure. Section 3 derives the wealth-revaluation mechanism and its implications for exchange-rate disconnect. Section 4 quantifies the baseline model. The augmented model further introduces average belief distortions and UIP shock. Section 5 presents evidence on survey beliefs, dealer positions, and subsequent currency volatility. Section 6 concludes. The appendices contain proofs and additional quantitative and empirical results.

2 Model

This section outlines a two-country general equilibrium model with segmented asset markets and risk averse international arbitrageurs who try to make profits from carry trades. This framework follows Itskhoki and Mukhin (2021); Kekre and Lenel (2024), but deviates from their setup by (i) assuming risk tolerance proportional to arbitrageur wealth and (ii) allowing arbitrageurs to hold heterogeneous, potentially subjective beliefs.

[2] Miyamoto et al. (2025) find that dominant real-exchange-rate shocks are largely distinct from dominant business-cycle shocks. This addresses a different notion of connection from responses to news about future productivity in Chahrour et al. (2024): such news can move currency prices before the corresponding macroeconomic changes materialize.

2.1 Environment

Time is discrete, $t = 0, 1, 2, \ldots$. There are two countries, Home and Foreign, comprised of a measure 1 households. The economy features three sets of agents: Home households, Foreign households, and arbitrageurs, whose roles are detailed below. Home and Foreign households share common beliefs, while arbitrageurs may hold heterogeneous and potentially subjective beliefs.

Home households The representative Home household chooses consumption and Home bond holdings to maximize expected lifetime utility,

$$\max_{\{c_{t+s},\, b_{t+s}\}_{s\geq 0}} \mathbb{E}_t \left[\sum_{s=0}^{\infty} \left(\prod_{k=0}^{s-1} \beta_{t+k} \right) \frac{c_{t+s}^{1-\sigma}}{1-\sigma} \right],$$

subject to its budget constraint below at every date and a no-Ponzi condition, given initial bond holdings. Here β_t is the time-varying subjective discount factor from date t to $t+1$, and $\sigma > 0$ is relative risk aversion. Expectations reflect the households' common beliefs, and the empty product for $s = 0$ equals one. The consumption basket is a CES aggregator over domestic Home goods and imported Foreign goods,

$$c_t = \left[(1-\gamma)^{\frac{1}{\theta}} \left(c_{Ht}\right)^{\frac{\theta-1}{\theta}} + \gamma^{\frac{1}{\theta}} \left(c_{Ft}\right)^{\frac{\theta-1}{\theta}} \right]^{\frac{\theta}{\theta-1}}.$$

Here, θ denotes the trade elasticity, which governs the elasticity of substitution across domestic goods and imports, and $\gamma \in (0, 1/2)$ is the import expenditure share at equal prices; a lower value means greater consumption home bias.

At the beginning of periods t, the representative Home household is endowed with predetermined position in Home bond holdings, b_{t-1}. This bond is denominated in the Home consumption bundle, paying a risk-free interest rate r_t at $t+1$. Intermediate production is exogenous in this economy. The Home household receives an endowment of the Home intermediate good, y_t, which it sells at producer price p_t measured in units of its consumption bundle. It also receives profit transfers π_t from the domestic final good firms it owns. It also receives net profits from intermediaries, π_t^a, which are described further below. It can either consume their wealth or save it in the form of Home bonds. Putting all together, the budget constraint of Home household is

$$c_t + \frac{b_t}{1+r_t} = p_t y_t + b_{t-1} + \pi_t + \pi_t^a. \tag{1}$$

The households minimize expenditure, resulting in the conventional constant elasticity demand schedules:

$$c_{Ht} = (1-\gamma)\,(p_{Ht})^{-\theta}\,c_t, \qquad c_{Ft} = \gamma\,(p_{Ft})^{-\theta}\,c_t.$$

Foreign households The representative Foreign household solves the symmetric intertemporal problem,

$$\max_{\{c^*_{t+s},\,b^*_{t+s}\}_{s\geq 0}} \mathbb{E}_t\left[\sum_{s=0}^{\infty}\left(\prod_{k=0}^{s-1}\beta^*_{t+k}\right)\frac{(c^*_{t+s})^{1-\sigma}}{1-\sigma}\right],$$

subject to its budget constraint below at every date and a no-Ponzi condition, given initial Foreign bond holdings. The one-period subjective discount factor is β^*_t, and Foreign counterparts are denoted by a superscript $*$. The consumption basket is a CES aggregator over domestic Foreign goods and imported Home goods,

$$c^*_t = \left[(1-\gamma)^{\frac{1}{\theta}}\,(c^*_{Ft})^{\frac{\theta-1}{\theta}} + \gamma^{\frac{1}{\theta}}\,(c^*_{Ht})^{\frac{\theta-1}{\theta}}\right]^{\frac{\theta}{\theta-1}}.$$

Denoting b^*_t the household's saving in its Foreign domestic bond denominated by its overall consumption bundle which pays r^*_t at $t+1$, y^*_t its endowment of the Foreign intermediate good, and π^*_t its transfers from domestic firms, its budget constraint is

$$c^*_t + \frac{b^*_t}{1+r^*_t} = p^*_t y^*_t + \pi^*_t + b^*_{t-1},$$

where p^*_{Ft} (p^*_{Ht}) is the price of the domestic (imported) bundle, and p^*_t the producer price of its endowment, relative to its overall consumption bundle. The CES demands are given by

$$c^*_{Ht} = \gamma\,(p^*_{Ht})^{-\theta}\,c^*_t,\; c^*_{Ft} = (1-\gamma)\,(p^*_{Ft})^{-\theta}\,c^*_t.$$

We have two Euler equations from households:

$$\beta_t \mathbb{E}_t\left[\left(\frac{c_{t+1}}{c_t}\right)^{-\sigma}\right](1+r_t) = 1$$

$$\beta^*_t \mathbb{E}_t\left[\left(\frac{c^*_{t+1}}{c^*_t}\right)^{-\sigma}\right](1+r^*_t) = 1$$

Arbitrageurs A third set of agents, arbitrageurs, trade in both Home and Foreign bonds. There is a continuum of arbitrageurs in the unit interval $\iota \in [0,1]$. Each individual arbitrageur is infinitely small, taking prices as exogenously given. For concreteness, they are

comprised of two groups, denoted by $i \in \{A, B\}$, whose only ex-ante heterogeneity comes from their expectations. We let $\eta_i \geq 0$ denote the mass of arbitrageurs of group $i \in \{A, B\}$, where $\eta_A + \eta_B = 1$.[3]

They invest their current wealth in the risk-free Home bond and take a self-financing carry position by going long the Foreign bond and short an equal value of the Home bond, or vice versa. We denote arbitrageur ι belonging to group i's carry positions in the corresponding bonds as $b_t^{*i}(\iota)$ and $b_t^i(\iota)$. The carry positions satisfy $\frac{b_t^i(\iota)}{1+r_t} = -\frac{e_t b_t^{*i}(\iota)}{1+r_t^*}$, where e_t is the price of the Foreign consumption bundle in units of the Home bundle (the real exchange rate). Thus, the increase of e_t represents the appreciation of Foreign currency. I denote the size of this position in Home consumption units as $n_t^i(\iota) = \frac{e_t b_t^{*i}(\iota)}{1+r_t^*}$. Arbitrageurs have a discrete time version of the perpetual youth structure, with a constant exit probability ξ. When they exit, their accumulated wealth is rebated back to Home households. While they survive, they retain their wealth and rebalance their carry positions each period.

Following Kekre and Lenel (2026a), arbitrageurs have a myopic mean-variance objective over next-period wealth. At each date t, arbitrageur ι in group i chooses its current position according to

$$\max_{n_t^i(\iota)} \left\{ E_t^i \left[w_{t+1}^i(\iota) \right] - \frac{\nu}{2(1+r_t) w_t^i(\iota)} \mathrm{Var}_t^i \left[w_{t+1}^i(\iota) \right] \right\},$$

where $\nu > 0$ governs risk aversion, gross current wealth $(1+r_t)w_t^i(\iota) > 0$ is taken as given, and conditional moments reflect group i's beliefs. Risk tolerance is proportional to current wealth. This objective can be motivated by a second-order certainty-equivalent approximation to CRRA utility over next-period wealth, with risk tolerance evaluated at risk-free next-period wealth. Myopic portfolio choice rules out intertemporal hedging demand by assumption. The choice is subject to the law of motion for individual wealth before exit,

$$w_{t+1}^i(\iota) = w_t^i(\iota)(1+r_t) + n_t^i(\iota) \left(\frac{e_{t+1}}{e_t}(1+r_t^*) - (1+r_t) \right).$$

Let W_t^i and N_t^i denote wealth and the Home-valued carry position per member of group i, measured after exit and entry at date t. Aggregate arbitrageur wealth and the aggregate carry position are $W_t = \sum_{i \in \{A,B\}} \eta_i W_t^i$ and $N_t = \sum_{i \in \{A,B\}} \eta_i N_t^i$. Averaging individual wealth

[3] Two groups suffice to study the wealth-disagreement mechanism. The portfolio aggregation argument extends to any finite number of groups; the analytical results below use two groups and the symmetry restrictions stated in Section 2.2.

within each group and accounting for exit and entry gives

$$W_{t+1}^i = (1-\xi)\left(W_t^i(1+r_t) + N_t^i\left(\frac{e_{t+1}}{e_t}(1+r_t^*) - (1+r_t)\right)\right) + \xi W_0^i.$$

Here, a fraction ξ of arbitrageurs exit the market at the beginning of each period, and they are replaced by the same mass of newborn arbitrageurs, each endowed with actual wealth W_0^i, where the constant endowment is $W_0^i = W_0$ common to both groups. The entrants' contribution to average group wealth, ξW_0^i, enters the law of motion for wealth, and is financed by Home households. Once they exit the market, the accumulated wealth is rebated back to Home household. On aggregate, in Home consumption units, exiting arbitrageurs rebate $\frac{\xi}{1-\xi}(W_t - \xi W_0)$ and new entrants receive ξW_0. Taken together, writing the entrant endowment averaged across groups as $W_0 = \sum_{i\in\{A,B\}} \eta_i W_0^i$, the net profits Home household receives each period is

$$\pi_t^a = \frac{\xi}{1-\xi}\left(W_t - W_0\right).$$

In the limiting case in which the exit probability approaches unity, $\xi \to 1$, the group wealth law resets wealth to the fixed entrant endowment each period, so the wealth distribution ceases to be a state variable.

Final good firms In each country, a continuum of monopolistically competitive final good firms, owned by domestic households, transforms the domestic intermediate good one-for-one into varieties sold globally at destination-specific flexible prices. As in Kekre and Lenel (2024), these firms exist in the model *only* to accommodate incomplete pass-through of the exchange rate into the prices consumers face, as in the data. This model block replaces the law of one price (LOP) assumption, motivated by empirically observed large deviations from LOP.

The bundles $c_{Ht}, c_{Ft}, c_{Ht}^*, c_{Ft}^*$ of are themselves CES aggregators over the varieties, with elasticity θ_v across *domestic* varieties. Following Alessandria (2009) and Kekre and Lenel (2024), the elasticity across *imported* varieties moves with the real exchange rate, that is, the Home households' aggregators over domestic and imported varieties are

$$c_{Ht} = \left(\int_0^1 c_{Ht}(i)^{\frac{\theta_v - 1}{\theta_v}} di\right)^{\frac{\theta_v}{\theta_v - 1}}, \quad c_{Ft} = \left(\int_0^1 c_{Ft}(i^*)^{\frac{\theta_{v,t} - 1}{\theta_{v,t}}} di^*\right)^{\frac{\theta_{v,t}}{\theta_{v,t} - 1}},$$

respectively, where the elasticity of substitution across imported varieties is $\theta_{v,t} = \theta_v \left(e_t/e\right)^{\xi_v}$, and throughout I denote steady-state values without time subscripts. Similarly, Foreign

households' aggregators over imported and domestic varieties are

$$c^*_{Ht} = \left(\int_0^1 c^*_{Ht}(i)^{\frac{\theta^*_{v,t}-1}{\theta^*_{v,t}}} \, di \right)^{\frac{\theta^*_{v,t}}{\theta^*_{v,t}-1}}, \quad c^*_{Ft} = \left(\int_0^1 c^*_{Ft}(i^*)^{\frac{\theta_v-1}{\theta_v}} \, di^* \right)^{\frac{\theta_v}{\theta_v-1}},$$

where the elasticity of substitution across imported varieties is $\theta^*_{v,t} \equiv \theta_v \left(e_t/e\right)^{-\xi_v}$. This captures the pricing-to-market in a reduced form. Here, pricing complementarities make markups fall in markets where a firm's relative cost has risen. $\xi_v = 0$ restores constant elasticities and the LOP.

A Home firm i buys the intermediate at p_t (from Home households) and chooses prices $\{p_{Ht}(i),\, p^*_{Ht}(i)\}$, facing the demand schedules in the two markets, to maximize profits

$$\pi_t(i) = \big(p_{Ht}(i) - p_t\big)\Big(\frac{p_{Ht}(i)}{p_{Ht}}\Big)^{-\theta_v} c_{Ht} + \big(e_t\, p^*_{Ht}(i) - p_t\big)\Big(\frac{p^*_{Ht}(i)}{p^*_{Ht}}\Big)^{-\theta^*_{v,t}} c^*_{Ht}, \tag{2}$$

where $e_t p^*_{Ht}(i)$ converts export revenue into Home real units and c^*_{Ht} is aggregate Foreign demand. A Foreign firm i^* solves the mirror problem,

$$\pi^*_t(i^*) = \Big(\frac{p_{Ft}(i^*)}{e_t} - p^*_t\Big)\Big(\frac{p_{Ft}(i^*)}{p_{Ft}}\Big)^{-\theta_{v,t}} c_{Ft} + \big(p^*_{Ft}(i^*) - p^*_t\big)\Big(\frac{p^*_{Ft}(i^*)}{p^*_{Ft}}\Big)^{-\theta_v} c^*_{Ft}. \tag{3}$$

Per-capita profit rebates are $\pi_t = \int_0^1 \pi_t(i)\, di$ to Home households and $\pi^*_t = \int_0^1 \pi^*_t(i^*)\, di^*$ to Foreign households.

Market clearing and driving forces There are four markets: Home goods, Foreign goods, Home bonds, and Foreign bonds. Both bonds are in zero net supply, so their market-clearing conditions are

$$\frac{b_t}{1+r_t} + W_t - N_t = 0, \qquad \frac{e_t b^*_t}{1+r^*_t} + N_t = 0.$$

Here b_t and b^*_t are the face values of household bond holdings; the two equations express all positions in Home consumption units.

The economy is perturbed by demand and supply forces, each of which are modeled as discount factor and endowment shocks. For each driving variable $x_t \in \{\beta_t, \beta^*_t, y_t, y^*_t\}$ with its steady-state value x, I assume $\hat{x}_t = \log x_t - \log x$ has a persistent AR(1) component, τ^x_t, with persistence parameter ρ_x, and mean zero short-run i.i.d. component, v^x_t, with standard deviation σ_{v^x}:

$$\hat{x}_t = \tau^x_t + v^x_t,$$

where $\tau_t^x = \rho_x \tau_{t-1}^x + \varepsilon_t^x$, with innovation standard deviation σ_{ε^x}. Persistent components of shocks are correlated across countries to capture clear international comovement patterns observed in the data. The correlation of shocks with another driving variable x' is denoted by $\rho_{xx'}$. Agents observe realizations of $\hat{x}_t$, but not τ_t^x and v_t^x separately. Therefore, agents need to filter out the persistent components to form expectations about future x. This is discussed in detail in a separate section below.

2.2 Information structure

The two groups of arbitrageurs interpret the same observed fundamentals differently. Following Caballero and Simsek (2022), they are opinionated: each group regards the other's interpretation as uninformative conditional on its own. Current interpretations and beliefs are mutually observed, but neither the other group's beliefs nor equilibrium prices provide additional information about the latent persistent component. Disagreement therefore persists despite common observations.

Group $i \in \{A, B\}$ treats $\hat{x}_t + \mu_t^{xi}$ as a noisy observation of τ_t^x, where $\mu_t^{xi} = \mu^{xi} + \widetilde{\mu}_t^{xi}$ consists of a constant interpretation and a transitory interpretation drawn from a joint normal distribution that is i.i.d. across periods. The groups have equal population shares, $\eta_A = \eta_B = 1/2$, so average interpretation bias is $\bar{\mu}_t^x = (\mu_t^{xA} + \mu_t^{xB})/2$.

Let $\tau_{t|t}^{xi} \equiv E_t^i[\tau_t^x]$ denote group i's filtered persistent state. Define consensus beliefs as $\bar{\tau}_{t|t}^x = (\tau_{t|t}^{xA} + \tau_{t|t}^{xB})/2$ and disagreement as $d_t^x = \tau_{t|t}^{xA} - \tau_{t|t}^{xB}$. With common steady-state Kalman gain G_x and $\varphi_x = \rho_x(1 - G_x)$, their laws of motion are

$$\begin{aligned} \bar{\tau}_{t|t}^x &= \varphi_x \bar{\tau}_{t-1|t-1}^x + G_x(\hat{x}_t + \bar{\mu}_t^x), \\ d_t^x &= (1 - \varphi_x) d_{ss}^x + \varphi_x d_{t-1}^x + G_x \varepsilon_{\mu,t}^x, \end{aligned}$$

where d_{ss}^x is the permanent belief gap and $\varepsilon_{\mu,t}^x = \widetilde{\mu}_t^{xA} - \widetilde{\mu}_t^{xB} \sim N(0, \sigma_{\mu,x}^2)$ is the transitory disagreement innovation.

In the zero-average-bias benchmark, $\bar{\mu}_t^x = 0$, consensus beliefs follow the rational Bayesian filter even though the groups disagree. Average bias and disagreement can therefore be varied separately. Both groups understand the same equilibrium mapping from fundamentals and the distribution of beliefs to prices and wealth, and they agree on the expected path of disagreement.

Remark 1. (Analytical benchmark.) The information structure nests the simpler case used in Section 3. With unbiased consensus beliefs, $\bar{\mu}_t^x = 0$, let transitory fundamental noise and disagreement innovations vanish, $\sigma_{v^x} \to 0$ and $\sigma_{\mu,x} \to 0$, for each fundamental x.

Then $G_x \to 1$ and $\varphi_x \to 0$. In the limiting benchmark, observed fundamentals follow $\hat{x}_t = \rho_x \hat{x}_{t-1} + \varepsilon_t^x$, and the fundamental belief gap is constant, $d_t^x = d_{ss}^x = \mu^{xA} - \mu^{xB}$. Wealth redistribution nevertheless changes the influence of these beliefs on equilibrium prices. Under the log-linear approximation used below, the total one-period exchange-rate forecast gap is also constant, $D_t = d_{ss}$, and hence $d_t = 0$. Appendix B.5 derives the consistency condition linking d_{ss} to the fundamental gaps d_{ss}^x.

Appendix B.1 gives the formal information structure, Proposition 4, and its proof. The next section shows how wealth weighting makes disagreement affect equilibrium pricing even when consensus beliefs are unbiased.

Remark 2. (Private information and price learning.) With heterogeneous subjective long-run models and common constant filtering gains, private signals and shared price signals can preserve the AR(1) form of centered fundamental belief gaps. Here a nonzero permanent gap comes from fixed subjective model differences. Appendix B.1.1 states the conditions and proves the result.

2.3 Equilibrium and solution

Appendix A collects the full equilibrium system and defines the units and approximation conventions used below.

First-order approximation We use the small noise approximation developed by Borovička, Hansen and Sargent (2023), as is common in recent segmented market models such as Itskhoki and Mukhin (2021) and Kekre and Lenel (2024). Let all stochastic deviations be scaled by δ and let risk aversion satisfy

$$\nu(\delta)\delta^2 \to \Gamma \in (0, \infty)$$

as $\delta \to 0$. In words, as variances become infinitely small, risk aversion becomes infinitely larger at the same rate as variance. In effect, the risk terms remain at the first order and portfolio positions remain bounded even as primitive risk vanishes. The portfolio rule below assumes a common conditional variance across groups and holds it fixed at its steady-state value in the first-order approximation. We normalize this steady-state conditional variance of $\Delta \log e_{t+1}$ to δ^2, absorbing its constant scale into Γ. All equilibrium conditions below are first order log-linear conditions around the possibly asymmetric deterministic steady state. The log deviation from the steady state is denoted as $\hat{z}_t = \log z_t - \log z$, where I denote steady state objects by dropping time subscript throughout the paper. For interest rates, hats denote log

deviations of gross rates: $\hat{r}_t = \log[(1+r_t)/(1+r)]$ and $\hat{r}^*_t = \log[(1+r^*_t)/(1+r^*)]$. Tildes denote additive deviations from steady state, with normalizations stated where needed.

Optimal portfolio choice Writing the arithmetic carry return as $R^{\text{ex}}_{t+1} \equiv \frac{e_{t+1}}{e_t}(1+r^*_t) - (1+r_t)$, the first-order condition for the myopic mean-variance problem is

$$E^i_t\left[R^{\text{ex}}_{t+1}\right] - \nu \frac{n^i_t(\iota)}{(1+r_t)w^i_t(\iota)} \text{Var}^i_t\left[R^{\text{ex}}_{t+1}\right] = 0.$$

The corresponding log carry return is $rx_{t+1} \equiv \Delta \log e_{t+1} + \log[(1+r^*_t)/(1+r_t)]$. Following an approximation approach similar to Campbell and Viceira (1999) and aggregating decisions within each type, we obtain the log-return approximation

$$\frac{N^i_t}{W^i_t} = \frac{E^i_t\left[\Delta \log e_{t+1}\right] + \log \frac{1+r^*_t}{1+r_t}}{\nu \text{Var}^i_t\left[\Delta \log e_{t+1}\right]},$$

which is the maintained myopic portfolio rule. Appendix B.2 derives the exact optimum, the log-return approximation, and the first-order position rule. Let $L_i \equiv E^i\left[\log e\right] + \log \frac{1+r^*}{1+r}$ denote group i's steady-state expected carry return and define $UIP^i_t \equiv E^i_t[\hat{e}_{t+1}] - \hat{e}_t + \hat{r}^*_t - \hat{r}_t$. For exchange-rate forecasts in first-order equations, $E^i_t[\hat{e}_{t+1}]$ denotes the change from group i's own steady-state forecast. At first order,

$$\widetilde{N}^i_t = \frac{L_i}{\Gamma}\hat{W}^i_t + \frac{1}{\Gamma}UIP^i_t,$$

where $\tilde{N}^i_t = \frac{N^i_t - N^i}{W^i}$. Heterogeneous expectations affect positions through UIP^i_t, while permanent disagreement affects both steady-state positions and their wealth elasticities through L_i.

Modified UIP Substituting the position rule into bond market clearing, we obtain a version of UIP in this economy.

Lemma 1. *The modified UIP condition in levels is*

$$\sum_{i\in\{A,B\}} \omega^i_t \frac{E^i_t\left[\Delta \log e_{t+1}\right] + \log \frac{1+r^*_t}{1+r_t}}{\nu var^i_t\left[\Delta \log e_{t+1}\right]} = 1 + \frac{b_t}{(1+r_t)W_t}, \tag{4}$$

where $\omega^i_t = \frac{W^i_t}{2W_t}$ denotes group i's share of aggregate arbitrageur wealth. For wealth shares, $\widetilde{\omega}^i_t \equiv \omega^i_t - \omega^i$ denotes a level deviation from steady state. In the benchmark with unbiased

consensus beliefs, the first-order condition is

$$\mathbb{E}_t\left[\hat{e}_{t+1}\right] - \hat{e}_t + \hat{r}_t^* - \hat{r}_t = -d_{ss}\widetilde{\omega}_t^A - \delta_w d_t + \frac{\beta\Gamma}{W}\left[y\widetilde{b}_t - b(\hat{W}_t + \hat{r}_t)\right], \tag{5}$$

where $\mathbb{E}_t$ denotes rational expectations. The one-period forecast gap is $D_t \equiv E_t^A[\log e_{t+1}] - E_t^B[\log e_{t+1}] = d_{ss} + d_t$, with steady-state level $d_{ss} = E^A[\log e] - E^B[\log e]$ and centered deviation $d_t = E_t^A[\hat{e}_{t+1}] - E_t^B[\hat{e}_{t+1}]$. Here $\delta_w = (\omega^A - \omega^B)/2$ and $\widetilde{b}_t = (b_t - b)/y$; b and W are steady-state Home household bond holdings and aggregate arbitrageur wealth.

The modified UIP condition (4) shows that exchange rates are priced by the wealth-weighted average of the two groups' expectations. To first order, this wealth-weighted belief decomposes into a wealth-share term, $-d_{ss}\widetilde{\omega}_t^A$ and a time-varying disagreement term $-\delta_w d_t$. From the viewpoint of an econometrician outside the model who holds objective beliefs, both act as mispricing forces, whose implications to be explored in Section 3. The remaining terms capture bond supply relative to arbitrageur wealth, as in segmented-asset-market models, such as Itskhoki and Mukhin (2021); Kekre and Lenel (2024).

Wealth–share dynamics One of the key components of the current model is the dynamics of the wealth share. They are characterized by the following equation:

Lemma 2. *The first-order dynamics of the wealth-share deviation of group A in the benchmark with unbiased consensus beliefs are given by*

$$\widetilde{\omega}_{t+1}^A = (1-\xi)\left\{\left[\frac{1}{\beta} + \frac{R^{\text{ex}}(\bar{L} - \delta_w d_{ss})}{\Gamma}\right]\widetilde{\omega}_t^A + \omega^A\omega^B\left[\frac{R^{\text{ex}}}{\Gamma}d_t + \frac{d_{ss}}{\Gamma}\left(R^{\text{ex}}\hat{W}_t + \widetilde{R}_{t+1}^{\text{ex}}\right)\right]\right\} \tag{6}$$

where $R^{\text{ex}} \equiv 1/\beta^* - 1/\beta$ is the steady-state excess return on Foreign bonds, $\bar{L} = \log\beta - \log\beta^*$, and $\widetilde{R}_{t+1}^{\text{ex}} = \frac{1}{\beta^*}(\hat{e}_{t+1} - \hat{e}_t + \hat{r}_t^*) - \frac{1}{\beta}\hat{r}_t$.

Appendix B.4 derives this law and its symmetric analytical limit.

Equation (6) makes the role of belief heterogeneity explicit. Absent disagreement, the two groups choose the same carry position per unit of wealth. Aggregate portfolio demand then depends only on total wealth, so a representative arbitrageur suffices to determine prices and the wealth share drops out of the pricing state vector. Disagreement components, d_{ss} and d_t, create differences in positions per unit of wealth and turns the share into a state that affects prices, driven by two forces.

The first is revaluation. Under the normalization $d_{ss} > 0$, group-A arbitrageurs are the optimists about Foreign bonds: they expect higher excess returns on Foreign bonds than

group-B does. In steady state their Foreign-bond exposure per unit of wealth is larger by d_{ss}/Γ, so any shock that delivers a positive carry return $\widetilde{R}^{\text{ex}}_{t+1}$—whatever its source—revalues wealth from group-B toward group-A, raising $\widetilde{\omega}^A_{t+1}$.

The second is redistribution through the steady-state excess return R^{ex}. Higher aggregate wealth $\hat{W}_t$ or wider disagreement d_t enlarges the optimists' Foreign-bond exposure relative to the pessimists'. Holding returns fixed, these changes raise group A's wealth share when $R^{\text{ex}} > 0$.

3 Analytical Results

This section isolates the model's core mechanism in a three-equation exchange-rate block. Fundamental news moves exchange rates and redistributes wealth across arbitrageurs. This redistribution changes the weight of their beliefs in market prices, feeding back into exchange rates and amplifying the initial response. This feedback alone can generate the key FX-disconnect moments emphasized in the recent literature.

3.1 Simplifying limits

I impose three restrictions to isolate this mechanism. The first follows Itskhoki and Mukhin (2021): I let the financial sector grow large relative to the real economy, $\frac{y}{W} \to 0$. The bond-supply terms drop out of the modified UIP condition (19), which becomes

$$\mathbb{E}_t\left[\Delta\hat{e}_{t+1}\right] + \hat{r}^*_t - \hat{r}_t = -d_{ss}\widetilde{\omega}^A_t - \delta_w d_t. \tag{7}$$

The second limit sends the steady-state discount factor differential to zero, $\beta/\beta^* \to 1$. With equal entrant endowments and $(1-\xi)/\beta < 1$, the steady-state wealth formula in Appendix A.2 gives $W^A = W^B$, so $\omega^A = \omega^B = 1/2$ and $\delta_w = 0$. The wedge collapses to $d_{ss}\widetilde{\omega}^A_t$, and the wealth-share dynamics reduce to

$$\widetilde{\omega}^A_{t+1} = (1-\xi)\left(\frac{1}{\beta}\widetilde{\omega}^A_t + \omega^A\omega^B\frac{d_{ss}}{\beta\Gamma}\left(\hat{e}_{t+1} - \hat{e}_t + \hat{r}^*_t - \hat{r}_t\right)\right).$$

The centered log carry return, $\widetilde{rx}_{t+1} \equiv rx_{t+1} - \bar{L} = \beta\widetilde{R}^{\text{ex}}_{t+1} = \hat{e}_{t+1} - \hat{e}_t + \hat{r}^*_t - \hat{r}_t$, decomposes into

$$\begin{aligned}\widetilde{rx}_{t+1} &= \widetilde{rx}_{t+1} - \mathbb{E}_t\widetilde{rx}_{t+1} + \mathbb{E}_t\Delta\hat{e}_{t+1} + \hat{r}^*_t - \hat{r}_t \\ &= \hat{e}_{t+1} - \mathbb{E}_t\hat{e}_{t+1} - d_{ss}\widetilde{\omega}^A_t,\end{aligned}$$

where the second line uses the reduced UIP condition (7). Substituting gives the law of motion

$$\widetilde{\omega}^A_{t+1} = \lambda \widetilde{\omega}^A_t + b_w \varepsilon^e_{t+1}, \tag{8}$$

in which $\lambda = \frac{1-\xi}{\beta}\left(1 - \frac{d_{ss}^2}{4\Gamma}\right)$ is the persistence of the wealth share, $b_w = \frac{(1-\xi)d_{ss}}{4\beta\Gamma}$ is its revaluation exposure to news, and $\varepsilon^e_{t+1} \equiv \hat{e}_{t+1} - \mathbb{E}_t\hat{e}_{t+1}$ is the equilibrium exchange rate innovation. Here Γ is the effective risk-cost parameter, so a larger value reduces positions per unit of wealth, ξ their exit rate, and ω^A, ω^B the steady-state wealth shares of the two groups.

Two sign conventions carry through the rest of the paper. A rise in $\hat{e}_t$ is a Home depreciation. Group-A is optimistic about foreign-currency returns and is long the Foreign bond, so a surprise depreciation revalues wealth toward the optimists and $b_w > 0$. A positive wedge holds the exchange rate above its fundamental level, and implies expected appreciation as the wedge decays[4].

The third limit adopts Remark 1: fundamental realizations follow AR(1) processes, $d^x_t = d^x_{ss}$, and the centered exchange-rate disagreement is $d_t = 0$. The constant d_{ss} still determines steady-state portfolio differences and the wealth-revaluation feedback. These restrictions are maintained throughout this section and its proofs. For each shock type, persistence and innovation variance are common across countries. The common persistence is ρ_β for discount-factor shocks and ρ_y for endowment shocks.

3.2 The disagreement–wealth multiplier

Equations (8) and (7) close the exchange rate block when combined with the household Euler equations and goods market clearing. The latter implies

$$\hat{c}_t - \hat{c}^*_t = \phi_y\left(\hat{y}_t - \hat{y}^*_t\right) - \chi\phi_e\hat{e}_t, \tag{9}$$

where $\phi_y = \frac{1}{1-2\gamma}$, $\phi_e = \frac{4\theta\gamma(1-\gamma)}{(1-2\gamma)^2}$, and $\chi = 1 - \frac{\xi_v}{\theta_v - 1}$ is the degree of incomplete pass-through. Combining the relative household Euler equation, $\hat{r}_t - \hat{r}^*_t = \sigma\left(\mathbb{E}_t\Delta\hat{c}_{t+1} - \mathbb{E}_t\Delta\hat{c}^*_{t+1}\right) - \left(\hat{\beta}_t - \hat{\beta}^*_t\right)$, with goods market clearing yields

$$\hat{r}_t - \hat{r}^*_t = -\left(\hat{\beta}_t - \hat{\beta}^*_t\right) - \sigma\phi_y(1-\rho_y)(\hat{y}_t - \hat{y}^*_t) - S\mathbb{E}_t\Delta\hat{e}_{t+1}, \tag{10}$$

where $S \equiv \sigma\chi\phi_e > 0$.

The composite $S = \sigma\chi\phi_e$ combines expenditure switching with intertemporal substitu-

[4]This limit also removes aggregate arbitrageur wealth from the state vector in Lemma 2. When the steady-state carry return is zero, a common expansion of both groups' balance sheets does not redistribute wealth at first order.

tion. An expected Home depreciation lowers future relative consumption through goods market clearing (9), and the Euler equation translates the resulting decline in expected relative consumption growth into a lower Home interest differential. Using modified UIP (7) to eliminate expected depreciation delivers the reduced system.

The equilibrium collapses to three equations in the exchange rate, the interest differential, and the optimists' wealth share.

Proposition 1. *Under Remark 1, the financial limits and cross-country shock symmetry above, stationary fundamental processes, and* $0 < \lambda < 1$*, the reduced equilibrium with* $d_t = 0$ *is characterized by:*

i. Interest rate differential: $\hat{r}_t - \hat{r}_t^* = i_t^f + \frac{S}{1+S} d_{ss} \widetilde{\omega}_t^A$, where

$$i_t^f = -\frac{1}{1+S}\left\{\hat{\beta}_t - \hat{\beta}_t^* + \sigma\phi_y(1-\rho_y)(\hat{y}_t - \hat{y}_t^*)\right\} \quad .$$

ii. Modified UIP: $\mathbb{E}_t \Delta\hat{e}_{t+1} + \hat{r}_t^* - \hat{r}_t = -d_{ss}\widetilde{\omega}_t^A$.

iii. Wealth share dynamics: $\widetilde{\omega}_{t+1}^A = \lambda\widetilde{\omega}_t^A + b_w \varepsilon_{t+1}^e$.

The system is block recursive: these equations determine exchange rates, interest differentials, and wealth shares; consumption and goods prices follow. Part (i) prices the wealth wedge into the observed interest differential, while part (ii) prices it into expected currency returns. The feedback arises in part (iii): the innovation ε_{t+1}^e that drives the wealth share is itself determined within the system.

Solving parts (i) and (ii) forward under stationarity maps the wealth share into the level of the exchange rate,

$$\hat{e}_t = -\mathbb{E}_t \sum_{j=0}^{\infty} i_{t+j}^f + \frac{d_{ss}}{(1+S)(1-\lambda)} \widetilde{\omega}_t^A, \tag{11}$$

where the frictionless fundamental component is

$$f_t \equiv -\mathbb{E}_t \sum_{j=0}^{\infty} i_{t+j}^f = \frac{1}{1+S}\left[\frac{\hat{\beta}_t - \hat{\beta}_t^*}{1-\rho_\beta} + \sigma\phi_y(\hat{y}_t - \hat{y}_t^*)\right].$$

The coefficient $\kappa_\omega \equiv d_{ss}/[(1+S)(1-\lambda)]$ is the price of one unit of wealth share. The wedge is d_{ss} per period, its expected duration is $1/(1-\lambda)$, and goods-market adjustment attenuates its exchange-rate effect by $1/(1+S)$.

Taking innovations in (11), direct fundamental news is

$$\varepsilon^f_{t+1} \equiv f_{t+1} - \mathbb{E}_t f_{t+1} = \frac{1}{1+S}\left[\frac{\varepsilon^{\beta}_{t+1} - \varepsilon^{\beta^*}_{t+1}}{1-\rho_\beta} + \sigma\phi_y(\varepsilon^{y}_{t+1} - \varepsilon^{y^*}_{t+1})\right].$$

By part (iii), the exchange-rate surprise revalues the wealth share by $b_w\varepsilon^e_{t+1}$. The same revaluation enters the exchange rate through κ_ω, giving the innovation fixed point $\varepsilon^e_{t+1} = \varepsilon^f_{t+1} + \kappa_\omega b_w \varepsilon^e_{t+1}$.

Solving it shows that disagreement multiplies every fundamental exchange rate surprise by a factor greater than one.

Corollary 1. *In the stationary benchmark of Proposition 1, with $d_{ss} > 0$, conditional on macro fundamental shocks $\hat{x}_t \in \left\{\hat{\beta}_t, \hat{\beta}^*_t, \hat{y}_t, \hat{y}^*_t\right\}$, the equilibrium exchange rate innovation ε^e_{t+1} and its counterpart in an economy without the wealth share wedge, ε^f_{t+1}, satisfy*

$$\varepsilon^e_{t+1} = \Xi\varepsilon^f_{t+1}, \qquad \Xi \equiv \frac{1}{1-g},$$

with

$$g = \frac{b_w d_{ss}}{(1+S)(1-\lambda)}.$$

Moreover,

$$1 < \Xi < \Xi_{\max} \equiv 1 + \frac{1}{S}.$$

The multiplier Ξ is strictly increasing in d_{ss}, and strictly decreasing in $\Gamma, \xi, \beta, \sigma, \chi, \theta$, and γ.

The proof is in Appendix B.5. The multiplier reflects a feedback between exchange rates and the distribution of wealth. A surprise Home depreciation enriches arbitrageurs who are optimistic about Foreign-currency returns and therefore hold larger Foreign positions. Their additional wealth gives their beliefs greater weight in equilibrium pricing, amplifying the initial depreciation. The strength of this feedback is

$$g = \underbrace{b_w}_{\text{wealth moved by news}} \quad \underbrace{\frac{d_{ss}}{1+S}}_{\text{pricing effect of that wealth}} \quad \underbrace{\frac{1}{1-\lambda}}_{\text{duration of that influence}}.$$

The first factor, b_w, measures how much an exchange-rate surprise changes the optimists' wealth share. The second measures how strongly that redistribution affects pricing: greater disagreement increases its effect, while consumption and interest-rate adjustment attenuate it through $1/(1+S)$. The third captures persistence. Because wealth shares decay at rate λ, prices capitalize the entire expected sequence of future belief wedges, $1 + \lambda + \lambda^2 + \cdots =$

$1/(1-\lambda)$. Each additional round of the feedback contributes a fraction g of the preceding round, giving

$$\Xi = 1 + g + g^2 + \cdots = \frac{1}{1-g}.$$

A Home depreciation lowers relative consumption through goods market clearing (9). The expected recovery in relative consumption raises the Home interest differential through the Euler equation (10), supporting the currency and placing an upper bound, $\Xi_{\max}$, on amplification.

Impact amplification is independent of shock volatility and persistence. It depends on preference and trade parameters, wealth persistence, and steady-state disagreement. Section 4 disciplines wealth persistence with intermediary net worth (He et al. (2017)) and d_{ss} with dispersion in the Blue Chip Financial Forecasts exchange rate survey.

Figures 1 and 2 plot the impulse responses to a positive Home discount-factor shock and a positive Home endowment shock, respectively[5]. A positive Home discount-factor shock lowers relative Home consumption and depreciates the currency. The depreciation revalues Foreign-bond positions toward optimists, amplifying their demand for Foreign assets and the initial exchange rate response. The resulting wealth wedge implies negative expected Foreign excess returns, which gradually unwind the wealth transfer. A positive Home endowment shock also depreciates the currency and activates this feedback. Its direct income effect raises Home relative consumption, while the depreciation lowers it through goods market clearing (9).

3.3 FX disconnect moments

The reduced system connects four disconnect moments: a near-random-walk exchange rate, small interest differentials relative to currency movements, a negative Fama coefficient, and a negative Backus–Smith correlation. To express these moments using a single fundamental state, I set $\rho_\beta = \rho_y \equiv \rho_x \in (0,1)$.

The fundamental component f_t in (11) then follows $f_t = \rho_x f_{t-1} + \varepsilon_t^f$, where ε_t^f is fundamental news. The fundamental interest differential is $i_t^f = \mathbb{E}_t \Delta f_{t+1} = -(1-\rho_x) f_t$. As $\rho_x \to 1$, predictable fundamental changes become small relative to fundamental news.

The first three moments emerge as fundamentals and the arbitrageurs' wealth share both become nearly permanent. In particular, the Fama coefficient, β_F, the coefficient from regressing $\Delta \hat{e}_{t+1}$ onto $\hat{r}_t - \hat{r}_t^*$, turns negative and the associated R^2 becomes arbitrarily small.

[5]Using the calibration in Section 4.1, the reduced equilibrium yields $\Xi = 1.83$ (upper bound $\Xi_{\max} = 1.95$), so wealth revaluation accounts for about 70% of exchange rate variance on impact. Additional states in the quantitative model can change this magnitude.

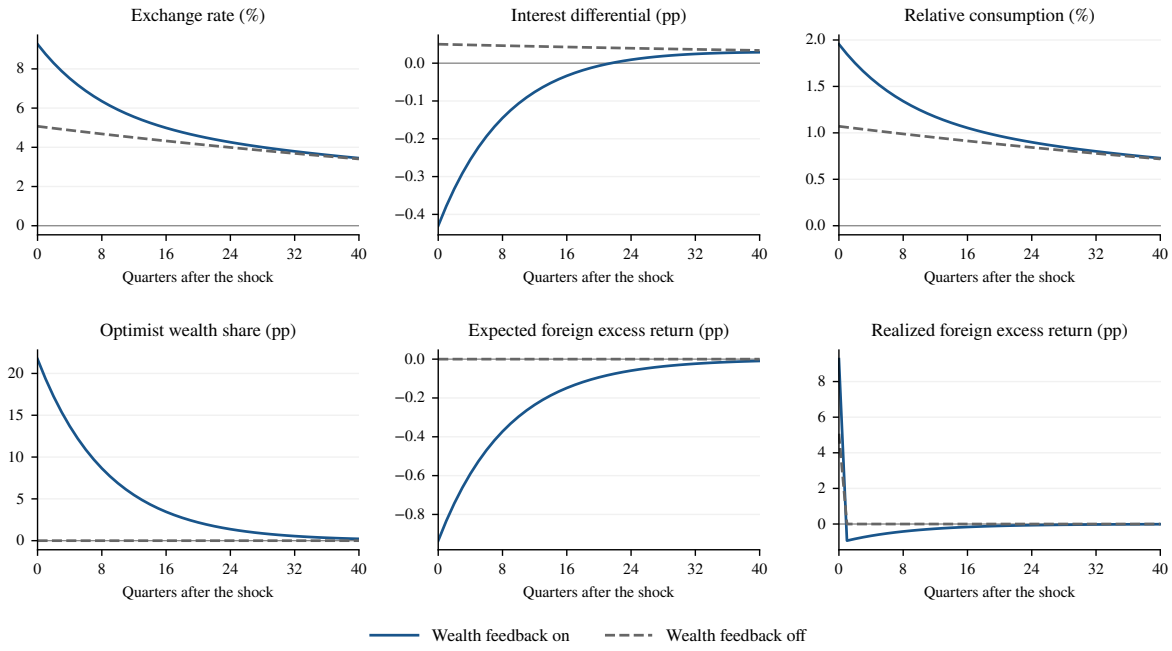


Figure 1: Reduced-equilibrium demand shock: wealth-revaluation feedback on versus off

Notes: Impulse responses in the no-NFA, no-transitory-noise reduced equilibrium to a one-standard-deviation Home demand (discount-factor) shock. The solid line retains the wealth-revaluation feedback; the dashed line shuts down only that feedback while holding fundamentals and the shock size fixed. Hence the ratio of the impact exchange-rate responses is exactly Ξ. Exchange rates and consumption are log deviations multiplied by 100; the remaining panels report percentage-point changes. Interest and consumption differentials are Foreign minus Home.

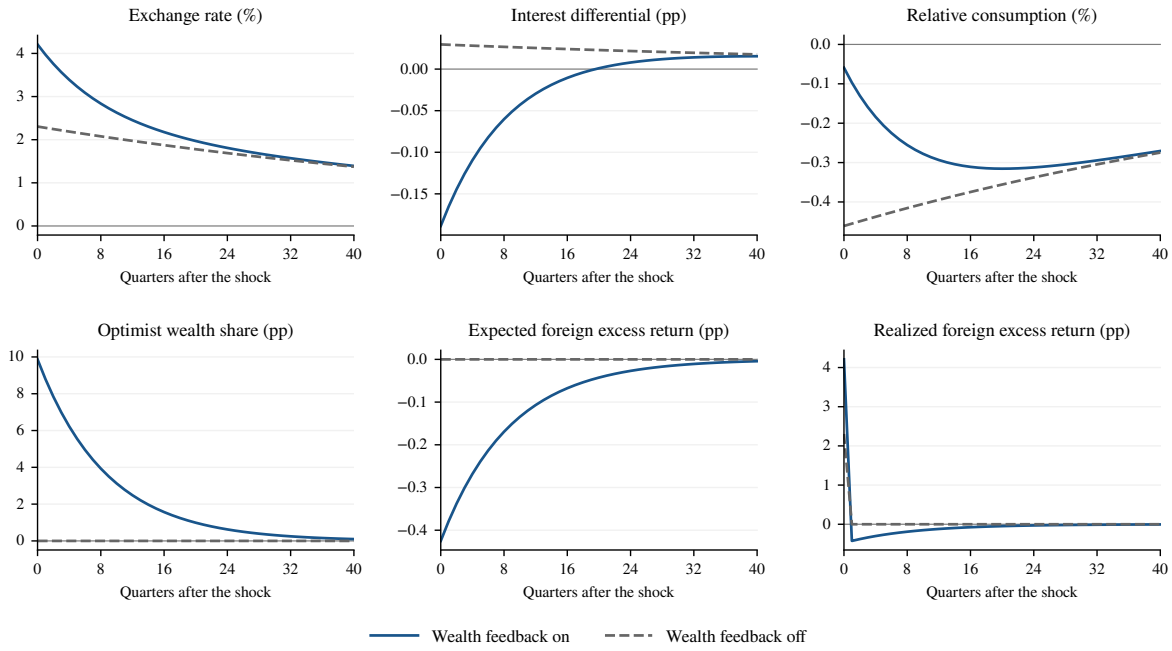


Figure 2: Reduced-equilibrium supply shock: wealth-revaluation feedback on versus off

Notes: Impulse responses in the no-NFA, no-transitory-noise reduced equilibrium to a one-standard-deviation Home endowment shock. The solid line retains the wealth-revaluation feedback; the dashed line shuts down only that feedback while holding fundamentals and the shock size fixed. Hence the ratio of the impact exchange-rate responses is exactly Ξ. Exchange rates and consumption are log deviations multiplied by 100; the remaining panels report percentage-point changes. Interest and consumption differentials are Foreign minus Home.

Proposition 2. *In the reduced equilibrium of Proposition 1, consider a sequence with* $0 < \rho_x, \lambda < 1$*, holding* $S > 0$ *and* $\Xi \in (1, \Xi_{\max})$ *fixed. As* $\rho_x \lambda \to 1$*,*

1. the exchange rate approaches a random walk;
2. the volatility of $\hat{r}_t - \hat{r}_t^*$ relative to that of $\Delta \hat{e}_t$ approaches zero; and
3. the Fama coefficient β_F turns negative provided ρ_x approaches one faster than λ, and the associated R^2 approaches zero along every nondegenerate joint path.

The proof is in Appendix B.6. The limit isolates persistence: fundamentals and wealth shares become nearly permanent at any fixed amplification factor $1 < \Xi < \Xi_{\max}$. The appendix shows how financial primitives support this limit within the constant-disagreement benchmark of Remark 1.[6]

Part 1 is the Meese and Rogoff (1983) result. The exchange rate combines the fundamental f_t and the priced wealth share, both driven by the same news. As $\rho_x, \lambda \to 1$, both components become persistent, predictable depreciation shrinks relative to its innovation, and autocorrelations of exchange rate changes vanish at every fixed horizon. Moreover, $\text{Corr}\left(\Delta \hat{e}_{t+1}, \Delta \hat{e}_t\right)$ is strictly negative along the analytical path. News that depreciates the Home currency enriches optimists, whose additional demand amplifies that depreciation. Subsequently, the shock decays and the optimists lose their wealth gradually. This generates a large initial depreciation followed by a gradual reversal toward its steady state. In other words, depreciation today tends to be followed by appreciation.

Part 2 combines weak transmission to consumption with persistent wealth redistribution. A surprise depreciation enriches optimists, whose increased demand for Foreign assets amplifies the currency movement. Home bias and incomplete pass-through limit the resulting expenditure switching, allowing a large depreciation with a modest adjustment in relative consumption. When wealth redistribution is highly persistent, households expect relative consumption to recover only gradually. The resulting expected consumption growth differential is small, so the interest-rate response to the redistribution is limited. Persistent fundamentals likewise generate lasting exchange rate responses, allowing currency movements to be large relative to interest rate movements.

Part 3 explains why a higher interest rate can predict currency appreciation. When optimists gain wealth, their increased demand for Foreign assets raises the Foreign currency's price and lowers its interest rate relative to Home. The higher Home interest rate therefore

[6]For example, hold β, Γ, and S fixed and let $d_{ss} \to 0$ and $\xi \to 1 - \beta$, keeping $(1 - \xi) d_{ss}^2 / [4\Gamma(\beta + \xi - 1)]$ fixed and positive. Then $\lambda \to 1$ while Ξ is unchanged: weaker disagreement reduces the feedback per period, while slower arbitrageurs' wealth replacement prolongs its influence. Disagreement remains constant over time within each economy, and its level changes across the sequence.

temporarily accompanies a weak Home currency. As the optimists' wealth advantage fades, their buying pressure eases and the Home currency appreciates. Under the persistence condition in Proposition 2, this predictable reversal dominates the usual fundamental channel. A higher Home interest differential then forecasts appreciation, producing a negative Fama coefficient.

Along the same joint-persistence sequence, $R^2 \to 0$ because predictable depreciation vanishes relative to its innovation. Thus the negative Fama coefficient coexists with little explained variation. Appendix B.6 provides the exact finite-parameter coefficient and regression fit.

The Backus–Smith correlation also depends on the strength of amplification.

Proposition 3. *In the reduced equilibrium of Proposition 1, with orthogonal demand and supply innovations, consider stationary parameterizations with fixed $S > 0$ and $0 < \rho_x, \lambda < 1$. As $\rho_x \lambda \to 1$,*

$$\mathrm{Corr}(\Delta(\hat{c}_t - \hat{c}_t^*), \Delta\hat{e}_t) \longrightarrow \frac{s_y - \frac{S\Xi}{1+S}}{\sqrt{\left(s_y - \frac{S\Xi}{1+S}\right)^2 + s_y\left(1 - s_y\right)}}.$$

where $s_y = \mathrm{Var}(\varepsilon_{t+1}^{f,y})/\sigma_f^2 \in (0,1)$ is the supply share of direct fundamental-news variance. Additionally, if $\Xi \to \Xi_{\max}\left(= 1 + \frac{1}{S}\right)$ together with $\rho_x \lambda \to 1$, it converges

$$\mathrm{Corr}(\Delta(\hat{c}_t - \hat{c}_t^*), \Delta\hat{e}_t) \longrightarrow -\sqrt{1 - s_y}.$$

The first limit reflects a balance between the income effect of supply shocks and expenditure switching induced by currency movements. A Home supply shock raises income and depreciates the currency, with opposing effects on relative consumption through goods market clearing (9). A demand shock that depreciates the Home currency also lowers relative consumption. Disagreement amplifies currency movements and the associated expenditure switching while leaving the direct income effect unchanged, strengthening negative comovement.

The correlation turns negative when $s_y < S\Xi/(1+S)$: demand news and expenditure switching outweigh the supply income effect. Amplification therefore allows negative comovement with a larger supply share. Without disagreement, a sufficiently large demand share can still make the correlation negative. Supply news alone produces positive impact comovement, although later wealth adjustment can generate a negative unconditional correlation at finite persistence (Appendix C.2).[7]

In the second limit, amplification makes the depreciation induced by supply news almost

[7] Section 4 assesses these forces at the reported calibration.

exactly offset its positive income effect on relative consumption. Consumption growth is then driven mainly by demand news, while exchange rates continue to respond to both demand and supply news. Any positive demand share therefore gives a negative correlation; as demand news dominates, the correlation approaches minus one. This cancellation also raises exchange rate volatility relative to consumption growth. Home bias and incomplete pass-through further limit the consumption adjustment required for a given currency movement.

Persistent fundamentals and wealth shares generate weak exchange rate predictability and small interest differentials, while the wealth feedback produces the negative Fama and Backus–Smith correlations. These results arise without an exogenous UIP shock. Section 4 evaluates their quantitative magnitudes and the required shock mix using survey disagreement, intermediary net worth, and the baseline and no-disagreement comparisons.

4 Quantitative Analysis

The baseline model reproduces the main exchange-rate disconnect moments with disagreement and conventional demand and supply shocks. It features rational consensus beliefs about fundamentals, but individual beliefs remain dispersed and agents infer persistent states from noisy observations. The augmented model adds average belief distortions and a reduced-form UIP shock. I first present the baseline calibration and results, then examine what these additions contribute. The central comparison removes disagreement: even after re-estimating fundamental shock processes, the resulting economy generates too little exchange-rate volatility and a positive Fama coefficient.

4.1 Baseline specification and calibration

The baseline sets $\bar{\mu}_t^x = 0$ for every fundamental and period. Population-average beliefs therefore follow the correctly specified Bayesian filter, while the two groups continue to disagree. Agents still infer latent states from noisy observations, and individual forecasts can differ systematically from the Bayesian forecast. Differences in wealth can still give greater pricing influence to one group's beliefs.

Time-varying disagreement affects both positions and their influence on prices. A wider forecast gap moves the groups' currency positions further apart, so a given exchange-rate surprise redistributes more wealth. The resulting redistribution also shifts the wealth-weighted forecast more strongly. The quantitative baseline retains the products $d_t\widetilde{\omega}_t^A$ in bond pricing and $d_t\widetilde{R}_{t+1}^{\text{ex}}$ in wealth accumulation. Perceived risk is fixed, and forecasts use the affine perceived transition described in Appendix E. These interactions generate state-dependent

volatility without adding parameters.

The calibration combines external parameter choices with joint matching of macroeconomic, yield, survey, and intermediary moments. I set openness to $\gamma = 0.05$, household risk aversion to $\sigma = 5$, and the trade elasticity to $\theta = 1.5$, following Chahrour et al. (2024); Candian and De Leo (2025); Itskhoki and Mukhin (2021). Exchange-rate pass-through to relative consumer prices is $\chi = 0.60$, following Kekre and Lenel (2024, 2026b), and the two belief groups have equal population shares. The discount factors $\beta = 0.9980$ and $\beta^* = 0.9975$ match average real rates, making the dollar the steady-state funding currency.

Table 1 reports the baseline parameters and their principal empirical disciplines. Data cover 1991Q2–2020Q1. Home is the United States; Foreign combines Australia, Canada, Switzerland, the euro area, the United Kingdom, Japan, and Sweden using the trade weights underlying the Fed's advanced-economy dollar index. Following Kekre and Lenel (2024), output is nominal GDP less investment and government spending, divided by the consumption deflator and population. Long-yield dynamics discipline discount-factor shocks. Relative-output persistence and growth volatility discipline endowment shocks; cross-country yield and consumption comovement help determine their international correlations. Exchange-rate moments also enter calibration, so the level of fit is not an independent test of the mechanism.

Broker-dealer equity relative to GDP disciplines arbitrageur wealth W and the exit rate ξ. I convert the capital ratio in He et al. (2017) into the model's wealth-to-output units using broker-dealer assets relative to quarterly GDP. The intermediation friction Γ is calibrated jointly with exchange-rate volatility: a lower value lets arbitrageurs take larger positions for a given expected return. The baseline reproduces intermediary equity persistence reasonably closely, at 0.924 against 0.922 in the data.

Blue Chip Financial Forecasts disciplines disagreement. The survey reports institution-level forecasts of a broad dollar index one quarter ahead. I measure the forecast gap as the spread between the most optimistic and most pessimistic forecasters' expected log dollar changes, using a common current exchange rate. Its mean, persistence, and innovation volatility discipline d_{ss}, the common Kalman gain G, and the common interpretation-innovation scale σ_μ. These moments anchor the strength and variation of disagreement; they are not chosen solely to fit exchange rates.

4.2 Baseline exchange-rate moments

Table 2 compares the baseline with two economies without disagreement. Both set $d_{ss} = 0$ and eliminate disagreement innovations, recomputing the steady state. The first holds the

Table 1: Baseline calibration and empirical disciplines

Parameter	Value	Principal empirical discipline	Data	Model
W	0.0707	Mean intermediary equity/GDP	0.0726	0.0707
ξ	9.61×10^{-3}	AR(1), intermediary equity/GDP	0.922	0.924
Γ	3.98×10^{-3}	SD, quarterly depreciation (%)	3.33	3.44
d_{ss}	0.0431	Mean forecast gap (%)	4.40	4.31
G	0.309	AR(1), forecast gap	0.648	0.660
σ_μ	1.78×10^{-3}	Innovation SD, forecast gap (%)	0.56	0.58
ρ_β	0.9901	AR(1), 10-year yield differential	0.940	0.981
$\sigma_{\varepsilon^\beta}$	1.03×10^{-3}	SD, change in annualized 10-year yield differential (%)	0.27	0.16
$\rho_{\beta\beta^*}$	0.393	Correlation, Home and Foreign 10-year yield changes	0.810	0.872
ρ_y	0.9871	AR(1), relative output	0.945	0.944
σ_{ε^y}	8.53×10^{-3}	SD, relative output growth (%)	0.77	0.82
ρ_{yy^*}	0.536	Correlation, Home and Foreign consumption growth	0.629	0.584

Notes: The baseline retains disagreement and endogenous wealth, sets average interpretation bias to zero, and has no exogenous UIP shock. Home is the United States; Foreign is its trade-weighted advanced-economy counterpart. Data cover 1991Q2–2020Q1. AR(1) summarizes first-order persistence and SD is standard deviation. The forecast gap is the difference between the two groups' expected next-quarter log exchange-rate changes, measured by Blue Chip survey dispersion. Relative output is Home minus Foreign log output. Yields have a 40-quarter maturity and are expressed in quarterly units unless explicitly annualized. W is steady-state arbitrageur wealth relative to output, ξ the exit probability, and Γ the effective risk-cost parameter. d_{ss} is steady-state exchange-rate disagreement, G the common Kalman gain, and σ_μ the common interpretation-innovation scale. For $x \in \{\beta, y\}$, ρ_x is persistence, σ_{ε^x} the persistent-innovation SD, and ρ_{xx^*} its cross-country correlation. Rows show principal disciplines within a joint calibration, rather than separate identification claims. Model moments average 1,000 simulations of 116 quarters after a 1,000-quarter burn-in. Parameters are rounded only for display.

other baseline parameters fixed. The second re-estimates the fundamental shock block, allowing conventional shocks to do more of the work. This comparison tests whether those shocks can generate the same joint pattern without the wealth–disagreement channel.

Table 2: Baseline exchange-rate moments and the role of disagreement

Moment	Data	Baseline	No disag. Fixed	No disag. Recal.
Autocorrelation, exchange-rate changes	0.013	-0.050	-0.009	-0.003
Autocorrelation, exchange-rate level	0.950	0.908	0.930	0.926
SD, quarterly depreciation (%)	3.33	3.44	1.31	1.96
FX/consumption-growth volatility	6.43	4.11	1.58	2.95
FX/output-growth volatility	5.42	4.03	1.54	2.94
Short-rate differential/FX volatility	0.110	0.166	0.247	0.309
Long-yield differential/FX volatility	0.060	0.074	0.240	0.295
Backus–Smith correlation	-0.169	-0.244	0.546	-0.131
Fama coefficient	-0.213	-0.480	1.023	0.906
Fama R^2 (%)	0.05	1.35	4.31	5.28

Notes: The baseline has rational consensus beliefs about fundamentals, dispersed individual beliefs, and no exogenous UIP shock. Both counterfactuals remove the permanent belief gap and disagreement innovations. Fixed retains the other baseline parameters; Recal. re-estimates the fundamental shock block. $\hat{e}_t$ is the log real exchange rate; an increase is Home depreciation. Δ denotes a quarterly change. FX volatility is $\mathrm{SD}(\Delta\hat{e}_t)$; consumption and output growth refer to Home. Rate ratios use the SD of the level of the Home-minus-Foreign yield differential, in quarterly units, divided by FX volatility; short and long maturities are one and 40 quarters. Backus–Smith is $\mathrm{Corr}(\Delta\hat{e}_t, \Delta(\hat{c}_t - \hat{c}^*_t))$. The Fama regression is $\Delta\hat{e}_{t+1} = \alpha + \beta^{\mathrm{Fama}}(\hat{r}_t - \hat{r}^*_t) + u_{t+1}$; its R^2 is reported in percent. Data are U.S.–advanced-foreign-economy moments over 1991Q2–2020Q1. Model entries average 1,000 simulations of 116 quarters after a 1,000-quarter burn-in.

Volatility and persistence The baseline generates persistent exchange-rate levels and little autocorrelation in quarterly changes. Removing disagreement preserves these features because fundamental shocks remain persistent. What changes sharply is the size of currency movements: quarterly depreciation volatility falls from 3.44 percent in the baseline to 1.31 percent, compared with 3.33 percent in the data. Re-estimating fundamental shocks raises it to 1.96 percent but does not recover the observed scale.

The relative volatilities show the same mechanism. Exchange-rate changes are 4.11 times as volatile as Home consumption growth in the baseline, against 6.43 in the data; without disagreement, the ratio falls to 1.58, or 2.95 after recalibration. Wealth redistribution amplifies asset demand and currency movements, while home bias and incomplete pass-through limit their transmission to consumption. The short-rate-differential/FX volatility ratio is consequently much smaller in the baseline: 0.166, against 0.247 and 0.309 without disagreement and 0.110 in the data. The model remains short of the observed separation between

exchange rates and quantities, but disagreement substantially narrows the gap.

Currency-return predictability The baseline Fama coefficient is −0.480, compared with −0.213 in the data. When wealth shifts toward investors optimistic about Foreign returns, their asset demand weakens the Home currency and raises the Home interest differential. The expected reversal of that redistribution then supports Home appreciation. A high interest differential therefore forecasts appreciation rather than the depreciation implied by UIP. Removing disagreement changes the coefficient to 1.023, or 0.906 after recalibration. Persistent fundamentals can generate persistent currencies, but they do not generate this reversal in return predictability.

Consumption comovement The baseline Backus–Smith correlation is −0.244, against −0.169 in the data. Removing disagreement at fixed fundamental shock parameters makes it positive, at 0.546. Disagreement strengthens demand-induced currency movements relative to supply's positive income effect, allowing negative comovement to dominate. Re-estimating the shock mix can recover a negative correlation without disagreement, but then the economy still misses volatility and the Fama sign. The mechanism's contribution is to reproduce these features together.

Supply-only correlations remain positive at the reported calibration. The possibility of a negative supply-only unconditional correlation established in Appendix C.2 therefore does not drive the quantitative fit. Appendix C isolates the demand and supply channels and shows that permanent disagreement alone can reverse the correlation at unchanged fundamental shock parameters.

Sources of exchange-rate variation Demand shocks account for 62–67 percent of baseline exchange-rate forecast-error variance across horizons from one to 40 quarters; supply shocks account for most of the remainder (Figure 3a). Signal-noise and disagreement innovations together contribute less than 3 percent. This small direct contribution does not measure the importance of disagreement: fundamental news changes currency values, redistributes wealth across opposing positions, and changes whose beliefs influence prices. Disagreement primarily changes the propagation of macroeconomic news.

4.3 The augmented model

The augmented model adds two wedges to the baseline: average belief distortions and an exogenous UIP shock. These additions allow the wealth–disagreement mechanism to be evaluated alongside two forces emphasized in recent international macroeconomics and finance

research (Du et al. (2021); Stavrakeva and Tang (2024); Candian and De Leo (2025)). The model retains the same financial interactions but is calibrated separately (Appendix Table D1). Comparisons across calibrations assess overall fit. To isolate each wedge's contribution, I also remove it while holding the remaining parameters at their augmented calibration values.

Average belief distortions Following Candian and De Leo (2025), I allow consensus beliefs to update with perceived gain κ_x and persistence $\widetilde{\rho}_x$. The interpretation bias is

$$\bar{\mu}_t^x = \left(\frac{\kappa_x}{G_x} - 1\right)\hat{x}_t + \frac{(1-\kappa_x)\tilde{\rho}_x - \rho_x(1-G_x)}{G_x}\bar{\tau}^x_{t-1|t-1}.$$

Using the same perceived persistence for forecasts reproduces their conditional-mean expectations at every horizon (Appendix D.2). The calibrated pairs $(\kappa_y, \widetilde{\rho}_y) = (0.259, 0.9995)$ and $(\kappa_\beta, \widetilde{\rho}_\beta) = (0.969, 0.9924)$ target the horizon profile of currency-return predictability. Forecast-error regressions following Coibion and Gorodnichenko (2015) and Kohlhas and Walther (2021) serve as external diagnostics. These distortions add no primitive shocks.

Exogenous UIP shock The shock ψ_t enters the modified UIP condition (5) and captures movements in risk-bearing capacity beyond those generated by endogenous wealth. Following Kekre and Lenel (2024), I set its AR(1) persistence to $\rho_\psi = 0.85$ and use the excess bond premium of Gilchrist and Zakrajšek (2012) as its empirical proxy. Its innovation scale, $\sigma_\psi = 0.0138$, is calibrated to exchange-rate level volatility. The added shock permits a direct comparison with explanations based on exogenous financial disturbances.

What the additions change Table 3 separates their effects at the augmented calibration. Removing both wedges still yields a negative Fama coefficient, −0.576. This differs from the baseline's −0.480 because the remaining parameters have not been recalibrated. The endogenous wealth–disagreement channel therefore produces the Fama sign in both calibrations, without an exogenous UIP shock.

Average belief distortions mainly change interest rates and predictability across horizons. Unlike disagreement around a rational consensus, a common forecast distortion survives population averaging and enters household Euler equations. The augmented model's short-rate-differential/FX volatility ratio rises to 0.513, compared with 0.166 in the baseline; removing average bias at the augmented calibration lowers it to 0.374. Removing the UIP shock lowers it only to 0.473. The added wedges thus worsen this aspect of the fit, even though long-yield volatility remains close to the data.

Table 3: Augmented model and removal of the additional wedges

Moment	Augmented	No UIP	No avg. bias	Neither
Autocorrelation, exchange-rate level	0.917	0.922	0.900	0.905
SD, quarterly depreciation (%)	4.03	3.64	3.24	2.75
FX/consumption-growth volatility	4.46	4.11	3.58	3.09
Short-rate differential/FX volatility	0.513	0.473	0.374	0.139
Long-yield differential/FX volatility	0.061	0.062	0.063	0.064
Backus–Smith correlation	-0.471	-0.320	-0.327	-0.071
Fama coefficient	-0.449	-0.208	-0.950	-0.576
40-quarter cumulative-return slope	0.36	1.19	-2.10	-7.75

Notes: All columns use the augmented calibration and retain disagreement and both financial interaction terms. No UIP sets exogenous UIP innovations to zero; No avg. bias restores Bayesian consensus beliefs about fundamentals; Neither applies both restrictions. Other parameters are held fixed, so Neither differs from the separately calibrated baseline. FX denotes the real exchange rate, rising with Home depreciation. FX volatility is the SD of quarterly log depreciation; consumption refers to Home. Rate ratios use quarterly-equivalent Home-minus-Foreign yield levels, at one- and 40-quarter maturities, divided by FX volatility. Backus–Smith is the correlation between depreciation and Home-minus-Foreign consumption growth. The Fama slope regresses next-quarter depreciation on today's Home-minus-Foreign short-rate differential. The cumulative-return slope regresses the sum of foreign-currency excess returns over the next 40 quarters on that same differential. Entries average 1,000 simulations of 116 quarters after a 1,000-quarter burn-in. Data benchmarks and the baseline appear in Table 2.

The augmented model also generates the reversal in future excess-return predictability documented by Engel (2016). The coefficient on today's Home-minus-Foreign interest differential starts at −1.45 for next-quarter Foreign excess returns, turns positive at six quarters, and reaches 0.22 at twelve quarters. Consensus beliefs initially underreact to news and then extrapolate it too persistently, changing both interest differentials and subsequent currency adjustment. Removing average bias delays the sign change to eight quarters. Its contribution is to the timing and cumulative sign: the 40-quarter cumulative-return slope is 0.36 in the augmented model, against −6.52 in the baseline. This is a qualitative improvement; it remains far below 20.65 in the data. Appendix D reports the profiles and fixed-parameter comparisons.

The exogenous UIP shock is a secondary source of currency variation. In the augmented model, demand and supply shocks account for 34 and 45 percent of one-quarter variance, respectively; their combined share rises to 87 percent at 40 quarters. The UIP shock's share falls from 18 to 13 percent (Figure 3b). The precise demand–supply split depends on the calibration. Supply shocks also operate through a demand channel: when agents overestimate the persistence of supply news, their spending responds to exaggerated expectations of future income, even though the underlying endowment process is unchanged. Across both models, conventional macroeconomic news drives most exchange-rate variation, while the exogenous

UIP shock plays a secondary role, consistent with the emphasis in Kekre and Lenel (2024).

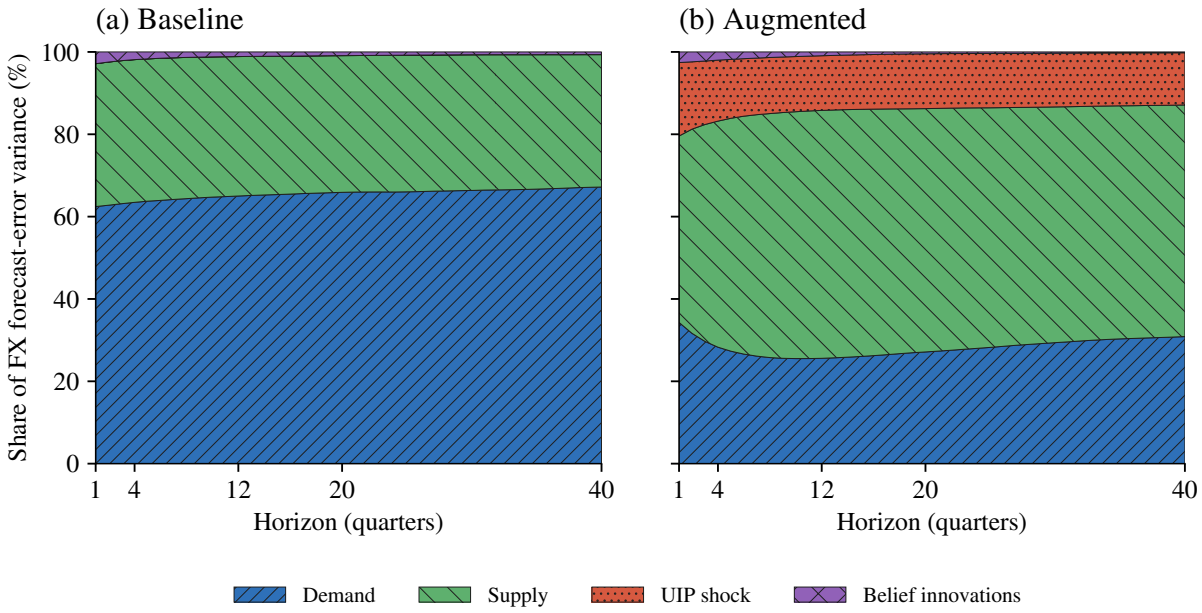


Figure 3: Exchange-rate variance decomposition: baseline and augmented models

Notes: Shapley shares of the forecast-error variance of cumulative log real exchange-rate changes at horizons of one to 40 quarters. Panel (a) uses the baseline calibration, with disagreement and Bayesian consensus beliefs but no exogenous UIP shock. Panel (b) uses the separately calibrated augmented model, which adds average belief distortions and UIP risk. Demand and supply refer to discount-factor and endowment innovations. Belief innovations combine transitory signal noise and interpretation innovations. Shares measure the origin of news, including its propagation through endogenous positions and wealth; they sum to 100 percent within each panel.

4.4 Historical comparison at the augmented calibration

The historical exercise uses the first-order augmented model, initially setting the exogenous UIP shock to zero. It retains average belief distortions and the augmented calibration to make the comparisons transparent. I infer two country demand shocks, two output shocks, and a discount-factor disagreement shock from long yields, detrended output per capita, and Blue Chip disagreement. Given that exchange rates are excluded from the input data, it is noteworthy that the simulated series tracks the overall trend and cyclical patterns fairly well. A representative-agent comparison removes disagreement, recomputes the steady state, and re-infers its four fundamental shocks from the same macroeconomic observations. Appendix F.3 gives the measurement and initialization details.

Without the UIP shock, the disagreement model generates 1.00 times the observed level standard deviation, compared with 0.40 for the representative-agent economy; their level

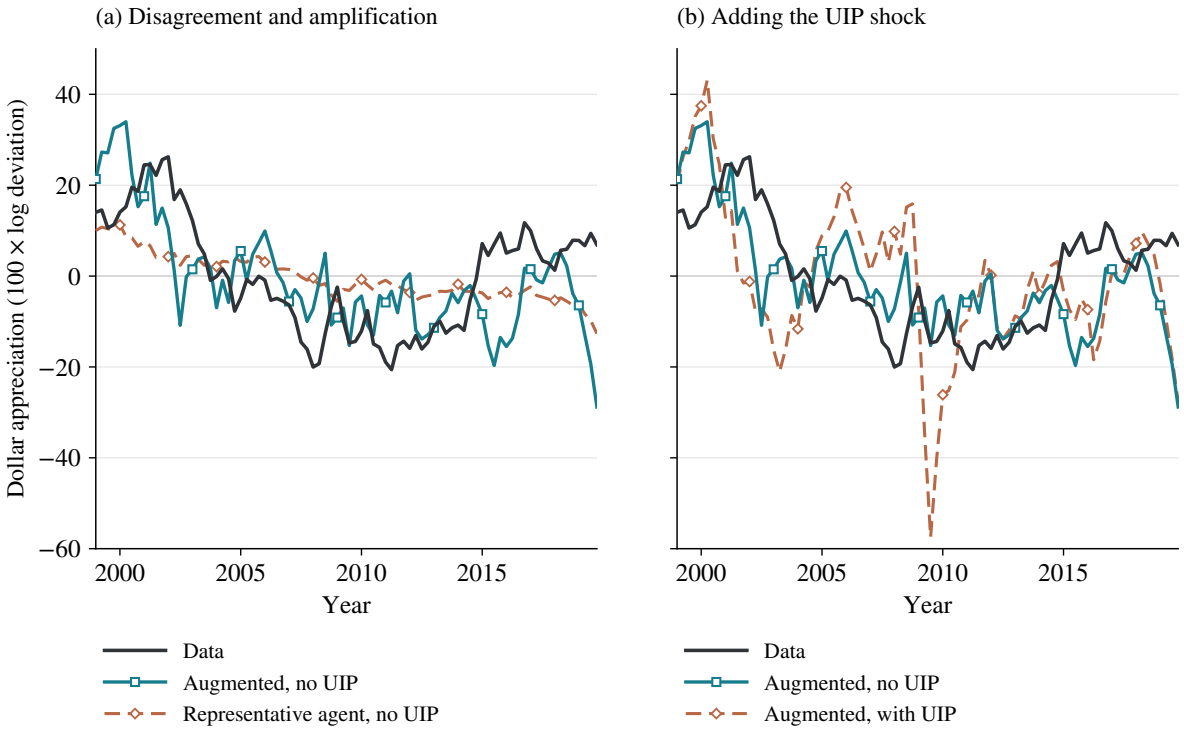


Figure 4: Historical exchange-rate inversion at the augmented calibration, 1999Q1–2019Q4

Notes: The nominal dollar index rises with dollar appreciation; plotted values are 100 times log deviations. Panel (a) compares the augmented model without UIP shocks with its representative-agent counterpart. Panel (b) adds the observed excess-bond-premium UIP path at the augmented calibration's scale and re-infers the remaining shocks. All use the same macroeconomic observations; the disagreement models additionally match survey dispersion. Observed relative consumer prices convert model real exchange rates into nominal paths. Mean alignment is fixed over 1999Q1–2019Q2, with no smoothing or volatility rescaling.

correlations are 0.50 and 0.43. The main gain is the scale of currency movements generated from macroeconomic observations. Output shocks account for about 43 percent of the model's real level variance, while direct disagreement innovations contribute little. This is consistent with disagreement amplifying fundamental news through positions and wealth rather than supplying most of the news itself.

Adding the excess-bond-premium UIP path reduces the level correlation to 0.32 and raises root-mean-squared error from 12.4 to 17.3 log percentage points. These results reflect a tradeoff in the calibration: the parameters jointly discipline macroeconomic, interest-rate, survey, and intermediary moments, and improving their overall fit requires some compromise on exchange-rate volatility. The UIP shock's scale is therefore constrained by the joint calibration and is not adjusted separately to improve the historical exchange-rate fit. At the reported parameters, the excess-bond-premium path adds currency movements that align poorly with observed exchange rates. Its inclusion consequently worsens the historical fit, although this result is specific to the proxy and calibration used here. Appendix F.3 reports the extended sample and a restricted demand-shock inversion; Appendix F.3.4 presents the complementary accounting exercise that conditions on realized survey news.

5 Empirical Supporting Evidence

The supporting evidence examines two implications of the mechanism: survey beliefs should be associated with intermediaries' currency positions, and disagreement should contain information about subsequent currency volatility. I use bilateral data for the advanced-foreign-economy (AFE) currencies represented in the quantitative analysis.[8]

5.1 Data and measurement

Currency coverage and quotation The AFE currency universe is AUD, CAD, CHF, EUR, GBP, JPY, and SEK against the U.S. dollar. The volatility exercise and the appendix exchange-rate regressions use this seven-currency universe. Dealer-position regressions use its six-currency subset with matched CFTC data: AUD, CAD, CHF, EUR, GBP, and JPY. NOK and NZD are excluded. The regressions use bilateral observations, with currency effects where observations are pooled; they do not construct a trade-weighted dollar index.

Let $e^N_{i,t}$ be the nominal U.S.-dollar price of one unit of currency i. An increase is dollar depreciation, following the direction of the model's real exchange rate e_t. I convert both

[8] Appendix G.1 also asks whether disagreement adds explanatory information to exchange-rate regressions with macroeconomic and financial controls.

spot rates and individual forecasts to this quotation before constructing log variables. For CAD, CHF, JPY, and SEK, conversion reverses the sign of the native log quotation. The empirical spot rates and survey forecasts are nominal; the model exchange rate is real. When a real exchange-rate control is needed, I construct $\log e_{i,t} = \log e^N_{i,t} + \log P_{i,t} - \log P_{US,t}$ using observed consumer price indexes.

Bloomberg forecasts Bloomberg reports individual exchange-rate level forecasts for calendar quarter-end targets from 2006 onward. I construct 90-, 180-, 270-, and 360-day forecasts by interpolating in logs between adjacent future quarter-end targets that bracket the required date. Each interpolation uses only forecasts reported by the same institution on the same survey date, with no extrapolation beyond available targets. I then average the resulting log forecasts within institution and calendar month. The consensus $\bar{f}^h_{i,t}$ is the equally weighted mean of these institution-level log forecasts. Disagreement $\mathcal{D}^h_{i,t}$ is their cross-sectional standard deviation. I require at least eight institutions for a currency–month–horizon observation. Thus dispersion is measured across log forecasts, rather than approximated by a standard deviation of raw quotes divided by their mean or spot rate. These survey statistics proxy log exchange-rate beliefs; they do not directly observe the model's belief groups or their wealth weights.

Spot rates are averaged in logs within each calendar month. The dealer regression uses monthly changes in $\bar{f}^{90}_{i,t}$ and $\log e^N_{i,t}$. Because the 90-day target rolls forward and the respondent set can change, a consensus revision is not a forecast revision for a fixed future date or a fixed cohort. The monthly appendix exercise uses the four-horizon term structure of $\log \mathcal{D}^h_{i,t}$ and macroeconomic controls available through February 2021; complete-case estimates end in January 2021. The volatility exercise uses the longer survey and spot-rate archive through 2024.

CFTC dealer positions I use weekly dealer long-minus-short positions from the CFTC Traders in Financial Futures data, retaining only the direct U.S.-dollar currency contracts. In particular, euro–yen and euro–sterling contracts are excluded from the EUR series. I average net positions and total open interest within each calendar month, then divide monthly net positions by the trailing 12-month mean of monthly open interest, requiring all twelve months. The resulting position $n^D_{i,t}$ is positive when dealers are net long the foreign currency. It can be negative and is therefore not logged.

Quarterly disagreement and realized volatility Quarterly disagreement $\text{Disp}_{i,t}$ is the average of $\log \mathcal{D}^{90}_{i,m}$ over available survey months m in quarter t. I do not fill missing surveys.

Daily log spot changes determine realized volatility: $RV_{i,t} = (\sum_{d \in t} (\Delta \log e^N_{i,d})^2)^{1/2}$. A month must contain at least 15 valid daily changes, and a quarter must contain all three valid months. Realized volatility is unchanged if the exchange-rate quotation is inverted. All changes, lags, and outcome leads use actual calendar periods, so missing surveys do not shorten the forecast horizon.

5.2 Survey forecasts and dealer positions

The model's portfolio rule predicts greater demand for Foreign assets when expected future Foreign-currency prices rise, holding current prices fixed. I estimate a standardized version of the corresponding dynamic relationship in dealer positions:

$$\begin{aligned} \Delta n^D_{i,t} = \alpha_i + \alpha_t + \varrho\, \Delta n^D_{i,t-1} + \theta_0\, \Delta \bar{f}^{90}_{i,t} + \theta_1\, \Delta \bar{f}^{90}_{i,t-1} \\ + \gamma_0\, \Delta \log e^N_{i,t} + \gamma_1\, \Delta \log e^N_{i,t-1} + u_{i,t}. \end{aligned} \tag{12}$$

All six displayed regression variables are standardized within currency on the complete estimation sample. Currency and month effects absorb persistent currency differences and common monthly movements. Pooled inference uses calendar-month score aggregation with Bartlett HAC standard errors and six lags; currency-specific regressions use the corresponding Newey–West estimator. The consensus is measured across Bloomberg respondents, who need not be the dealers represented in the CFTC data, so this is a test of whether measured beliefs and dealer positions move together.

Table 4 reports the pooled estimates. The sum of the current and lagged consensus-revision coefficients is 0.391 (standard error 0.092). The contemporaneous depreciation coefficient is -0.725. A positive forecast-revision coefficient and a negative spot coefficient have the signs of the model's portfolio rule: a higher expected future Foreign-currency price raises desired Foreign exposure, while a higher current price lowers its expected return. The regression does not separately identify this portfolio channel from dealers' inventory management or other common determinants of forecasts and positions.

The currency-specific estimates in Table 5 have positive sums of forecast-revision coefficients in six of the six currencies, with four significant at the unadjusted five-percent level. The separate regressions allow the strength of the forecast–position relationship to differ across currencies. The magnitudes differ substantially across currencies.

5.3 Disagreement and conditional exchange-rate volatility

The financial interactions in the quantitative model allow disagreement to affect subsequent conditional volatility. A wider forecast gap raises the groups' opposing currency exposures,

Table 4: Survey consensus revisions and dealer positions

Regressor	Standardized coefficient
Lagged position change $\Delta n^D_{i,t-1}$	0.101^{**} (0.044)
Current consensus revision $\Delta \bar{f}^{90}_{i,t}$	0.230^{***} (0.076)
Lagged consensus revision $\Delta \bar{f}^{90}_{i,t-1}$	0.161^{***} (0.046)
Current depreciation $\Delta \log e^N_{i,t}$	-0.725^{***} (0.126)
Lagged depreciation $\Delta \log e^N_{i,t-1}$	-0.196^{**} (0.089)
Observations	1,242
Within R^2	0.345
Joint p-value: consensus revisions	< 0.001

Notes: The dependent variable is the monthly change in n^D, where n^D is dealer long-minus-short contracts divided by trailing 12-month mean open interest. The six AFE currencies are AUD, CAD, CHF, EUR, GBP, and JPY; only direct USD futures contracts are included. The sample spans 2007-07–2024-09, with missing observations omitted. Consensus revisions are changes in the mean of institution-level log 90-day forecasts; depreciation is the change in the monthly mean log USD price of Foreign currency. Variables are standardized separately within each currency on the estimation sample. Currency and month effects are included. Within R^2 measures fit after removing both effects. Parentheses contain Bartlett HAC standard errors with six calendar-month lags, scores summed across currencies by month, and an $n/(n - \text{rank})$ correction; tests use t_{T-1} or the corresponding joint Wald F statistic, where T is the number of months. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Dealer-position regressions by AFE currency

Currency	Consensus revision $\theta_0 + \theta_1$	Depreciation γ_0	Lagged change ϱ	Within R^2	N
AUD	0.211 (0.239)	-0.566^{***} (0.181)	0.122^{**} (0.060)	0.338	207
CAD	0.362^{***} (0.121)	-0.580^{***} (0.115)	0.098 (0.081)	0.393	207
CHF	0.458^{***} (0.109)	-0.685^{***} (0.108)	0.022 (0.051)	0.458	207
EUR	0.225^{*} (0.128)	-0.531^{***} (0.101)	0.130 (0.085)	0.474	207
GBP	0.316^{***} (0.085)	-0.684^{***} (0.110)	0.200^{***} (0.063)	0.394	207
JPY	0.383^{***} (0.123)	-0.765^{***} (0.102)	0.107 (0.085)	0.467	207

Notes: The dependent variable is the within-currency standardized monthly change in n^D, where n^D is dealer long-minus-short contracts divided by a complete trailing 12-month mean of open interest. Each regression includes current and lagged changes in the mean log 90-day Bloomberg forecast, current and lagged monthly log USD/Foreign spot changes, and a lagged position change. All regression variables are standardized on the currency's estimation sample. Positive spot changes are dollar depreciation; positive positions are long Foreign currency. An intercept is included. Consensus revision reports the sum of the two forecast coefficients; Lagged change reports the autoregressive coefficient. Parentheses contain Newey–West standard errors with six calendar-month lags, a finite-sample correction, and t_{T-1} tests. The sample is the direct-USD-contract subset of the AFE basket; periods span 2007-07–2024-09. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

making wealth redistribution more sensitive to exchange-rate surprises. I assess this prediction using future realized volatility, which is not a calibration target.

For horizons $h = 1, \ldots, 12$ quarters, I estimate the panel local projection

$$\log RV_{i,t+h} = \alpha_i^h + \alpha_t^h + b_h \text{Disp}_{i,t} + \boldsymbol{\gamma}_h' X_{i,t} + u_{i,t+h}.$$

The controls $X_{i,t}$ include current log realized volatility, two quarterly lags of log realized volatility and disagreement, current and lagged changes in log spot rates, and the quarterly mean log forecaster count. Currency and quarter effects absorb persistent currency differences and common quarterly movements. I scale b_h by the residual standard deviation of disagreement on the full predictor sample, held fixed across horizons: $s_{\text{Disp}} = 0.146$. The reported effect is $100[\exp(b_h s_{\text{Disp}}) - 1]$. Inference uses Bartlett HAC with four quarterly lags on scores summed across currencies by calendar quarter and t_{T-1} critical values. The following-year outcome uses eight lags to allow for overlapping windows.

Figure 5 compares the augmented-model response with the AFE estimates. In the model, a one-standard-deviation disagreement shock raises next-quarter conditional volatility by 9.05 percent, declining to 2.57 percent at four quarters. In the data, a one-residual-standard-deviation increase in disagreement predicts a 2.41 percent change in next-quarter realized volatility, with a 95 percent pointwise interval of [1.16, 3.67]. This regression uses 497 currency–quarter observations and predictor quarters 2006Q4–2024Q2. The two panels use different normalizations, so they assess the direction and persistence of the relationship without establishing a quantitative match.

I also measure risk over the entire following year using $RV_{i,t}^{\text{year}} = (\sum_{j=1}^{4} RV_{i,t+j}^2)^{1/2}$. The estimated effect is 1.74 percent, with a 95 percent interval of [0.97, 2.52], using 476 observations and predictor quarters 2006Q4–2023Q3. The following-year estimate is 2.06 percent using the interquartile range of log forecasts and 1.55 percent after excluding windows overlapping the global financial crisis or the onset of COVID-19. Omitting one currency at a time gives point estimates from 1.08 to 2.18 percent.

Quarter effects make these conditional comparisons across currencies within the same quarter. The estimates need not describe a common time-series response across all currencies, and survey disagreement is not an externally identified shock. The evidence therefore evaluates a predictive implication of the mechanism, with pointwise uncertainty reported explicitly.

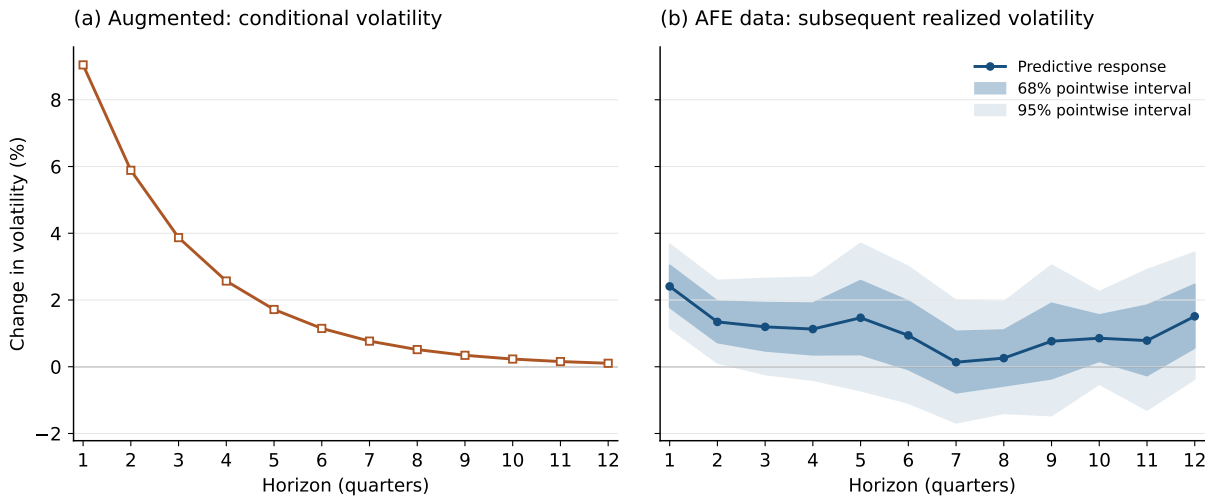


Figure 5: Disagreement and subsequent volatility: model and AFE data

Notes: Left: generalized response of the mean conditional standard deviation of exchange-rate changes to a one-standard-deviation disagreement shock in the augmented model, relative to its unconditional mean. Right: bilateral AUD, CAD, CHF, EUR, GBP, JPY, and SEK panel local projections. Disagreement is the quarterly mean log cross-sectional SD of individual log 90-day forecasts; at least eight institutions are required per survey month. Outcomes are log realized volatility from daily log nominal spot changes in entirely subsequent quarters. Controls are current log volatility, two lags of log volatility and disagreement, current and lagged depreciation, and mean log respondent count, with currency and quarter effects. Effects are scaled by a fixed predictor-sample residual SD of 0.146 and transformed into percent. Dark and light bands are 68 and 95 percent pointwise intervals using calendar-quarter score Bartlett HAC(4), finite-sample covariance correction, and t_{T-1} critical values. Missing observations are retained in the calendar when constructing horizons. Shock normalizations differ across panels.

6 Conclusion

This paper shows that disagreement among arbitrageurs can generate exchange rate disconnect through endogenous wealth redistribution. Heterogeneous beliefs and wealth dynamics create a wealth–disagreement wedge: a time-varying deviation from UIP that reflects who holds wealth when beliefs differ. This wedge provides a channel through which standard macroeconomic shocks generate large exchange rate movements, distinct from conventional risk compensation and the effects of external imbalances.

With disagreement disciplined by survey data, the calibrated model jointly reproduces high relative exchange rate volatility, negative Fama and Backus–Smith coefficients, and near-random-walk exchange rates. The baseline imposes rational consensus beliefs about fundamentals while retaining individual disagreement. Holding all parameters fixed, eliminating disagreement reduces quarterly exchange rate volatility by 62 percent and reverses the signs of both the Fama and Backus–Smith coefficients. Re-estimating the conventional shocks does not restore the model's joint fit. Conventional demand and supply shocks nevertheless account for most exchange rate variation. Disagreement therefore matters primarily by altering the propagation of these shocks: currency positions redistribute wealth, thereby changing whose expectations exert the greatest influence on equilibrium prices. Historical and empirical exercises provide complementary evidence for this mechanism.

The broader implication is that an apparent disconnect from macroeconomic fundamentals can arise because their effects depend on the history of trading gains and losses within financial markets. The distribution of wealth across investors with different beliefs changes the exchange rate response to fundamental news, making that relationship difficult to recover from aggregate observables alone. This perspective complements work emphasizing currency demand and financial intermediation in exchange rate determination, including Dao, Gourinchas and Itskhoki (2025). A direct test of the wealth–disagreement mechanism requires investor-level data linking expectations, currency positions, wealth, and risk-bearing capacity. Such data would allow researchers to examine whether trading gains increase the pricing influence of investors whose beliefs differ from the consensus.

These findings suggest several directions for future research. First, the current model abstracts from bond risk premia: it includes only short-term bonds, while long-term bonds are priced according to the expectations hypothesis. A substantial body of evidence indicates that long-term yields are more closely connected to exchange rates and that UIP performs better at longer horizons. Incorporating long-term bonds could therefore link exchange rates to both conventional and unconventional monetary policy and shed light on the global financial cycle and the special role of the U.S. dollar. Second, extending the model to a

policy setting could clarify the implications of a UIP wedge whose fluctuations are driven by conventional macroeconomic shocks. Finally, directly measuring financial flows among arbitrageurs would provide an important test of the proposed mechanism and contribute more broadly to the literature on heterogeneous beliefs in finance. I leave these extensions for future research.

References

Alessandria, George, "Consumer search, price dispersion, and international relative price fluctuations," *International Economic Review*, 2009, *50* (3), 803–829.

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

_ **and** _ , "Puzzling exchange rate dynamics and delayed portfolio adjustment," *Journal of International Economics*, 2021, *131*, 103460.

Borovička, Jaroslav, Lars Peter Hansen, and Thomas J Sargent, "Exploring recursive utility," Working paper 2023.

Caballero, Ricardo J and Alp Simsek, "Monetary policy with opinionated markets," *American Economic Review*, 2022, *112* (7), 2353–2392.

Campbell, John Y and Luis M Viceira, "Consumption and portfolio decisions when expected returns are time varying," *The Quarterly Journal of Economics*, 1999, *114* (2), 433–495.

Candian, Giacomo and Pierre De Leo, "Imperfect Exchange Rate Expectations," *The Review of Economics and Statistics*, 09 2025, *107* (5), 1406–1423.

Caramp, Nicolas and Dejanir H Silva, "Monetary policy and wealth effects: The role of risk and heterogeneity," *The Journal of Finance*, 2026, *81* (2), 1011–1052.

Chahrour, Ryan, Vito Cormun, Pierre De Leo, Pablo A Guerrón-Quintana, and Rosen Valchev, "Exchange rate disconnect revisited," Technical Report, National Bureau of Economic Research 2024.

Chang, Chun, Zheng Liu, and Mark M. Spiegel, "Capital controls and optimal Chinese monetary policy," *Journal of Monetary Economics*, 2015, *74*, 1–15.

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

Dao, Mai C, Pierre-Olivier Gourinchas, and Oleg Itskhoki, "Breaking Parity: Equilibrium Exchange Rates and Currency Premia," Technical Report, National Bureau of Economic Research 2025.

Detemple, Jerome and Shashidhar Murthy, "Intertemporal asset pricing with heterogeneous beliefs," *Journal of Economic Theory*, 1994, *62* (2), 294–320.

Du, Qingyuan, Stefano Eusepi, and Bruce Preston, "Non-rational beliefs in an open economy," *Review of Economic Dynamics*, 2021, *41*, 174–204.

Engel, Charles, "Exchange rates, interest rates, and the risk premium," *American Economic Review*, 2016, *106* (2), 436–474.

_ **and Steve Pak Yeung Wu**, "Exchange Rate Models are Better than You Think, and Why They Didn't Work in the Old Days," Technical Report, National Bureau of Economic Research 2024.

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

Gabaix, Xavier and Matteo Maggiori, "International liquidity and exchange rate dynamics," *The Quarterly Journal of Economics*, 2015, *130* (3), 1369–1420.

Gilchrist, Simon and Egon Zakrajšek, "Credit spreads and business cycle fluctuations," *American economic review*, 2012, *102* (4), 1692–1720.

Gourinchas, Pierre-Olivier and Aaron Tornell, "Exchange rate puzzles and distorted beliefs," *Journal of International Economics*, 2004, *64* (2), 303–333.

Hassan, Tarek A, Thomas M Mertens, and Tony Zhang, "Not so disconnected: Exchange rates and the capital stock," *Journal of International Economics*, 2016, *99*, S43–S57.

He, Zhiguo, Bryan Kelly, and Asaf Manela, "Intermediary asset pricing: New evidence from many asset classes," *Journal of Financial Economics*, 2017, *126* (1), 1–35.

Itskhoki, Oleg and Dmitry Mukhin, "Exchange rate disconnect in general equilibrium," *Journal of Political Economy*, 2021, *129* (8), 2183–2232.

_ **and** _ , "What drives the exchange rate?," *IMF Economic Review*, 2025, *73* (1), 86–117.

Jiang, Zhengyang, Arvind Krishnamurthy, Hanno Lustig, and Jialu Sun, "Convenience yields and exchange rate puzzles," Technical Report, National Bureau of Economic Research 2024.

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

Kekre, Rohan and Moritz Lenel, "Exchange rates, natural rates, and the price of risk," *University of Chicago, Becker Friedman Institute for Economics Working Paper*, 2024, (2024-114).

_ **and** _ , "A model of us monetary policy and the global financial cycle," *Working Paper*, 2026.

_ **and** _ , "What Do Asset Prices in April 2025 Say About Demand for the Dollar?," Working Paper 35466, National Bureau of Economic Research July 2026.

Kohlhas, Alexandre N. and Ansgar Walther, "Asymmetric Attention," *American Economic Review*, September 2021, *111* (9), 2879–2925.

Lilley, Andrew, Matteo Maggiori, Brent Neiman, and Jesse Schreger, "Exchange rate reconnect," *Review of Economics and Statistics*, 2022, *104* (4), 845–855.

Meese, Richard A and Kenneth Rogoff, "Empirical exchange rate models of the seventies: Do they fit out of sample?," *Journal of international economics*, 1983, *14* (1-2), 3–24.

Miyamoto, Wataru, Thuy Lan Nguyen, and Hyunseung Oh, "In search of dominant drivers of the real exchange rate," *Review of Economics and Statistics*, 2025, *107* (4), 1059–1073.

Molavi, Pooya, Alireza Tahbaz-Salehi, and Andrea Vedolin, "Model complexity, expectations, and asset prices," *Review of Economic Studies*, 2024, *91* (4), 2462–2507.

Stavrakeva, Vania and Jenny Tang, "A fundamental connection: Exchange rates and macroeconomic expectations," *Review of Economics and Statistics*, 2024, pp. 1–49.

Valchev, Rosen, "Bond convenience yields and exchange rate dynamics," *American Economic Journal: Macroeconomics*, 2020, *12* (2), 124–166.

Xiong, Wei and Hongjun Yan, "Heterogeneous expectations and bond markets," *The Review of Financial Studies*, 2010, *23* (4), 1433–1466.

Appendix contents

A Full first-order equilibrium conditions **46**

A.1 Notation . . . 46

A.2 Steady-state quantities . . . 47

A.3 First-order equilibrium system . . . 48

A.4 Fundamentals, filtering, and expectations . . . 51

B Omitted Proofs and Derivations **53**

B.1 Information structure and belief dynamics . . . 53

B.2 Optimal portfolio choice . . . 57

B.3 Proof of Lemma 1 . . . 58

B.4 Proof of Lemma 2 . . . 59

B.5 Proof of Proposition 1 and Corollary 1 . . . 61

B.6 Proof of Proposition 2 . . . 62

B.7 Proof of Proposition 3 . . . 66

C Disagreement and the Backus–Smith correlation **69**

C.1 The calibrated economies and the fixed-shock comparison . . . 69

C.2 How supply shocks alone can produce a negative correlation . . . 71

C.3 A numerical illustration in the reduced equilibrium . . . 73

D Augmented model: calibration and currency-return predictability **73**

D.1 Augmented calibration . . . 73

D.2 Equivalence to the Candian–De Leo belief process . . . 75

D.3 Predictability across horizons . . . 77

E Time-varying disagreement in the quantitative model **80**

E.1 Beliefs and portfolio demand . . . 80

E.2 The two financial interactions . . . 80

E.3 Forecasting and solution . . . 81

F Additional quantitative results **83**

F.1 Full comparison of exchange-rate moments across model variants . . . 83

F.2 Full impulse responses across model variants . . . 88

F.3 Historical inversion and recursive accounting . . . 95

G Additional Empirical Results **103**

G.1 Disagreement and exchange-rate changes 103

A Full first-order equilibrium conditions

This appendix collects the full first-order system of the structural model around the possibly asymmetric steady state. It retains the Home-bond UIP formulation, pricing to market, both groups' wealth dynamics, and the information structure in which current fundamentals are observed. Coefficients are written in terms of primitives and steady-state quantities; Γ is the effective risk-cost parameter defined in the portfolio approximation.

A.1 Notation

The Home and Foreign intermediate-good endowments are y_t and y_t^*, respectively, with steady-state values y and y^*. The producer prices are p_t, p_t^*, whereas $p_{Ht}, p_{Ft}, p_{Ht}^*, p_{Ft}^*$ are consumer prices. A star denotes the Foreign counterpart; for consumption quantities and consumer prices, it identifies the country of consumption. A subscript H or F on consumption quantities and consumer prices identifies the good's origin. The exchange rate e_t is Home consumption units per Foreign consumption unit. We use equal group masses, $\eta_A = \eta_B = 1/2$, and steady-state values have no time subscript.

For positive quantities and prices, $\hat{z}_t = \log(z_t/z)$. Interest-rate hats refer to gross returns:

$$\hat{r}_t = \log \frac{1+r_t}{1+r}, \qquad \hat{r}_t^* = \log \frac{1+r_t^*}{1+r^*}.$$

Bond positions, carry positions, and wealth shares can be zero or signed, so they use the following level deviations:

$$\widetilde{b}_t = \frac{b_t - b}{y}, \qquad \widetilde{b}_t^* = \frac{b_t^* - b^*}{y^*}, \qquad \widetilde{N}_t^i = \frac{N_t^i - N^i}{W^i}, \qquad \widetilde{\omega}_t^i = \omega_t^i - \omega^i.$$

Thus each household bond deviation is normalized by its own country's steady-state endowment, and each carry-position deviation is normalized by its own group's steady-state wealth. The latent persistent fundamental is denoted by τ_t^x, as in the information section.

Aggregate arbitrageur wealth and the Home-valued carry position are $W_t = (W_t^A + W_t^B)/2$ and $N_t = (N_t^A + N_t^B)/2$, measured after exit and entry. The bond markets clear when

$$\frac{b_t}{1+r_t} + W_t - N_t = 0, \qquad \frac{e_t b_t^*}{1+r_t^*} + N_t = 0.$$

The aggregate wealth law implies that exiting arbitrageurs rebate $\frac{\xi}{1-\xi}(W_t - \xi W_0)$, while new

entrants receive ξW_0. Net transfers and their first-order deviations are

$$\pi_t^a = \frac{\xi}{1-\xi}(W_t - W_0), \qquad \pi_t^a - \pi^a = \frac{\xi}{1-\xi} W \hat{W}_t.$$

A.2 Steady-state quantities

Normalize $e = 1$ and choose endowments so that all four consumer prices are one. The producer prices are $p = p^* = (\theta_v - 1)/\theta_v$, rather than one, because domestic firms charge a constant markup. Household Euler equations give $1 + r = 1/\beta$ and $1 + r^* = 1/\beta^*$. In the benchmark with offsetting permanent forecasts,

$$\frac{N^A}{W^A} = \frac{\log(\beta/\beta^*) + d_{ss}/2}{\Gamma}, \qquad \frac{N^B}{W^B} = \frac{\log(\beta/\beta^*) - d_{ss}/2}{\Gamma}.$$

With the same entrant endowment in each group,

$$W^i = \frac{\xi W_0}{1 - (1-\xi)\left[\frac{1}{\beta} + \left(\frac{1}{\beta^*} - \frac{1}{\beta}\right)\frac{N^i}{W^i}\right]}, \qquad i \in \{A, B\}.$$

The position ratios in this expression are given explicitly above. Positive finite wealth requires positive denominators. Aggregate wealth, positions, and shares satisfy

$$W = \frac{W^A + W^B}{2}, \quad N = \frac{N^A + N^B}{2}, \quad \omega^i = \frac{W^i}{2W}, \quad b = -\frac{W - N}{\beta}, \quad b^* = -\frac{N}{\beta^*}.$$

Offsetting permanent forecasts do not imply equal wealth shares when the steady-state carry return is nonzero.

The remaining steady-state quantities satisfy

$$\begin{aligned}
&c_H = (1-\gamma)c, && c_F = \gamma c, && c_H^* = \gamma c^*, \quad c_F^* = (1-\gamma)c^*,\\
&y = c_H + c_H^*, && y^* = c_F + c_F^*, &&\\
&c = y + (1-\beta)b + \frac{\xi}{1-\xi}(W - W_0), && c^* = y^* + (1-\beta^*)b^*. &&
\end{aligned}$$

Firm profits are already included in the last two equations: producer income plus firm profits equals sales revenue. One may normalize $y = 1$ and solve these conditions jointly for the remaining steady-state quantities. Only an independent subset is needed because the aggregate budget identity is implied by market clearing.

A.3 First-order equilibrium system

The following equations apply for $i \in \{A, B\}$. Subjective forecast terms are first-order changes around each group's own steady-state forecast. The permanent forecast difference is already separated into d_{ss} and must not be added again to $E_t^A \hat{e}_{t+1} - E_t^B \hat{e}_{t+1}$.

Household Euler equations Because current discount factors are observed under the stated information structure,

$$\sigma(\mathbb{E}_t \hat{c}_{t+1} - \hat{c}_t) = \hat{\beta}_t + \hat{r}_t,$$
$$\sigma(\mathbb{E}_t \hat{c}^*_{t+1} - \hat{c}^*_t) = \hat{\beta}^*_t + \hat{r}^*_t.$$

Filtering affects forecasts of future fundamentals. The current discount factor is not replaced by the posterior estimate of its persistent component in these equations.

Consumer demand and price indices CES demand and the two consumption-price normalizations give

$$\hat{c}_{Ht} = \hat{c}_t - \theta \hat{p}_{Ht}, \qquad \hat{c}_{Ft} = \hat{c}_t - \theta \hat{p}_{Ft},$$
$$\hat{c}^*_{Ht} = \hat{c}^*_t - \theta \hat{p}^*_{Ht}, \qquad \hat{c}^*_{Ft} = \hat{c}^*_t - \theta \hat{p}^*_{Ft},$$
$$0 = (1-\gamma)\hat{p}_{Ht} + \gamma \hat{p}_{Ft}, \qquad 0 = \gamma \hat{p}^*_{Ht} + (1-\gamma)\hat{p}^*_{Ft}.$$

Firm pricing The elasticity across domestic varieties is $\theta_v > 1$, distinct from the trade elasticity θ. The imported-variety elasticities are $\theta_{v,t} = \theta_v (e_t/e)^{\xi_v}$ and $\theta^*_{v,t} = \theta_v (e_t/e)^{-\xi_v}$. Differentiating the markup $\theta_{v,t}/(\theta_{v,t}-1)$ and the firms' optimal prices gives

$$\hat{p}_{Ht} = \hat{p}_t, \qquad \hat{p}^*_{Ft} = \hat{p}^*_t,$$
$$\hat{p}^*_{Ht} = \hat{p}_t - \left(1 - \frac{\xi_v}{\theta_v - 1}\right)\hat{e}_t, \qquad \hat{p}_{Ft} = \hat{p}^*_t + \left(1 - \frac{\xi_v}{\theta_v - 1}\right)\hat{e}_t.$$

Thus the law of one price holds when $\xi_v = 0$. The pass-through coefficient called χ elsewhere in the paper is $1 - \xi_v/(\theta_v - 1)$; it is written out here. Solving with the price-index equations gives

$$\hat{p}_{Ht} = -\frac{\gamma}{1-2\gamma}\left(1 - \frac{\xi_v}{\theta_v - 1}\right)\hat{e}_t, \qquad \hat{p}_{Ft} = \frac{1-\gamma}{1-2\gamma}\left(1 - \frac{\xi_v}{\theta_v - 1}\right)\hat{e}_t,$$
$$\hat{p}^*_{Ht} = -\frac{1-\gamma}{1-2\gamma}\left(1 - \frac{\xi_v}{\theta_v - 1}\right)\hat{e}_t, \qquad \hat{p}^*_{Ft} = \frac{\gamma}{1-2\gamma}\left(1 - \frac{\xi_v}{\theta_v - 1}\right)\hat{e}_t.$$

These solved prices can replace the pricing and index equations; they are not additional independent restrictions.

Goods-market clearing Market clearing in physical quantities is

$$y\hat{y}_t = c_H\hat{c}_{Ht} + c_H^*\hat{c}_{Ht}^*, \qquad y^*\hat{y}_t^* = c_F\hat{c}_{Ft} + c_F^*\hat{c}_{Ft}^*.$$

Equivalently, substituting demand and prices gives a system directly in aggregate consumption and the exchange rate:

$$\begin{aligned} y\hat{y}_t &= c_H\hat{c}_t + c_H^*\hat{c}_t^* + \frac{\theta[\gamma c_H + (1-\gamma)c_H^*]}{1-2\gamma}\left(1 - \frac{\xi_v}{\theta_v - 1}\right)\hat{e}_t, \\ y^*\hat{y}_t^* &= c_F\hat{c}_t + c_F^*\hat{c}_t^* - \frac{\theta[(1-\gamma)c_F + \gamma c_F^*]}{1-2\gamma}\left(1 - \frac{\xi_v}{\theta_v - 1}\right)\hat{e}_t. \end{aligned}$$

Household budgets, including firm profits Consolidating producer income and the profits of domestically owned firms gives the exact revenue identities

$$\begin{aligned} p_t y_t + \pi_t &= p_{Ht}c_{Ht} + e_t p_{Ht}^* c_{Ht}^*, \\ p_t^* y_t^* + \pi_t^* &= \frac{p_{Ft}}{e_t}c_{Ft} + p_{Ft}^* c_{Ft}^*. \end{aligned}$$

At first order, use goods clearing and the pricing equations to obtain

$$\begin{aligned} \beta y\widetilde{b}_t - \beta b\hat{r}_t - y\widetilde{b}_{t-1} &= y(\hat{y}_t + \hat{p}_{Ht}) - c\hat{c}_t + \frac{\xi_v}{\theta_v - 1}c_H^*\hat{e}_t + \frac{\xi}{1-\xi}W\hat{W}_t, \\ \beta^* y^*\widetilde{b}_t^* - \beta^* b^*\hat{r}_t^* - y^*\widetilde{b}_{t-1}^* &= y^*(\hat{y}_t^* + \hat{p}_{Ft}^*) - c^*\hat{c}_t^* - \frac{\xi_v}{\theta_v - 1}c_F\hat{e}_t. \end{aligned}$$

The terms proportional to $\xi_v/(\theta_v - 1)$ reflect export revenues when firms price differently across destinations. They vanish with full pass-through. Producer income alone would omit these profit-rebate terms. The two budgets are shown for completeness; once goods clearing, both bond markets, and the aggregate wealth law hold, only one household budget is independent.

Carry positions and group wealth The position rules are

$$\begin{aligned} \widetilde{N}_t^A &= \frac{\log(\beta/\beta^*) + d_{ss}/2}{\Gamma}\hat{W}_t^A + \frac{E_t^A\hat{e}_{t+1} - \hat{e}_t + \hat{r}_t^* - \hat{r}_t}{\Gamma}, \\ \widetilde{N}_t^B &= \frac{\log(\beta/\beta^*) - d_{ss}/2}{\Gamma}\hat{W}_t^B + \frac{E_t^B\hat{e}_{t+1} - \hat{e}_t + \hat{r}_t^* - \hat{r}_t}{\Gamma}. \end{aligned}$$

Differentiating the group wealth law with fixed entrant endowment gives

$$\hat{W}^i_{t+1} = (1-\xi)\Bigg\{ \frac{\hat{W}^i_t + \hat{r}_t}{\beta} + \left(\frac{1}{\beta^*} - \frac{1}{\beta}\right)\widetilde{N}^i_t + \frac{N^i}{W^i}\left[\frac{\hat{e}_{t+1} - \hat{e}_t + \hat{r}^*_t}{\beta^*} - \frac{\hat{r}_t}{\beta}\right]\Bigg\}.$$

The steady-state ratios N^i/W^i are given above; they are not additional parameters. This expression keeps the level deviation of the arithmetic carry return explicit, including both gross-return factors.

Wealth aggregation and Home-bond UIP The first-order aggregation identities are

$$\hat{W}_t = \omega^A \hat{W}^A_t + \omega^B \hat{W}^B_t, \qquad \widetilde{\omega}^A_t = \omega^A\omega^B(\hat{W}^A_t - \hat{W}^B_t), \qquad \widetilde{\omega}^B_t = -\widetilde{\omega}^A_t.$$

Substituting the portfolio rules into Home-bond clearing gives the full UIP condition with its coefficients expanded:

$$\omega^A E^A_t \hat{e}_{t+1} + \omega^B E^B_t \hat{e}_{t+1} - \hat{e}_t + \hat{r}^*_t - \hat{r}_t = -d_{ss}\widetilde{\omega}^A_t + \frac{\beta\Gamma}{W}\left[y\widetilde{b}_t - b(\hat{W}_t + \hat{r}_t)\right]. \tag{13}$$

In the unbiased-consensus benchmark, the weighted forecast on the left equals $\mathbb{E}_t\hat{e}_{t+1} + (\omega^A - \omega^B)(E^A_t\hat{e}_{t+1} - E^B_t\hat{e}_{t+1})/2$, recovering Lemma 1 without introducing absorption coefficients. The group wealth equations and aggregation identities already imply Lemma 2; its share law is not an additional equilibrium equation.

Foreign bonds and accounting closure Foreign-bond clearing recovers the Foreign household's bond position:

$$y^*\widetilde{b}^*_t = -\frac{1}{\beta^*}\left[W\sum_i \omega^i \widetilde{N}^i_t + N(\hat{r}^*_t - \hat{e}_t)\right].$$

For comparison, Home-bond clearing before substituting the position rules is

$$\beta y\widetilde{b}_t - \beta b\hat{r}_t + W\left(\hat{W}_t - \sum_i \omega^i \widetilde{N}^i_t\right) = 0.$$

Use either this Home-clearing equation or the UIP equation, not both as independent restrictions. The exact consolidated financial identity behind the household-budget redundancy is

$$\frac{b_t}{1+r_t} + \frac{e_t b_t^*}{1+r_t^*} - b_{t-1} - e_t b_{t-1}^* = \frac{\xi}{1-\xi}(W_t - W_0) = \pi_t^a.$$

It follows directly from the two bond-clearing equations and the aggregate wealth law. Together with goods clearing and the firms' revenue identities, it verifies the Foreign budget whenever the Home budget holds.

A.4 Fundamentals, filtering, and expectations

For each $x \in \{\beta, \beta^*, y, y^*\}$, the observable realization and its latent persistent component satisfy

$$\hat{x}_t = \tau_t^x + v_t^x, \qquad \tau_t^x = \rho_x \tau_{t-1}^x + \varepsilon_t^x.$$

The transitory component and persistent innovation have standard deviations σ_{v^x} and σ_{ε^x}. Their cross-country correlations are those specified in the calibration. The realization $\hat{x}_t$ is observed at t; filtering concerns its decomposition, rather than an additional noisy observation of $\hat{x}_t$.

Under the maintained scalar filtering rule for each fundamental, group i's posterior $\tau_{t|t}^{xi} = E_t^i[\tau_t^x]$ evolves as

$$\tau_{t|t}^{xi} = \rho_x(1 - G_x)\tau_{t-1|t-1}^{xi} + G_x(\hat{x}_t + \mu^{xi} + \widetilde{\mu}_t^{xi}).$$

The gain in the information section is determined by

$$G_x = \frac{\Sigma_{\tau,x}}{\Sigma_{\tau,x} + \sigma_{v^x}^2}, \qquad \Sigma_{\tau,x} = \rho_x^2(1 - G_x)\Sigma_{\tau,x} + \sigma_{\varepsilon^x}^2.$$

Here $\Sigma_{\tau,x}$ is the prior variance of τ_t^x, not the posterior variance. The steady-state posterior bias of group i is $G_x\mu^{xi}/[1 - \rho_x(1 - G_x)]$. The interpretation differences $\widetilde{\mu}_t^{xA} - \widetilde{\mu}_t^{xB}$ are the mean-zero shocks with variance $\sigma_{\mu,x}^2$.

Define consensus and the fundamental-specific disagreement gap by $\bar{\tau}_{t|t}^x = (\tau_{t|t}^{xA} + \tau_{t|t}^{xB})/2$ and $d_t^x = \tau_{t|t}^{xA} - \tau_{t|t}^{xB}$. Averaging and differencing the individual rules gives

$$\begin{aligned}
\bar{\tau}_{t|t}^x &= \rho_x(1 - G_x)\bar{\tau}_{t-1|t-1}^x + G_x\hat{x}_t + \frac{G_x}{2}(\mu_t^{xA} + \mu_t^{xB}), \\
d_t^x - d_{ss}^x &= \rho_x(1 - G_x)(d_{t-1}^x - d_{ss}^x) + G_x(\widetilde{\mu}_t^{xA} - \widetilde{\mu}_t^{xB}), \\
d_{ss}^x &= \frac{G_x(\mu^{xA} - \mu^{xB})}{1 - \rho_x(1 - G_x)}.
\end{aligned}$$

These are consequences of the individual filtering equations, not extra independent restrictions. The average-interpretation term is zero in the benchmark and is retained for extensions with average bias. The gap d_t^x concerns a particular latent fundamental; it is distinct from exchange-rate forecast disagreement $d_t = D_t - d_{ss} = UIP_t^A - UIP_t^B$ and from the steady-state exchange-rate forecast gap d_{ss}.

Current posteriors and interpretations are mutually observed. In the benchmark, subjective future beliefs follow the higher-order expectation formulas in Proposition 4, derived in Appendix B.1. Both groups apply those beliefs to the same equilibrium price and wealth laws. In particular, $E_t^i \hat{e}_{t+1}$ is the conditional forecast of the endogenous exchange-rate solution, not an independent shock or a freely chosen forecast coefficient. Given the exogenous innovation laws, initial wealth, inherited bonds, and initial posteriors, the equations above and these conditional expectation rules close the system. The solution is selected to satisfy household transversality and the maintained non-explosive equilibrium restriction.

Average-bias extension The zero-average-bias benchmark sets $\bar{\mu}_t^x = (\mu_t^{xA} + \mu_t^{xB})/2 = 0$. For the extension specified in the quantitative section, substitute

$$\bar{\mu}_t^x = \left(\frac{\kappa_x}{G_x} - 1 \right) \hat{x}_t + \frac{(1-\kappa_x)\tilde{\rho}_x - \rho_x(1-G_x)}{G_x} \bar{\tau}_{t-1|t-1}^x.$$

The consensus recursion then simplifies to

$$\bar{\tau}_{t|t}^x = (1-\kappa_x)\tilde{\rho}_x \bar{\tau}_{t-1|t-1}^x + \kappa_x \hat{x}_t.$$

Here κ_x is the perceived information gain and $\tilde{\rho}_x$ the perceived persistence, as in the main text. This extension uses the existing fundamental shocks. When households share the common biased beliefs, the conditional expectations in both household Euler equations are evaluated under that same subjective law. The current observed discount factors on the right-hand sides remain $\hat{\beta}_t$ and $\hat{\beta}_t^*$.

The structural system above sets the additional reduced-form UIP shock ψ_t to zero. The quantitative pricing-wedge extension specifies $\psi_t = \rho_\psi \psi_{t-1} + \varepsilon_t^\psi$ and adds ψ_t to the right-hand side of UIP. That specification replaces the structural pricing restriction; it is not an additional restriction to impose alongside unmodified Home-bond clearing and portfolio demand. A structural interpretation of this extra shock requires specifying its counterpart in portfolio demand or outside bond demand.

B Omitted Proofs and Derivations

B.1 Information structure and belief dynamics

Subjective interpretation of public information For each driving variable, agents observe $\hat{x}_t = \tau_t^x + v_t^x$ but not its persistent and transitory components separately. Under group i's subjective probability measure, the observation satisfies

$$\hat{x}_t =^i \tau_t^x - \mu_t^{xi} + v_t^x. \tag{14}$$

The notation $=^i$ denotes equality under group i's perceived data-generating process. Group i thus treats $\hat{x}_t + \mu_t^{xi}$ as a noisy but unbiased observation of τ_t^x.

Each group considers the other's interpreted observation a garbled version of its own:

$$\hat{x}_t + \mu_t^{xj} =^i \tau_t^x + v_t^x + \mu_t^{xj} - \mu_t^{xi}, \qquad j \neq i.$$

Interpretations and current posteriors are mutually observed. Conditional on its own interpretation, group i regards the other group's interpretation as uninformative about τ_t^x. Equilibrium prices therefore reveal no additional information.

Interpretations have permanent and transitory components, $\mu_t^{xi} = \mu^{xi} + \widetilde{\mu}_t^{xi}$. The transitory components are jointly normal and i.i.d. across periods, with

$$\widetilde{\mu}_t^{xA} - \widetilde{\mu}_t^{xB} \sim N(0, \sigma_{\mu,x}^2). \tag{15}$$

The groups have equal population shares, and the zero-average-bias benchmark imposes $\bar{\mu}_t^x = (\mu_t^{xA} + \mu_t^{xB})/2 = 0$. The following proposition collects the resulting filtering and higher-order expectation formulas.

Proposition 4. *Denote the filtered time-t persistent state by group $i \in \{A, B\}$ given time t information as $\tau_{t|t}^{xi} \equiv E_t^i[\tau_t^x]$. Let the steady state Kalman gain be $G_x = \frac{\left(\sigma_{v^x}^2\right)^{-1}}{(\Sigma_{\tau,x})^{-1} + \left(\sigma_{v^x}^2\right)^{-1}}$, where $\Sigma_{\tau,x}$ is the steady-state prior variance of τ_t^x, which solves $\Sigma_{\tau,x} = \rho_x^2 (1 - G_x) \Sigma_{\tau,x} + \sigma_{\varepsilon^x}^2$. The individual beliefs, consensus beliefs, $\bar{\tau}_{t|t}^x \equiv \frac{1}{2}\left(\tau_{t|t}^{xA} + \tau_{t|t}^{xB}\right)$, and disagreement, $d_t^x \equiv \tau_{t|t}^{xA} - \tau_{t|t}^{xB}$, evolve according to*

$$\begin{aligned}
\tau_{t|t}^{xi} &= G_x \mu^{xi} + \varphi_x \tau_{t-1|t-1}^{xi} + G_x \left(\hat{x}_t + \widetilde{\mu}_t^{xi}\right), \\
\bar{\tau}_{t|t}^x &= G_x \bar{\mu}_t^x + \varphi_x \bar{\tau}_{t-1|t-1}^x + G_x \hat{x}_t, \\
d_t^x &= (1 - \varphi_x)\, d_{ss}^x + \varphi_x d_{t-1}^x + G_x (\widetilde{\mu}_t^{xA} - \widetilde{\mu}_t^{xB}),
\end{aligned}$$

respectively, where $\varphi_x = \rho_x (1 - G_x)$*, and* $d^x_{ss} = \frac{G_x}{1-\varphi_x} \left(\mu^{xA} - \mu^{xB}\right)$*. Moreover, current beliefs are mutually observed, and for* $h \geq 1$*,*

$$\begin{aligned}
E^A_t \left[\tau^{xA}_{t+h|t+h}\right] &= \rho^h_x \tau^{xA}_{t|t},\ E^B_t \left[\tau^{xB}_{t+h|t+h}\right] = \rho^h_x \tau^{xB}_{t|t}, \\
E^A_t \left[\tau^{xB}_{t+h|t+h}\right] &= \varphi^h_x \tau^{xB}_{t|t} + \left(\rho^h_x - \varphi^h_x\right) \tau^{xA}_{t|t} - \left(1 - \varphi^h_x\right) d^x_{ss}, \\
E^B_t \left[\tau^{xA}_{t+h|t+h}\right] &= \varphi^h_x \tau^{xA}_{t|t} + \left(\rho^h_x - \varphi^h_x\right) \tau^{xB}_{t|t} + \left(1 - \varphi^h_x\right) d^x_{ss}.
\end{aligned}$$

Consequently, both groups agree on the path for disagreement:

$$E^A_t \left[d^x_{t+h}\right] = E^B_t \left[d^x_{t+h}\right] = d^x_{ss} + \varphi^h_x \left(d^x_t - d^x_{ss}\right).$$

Proof

Individual belief updating Each agent i observes the current realization $\hat{x}_t$ but not its persistent and transitory components separately. Group i treats $\hat{x}_t + \mu^{xi}_t$ as a noisy observation of the latent persistent component. Combining the prior mean $\rho_x \tau^{xi}_{t-1|t-1}$ with this interpreted observation gives the steady-state Kalman update

$$\tau^{xi}_{t|t} = (1 - G_x)\rho_x \tau^{xi}_{t-1|t-1} + G_x(\hat{x}_t + \mu^{xi}_t) = \varphi_x \tau^{xi}_{t-1|t-1} + G_x(\hat{x}_t + \mu^{xi}_t),$$

where $\varphi_x \equiv \rho_x(1 - G_x)$. There is no second signal for the realized innovation: adding one would double count information already contained in the observed $\hat{x}_t$.

Consensus beliefs Taking the weighted average across agents with weights η_i:

$$\bar{\tau}^x_{t|t} = \sum_i \eta_i \tau^{xi}_{t|t} = \varphi_x \bar{\tau}^x_{t-1|t-1} + G_x \left(\hat{x}_t + \sum_i \eta_i \mu^{xi}_t\right).$$

Since $\sum_i \eta_i \mu^{xi}_t = 0$ in the zero-average-bias benchmark, this reduces to the benchmark consensus equation. With average bias, the same averaging identity retains the interpretation term shown in the main text.

Disagreement Taking the difference between agents i and j:

$$d^x_t = \varphi_x d^x_{t-1} + G_x \Delta\mu^x_t = (1 - \varphi_x) d^x_{ss} + \varphi_x d^x_{t-1} + G_x \varepsilon^x_{\mu,t}.$$

The second equality uses $\Delta\mu^x_t \equiv \mu^{xA}_t - \mu^{xB}_t = \mu^{xA} - \mu^{xB} + \varepsilon^x_{\mu,t}$ and $d^x_{ss} = G_x(\mu^{xA} - \mu^{xB})/(1 - \varphi_x)$.

Higher-order expectations For $h = 0$, agent i observes agent j's current belief directly, so $E_t^i[\tau_{t|t}^{xj}] = \tau_{t|t}^{xj}$.

For one period ahead, the subjective conditional means implied by the common transition law and the interpretation convention are

$$E_t^A[\tau_{t+1|t+1}^{xA}] = \rho_x \tau_{t|t}^{xA},$$
$$E_t^A[\tau_{t+1|t+1}^{xB}] = \varphi_x \tau_{t|t}^{xB} + (\rho_x - \varphi_x)\tau_{t|t}^{xA} - (1 - \varphi_x)d_{ss}^x,$$

with the symmetric expressions under group B. Iterating these affine recursions proves the higher-order-expectation formulas in Proposition 4. In particular,

$$E_t^A[\tau_{t+h|t+h}^{xB}] = \varphi_x^h \tau_{t|t}^{xB} + (\rho_x^h - \varphi_x^h)\tau_{t|t}^{xA} - (1 - \varphi_x^h)d_{ss}^x,$$
$$E_t^A[d_{t+h}^x] = d_{ss}^x + \varphi_x^h(d_t^x - d_{ss}^x),$$

and the corresponding formulas for group B follow by exchanging the group labels and reversing the sign of the permanent gap. This also shows formally that the groups agree on the future path of disagreement even though their forecasts of the persistent component differ.

B.1.1 Private information and price learning

This subsection proves Remark 2 conditional on a candidate stationary linear price signal. For each fundamental x, retain the objective state τ_t^x and let group i use a fixed subjective long-run mean ℓ^{xi}. Its prior mean is $\tau_{t|t-1}^{xi} = \rho_x \tau_{t-1|t-1}^{xi} + (1 - \rho_x)\ell^{xi}$. Suppose the public observation is $\hat{x}_t = \tau_t^x + v_t^x$, the private signal is $s_t^{xi} = \tau_t^x + \zeta_t^{xi}$, and a common, normalized price signal is $q_t^{x,p} = \tau_t^x + \zeta_t^{x,p}$. Within each subjective model, the signal errors are jointly Gaussian and centered conditional on the state and prior information, with a common time-invariant covariance matrix. The private-signal difference is objectively mean zero and iid over time. Signal precision and prior covariance assumptions are symmetric across groups.

Filtering and disagreement Let $K_{0,x}$, $K_{s,x}$, and $K_{p,x}$ be the common constant gains, with $k_x \equiv K_{0,x} + K_{s,x} + K_{p,x}$, $\psi_x \equiv \rho_x(1 - k_x)$, and $|\psi_x| < 1$. Bayesian updating within each subjective model gives

$$\tau_{t|t}^{xi} = \psi_x \tau_{t-1|t-1}^{xi} + (1 - k_x)(1 - \rho_x)\ell^{xi} + K_{0,x}\hat{x}_t + K_{s,x}s_t^{xi} + K_{p,x}q_t^{x,p}.$$

Subtracting cancels the public and price signals. With $\Delta\ell^x \equiv \ell^{xA} - \ell^{xB}$, the posterior gap and its stationary mean satisfy

$$\begin{aligned} d_t^x &= \psi_x d_{t-1}^x + (1-k_x)(1-\rho_x)\Delta\ell^x + K_{s,x}(\zeta_t^{xA} - \zeta_t^{xB}), \\ d_{ss}^x &= \frac{(1-k_x)(1-\rho_x)}{1-\psi_x}\Delta\ell^x, \\ d_t^x - d_{ss}^x &= \psi_x(d_{t-1}^x - d_{ss}^x) + K_{s,x}(\zeta_t^{xA} - \zeta_t^{xB}). \end{aligned}$$

The one-period fundamental forecast gap is $F_t^x = \rho_x d_t^x + (1-\rho_x)\Delta\ell^x$. Therefore

$$F_{ss}^x = \frac{1-\rho_x}{1-\psi_x}\Delta\ell^x, \qquad F_t^x - F_{ss}^x = \psi_x(F_{t-1}^x - F_{ss}^x) + \rho_x K_{s,x}(\zeta_t^{xA} - \zeta_t^{xB}).$$

Both centered gaps have the AR(1) form of the baseline, with gains determined by the alternative information structure. With zero private-signal differences, constancy also requires initialization at the stationary gap unless $\psi_x = 0$.

Wealth weighting and scope For any two one-period exchange-rate forecasts F_t^A, F_t^B, write $\bar{F}_t = (F_t^A + F_t^B)/2$, $D_t = F_t^A - F_t^B = d_{ss} + d_t$, and $\omega_t^A = 1/2 + \delta_w + \widetilde{\omega}_t^A$. Direct aggregation gives

$$\begin{aligned} F_t^W &= \bar{F}_t + (\delta_w + \widetilde{\omega}_t^A)(d_{ss} + d_t), \\ F_t^W - F_{ss}^W &= (\bar{F}_t - \bar{F}_{ss}) + d_{ss}\widetilde{\omega}_t^A + \delta_w d_t + \widetilde{\omega}_t^A d_t. \end{aligned}$$

Dropping the product of deviations yields the same first-order pricing wedge as in Lemma 1, irrespective of how the forecasts are formed.

The result is conditional, rather than a solution of a price-learning equilibrium. Price coefficients, signal covariances, and filtering gains must satisfy their joint equilibrium fixed point. Averaging the filters retains the average private signal and the price signal, so $\ell^{xA} + \ell^{xB} = 0$ alone does not make consensus equal the objective public-information filter. Nor does this argument recover arbitrary average-bias coefficients or all the higher-order beliefs in Proposition 4. The fixed ℓ^{xi} are persistent subjective model differences, not unknown means that agents eventually learn under a common prior. With a common prior and objectively correct forecasts, the law of iterated expectations gives $\mathbb{E}(F_t^A - F_t^B) = 0$; it therefore does not provide the nonzero deterministic d_{ss} used in the symmetric analytical steady state. □

B.2 Optimal portfolio choice

Fix a date, group, and arbitrageur, and suppress their indices temporarily. Let $R = 1+r_t > 0$, $w = w_t^i(\iota) > 0$, $\mu = E_t^i[R_{t+1}^{\text{ex}}]$, and $v = \text{Var}_t^i(R_{t+1}^{\text{ex}}) > 0$. Substituting $w' = Rw + nR^{\text{ex}}$ into the stated mean-variance objective gives

$$J(n) = Rw + n\mu - \frac{\nu n^2 v}{2Rw}, \qquad J''(n) = -\frac{\nu v}{Rw} < 0.$$

Thus the first-order condition is sufficient for the unique unconstrained optimum:

$$\mu - \frac{\nu n v}{Rw} = 0, \qquad \frac{n}{w} = \frac{R\mu}{\nu v} = \frac{E_t^i[z_{t+1}]}{\nu \,\text{var}_t^i(z_{t+1})}, \quad z_{t+1} \equiv \frac{R_{t+1}^{\text{ex}}}{R}. \tag{16}$$

The gross Home return cancels when both the mean and variance are expressed in terms of the normalized arithmetic carry return z_{t+1}.

For the log-return approximation, define

$$rx_{t+1} \equiv \Delta \log e_{t+1} + \log \frac{1+r_t^*}{1+r_t} = \log(1+z_{t+1}).$$

The maintained approximation replaces $z_{t+1} = \exp(rx_{t+1})-1$ by rx_{t+1} in its conditional mean and variance. This step is an additional return approximation; the exact mean-variance optimum is the preceding equation. Since interest rates are known at t, $\text{var}_t^i(rx_{t+1}) = \text{var}_t^i(\Delta \log e_{t+1})$. Members of a group share conditional moments, so their optimal positions are proportional to wealth. Averaging within the group yields

$$\frac{N_t^i}{W_t^i} = \frac{M_t^i}{\nu \,\text{var}_t^i(rx_{t+1})}, \qquad M_t^i \equiv E_t^i[rx_{t+1}]. \tag{17}$$

Equality here and below refers to the maintained log-return approximation. Under the common, fixed-variance normalization, $\nu \,\text{var}_t^i(rx_{t+1}) = \Gamma$ at the relevant order. In steady state $M^i = L_i$ and $N^i/W^i = L_i/\Gamma$. Since the first-order deviation of M_t^i is UIP_t^i, differentiating $N_t^i = W_t^i M_t^i/\Gamma$ gives

$$\widetilde{N}_t^i \equiv \frac{N_t^i - N^i}{W^i} = \frac{L_i}{\Gamma}\hat{W}_t^i + \frac{1}{\Gamma}UIP_t^i. \tag{18}$$

The objective also has a local certainty-equivalent interpretation. For CRRA utility u, let $w' = Rw + \epsilon Y$ and define $u(CE) = E_t^i[u(w')]$. A second-order expansion at Rw, holding ν fixed, gives

$$CE = Rw + \epsilon E_t^i[Y] - \frac{\nu \epsilon^2}{2Rw}\,\text{var}_t^i(Y) + o(\epsilon^2).$$

Here $-u''(Rw)/u'(Rw) = \nu/(Rw)$, including the log-utility case $\nu = 1$. We take the per-

period mean-variance objective as primitive; this expansion supplies its local CRRA interpretation. The choice takes current wealth and conditional moments as given and contains no continuation value, so intertemporal hedging demand is excluded by the myopia assumption. □

B.3 Proof of Lemma 1

Use the portfolio rule in Appendix B.2. Bond-market clearing gives

$$\frac{b_t}{1+r_t} + W_t - N_t = 0, \qquad \frac{e_t b_t^*}{1+r_t^*} + N_t = 0. \tag{19}$$

With M_t^i defined in the portfolio derivation, Home clearing implies

$$\sum_i \omega_t^i \frac{M_t^i}{\nu \operatorname{var}_t^i(rx_{t+1})} = \frac{N_t}{W_t} = 1 + \frac{b_t}{(1+r_t)W_t}.$$

This proves the level condition in Lemma 1.

At first order $\nu \operatorname{var}_t^i(rx_{t+1}) = \Gamma$. The deviation of $\sum_i \omega_t^i M_t^i$ is

$$\sum_i \omega^i UIP_t^i + \sum_i L_i \widetilde{\omega}_t^i = \mathbb{E}_t \hat{e}_{t+1} - \hat{e}_t + \hat{r}_t^* - \hat{r}_t + \delta_w d_t + d_{ss} \widetilde{\omega}_t^A.$$

This uses unbiased consensus beliefs, $\widetilde{\omega}_t^B = -\widetilde{\omega}_t^A$, $d_{ss} = L_A - L_B$, and $\delta_w = (\omega^A - \omega^B)/2$. The product of changing wealth shares and changing disagreement is second order and drops out. With average forecast bias, replace $\mathbb{E}_t \hat{e}_{t+1}$ by the simple average of the groups' forecasts.

The constant one in the Home-clearing identity has zero deviation. Differentiating the bond ratio, with $1 + r = 1/\beta$, gives

$$\Gamma \, \mathrm{d}\left[\frac{b_t}{(1+r_t)W_t}\right] = \frac{\beta\Gamma}{W}\left[y\widetilde{b}_t - b(\hat{W}_t + \hat{r}_t)\right]. \tag{20}$$

Equating the two deviations proves the first-order Home-bond UIP condition. The factor y converts the normalized bond deviation $\widetilde{b}_t$ back to Home bond face value, whereas $\hat{r}_t$ is the log deviation of the gross Home return.

Equivalence with the Foreign-bond version Foreign clearing implies the exact identity

$$1 + \frac{b_t}{(1+r_t)W_t} = \frac{N_t}{W_t} = -\frac{e_t b_t^*}{(1+r_t^*)W_t}.$$

The factor $e_t/(1+r_t^*)$ converts Foreign bond face value to Home-valued investment. Since the two ratios differ only by the constant one, their first-order deviations agree. Using $e=1$ and $\widetilde{b}_t^* = (b_t^* - b^*)/y^*$ gives the explicit equality

$$\begin{aligned}\frac{\beta\Gamma}{W}\left[y\widetilde{b}_t - b(\hat{W}_t + \hat{r}_t)\right] \\ = -\frac{\beta^*\Gamma}{W}\left[y^*\widetilde{b}_t^* + b^*(\hat{e}_t - \hat{r}_t^* - \hat{W}_t)\right].\end{aligned}$$

Consequently the equivalent Foreign-bond UIP condition is

$$\begin{aligned}\mathbb{E}_t\hat{e}_{t+1} - \hat{e}_t + \hat{r}_t^* - \hat{r}_t = -d_{ss}\widetilde{\omega}_t^A - \delta_w d_t \\ - \frac{\beta^*\Gamma}{W}\left[y^*\widetilde{b}_t^* + b^*(\hat{e}_t - \hat{r}_t^* - \hat{W}_t)\right].\end{aligned}$$

The equality follows from market clearing; the Home and Foreign bond positions cannot be varied independently in equilibrium. Both derivations remain valid at zero steady-state bond holdings because bond deviations are in levels. □

B.4 Proof of Lemma 2

At the steady state, $1+r = 1/\beta$, $1+r^* = 1/\beta^*$, $R^{\text{ex}} = r^* - r = 1/\beta^* - 1/\beta$, and $N^i/W^i = L_i/\Gamma$. Differentiate the group wealth law, holding the entrant endowment fixed, and divide by steady-state group wealth:

$$\hat{W}_{t+1}^i = (1-\xi)\left[\frac{\hat{W}_t^i + \hat{r}_t}{\beta} + R^{\text{ex}}\widetilde{N}_t^i + \frac{L_i}{\Gamma}\widetilde{R}_{t+1}^{\text{ex}}\right], \tag{21}$$

where

$$\widetilde{R}_{t+1}^{\text{ex}} = \frac{\hat{e}_{t+1} - \hat{e}_t + \hat{r}_t^*}{\beta^*} - \frac{\hat{r}_t}{\beta}.$$

The tilde on R^{ex} denotes the level deviation of the arithmetic carry return. Its coefficient is the steady-state carry position per unit of wealth, L_i/Γ, with no additional R^{ex} factor.

Since population weights are fixed, first-order aggregation and differentiation of the wealth share give

$$\hat{W}_t = \omega^A\hat{W}_t^A + \omega^B\hat{W}_t^B, \qquad \widetilde{\omega}_t^A = \omega^A\omega^B(\hat{W}_t^A - \hat{W}_t^B).$$

In particular, the wealth-share deviation is measured in levels. Set $\Delta W_t = \hat{W}_t^A - \hat{W}_t^B$. Then $\hat{W}_t^A = \hat{W}_t + \omega^B\Delta W_t$ and $\hat{W}_t^B = \hat{W}_t - \omega^A\Delta W_t$. Substituting these identities into the difference

of portfolio rules yields

$$\widetilde{N}_t^A - \widetilde{N}_t^B = \frac{d_{ss}}{\Gamma}\hat{W}_t + \frac{\omega^B L_A + \omega^A L_B}{\Gamma}\Delta W_t + \frac{d_t}{\Gamma}.$$

Under unbiased consensus beliefs in steady state, $L_A = \bar{L} + d_{ss}/2$ and $L_B = \bar{L} - d_{ss}/2$. The cross-weighted return is therefore

$$\omega^B L_A + \omega^A L_B = \bar{L} - \delta_w d_{ss}.$$

Subtract the group wealth equations and multiply by $\omega^A\omega^B$. The common risk-free-rate term cancels, leaving

$$\widetilde{\omega}_{t+1}^A = (1-\xi)\Bigg\{ \left[\frac{1}{\beta} + \frac{R^{\text{ex}}(\bar{L} - \delta_w d_{ss})}{\Gamma}\right]\widetilde{\omega}_t^A + \omega^A\omega^B\left[\frac{R^{\text{ex}}}{\Gamma}d_t + \frac{d_{ss}}{\Gamma}\left(R^{\text{ex}}\hat{W}_t + \widetilde{R}_{t+1}^{\text{ex}}\right)\right]\Bigg\}.$$

This is the general wealth-share law in the lemma. The term $-\delta_w d_{ss}$ in its persistence coefficient vanishes under symmetric steady-state wealth shares.

For the analytical limit, impose $\beta^* = \beta$, the symmetric steady-state shares $\omega^A = \omega^B = 1/2$, and the zero-NFA benchmark used for reduced UIP. Then $R^{\text{ex}} = \bar{L} = \delta_w = 0$, and Remark 1 gives $d_t = 0$. Define the centered first-order log carry return $\widetilde{rx}_{t+1} \equiv rx_{t+1} - \bar{L} = \beta\widetilde{R}_{t+1}^{\text{ex}} = \hat{e}_{t+1} - \hat{e}_t + \hat{r}_t^* - \hat{r}_t$. For this derivation, write $a \equiv (1-\xi)/\beta$. The share law becomes

$$\widetilde{\omega}_{t+1}^A = a\widetilde{\omega}_t^A + b_w\widetilde{rx}_{t+1}, \qquad b_w = \frac{(1-\xi)\omega^A\omega^B d_{ss}}{\beta\Gamma}.$$

Reduced UIP gives $\mathbb{E}_t\widetilde{rx}_{t+1} = -d_{ss}\widetilde{\omega}_t^A$. Since rates at t are known, the innovation in $\widetilde{rx}_{t+1}$ is $\varepsilon_{t+1}^e = \hat{e}_{t+1} - \mathbb{E}_t\hat{e}_{t+1}$. Consequently,

$$\widetilde{\omega}_{t+1}^A = \lambda\widetilde{\omega}_t^A + b_w\varepsilon_{t+1}^e, \qquad \lambda = a - b_w d_{ss} = \frac{1-\xi}{\beta}\left(1 - \frac{\omega^A\omega^B d_{ss}^2}{\Gamma}\right).$$

These are the reduced-law coefficients in the text, evaluated at $\beta^* = \beta$. Symmetry is part of the maintained analytical benchmark; it is not inferred solely from a zero steady-state carry return. □

B.5 Proof of Proposition 1 and Corollary 1

Maintain Remark 1, common persistence and innovation variance across countries for each shock type, zero Foreign bond-supply terms, $\beta^* = \beta$, equal entrant endowments, and $a = (1-\xi)/\beta < 1$. Then $\omega^A = \omega^B = 1/2$. The terms multiplying d_t vanish from the first-order UIP and wealth equations because $\delta_w = R^{\text{ex}} = 0$. Thus the reduced financial equations follow even before verifying that the centered gap is zero.

Reduced equilibrium and forward solution Household Euler equations and goods-market clearing imply $\hat{r}_t - \hat{r}_t^* = (1+S)i_t^f - S\mathbb{E}_t\Delta\hat{e}_{t+1}$, where i_t^f is the fundamental interest differential defined in Proposition 1. Combining this with reduced UIP gives

$$\hat{r}_t - \hat{r}_t^* = i_t^f + \frac{S}{1+S}d_{ss}\widetilde{\omega}_t^A, \qquad \mathbb{E}_t\Delta\hat{e}_{t+1} = i_t^f - \frac{d_{ss}}{1+S}\widetilde{\omega}_t^A.$$

Appendix B.4 gives $\widetilde{\omega}_{t+1}^A = \lambda\widetilde{\omega}_t^A + b_w\varepsilon_{t+1}^e$, with $b_w = ad_{ss}/(4\Gamma)$ and $\lambda = a - b_w d_{ss}$. For $0 < \rho_\beta, \rho_y, \lambda < 1$, the stationary forward solution is

$$\hat{e}_t = f_t + \kappa_\omega\widetilde{\omega}_t^A, \qquad f_t = -\mathbb{E}_t\sum_{j=0}^{\infty} i_{t+j}^f, \qquad \kappa_\omega = \frac{d_{ss}}{(1+S)(1-\lambda)}.$$

The explicit f_t in the text follows by summing the AR(1) processes for the discount-factor and endowment differentials. In particular, its supply component is $\sigma\phi_y(\hat{y}_t - \hat{y}_t^*)/(1+S)$, as used in the Backus--Smith proof.

Multiplier and parameter restrictions Taking innovations in the forward solution yields $\varepsilon_{t+1}^e = \varepsilon_{t+1}^f + \kappa_\omega b_w\varepsilon_{t+1}^e$. Let $h = d_{ss}^2/(4\Gamma)$, so $\lambda = a(1-h)$ and $b_w d_{ss} = ah$. For $0 < a < 1$ and $0 < \lambda < 1$,

$$g = \frac{ah}{(1+S)(1-a+ah)}, \qquad \Xi = \frac{(1+S)(1-a+ah)}{(1+S)(1-a)+Sah}, \qquad 0 < g < \frac{1}{1+S}.$$

Hence $1 < \Xi < (1+S)/S$ when $d_{ss} > 0$; at $d_{ss} = 0$, $\Xi = 1$. Writing $R = ah/(1-a)$ gives $\Xi = (1+S)(1+R)/(1+S+SR)$, which increases with R and decreases with S. Since R increases with d_{ss} and a, it decreases with Γ, ξ, and β. The goods-market comparative statics follow from $S = \sigma\chi\phi_e$ in the home-bias region $0 < \gamma < 1/2$. This proves the corollary within the stationary parameter region.

Constant exchange-rate forecast disagreement It remains to verify that the reduced solution is consistent with the information benchmark. At $G_x = 1$, $\varphi_x = 0$, and zero

interpretation innovations, $\tau^{xi}_{t|t} = \hat{x}_t + \mu^{xi}$. Under group i's perceived observation equation, $\mathbb{E}^i_t \hat{x}_{t+1} = \rho_x \tau^{xi}_{t|t} - \mu^{xi} = \rho_x \hat{x}_t + (\rho_x - 1)\mu^{xi}$. The fundamental forecast gap is therefore

$$\begin{aligned} C_f &\equiv E^A_t[f_{t+1}] - E^B_t[f_{t+1}] \\ &= -\frac{1}{1+S}\left[d^{\beta}_{ss} - d^{\beta^*}_{ss} + \sigma\phi_y(1-\rho_y)(d^{y}_{ss} - d^{y^*}_{ss})\right]. \end{aligned}$$

This gap is constant. The current wealth share and interest rates are observed by both groups. Taking their forecast difference in the realized share law $\widetilde{\omega}^A_{t+1} = a\widetilde{\omega}^A_t + b_w \widetilde{rx}_{t+1}$ therefore gives $E^A_t[\widetilde{\omega}^A_{t+1}] - E^B_t[\widetilde{\omega}^A_{t+1}] = b_w D_t$, where D_t is the total exchange-rate forecast gap. Taking the forecast difference in the affine exchange-rate solution now yields

$$D_t = C_f + \kappa_\omega b_w D_t, \qquad D_t = \Xi C_f.$$

Since $g = \kappa_\omega b_w < 1$, the total gap is constant, and its steady-state level must satisfy $d_{ss} = \Xi C_f$. Thus $d_t = D_t - d_{ss} = 0$. This is the consistency condition between the fundamental biases and the steady-state carry-return gap; they are not independent primitive parameters. It verifies Proposition 1 and Remark 1 for the maintained affine solution. Away from the joint zero-noise limit, zero disagreement innovations alone only make $d^x_t - d^x_{ss}$ decay at φ_x; exact constancy requires initialization at d^x_{ss}. □

B.6 Proof of Proposition 2

Maintain Remark 1, Proposition 1, and common shock persistence $\rho_\beta = \rho_y \equiv \rho_x \in (0,1)$. In particular, $d_t = 0$ and d_{ss} is constant over time at each interior point of the stationary parameter sequence.

For this proof, write $a \equiv (1-\xi)/\beta$. Fixed amplification is compatible with this limit. Fix $S > 0$ and $1 < \Xi < \Xi_{\max}$, and write $c_\Xi \equiv (1+S)(1-\Xi^{-1}) \in (0,1)$. For each $0 < \lambda < 1$, choose

$$a = \lambda + c_\Xi(1-\lambda), \qquad h \equiv \frac{d^2_{ss}}{4\Gamma} = \frac{c_\Xi(1-\lambda)}{a}.$$

Then $0 < a, h < 1$, $\lambda = a(1-h)$, $b_w d_{ss} = ah$, and the multiplier remains Ξ. For example, fix $\beta \in (0,1)$ and $d_{ss} > 0$, and set $\xi = 1 - a\beta$ and $\Gamma = d^2_{ss}/(4h)$, with fundamental biases satisfying the consistency condition in Appendix B.5. Thus λ can approach one with amplification and the steady-state forecast gap held fixed.

Proof. Let

$$\begin{aligned} e^w_t &\equiv \kappa_\omega \widetilde{\omega}^A_t, \qquad u^w_t \equiv (1-\lambda)e^w_t, \\ i^f_t &= \mathbb{E}_t \Delta f_{t+1} = -(1-\rho_x) f_t. \end{aligned}$$

With common shock persistence,

$$f_{t+1} = \rho_x f_t + \varepsilon^f_{t+1}, \qquad \mathbb{E}_t \varepsilon^f_{t+1} = 0, \qquad \mathrm{Var}(\varepsilon^f_{t+1}) = \sigma_f^2.$$

The reduced wealth-share law and the exchange-rate-innovation fixed point are

$$\widetilde{\omega}^A_{t+1} = \lambda \widetilde{\omega}^A_t + b_w \varepsilon^e_{t+1}, \qquad \varepsilon^e_{t+1} = \Xi \varepsilon^f_{t+1}.$$

Because

$$\Xi = \frac{1}{1 - \kappa_\omega b_w}, \qquad \kappa_\omega b_w \Xi = \Xi - 1,$$

the priced wealth state obeys

$$e^w_{t+1} = \lambda e^w_t + (\Xi - 1)\varepsilon^f_{t+1}. \tag{22}$$

Moreover,

$$\kappa_\omega = \frac{d_{ss}}{(1+S)(1-\lambda)}.$$

Writing $i_t = \widehat{r}_t - \widehat{r}^*_t$, part (i) of the reduced system gives

$$i_t = i^f_t + \frac{S}{1+S} d_{ss} \widetilde{\omega}^A_t = i^f_t + S u^w_t, \tag{23}$$

$$\widehat{e}_t = f_t + e^w_t. \tag{24}$$

Combining the last identity with the laws of motion for f_t and e^w_t yields

$$\Delta \widehat{e}_{t+1} = i^f_t - u^w_t + \Xi \varepsilon^f_{t+1}. \tag{25}$$

Thus $i^f_t - u^w_t$ is expected depreciation and $\Xi \varepsilon^f_{t+1}$ is the date-$(t+1)$ exchange-rate innovation.

For $0 < \rho_x, \lambda < 1$, the two predetermined components follow

$$i^f_{t+1} = \rho_x i^f_t - (1 - \rho_x)\varepsilon^f_{t+1},$$
$$u^w_{t+1} = \lambda u^w_t + (1 - \lambda)(\Xi - 1)\varepsilon^f_{t+1}.$$

Their stationary second moments are therefore

$$\frac{\mathrm{Var}(i_t^f)}{\sigma_f^2} = \frac{(1-\rho_x)^2}{1-\rho_x^2} = \frac{1-\rho_x}{1+\rho_x}, \tag{26}$$

$$\frac{\mathrm{Var}(u_t^w)}{\sigma_f^2} = \frac{(\Xi-1)^2(1-\lambda)^2}{1-\lambda^2} = (\Xi-1)^2\frac{1-\lambda}{1+\lambda}, \tag{27}$$

$$\frac{\mathrm{Cov}(i_t^f, u_t^w)}{\sigma_f^2} = -\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda}. \tag{28}$$

Because (i_t^f, u_t^w) are predetermined relative to ε_{t+1}^f, equation (25) and (26)–(28) give

$$\begin{aligned}\frac{\mathrm{Var}(\Delta\widehat{e}_{t+1})}{\sigma_f^2} &= \Xi^2 + \frac{1-\rho_x}{1+\rho_x} + (\Xi-1)^2\frac{1-\lambda}{1+\lambda} \\ &\quad + 2\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda}.\end{aligned} \tag{29}$$

The exact first autocovariance is

$$\begin{aligned}&\frac{\mathrm{Cov}(\Delta\widehat{e}_{t+1}, \Delta\widehat{e}_t)}{\sigma_f^2} \\ &= -\frac{1-\rho_x}{1+\rho_x} - (\Xi-1)^2\frac{1-\lambda}{1+\lambda} - (\Xi-1)\frac{(1-\rho_x)^2+(1-\lambda)^2}{1-\rho_x\lambda}.\end{aligned} \tag{30}$$

Consequently,

$$\mathrm{AC}_1(\Delta\widehat{e}) = -\frac{\frac{1-\rho_x}{1+\rho_x} + (\Xi-1)^2\frac{1-\lambda}{1+\lambda} + (\Xi-1)\frac{(1-\rho_x)^2+(1-\lambda)^2}{1-\rho_x\lambda}}{\Xi^2 + \frac{1-\rho_x}{1+\rho_x} + (\Xi-1)^2\frac{1-\lambda}{1+\lambda} + 2\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda}}. \tag{31}$$

As $\rho_x\lambda \to 1$, both $\rho_x \to 1$ and $\lambda \to 1$. Every term in the numerator of (31) then converges to zero, whereas its denominator converges to $\Xi^2 > 0$. More strongly,

$$\begin{aligned}\frac{\mathrm{Var}\left(\mathbb{E}_t\Delta\widehat{e}_{t+1}\right)}{\sigma_f^2} &= \frac{\mathrm{Var}(i_t^f - u_t^w)}{\sigma_f^2} \\ &= \frac{1-\rho_x}{1+\rho_x} + (\Xi-1)^2\frac{1-\lambda}{1+\lambda} + 2\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda} \longrightarrow 0.\end{aligned}$$

Hence $(\Delta\widehat{e}_{t+1} - \Xi\varepsilon_{t+1}^f)/\sigma_f \to 0$ in mean square. After normalization by the fundamental-news standard deviation, exchange-rate changes therefore approach iid innovations at every fixed horizon, which is the precise sense in which the exchange-rate level approaches a random

walk. This proves part 1.

Next, equation (23) gives

$$\frac{\mathrm{Var}(\widehat{r}_t - \widehat{r}_t^*)}{\sigma_f^2} = \frac{1-\rho_x}{1+\rho_x} + S^2(\Xi-1)^2\frac{1-\lambda}{1+\lambda} - 2S\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda}. \tag{32}$$

Combining (32) with (29) produces the exact ratio

$$\frac{\mathrm{sd}(\widehat{r}_t - \widehat{r}_t^*)}{\mathrm{sd}(\Delta\widehat{e}_{t+1})} = \left[\frac{\frac{1-\rho_x}{1+\rho_x} + S^2(\Xi-1)^2\frac{1-\lambda}{1+\lambda} - 2S\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda}}{\Xi^2 + \frac{1-\rho_x}{1+\rho_x} + (\Xi-1)^2\frac{1-\lambda}{1+\lambda} + 2\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda}}\right]^{1/2}. \tag{33}$$

Its numerator converges to zero and its denominator to Ξ^2 under joint persistence. This proves part 2.

For part 3, define the population Fama regression by

$$\Delta\widehat{e}_{t+1} = \alpha + \beta_F(\widehat{r}_t - \widehat{r}_t^*) + u_{t+1}.$$

Equations (23) and (25) imply

$$\frac{\mathrm{Cov}(\Delta\widehat{e}_{t+1}, \widehat{r}_t - \widehat{r}_t^*)}{\sigma_f^2} = \frac{1-\rho_x}{1+\rho_x} - (S-1)\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda} - S(\Xi-1)^2\frac{1-\lambda}{1+\lambda}. \tag{34}$$

Dividing (34) by (32) yields the exact coefficient

$$\boxed{\begin{aligned}\beta_F = &\left[\frac{1-\rho_x}{1+\rho_x} - (S-1)\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda} - S(\Xi-1)^2\frac{1-\lambda}{1+\lambda}\right] \\ &\times \left[\frac{1-\rho_x}{1+\rho_x} + S^2(\Xi-1)^2\frac{1-\lambda}{1+\lambda} - 2S\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda}\right]^{-1}.\end{aligned}} \tag{35}$$

The exact coefficient of determination is

$$R^2 = \left[\frac{1-\rho_x}{1+\rho_x} - (S-1)\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda} - S(\Xi-1)^2\frac{1-\lambda}{1+\lambda}\right]^2 \times \left[\frac{1-\rho_x}{1+\rho_x} + S^2(\Xi-1)^2\frac{1-\lambda}{1+\lambda} - 2S\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda}\right]^{-1} \times \left[\Xi^2 + \frac{1-\rho_x}{1+\rho_x} + (\Xi-1)^2\frac{1-\lambda}{1+\lambda} + 2\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda}\right]^{-1}. \tag{36}$$

Suppose $1-\rho_x = o(1-\lambda)$; this is the statement that ρ_x approaches one faster than λ. Then

$$\frac{1-\rho_x}{1+\rho_x} = o(1-\lambda), \qquad \frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda} = o(1-\lambda).$$

The terms proportional to $(\Xi-1)^2(1-\lambda)/(1+\lambda)$ therefore dominate both the numerator and denominator of (35), giving

$$\beta_F \longrightarrow -\frac{1}{S} < 0.$$

Thus β_F is eventually negative at any fixed amplification factor $\Xi > 1$.

Finally, because $\widehat{r}_t - \widehat{r}_t^*$ is date-t measurable, the variance explained by the Fama regression cannot exceed the variance of expected depreciation. Hence

$$R^2 \le \frac{\mathrm{Var}(\mathbb{E}_t \Delta\widehat{e}_{t+1})}{\mathrm{Var}(\Delta\widehat{e}_{t+1})} = \frac{\frac{1-\rho_x}{1+\rho_x} + (\Xi-1)^2\frac{1-\lambda}{1+\lambda} + 2\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda}}{\Xi^2 + \frac{1-\rho_x}{1+\rho_x} + (\Xi-1)^2\frac{1-\lambda}{1+\lambda} + 2\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda}} \longrightarrow 0.$$

Therefore $R^2 \to 0$ along every joint-persistence path for which the Fama regressor is non-degenerate. At the exact endpoint where its variance is zero, the regression itself is undefined. □

B.7 Proof of Proposition 3

Maintain the same constant-disagreement benchmark as in Appendix B.6. Here $S > 0$ is fixed, $0 < \rho_x, \lambda < 1$, and the sequence satisfies $\rho_x\lambda \to 1$ and $\Xi \to \Xi_{\max}$. The primitive construction and exact moment identities in that appendix apply at every interior parameterization. Direct news below is ε_{t+1}^f, the innovation in the normalized fundamental component f_t.

Proof. Decompose direct fundamental news as

$$\varepsilon^f_{t+1} = \varepsilon^{f,y}_{t+1} + \varepsilon^{f,\perp}_{t+1}, \qquad \mathrm{Cov}(\varepsilon^{f,y}_{t+1}, \varepsilon^{f,\perp}_{t+1}) = 0,$$

where

$$s_y \equiv \frac{\mathrm{Var}(\varepsilon^{f,y}_{t+1})}{\sigma_f^2}, \qquad 0 < s_y < 1.$$

Let f^y_t and $i^{f,y}_t = -(1-\rho_x) f^y_t$ denote the supply components of f_t and i^f_t. The supply block satisfies

$$f^y_t = \frac{\sigma \phi_y}{1+S}(\widehat{y}_t - \widehat{y}^*_t).$$

Since $f^y_{t+1} = \rho_x f^y_t + \varepsilon^{f,y}_{t+1}$, it follows that

$$\phi_y \Delta(\widehat{y} - \widehat{y}^*)_{t+1} = \frac{S}{\sigma} \Xi_{\max}(i^{f,y}_t + \varepsilon^{f,y}_{t+1}), \qquad \Xi_{\max} = \frac{1+S}{S}. \tag{37}$$

Goods-market clearing then gives

$$\begin{aligned} \Delta(\widehat{c} - \widehat{c}^*)_{t+1} &= \phi_y \Delta(\widehat{y} - \widehat{y}^*)_{t+1} - \frac{S}{\sigma} \Delta \widehat{e}_{t+1} \\ &= \frac{S}{\sigma} \left[\Xi_{\max}(i^{f,y}_t + \varepsilon^{f,y}_{t+1}) - i^f_t + u^w_t - \Xi \varepsilon^f_{t+1} \right]. \end{aligned} \tag{38}$$

Common persistence and innovation orthogonality imply

$$\begin{aligned} \frac{\mathrm{Var}(i^{f,y}_t)}{\sigma_f^2} &= s_y \frac{1-\rho_x}{1+\rho_x}, \\ \frac{\mathrm{Cov}(i^{f,y}_t, u^w_t)}{\sigma_f^2} &= -s_y \frac{(\Xi - 1)(1-\rho_x)(1-\lambda)}{1 - \rho_x \lambda}. \end{aligned}$$

Using these identities, equations (25) and (38) give the following three exact moments:

$$\begin{aligned}
&\frac{\operatorname{Var}(\Delta\widehat{e}_{t+1})}{\sigma_f^2} \\
&\quad = \Xi^2 + \frac{1-\rho_x}{1+\rho_x} + (\Xi-1)^2\frac{1-\lambda}{1+\lambda} + 2\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda}, \qquad (39)\\
&\frac{\sigma}{S\sigma_f^2}\operatorname{Cov}(\Delta(\widehat{c}-\widehat{c}^*)_{t+1}, \Delta\widehat{e}_{t+1}) \\
&\quad = \Xi_{\max}s_y\left[\Xi + \frac{1-\rho_x}{1+\rho_x} + \frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda}\right] \\
&\qquad - \left[\Xi^2 + \frac{1-\rho_x}{1+\rho_x} + (\Xi-1)^2\frac{1-\lambda}{1+\lambda} + 2\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda}\right], \qquad (40)\\
&\frac{\sigma^2}{S^2\sigma_f^2}\operatorname{Var}(\Delta(\widehat{c}-\widehat{c}^*)_{t+1}) \\
&\quad = \Xi^2 + \frac{1-\rho_x}{1+\rho_x} + (\Xi-1)^2\frac{1-\lambda}{1+\lambda} + 2\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda} \\
&\qquad + \Xi_{\max}^2 s_y\left(1 + \frac{1-\rho_x}{1+\rho_x}\right) \\
&\qquad - 2\Xi_{\max}s_y\left[\Xi + \frac{1-\rho_x}{1+\rho_x} + \frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda}\right]. \qquad (41)
\end{aligned}$$

Therefore the exact Backus–Smith correlation is

$$\begin{aligned}
&\operatorname{Corr}(\Delta(\widehat{c}-\widehat{c}^*)_{t+1}, \Delta\widehat{e}_{t+1}) \\
&= \left\{\Xi_{\max}s_y\left[\Xi + \frac{1-\rho_x}{1+\rho_x} + \frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda}\right]\right. \\
&\qquad \left. - \left[\Xi^2 + \frac{1-\rho_x}{1+\rho_x} + (\Xi-1)^2\frac{1-\lambda}{1+\lambda} + 2\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda}\right]\right\} \\
&\quad \times \left\{\left[\Xi^2 + \frac{1-\rho_x}{1+\rho_x} + (\Xi-1)^2\frac{1-\lambda}{1+\lambda} + 2\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda}\right]\right. \\
&\qquad \times \left[\Xi^2 + \frac{1-\rho_x}{1+\rho_x} + (\Xi-1)^2\frac{1-\lambda}{1+\lambda} + 2\frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda}\right. \\
&\qquad\quad + \Xi_{\max}^2 s_y\left(1 + \frac{1-\rho_x}{1+\rho_x}\right) \\
&\qquad\quad \left.\left. - 2\Xi_{\max}s_y\left[\Xi + \frac{1-\rho_x}{1+\rho_x} + \frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda}\right]\right]\right\}^{-1/2}. \qquad (42)
\end{aligned}$$

Along the joint-persistence sequence,

$$\frac{1-\rho_x}{1+\rho_x} \to 0, \qquad (\Xi-1)^2\frac{1-\lambda}{1+\lambda} \to 0, \qquad \frac{(\Xi-1)(1-\rho_x)(1-\lambda)}{1-\rho_x\lambda} \to 0.$$

Before imposing the multiplier ceiling, the three moments above converge to

$$\begin{aligned}
\frac{\mathrm{Var}(\Delta\widehat{e})}{\sigma_f^2} &\longrightarrow \Xi^2, \\
\frac{\sigma}{S\sigma_f^2}\,\mathrm{Cov}(\Delta(\widehat{c}-\widehat{c}^*),\Delta\widehat{e}) &\longrightarrow \Xi\Xi_{\max}s_y - \Xi^2, \\
\frac{\sigma^2}{S^2\sigma_f^2}\,\mathrm{Var}(\Delta(\widehat{c}-\widehat{c}^*)) &\longrightarrow \Xi^2 - 2\Xi\Xi_{\max}s_y + \Xi_{\max}^2 s_y.
\end{aligned}$$

Hence

$$\mathrm{Corr}(\Delta(\widehat{c}-\widehat{c}^*),\Delta\widehat{e}) \longrightarrow \frac{\Xi_{\max}s_y - \Xi}{\sqrt{\Xi^2 - 2\Xi\Xi_{\max}s_y + \Xi_{\max}^2 s_y}}. \tag{43}$$

Finally set $\Xi \to \Xi_{\max}$. For fixed $0 < s_y < 1$, the numerator of (43) converges to $-\Xi_{\max}(1-s_y)$, while its denominator converges to $\Xi_{\max}\sqrt{1-s_y}$. Therefore

$$\mathrm{Corr}(\Delta(\widehat{c}-\widehat{c}^*),\Delta\widehat{e}) \longrightarrow -\sqrt{1-s_y}. \qquad \square$$

C Disagreement and the Backus–Smith correlation

This appendix separates two results. At the reported calibration, disagreement strengthens the negative covariance generated by demand shocks enough to reverse the Backus–Smith correlation in the baseline model. In the reduced equilibrium, sufficiently strong feedback can also generate a negative unconditional correlation with supply shocks alone, through the subsequent adjustment of the wealth share. The second result is a possibility result, not a property of the reported calibration. Throughout, the statistic is

$$\mathrm{Corr}(\Delta\hat{e}_t, \Delta(\hat{c}_t - \hat{c}_t^*)),$$

where an increase in e_t is a Home depreciation.

C.1 The calibrated economies and the fixed-shock comparison

Table C1 compares the reported correlations with their supply-conditional counterparts. The supply block retains output fundamentals and their associated signal and interpretation

innovations; demand and UIP innovations are set to zero. Every supply-conditional correlation is positive. Thus these calibrated economies do not explain the negative sign using the supply block alone.

Table C1: Backus–Smith correlations in the reported economies

Specification	All shocks	Supply block only
Baseline	−0.244	0.982
No disagreement, fixed parameters	0.546	0.998
No disagreement, re-estimated shocks	−0.131	0.997
Augmented model	−0.471	0.739
No UIP risk	−0.320	0.739
No average bias	−0.327	0.980
No average bias or UIP risk	−0.071	0.980

Notes: Entries are correlations between quarterly Home depreciation and Home-minus-Foreign consumption growth. Baseline has Bayesian consensus beliefs and no exogenous UIP shock. The augmented model adds average belief distortions and UIP risk; its ablations hold other augmented parameters fixed. No-disagreement economies remove permanent and transitory disagreement from the baseline, holding other parameters fixed or re-estimating the fundamental shock block as indicated. Supply-only activates endowment fundamentals and their associated signal-noise and interpretation innovations. Values average 1,000 simulations of 116 quarters after a 1,000-quarter burn-in. Removing UIP risk has no effect in the supply-only experiment.

To isolate disagreement, Table C2 compares each economy with its own no-disagreement counterpart. Within each pair, preferences, technology, aggregate steady-state arbitrageur wealth, the absolute risk-cost parameter Γ, fundamental shock processes, signal noise, average beliefs, and the UIP process are held fixed. The counterfactual sets $d_{ss} = 0$ and the transitory disagreement innovations to zero. Steady-state objects are recomputed; as in the main counterfactual, the entrant endowment adjusts to preserve aggregate steady-state wealth. No parameters are re-estimated.

Table C2: The effect of disagreement at fixed shock parameters

Parameterization	Without disagreement	With disagreement	Change
Baseline	0.546	−0.244	−0.790
Augmented parameters, no UIP risk	0.657	−0.320	−0.977
Augmented parameters, with UIP risk	−0.599	−0.471	0.128

Notes: Entries are correlations between quarterly Home depreciation and Home-minus-Foreign consumption growth, averaging 1,000 simulations of 116 quarters after a 1,000-quarter burn-in. Each row compares an economy with its own no-disagreement counterpart at unchanged preferences, technology, aggregate arbitrageur wealth, risk-cost parameter, and shock processes. The no-disagreement counterpart removes both the permanent belief gap and its innovations and recomputes the steady state. Change is With minus Without disagreement. The augmented rows use the augmented calibration, which differs from the baseline.

In the baseline, disagreement raises demand-driven exchange-rate volatility by a factor of

3.13, compared with 2.05 for supply-driven volatility. The magnitude of the isolated demand covariance rises by a factor of 9.83, whereas the positive supply covariance rises by only 1.41. Demand-only Backus–Smith is already -1 in both economies: with output fixed, goods-market clearing makes relative consumption growth a negative multiple of depreciation. The effect therefore concerns the relative strength of demand and supply in the mixed-shock economy.

A further experiment turns off transitory disagreement in both economies and retains the permanent belief gap in one. The resulting linear system permits an exact population covariance decomposition across the independent demand and supply blocks. In units of 10^{-4}, the demand covariance changes from -0.163 without disagreement to -1.509 with disagreement; the supply covariance changes from 0.690 to 0.976. Their sum changes from 0.527 to -0.533, and the population correlation falls from 0.548 to -0.212. Permanent disagreement alone therefore suffices for the sign reversal at unchanged fundamental shock parameters. This additive decomposition is separate from the preceding simulations, which retain the wealth–disagreement cross terms.

The last row of Table C2 also identifies the scope of the result. With UIP risk active, disagreement reduces UIP-driven exchange-rate volatility from 3.682% to 1.678%. Because UIP-only Backus–Smith is also -1, this reduces an additional source of negative covariance. The augmented model's overall correlation then becomes less negative despite stronger demand amplification. Disagreement need not lower the correlation for every shock composition.

C.2 How supply shocks alone can produce a negative correlation

Consider the reduced equilibrium with a constant belief gap and supply news only, so $s_y = 1$. Goods-market clearing gives

$$\Delta(\hat{c} - \hat{c}^*)_{t+1} = \phi_y \Delta(\hat{y} - \hat{y}^*)_{t+1} - \frac{S}{\sigma}\Delta\hat{e}_{t+1}.$$

The first term is the direct income effect; the second is expenditure switching induced by exchange-rate movements. On impact, their covariance with the exchange-rate innovation is

$$\mathrm{Cov}_t(\Delta(\hat{c} - \hat{c}^*)_{t+1}, \Delta\hat{e}_{t+1}) = \frac{S\sigma_f^2}{\sigma}\Xi(\Xi_{\max} - \Xi) > 0.$$

Thus amplification cannot reverse the impact comovement under the multiplier ceiling. The unconditional correlation also incorporates adjustments following past shocks.

Setting $s_y = 1$ in equation (40) gives the exact unconditional covariance:

$$\begin{aligned}
&\frac{\sigma}{S\sigma_f^2}\operatorname{Cov}(\Delta(\hat{c}-\hat{c}^*),\Delta\hat{e}) \\
&\quad = \Xi(\Xi_{\max}-\Xi) + (\Xi_{\max}-1)\frac{1-\rho_y}{1+\rho_y} \\
&\qquad - (\Xi-1)^2\frac{1-\lambda}{1+\lambda} + (\Xi_{\max}-2)\frac{(\Xi-1)(1-\rho_y)(1-\lambda)}{1-\rho_y\lambda}.
\end{aligned} \tag{44}$$

Here $\rho_x = \rho_y$ because supply is the only driving force. The correlation is negative if and only if the right-hand side is negative. Its sign depends on both amplification and the relative persistence of supply and wealth, not on Ξ alone.

To show that this region is nonempty, recall the restrictions imposed by the symmetric reduced equilibrium:

$$\lambda = a\left(1-\frac{d_{ss}^2}{4\Gamma}\right), \qquad \Xi = \frac{(1+S)(1-\lambda)}{1-a+S(1-\lambda)}, \qquad a = \frac{1-\xi}{\beta} < 1.$$

Fix $S > 0$ and $0 < \lambda < 1$, and take $a \uparrow 1$ and $\rho_y \uparrow 1$, with $d_{ss}^2/(4\Gamma) = 1 - \lambda/a$. This respects the coefficient restrictions at every interior point and implies $\Xi \uparrow \Xi_{\max}$. Equation (44) then converges to

$$-(\Xi_{\max}-1)^2\frac{1-\lambda}{1+\lambda} < 0.$$

By continuity, stationary economies sufficiently close to this limit have negative supply-only Backus–Smith correlations. This is a different limit from Proposition 3: supply becomes highly persistent while wealth-share persistence remains strictly below one.

The adjustment mechanism is visible in the impulse responses. Following a unit supply innovation $\varepsilon_t^f = 1$ from steady state, the responses at horizon $j \geq 0$ are

$$\hat{e}_{t+j} = \rho_y^j + (\Xi-1)\lambda^j, \qquad \frac{\sigma}{S}(\hat{c}-\hat{c}^*)_{t+j} = (\Xi_{\max}-1)\rho_y^j - (\Xi-1)\lambda^j.$$

When wealth adjusts faster than supply, the amplified exchange-rate component unwinds while the income effect remains. This unwinding moves the exchange rate and relative consumption in opposite directions. If feedback is strong enough, that negative covariance outweighs the positive impact covariance. Expenditure switching thus creates demand-like movements following supply news, without an additional exogenous demand shock.

C.3 A numerical illustration in the reduced equilibrium

Table C3 varies only the permanent belief gap within a separate numerical illustration. It holds fixed $\beta = \beta^* = 0.995$, $\xi = 0.01495$, $\Gamma = 0.001$, $\sigma = 5$, $\chi = 0.6$, $\theta = 1.5$, $\gamma = 0.05$, and $\rho_y = 0.95$, with symmetric wealth shares. These imply $a = 0.99$, $S = 1.0556$, and $\Xi_{\max} = 1.9474$. Only fundamental supply innovations are active; their variance is unchanged and cancels from the correlation. There are no demand, signal-noise, transitory disagreement, or UIP innovations.

Table C3: Supply-only Backus–Smith under stronger feedback

d_{ss}	λ	Ξ	Impact correlation	Unconditional correlation
0.000	0.990	1.000	1.000	1.000
0.015	0.934	1.702	1.000	0.994
0.030	0.767	1.871	1.000	0.140
0.040	0.594	1.903	1.000	−0.130
0.050	0.371	1.918	1.000	−0.265
0.060	0.099	1.927	1.000	−0.378

Exact population moments of the reduced equilibrium, calculated from equations (39)–(41) with $s_y = 1$ and independently verified using impulse-response sums. Every row satisfies $0 < \lambda < 1$ and $1 \leq \Xi < \Xi_{\max}$. This illustration uses a separate parameterization from the quantitative baseline and augmented model.

Increasing d_{ss} strengthens amplification and accelerates wealth-share adjustment. The unconditional correlation changes sign between $d_{ss} = 0.03$ and 0.04, while the impact correlation remains positive throughout. The example establishes that endogenous expenditure switching can overturn supply-driven comovement over time. Table C1 shows that the reported calibration instead reproduces the negative sign through the demand channel.

D Augmented model: calibration and currency-return predictability

D.1 Augmented calibration

Table D1 records the separate augmented calibration. The fixed-parameter comparisons below remove the additional wedges at these values; they do not substitute the baseline calibration.

Table D1: Augmented-model calibration

Parameter	Value	Economic role
W	0.0777	Steady-state arbitrageur wealth/output
ξ	9.73×10^{-3}	Exit probability
Γ	4.11×10^{-3}	Effective risk-cost parameter
d_{ss}	0.0448	Steady-state one-quarter exchange-rate forecast gap
G	0.318	Common Kalman gain
σ_μ	1.41×10^{-3}	Common interpretation-innovation SD
ρ_β	0.9861	Discount-factor persistence
$\sigma_{\varepsilon^\beta}$	7.17×10^{-4}	Discount-factor persistent-innovation SD
$\rho_{\beta\beta^*}$	0.219	Cross-country discount-factor innovation correlation
ρ_y	0.9867	Endowment persistence
σ_{ε^y}	9.15×10^{-3}	Endowment persistent-innovation SD
ρ_{yy^*}	0.691	Cross-country endowment innovation correlation
$\kappa_y, \widetilde{\rho}_y$	0.259, 0.9995	Perceived gain and persistence for endowments
$\kappa_\beta, \widetilde{\rho}_\beta$	0.969, 0.9924	Perceived gain and persistence for discount factors
ρ_ψ	0.85	Exogenous UIP-shock persistence
σ_ψ	0.0138	Exogenous UIP innovation SD

Notes: The augmented model adds average belief distortions and an exogenous uncovered-interest-parity (UIP) shock to the baseline, with parameters calibrated separately. Home is the United States; Foreign is its trade-weighted advanced-economy counterpart. The common external choices are $\gamma = 0.05$, $\sigma = 5$, $\theta = 1.5$, $\chi = 0.60$, $\beta = 0.9980$, $\beta^* = 0.9975$, and equal population shares. Baseline estimates and the common macroeconomic and survey disciplines appear in Table 1. Average-belief parameters discipline the currency-return predictability profile; the UIP innovation scale matches exchange-rate level volatility and its persistence is fixed from the excess bond premium. SD means standard deviation; all shock scales and forecast gaps are in log units. Values are rounded for display, including persistence parameters whose closeness to one is economically relevant.

D.2 Equivalence to the Candian–De Leo belief process

This subsection establishes the equivalence of consensus posterior means and conditional-mean forecasts to the belief process in Candian and De Leo (2025). Fix a fundamental x and write $m_t^x \equiv \bar{\tau}_{t|t}^x$ for the simple average posterior in this paper. Averaging the two groups' filtering rules gives

$$m_t^x = \varphi_x m_{t-1}^x + G_x(\hat{x}_t + \bar{\mu}_t^x), \qquad \varphi_x \equiv \rho_x(1 - G_x). \tag{45}$$

In the Candian–De Leo perceived model, the observed fundamental is the sum of a persistent component z_t^x and an unpredictable transitory component $\epsilon_t^{x,\mathrm{CDL}}$:

$$\hat{x}_t = z_t^x + \epsilon_t^{x,\mathrm{CDL}}, \qquad z_t^x = \tilde{\rho}_x z_{t-1}^x + u_t^x.$$

With steady-state Kalman gain κ_x, the posterior mean $m_t^{x,\mathrm{CDL}}$ and forecasts satisfy

$$m_t^{x,\mathrm{CDL}} = \tilde{\rho}_x m_{t-1}^{x,\mathrm{CDL}} + \kappa_x\big(\hat{x}_t - \tilde{\rho}_x m_{t-1}^{x,\mathrm{CDL}}\big) = \delta_x m_{t-1}^{x,\mathrm{CDL}} + \kappa_x \hat{x}_t, \tag{46}$$

$$E_t^{\mathrm{CDL}} \hat{x}_{t+h} = \tilde{\rho}_x^h m_t^{x,\mathrm{CDL}}, \qquad h \geq 1, \qquad \delta_x \equiv (1 - \kappa_x)\tilde{\rho}_x. \tag{47}$$

Proposition 5. *Suppose $G_x \in (0, 1]$, $\kappa_x \in (0, 1]$, and the two representations start from the same consensus posterior. The average interpretation rule*

$$\bar{\mu}_t^x = \left(\frac{\kappa_x}{G_x} - 1\right)\hat{x}_t + \frac{(1 - \kappa_x)\tilde{\rho}_x - \rho_x(1 - G_x)}{G_x} m_{t-1}^x \tag{48}$$

uniquely implements the Candian–De Leo consensus posterior for every observation history. If consensus forecasts also use the perceived persistence, $\bar{E}_t \hat{x}_{t+h} = \tilde{\rho}_x^h m_t^x$, the two representations generate identical conditional-mean forecasts and realized forecast errors at every horizon.

Proof. Substituting (48) into (45) yields

$$\begin{aligned} m_t^x &= \varphi_x m_{t-1}^x + G_x \hat{x}_t + (\kappa_x - G_x)\hat{x}_t + (\delta_x - \varphi_x) m_{t-1}^x \\ &= \delta_x m_{t-1}^x + \kappa_x \hat{x}_t. \end{aligned}$$

This is exactly (46). Equality of the initial posteriors therefore implies $m_t^x = m_t^{x,\mathrm{CDL}}$ at every date by induction, for any common sequence of observations. Conversely, equating the desired posterior update with (45) requires

$$G_x \bar{\mu}_t^x = (\kappa_x - G_x)\hat{x}_t + (\delta_x - \varphi_x) m_{t-1}^x.$$

Because $G_x > 0$, this uniquely determines the average interpretation rule along every such history. Applying the same perceived transition to the equal posterior means gives $\bar{E}_t \hat{x}_{t+h} = E_t^{\mathrm{CDL}} \hat{x}_{t+h}$ for all $h \geq 1$. Subtracting these equal forecasts from the same subsequent realization establishes equality of forecast errors. If $|\delta_x| < 1$, any initial posterior discrepancy instead decays geometrically at rate δ_x. □

The corresponding forecast-error dynamics can also be seen directly. Let $f_t^x \equiv \bar{E}_t \hat{x}_{t+1} = \tilde{\rho}_x m_t^x$ and $q_t^x \equiv \hat{x}_t - f_{t-1}^x$. The posterior recursion implies

$$\begin{aligned} f_t^x &= \delta_x f_{t-1}^x + \kappa_x \tilde{\rho}_x \hat{x}_t, \\ q_t^x - \delta_x q_{t-1}^x &= \hat{x}_t - \tilde{\rho}_x \hat{x}_{t-1}. \end{aligned}$$

These identities hold for any observation process. In the special case in which the true observed fundamental follows $\hat{x}_t = \rho_x \hat{x}_{t-1} + \varepsilon_t^x$, applying $1 - \rho_x L$, where L is the lag operator, gives the Candian–De Leo ARMA$(2,1)$ representation:

$$q_t^x = (\rho_x + \delta_x) q_{t-1}^x - \rho_x \delta_x q_{t-2}^x + \varepsilon_t^x - \tilde{\rho}_x \varepsilon_{t-1}^x. \tag{49}$$

The construction leaves a separate role for disagreement. For any prescribed interpretation difference $\Delta\mu_t^x \equiv \mu_t^{xA} - \mu_t^{xB}$, set

$$\mu_t^{xA} = \bar{\mu}_t^x + \tfrac{1}{2}\Delta\mu_t^x, \qquad \mu_t^{xB} = \bar{\mu}_t^x - \tfrac{1}{2}\Delta\mu_t^x.$$

Their average is (48), while their difference remains $\Delta\mu_t^x$. The common component cancels from the posterior-gap recursion:

$$d_t^x = \varphi_x d_{t-1}^x + G_x \Delta\mu_t^x, \qquad d_t^x \equiv \tau_{t|t}^{xA} - \tau_{t|t}^{xB}.$$

Thus the equivalence concerns consensus posterior means and forecasts, conditional on the same observations; it neither eliminates heterogeneous beliefs nor asserts equivalence of the two general-equilibrium models. The average interpretation wedge is a function of current observations and the inherited consensus posterior, so it introduces no new primitive shock. Matching posterior updating alone would not suffice for forecast equivalence if forecasts continued to use ρ_x in place of $\tilde{\rho}_x$.

D.3 Predictability across horizons

This appendix compares the augmented model with three ablations: removing the UIP shock, removing average bias, and removing both. All other parameters are held at the augmented-model calibration, and both wealth–disagreement cross terms are retained in the simulations. These comparisons differ from the separately calibrated baseline model. Each regression coefficient is averaged over 1,000 samples of 116 quarters following a 1,000-quarter burn-in, using common random innovations across specifications.

Let $i_t \equiv \hat{r}_t - \hat{r}_t^*$ denote the quarterly Home-minus-Foreign interest differential and define the foreign-currency excess return by $\widetilde{rx}_{t+1} \equiv \Delta\hat{e}_{t+1} - i_t$. For each horizon, I estimate

$$\widetilde{rx}_{t+h} = a_h + b_h i_t + u_{t+h}, \qquad h \geq 1,$$
$$\sum_{j=1}^{H} \widetilde{rx}_{t+j} = A_H + B_H i_t + v_{t,H}, \qquad H \geq 1.$$

The first regression measures the predictability of a one-quarter return at horizon h; the second measures the predictability of returns accumulated over H quarters. With this sign convention, a negative coefficient means that a high Home interest rate predicts Home currency outperformance. Each cumulative regression uses the same set of starting dates for all returns entering its dependent variable. It therefore need not equal the sum of the reported horizon-specific coefficients, which use different numbers of observations. These are predictive regressions, distinct from the structural shock responses in Figure 7.

Figure 6 shows that the augmented-model one-quarter coefficient first becomes positive at six quarters, compared with eight quarters without average bias and 28 quarters without either average bias or the UIP shock. The augmented-model cumulative coefficient first becomes positive at 29 quarters; without the UIP shock it does so at 16 quarters. Both specifications without average bias retain negative cumulative coefficients throughout the first 40 quarters. Average bias therefore advances the reversal and changes its cumulative sign over this horizon. The result is a finite-horizon, finite-sample comparison; it does not establish a positive infinite-horizon cumulative coefficient.

The effect is strongest for output news, for which the calibrated information gain is 0.26. Figure 7 illustrates the mechanism. Following a positive Home endowment shock, the Home interest differential falls and the Home currency continues to depreciate for two quarters. Foreign deposits initially earn positive excess returns, but their return response becomes negative four quarters after the shock. Without average bias, the initial interest-rate and excess-return responses have the opposite signs, and the initial depreciation starts unwinding

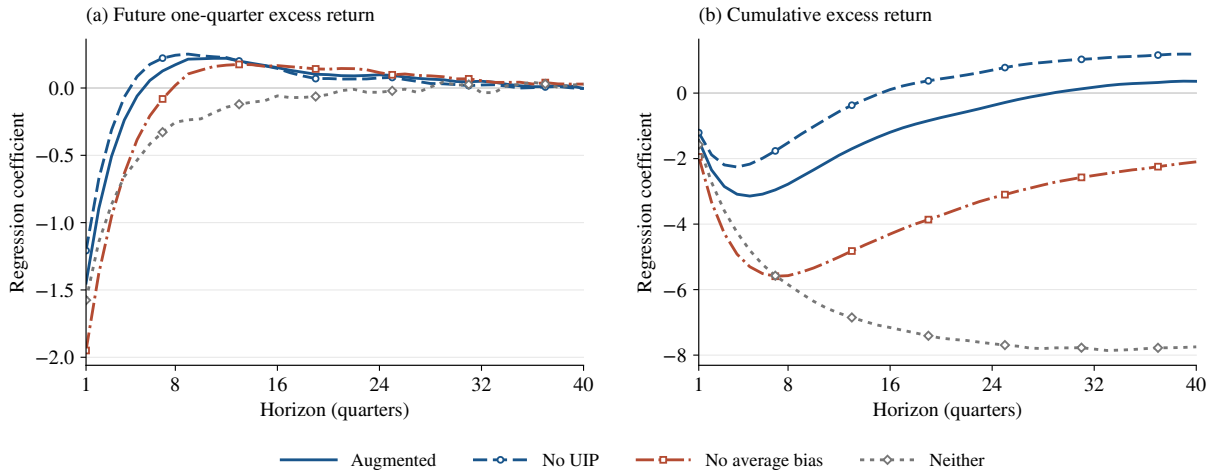


Figure 6: Average bias and the horizon profile of excess-return predictability

Notes: Lines are mean regression slopes across 1,000 simulations of 116 quarters after a 1,000-quarter burn-in. The left panel's b_h regresses the one-quarter Foreign excess return at horizon h on today's Home-minus-Foreign short-rate differential. The right panel's B_H uses cumulative Foreign excess returns over the next H quarters. A Foreign excess return is depreciation less the Home-minus-Foreign interest differential. All curves use the augmented calibration: No UIP removes exogenous UIP innovations, No average bias restores Bayesian consensus beliefs about fundamentals, and Neither imposes both restrictions, holding remaining parameters fixed. Neither is therefore distinct from the separately calibrated baseline. The augmented 40-quarter cumulative slope is 0.36, compared with 20.65 in the data.

immediately. Average bias has a smaller effect on the demand-shock responses.

Average bias also improves the cumulative-return result. The coefficient from regressing foreign excess returns accumulated over the next 40 quarters on today's Home-minus-Foreign interest differential is 0.36 in the augmented model, compared with -2.10 when average bias is removed at the same calibration and -6.52 in the separately calibrated baseline model. Removing only the UIP shock leaves a positive coefficient of 1.19. This improvement is qualitative: the augmented-model coefficient remains much smaller than the data estimate of 20.65, and its simulated 5th–95th percentile range, $[-3.08, 3.99]$, includes zero. Average bias helps reproduce the cumulative sign and the timing of the reversal, while the wealth–disagreement channel accounts for the benchmark's core FX-disconnect results.

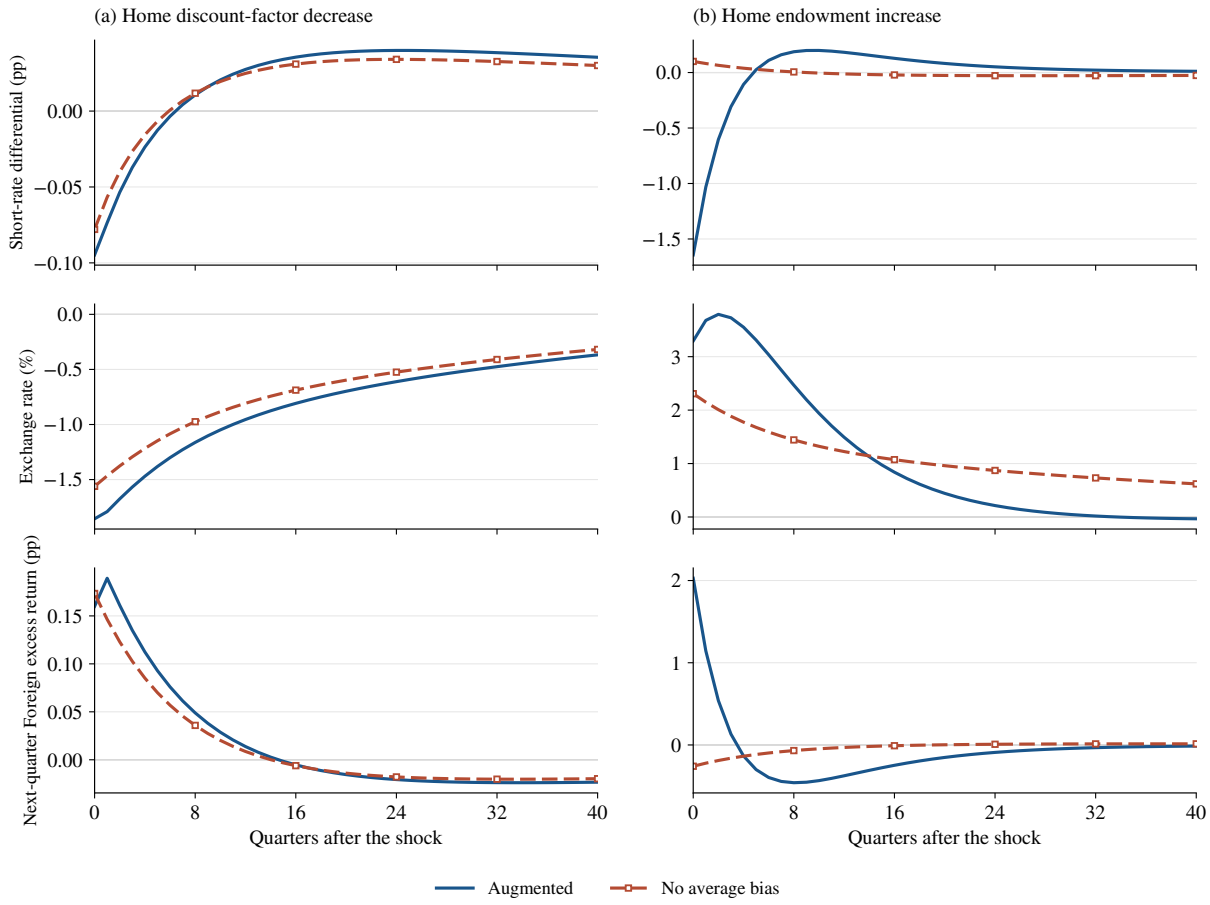


Figure 7: Average bias and the response to fundamental news

Notes: Responses to a one-standard-deviation decrease in the Home discount factor (left) and increase in the Home endowment (right), with the Foreign innovation set to zero. The paths start at the deterministic steady state and set subsequent innovations to zero. The comparison removes average bias at the augmented-model calibration. Rates and returns are in percentage points per quarter; an increase in the exchange rate denotes Home depreciation. At horizon h, the bottom panels report the foreign excess return earned from $t+h$ to $t+h+1$, $\hat{e}_{t+h+1} - \hat{e}_{t+h} - (\hat{r}_{t+h} - \hat{r}^*_{t+h})$.

E Time-varying disagreement in the quantitative model

The quantitative specification retains two financial interactions through which the current belief gap changes exchange-rate transmission. This appendix states the beliefs, portfolio rule, and approximation that define this specification. The interactions use existing parameters and disagreement states.

E.1 Beliefs and portfolio demand

Assumption 1: persistent types and changing beliefs. Let M_t^i denote group i's expected carry return in the maintained return normalization. The two groups' conditional means satisfy

$$M_t^A = \bar{M}_t + \frac{D_t}{2}, \qquad M_t^B = \bar{M}_t - \frac{D_t}{2}, \qquad D_t = d_{ss} + d_t. \tag{50}$$

Here $\bar{M}_t$ is the arithmetic mean of subjective forecasts, including any average bias. The permanent difference d_{ss} reflects fixed types, while d_t is generated by the existing signal-interpretation process, with its first-order forecast loadings held fixed. The groups retain their interpretation rules as new information arrives. Thus time variation in the gap introduces no new primitive shock. Group labels remain fixed, so D_t is signed; with equal population weights, forecast dispersion is $|D_t|/2$. The interpretation that a larger gap means greater dispersion applies on the region $D_t > 0$.

Assumption 2: common, fixed perceived risk. The mean–variance rule in the maintained return normalization is

$$\frac{N_t^i}{W_t^i} = \frac{M_t^i}{\Gamma}, \qquad \frac{N_t^A}{W_t^A} - \frac{N_t^B}{W_t^B} = \frac{d_{ss} + d_t}{\Gamma}, \qquad \Gamma > 0. \tag{51}$$

Myopic demand and the common effective risk cost Γ are maintained primitives of this specification. Perceived risk remains fixed when the financial interactions are introduced. Realized conditional exchange-rate variance can vary because the transmission of primitive shocks changes with beliefs. The distinction between fixed perceived risk and changing realized volatility is part of the specified belief system.

E.2 The two financial interactions

Assumption 3: retain the disagreement interactions. The quantitative model retains the products of disagreement with the wealth share in bond pricing and with the

subsequent carry return in wealth accumulation. Other equilibrium coefficients and wealth-share normalizations retain their first-order forms.

First, aggregate portfolio demand weights each group's forecast by its current wealth. With $\widetilde{\omega}_t^A = \omega_t^A - \omega^A$ and $\delta_w = \omega^A - 1/2$, the disagreement component, net of its steady-state value, is

$$\begin{aligned}\Psi_t^d &= \left(\delta_w + \widetilde{\omega}_t^A\right)(d_{ss} + d_t) - \delta_w d_{ss} \\ &= d_{ss}\widetilde{\omega}_t^A + \delta_w d_t + d_t\widetilde{\omega}_t^A.\end{aligned} \tag{52}$$

This identity follows by substituting (50) into wealth-weighted forecasts. In the modified UIP condition, the disagreement contribution is $-\Psi_t^d$; the capacity terms retain their existing first-order form. The additional product therefore makes the pricing effect of a wealth transfer equal to the current gap:

$$\frac{\partial \Psi_t^d}{\partial \widetilde{\omega}_t^A} = d_{ss} + d_t. \tag{53}$$

Second, wealth accumulation multiplies the position chosen at t by the return realized at $t+1$. In the quantitative model's coordinates, write the first-order simple-carry deviation as

$$\widetilde{R}_{t+1}^{\text{ex}} = \frac{\widehat{e}_{t+1} - \widehat{e}_t + \widehat{r}_t^*}{\beta^*} - \frac{\widehat{r}_t}{\beta}. \tag{54}$$

The changing relative exposure in (51) gives the following additions to the existing group log-wealth equations:

$$q_{t+1}^A = \frac{1-\xi}{2\Gamma} d_t \widetilde{R}_{t+1}^{\text{ex}}, \qquad q_{t+1}^B = -\frac{1-\xi}{2\Gamma} d_t \widetilde{R}_{t+1}^{\text{ex}}. \tag{55}$$

Since the retained share coordinate is $\widetilde{\omega}_t^A = \omega^A\omega^B(\widehat{W}_t^A - \widehat{W}_t^B)$, these terms add $(1-\xi)\omega^A\omega^B d_t \widetilde{R}_{t+1}^{\text{ex}}/\Gamma$ to wealth-share dynamics. Combined with the existing term proportional to d_{ss}, the return loading is proportional to $d_{ss} + d_t$. The coefficient is fixed by portfolio demand and survival. At date t, pricing uses current disagreement d_t, while realized wealth gains use the disagreement d_{t-1} prevailing when the position was opened.

E.3 Forecasting and solution

Assumption 4: affine forecasts and current financial feedback. Agents forecast using the first-order affine perceived transition. Current realized prices and wealth enter the state carried into the next period. At the forecasting stage, the additional future interaction terms are evaluated at their reference value, zero. Thus today's wealth changes affect future forecasts through the first-order transition, while additional forecast corrections

for future bilinear interactions are excluded. This restriction specifies how expectations are approximated separately from current financial accounting.

At each date the two current interactions are solved jointly. Let $\mathbf{Y}_t^0$ be the first-order decision rule evaluated at the inherited state and current information, and let $\mathbf{Y}_t$ collect the corresponding quantitative-model variables. The implemented rule is

$$\mathbf{Y}_t = \mathbf{Y}_t^0 + \mathbf{H}\mathbf{Q}_t, \qquad \mathbf{Q}_t = \begin{pmatrix} d_t \widetilde{\omega}_t^A \\ (1-\xi) d_{t-1} \widetilde{R}_t^{\text{ex}} / (2\Gamma) \\ -(1-\xi) d_{t-1} \widetilde{R}_t^{\text{ex}} / (2\Gamma) \end{pmatrix}. \tag{56}$$

The fixed matrix $\mathbf{H}$ contains the first-order current responses to additional terms in the UIP and two wealth equations. The entries of $\mathbf{Q}_t$ are determined by the displayed products and introduce no independent innovations. The solved variables are then carried into the next period's state. In the implementation, these additional inputs do not change d_t directly. Conditional on d_t and d_{t-1}, the two distinct interaction terms therefore satisfy a linear two-by-two system.

Assumption 5: local regularity. The contemporaneous feedback system is nonsingular on the neighborhood under consideration, with its inverse bounded, and the underlying wealth and share interpretation is admissible there. At the reference state, all added products and their first derivatives are zero. Hence the Jacobian of $\mathbf{Y}_t - \mathbf{Y}_t^0 - \mathbf{H}\mathbf{Q}_t$ with respect to $\mathbf{Y}_t$ is the identity. The implicit function theorem gives a unique smooth local solution. Scaling state and innovation deviations by ϵ, with d_{ss} and the first-order coefficients fixed, yields

$$\mathbf{Y}_t = \mathbf{Y}_t^0 + O(\epsilon^2). \tag{57}$$

This comparison holds over a fixed finite horizon while the path remains in the stated neighborhood. The specification preserves the analytical first-order solution and adds the two selected financial interactions. Solving their current feedback jointly retains the higher powers generated by those interactions. The approximation is defined by affine forecasts and bilinear financial feedback; a complete second-order expansion of all equilibrium equations is not imposed.

These two margins make financial transmission depend on beliefs: greater disagreement raises relative exposure and the pricing effect of wealth redistribution. The local response to news can therefore vary even with constant primitive shock variances. The quantitative exercise measures this state dependence and its contribution to unconditional moments.

F Additional quantitative results

F.1 Full comparison of exchange-rate moments across model variants

Table F1 reports the full moment comparison. Tables F2–F4 report the same moments with one innovation block active at a time.

Table F1: Full comparison of exchange-rate moments

Moment	Data	Baseline	No disag. Fixed	No disag. Recal.	Augmented	No UIP	No avg. bias	Neither
A. Exchange rates and macroeconomic variables								
$\rho_1(\Delta\hat{e}_t)$	0.013	-0.050	-0.009	-0.003	0.064	0.084	-0.043	-0.051
$\rho_1(\hat{e}_t)$	0.950	0.908	0.930	0.926	0.917	0.922	0.900	0.905
$\sigma(\Delta\hat{e}_t)$	0.033	0.034	0.013	0.020	0.040	0.036	0.032	0.027
$\sigma(\Delta\hat{e}_t)/\sigma(\Delta\hat{c}_t)$	6.43	4.11	1.58	2.95	4.46	4.11	3.58	3.09
$\sigma(\Delta\hat{e}_t)/\sigma(\Delta\hat{y}_t)$	5.42	4.03	1.54	2.94	4.40	3.98	3.55	3.00
$\sigma(\Delta(\hat{r}_t^{(1)} - \hat{r}_t^{(1)*}))/\sigma(\Delta\hat{e}_t)$	0.032	0.117	0.041	0.053	0.426	0.407	0.278	0.093
$\sigma(\hat{r}_t^{(1)} - \hat{r}_t^{(1)*})/\sigma(\Delta\hat{e}_t)$	0.110	0.166	0.247	0.309	0.513	0.473	0.374	0.139
$\sigma(\Delta(\hat{r}_t^{(40)} - \hat{r}_t^{(40)*}))/\sigma(\Delta\hat{e}_t)$	0.021	0.012	0.055	0.068	0.016	0.013	0.016	0.010
$\sigma(\hat{r}_t^{(40)} - \hat{r}_t^{(40)*})/\sigma(\Delta\hat{e}_t)$	0.060	0.074	0.240	0.295	0.061	0.062	0.063	0.064
$\mathrm{Corr}(\Delta\hat{e}_t, \Delta(\hat{c}_t - \hat{c}_t^*))$	-0.169	-0.244	0.546	-0.131	-0.471	-0.320	-0.327	-0.071
B. Excess-return predictability								
β^{Fama}	-0.213	-0.480	1.023	0.906	-0.449	-0.208	-0.950	-0.576
R^2_{Fama}	0.0005	0.0135	0.0431	0.0528	0.0606	0.0184	0.1261	0.0120
$\beta^{\mathrm{Fama,return}}$	-1.213	-1.480	0.023	-0.094	-1.449	-1.208	-1.950	-1.576
$R^2_{\mathrm{Fama,return}}$	0.017	0.062	0.024	0.034	0.369	0.252	0.373	0.047
$\beta^{\mathrm{Engel,cumul}}_{40q}$	20.65	-6.52	-3.77	-5.72	0.36	1.19	-2.10	-7.75
90% interval	[6.24, 35.06]	[−18.54, 3.27]	[−36.67, 26.99]	[−35.43, 21.28]	[−3.08, 3.99]	[−2.43, 4.65]	[−6.51, 2.63]	[−19.65, 2.36]
C. Changes-on-changes regressions								
$\Delta\hat{e}_t$ on $4\Delta(\hat{r}_t^{(1)} - \hat{r}_t^{(1)*})$	-0.50	1.62	-4.10	-2.90	-0.13	-0.33	0.56	2.08
90% interval	[−2.03, 1.03]	[0.35, 3.38]	[−4.78, −3.44]	[−3.43, −2.36]	[−0.21, −0.04]	[−0.41, −0.26]	[0.43, 0.68]	[0.57, 3.61]
Adjusted R^2	-0.005	0.358	0.444	0.364	0.046	0.283	0.381	0.457
$\Delta\hat{e}_t$ on $4\Delta(\hat{r}_t^{(40)} - \hat{r}_t^{(40)*})$	-1.42	-15.03	-4.16	-3.57	1.24	-6.82	2.57	-16.78
90% interval	[−3.37, 0.52]	[−18.02, −11.13]	[−4.45, −3.85]	[−3.72, −3.41]	[−1.75, 4.08]	[−10.70, −2.84]	[−0.28, 5.28]	[−20.04, −12.75]
Adjusted R^2	0.005	0.487	0.825	0.928	0.010	0.124	0.030	0.436
D. Level regressions								
$\hat{e}_t$ on $4(\hat{r}_t^{(1)} - \hat{r}_t^{(1)*})$	-3.47	0.55	-1.79	-1.23	-0.22	-0.45	0.34	0.90
90% interval	[−6.81, −0.12]	[−1.48, 2.42]	[−4.06, 0.24]	[−2.60, 0.19]	[−0.55, 0.07]	[−0.80, −0.13]	[−0.22, 0.83]	[−1.39, 2.94]
Adjusted R^2	0.222	0.084	0.344	0.276	0.038	0.104	0.084	0.106
$\hat{e}_t$ on $4(\hat{r}_t^{(40)} - \hat{r}_t^{(40)*})$	-7.79	-6.65	-2.23	-1.70	1.98	0.42	-2.75	-7.62
90% interval	[−12.81, −2.77]	[−10.66, −2.80]	[−4.27, −0.24]	[−3.09, −0.21]	[−4.70, 8.96]	[−6.72, 7.90]	[−7.94, 3.44]	[−11.62, −3.77]
Adjusted R^2	0.336	0.521	0.471	0.448	0.126	0.111	0.169	0.527

Notes: The baseline and its two no-disagreement counterfactuals appear first. Fixed keeps other baseline parameters; Recal. re-estimates the fundamental shock block. The last four columns share the augmented calibration: Augmented retains average belief distortions and UIP shocks; No UIP, No avg. bias, and Neither remove the indicated wedges. All disagreement models retain both financial interactions. Entries average 1,000 simulations of 116 quarters after a 1,000-quarter burn-in. ρ_1 denotes first-order autocorrelation and $\hat{e}_t$ rises with Home depreciation. $\Delta\hat{c}_t$ and $\Delta\hat{y}_t$ are Home consumption and output growth. Yields $\hat{r}_t^{(1)}$ and $\hat{r}_t^{(40)}$ have one- and 40-quarter maturities and are in quarterly units. Rate ratios divide yield-level or yield-change SDs, as indicated, by depreciation SD. Backus–Smith correlates depreciation with relative consumption growth. Fama regresses next-quarter depreciation on today's Home-minus-Foreign short-rate differential; its return version subtracts that differential from depreciation. Engel uses the next 40 quarters' cumulative Foreign excess return. Panels C and D report slopes and adjusted R^2 for the displayed regressions; factors of four annualize yield regressors. Data cover 1991Q2–2020Q1. Data intervals are 90% confidence intervals; model intervals are simulation 5th–95th percentiles. Calibration fit is not an independent validation test.

Table F2: Exchange-rate moments conditional on demand innovations

Moment	Data	Baseline	No disag. Fixed	No disag. Recal.	Augmented	No UIP	No avg. bias
A. Exchange rates and macroeconomic variables							
$\rho_1(\Delta\hat{e}_t)$	0.013	-0.048	0.005	0.005	-0.023	-0.023	-0.048
$\rho_1(\hat{e}_t)$	0.950	0.911	0.931	0.931	0.911	0.911	0.909
$\sigma(\Delta\hat{e}_t)$	0.033	0.028	0.009	0.017	0.024	0.024	0.020
$\sigma(\Delta\hat{e}_t)/\sigma(\Delta\hat{c}_t)$	6.43	9.51	9.48	9.48	9.51	9.51	9.51
$\sigma(\Delta\hat{e}_t)/\sigma(\Delta\hat{y}_t)$	5.42	–	–	–	–	–	–
$\sigma(\Delta(\hat{r}_t^{(1)} - \hat{r}_t^{(1)*}))/\sigma(\Delta\hat{e}_t)$	0.032	0.132	0.051	0.051	0.104	0.104	0.106
$\sigma(\hat{r}_t^{(1)} - \hat{r}_t^{(1)*})/\sigma(\Delta\hat{e}_t)$	0.110	0.187	0.335	0.335	0.158	0.158	0.157
$\sigma(\Delta(\hat{r}_t^{(40)} - \hat{r}_t^{(40)*}))/\sigma(\Delta\hat{e}_t)$	0.021	0.014	0.074	0.074	0.012	0.012	0.012
$\sigma(\hat{r}_t^{(40)} - \hat{r}_t^{(40)*})/\sigma(\Delta\hat{e}_t)$	0.060	0.086	0.330	0.330	0.077	0.077	0.077
$\mathrm{Corr}(\Delta\hat{e}_t, \Delta(\hat{c}_t - \hat{c}_t^*))$	-0.169	-1.000	-1.000	-1.000	-1.000	-1.000	-1.000
B. Excess-return predictability							
β^{Fama}	-0.213	-0.366	0.840	0.840	-0.388	-0.388	-0.385
R^2_{Fama}	0.0005	0.0139	0.0531	0.0531	0.0106	0.0106	0.0108
$\beta^{\mathrm{Fama,return}}$	-1.213	-1.366	-0.160	-0.160	-1.388	-1.388	-1.385
$R^2_{\mathrm{Fama,return}}$	0.017	0.071	0.042	0.042	0.049	0.049	0.050
$\beta^{\mathrm{Engel,cumul}}_{40q}$	20.65	-5.87	-7.26	-7.26	-7.29	-7.29	-6.53
90% interval	$[6.24, 35.06]$	$[-17.35, 3.20]$	$[-34.50, 20.04]$	$[-34.50, 20.04]$	$[-17.39, 1.92]$	$[-17.39, 1.92]$	$[-17.13, 3.40]$
C. Changes-on-changes regressions							
$\Delta\hat{e}_t$ on $4\Delta(\hat{r}_t^{(1)} - \hat{r}_t^{(1)*})$	-0.50	1.44	-2.31	-2.31	1.87	1.87	1.82
90% interval	$[-2.03, 1.03]$	$[0.28, 3.15]$	$[-2.89, -1.73]$	$[-2.89, -1.73]$	$[0.48, 3.39]$	$[0.48, 3.39]$	$[0.45, 3.33]$
Adjusted R^2	-0.005	0.335	0.218	0.218	0.445	0.445	0.426
$\Delta\hat{e}_t$ on $4\Delta(\hat{r}_t^{(40)} - \hat{r}_t^{(40)*})$	-1.42	-14.66	-3.32	-3.32	-16.95	-16.95	-16.46
90% interval	$[-3.37, 0.52]$	$[-17.24, -11.19]$	$[-3.40, -3.23]$	$[-3.40, -3.23]$	$[-19.74, -13.88]$	$[-19.74, -13.88]$	$[-18.98, -13.65]$
Adjusted R^2	0.005	0.651	0.969	0.969	0.634	0.634	0.666
D. Level regressions							
$\hat{e}_t$ on $4(\hat{r}_t^{(1)} - \hat{r}_t^{(1)*})$	-3.47	0.32	-1.02	-1.02	0.76	0.76	0.54
90% interval	$[-6.81, -0.12]$	$[-1.73, 2.17]$	$[-2.32, 0.33]$	$[-2.32, 0.33]$	$[-1.37, 2.59]$	$[-1.37, 2.59]$	$[-1.64, 2.41]$
Adjusted R^2	0.222	0.072	0.228	0.228	0.097	0.097	0.082
$\hat{e}_t$ on $4(\hat{r}_t^{(40)} - \hat{r}_t^{(40)*})$	-7.79	-6.04	-1.55	-1.55	-6.70	-6.70	-6.82
90% interval	$[-12.81, -2.77]$	$[-9.06, -2.78]$	$[-2.78, -0.20]$	$[-2.78, -0.20]$	$[-9.78, -3.49]$	$[-9.78, -3.49]$	$[-9.83, -3.67]$
Adjusted R^2	0.336	0.558	0.439	0.439	0.560	0.560	0.585

Notes: Only the demand block—discount-factor fundamentals and their associated signal-noise and interpretation innovations—is active; other innovations are zero. The data column reports unconditional U.S.–advanced-foreign-economy moments over 1991Q2–2020Q1. Baseline and the two no-disagreement columns use their own calibrations; the final three columns use the augmented calibration. Removing UIP risk has no effect when its innovation is inactive. Model entries average 1,000 simulations of 116 quarters after a 1,000-quarter burn-in. ρ_1 is first-order autocorrelation; $\hat{e}_t$ rises with Home depreciation. Consumption and output growth are Home quantities; stars denote Foreign. Yields have one- or 40-quarter maturity and are in quarterly units; factors of four annualize regression covariates. SD ratios divide the indicated quantity by depreciation SD. Backus–Smith correlates depreciation with relative consumption growth. Fama regresses next-quarter depreciation on today's Home-minus-Foreign short rate; the return version subtracts the differential from depreciation. Engel uses cumulative Foreign excess returns over 40 quarters. Other regression rows state the dependent variable and covariate. Data intervals are 90% confidence intervals and model intervals are simulation 5th–95th percentiles. A dash denotes an undefined ratio when the denominator has zero conditional variance.

Table F3: Exchange-rate moments conditional on supply innovations

Moment	Data	Baseline	No disag. Fixed	No disag. Recal.	Augmented	No UIP	No avg. bias
A. Exchange rates and macroeconomic variables							
$\rho_1(\Delta\hat{e}_t)$	0.013	-0.045	-0.018	-0.026	0.165	0.165	-0.046
$\rho_1(\hat{e}_t)$	0.950	0.907	0.930	0.909	0.932	0.932	0.906
$\sigma(\Delta\hat{e}_t)$	0.033	0.020	0.010	0.010	0.027	0.027	0.018
$\sigma(\Delta\hat{e}_t)/\sigma(\Delta\hat{c}_t)$	6.43	2.55	1.18	1.50	3.16	3.16	2.11
$\sigma(\Delta\hat{e}_t)/\sigma(\Delta\hat{y}_t)$	5.42	2.43	1.18	1.56	3.71	3.71	2.53
$\sigma(\Delta(\hat{r}_t^{(1)} - \hat{r}_t^{(1)*}))/\sigma(\Delta\hat{e}_t)$	0.032	0.044	0.030	0.057	0.544	0.544	0.046
$\sigma(\hat{r}_t^{(1)} - \hat{r}_t^{(1)*})/\sigma(\Delta\hat{e}_t)$	0.110	0.089	0.136	0.197	0.628	0.628	0.091
$\sigma(\Delta(\hat{r}_t^{(40)} - \hat{r}_t^{(40)*}))/\sigma(\Delta\hat{e}_t)$	0.021	0.005	0.031	0.040	0.013	0.013	0.004
$\sigma(\hat{r}_t^{(40)} - \hat{r}_t^{(40)*})/\sigma(\Delta\hat{e}_t)$	0.060	0.045	0.123	0.133	0.046	0.046	0.043
$\mathrm{Corr}(\Delta\hat{e}_t, \Delta(\hat{c}_t - \hat{c}_t^*))$	-0.169	0.982	0.998	0.997	0.739	0.739	0.980
B. Excess-return predictability							
β^{Fama}	-0.213	-1.000	1.792	1.374	-0.200	-0.200	-1.078
R^2_{Fama}	0.0005	0.0119	0.0407	0.0555	0.0259	0.0259	0.0132
$\beta^{\mathrm{Fama,return}}$	-1.213	-2.000	0.792	0.374	-1.200	-1.200	-2.078
$R^2_{\mathrm{Fama,return}}$	0.017	0.029	0.015	0.014	0.369	0.369	0.032
$\beta^{\mathrm{Engel,cumul}}_{40q}$	20.65	-15.11	13.14	6.33	1.54	1.54	-15.14
90% interval	$[6.24, 35.06]$	$[-28.88, 1.07]$	$[-30.40, 50.33]$	$[-19.08, 25.94]$	$[-1.34, 4.62]$	$[-1.34, 4.62]$	$[-27.76, -0.17]$
C. Changes-on-changes regressions							
$\Delta\hat{e}_t$ on $4\Delta(\hat{r}_t^{(1)} - \hat{r}_t^{(1)*})$	-0.50	5.44	-8.21	-4.36	-0.38	-0.38	5.26
90% interval	$[-2.03, 1.03]$	$[4.76, 5.83]$	$[-8.32, -8.04]$	$[-4.40, -4.30]$	$[-0.39, -0.38]$	$[-0.39, -0.38]$	$[5.00, 5.44]$
Adjusted R^2	-0.005	0.906	0.988	0.990	0.690	0.690	0.939
$\Delta\hat{e}_t$ on $4\Delta(\hat{r}_t^{(40)} - \hat{r}_t^{(40)*})$	-1.42	-41.36	-8.06	-6.17	-2.05	-2.05	-43.14
90% interval	$[-3.37, 0.52]$	$[-50.40, -32.52]$	$[-8.12, -7.97]$	$[-6.21, -6.09]$	$[-5.90, 1.47]$	$[-5.90, 1.47]$	$[-54.34, -32.29]$
Adjusted R^2	0.005	0.620	0.995	0.992	0.014	0.014	0.540
D. Level regressions							
$\hat{e}_t$ on $4(\hat{r}_t^{(1)} - \hat{r}_t^{(1)*})$	-3.47	2.34	-5.38	-3.01	-0.49	-0.49	2.52
90% interval	$[-6.81, -0.12]$	$[-1.85, 4.99]$	$[-6.95, -3.69]$	$[-3.88, -2.07]$	$[-0.62, -0.40]$	$[-0.62, -0.40]$	$[-1.47, 4.91]$
Adjusted R^2	0.222	0.197	0.829	0.802	0.187	0.187	0.218
$\hat{e}_t$ on $4(\hat{r}_t^{(40)} - \hat{r}_t^{(40)*})$	-7.79	-11.43	-6.08	-4.52	13.25	13.25	-11.58
90% interval	$[-12.81, -2.77]$	$[-15.67, -7.04]$	$[-7.35, -4.56]$	$[-5.64, -3.25]$	$[11.30, 14.91]$	$[11.30, 14.91]$	$[-16.01, -7.30]$
Adjusted R^2	0.336	0.565	0.897	0.838	0.719	0.719	0.545

Notes: Only the supply block—endowment fundamentals and their associated signal-noise and interpretation innovations—is active; other innovations are zero. The data column reports unconditional U.S.–advanced-foreign-economy moments over 1991Q2–2020Q1. Baseline and the two no-disagreement columns use their own calibrations; the final three columns use the augmented calibration. Removing UIP risk has no effect when its innovation is inactive. Model entries average 1,000 simulations of 116 quarters after a 1,000-quarter burn-in. ρ_1 is first-order autocorrelation; $\hat{e}_t$ rises with Home depreciation. Consumption and output growth are Home quantities; stars denote Foreign. Yields have one- or 40-quarter maturity and are in quarterly units; factors of four annualize regression covariates. SD ratios divide the indicated quantity by depreciation SD. Backus–Smith correlates depreciation with relative consumption growth. Fama regresses next-quarter depreciation on today's Home-minus-Foreign short rate; the return version subtracts the differential from depreciation. Engel uses cumulative Foreign excess returns over 40 quarters. Other regression rows state the dependent variable and covariate. Data intervals are 90% confidence intervals and model intervals are simulation 5th–95th percentiles. A dash denotes an undefined ratio when the denominator has zero conditional variance.

Table F4: Exchange-rate moments conditional on the UIP-risk innovation

Moment	Data	Augmented model	No avg. bias
A. Exchange rates and macroeconomic variables			
$\rho_1(\Delta\hat{e}_t)$	0.013	-0.021	-0.021
$\rho_1(\hat{e}_t)$	0.950	0.886	0.886
$\sigma(\Delta\hat{e}_t)$	0.033	0.017	0.017
$\sigma(\Delta\hat{e}_t)/\sigma(\Delta\hat{c}_t)$	6.43	9.51	9.51
$\sigma(\Delta\hat{e}_t)/\sigma(\Delta\hat{y}_t)$	5.42	–	–
$\sigma(\Delta(\hat{r}_t^{(1)} - \hat{r}_t^{(1)*}))/\sigma(\Delta\hat{e}_t)$	0.032	0.502	0.502
$\sigma(\hat{r}_t^{(1)} - \hat{r}_t^{(1)*})/\sigma(\Delta\hat{e}_t)$	0.110	0.681	0.681
$\sigma(\Delta(\hat{r}_t^{(40)} - \hat{r}_t^{(40)*}))/\sigma(\Delta\hat{e}_t)$	0.021	0.026	0.026
$\sigma(\hat{r}_t^{(40)} - \hat{r}_t^{(40)*})/\sigma(\Delta\hat{e}_t)$	0.060	0.061	0.061
$\mathrm{Corr}(\Delta\hat{e}_t, \Delta(\hat{c}_t - \hat{c}_t^*))$	-0.169	-1.000	-1.000
B. Excess-return predictability			
β^{Fama}	-0.213	-0.983	-0.983
R^2_{Fama}	0.0005	0.4385	0.4385
$\beta^{\mathrm{Fama,return}}$	-1.213	-1.983	-1.983
$R^2_{\mathrm{Fama,return}}$	0.017	0.762	0.762
$\beta^{\mathrm{Engel,cumul}}_{40q}$	20.65	-1.66	-1.66
90% interval	$[6.24, 35.06]$	$[-5.29, 2.19]$	$[-5.29, 2.19]$
C. Changes-on-changes regressions			
$\Delta\hat{e}_t$ on $4\Delta(\hat{r}_t^{(1)} - \hat{r}_t^{(1)*})$	-0.50	0.47	0.47
90% interval	$[-2.03, 1.03]$	$[0.46, 0.47]$	$[0.46, 0.47]$
Adjusted R^2	-0.005	0.879	0.879
$\Delta\hat{e}_t$ on $4\Delta(\hat{r}_t^{(40)} - \hat{r}_t^{(40)*})$	-1.42	9.73	9.73
90% interval	$[-3.37, 0.52]$	$[9.68, 9.78]$	$[9.68, 9.78]$
Adjusted R^2	0.005	0.999	0.999
D. Level regressions			
$\hat{e}_t$ on $4(\hat{r}_t^{(1)} - \hat{r}_t^{(1)*})$	-3.47	0.31	0.31
90% interval	$[-6.81, -0.12]$	$[0.16, 0.42]$	$[0.16, 0.42]$
Adjusted R^2	0.222	0.174	0.174
$\hat{e}_t$ on $4(\hat{r}_t^{(40)} - \hat{r}_t^{(40)*})$	-7.79	9.00	9.00
90% interval	$[-12.81, -2.77]$	$[8.47, 9.44]$	$[8.47, 9.44]$
Adjusted R^2	0.336	0.988	0.988

Notes: Only the UIP-risk innovation block is active; other innovations are zero. The data column reports unconditional U.S.–advanced-foreign-economy moments over 1991Q2–2020Q1. Both columns use the augmented calibration; No avg. bias removes average belief distortions. The baseline has no exogenous UIP shock and is therefore omitted. Model entries average 1,000 simulations of 116 quarters after a 1,000-quarter burn-in. ρ_1 is first-order autocorrelation; $\hat{e}_t$ rises with Home depreciation. Consumption and output growth are Home quantities; stars denote Foreign. Yields have one- or 40-quarter maturity and are in quarterly units; factors of four annualize regression covariates. SD ratios divide the indicated quantity by depreciation SD. Backus–Smith correlates depreciation with relative consumption growth. Fama regresses next-quarter depreciation on today's Home-minus-Foreign short rate; the return version subtracts the differential from depreciation. Engel uses cumulative Foreign excess returns over 40 quarters. Other regression rows state the dependent variable and covariate. Data intervals are 90% confidence intervals and model intervals are simulation 5th–95th percentiles. A dash denotes an undefined ratio when the denominator has zero conditional variance.

F.2 Full impulse responses across model variants

Each panel compares all seven specifications in Table F1. Responses start at the deterministic steady state and follow a one-standard-deviation Home discount-factor decrease or endowment increase under each calibration; the Foreign innovation and all subsequent innovations are zero. Augmented-model variants share a calibration, while the baseline and the no-disagreement variants follow the table. UIP-shutdown variants coincide with their counterparts for these shocks; separate markers identify them.

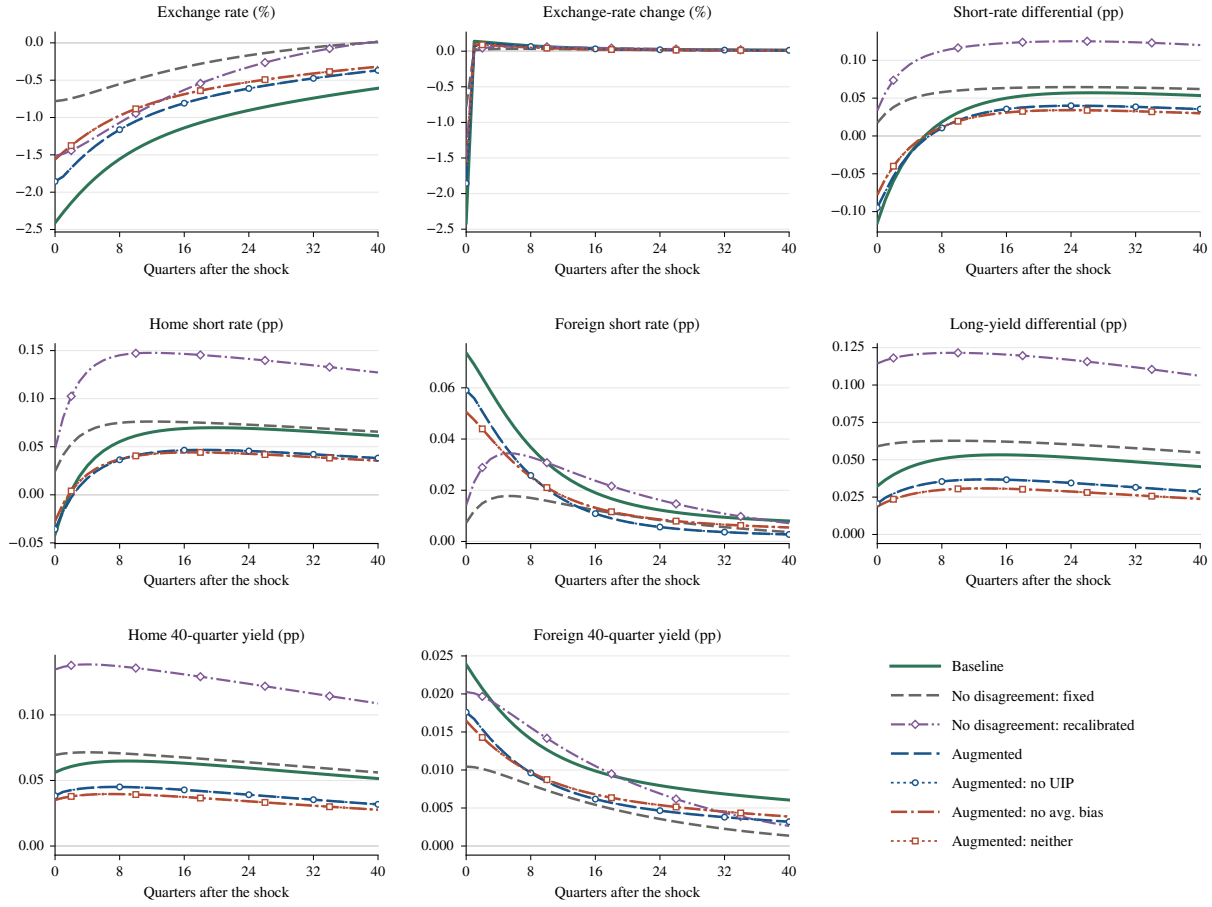


Figure F1: Home discount-factor decrease: exchange rates and yields

Notes: Responses to a one-standard-deviation Home discount-factor decrease at quarter zero, using each model's own innovation scale; Foreign and subsequent innovations are zero. Baseline and its no-disagreement counterfactuals use their respective calibrations. Augmented ablations hold its remaining parameters fixed. The augmented/no-UIP paths coincide, as do the no-average-bias/neither paths. An exchange-rate increase is Home depreciation. Rates and yields are in percentage points per quarter; long yields have a 40-quarter maturity.

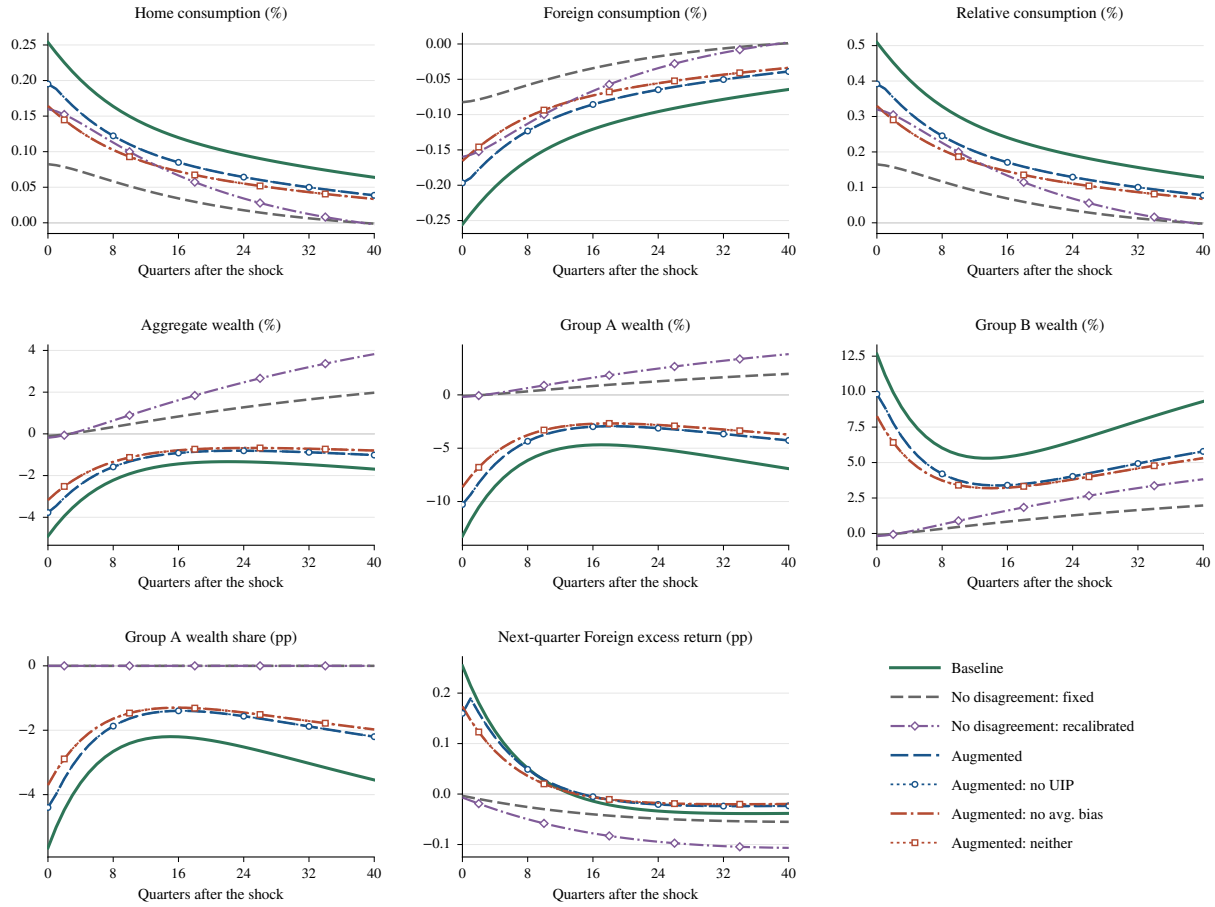


Figure F2: Home discount-factor decrease: consumption, wealth, and returns

Notes: Responses to a one-standard-deviation Home discount-factor decrease at quarter zero, using each model's own innovation scale; Foreign and subsequent innovations are zero. Baseline and its no-disagreement counterfactuals use their respective calibrations. Augmented ablations hold its remaining parameters fixed. The augmented/no-UIP paths coincide, as do the no-average-bias/neither paths. Consumption and wealth are percent deviations from steady state; the Group A wealth share and next-quarter Foreign excess return are in percentage points. The return excludes the impact valuation gain.

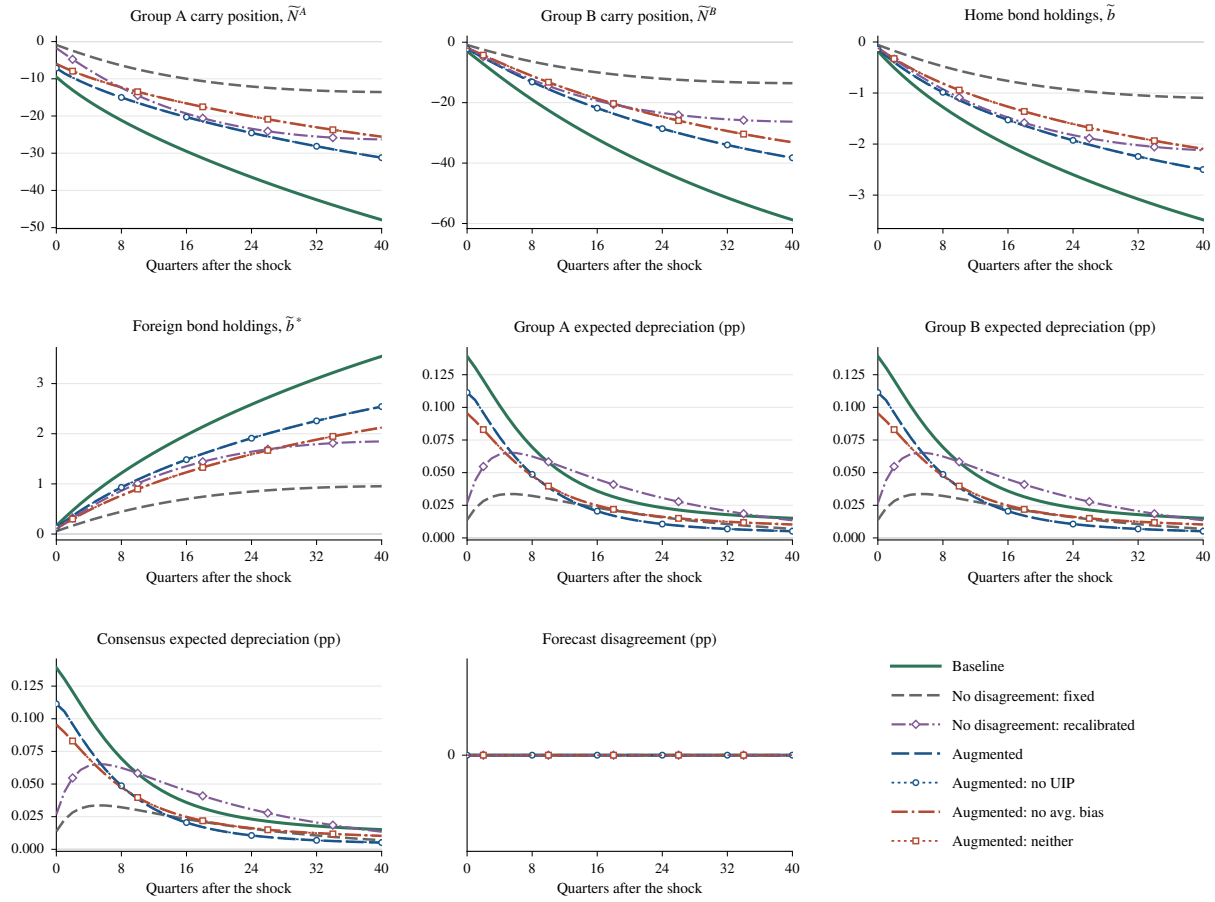


Figure F3: Home discount-factor decrease: positions and exchange-rate expectations

Notes: Responses to a one-standard-deviation Home discount-factor decrease at quarter zero, using each model's own innovation scale; Foreign and subsequent innovations are zero. Baseline and its no-disagreement counterfactuals use their respective calibrations. Augmented ablations hold its remaining parameters fixed. The augmented/no-UIP paths coincide, as do the no-average-bias/neither paths. Carry positions $\widetilde{N}^A, \widetilde{N}^B$ are normalized by group steady-state wealth; household bond deviations $\widetilde{b}, \widetilde{b}^*$ by own-country steady-state output. Forecasts are changes in expected next-quarter depreciation, in percentage points. Group A is optimistic about Foreign returns.

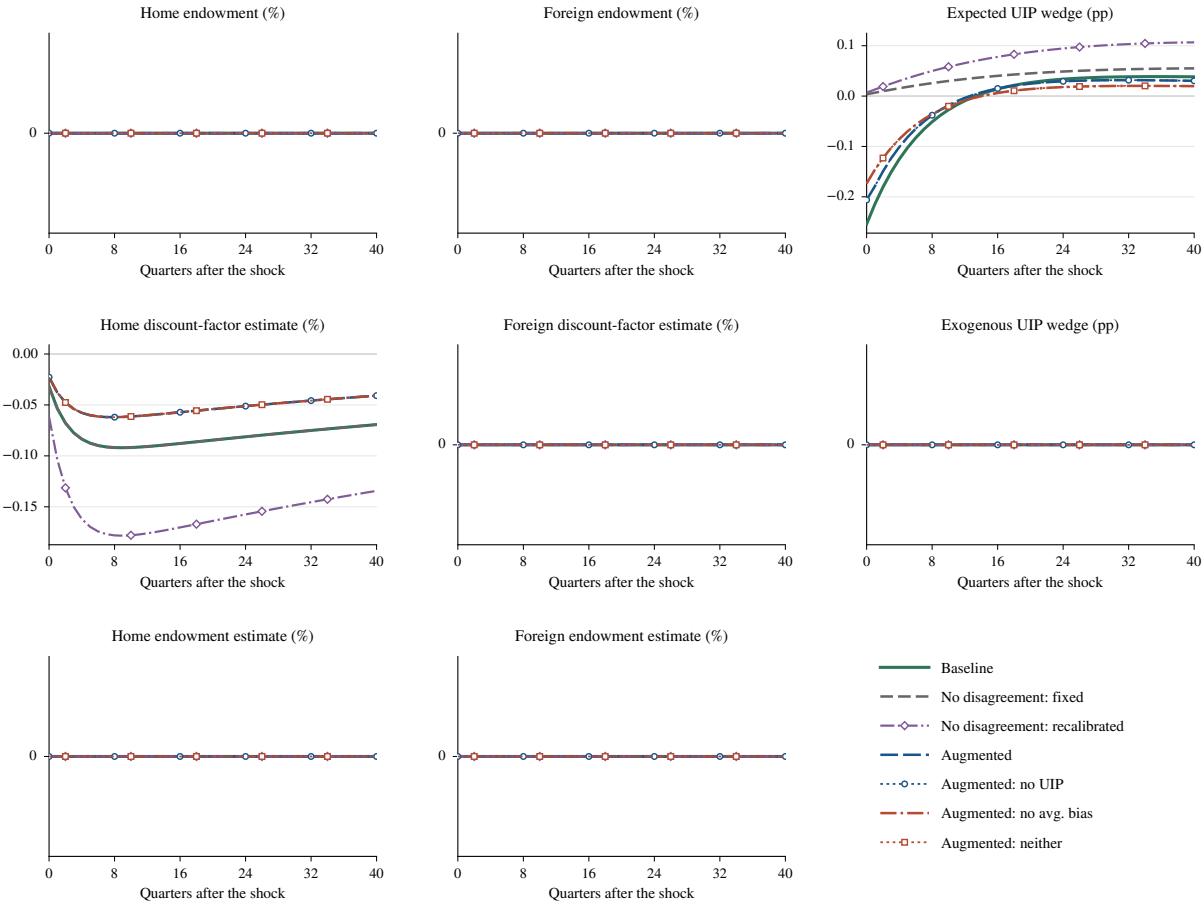


Figure F4: Home discount-factor decrease: endowments, fundamental estimates, and UIP wedges

Notes: Responses to a one-standard-deviation Home discount-factor decrease at quarter zero, using each model's own innovation scale; Foreign and subsequent innovations are zero. Baseline and its no-disagreement counterfactuals use their respective calibrations. Augmented ablations hold its remaining parameters fixed. The augmented/no-UIP paths coincide, as do the no-average-bias/neither paths. Fundamental estimates are posterior estimates of persistent states, in percent deviations. The UIP wedge is the Home-minus-Foreign short rate less consensus expected depreciation, in percentage points. The exogenous UIP innovation is zero.

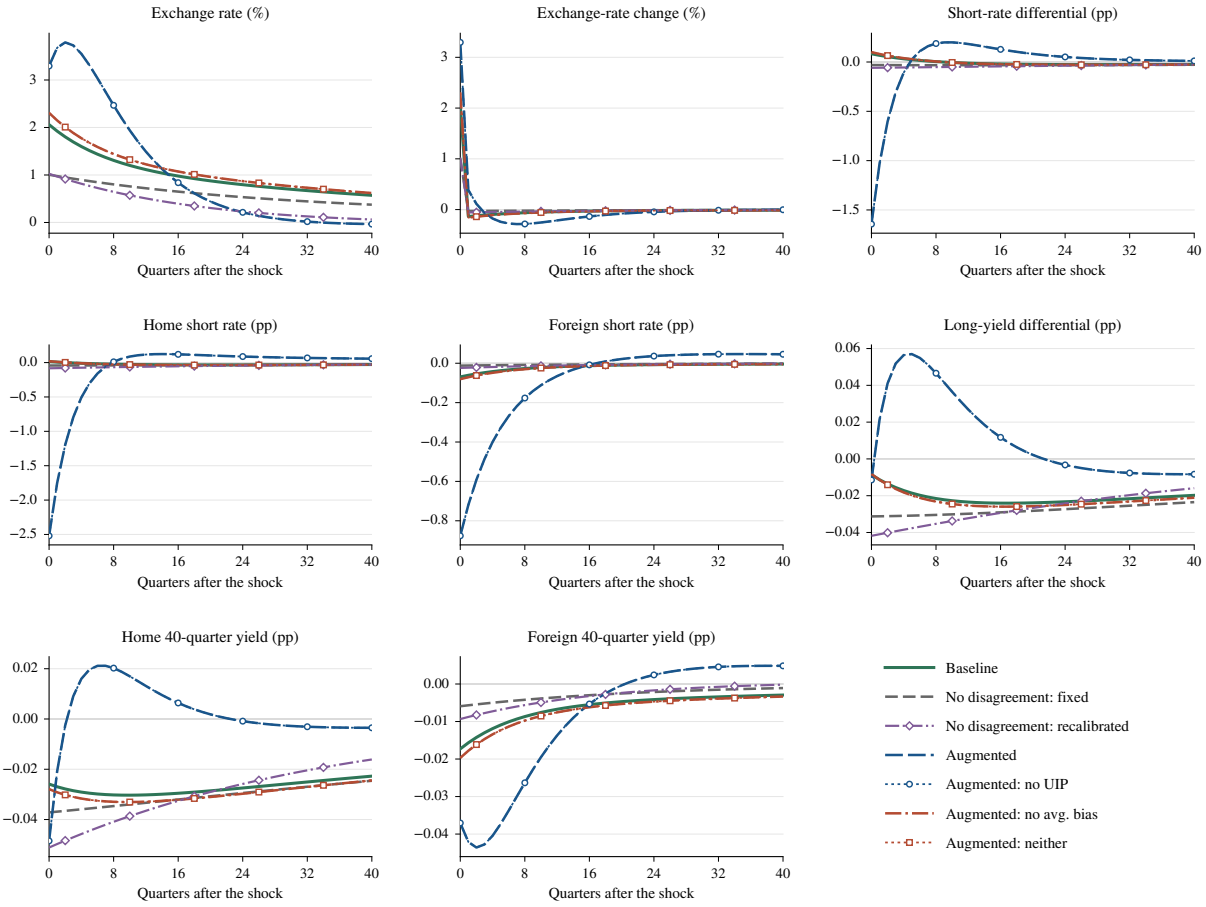


Figure F5: Home endowment increase: exchange rates and yields

Notes: Responses to a one-standard-deviation Home endowment increase at quarter zero, using each model's own innovation scale; Foreign and subsequent innovations are zero. Baseline and its no-disagreement counterfactuals use their respective calibrations. Augmented ablations hold its remaining parameters fixed. The augmented/no-UIP paths coincide, as do the no-average-bias/neither paths. An exchange-rate increase is Home depreciation. Rates and yields are in percentage points per quarter; long yields have a 40-quarter maturity.

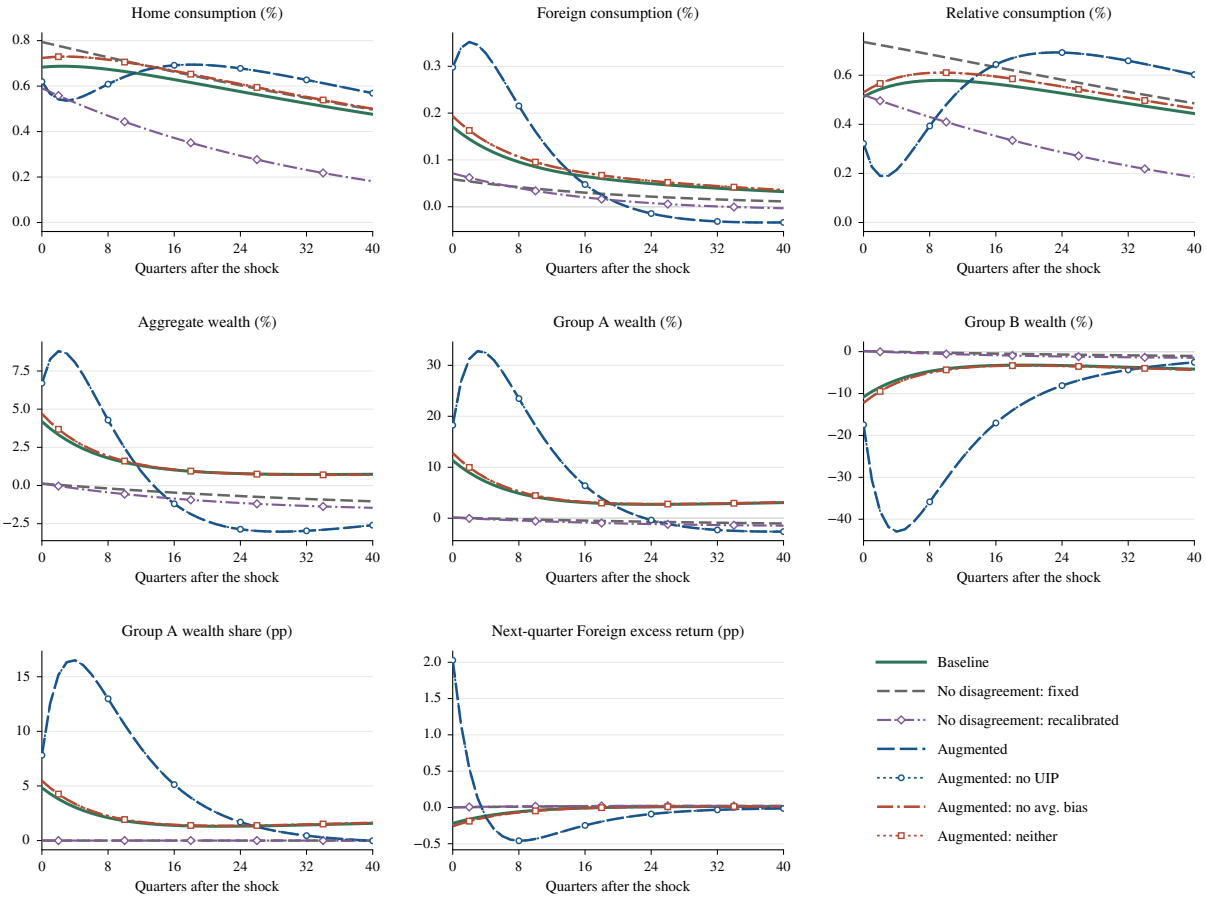


Figure F6: Home endowment increase: consumption, wealth, and returns

Notes: Responses to a one-standard-deviation Home endowment increase at quarter zero, using each model's own innovation scale; Foreign and subsequent innovations are zero. Baseline and its no-disagreement counterfactuals use their respective calibrations. Augmented ablations hold its remaining parameters fixed. The augmented/no-UIP paths coincide, as do the no-average-bias/neither paths. Consumption and wealth are percent deviations from steady state; the Group A wealth share and next-quarter Foreign excess return are in percentage points. The return excludes the impact valuation gain.

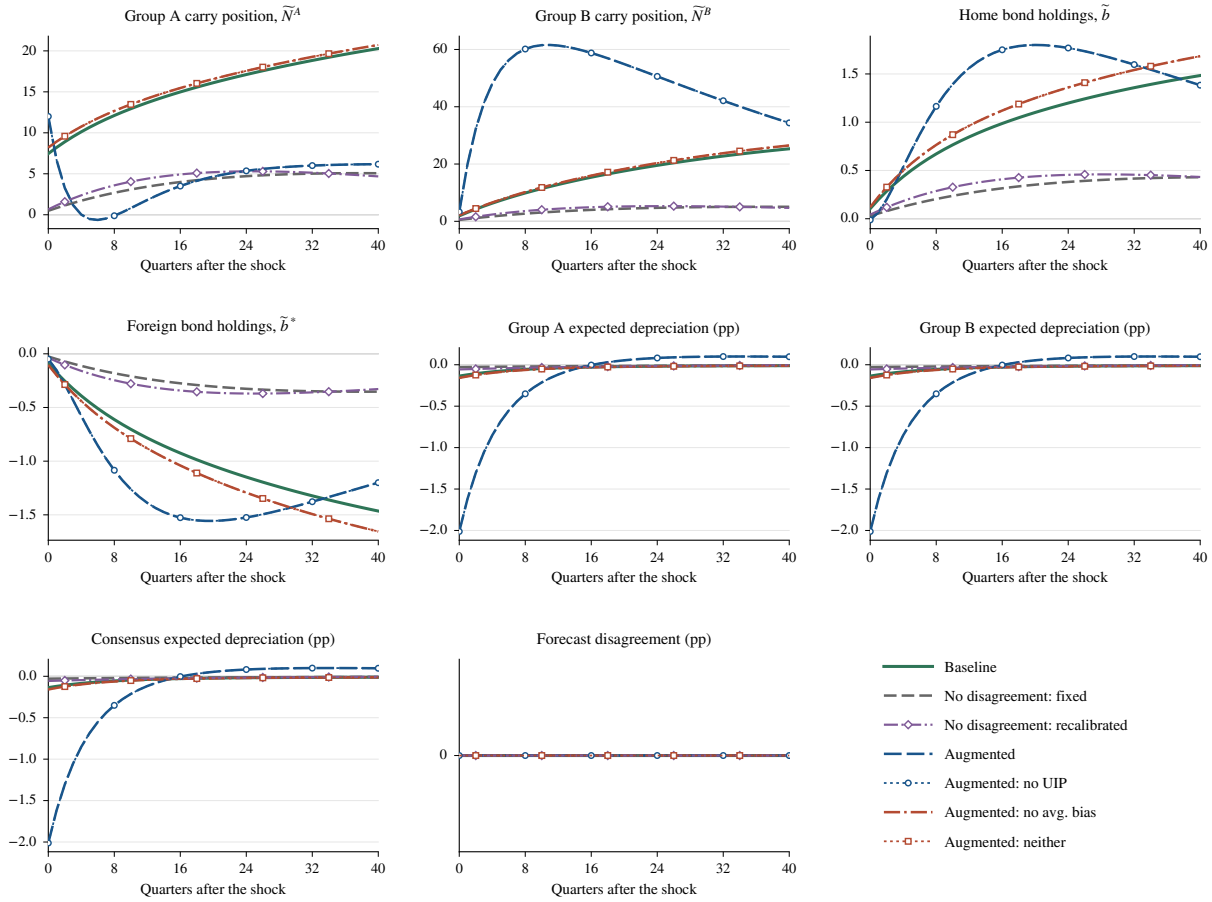


Figure F7: Home endowment increase: positions and exchange-rate expectations

Notes: Responses to a one-standard-deviation Home endowment increase at quarter zero, using each model's own innovation scale; Foreign and subsequent innovations are zero. Baseline and its no-disagreement counterfactuals use their respective calibrations. Augmented ablations hold its remaining parameters fixed. The augmented/no-UIP paths coincide, as do the no-average-bias/neither paths. Carry positions $\widetilde{N}^A, \widetilde{N}^B$ are normalized by group steady-state wealth; household bond deviations $\widetilde{b}, \widetilde{b}^*$ by own-country steady-state output. Forecasts are changes in expected next-quarter depreciation, in percentage points. Group A is optimistic about Foreign returns.

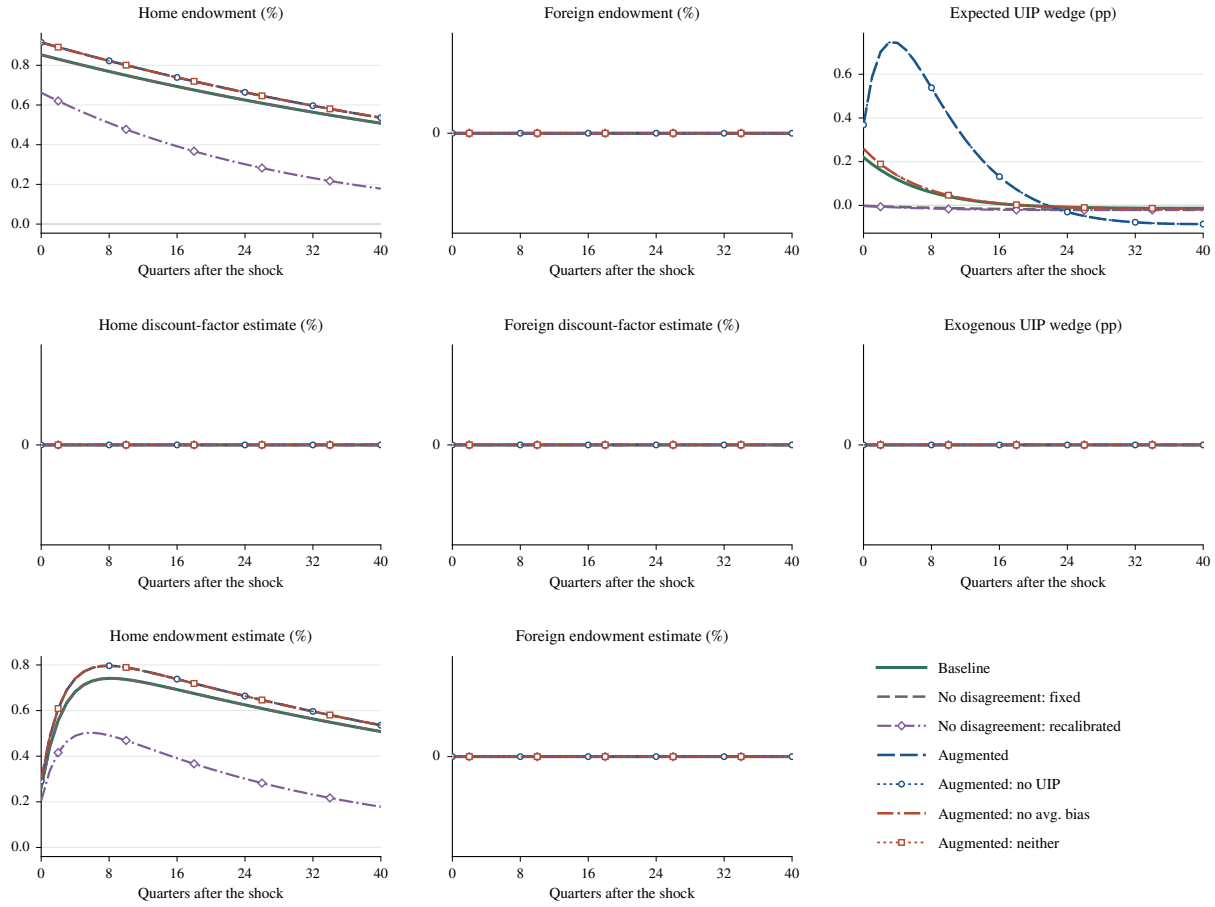


Figure F8: Home endowment increase: endowments, fundamental estimates, and UIP wedges
Notes: Responses to a one-standard-deviation Home endowment increase at quarter zero, using each model's own innovation scale; Foreign and subsequent innovations are zero. Baseline and its no-disagreement counterfactuals use their respective calibrations. Augmented ablations hold its remaining parameters fixed. The augmented/no-UIP paths coincide, as do the no-average-bias/neither paths. Fundamental estimates are posterior estimates of persistent states, in percent deviations. The UIP wedge is the Home-minus-Foreign short rate less consensus expected depreciation, in percentage points. The exogenous UIP innovation is zero.

F.3 Historical inversion and recursive accounting

F.3.1 Measurement and historical decomposition

The historical inversion uses the same U.S. and Fed-weighted advanced-foreign-economy data as the calibration. Both country output series are in logs per capita. I remove a common linear trend formed from the two fitted country trends, using the augmented economy's steady-state output shares, approximately 0.502 and 0.498. The trend is estimated over the original 1991Q2-2020Q1 calibration window and held fixed. Long yields are expressed in quarterly units. Macro observations and disagreement retain their original 1999Q1-2019Q2 centering.

The main inversion begins in 1999Q1 with inherited endogenous states at steady state. Initial shocks match the first observations, and the inferred states then evolve through 2019Q4. Blue Chip disagreement is the mean of the ten highest minus the mean of the ten lowest forecaster-quarter averages of the within-issue log ratio of next-quarter to current-quarter dollar forecasts. The two demand innovations, two output innovations, and one antisymmetric beta-belief innovation match the five observations in the disagreement model. The representative-agent economy uses only the four macro observations. Exchange-rate realizations and survey forecast errors are excluded from this step.

The figures plot 100 times log deviations under a dollar-appreciation convention. Let s_t^{app} denote the log nominal dollar exchange rate in this convention and q_t^{app} its real counterpart. I construct the nominal path as $s_t^{\text{app,model}} = q_t^{\text{app,model}} + (\log P_t^* - \log P_t)^{\text{data}}$, up to the fixed index normalization, where P_t and P_t^* are observed CPI levels. The observed relative CPI path is added identically to the real data and each model path. Consequently, nominal and real pointwise errors are identical. The nominal presentation is conditional on observed prices; the model does not separately predict the two nominal price levels. Every plotted path retains its 1999Q1-2019Q2 mean alignment, with no smoothing or volatility rescaling.

The historical shock decomposition replays the augmented model's inferred innovations one block at a time, holding the remaining innovations at zero. For an additive real exchange-rate component $q_t^{\text{app},j}$, I report its covariance allocation $\text{Cov}(q_t^{\text{app},j}, q_t^{\text{app}})/\text{Var}(q_t^{\text{app}})$. These allocations sum to one and allow the inferred shock blocks to be correlated. Demand and output account for 56.5 and 43.5 percent of model real level variance, respectively, while the disagreement contribution is approximately zero. This is a decomposition of the model's fitted history, separate from the simulated forecast-error variance decomposition. In quarterly changes, the corresponding demand and output allocations are 92.2 and 7.6 percent, so the large historical output contribution concerns levels. The observed CPI component is kept separate from the structural shocks when decomposing the conditional nominal path.

The EBP comparison imposes the same externally measured UIP path at the calibrated persistence and innovation scale, then re-infers the other shocks. In the main sample, adding EBP worsens the disagreement model's level fit, although its quarterly-change correlation rises from approximately -0.009 to 0.099. The conclusion is specific to this proxy and the level comparison.

F.3.2 Extended sample

Figure F9 extends the augmented model without UIP shocks through 2024Q4, the last quarter with the joint macro and survey inputs. Data begin in 1991Q1, using an initial index anchor recovered from the saved 1991Q2 level and quarterly growth. Model inference begins

in 1991Q2, the first quarter with the calibrated yield source, and continues through 2024Q4 without a reset in 1999. The main-sample paths retain their original 1999 initialization and are not extracted from this longer history.

The shaded period before 1999 uses reported legacy trade-weighted dollar forecasts. I apply the same within-issue log forecast ratio and top-ten-minus-bottom-ten construction. The usable cross-sections contain 38-47 respondents per quarter. Early analyst codes identify respondents within each quarter; ambiguous duplicate identifiers and invalid index entries are excluded according to the archive rules. The legacy dollar index uses a different currency basket, so its disagreement series is a historical proxy for the calibrated measure. No loading is fitted to realized exchange rates, and this period is not separately recentered.

From 2022Q4 onward, also shaded, some quarters have fewer than 20 forecasters. I replace the unavailable top-minus-bottom measure with cross-sectional standard deviation mapped into its original units using a regression estimated over 1999Q1-2019Q2:

$$D_t^{\text{fit}} = 0.0000249799 + 2.10477218 \ \text{SD}_i[\log(F_{i,t,t+1}/F_{i,t,t})] \, .$$

Here $F_{i,t,t+1}$ and $F_{i,t,t}$ are respondent i's next- and current-quarter forecasts, with log ratios averaged within respondent and quarter before measuring dispersion. The frozen mapping applies throughout 2022Q4-2024Q4, including quarters in which the original statistic is available. This is a measurement extension of the survey input.

I retain calibration observations through 2020Q1 and append the same Fed-weighted AFE macro series thereafter. Index constants are aligned over the overlap. The foreign-yield source switches from the saved v3 to v4 spread vintage after 2020Q1; their endpoint gap is approximately 0.0048 basis points. The output trend is extrapolated without re-estimation. Parameters, EBP scaling, CPI conversion, and centering remain fixed. The large Covid-era level miss limits the augmented model's performance over the extended sample, and the EBP comparison varies across windows and fit measures. These exercises are conditional historical accounts, not real-time forecasts.

F.3.3 Demand and beta-disagreement shocks

The heterogeneous-agent restriction retains three innovations: Home demand, Foreign demand, and beta-disagreement. All supply and signal-noise innovations are zero. I infer this history anew from the two long yields and survey disagreement; output is held out. The representative-agent benchmark retains only its two demand innovations and matches the same two yields. The right panel separately adds the EBP-implied UIP path.

Figure F10 reports the 1999Q1-2019Q4 comparison. Restricting the shock set lowers

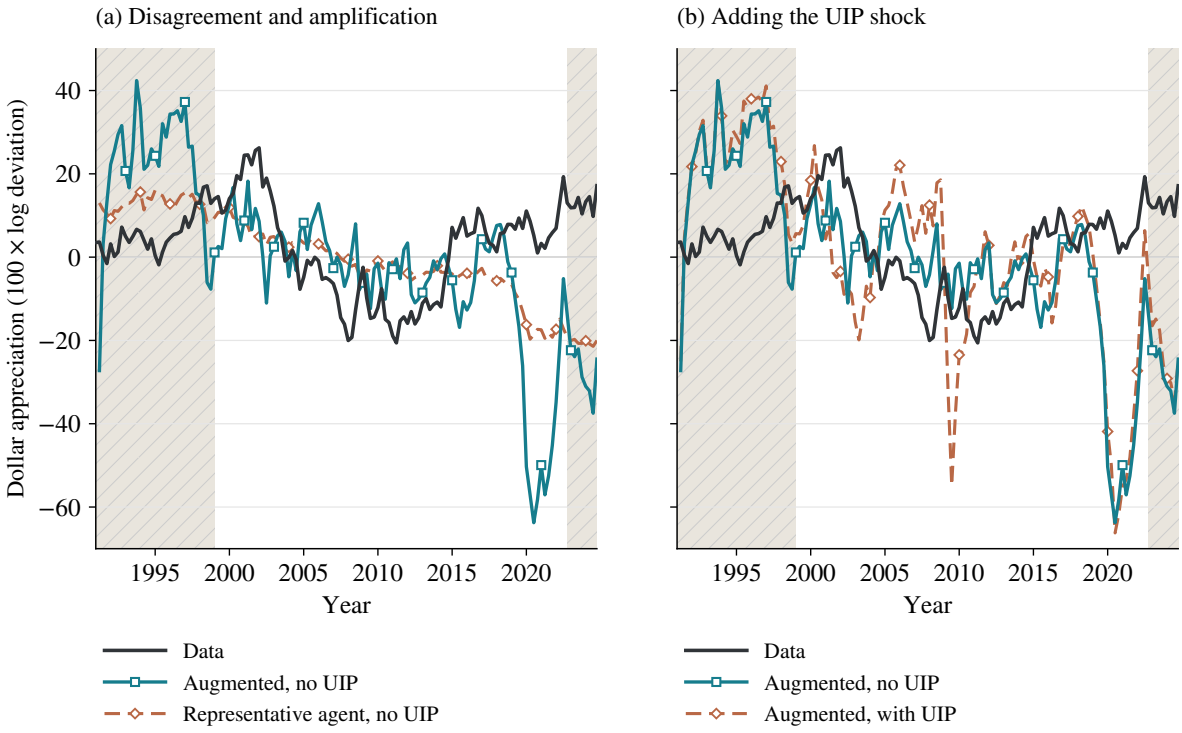


Figure F9: Historical inversion at the augmented calibration, 1991Q1–2024Q4

Notes: The historical exercise uses the first-order augmented calibration. Panel (a) compares the augmented model without UIP shocks with its representative-agent counterpart; panel (b) adds the observed excess-bond-premium UIP path and re-infers other shocks. The nominal dollar index rises with dollar appreciation. Paths are 100 times log deviations, mean-aligned over 1999Q1–2019Q2 without smoothing or volatility rescaling. All models condition on the same observed relative consumer-price path. Model inference begins in 1991Q2 and continues without a reset in 1999. Hatching marks the pre-1999 legacy survey and the dispersion-proxy continuation from 2022Q4.

the disagreement model's level correlation from 0.497 to 0.298 and its model-to-data level standard-deviation ratio from 1.00 to 0.61. The representative-agent model has a higher correlation of 0.592 but a standard-deviation ratio of only 0.28. Adding EBP lowers the disagreement model's level correlation to 0.103. The restriction therefore preserves the representative-agent shortfall in amplitude, but it does not strengthen the disagreement model's level comovement. Figure F11 gives the corresponding extended history, subject to the same measurement qualifications as the unrestricted historical exercise. Full and restricted figures share the same axes within each sample.

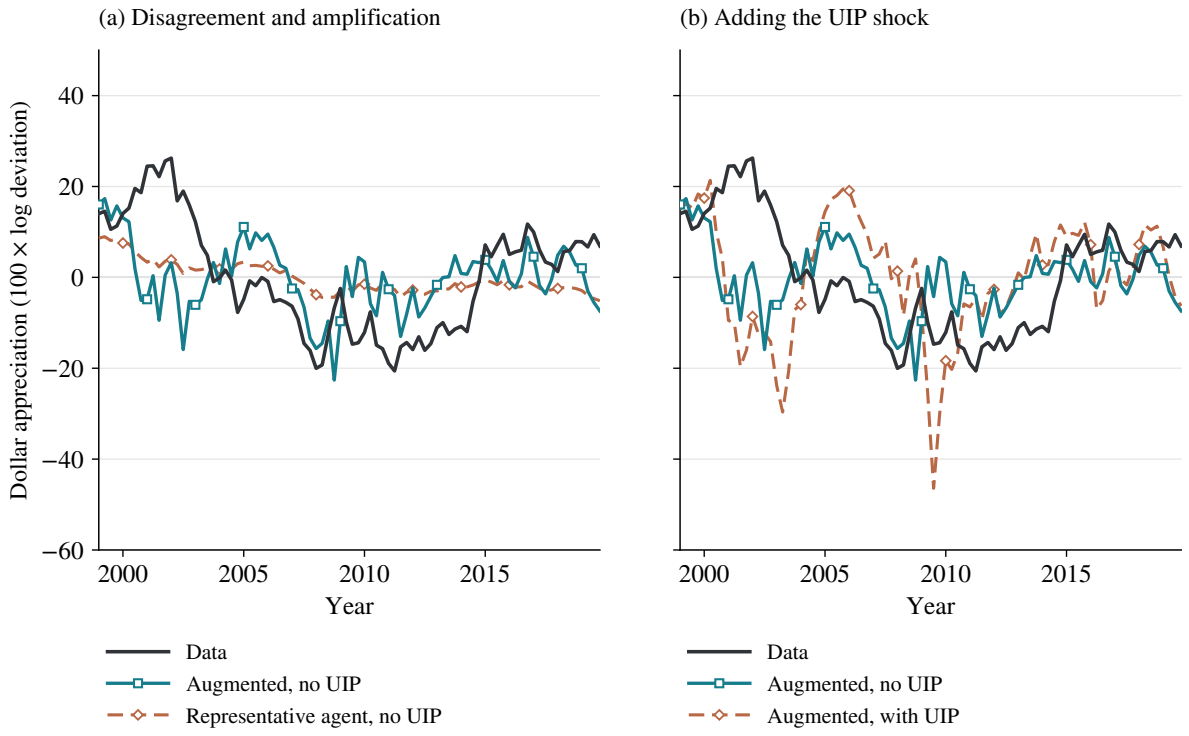


Figure F10: Demand and disagreement inversion at the augmented calibration, 1999Q1–2019Q4

Notes: The historical exercise uses the first-order augmented calibration. Panel (a) compares the augmented model without UIP shocks with its representative-agent counterpart; panel (b) adds the observed excess-bond-premium UIP path and re-infers other shocks. The nominal dollar index rises with dollar appreciation. Paths are 100 times log deviations, mean-aligned over 1999Q1–2019Q2 without smoothing or volatility rescaling. All models condition on the same observed relative consumer-price path. Only Home and Foreign discount-factor innovations and a discount-factor disagreement innovation are inferred from long yields and survey disagreement; output is held out. The representative-agent economy uses only its two discount-factor innovations.

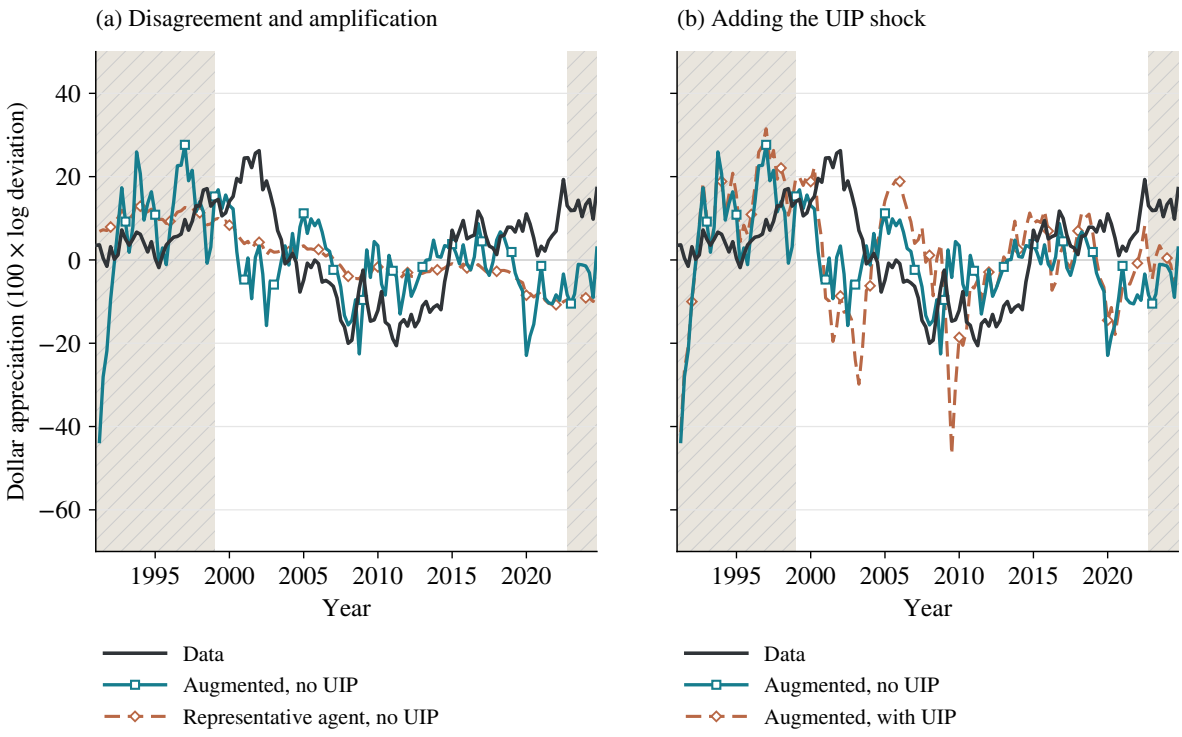


Figure F11: Demand and disagreement inversion at the augmented calibration, 1991Q1–2024Q4

Notes: The historical exercise uses the first-order augmented calibration. Panel (a) compares the augmented model without UIP shocks with its representative-agent counterpart; panel (b) adds the observed excess-bond-premium UIP path and re-infers other shocks. The nominal dollar index rises with dollar appreciation. Paths are 100 times log deviations, mean-aligned over 1999Q1–2019Q2 without smoothing or volatility rescaling. All models condition on the same observed relative consumer-price path. Only the two country discount-factor innovations and a discount-factor disagreement innovation are inferred; output is held out. Inference begins in 1991Q2 without a 1999 reset. Hatching marks the pre-1999 legacy survey and the dispersion-proxy continuation from 2022Q4.

F.3.4 Recursive historical accounting

This complementary accounting exercise feeds observed output, interest rates, disagreement, and realized survey forecast errors into the model's core recursion. It examines how positions and wealth propagate observed news. The representative-agent/no-wealth recursion below is the earlier accounting benchmark, distinct from the solved representative-agent economy in Section 4.4. Because realized survey news contains exchange-rate realizations, the fit reported here is conditional and does not identify the origin of that news.

Specifically, it feeds observed Home and Foreign output, the rate differential, survey disagreement, and survey forecast errors into the model's core recursion[9]. In the augmented model without the exogenous UIP shock, the modified UIP is

$$\Delta\hat{e}^{\text{aug}}_{t+1} = \hat{r}_t - \hat{r}^*_t + \varepsilon^e_{t+1} + \mathcal{N}_t - d_{ss}\widetilde{\omega}^A_t - \delta_w d_t, \tag{58}$$

and the wealth share $\widetilde{\omega}^A_t$ and the NFA terms $\mathcal{N}_t$ are implied from the model's equilibrium conditions, given the data.

The UIP risk version adds excess bond premium ψ_t additively in Equation 58, as in the augmented model and recent segmented asset market models with UIP shocks such as Itskhoki and Mukhin (2021); Kekre and Lenel (2024). Figure F12(a) first asks whether such a UIP-shock proxy improves this conditional historical fit. It does not. Without the proxy, the correlation between model and data in levels is 0.90; including it lowers the correlation to 0.63. The change correlation falls from 0.929 to 0.732 as well. The proxy helps selected unconditional and forecast-horizon moments, but it adds noise to this realized historical path.

The representative-agent/no-wealth comparison removes the wealth-share and disagreement:

$$\Delta\hat{e}^{\text{RA}}_{t+1} = \hat{r}_t - \hat{r}^*_t + \varepsilon^e_{t+1} + \mathcal{N}_t + \psi_t. \tag{59}$$

Figure F12(b) then compares the preferred imputation, which omits the UIP-risk proxy, with the representative-agent/no-wealth recursion in equation (59). The augmented model raises the level correlation from 0.40 to 0.90 and lowers the level root-mean-squared error from 0.377 to 0.082, a 78 percent reduction. The corresponding correlations in quarterly changes are already 0.848 and 0.929 because both recursions share the realized survey-news term. The informative gain is instead in the low-frequency path: endogenous wealth and external positions prevent the cumulated exchange rate from drifting away from the data.

[9]Since I feed survey forecast errors into the model's recursion, this exercise must be seen as a validation of the model's core mechanism, rather than the model's ability to fit the exchange rate given non-exchange rate variables.

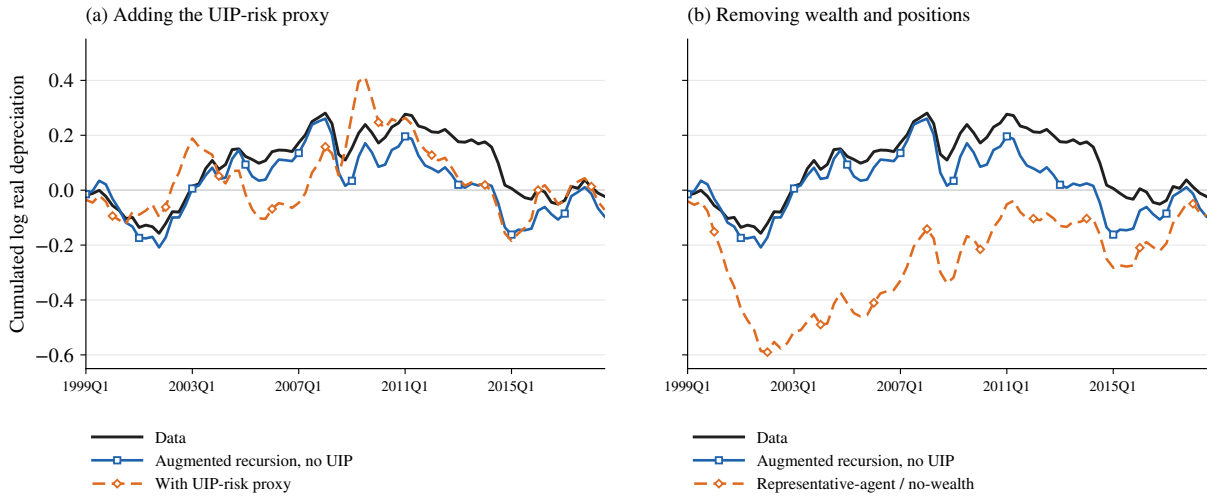


Figure F12: Conditional historical exchange-rate accounting

Notes: Conditional historical accounting for the trade-weighted real dollar exchange rate, cumulated relative to 1998Q4; an increase denotes dollar depreciation. All recursions use observed Home and Foreign output, interest rates, disagreement, and realized survey forecast news, which contains exchange-rate realizations. Square markers identify the augmented recursion without the UIP-risk proxy. The dashed line adds the excess-bond-premium proxy in panel (a) and removes wealth-share and net-foreign-asset accumulation in panel (b). The latter is an accounting benchmark, distinct from the solved representative-agent economy in the historical inversion.

The historical accounting exercise gives the same message as the simulations. Realized survey news anchors much of the quarter-to-quarter movement in both economies. What disagreement adds is an endogenous stock-flow correction: positions redistribute wealth, wealth changes the marginal investor, and the resulting NFA and wealth-share terms keep the cumulated exchange rate from drifting away from the data. Disagreement does not matter primarily by adding another shock. It is the indispensable state through which ordinary news becomes a persistent stock–flow adjustment in the exchange rate.

The recursion figure uses the trade-weighted real dollar exchange rate. The solid line with square markers is the augmented recursion excluding the UIP-risk proxy; the dashed line adds the proxy in panel (a) and removes wealth and NFA accumulation in panel (b). All specifications share the observed rates and survey news, and paths are cumulated relative to 1998Q4. These conventions and the earlier accounting benchmark are retained for this complementary exercise.

G Additional Empirical Results

G.1 Disagreement and exchange-rate changes

The belief component of the model's UIP wedge is $-d_{ss}\widetilde{\omega}_t^A - \delta_w d_t$. Prices depend on both wealth redistribution and the forecast gap; the direct disagreement term vanishes at equal steady-state wealth shares. This motivates asking whether survey dispersion adds explanatory information beyond observed fundamentals, without imposing a common coefficient sign across currencies. Section 5 defines the seven-currency AFE sample, nominal quotation $e_{i,t}^N$, and disagreement $\mathcal{D}_{i,t}^h$ computed across individual log forecasts.

Following the monthly explanatory-regression approach of Engel and Wu (2024), I estimate separately for each AFE currency

$$\Delta \log e_{i,t}^N = a_i + B_i' X_{i,t} + \theta_{1i}\Delta PC_{1,i,t} + \theta_{2i}\Delta PC_{2,i,t} + u_{i,t}. \tag{60}$$

The two principal components summarize the correlation matrix of $\log \mathcal{D}_{i,t}^h$ at $h = 90, 180, 270, 360$ days. I estimate their loadings separately by currency using complete four-horizon observations through February 2021, then take exact calendar-month changes. Forecast quotes are logged individually before calculating their cross-sectional standard deviations. There is no division of raw quote dispersion by a consensus or spot rate. PCA loadings use the full sample, so the exercise is an explanatory regression rather than a real-time forecasting test.

The controls $X_{i,t}$ are the change in the three-month real Home-minus-Foreign interest differential, Foreign and U.S. year-on-year inflation, the lagged log real exchange rate, changes in VIX and the three-month government-bond CIP deviation, and the U.S. goods-and-services trade balance divided by industrial production. Nominal yield differentials are converted from percentage points to decimal units before subtracting inflation differences. Inflation is computed from exact twelve-month log CPI changes. The real-rate proxy uses realized inflation; it is not a directly observed expected real return. The real exchange rate is $\log e_{i,t}^N + \log P_{i,t} - \log P_{US,t}$.

I compare regressions with and without disagreement on identical currency–month observations. Monthly spot changes use differences of monthly mean log rates. All lags refer to calendar months, including periods without valid forecasts. Inference uses Bartlett HAC with six calendar-month lags, zero scores in missing months, and a finite-sample covariance correction. Joint tests condition on estimated PCA loadings; they do not incorporate loading-estimation uncertainty. The table reports unadjusted tests, and I also apply a Holm correction across the seven AFE currencies.

Table G1: Survey disagreement in monthly AFE exchange-rate regressions

Currency	N	Controls only R^2	With dispersion R^2	Δ adjusted R^2 (pp)	Joint p-value
AUD	103	0.338	0.353	0.14	0.252
CAD	97	0.233	0.244	-0.71	0.261
CHF	99	0.130	0.152	0.40	0.371
EUR	104	0.336	0.342	-0.93	0.661
GBP	101	0.318	0.334	0.17	0.372
JPY	104	0.212	0.218	-1.09	0.622
SEK	86	0.287	0.316	1.17	0.183

Notes: Separate contemporaneous OLS regressions for AUD, CAD, CHF, EUR, GBP, JPY, and SEK. The dependent variable is monthly log nominal depreciation against USD, positive when the dollar depreciates. Both specifications use identical observations, an intercept, changes in the three-month real Home-minus-Foreign interest differential, Foreign and U.S. annual log CPI inflation, the lagged log real exchange rate, VIX and government-bond CIP changes, and the U.S. trade-balance/industrial-production control. The added regressors are changes in two currency-specific principal components of log SDs of individual log forecasts at four horizons (90–360 days), requiring at least eight institutions per currency–month–horizon. Complete-case samples contain 86–104 observations per currency, end in January 2021, and contain gaps; the macroeconomic window is restricted through February 2021. The penultimate column is the change in adjusted R^2 in percentage points. Joint p-values use HAC Wald F tests with six calendar-month lags and an $n/(n-k)$ covariance correction, conditional on estimated PCA loadings and unadjusted for multiple tests.

Table G1 reports a median increase in unadjusted R^2 of 1.53 percentage points. No currency's joint test rejects at the unadjusted five-percent level. None of the seven joint tests rejects at five percent after the Holm correction. Adjusted R^2 can fall when the added regressors explain little variation. The results are a supporting diagnostic of contemporaneous association, not a structural estimate of the disagreement wedge, a causal test, or evidence of out-of-sample predictability.