

Exchange-Rate Limit Cycles: A New Forecasting Benchmark*

José Lores-Diz[†] Shengliang Ou[‡] Donghai Zhang[§]

October, 2026

Abstract

We document that exchange-rate cyclical components repeatedly align with their eight-year lags, consistent with limit-cycle theory: without shocks, the economy returns to the same state after a fixed interval. This motivates a limit-cycle benchmark that forecasts exchange-rate changes using the same-horizon change at the corresponding phase of the preceding cycle. Like the random walk, a benchmark long considered difficult to outperform, it uses only exchange-rate history, without macroeconomic or financial predictors. It lowers out-of-sample mean squared forecast errors relative to the random walk for most of the twenty currencies in our sample at twelve-, sixteen-, and twenty-quarter horizons.

Keywords: Exchange-rate forecasting; limit cycles; random walk; out-of-sample predictability.

JEL Classification: F31, F37, F41, E32.

*We are grateful to Guangyu Pei for conversations that inspired this project. We thank Ricardo Duque Gabriel and Lorenzo Zecca for excellent research assistance.

[†]Universitat Pompeu Fabra, jose.loresdiz@gmail.com.

[‡]Shanghai Jiao Tong University, shengliang.ou@gmail.com.

[§]National University of Singapore, donghai.d.zhang@gmail.com.

1 Introduction

A longstanding finding in the exchange-rate forecasting literature is that the random walk is difficult to beat out of sample (Meese and Rogoff, 1983a,b). This finding has been documented across currencies, sample periods, and forecasting methods (Rossi, 2013). The random walk therefore remains the default benchmark for empirical and structural exchange-rate models. Its appeal also rests on its simplicity: it requires no macroeconomic predictors and uses the current exchange rate as its forecast of the future exchange rate.

We document a recurrent pattern in exchange rates that the random-walk forecast does not exploit. In quarterly nominal and real bilateral exchange rates for twenty advanced economies, the cyclical component repeatedly aligns with the same component eight years earlier. Several large peaks and troughs occur at similar points in successive cycles. This evidence suggests that an exchange rate’s own history contains information about its future movement.

To explain this recurrence, we introduce household credit frictions and durable accumulation into the small-open-economy framework of Galí and Monacelli (2005), following Beaudry et al. (2020). An increase in employment lowers default risk and borrowing costs, encouraging households to spend more and reinforcing the expansion. The expansion also increases purchases of new durables. As these purchases build up the durable stock, demand for additional durables weakens, reducing production and employment. During the resulting contraction, the durable stock depreciates, restoring demand for new purchases and initiating another expansion. Together, these forces produce a limit cycle: the economy returns to the same state after a fixed interval, so movements repeat at corresponding phases of successive cycles. The associated movements in domestic consumption generate exchange-rate fluctuations through trade adjustment: higher consumption raises import demand and requires a real depreciation to increase exports.

We use this recurrence to forecast future exchange-rate changes. For a horizon shorter than one cycle, we predict the future log nominal exchange-rate change from the same-horizon change at the corresponding phase of the preceding cycle. We call this the limit-cycle benchmark.

We evaluate the limit-cycle benchmark’s out-of-sample forecasting performance relative to the random walk. At twelve-, sixteen-, and twenty-quarter horizons, it delivers lower mean squared forecast error (MSFE) for fifteen, fourteen, and fourteen of the twenty

currencies. Across all twenty currencies, the median MSFE reductions are approximately 7, 12, and 19 percent, respectively. Both conventional Clark–West (2007) inference and finite-sample wild-bootstrap inference support statistically significant forecast improvements across multiple currencies.

A closer examination relates the forecast improvements over the random walk to the size of the limit-cycle prediction. When the predicted exchange-rate change is close to zero, the limit-cycle forecast is close to the random walk’s no-change forecast, and their squared forecast errors are similar. Among currencies where the limit-cycle benchmark has lower MSFE than the random walk, larger predicted appreciations or depreciations are associated with larger reductions in squared forecast error relative to the random walk. Turning to forecast horizons, out-of-sample improvements over the random walk are more limited at one, four, and eight quarters: fewer currencies have lower MSFE than the random walk, and fewer improvements are statistically significant. Even where the limit-cycle benchmark does not outperform the random walk at these horizons, its MSFE is generally close to that of the random walk.

Lastly, we test whether the gains at longer horizons could arise by chance if exchange rates followed a driftless random walk. We repeat the forecasting exercise on 10,000 simulated random-walk panels. Only 2.89 percent produce at least as many MSFE improvements as observed in the data. The observed forecasting success of the limit-cycle benchmark is therefore unlikely to arise by chance under the random-walk null.

Literature Review Rossi (2013) surveys the macroeconomic predictors considered in the exchange-rate forecasting literature. These predictors include interest-rate differentials across countries motivated by uncovered interest parity (Meese and Rogoff, 1988); price and inflation differentials motivated by purchasing power parity (Rogoff, 1996; Cheung et al., 2005); and money and output differentials in monetary models (Mark, 1995; Mark and Sul, 2001). Productivity-based models use relative productivity differentials as the predictor, following Balassa (1964) and Samuelson (1964); Cheung et al. (2005) evaluate their forecasting performance. Other predictors include external positions in portfolio-balance models (Hooper and Morton, 1982), inflation and output gaps in Taylor-rule models (Engel and West, 2006; Molodtsova and Papell, 2009), and external imbalances in models of international financial adjustment (Gourinchas and Rey, 2007; Della Corte et al., 2012). Related work examines commodity-price and exchange-rate relationships (Chen and Rogoff, 2003; Chen et al., 2010; Ferraro et al., 2015). The difficulty of consistently outperforming the random walk with macroeconomic predictors in out-of-sample

forecasts is known as the Meese–Rogoff puzzle (Rossi, 2013).

Our paper proposes a new forecasting benchmark that is harder to beat than the random walk. Its role as a benchmark rests on its informational simplicity: it uses only the exchange rate’s own history, without macroeconomic or financial predictors. Models that introduce such predictors should therefore be evaluated against the limit-cycle benchmark in addition to the random walk.

Our approach is closely related to research using simple exchange-rate-based predictors. Engel and Wu (2023) argue that economic predictors should add forecasting value beyond the exchange-rate level itself. Eichenbaum et al. (2020) show that the real exchange rate predicts nominal adjustment in inflation-targeting economies, while Kilian and Taylor (2003) explain how nonlinear mean reversion in real exchange rates can generate predictability in nominal exchange rates at longer horizons.¹ Our benchmark uses the same-horizon exchange-rate change at the corresponding phase of the preceding cycle.

Our economic explanation builds on the limit-cycle model of Beaudry et al. (2020). We incorporate its mechanism into the small-open-economy framework of Galí and Monacelli (2005) to explain exchange-rate determination. Existing work in the open-economy literature explains fluctuations through responses to exogenous shocks (Backus et al., 1992; Stockman and Tesar, 1995; Obstfeld and Rogoff, 1995; Chari et al., 2002; and, more recently, Gabaix and Maggiori, 2015; Itskhoki and Mukhin, 2021; Kekre and Lenel, 2024). Our model complements this literature by generating a limit cycle that persists in the absence of shocks. Along this deterministic cycle, exchange-rate movements repeat at corresponding phases of successive cycles; shocks disturb their timing and amplitude.

The paper proceeds as follows. Section 2 defines limit cycles and documents recurrent movements in nominal and real exchange rates. Section 3 develops the small-open-economy model and explains the mechanism generating exchange-rate limit cycles. Section 4 introduces the forecasting benchmark and evaluates it using model-generated data. Section 5 assesses its out-of-sample performance across currencies. The final section concludes.

¹Earlier research on predictability at longer horizons also examines persistent exchange-rate swings and the role of macroeconomic fundamentals. Engel and Hamilton (1990) model prolonged appreciation and depreciation phases as stochastic shifts between trends. Mark (1995) and Engel, Mark and West (2007) find evidence of longer-horizon predictability using macroeconomic fundamentals.

2 Exchange-Rate Cycles and Their Recurrence

This section documents the recurrent exchange-rate movements that motivate our forecasting benchmark.

The Concept of Limit Cycles Like business cycles, limit cycles involve expansions and contractions. They entail a stronger form of recurrence, however: along a deterministic limit cycle, the economy returns to the same state after a fixed interval (Goodwin, 1951; Benhabib and Nishimura, 1979; Beaudry et al., 2020). Movements therefore repeat at corresponding phases of successive cycles. Let x_t be an observable along a limit cycle, and let the positive integer p denote the number of periods needed to complete one full cycle. Then, in a deterministic limit cycle, $x_t = x_{t-p}$: the current series and its p -period lag coincide.

In practice, business-cycle shocks can disturb movements around an underlying limit cycle. These shocks change the timing and amplitude of observed cycles, making exact repetition unlikely. We therefore look for approximate recurrence: repeated alignment of the cyclical component with the same component one cycle earlier.


Data We use quarterly bilateral nominal exchange rates against the U.S. dollar from Bank for International Settlements (2020). The sample spans 1973Q1–2020Q1, ending at the onset of the COVID-19 pandemic and excluding the subsequent pandemic period.² Let $\mathcal{E}_{i,t}$ denote country i 's nominal exchange rate. Using consumer price indices from International Monetary Fund (2020), we construct the corresponding real exchange rate as $Q_{i,t} \equiv \mathcal{E}_{i,t}P_t^* / P_{i,t}$, where $P_{i,t}$ is the domestic consumer price index and P_t^* is the U.S. consumer price index.

We study 20 countries: the 18 economies in Rossi (2013) (Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Greece, Ireland, Italy, Japan, New Zealand, Portugal, Spain, Sweden, Switzerland, and the United Kingdom), plus South Korea and Norway.

Exchange-Rate Limit-Cycle Behavior We examine whether short- to medium-term exchange-rate movements recur at corresponding phases of successive cycles. To iso-

²For the band-pass filter, we also retain observations from 1971Q1–1972Q4 as initial input. The displayed sample begins in 1973Q1. At the other end of the sample, the pandemic period is unusual because it combines exceptional disruptions to economic activity with large policy interventions. Such shocks can alter the timing and amplitude of movements around an underlying limit cycle, weakening their alignment with the preceding cycle and the predictive value of preceding-cycle movements.

Figure 1: Exchange-Rate Limit-Cycle Behavior in the British Pound



Notes: The blue lines plot the 2–12-year band-pass-filtered components of the log nominal and real exchange rates against the U.S. dollar. The red lines plot the same components lagged by 32 quarters (eight years).

late these movements from higher-frequency fluctuations and longer-term trends, we extract the 2–12-year component of each exchange-rate series. Specifically, we apply the band-pass filter of [Christiano and Fitzgerald \(2003\)](#) to $\log \mathcal{E}_{i,t}$ and $\log Q_{i,t}$, retaining periodicities between 8 and 48 quarters. We use eight years as a reference interval, corresponding to the upper end of the conventional business-cycle periodicities in [Baxter and King \(1999\)](#). We then compare each filtered component with its eight-year lag to assess whether movements recur at corresponding phases of successive intervals. These filtered exchange-rate series are used only in this section to examine limit-cycle behavior. The forecasting exercises in Section 5 use the original, unfiltered exchange-rate series.

Figure 1 illustrates this comparison for the British pound. In each panel, let x_t denote the filtered exchange-rate component. The blue line plots x_t , and the red line plots x_{t-32} , its eight-year lag. If the component repeated exactly every eight years, the two lines would coincide.

Consistent with limit-cycle behavior, there is a large overlap between the blue and red lines, indicating that several large peaks and troughs occur at similar points in successive eight-year intervals. This pattern appears in both nominal and real exchange rates. Figures A.1 and A.2 in the appendix show that this recurrence extends beyond the United Kingdom.

To quantify the alignment, we calculate, for each currency, the correlation between the filtered nominal exchange rate and its 32-quarter lag. Table 1 reports the average across

Table 1: Eight-Year Exchange-Rate Alignment against the Random-Walk Null

Currency panel	Observed data	Monte Carlo wild bootstrap under the random-walk null	
	Mean correlation	Mean correlation	p -value
All twenty currencies	0.505	0.166	0.0029

Notes: Alignment is the correlation between the 2–12-year component of a log nominal exchange rate and its 32-quarter lag, averaged equally across currencies. Correlations use 165 paired observations from 1979Q1–2020Q1; filter input begins in 1971Q1. The random-walk mean and upper-tail p -value use 10,000 wild-bootstrap panels with the same detrending, filter, and comparison dates. Appendix A.1 gives the details.

all twenty currencies, 0.505.

Could this alignment arise simply from applying the band-pass filter to exchange rates whose true data-generating process is a random walk? Using a wild bootstrap based on the observed quarterly exchange-rate changes, we generate 10,000 panels in which each log exchange rate follows a driftless random walk. We first generate the unfiltered random-walk series, then apply the same detrending and 2–12-year band-pass filter used for the observed exchange rates. For each filtered panel, we calculate the correlations with the 32-quarter lags and average them across currencies, exactly as in the data. The mean correlation across these simulated panels is 0.166. Only about 0.29 percent of the panels produce an average correlation at least as high as the observed 0.505. We therefore reject the hypothesis that the observed alignment can be explained by applying the band-pass filter to exchange rates generated by a random walk ($p = 0.0029$). Appendix A.1 describes the statistic and simulations.

This repeated alignment motivates using the exchange-rate change at the corresponding phase of the preceding cycle as a predictor. Before evaluating the forecasting performance of this approach against the random-walk benchmark, we develop a theoretical model in Section 3 that explains how exchange-rate limit cycles can arise.

3 A Model of Exchange-Rate Limit Cycles

To explain the recurrent exchange-rate movements documented in Section 2, we extend the small-open-economy framework of Galí and Monacelli (2005) with the credit-market complementarity and durable accumulation in Beaudry et al. (2020). Section 3.1 presents the model environment. Section 3.2 derives the key equilibrium dynamics, explains the

limit-cycle mechanism, and illustrates the resulting exchange-rate recurrence in Figure 2. Appendix B reports the calibration details, and Appendix D provides the equilibrium conditions and derivations.

3.1 Environment

Time is discrete. The home economy contains households, firms, competitive banks, and a central bank. The economy is small relative to the rest of the world, so foreign output and foreign prices are taken as given.

3.1.1 Households.

There is a unit mass of identical households indexed by $h \in [0, 1]$. Each household is a large family composed of a continuum of workers. Family members pool consumption, labor income, firm profits, and durable holdings. A fraction $e_{h,t}$ of household h 's members is employed in period t , so $1 - e_{h,t}$ is both the unemployment rate within the family and the family's leisure. The lifetime expected utility of household h is

$$U_h = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \xi_{t-1} [U(C_{h,t} - \gamma C_{t-1}) + \eta_L \log(1 - e_{h,t})], \quad (3.1)$$

where $\mathbb{E}_0$ is the expectation conditional on information at date zero, $\beta \in (0, 1)$ is the discount factor, and ξ_{t-1} is an intertemporal preference shifter. The variable $C_{h,t}$ denotes household consumption of non-durable goods, while $C_t \equiv \int_0^1 C_{h,t} dh$ is aggregate consumption. The parameter η_L weights leisure relative to consumption, and $\gamma \in [0, 1)$ governs external habit formation, so current utility depends on consumption relative to the previous period's aggregate consumption. We use $U(a) = (a^{1-\omega} - 1)/(1 - \omega)$, where $\omega > 0$ is the inverse of the intertemporal elasticity of substitution. The preference innovation is $\mu_t \equiv \xi_t/\xi_{t-1}$. It acts as an exogenous demand shock, and we assume $\log \mu_t \sim \mathcal{N}(0, \sigma_u^2)$, where σ_u^2 is the shock variance.³

The consumption bundle is $C_{h,t} = \frac{C_{H,h,t}^{1-\nu} C_{F,h,t}^\nu}{(1-\nu)^{1-\nu} \nu^\nu}$, where $0 < \nu < 1$ and $C_{H,h,t}$ and $C_{F,h,t}$ are household h 's consumption of home and foreign goods. Let $P_{H,t}$ denote the home-currency price index for home goods and $P_{F,t}$ the home-currency price of the foreign composite good. The associated consumer price index is $P_t = P_{H,t}^{1-\nu} P_{F,t}^\nu$. Cost minimiza-

³The model generates a deterministic limit cycle with $\mu_t = 1$; the shock adds business-cycle noise for evaluating the proposed forecasting benchmark.

tion gives home- and foreign-good demand:

$$C_{H,h,t} = (1 - \nu) \frac{P_t C_{h,t}}{P_{H,t}}, \quad C_{F,h,t} = \nu \frac{P_t C_{h,t}}{P_{F,t}}. \quad (3.2)$$

Thus, ν is the expenditure share on imports.

Home goods are themselves a constant-elasticity aggregate of differentiated varieties $i \in [0, 1]$, with $C_{H,h,t} = \left[\int_0^1 C_{H,h,t}(i)^{(\epsilon-1)/\epsilon} di \right]^{\epsilon/(\epsilon-1)}$. The parameter $\epsilon > 1$ is the elasticity of substitution across home varieties. Producer i charges $P_{H,t}(i)$. The corresponding home-good price index is $P_{H,t} = \left[\int_0^1 P_{H,t}(i)^{1-\epsilon} di \right]^{1/(1-\epsilon)}$. Cost minimization gives demand for producer i 's variety:

$$C_{H,h,t}(i) = \left(\frac{P_{H,t}(i)}{P_{H,t}} \right)^{-\epsilon} C_{H,h,t}. \quad (3.3)$$

Households also own the stock of durable goods. These goods are productive inputs used by firms to produce non-durable home consumption goods. Let $X_{h,t}$ denote household h 's stock available for rental in period t , and let $I_{h,t}^X$ denote its purchases of new durables during that period. The stock evolves according to

$$X_{h,t+1} = (1 - \delta)X_{h,t} + I_{h,t}^X, \quad (3.4)$$

where $\delta \in (0, 1)$ is the depreciation rate. Durables purchased in period t become part of the stock available in period $t + 1$. Durable goods do not enter household utility. Households rent the current stock to firms at the nominal rental rate R_t^X , and new durables trade at the nominal price P_t^X .

Financial Frictions and the Household's Intertemporal Budget Constraint. The financial friction operates at the level of individual household members. Each worker must borrow to repay outstanding debt and finance the family's consumption and durable purchases before employment is realized and household income is received. The household head pools members' funds and makes the family's aggregate consumption and durable investment decisions. However, loan obligations remain individual: banks can recover an unemployed member's debt from the family only with probability $\phi \in [0, 1]$. For simplicity, we assume away international financial markets. Household members therefore obtain one-period loans from domestic banks.

Consider worker j in household h . Let $D_{j,h,t}$ denote the worker's outstanding debt

at the beginning of period t , and let $L_{j,h,t}$ denote the new loan. The worker first uses the loan to repay $D_{j,h,t}$ and deposits the remainder, $L_{j,h,t} - D_{j,h,t}$, at the safe deposit rate i_t . The head of the household pools these deposits and pays for the family's consumption and durable purchases. The payments are $P_t C_{h,t}$ for consumption and $P_t^X I_{h,t}^X$ for new durables, where P_t^X is the nominal price of a new durable. Firms receive the corresponding bank transfers before they pay wages and rental income.

Since the lending rate r_t is at least as high as the deposit rate i_t , household members borrow only to meet this liquidity requirement. Aggregating across its members, household h 's new loan therefore satisfies $L_{h,t} = D_{h,t} + P_t C_{h,t} + P_t^X I_{h,t}^X$, where $D_{h,t}$ and $L_{h,t}$ are aggregate beginning-of-period debt and new borrowing.

Employment is realized after borrowing and expenditure decisions. A fraction e_t of workers is employed. Each employed worker earns the nominal wage W_t and repays the loan at the gross lending rate $1 + r_t$. An unemployed worker has no labor income and cannot repay directly. Recovery costs the bank Φ_t per unit of lending. With probability $1 - \phi$, the loan defaults. The probability that a loan is either repaid by the worker or recovered from the family is therefore $e_t + (1 - e_t)\phi$.

These flows determine the household's intertemporal budget constraint. Let $\mathcal{P}_t$ denote nominal firm profits distributed to each household. The household's debt at the beginning of period $t + 1$ is

$$D_{h,t+1} = [e_t + (1 - e_t)\phi] (1 + r_t) (D_{h,t} + P_t C_{h,t} + P_t^X I_{h,t}^X) - (1 + i_t) (e_t W_t + R_t^X X_{h,t} + \mathcal{P}_t). \quad (3.5)$$

The first term is the value of loan obligations that are repaid or recovered. The second term is household income carried at the safe rate: labor income, the rental return $R_t^X X_{h,t}$ on durable holdings, and firm profits.

Optimality Conditions. Given these preferences and the budget constraint (3.5), the household head chooses non-durable consumption and durable purchases and coordinates members' labor-supply decisions to maximize expected lifetime utility in equation (3.1), subject to the durable accumulation equation (3.4). We now discuss the three optimality conditions associated with these choices. Because households are identical, we suppress the household subscript below.

The Euler equation for non-durable consumption is

$$U'(C_t - \gamma C_{t-1}) = \beta \mu_t [e_t + (1 - e_t)\phi] (1 + r_t) \mathbb{E}_t \left[\frac{U'(C_{t+1} - \gamma C_t)}{\Pi_{t+1}} \right]. \quad (3.6)$$

Let $\Pi_{t+1} \equiv P_{t+1}/P_t$ denote gross consumer-price inflation. As in the standard consumption Euler equation, households compare the marginal utility of consumption today with its discounted future value. Here, the relevant intertemporal price depends on the gross lending rate $1 + r_t$, adjusted by the probability that a loan is repaid or recovered, $e_t + (1 - e_t)\phi$. Lower employment raises lending rates and increases the effective cost of current consumption.

The Euler equation for durable purchases is

$$U'(C_t - \gamma C_{t-1}) = \beta \mu_t \mathbb{E}_t \left[\frac{U'(C_{t+1} - \gamma C_t)}{\Pi_{t+1} P_t^X} \left\{ \frac{R_{t+1}^X (1 + i_{t+1})}{[e_{t+1} + (1 - e_{t+1})\phi] (1 + r_{t+1})} + (1 - \delta) P_{t+1}^X \right\} \right]. \quad (3.7)$$

The right-hand side captures the two payoffs from a new durable: the rental return R_{t+1}^X and the remaining resale value $(1 - \delta) P_{t+1}^X$. The rental return is adjusted by the financing wedge between deposits and loans and by the probability of repayment or recovery. Let $w_t \equiv W_t/P_t$ denote the real wage. The household's labor-supply condition is

$$w_t = \frac{\eta_L}{(1 - e_t) U'(C_t - \gamma C_{t-1})} \frac{[e_t + (1 - e_t)\phi] (1 + r_t)}{1 + i_t} + \frac{1 + r_t}{1 + i_t} (1 - \phi) \left(C_t + \frac{P_t^X}{P_t} I_t^X \right). \quad (3.8)$$

In this condition, the first term on the right-hand side compensates a worker for forgone leisure. Because accepting employment also requires repaying the loan, the second term compensates the worker for giving up the possibility of default when unemployed.

3.1.2 Banks with Free Entry and the Lending Premium.

Competitive banks fund household loans at the one-period safe deposit rate i_t and charge the lending rate r_t . Their lending income depends on borrowers' employment outcomes: a loan is repaid by an employed worker with probability e_t and recovered from the family of an unemployed worker with probability $(1 - e_t)\phi$. Recovering a loan from an unemployed worker's family, however, incurs a cost proportional to the loan amount, with per-unit cost Φ_t .⁴ With free entry, the lending rate must therefore allow banks to cover their funding and recovery costs while earning zero expected profits. Defining the gross lending premium as the gross lending rate relative to the gross deposit rate, the free-entry

⁴Section 3.2 parameterizes the recovery cost Φ_t as a cubic polynomial in log employment.

condition therefore gives:

$$R_t^P \equiv \frac{1 + r_t}{1 + i_t} = \frac{1 + (1 - e_t)\phi\Phi_t}{e_t + (1 - e_t)\phi}. \quad (3.9)$$

A decline in employment shifts more loans from direct repayment to costly recovery or default. Banks therefore require a higher lending premium to break even.

3.1.3 Firms.

A continuum of monopolistically competitive firms indexed by $i \in [0, 1]$ produces differentiated non-durable home goods using durable goods as the sole direct input. Each firm chooses its price and the combination of rented durables and labor used in production to maximize profits, subject to demand for its variety. Firm i obtains durables either by renting $X_{i,t}$ units from the existing stock owned by households or by employing $e_{i,t}$ workers to produce $F(e_{i,t})$ new units. Both sources are available for current production, so output is

$$Y_{H,t}(i) = X_{i,t} + F(e_{i,t}), \quad F(e_{i,t}) = e_{i,t}^\alpha. \quad (3.10)$$

The parameter $\alpha \in (0, 1]$ governs returns to labor in the production of new durables. The firm pays the nominal wage W_t for labor and the nominal rental rate R_t^X on rented durables. After current production, it can sell a fraction $\psi \in (0, 1]$ of the newly produced durables at price P_t^X .

For a given quantity of output, the firm therefore chooses rented durables and employment to minimize production costs net of revenue from selling durables. At an interior optimum, the marginal cost of renting an existing durable equals the marginal cost of producing a new one, less the revenue from its subsequent sale:

$$R_t^X = \frac{W_t}{F_e(e_{i,t})} - \psi P_t^X = \frac{W_t}{\alpha e_{i,t}^{\alpha-1}} - \psi P_t^X, \quad (3.11)$$

where F_e is the marginal product of labor in producing new durables. We consider a symmetric equilibrium with $e_{i,t} = e_t$. Defining $I_t^X \equiv \int_0^1 I_{h,t}^X dh$, aggregate purchases of new durables therefore satisfy $I_t^X = \psi F(e_t) = \psi e_t^\alpha$. Thus, ψ is the share of newly produced durables that enters the household-owned stock.

The firm also chooses the price of its variety, $P_{H,t}(i)$, to maximize profits, taking into account how that price affects demand and hence the quantity it must produce. Profits

are distributed equally across households. Changing its price entails the Rotemberg cost $\frac{\chi}{2}[P_{H,t}(i)/P_{H,t-1}(i) - 1]^2 Y_{H,t}(i)$, where $\chi > 0$ governs nominal price rigidity. Let $\Pi_{H,t} \equiv P_{H,t}/P_{H,t-1}$ denote gross home-goods inflation, and let hats denote log deviations from the nonstochastic steady state. Aggregating the firms' price-setting conditions and log-linearizing yields the aggregate New Keynesian Phillips curve:

$$\hat{\pi}_{H,t} = \beta \mathbb{E}_t \hat{\pi}_{H,t+1} + \kappa \widehat{mc}_t, \quad (3.12)$$

where $mc_t \equiv R_t^X/P_{H,t}$ denotes marginal cost in units of home goods and $\widehat{mc}_t \equiv \log(mc_t/\bar{mc})$. The variable $\hat{\pi}_{H,t}$ is the log deviation of $\Pi_{H,t}$, and κ is the Phillips-curve slope implied by χ .

Consumer-price inflation also reflects changes in the terms of trade. Since $P_t = P_{H,t} \mathcal{S}_t^\nu$, gross consumer-price inflation satisfies

$$\Pi_t = \Pi_{H,t} \left(\frac{\mathcal{S}_t}{\mathcal{S}_{t-1}} \right)^\nu. \quad (3.13)$$

In log-linear form, $\hat{\pi}_t = \hat{\pi}_{H,t} + \nu \Delta \hat{s}_t$, where $\hat{\pi}_t$ is the log deviation of Π_t and $\Delta \hat{s}_t \equiv \hat{s}_t - \hat{s}_{t-1}$.

3.1.4 Monetary policy.

The central bank sets i_t , the safe nominal rate on one-period deposits. Following [Beaudry et al. \(2020\)](#), monetary policy follows the Taylor-type rule⁵

$$1 + i_t \approx \Theta \mathbb{E}_t [e_{t+1}^{\varphi_e} \Pi_{t+1}]. \quad (3.14)$$

The scale parameter $\Theta > 0$ helps determine the steady-state nominal rate, while $\varphi_e \geq 0$ governs the response to expected employment. The central bank therefore raises the safe rate when expected employment or expected consumer-price inflation increases.

3.1.5 International Prices and Trade.

Let $\mathcal{E}_t$ denote the nominal exchange rate, measured as units of home currency per unit of foreign currency, and let P_t^* denote the foreign price level. The law of one price implies $P_{F,t} = \mathcal{E}_t P_t^*$. Together with the consumer price index $P_t = P_{H,t}^{1-\nu} P_{F,t}^\nu$, this gives the real

⁵The nonlinear rule and its first-order approximation are reported in [Appendix D](#).

exchange rate:

$$Q_t = \frac{\mathcal{E}_t P_t^*}{P_t} = \frac{P_{F,t}}{P_t} = \mathcal{S}_t^{1-\nu}. \quad (3.15)$$

Increases in $\mathcal{E}_t$ and Q_t represent nominal and real depreciations, respectively.

These relative prices determine export demand. With symmetric preferences and foreign consumption equal to the constant foreign output Y^* , foreign demand for home goods is $\nu \mathcal{S}_t Y^*$. Home households, in turn, spend the fraction ν of consumption expenditure on imports, so $P_{F,t} C_{F,t} = \nu P_t C_t$, where $C_{F,t}$ denotes aggregate imports.

Because households cannot borrow or lend internationally, export revenue must equal import expenditure in every period:

$$P_{H,t} (\nu \mathcal{S}_t Y^*) = P_{F,t} C_{F,t} = \nu P_t C_t. \quad (3.16)$$

Using $P_{H,t} \mathcal{S}_t = P_{F,t}$, this condition yields

$$Q_t = \frac{C_t}{Y^*}. \quad (3.17)$$

The intuition behind this relationship is straightforward. An increase in home consumption raises import expenditure. To finance these imports, the economy must export more, requiring a real depreciation that increases foreign demand for home goods. Domestic consumption cycles therefore generate real exchange-rate cycles. Given domestic and foreign price levels, the corresponding nominal exchange rate is $\mathcal{E}_t = Q_t P_t / P_t^*$.

3.2 Equilibrium Dynamics and the Limit-Cycle Mechanism

Appendix D collects the equilibrium conditions, including durable-rental and home-goods market clearing, and derives the reduced equilibrium dynamics and their coefficient formulas. Here, we focus on four equations that describe the limit-cycle mechanism through the dynamics of the lending premium, the durable stock, employment, and the real exchange rate.

Let hats denote log deviations from the nonstochastic steady state. We parameterize Φ_t as a cubic polynomial in $\hat{e}_t$. Substituting this function into the bank zero-profit condition and taking a third-order approximation of the log lending premium around the steady state gives

$$\hat{R}_t^P = \rho_1 \hat{e}_t + \rho_2 \hat{e}_t^2 + \rho_3 \hat{e}_t^3, \quad (3.18)$$

where $\rho_1 < 0$, while ρ_2 and ρ_3 capture higher-order curvature.⁶

The cubic lending premium is the only nonlinearity retained in the solution; all other equilibrium conditions are log-linearized around the nonstochastic steady state.

Define $\hat{X}_t^* \equiv [\psi/(\alpha\delta)]\hat{X}_t$, where $\hat{X}_t$ is the log deviation of the durable stock. The household and firm optimality conditions, together with bank pricing and market clearing, yield the remaining three equations that are key for the mechanism:

$$\hat{X}_{t+1}^* = (1 - \delta)\hat{X}_t^* + \psi\hat{e}_t, \quad (3.19)$$

$$\hat{e}_t = \tilde{\alpha}_1\hat{X}_t^* + \tilde{\alpha}_2\hat{e}_{t-1} + \tilde{\alpha}_3\mathbb{E}_t\hat{e}_{t+1} + \tilde{\alpha}_4\hat{R}_t^P + \tilde{\alpha}_5\hat{\mu}_t, \quad (3.20)$$

$$\hat{q}_t = (1 - \nu)\frac{\alpha\delta}{\delta + \psi}(\hat{X}_t^* + \hat{e}_t). \quad (3.21)$$

Here $\hat{e}_t$ is the log deviation of employment, $\hat{\mu}_t$ is the preference innovation, and $\hat{q}_t$ is the log deviation of the real exchange rate. The coefficients $\tilde{\alpha}_1, \dots, \tilde{\alpha}_5$ are functions of the structural parameters.

Limit-Cycle Mechanism. We now discuss the intuition for the limit-cycle mechanism. Consider an expansion in which higher production increases the stock of durables available in subsequent periods, as described by equation (3.19). The larger stock reduces the incentive to purchase additional durables, as captured by $\tilde{\alpha}_1 < 0$ in the employment equation (3.20). Eventually, purchases and employment decline, tightening credit conditions and reinforcing the contraction. During the contraction, depreciation gradually reduces the stock and restores demand for new durables, initiating the next expansion. Habit formation slows these adjustments and lengthens each phase of the cycle.


With only a linear premium, the reduced system could not generate a limit cycle, which requires nonlinear dynamics. Under the benchmark calibration, the cubic term in equation (3.18) moderates financial feedback away from the steady state and supports a stable limit cycle.⁷

Through equation (3.21), the limit cycle in durable holdings and employment generates a real exchange-rate cycle. Along the deterministic cycle, exchange-rate movements repeat at corresponding phases of successive cycles. Demand shocks disturb this recurrence by changing the timing and amplitude of observed movements.

⁶Appendix D derives the coefficient formulas.

⁷Under the benchmark calibration, the linear complementarity makes the steady state locally unstable, so oscillations would be explosive without the cubic term. With $\rho_3 > 0$, the cubic term raises the premium relative to its linear value in a boom and lowers it relative to its linear value in a recession.

Figure 2: Exchange-Rate Limit-Cycle Behavior in the Model



Notes: The blue line plots the simulated exchange rate q_t , and the red line plots its 32-quarter (eight-year) lag. The figure shows $T = 200$ observations after discarding the first 1,000 observations.

Figure 2 illustrates this recurrence in the simulated exchange rate. To produce this figure, we follow the benchmark calibration of [Beaudry et al. \(2020\)](#) and calibrate the nonlinear lending premium to produce a cycle of approximately eight years; Appendix B reports the parameter choices and solution method. We solve the resulting nonlinear rational-expectations model using the method of [Galizia \(2021\)](#).

Figure 2 compares the simulated exchange rate q_t with its 32-quarter lag. Peaks and troughs recur at similar phases of successive cycles, resembling the repeated alignment documented in Figure 1 in Section 2.

To quantify this resemblance, we apply the empirical alignment measure to model-generated nominal exchange rates. We draw 10,000 windows of 200 quarters from a long simulation with demand shocks, apply the same 2–12-year band-pass filter, and calculate each component's correlation with its 32-quarter lag. Table 2 reports a mean correlation of 0.630, compared with the observed twenty-currency mean of 0.505 in Table 1. The model therefore generates positive alignment of similar magnitude to the empirical observation.

Table 2: Eight-Year Exchange-Rate Alignment in the Model

	Mean correlation	Central 95% range of model-generated correlations
Limit-cycle model: simulated windows	0.630	[0.205, 0.890]

Notes: Alignment is the correlation between the 2–12-year component of a log nominal exchange rate and its 32-quarter lag. The table summarizes individual-series correlations across 10,000 windows of 200 quarters, drawn with replacement from a 100,000-quarter simulation with demand shocks. Each window supplies 168 paired observations, compared with 165 per currency in the data. The range reports the 2.5th and 97.5th percentiles of the model-generated correlations, not a confidence interval for their mean. Appendix B.2 describes the calculation.

4 Forecasting Exchange Rates Using Limit Cycles

The recurrence documented in the data and generated by the model suggests using exchange-rate movements from the preceding cycle to predict future changes. We call this forecasting approach the limit-cycle benchmark. It uses only the exchange rate’s own history, preserving the random walk’s informational simplicity.

4.1 The Limit-Cycle Benchmark

Random-walk benchmark. Under the random walk,

$$\log \mathcal{E}_t = \log \mathcal{E}_{t-1} + v_t, \quad (4.1)$$

where v_t is an unpredictable innovation with mean zero. Our forecasting target is the cumulative exchange-rate change over the next h quarters:

$$\Delta_h \log \mathcal{E}_{t+h} \equiv \log \mathcal{E}_{t+h} - \log \mathcal{E}_t. \quad (4.2)$$

Because future innovations are unpredictable, the random walk forecasts no change over this horizon. Let $\mathbb{E}_t^m[\cdot]$ denote the forecast under model $m \in \{LC, RW\}$, using information available at date t . The random-walk forecast is

$$\mathbb{E}_t^{RW}[\Delta_h \log \mathcal{E}_{t+h}] = 0. \quad (4.3)$$

Limit-cycle benchmark. The limit-cycle benchmark instead uses the exchange-rate change at the corresponding phase of the preceding cycle. For a deterministic limit

cycle of length p , this change repeats exactly:

$$\Delta_h \log \mathcal{E}_{t+h} = \Delta_h \log \mathcal{E}_{t-p+h}. \quad (4.4)$$

where $\Delta_h \log \mathcal{E}_{t-p+h} = \log \mathcal{E}_{t-p+h} - \log \mathcal{E}_{t-p}$. This preceding-cycle change is observed at forecast date t when $p > h$.

Shocks disturb the timing and amplitude of individual cycles, making exact repetition unlikely. We therefore estimate how much of the preceding-cycle movement recurs:

$$\Delta_h \log \mathcal{E}_{t+h} = \beta_h \Delta_h \log \mathcal{E}_{t-p+h} + \varepsilon_{t+h}, \quad (4.5)$$

The coefficient β_h determines how much of the preceding-cycle exchange-rate change enters the forecast. A value of one predicts that the earlier movement will repeat in full. We omit an intercept so that a value of zero gives the random walk's no-change forecast.

Exact repetition implies a coefficient of one for the deterministic cyclical component. Across countries, the estimated coefficient averages 0.2, so the forecast incorporates, on average, one fifth of the observed preceding-cycle change. This smaller coefficient is consistent with the recurrence documented in Figure 1 and Table 1: those results concern the filtered cyclical component, whereas the raw exchange rate also contains movements that need not recur. For the evaluation of forecasting performance, we use raw, unfiltered exchange-rate data because the target is the actual exchange-rate change, including movements outside cyclical frequencies.

Real-time implementation. To update the forecast as new observations become available, we re-estimate β_h at each forecast date using a rolling window of fixed size L quarterly observations. We denote the resulting estimate by $\hat{\beta}_{t,h}$, where t identifies the estimation date. At date t , the available information is the exchange-rate history $\mathcal{I}_t = \{\mathcal{E}_s : s \leq t\}$. A past observation dated τ can enter the estimation only if its subsequent exchange-rate change, $\log \mathcal{E}_{\tau+h} - \log \mathcal{E}_\tau$, is already observed. This requires $\tau + h \leq t$. The estimation dates are therefore $\mathcal{W}_{t,h} = \{t - h - L + 1, \dots, t - h\}$. Using these observations, we estimate

$$\hat{\beta}_{t,h} = \arg \min_{\beta} \sum_{\tau \in \mathcal{W}_{t,h}} (\Delta_h \log \mathcal{E}_{\tau+h} - \beta \Delta_h \log \mathcal{E}_{\tau-p+h})^2. \quad (4.6)$$

The resulting real-time forecast is

$$\mathbb{E}_t^{LC}[\Delta_h \log \mathcal{E}_{t+h}] = \hat{\beta}_{t,h} \Delta_h \log \mathcal{E}_{t-p+h}. \quad (4.7)$$

Forecast evaluation. Once the exchange-rate change over the next h quarters is observed, we compare it with each model's forecast. For $m \in \{LC, RW\}$, the out-of-sample forecast error is

$$e_{t+h|t}^m = \Delta_h \log \mathcal{E}_{t+h} - \mathbb{E}_t^m[\Delta_h \log \mathcal{E}_{t+h}].$$

Let $\mathcal{T}_h$ denote the set of forecast dates used for evaluation. The mean squared forecast error across these dates is

$$\widehat{MSFE}_h^m = \frac{1}{|\mathcal{T}_h|} \sum_{t \in \mathcal{T}_h} \left(e_{t+h|t}^m \right)^2. \quad (4.8)$$

To compare the two benchmarks, we divide the limit-cycle forecast's MSFE by the random walk's:

$$RelMSFE_h = \frac{\widehat{MSFE}_h^{LC}}{\widehat{MSFE}_h^{RW}}. \quad (4.9)$$

A value below one indicates that the limit-cycle benchmark forecasts more accurately than the random walk.

4.2 Evaluation Using Model-Generated Data

Before evaluating the limit-cycle benchmark using observed exchange rates in Section 5, we examine its forecasting performance using data generated by the model. We ask whether the preceding-cycle predictor improves forecasts when demand shocks disturb the underlying recurrence, and how these gains depend on the estimation-window length and sample volatility.

We simulate the nominal exchange rate from the calibrated model with demand shocks and apply the forecasting procedure in Section 4.1.⁸ We set the cycle length to $p = 32$ quarters and the forecast horizon to $h = 12$ quarters. Thus, the forecast uses the exchange-rate change from $t - 32$ to $t - 20$ to predict the change from t to $t + 12$. We consider rolling estimation windows of $L = 64, 128$, and 256 quarters, corresponding to two, four, and eight cycles.

⁸The log demand shock has mean zero: $\mathbb{E}[\log \mu_t] = 0$. In the third-order solution, inflation and other endogenous variables depend on polynomial terms in the log shock and the state variables, up to degree three. These nonlinear terms can have nonzero means even when the log shock has mean zero; for example, $\mathbb{E}[(\log \mu_t)^2] = \sigma_u^2$. Mean inflation therefore need not be zero. Accumulating this constant component of inflation generates a deterministic trend in the nominal exchange rate. We remove this trend before forecasting so that the average change in the detrended series is zero, as under the random-walk benchmark.

Table 3: Relative MSFE by Rolling-Window Length and Sample Volatility

	Rolling estimation window		
	64 quarters (2 cycles)	128 quarters (4 cycles)	256 quarters (8 cycles)
Long simulated series			
With demand shocks	0.722	0.686	0.672
Without demand shocks	0.075	0.074	0.077
Evaluation blocks with demand shocks, ordered by volatility			
1st–20th volatility percentiles	0.895	0.838	0.823
20th–40th percentiles	0.804	0.766	0.744
40th–60th percentiles	0.746	0.706	0.690
60th–80th percentiles	0.699	0.664	0.649
80th–100th percentiles	0.641	0.613	0.606

Notes: The table reports relative mean squared forecast errors (MSFEs) using model-generated data. Relative MSFE is the limit-cycle benchmark’s MSFE divided by the random walk’s MSFE. Values below one favor the limit-cycle benchmark. The forecasts use $p = 32$ quarters and $h = 12$ quarters. The first two rows use simulated series of 100,000 observations, with and without demand shocks. The remaining rows average relative MSFE across 20,000 evaluation blocks, each containing 200 valid forecasts from the simulation with shocks. All window lengths use the same evaluation dates within a block. The preceding estimation observations are drawn from the long series, and every regression uses only outcomes known at its forecast date. Volatility is the standard deviation of the log nominal exchange-rate level within a block. The reported groups run from the 1st–20th percentiles through the 80th–100th percentiles; the lowest one percent is excluded from the first group.

Table 3 reports two exercises. The first evaluates forecasting performance using long simulated series of 100,000 observations, with and without demand shocks. The second draws 20,000 evaluation blocks from the simulation with shocks, each containing 200 valid forecasts. Blocks are drawn from dates for which all three estimation-window lengths produce valid forecasts, so the specifications are evaluated on the same dates. The estimation windows use preceding observations from the long series and need not fit inside the evaluation block. We group the blocks by nominal exchange-rate volatility and report the average relative MSFE within each percentile range.

The limit-cycle benchmark delivers lower MSFE than the random walk at every estimation-window length and in every volatility group reported in Table 3. With demand shocks, relative MSFE declines from 0.722 with a 64-quarter window to 0.672 with a 256-quarter window, corresponding to MSFE reductions of 27.8 and 32.8 percent relative to the random walk. Without demand shocks, relative MSFE ranges from 0.074 to 0.077. Longer windows also improve average relative performance within every volatility

group, suggesting that observing more cycles helps distinguish recurrence from stochastic variation. Gains are larger in the more volatile evaluation blocks at all three window lengths.

5 Empirical Performance

We now examine whether the forecasting gains in model-generated data extend to observed exchange rates. Section 5.1 describes the out-of-sample forecasting exercise and statistical inference. Section 5.2 reports the forecast comparisons with the random walk and examines when the limit-cycle benchmark performs better. Finally, Section 5.3 tests whether the observed pattern of forecast gains could arise by chance if exchange rates followed a random walk.

5.1 Empirical Implementation

We apply the limit-cycle benchmark in Section 4 separately to the twenty bilateral nominal exchange rates, evaluating out-of-sample forecasts at twelve-, sixteen-, and twenty-quarter horizons. We use a common cycle length of $p = 32$ quarters across currencies and forecast horizons. This eight-year period corresponds to the recurrence examined in Section 2 and the model calibration in Section 3.

At each forecast date, we estimate the recurrence coefficient using the $L = 64$ most recent observations for which the subsequent h -quarter exchange-rate change is already known, following Section 4.1. This window covers two eight-year cycles. A 32-quarter window covers only one cycle and may provide too little information to estimate recurrence reliably, whereas a window longer than 64 quarters leaves fewer observations for out-of-sample forecast evaluation. We use 64 quarters to balance these considerations. The procedure yields 70, 62, and 54 out-of-sample forecasts per currency at the three horizons, respectively.

We compare forecast accuracy using the relative MSFE defined in equation (4.9). The random-walk forecast is nested in the limit-cycle benchmark when $\beta_h = 0$. To assess statistical significance, we therefore use the Clark and West (2007) test, which is designed to compare the predictive accuracy of nested forecasting models.

Persistence and estimation error can affect inference about predictive ability at longer horizons (Rossi, 2005). Alongside the conventional Clark–West p -value, we therefore report a finite-sample p -value that compares the same statistic with its distribution across

10,000 wild-bootstrap simulations under the random-walk null. Following [Engel and Wu \(2023\)](#) and drawing on [Gonçalves and Kilian \(2004\)](#), we construct these simulations from the observed quarterly exchange-rate changes. We first subtract each currency’s average quarterly change. At each date, we multiply all currencies’ demeaned changes by the same randomly drawn sign, either +1 or −1 with equal probability. Signs are drawn independently across dates. This preserves the magnitudes of the demeaned changes and whether currencies move in the same or opposite directions at each date, while imposing unpredictable, zero-mean changes. Accumulating these changes produces the simulated random-walk series. We repeat the complete rolling forecasting procedure on each simulated panel. Appendix [E](#) provides the details.

5.2 Empirical Results

The Limit-Cycle Benchmark Outperforms the Random Walk Table [4](#) reports relative MSFEs, conventional Clark–West p -values, and wild-bootstrap p -values for 12-, 16-, and 20-quarter forecasts. At 12 quarters, the limit-cycle benchmark has a lower MSFE than the random walk for 15 of the 20 currencies, and the Clark–West p -value is below 0.05 for 11. At 16 quarters, the corresponding counts are 14 and 13. At 20 quarters, they are 14 and 11.⁹ In terms of magnitude, the median MSFE reductions across all twenty currencies are approximately 7, 12, and 19 percent at the three horizons, respectively. Accounting for small-sample distortions through the wild bootstrap generally yields larger p -values than conventional Clark–West inference. The bootstrap results continue to support statistically significant forecast gains for several currencies at all three forecast horizons. Restricting attention to the ten non-euro currencies,¹⁰ the limit-cycle benchmark delivers lower MSFE in seven, six, and six cases at the three horizons, respectively.

When Does the Limit-Cycle Benchmark Outperform? The limit-cycle forecast departs from the random walk when it predicts a nonzero exchange-rate change. Its distance from the no-change forecast is the absolute value of the prediction in equation (4.7), $|\hat{\beta}_{t,h}\Delta_h \log \mathcal{E}_{t-p+h}|$. Both the preceding-cycle movement and the rolling coefficient determine this distance. When the prediction is close to zero, the two forecasts and their forecast errors are similar.

⁹Greece has the largest relative forecast losses at all three horizons. Its exchange-rate history includes a controlled depreciation against the European Currency Unit during 1983–1997 ([Lazaretou, 2004](#)).


¹⁰Ten countries in the sample adopted the euro, so their exchange-rate histories combine distinct national-currency movements before adoption with shared euro–dollar movements afterward.

Table 4: Out-of-Sample Forecast Performance against the Random Walk

Country	12 quarters			16 quarters			20 quarters		
	Rel. MSFE	CW <i>p</i> -value	Wild bootstrap <i>p</i> -value	Rel. MSFE	CW <i>p</i> -value	Wild bootstrap <i>p</i> -value	Rel. MSFE	CW <i>p</i> -value	Wild bootstrap <i>p</i> -value
Australia	1.07	0.64	0.52	1.02	0.45	0.44	0.96	0.28	0.38
Austria	0.90	0.02	0.10	0.86	0.00	0.06	0.74	0.00	0.07
Belgium	0.89	0.00	0.04	0.80	0.00	0.00	0.64	0.00	0.00
Canada	0.96	0.11	0.19	0.84	0.04	0.13	0.73	0.02	0.13
Denmark	0.93	0.02	0.10	0.79	0.00	0.00	0.65	0.00	0.00
Finland	1.23	0.88	0.74	1.16	0.98	0.89	1.02	0.51	0.48
France	0.90	0.01	0.07	0.79	0.00	0.00	0.62	0.00	0.00
Germany	0.90	0.02	0.09	0.85	0.00	0.05	0.74	0.00	0.05
Greece	1.73	0.99	0.92	1.94	0.98	0.87	2.19	0.99	0.89
Ireland	0.89	0.00	0.04	0.81	0.00	0.00	0.62	0.00	0.01
Italy	0.93	0.00	0.04	0.83	0.00	0.01	0.69	0.00	0.05
Japan	1.22	0.95	0.85	1.07	0.85	0.68	1.04	0.79	0.62
Korea, Rep. of	0.97	0.18	0.32	1.05	0.25	0.38	1.24	0.44	0.47
New Zealand	0.90	0.00	0.04	0.92	0.01	0.06	1.06	0.42	0.44
Norway	0.98	0.10	0.18	0.89	0.02	0.11	0.73	0.01	0.10
Portugal	0.96	0.14	0.22	0.90	0.11	0.23	0.87	0.14	0.29
Spain	0.84	0.00	0.00	0.73	0.00	0.00	0.57	0.00	0.01
Sweden	1.02	0.52	0.44	0.99	0.04	0.14	0.96	0.02	0.15
Switzerland	0.91	0.05	0.13	1.09	0.81	0.65	1.17	0.86	0.66
United Kingdom	0.86	0.03	0.10	0.87	0.04	0.14	0.98	0.17	0.31

Notes: The table reports forecasts based on a 64-quarter rolling estimation window and a common cycle length of $p = 32$ quarters. Relative mean squared forecast error (MSFE) is the limit-cycle benchmark's MSFE divided by the random walk's MSFE, as defined in equation (4.9). Values below one favor the limit-cycle benchmark. The CW columns report conventional one-sided Clark–West p -values for the alternative that the limit-cycle benchmark has a lower expected forecast loss. The Wild bootstrap columns report one-sided wild-bootstrap p -values for the same Clark–West statistic, obtained from 10,000 simulations under the driftless random-walk null as described in Appendix E. Values are reported to two decimal places; significance counts in the text use the unrounded p -values, and entries shown as 0.00 are positive values below 0.005 after rounding. The baseline sample ends in 2020Q1. The forecast dates span 1999Q4–2017Q1, 2000Q4–2016Q1, and 2001Q4–2015Q1 for the 12-, 16-, and 20-quarter horizons, yielding 70, 62, and 54 forecasts per currency, respectively.

Figure 3: Forecast Gains and the Absolute Limit-Cycle Forecast



Notes: The figure reports forecasts for the United Kingdom. Each point represents a forecast at the 12-quarter horizon. The vertical axis is the random walk's squared forecast error minus the limit-cycle benchmark's squared forecast error. Positive values favor the limit-cycle benchmark. The horizontal axis is the absolute predicted log exchange-rate change, $\left| \hat{\beta}_{t,12}(\ln \varepsilon_{t-p+12} - \ln \varepsilon_{t-p}) \right|$, with $p = 32$. The red line shows the fitted linear relationship.

Figure 3 relates forecast gains to the absolute limit-cycle prediction for the United Kingdom at the twelve-quarter horizon. For each forecast date, the vertical axis is the random walk's squared forecast error minus the limit-cycle benchmark's squared forecast error. Positive values therefore measure gains relative to the random walk. To examine the source of gains where the benchmark outperforms the random walk, Appendix Figure C.1 reports the fifteen currencies with lower MSFE at this horizon. The fitted relationship is positive for all fifteen currencies, including the United Kingdom. A larger predicted change can produce either a larger gain or a larger loss relative to the random walk. Among these successful currencies, the positive fitted relationships indicate that larger limit-cycle predictions track realized exchange-rate changes sufficiently well to deliver larger forecast gains on average.

Forecast Horizon and Data Frequency Forecast gains are concentrated at medium and long horizons. Appendix Table C.1 reports results at 1-, 4-, and 8-quarter horizons.

RelMSFE is below one in 7, 2, and 11 of the 20 currencies, respectively, while Clark–West p -values are below 0.05 in 2, 2, and 5 cases. The corresponding wild-bootstrap counts are 2, 1, and 2. This weaker performance is consistent with short-run exchange-rate variation obscuring the recurrent component at shorter horizons. Even where the limit-cycle benchmark does not outperform the random walk at these horizons, its MSFE is generally close to that of the random walk.

Using monthly data produces a similar pattern. At the corresponding 36-, 48-, and 60-month horizons, the benchmark delivers lower MSFE for fifteen, fourteen, and fourteen currencies, respectively (Appendix Table C.2). The numbers of conventional Clark–West rejections at the five-percent level differ from the quarterly results by at most one. Wild-bootstrap p -values are below 0.05 for four, six, and five currencies, respectively.

5.3 Joint Evidence against the Random-Walk Benchmark

Table 4 evaluates forecast performance separately for each currency and horizon. We now ask whether the overall pattern of forecast gains could arise by chance if exchange rates followed a driftless random walk. We evaluate the sixty currency–horizon comparisons jointly, accounting for shared exchange-rate movements and overlapping forecast periods.

We generate 10,000 wild-bootstrap panels under the driftless random-walk null using the procedure described in Section 5.1. Each panel matches the sample size and forecast horizons of the empirical exercise. We repeat the complete forecasting procedure, re-estimating the model at every forecast date. Table 5 reports how often these random-walk panels produce results at least as favorable as those observed, using three summary statistics. Appendix E provides the details.

The first row of Table 5 evaluates whether the observed MSFE improvements in 43 of the 60 currency–horizon comparisons could arise by chance. Only 2.89 percent of simulated random-walk panels produce at least as many improvements. We therefore reject the random-walk null at the five-percent level based on the observed number of forecast gains.

The second row reduces the repeated representation of the euro area by retaining Germany alongside the ten non-euro currencies. The limit-cycle benchmark delivers lower MSFE in 22 of these 33 comparisons. Only 1.25 percent of simulated panels produce at least as many improvements, so the rejection also holds when the euro area is represented once.

Table 5: Joint Evidence against the Random-Walk Benchmark

Joint statistic	Sample	Observed	Wild bootstrap tail frequency (%)
Comparisons with Rel. MSFE < 1	All 20 currencies	43 of 60	2.89
Comparisons with Rel. MSFE < 1	One euro-area representative	22 of 33	1.25
Maximum CW statistic	All 20 currencies	11.38	0.49

Notes: The table reports joint statistics based on relative mean squared forecast errors (MSFEs) and Clark–West (CW) statistics. A comparison is one currency at one forecast horizon, with $h = 12, 16, 20$ quarters. The final column reports the percentage of 10,000 wild-bootstrap panels generated under the driftless random-walk null that match or exceed the observed statistic. For example, 2.89 corresponds to a tail probability of 0.0289. All statistics use the same joint wild-bootstrap simulations described in Appendix E. The first two rows count comparisons with lower realized MSFE. The second row retains Germany and all ten non-euro currencies, leaving 11 series and 33 comparisons. The final row compares maxima over all 60 comparisons; the observed maximum is attained by Denmark at $h = 16$. The sample and forecast dates match Table 4.

The third row asks whether the strongest individual forecast result could arise by chance when searching across all 60 comparisons. We use the Clark–West statistic underlying the p -values in Table 4, which measures the adjusted forecast-loss improvement in estimated standard-error units. Its observed maximum is 11.38, attained by Denmark at the sixteen-quarter horizon. We select the largest statistic in every simulated panel, allowing each panel the same search across currencies and horizons. Only 0.49 percent produce a maximum at least as large, yielding rejection at the one-percent level.

Taken together, the simulations show that the success of the limit-cycle forecasts is unlikely to arise by chance under the random-walk null.

6 Conclusion

We propose a forecasting benchmark based on the tendency of exchange-rate movements to repeat at corresponding phases of successive cycles. In the data, exchange-rate cyclical components repeatedly align with their eight-year lags. Our model explains how durable accumulation and lending conditions can generate a limit cycle, along which the economy returns to the same state after a fixed interval, causing exchange-rate movements to repeat. Business-cycle shocks disturb this repetition, but movements from the preceding cycle can still help predict subsequent changes. Our limit-cycle benchmark therefore uses the exchange-rate change at the corresponding phase of the preceding cycle to forecast

the change over the same horizon at the current forecast date.

Using only exchange-rate history, the proposed benchmark delivers lower out-of-sample MSFE than the random walk for a majority of the twenty currencies at twelve-, sixteen-, and twenty-quarter horizons. Median MSFE reductions across currencies are approximately 7, 12, and 19 percent, respectively. Among currencies with lower MSFE, gains tend to increase with the absolute limit-cycle forecast. Gains are more limited at shorter horizons. Joint wild-bootstrap tests indicate that the observed pattern of forecast success is unlikely to arise by chance under the random-walk null.

These findings give exchange-rate history a larger role in evaluating forecasting models. Like the random walk, the limit-cycle benchmark requires no macroeconomic or financial predictors, but it captures recurrent movements that the random walk omits. Models that introduce such predictors should therefore be evaluated against the limit-cycle benchmark in addition to the random walk. The relevant question is whether those predictors improve forecasts beyond the information already contained in past exchange-rate movements.

References

- Backus, David K., Patrick J. Kehoe, and Finn E. Kydland**, "International Real Business Cycles," *Journal of Political Economy*, 1992, 100 (4), 745–775.
- Balassa, Bela**, "The Purchasing-Power Parity Doctrine: A Reappraisal," *Journal of Political Economy*, 1964, 72 (6), 584–596.
- Bank for International Settlements**, "BIS nominal exchange rate data set," 2020. Accessed: 2020.
- Baxter, Marianne and Robert G. King**, "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series," *The Review of Economics and Statistics*, 1999, 81 (4), 575–593.
- Beaudry, Paul, Dana Galizia, and Franck Portier**, "Putting the Cycle Back into Business Cycle Analysis," *American Economic Review*, January 2020, 110 (1), 1–47.
- Benhabib, Jess and Kazuo Nishimura**, "The Hopf Bifurcation and the Existence and Stability of Closed Orbits in Multisector Models of Optimal Economic Growth," *Journal of Economic Theory*, 1979, 21 (3), 421–444.
- Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan**, "Can Sticky Price Models Generate Volatile and Persistent Real Exchange Rates?," *The Review of Economic Studies*, 2002,

69 (3), 533–563.

Cheung, Yin-Wong, Menzie D. Chinn, and Antonio Garcia Pascual, “Empirical Exchange Rate Models of the Nineties: Are Any Fit to Survive?,” *Journal of International Money and Finance*, 2005, 24 (7), 1150–1175.

Chin Chen, Yu and Kenneth S. Rogoff, “Commodity Currencies,” *Journal of International Economics*, 2003, 60 (1), 133–160.

—, —, and **Barbara Rossi**, “Can Exchange Rates Forecast Commodity Prices?,” *Quarterly Journal of Economics*, 2010, 125 (3), 1145–1194.

Christiano, Lawrence J. and Terry J. Fitzgerald, “The Band Pass Filter,” *International Economic Review*, 2003, 44 (2), 435–465.

Clark, Todd E. and Kenneth D. West, “Approximately normal tests for equal predictive accuracy in nested models,” *Journal of Econometrics*, 2007, 138 (1), 291–311. 50th Anniversary Econometric Institute.

Della Corte, Pasquale, Lucio Sarno, and Giulia Sestieri, “The Predictive Information Content of External Imbalances for Exchange Rate Returns: How Much Is It Worth?,” *Review of Economics and Statistics*, 2012, 94 (1), 100–115.

Eichenbaum, M S, B K Johannsen, and S T Rebelo, “Monetary Policy and the Predictability of Nominal Exchange Rates,” *The Review of Economic Studies*, 05 2020, 88 (1), 192–228.

Engel, Charles and James D. Hamilton, “Long Swings in the Dollar: Are They in the Data and Do Markets Know It?,” *American Economic Review*, 1990, 80 (4), 689–713.

— and **Kenneth D. West**, “Taylor Rules and the Deutschmark–Dollar Real Exchange Rate,” *Journal of Money, Credit and Banking*, 2006, 38 (5), 1175–1194.

— and **Steve Pak Yeung Wu**, “Forecasting the U.S. Dollar in the 21st Century,” *Journal of International Economics*, 2023, 141, 103715.

—, **Nelson C. Mark, and Kenneth D. West**, “Exchange Rate Models Are Not as Bad as You Think,” *NBER Macroeconomics Annual*, 2007, 22, 381–441.

Ferraro, Domenico, Kenneth S. Rogoff, and Barbara Rossi, “Can Oil Prices Forecast Exchange Rates? An Empirical Analysis of the Relationship between Commodity Prices and Exchange Rates,” *Journal of International Money and Finance*, 2015, 54, 116–141.

Gabaix, Xavier and Matteo Maggiori, “International Liquidity and Exchange Rate Dynamics,” *The Quarterly Journal of Economics*, 2015, 130 (3), 1369–1420.

Galizia, Dana, “Saddle cycles: Solving rational expectations models featuring limit cycles (or chaos) using perturbation methods,” *Quantitative Economics*, 2021, 12 (3), 869–901.

- Galí, Jordi and Tommaso Monacelli**, "Monetary Policy and Exchange Rate Volatility in a Small Open Economy," *The Review of Economic Studies*, 07 2005, 72 (3), 707–734.
- Gonçalves, Sílvia and Lutz Kilian**, "Bootstrapping autoregressions with conditional heteroskedasticity of unknown form," *Journal of Econometrics*, 2004, 123 (1), 89–120.
- Goodwin, Richard M.**, "The Nonlinear Accelerator and the Persistence of Business Cycles," *Econometrica*, 1951, 19 (1), 1–17.
- Gourinchas, Pierre-Olivier and Hélène Rey**, "International Financial Adjustment," *Journal of Political Economy*, 2007, 115 (4), 665–703.
- Hooper, Peter and John Morton**, "Fluctuations in the Dollar: A Model of Nominal and Real Exchange Rate Determination," *Journal of International Money and Finance*, 1982, 1, 39–56.
- International Monetary Fund**, "IMF Data," 2020. Accessed: 2020.
- Itskhoki, Oleg and Dmitry Mukhin**, "Exchange Rate Disconnect in General Equilibrium," *Journal of Political Economy*, 2021, 129 (8), 2183–2232.
- Kekre, Rohan and Moritz Lenel**, "Exchange Rates, Natural Rates, and the Price of Risk," Working Paper 32976, National Bureau of Economic Research September 2024.
- Kilian, Lutz and Mark P. Taylor**, "Why Is It So Difficult to Beat the Random Walk Forecast of Exchange Rates?," *Journal of International Economics*, 2003, 60 (1), 85–107.
- Lazaretou, Sophia**, "Greek Monetary Economics in Retrospect: The Adventures of the Drachma," *Economic Research Department Working Paper*, 2004, (2).
- Mark, Nelson C.**, "Exchange Rates and Fundamentals: Evidence on Long-Horizon Predictability," *American Economic Review*, 1995, 85 (1), 201–218.
- **and Donggyu Sul**, "Nominal Exchange Rates and Monetary Fundamentals: Evidence from a Small Post-Bretton Woods Panel," *Journal of International Economics*, 2001, 53 (1), 29–52.
- 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), 3–24.
- **and Kenneth S. Rogoff**, "The Out-of-Sample Failure of Empirical Exchange Rate Models: Sampling Error or Misspecification?," in Jacob A. Frenkel, ed., *Exchange Rates and International Macroeconomics*, Chicago and London: University of Chicago Press, 1983, pp. 67–112.
- **and —**, "Was It Real? The Exchange Rate–Interest Differential Relation over the Modern Floating-Rate Period," *Journal of Finance*, 1988, 43 (4), 933–948.
- Molodtsova, Tanya and David H. Papell**, "Out-of-Sample Exchange Rate Predictability

- with Taylor Rule Fundamentals," *Journal of International Economics*, 2009, 77 (2), 167–180.
- Obstfeld, Maurice and Kenneth Rogoff**, "Exchange Rate Dynamics Redux," *Journal of Political Economy*, 1995, 103 (3), 624–660.
- Rogoff, Kenneth S.**, "The Purchasing Power Parity Puzzle," *Journal of Economic Literature*, 1996, 34 (2), 647–668.
- Rossi, Barbara**, "Testing Long-Horizon Predictive Ability with High Persistence, and the Meese–Rogoff Puzzle," *International Economic Review*, 2005, 46 (1), 61–92.
- , "Exchange Rate Predictability," *Journal of Economic Literature*, December 2013, 51 (4), 1063–1119.
- Samuelson, Paul A.**, "Theoretical Notes on Trade Problems," *Review of Economics and Statistics*, 1964, 46 (2), 145–154.
- Stockman, Alan C. and Linda L. Tesar**, "Tastes and Technology in a Two-Country Model of the Business Cycle: Explaining International Comovements," *American Economic Review*, 1995, 85 (1), 168–185.

Appendix A Exchange-Rate Recurrence in the Data

Figures A.1 and A.2 extend the comparison in the main text to all twenty currencies. Each panel compares the 2–12-year component of a log exchange rate with the same component eight years earlier. The nominal and real series are shown separately, using the same filter and lag as in Section 2. The filter input covers 1971Q1–2020Q1, and the displayed sample begins in 1973Q1. The simulation test evaluates whether the observed alignment could arise from filtering random walks.

A.1 Simulation Test of Eight-Year Alignment

The test in Table 1 asks whether the observed alignment could arise from applying the same filter to random walks. Let $x_{i,t}$ denote the filtered log nominal exchange rate. For each currency, we calculate the ordinary Pearson correlation

$$A_i = \text{Corr}_{t \in \mathcal{P}}(x_{i,t}, x_{i,t-32}), \quad A_{\mathcal{I}} = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} A_i, \quad (\text{A.1})$$

where $\mathcal{P}$ contains the 165 paired dates from 1979Q1–2020Q1. The two paired series are centered separately in each correlation. The set $\mathcal{I}$ contains all twenty currencies. Correlation measures alignment while allowing amplitudes to differ.


We generate $B = 10,000$ driftless random-walk panels using the common-sign wild procedure in Appendix E.2. We apply the same endpoint detrending and Christiano–Fitzgerald filter as in the data, retaining periodicities of 8–48 quarters. Every panel uses filter input from 1971Q1–2020Q1 and the same paired comparison dates. The raw simulated series are filtered before calculating their 32-quarter correlations.

The upper-tail simulation probability is

$$\hat{p}_{\mathcal{I}} = \frac{1 + \sum_{b=1}^B \mathbf{1}\{A_{\mathcal{I}}^{*(b)} \geq A_{\mathcal{I}}^{obs}\}}{B + 1}. \quad (\text{A.2})$$


Each replication recomputes the twenty-currency mean, accounting for cross-currency dependence and the serial dependence introduced by filtering. The test fixes the 32-quarter lag and filter band; it does not adjust for a search over either. It assesses alignment relative to the specified random-walk null. Other predictable processes can also generate alignment. No forecasting regression or Clark–West statistic enters the exercise.

Figure A.1: Nominal Exchange-Rate Limit-Cycle Behavior across Countries



Notes: The blue lines plot the band-pass-filtered components of log nominal exchange rates against the U.S. dollar. The red lines plot the same components lagged by 32 quarters (eight years). The filter retains periodicities between 8 and 48 quarters (2–12 years). Exact repetition after eight years would make the two lines coincide.

Figure A.2: Real Exchange-Rate Limit-Cycle Behavior across Countries



Notes: The blue lines plot the band-pass-filtered components of log real exchange rates against the U.S. dollar. The red lines plot the same components lagged by 32 quarters (eight years). The filter retains periodicities between 8 and 48 quarters (2–12 years). Exact repetition after eight years would make the two lines coincide.

Appendix B Model Calibration and Simulation

This appendix reports the calibration and solution method used for the exchange-rate simulation in Figure 2, the model-generated alignment measure in Table 2, and the corresponding employment dynamics.

B.1 Calibration and Solution Method

The quantitative experiment follows the benchmark calibration of [Beaudry et al. \(2020\)](#). We set the habit parameter to $\gamma = 0.53$, the depreciation rate to $\delta = 0.05099$, and the inverse intertemporal elasticity of substitution to $\omega = 0.30$. The probability of recovering a loan is $\phi = 0.87$, the monetary-policy response coefficient is $\varphi_e = 0.04$, and the steady-state unemployment target is 0.058. We set the discount factor to $\beta = 0.99$, the Calvo-equivalent price-stickiness target to $\theta = 0.75$, the elasticity of substitution across domestic goods to $\epsilon = 6$, and openness to $\nu = 0.40$. The Rotemberg adjustment-cost parameter matches the Phillips-curve slope implied by the Calvo target.

We calibrate the nonlinear block to the recurrence documented in Section 2. We set $\rho_1 = -0.1565$, $\rho_2 = 0$, and $\rho_3 = 1$. The bank zero-profit condition then implies a steady-state recovery cost of $\Phi = 0.04048$. Together with the depreciation rate, these values generate a deterministic cycle of approximately 32 quarters. We solve the resulting nonlinear rational-expectations model using the method of [Galizia \(2021\)](#).


B.2 Recurrence in Model-Generated Exchange Rates

For the model comparison in Table 2, we draw 10,000 contiguous 200-quarter windows, with replacement, from a 100,000-quarter detrended log nominal exchange-rate simulation. We apply the endpoint detrending and 8–48-quarter Christiano–Fitzgerald filter used in Appendix A.1, then calculate each filtered series’ correlation with its 32-quarter lag. Each window yields 168 paired observations. These are single-series correlations, unlike the twenty-currency panel means. Their mean and central 95 percent range are descriptive statistics, without a test p -value.

B.3 Simulated Employment Dynamics

Figure B.1 reports the simulated employment series that accompanies the exchange-rate simulation in the main text. Employment is the domestic activity variable through which durable accumulation and lending conditions generate the cycle.

Figure B.1: Employment Dynamics in the Model



Notes: The blue line plots the log deviation of employment from its steady state. The figure uses observations 1,001–1,200 of the stochastic simulation, matching the window in Figure 2.

Appendix C Additional Forecasting Results

These exercises examine how forecast performance changes with the horizon and data frequency. The final exercise relates forecast gains to the absolute limit-cycle prediction.

C.1 Shorter Forecast Horizons

Table C.1 evaluates forecasts at one-, four-, and eight-quarter horizons. Relative MSFE is below one for seven, two, and eleven currencies, respectively. Wild-bootstrap p -values are below 0.05 for two, one, and two currencies. Improvements are less common than at the longer horizons in Table 4.

Table C.1: Out-of-Sample Forecast Performance at Shorter Horizons

Country	1 quarter			4 quarters			8 quarters		
	Rel. MSFE	CW p -value	Wild bootstrap p -value	Rel. MSFE	CW p -value	Wild bootstrap p -value	Rel. MSFE	CW p -value	Wild bootstrap p -value
Australia	1.04	0.45	0.36	1.04	0.20	0.22	1.04	0.24	0.28
Austria	0.99	0.07	0.08	1.02	0.31	0.29	0.96	0.17	0.22
Belgium	1.00	0.10	0.10	1.04	0.54	0.45	0.98	0.17	0.22
Canada	1.02	0.14	0.12	1.04	0.67	0.48	1.02	0.44	0.38
Denmark	1.00	0.10	0.10	1.04	0.57	0.47	1.00	0.33	0.33
Finland	1.00	0.10	0.09	1.01	0.20	0.21	1.08	0.31	0.33
France	1.00	0.13	0.12	1.05	0.82	0.67	0.99	0.14	0.20
Germany	0.99	0.07	0.08	1.02	0.30	0.28	0.96	0.17	0.22
Greece	0.99	0.17	0.15	1.09	0.29	0.27	1.19	0.43	0.38
Ireland	1.00	0.18	0.16	1.02	0.70	0.56	0.95	0.00	0.04
Italy	1.01	0.42	0.34	1.01	0.41	0.37	0.96	0.09	0.16
Japan	0.99	0.18	0.19	1.06	0.64	0.54	1.25	0.88	0.75
Korea, Rep. of	1.02	1.00	0.99	1.01	0.60	0.49	1.00	0.19	0.32
New Zealand	1.05	0.29	0.24	1.10	0.22	0.24	0.97	0.02	0.05
Norway	1.01	0.85	0.68	1.04	0.85	0.70	1.02	0.56	0.47
Portugal	0.99	0.08	0.09	1.05	0.99	0.95	1.00	0.20	0.24
Spain	1.00	0.21	0.19	1.03	0.76	0.62	0.98	0.01	0.07
Sweden	0.97	0.02	0.05	1.00	0.11	0.14	1.07	0.41	0.38
Switzerland	1.01	0.13	0.14	0.97	0.05	0.08	0.87	0.05	0.11
United Kingdom	0.98	0.02	0.03	0.93	0.01	0.02	0.85	0.00	0.02

Notes: The table reports forecasts at the 1-, 4-, and 8-quarter horizons for the quarterly sample ending in 2020Q1. The rolling window is 64 quarters and the cycle length is 32 quarters. Relative MSFE is the limit-cycle benchmark's MSFE divided by the random walk's MSFE; values below one favor the limit-cycle benchmark. The CW columns report conventional one-sided Clark–West p -values. The wild-bootstrap columns compare the same statistics with their distributions across 10,000 complete random-walk panels, repeating the rolling forecasting procedure in each panel. Appendix E.5 describes the design-specific simulations. Entries are rounded to two decimal places; counts use unrounded results, and 0.00 denotes a p -value below 0.005.

C.2 Monthly Data

Table C.2 repeats the exercise using monthly data at 36-, 48-, and 60-month horizons. The benchmark has lower MSFE for fifteen, fourteen, and fourteen currencies, respectively, matching the quarterly counts at the corresponding horizons. Wild-bootstrap p -values are below 0.05 for four, six, and five currencies.

Table C.2: Out-of-Sample Forecast Performance Using Monthly Data


Country	36 months			48 months			60 months		
	Rel. MSFE	CW p -value	Wild bootstrap p -value	Rel. MSFE	CW p -value	Wild bootstrap p -value	Rel. MSFE	CW p -value	Wild bootstrap p -value
Australia	1.07	0.60	0.49	1.03	0.50	0.46	0.96	0.29	0.39
Austria	0.91	0.03	0.11	0.86	0.00	0.07	0.74	0.00	0.07
Belgium	0.90	0.00	0.04	0.80	0.00	0.00	0.63	0.00	0.00
Canada	0.96	0.11	0.20	0.84	0.04	0.14	0.73	0.03	0.15
Denmark	0.93	0.02	0.10	0.79	0.00	0.00	0.63	0.00	0.00
Finland	1.23	0.89	0.75	1.17	0.98	0.90	1.03	0.57	0.51
France	0.91	0.01	0.07	0.79	0.00	0.00	0.61	0.00	0.00
Germany	0.91	0.02	0.10	0.86	0.00	0.05	0.74	0.00	0.06
Greece	1.73	0.99	0.93	1.95	0.98	0.88	2.22	0.99	0.90
Ireland	0.90	0.00	0.05	0.81	0.00	0.00	0.62	0.00	0.01
Italy	0.93	0.00	0.06	0.82	0.00	0.01	0.68	0.00	0.05
Japan	1.20	0.96	0.86	1.07	0.85	0.69	1.04	0.82	0.64
Korea, Rep. of	0.97	0.18	0.34	1.05	0.26	0.39	1.27	0.52	0.50
New Zealand	0.88	0.00	0.04	0.94	0.01	0.09	1.05	0.38	0.43
Norway	0.98	0.11	0.20	0.89	0.02	0.12	0.73	0.01	0.11
Portugal	0.96	0.14	0.23	0.89	0.09	0.23	0.87	0.15	0.31
Spain	0.85	0.00	0.00	0.72	0.00	0.00	0.56	0.00	0.01
Sweden	1.02	0.47	0.43	0.99	0.06	0.19	0.96	0.01	0.14
Switzerland	0.92	0.05	0.16	1.10	0.82	0.66	1.18	0.85	0.67
United Kingdom	0.86	0.03	0.11	0.87	0.03	0.15	0.98	0.17	0.32

Notes: The table reports forecasts at the 36-, 48-, and 60-month horizons for the sample ending in January 2020. The rolling window is 192 months and the cycle length is 96 months. Relative MSFE is the limit-cycle benchmark's MSFE divided by the random walk's MSFE; values below one favor the limit-cycle benchmark. The CW columns report conventional one-sided Clark–West p -values. The wild-bootstrap columns compare the same statistics with their distributions across 10,000 complete monthly random-walk panels. Every panel reconstructs the predictors and re-estimates the rolling regressions. Appendix E.5 provides the details. Entries are rounded to two decimal places; counts use unrounded results, and 0.00 denotes a p -value below 0.005.

C.3 Forecast Gains and the Absolute Limit-Cycle Forecast

Figure C.1 examines the source of gains for the fifteen currencies with lower twelve-quarter MSFE than the random walk. Each panel includes all forecast dates and uses the axes in Figure 3. For these currencies, forecasts near zero perform similarly to the random walk, while larger absolute predictions are associated with larger gains. All fifteen fitted slopes are positive.

Figure C.1: Forecast Gains among Currencies with Lower Limit-Cycle MSFE



Notes: The figure includes the fifteen currencies with relative MSFE below one at the 12-quarter horizon in Table 4. Each panel includes all forecast dates for that currency. The vertical axis is the random walk's squared forecast error minus the limit-cycle benchmark's squared forecast error. Positive values favor the limit-cycle benchmark. The horizontal axis is the absolute limit-cycle prediction, including the rolling coefficient. The red lines show fitted linear relationships.

Appendix D Model Equilibrium and Derivations

This appendix collects the equilibrium conditions and derives the reduced dynamics used to explain the limit-cycle mechanism. The derivations express consumption, employment, and prices in terms of the durable stock, lending conditions, and demand shocks. Coefficient formulas are collected at the end for use in numerical implementation.

D.1 Equilibrium Conditions

Households. The consumption and durable-purchase Euler equations are

$$U'(C_t - \gamma C_{t-1}) = \beta \mu_t [e_t + (1 - e_t) \phi] (1 + r_t) E_t \left[\frac{U'(C_{t+1} - \gamma C_t)}{\Pi_{t+1}} \right] \quad (\text{D.1})$$

$$U'(C_t - \gamma C_{t-1}) = \beta \mu_t E_t \left[\frac{U'(C_{t+1} - \gamma C_t)}{\Pi_{t+1} P_t^X} \left\{ \frac{R_{t+1}^X (1 + i_{t+1})}{[e_{t+1} + (1 - e_{t+1}) \phi] (1 + r_{t+1})} + (1 - \delta) P_{t+1}^X \right\} \right] \quad (\text{D.2})$$

The labor-supply condition determines the real reservation wage:

$$w_t = \frac{W_t}{P_t} = \frac{\eta_L}{(1 - e_t) U'(C_t - \gamma C_{t-1})} \frac{[e_t + (1 - e_t) \phi] (1 + r_t)}{1 + i_t} + \frac{1 + r_t}{1 + i_t} (1 - \phi) \left(C_t + \frac{P_t^X}{P_t} I_t^X \right) \quad (\text{D.3})$$

Banks and the lending premium. Free entry equates banks' expected receipts with funding and recovery costs, giving the gross lending premium

$$R_t^P = \frac{1 + r_t}{1 + i_t} = \frac{1 + (1 - e_t) \phi \Phi_t}{e_t + (1 - e_t) \phi} \quad (\text{D.4})$$

To obtain the premium approximation in equation (3.18), let $x = \log(e/\bar{e})$, write $\Phi(x)$ for the recovery-cost function, and define

$$g(x) = \log R^P(\bar{e} e^x, \Phi(x)) - \log \bar{R}^P.$$

The polynomial coefficients satisfy $\rho_j = g^{(j)}(0)/j!$ for $j = 1, 2, 3$. If $\Phi'(0) = 0$, the bank zero-profit condition implies

$$\rho_1 = -\frac{(1 - \phi + \phi\Phi)\bar{e}}{[\bar{e} + (1 - \bar{e})\phi][1 + (1 - \bar{e})\phi\Phi]}.$$

The coefficients ρ_2 and ρ_3 capture the higher-order curvature of bank pricing and the recovery-cost function with respect to log employment.

Production and market clearing. In the symmetric equilibrium, firms rent the entire stock of durables supplied by households. Durable-market clearing and aggregate home-goods production therefore satisfy

$$X_t \equiv \int_0^1 X_{h,t} dh = \int_0^1 X_{i,t} di, \quad Y_{H,t} = X_t + e_t^\alpha. \quad (\text{D.5})$$

Households' aggregate purchases of new durables equal firms' sales:

$$I_t^X \equiv \int_0^1 I_{h,t}^X dh = \psi F(e_t) = \psi e_t^\alpha. \quad (\text{D.6})$$

The home-goods market clears when output net of price-adjustment costs equals domestic demand and exports:

$$Y_{H,t} \left[1 - \frac{\chi}{2} (\Pi_{H,t} - 1)^2 \right] = (1 - \nu) \frac{P_t}{P_{H,t}} C_t + \nu S_t Y^*. \quad (\text{D.7})$$

Balanced trade in equation (3.16) implies $\nu S_t Y^* = \nu P_t C_t / P_{H,t}$. Substituting into the home-goods market-clearing condition gives the resource constraint

$$NX_t = Y_{H,t} \left[1 - \frac{\chi}{2} (\Pi_{H,t} - 1)^2 \right] - \frac{P_t C_t}{P_{H,t}} = 0, \quad (\text{D.8})$$

where NX_t denotes net exports measured in units of home goods.

Durable purchases add to the stock available in the following period:

$$X_{h,t+1} = (1 - \delta) X_{h,t} + I_{h,t}^X. \quad (\text{D.9})$$

Monetary policy. The nonlinear interest-rate rule is

$$1 + i_t = \Theta \frac{E_t[e_{t+1}^{\varphi_e} U'(C_{t+1} - \gamma C_t)]}{E_t[U'(C_{t+1} - \gamma C_t) / \Pi_{t+1}]} \quad (\text{D.10})$$

To first order, this nonlinear rule is equivalent to $1 + i_t \approx \Theta \mathbb{E}_t[e_{t+1}^{\varphi_e} \Pi_{t+1}]$, as reported in equation (3.14).

Exchange rates. The consumption price index links the real exchange rate to the terms of trade:

$$Q_t = \frac{P_{F,t}}{P_t} = S_t^{1-\nu} \quad (\text{D.11})$$

Balanced trade then relates the real exchange rate to domestic consumption:

$$Q_t = \frac{C_t}{Y^*} \quad (\text{D.12})$$

Rental returns and price setting. Firms' choice between labor and rented durables gives

$$r_t^X \equiv \frac{R_t^X}{P_t} = \frac{w_t}{\alpha e_t^{\alpha-1}} - \psi \frac{P_t^X}{P_t} \quad (\text{D.13})$$

The first-order price-setting condition is

$$\widehat{\pi}_{H,t} = \beta E_t \widehat{\pi}_{H,t+1} + \kappa \widehat{mc}_t \quad (\text{D.14})$$

$$\widehat{mc}_t = \widehat{r}_t^X + \nu \widehat{s}_t \quad (\text{D.15})$$

Here $mc_t = R_t^X / P_{H,t} = r_t^X S_t^\nu$ is marginal cost in units of home goods, and $\widehat{mc}_t$ is its log deviation from $\overline{mc} = (\epsilon - 1)/\epsilon$. Rental-cost minimization already accounts for revenue from sales of newly produced durables.

Consumer-price inflation also reflects changes in import prices:

$$\pi_t = \pi_{H,t} + \nu \Delta s_t \quad (\text{D.16})$$

D.2 Consumption, Durable Accumulation, and Employment

The reduced dynamics follow by combining household optimality with monetary policy and market clearing. The utility function is given by

$$U(C_t) = \frac{(C_t - \gamma C_{t-1})^{(1-\omega)} - 1}{1-\omega} + \eta_L \log(1 - e_t)$$

D.2.1 Consumption Euler Equation

Combining equations (D.1), (D.10) and (D.4), we use the bank zero-profit condition to write

$$\begin{aligned} [e_t + (1 - e_t) \phi] (1 + r_t) &= [e_t + (1 - e_t) \phi] R_t^P (1 + i_t) \\ &= [1 + (1 - e_t) \phi \Phi_t] (1 + i_t), \end{aligned} \quad (\text{D.17})$$

or, equivalently,

$$1 + (1 - e_t) \phi \Phi_t = R_t^P [e_t + (1 - e_t) \phi]. \quad (\text{D.18})$$

Substituting this expression into (D.1) gives

$$U'(C_t - \gamma C_{t-1}) = \beta \mu_t R_t^P [e_t + (1 - e_t) \phi] (1 + i_t) E_t \left[\frac{U'(C_{t+1} - \gamma C_t)}{\Pi_{t+1}} \right]. \quad (\text{D.19})$$

Finally, substituting the monetary policy rule (D.10) eliminates $(1 + i_t)$ and gives

$$U'(C_t - \gamma C_{t-1}) = \beta \Theta \mu_t R_t^P [e_t + (1 - e_t) \phi] E_t [e_{t+1}^{\varphi_e} U'(C_{t+1} - \gamma C_t)]. \quad (\text{D.20})$$

Let hatted variables denote log deviations from the nonstochastic steady state and define

$$m_t \equiv U'(\widehat{C_t - \gamma C_{t-1}}).$$

The log-linearized version of (D.20) is

$$m_t = \hat{\mu}_t + \hat{R}_t^P + \frac{(1 - \phi) \bar{e}}{\bar{e} + (1 - \bar{e}) \phi} \hat{e}_t + \varphi_e E_t \hat{e}_{t+1} + E_t m_{t+1}, \quad (\text{D.21})$$

Using

$$U'(C_t - \gamma C_{t-1}) = (C_t - \gamma C_{t-1})^{-\omega},$$

we have

$$m_t = -\frac{\omega}{1 - \gamma} (\hat{c}_t - \gamma \hat{c}_{t-1}). \quad (\text{D.22})$$

Substituting this expression into (D.21) gives

$$E_t \hat{c}_{t+1} - (1 + \gamma) \hat{c}_t + \gamma \hat{c}_{t-1} = \frac{1 - \gamma}{\omega} \left(\hat{\mu}_t + \hat{R}_t^P + \frac{(1 - \phi) \bar{e}}{\bar{e} + (1 - \bar{e}) \phi} \hat{e}_t + \varphi_e E_t \hat{e}_{t+1} \right), \quad (\text{D.23})$$

D.2.2 Durable Accumulation and Consumption

To express durable accumulation in the normalization used in the main text, define

$$\hat{X}_t^* \equiv \frac{\psi}{\alpha \delta} \hat{X}_t, \quad \hat{X}_t = \frac{\alpha \delta}{\psi} \hat{X}_t^*. \quad (\text{D.24})$$

The law of motion of durable goods in (D.9) is given by

$$X_{t+1} = (1 - \delta) X_t + I_t^X, \quad I_t^X = \psi e_t^\alpha,$$

Investment chosen in period t affects the stock available in period $t + 1$. Using the steady-state relation $\bar{I}^X = \delta \bar{X}$, the normalized stock evolves as

$$\hat{X}_{t+1}^* = (1 - \delta) \hat{X}_t^* + \psi \hat{e}_t. \quad (\text{D.25})$$

This is the durable-stock equation (3.19) in the main text.

Now use (D.12), (D.11), (D.8), and $Y_{H,t} = X_t + e_t^\alpha$. Since the Rotemberg cost is second order around the zero-inflation steady state, (D.8) implies $\hat{y}_{H,t} = \nu \hat{s}_t + \hat{c}_t$. Combining this with $\hat{c}_t = (1 - \nu) \hat{s}_t$ gives $\hat{y}_{H,t} = \hat{s}_t$. Moreover,

$$\frac{\bar{X}}{\bar{Y}_H} = \frac{\psi}{\delta + \psi}, \quad \frac{\bar{e}^\alpha}{\bar{Y}_H} = \frac{\delta}{\delta + \psi}.$$

Therefore, after applying the normalization (D.24),

$$\begin{aligned} \hat{c}_t &= (1 - \nu) \hat{y}_{H,t} \\ &= (1 - \nu) \frac{\alpha \delta}{\delta + \psi} (\hat{X}_t^* + \hat{e}_t). \end{aligned} \quad (\text{D.26})$$

Using (D.25), the lagged normalized stock is

$$\hat{X}_{t-1}^* = \frac{\hat{X}_t^* - \psi \hat{e}_{t-1}}{1 - \delta}, \quad (\text{D.27})$$

and the expected future normalized stock is

$$E_t \hat{X}_{t+1}^* = (1 - \delta) \hat{X}_t^* + \psi \hat{e}_t. \quad (\text{D.28})$$

D.2.3 Employment Dynamics

Substituting consumption into its Euler equation yields an equation for employment. The three consumption terms entering (D.23) follow from (D.26), (D.27), and (D.28),

$$\begin{aligned} \hat{c}_t &= (1 - \nu) \frac{\alpha \delta}{\delta + \psi} (\hat{X}_t^* + \hat{e}_t), \\ \hat{c}_{t-1} &= (1 - \nu) \frac{\alpha \delta}{\delta + \psi} \left[\frac{\hat{X}_t^*}{1 - \delta} + \left(1 - \frac{\psi}{1 - \delta} \right) \hat{e}_{t-1} \right], \\ E_t \hat{c}_{t+1} &= (1 - \nu) \frac{\alpha \delta}{\delta + \psi} [(1 - \delta) \hat{X}_t^* + \psi \hat{e}_t + E_t \hat{e}_{t+1}]. \end{aligned}$$

Therefore the left-hand side of (D.23) becomes

$$\begin{aligned} &E_t \hat{c}_{t+1} - (1 + \gamma) \hat{c}_t + \gamma \hat{c}_{t-1} \\ &= (1 - \nu) \frac{\alpha \delta}{\delta + \psi} \left[- \left(\delta + \gamma - \frac{\gamma}{1 - \delta} \right) \hat{X}_t^* - (1 + \gamma - \psi) \hat{e}_t \right. \\ &\quad \left. + \gamma \left(1 - \frac{\psi}{1 - \delta} \right) \hat{e}_{t-1} + E_t \hat{e}_{t+1} \right]. \end{aligned}$$

Substituting this expression into (D.23) gives

$$\begin{aligned} &(1 - \nu) \frac{\alpha \delta}{\delta + \psi} \left[- \left(\delta + \gamma - \frac{\gamma}{1 - \delta} \right) \hat{X}_t^* - (1 + \gamma - \psi) \hat{e}_t \right. \\ &\quad \left. + \gamma \left(1 - \frac{\psi}{1 - \delta} \right) \hat{e}_{t-1} + E_t \hat{e}_{t+1} \right] \\ &= \frac{1 - \gamma}{\omega} \left[\hat{\mu}_t + \hat{R}_t^P + \frac{(1 - \phi) \bar{e}}{\bar{e} + (1 - \bar{e}) \phi} \hat{e}_t + \varphi_e E_t \hat{e}_{t+1} \right]. \end{aligned}$$

Collecting current employment on the left and the remaining terms on the right gives

$$m_1 \hat{e}_t = m_3 \hat{X}_t^* + m_2 \hat{e}_{t-1} + m_4 E_t \hat{e}_{t+1} + m_5 \hat{R}_t^P + m_6 \hat{\mu}_t, \quad (\text{D.29})$$

where

$$\begin{aligned}
m_1 &= (1 - \nu) \frac{\alpha\delta}{\delta + \psi} (1 + \gamma - \psi) + \frac{1 - \gamma}{\omega} \frac{(1 - \phi) \bar{e}}{\bar{e} + (1 - \bar{e}) \phi}, \\
m_2 &= (1 - \nu) \frac{\gamma\alpha\delta}{\delta + \psi} \left(1 - \frac{\psi}{1 - \delta} \right), \\
m_3 &= - (1 - \nu) \frac{\alpha\delta}{\delta + \psi} \left(\delta + \gamma - \frac{\gamma}{1 - \delta} \right), \\
m_4 &= (1 - \nu) \frac{\alpha\delta}{\delta + \psi} - \frac{1 - \gamma}{\omega} \varphi_e, \\
m_5 &= -\frac{1 - \gamma}{\omega}, \quad m_6 = m_5.
\end{aligned} \tag{D.30}$$

Thus, before substituting out the linear part of the risk premium,

$$\hat{e}_t = \alpha_{10} \hat{X}_t^* + \alpha_{20} \hat{e}_{t-1} + \alpha_{30} E_t \hat{e}_{t+1} + \alpha_{40} \hat{R}_t^P + \alpha_{50} \hat{\mu}_t, \tag{D.31}$$

with

$$\alpha_{10} = \frac{m_3}{m_1}, \quad \alpha_{20} = \frac{m_2}{m_1}, \quad \alpha_{30} = \frac{m_4}{m_1}, \quad \alpha_{40} = \frac{m_5}{m_1}, \quad \alpha_{50} = \frac{m_6}{m_1}. \tag{D.32}$$

These are the coefficients denoted by $\tilde{\alpha}_1, \dots, \tilde{\alpha}_5$ in equation (3.20), which retains the full lending premium. Using the premium approximation,

$$\hat{R}_t^P = \rho_1 \hat{e}_t + \hat{r}_t^P, \quad \hat{r}_t^P \equiv \rho_2 \hat{e}_t^2 + \rho_3 \hat{e}_t^3, \tag{D.33}$$

and moving the $\rho_1 \hat{e}_t$ term to the left-hand side gives the employment equation with the linear premium component absorbed into its coefficients:

$$\hat{e}_t = \alpha_1 \hat{X}_t^* + \alpha_2 \hat{e}_{t-1} + \alpha_3 E_t \hat{e}_{t+1} + \alpha_4 \hat{r}_t^P + \alpha_5 \hat{\mu}_t, \tag{D.34}$$

where

$$\begin{aligned}
\alpha_1 &= \frac{\alpha_{10}}{1 - \alpha_{40} \rho_1}, & \alpha_2 &= \frac{\alpha_{20}}{1 - \alpha_{40} \rho_1}, & \alpha_3 &= \frac{\alpha_{30}}{1 - \alpha_{40} \rho_1}, \\
\alpha_4 &= \frac{\alpha_{40}}{1 - \alpha_{40} \rho_1}, & \alpha_5 &= \frac{\alpha_{50}}{1 - \alpha_{40} \rho_1}.
\end{aligned} \tag{D.35}$$

D.3 Inflation, Rental Returns, Durable Prices, and Wages

These conditions determine prices and wages alongside the employment and durable-stock dynamics. Each equation is expressed using the same stock normalization and lending-premium decomposition.

D.3.1 Domestic Inflation

Combining the price-setting condition (D.14) with marginal cost in (D.15) gives the domestic inflation equation. Under Rotemberg pricing with marginal cost common across firms, its slope is

$$\kappa = \frac{\epsilon - 1}{\chi}.$$

We choose the adjustment-cost parameter χ to match the following benchmark slope parameterization

$$\kappa \equiv \frac{(1 - \theta)(1 - \beta\theta)}{\theta} \frac{\alpha}{\alpha + (1 - \alpha)\epsilon}. \quad (\text{D.36})$$

Using $Y_{H,t} = X_t + e_t^\alpha$, the log deviation of domestic output is

$$\begin{aligned} \hat{y}_{H,t} &= \frac{\bar{X}}{\bar{Y}_H} \hat{X}_t + \frac{\alpha \bar{e}^\alpha}{\bar{Y}_H} \hat{e}_t \\ &= \frac{\alpha \delta}{\delta + \psi} (\hat{X}_t^* + \hat{e}_t), \end{aligned} \quad (\text{D.37})$$

where the second equality uses (D.24). Since $\hat{c}_t = (1 - \nu)\hat{s}_t$ and $\hat{y}_{H,t} = \hat{s}_t$, we also have

$$\hat{s}_t = \frac{\alpha \delta}{\delta + \psi} (\hat{X}_t^* + \hat{e}_t).$$

Therefore

$$\begin{aligned} \widehat{mc}_t &= \hat{r}_t^X + \nu \hat{s}_t \\ &= \hat{r}_t^X + \nu \frac{\alpha \delta}{\delta + \psi} \hat{X}_t^* + \nu \frac{\alpha \delta}{\delta + \psi} \hat{e}_t. \end{aligned} \quad (\text{D.38})$$

Substituting (D.38) into (D.14) gives

$$\hat{\pi}_{H,t} = \beta E_t \hat{\pi}_{H,t+1} + \lambda_1 \hat{r}_t^X + \lambda_2 \hat{X}_t^* + \lambda_3 \hat{e}_t, \quad (\text{D.39})$$

with

$$\lambda_1 = \kappa, \quad \lambda_2 = \kappa \nu \frac{\alpha \delta}{\delta + \psi}, \quad \lambda_3 = \kappa \nu \frac{\alpha \delta}{\delta + \psi}. \quad (\text{D.40})$$

D.3.2 Rental Return on Durables

To log-linearize the rental-rate condition (D.13), express prices and returns in consumption units. Let $p_t^X \equiv P_t^X / P_t$ and retain $r_t^X \equiv R_t^X / P_t$ for the real rental return, with R_t^X denoting the nominal rental rate throughout. Define $\hat{r}_t^X \equiv \log(r_t^X / \bar{r}^X)$. The rental-rate condition becomes:

$$r_t^X = \frac{w_t}{\alpha e_t^{\alpha-1}} - \psi p_t^X.$$

Around the steady state, this gives

$$\begin{aligned} \hat{r}_t^X &= \frac{\bar{w}}{\alpha \bar{e}^{\alpha-1} \bar{r}^X} (\hat{w}_t - (\alpha - 1) \hat{e}_t) - \frac{\psi \bar{p}^X}{\bar{r}^X} \hat{p}_t^X \\ &= \lambda_4 \hat{w}_t + \lambda_5 \hat{e}_t + \lambda_6 \hat{p}_t^X, \end{aligned} \quad (\text{D.41})$$

where

$$\lambda_4 = \frac{\bar{w}}{\alpha \bar{e}^{\alpha-1} \bar{r}^X}, \quad \lambda_5 = \lambda_4 (1 - \alpha), \quad \lambda_6 = -\frac{\psi \bar{p}^X}{\bar{r}^X}. \quad (\text{D.42})$$

D.3.3 Durable-Price Euler Equation

The durable-price Euler equation is obtained from (D.2). After using the bank rate at date $t + 1$, it implies

$$\begin{aligned} (C_t - \gamma C_{t-1})^{-\omega} &= \beta E_t \left[\mu_t (C_{t+1} - \gamma C_t)^{-\omega} \right. \\ &\quad \left. \left(\frac{r_{t+1}^X}{p_t^X R_{t+1}^P [e_{t+1} + (1 - e_{t+1}) \phi]} + (1 - \delta) \frac{p_{t+1}^X}{p_t^X} \right) \right]. \end{aligned} \quad (\text{D.43})$$

Let

$$\Gamma_1 \equiv \frac{\bar{r}^X}{\bar{p}^X \bar{R}^P [\bar{e} + (1 - \bar{e}) \phi]}. \quad (\text{D.44})$$

Using

$$m_{t+1} - m_t = \frac{\omega}{1 - \gamma} [(1 + \gamma) \hat{c}_t - \gamma \hat{c}_{t-1} - E_t \hat{c}_{t+1}],$$

the log-linearized durable-price Euler equation is

$$\begin{aligned} \hat{p}_t^X &= \frac{\omega}{(1 - \gamma) \beta (\Gamma_1 + 1 - \delta)} [(1 + \gamma) \hat{c}_t - \gamma \hat{c}_{t-1} - E_t \hat{c}_{t+1}] \\ &\quad + \frac{\Gamma_1}{\Gamma_1 + 1 - \delta} \left(E_t \hat{r}_{t+1}^X - E_t \hat{R}_{t+1}^P - \frac{(1 - \phi) \bar{e}}{\bar{e} + (1 - \bar{e}) \phi} E_t \hat{e}_{t+1} \right) \\ &\quad + \frac{1 - \delta}{\Gamma_1 + 1 - \delta} E_t \hat{p}_{t+1}^X + \frac{1}{\beta (\Gamma_1 + 1 - \delta)} \hat{\mu}_t. \end{aligned} \quad (\text{D.45})$$

From (D.26), (D.27), and (D.28),

$$\begin{aligned} (1 + \gamma) \hat{c}_t - \gamma \hat{c}_{t-1} - E_t \hat{c}_{t+1} &= (1 - \nu) \frac{\alpha \delta}{\delta + \psi} \left[\left(\delta + \gamma - \frac{\gamma}{1 - \delta} \right) \hat{X}_t^* \right. \\ &\quad \left. + (1 + \gamma - \psi) \hat{e}_t + \frac{\gamma (\psi + \delta - 1)}{1 - \delta} \hat{e}_{t-1} - E_t \hat{e}_{t+1} \right]. \end{aligned} \quad (\text{D.46})$$

Using

$$E_t \hat{R}_{t+1}^P = \rho_1 E_t \hat{e}_{t+1} + E_t \hat{r}_{t+1}^P,$$

substitution of (D.46) into (D.45) gives the pre-coefficient expression

$$\begin{aligned} \hat{p}_t^X &= \frac{\omega}{1 - \gamma} \frac{1}{\beta (\Gamma_1 + 1 - \delta)} (1 - \nu) \frac{\alpha \delta}{\delta + \psi} \left[\left(\delta + \gamma - \frac{\gamma}{1 - \delta} \right) \hat{X}_t^* \right. \\ &\quad \left. + (1 + \gamma - \psi) \hat{e}_t + \frac{\gamma (\psi + \delta - 1)}{1 - \delta} \hat{e}_{t-1} - E_t \hat{e}_{t+1} \right] \\ &\quad + \frac{\Gamma_1}{\Gamma_1 + 1 - \delta} E_t \hat{r}_{t+1}^X - \frac{\Gamma_1}{\Gamma_1 + 1 - \delta} E_t \hat{r}_{t+1}^P \\ &\quad - \frac{\Gamma_1}{\Gamma_1 + 1 - \delta} \left[\frac{(1 - \phi) \bar{e}}{\bar{e} + (1 - \bar{e}) \phi} + \rho_1 \right] E_t \hat{e}_{t+1} \\ &\quad + \frac{1 - \delta}{\Gamma_1 + 1 - \delta} E_t \hat{p}_{t+1}^X + \frac{1}{\beta (\Gamma_1 + 1 - \delta)} \hat{\mu}_t. \end{aligned}$$

Therefore

$$\begin{aligned}\hat{p}_t^X &= \lambda_7 \hat{X}_t^* + \lambda_8 \hat{e}_t + \lambda_9 \hat{e}_{t-1} + \lambda_{10} E_t \hat{r}_{t+1}^X + \lambda_{11} E_t \hat{r}_{t+1}^P \\ &\quad + \lambda_{12} E_t \hat{p}_{t+1}^X + \lambda_{13} E_t \hat{e}_{t+1} + \lambda_{14} \hat{\mu}_t,\end{aligned}\tag{D.47}$$

where

$$\begin{aligned}\lambda_7 &= \frac{\omega}{1-\gamma} \frac{1}{\beta(\Gamma_1+1-\delta)} (1-\nu) \frac{\alpha\delta}{\delta+\psi} \left(\delta + \gamma - \frac{\gamma}{1-\delta} \right), \\ \lambda_8 &= \frac{\omega}{1-\gamma} \frac{1}{\beta(\Gamma_1+1-\delta)} (1-\nu) \frac{\alpha\delta}{\delta+\psi} (1+\gamma-\psi), \\ \lambda_9 &= \frac{\omega}{1-\gamma} \frac{1}{\beta(\Gamma_1+1-\delta)} (1-\nu) \frac{\alpha\delta}{\delta+\psi} \frac{\gamma(\psi+\delta-1)}{1-\delta}, \\ \lambda_{10} &= \frac{\Gamma_1}{\Gamma_1+1-\delta}, \quad \lambda_{11} = -\frac{\Gamma_1}{\Gamma_1+1-\delta}, \quad \lambda_{12} = \frac{1-\delta}{\Gamma_1+1-\delta}, \\ \lambda_{13} &= -\frac{\omega}{1-\gamma} \frac{1}{\beta(\Gamma_1+1-\delta)} (1-\nu) \frac{\alpha\delta}{\delta+\psi} \\ &\quad - \frac{\Gamma_1}{\Gamma_1+1-\delta} \left[\frac{(1-\phi)\bar{e}}{\bar{e}+(1-\bar{e})\phi} + \rho_1 \right], \quad \lambda_{14} = \frac{1}{\beta(\Gamma_1+1-\delta)}.\end{aligned}\tag{D.48}$$

D.3.4 Reservation Wage

Using the lending premium in the labor-supply condition (D.3), the reservation-wage equation becomes

$$\begin{aligned}w_t (1-e_t) (C_t - \gamma C_{t-1})^{-\omega} &= \left[\eta_L [e_t + (1-e_t)\phi] \right. \\ &\quad \left. + (1-e_t) (C_t - \gamma C_{t-1})^{-\omega} (1-\phi) \left(C_t + p_t^X I_t^X \right) \right] R_t^P\end{aligned}\tag{D.49}$$

To linearize this expression, evaluate its components at the steady state. The left-hand side of (D.49), the first bracket component, and the second bracket component evaluate to

$$\begin{aligned}\text{left-hand side :} &\quad \bar{w} (1-\bar{e}) [(1-\gamma)\bar{C}]^{-\omega}, \\ \text{first bracket component :} &\quad \eta_L [\bar{e} + (1-\bar{e})\phi], \\ \text{second bracket component :} &\quad (1-\phi) (1-\bar{e}) [(1-\gamma)\bar{C}]^{-\omega} \left(\bar{C} + \bar{p}^X \bar{I}^X \right).\end{aligned}$$

Hence the steady-state value of the full bracket multiplying R_t^P is

$$\Omega_w \equiv \eta_L [\bar{e} + (1 - \bar{e}) \phi] + (1 - \phi) (1 - \bar{e}) [(1 - \gamma) \bar{C}]^{-\omega} (\bar{C} + \bar{p}^X \bar{I}^X),$$

and the steady-state wage equation is

$$\bar{w} (1 - \bar{e}) [(1 - \gamma) \bar{C}]^{-\omega} = \Omega_w \bar{R}^P.$$

Now linearize the bracket step by step. For the first bracket component,

$$\begin{aligned} \eta_L [e_t + (1 - e_t) \phi] &= \eta_L \{ \bar{e} (1 + \hat{e}_t) + [1 - \bar{e} (1 + \hat{e}_t)] \phi \} \\ &= \eta_L [\bar{e} + (1 - \bar{e}) \phi] + \eta_L (1 - \phi) \bar{e} \hat{e}_t. \end{aligned}$$

After subtracting its steady-state value and dividing by the full steady-state bracket Ω_w , this first component contributes

$$\frac{\eta_L (1 - \phi) \bar{e}}{\Omega_w} \hat{e}_t.$$

For the second bracket component, first expand each factor around the steady state:

$$\begin{aligned} 1 - e_t &= (1 - \bar{e}) \left[1 - \frac{\bar{e}}{1 - \bar{e}} \hat{e}_t \right], \\ C_t - \gamma C_{t-1} &= (1 - \gamma) \bar{C} \left[1 + \frac{\hat{C}_t - \gamma \hat{C}_{t-1}}{1 - \gamma} \right], \\ (C_t - \gamma C_{t-1})^{-\omega} &= [(1 - \gamma) \bar{C}]^{-\omega} \left[1 - \frac{\omega}{1 - \gamma} (\hat{C}_t - \gamma \hat{C}_{t-1}) \right], \\ C_t + p_t^X I_t^X &= (\bar{C} + \bar{p}^X \bar{I}^X) \left[1 + \frac{\bar{C} \hat{C}_t + \bar{p}^X \bar{I}^X (\hat{p}_t^X + \hat{I}_t^X)}{\bar{C} + \bar{p}^X \bar{I}^X} \right]. \end{aligned}$$

Multiplying these first-order terms and dropping products of deviations gives

$$\begin{aligned} &(1 - e_t) (C_t - \gamma C_{t-1})^{-\omega} (1 - \phi) (C_t + p_t^X I_t^X) \\ &= (1 - \phi) (1 - \bar{e}) [(1 - \gamma) \bar{C}]^{-\omega} (\bar{C} + \bar{p}^X \bar{I}^X) \left[1 - \frac{\bar{e}}{1 - \bar{e}} \hat{e}_t - \frac{\omega}{1 - \gamma} (\hat{C}_t - \gamma \hat{C}_{t-1}) \right. \\ &\quad \left. + \frac{\bar{C} \hat{C}_t + \bar{p}^X \bar{I}^X (\hat{p}_t^X + \hat{I}_t^X)}{\bar{C} + \bar{p}^X \bar{I}^X} \right]. \end{aligned}$$

Therefore, after subtracting the second component's steady-state value and dividing by Ω_w , its contribution to the full bracket is first

$$\frac{(1-\phi)(1-\bar{e})[(1-\gamma)\bar{C}]^{-\omega}(\bar{C}+\bar{p}^X\bar{I}^X)}{\Omega_w} \left[-\frac{\bar{e}}{1-\bar{e}}\hat{e}_t - \frac{\omega}{1-\gamma}(\hat{c}_t - \gamma\hat{c}_{t-1}) + \frac{\bar{C}\hat{c}_t + \bar{p}^X\bar{I}^X(\hat{p}_t^X + \hat{I}_t^X)}{\bar{C} + \bar{p}^X\bar{I}^X} \right].$$

Keeping the internal denominator visible, this equals

$$\begin{aligned} & \frac{(1-\phi)(1-\bar{e})[(1-\gamma)\bar{C}]^{-\omega}(\bar{C}+\bar{p}^X\bar{I}^X)}{\Omega_w} \left[-\frac{\bar{e}}{1-\bar{e}}\hat{e}_t - \frac{\omega}{1-\gamma}(\hat{c}_t - \gamma\hat{c}_{t-1}) \right] \\ & + \frac{(1-\phi)(1-\bar{e})[(1-\gamma)\bar{C}]^{-\omega}(\bar{C}+\bar{p}^X\bar{I}^X)}{\Omega_w} \frac{\bar{C}}{\bar{C} + \bar{p}^X\bar{I}^X} \hat{c}_t \\ & + \frac{(1-\phi)(1-\bar{e})[(1-\gamma)\bar{C}]^{-\omega}(\bar{C}+\bar{p}^X\bar{I}^X)}{\Omega_w} \frac{\bar{p}^X\bar{I}^X}{\bar{C} + \bar{p}^X\bar{I}^X} (\hat{p}_t^X + \hat{I}_t^X). \end{aligned}$$

The $\bar{C} + \bar{p}^X\bar{I}^X$ factors cancel only in the last two terms, so the same expression can be written as

$$\begin{aligned} & \frac{(1-\phi)(1-\bar{e})[(1-\gamma)\bar{C}]^{-\omega}(\bar{C}+\bar{p}^X\bar{I}^X)}{\Omega_w} \left[-\frac{\bar{e}}{1-\bar{e}}\hat{e}_t - \frac{\omega}{1-\gamma}(\hat{c}_t - \gamma\hat{c}_{t-1}) \right] \\ & + \frac{(1-\phi)(1-\bar{e})[(1-\gamma)\bar{C}]^{-\omega}\bar{C}}{\Omega_w} \hat{c}_t \\ & + \frac{(1-\phi)(1-\bar{e})[(1-\gamma)\bar{C}]^{-\omega}\bar{p}^X\bar{I}^X}{\Omega_w} (\hat{p}_t^X + \hat{I}_t^X). \end{aligned}$$

The expansion can be summarized using the wage coefficients

$$\begin{aligned} q_{w0} &\equiv \frac{\eta_L [\bar{e} + (1-\bar{e})\phi]}{\Omega_w}, \\ q_{wC} &\equiv \frac{(1-\phi)(1-\bar{e})[(1-\gamma)\bar{C}]^{-\omega}\bar{C}}{\Omega_w}, \quad q_{wI} \equiv \frac{(1-\phi)(1-\bar{e})[(1-\gamma)\bar{C}]^{-\omega}\bar{p}^X\bar{I}^X}{\Omega_w}. \end{aligned}$$

These definitions imply

$$q_{wC} + q_{wI} = \frac{(1-\phi)(1-\bar{e})[(1-\gamma)\bar{C}]^{-\omega}(\bar{C} + \bar{p}^X\bar{I}^X)}{\Omega_w},$$

$$q_{w0} + q_{wC} + q_{wI} = 1,$$

The coefficients q_{wC} and q_{wI} are the normalized weights on consumption and durable expenditure in the reservation-wage condition.

Adding the log deviations of the three left-hand-side factors gives

$$\text{left-hand-side log-deviation} = \hat{w}_t - \frac{\bar{e}}{1 - \bar{e}} \hat{e}_t - \frac{\omega}{1 - \gamma} (\hat{c}_t - \gamma \hat{c}_{t-1}).$$

Therefore the whole right-hand-side bracket, after subtracting Ω_w and dividing by Ω_w , log-linearizes to

$$\begin{aligned} & \frac{\eta_L (1 - \phi) \bar{e}}{\Omega_w} \hat{e}_t + (q_{wC} + q_{wI}) \left[-\frac{\bar{e}}{1 - \bar{e}} \hat{e}_t - \frac{\omega}{1 - \gamma} (\hat{c}_t - \gamma \hat{c}_{t-1}) \right] \\ & + q_{wC} \hat{c}_t + q_{wI} (\hat{p}_t^X + \hat{I}_t^X). \end{aligned}$$

Multiplication by R_t^P adds $\hat{R}_t^P$. Combining the two sides of (D.49) gives

$$\begin{aligned} \hat{w}_t - \frac{\bar{e}}{1 - \bar{e}} \hat{e}_t - \frac{\omega}{1 - \gamma} (\hat{c}_t - \gamma \hat{c}_{t-1}) &= \frac{\eta_L (1 - \phi) \bar{e}}{\Omega_w} \hat{e}_t \\ &+ (q_{wC} + q_{wI}) \left[-\frac{\bar{e}}{1 - \bar{e}} \hat{e}_t - \frac{\omega}{1 - \gamma} (\hat{c}_t - \gamma \hat{c}_{t-1}) \right] \\ &+ q_{wC} \hat{c}_t + q_{wI} (\hat{p}_t^X + \hat{I}_t^X) + \hat{R}_t^P. \end{aligned}$$

Because $q_{w0} + q_{wC} + q_{wI} = 1$, the terms moved from the left-hand side combine with the negative product-share terms to leave q_{w0} multiplying them. Solving for $\hat{w}_t$ gives

$$\begin{aligned} \hat{w}_t &= q_{w0} \frac{\bar{e}}{1 - \bar{e}} \hat{e}_t + q_{w0} \frac{\omega}{1 - \gamma} (\hat{c}_t - \gamma \hat{c}_{t-1}) \\ &+ \frac{\eta_L (1 - \phi) \bar{e}}{\Omega_w} \hat{e}_t + q_{wC} \hat{c}_t + q_{wI} (\hat{p}_t^X + \hat{I}_t^X) + \hat{R}_t^P. \end{aligned}$$

The consumption terms are

$$\begin{aligned} \hat{c}_t &= (1 - \nu) \frac{\alpha \delta}{\delta + \psi} (\hat{X}_t^* + \hat{e}_t), \\ \hat{c}_t - \gamma \hat{c}_{t-1} &= (1 - \nu) \frac{\alpha \delta}{\delta + \psi} \left[\left(1 - \frac{\gamma}{1 - \delta} \right) \hat{X}_t^* + \hat{e}_t + \frac{\gamma (\psi + \delta - 1)}{1 - \delta} \hat{e}_{t-1} \right], \end{aligned}$$

where (D.26) and (D.27) have been used. Also,

$$\hat{I}_t^X = \alpha \hat{e}_t, \quad \hat{R}_t^P = \rho_1 \hat{e}_t + \hat{r}_t^P.$$

Substituting these objects into the solved wage equation and collecting by state, price, and premium terms gives

$$\begin{aligned} \hat{w}_t = & q_{w0} \frac{\omega}{1-\gamma} \left(1 - \frac{\gamma}{1-\delta} \right) (1-\nu) \frac{\alpha\delta}{\delta+\psi} \hat{X}_t^* + q_{wC} (1-\nu) \frac{\alpha\delta}{\delta+\psi} \hat{X}_t^* \\ & + q_{w0} \frac{\bar{e}}{1-\bar{e}} \hat{e}_t + \frac{\eta_L (1-\phi) \bar{e}}{\Omega_w} \hat{e}_t \\ & + q_{w0} \frac{\omega}{1-\gamma} (1-\nu) \frac{\alpha\delta}{\delta+\psi} \hat{e}_t + q_{wC} (1-\nu) \frac{\alpha\delta}{\delta+\psi} \hat{e}_t \\ & + q_{wI} \alpha \hat{e}_t + \rho_1 \hat{e}_t \\ & + q_{w0} \frac{\omega\gamma}{1-\gamma} (1-\nu) \frac{\alpha\delta}{\delta+\psi} \frac{\psi+\delta-1}{1-\delta} \hat{e}_{t-1} \\ & + q_{wI} \hat{p}_t^X + \hat{r}_t^P. \end{aligned}$$

Reading off these coefficients gives the reduced-form wage equation:

$$\hat{w}_t = \lambda_{15} \hat{X}_t^* + \lambda_{16} \hat{e}_t + \lambda_{17} \hat{e}_{t-1} + \lambda_{18} \hat{p}_t^X + \lambda_{19} \hat{r}_t^P, \quad (\text{D.50})$$

where

$$\begin{aligned} \lambda_{15} &= q_{w0} \left(\frac{\omega}{1-\gamma} - \frac{\omega\gamma}{(1-\gamma)(1-\delta)} \right) (1-\nu) \frac{\alpha\delta}{\delta+\psi} + q_{wC} (1-\nu) \frac{\alpha\delta}{\delta+\psi}, \\ \lambda_{16} &= q_{w0} \frac{\bar{e}}{1-\bar{e}} + \frac{\eta_L (1-\phi) \bar{e}}{\Omega_w} + q_{w0} \frac{\omega}{1-\gamma} (1-\nu) \frac{\alpha\delta}{\delta+\psi} + q_{wC} (1-\nu) \frac{\alpha\delta}{\delta+\psi} + q_{wI} \alpha + \rho_1, \\ \lambda_{17} &= q_{w0} \frac{\omega\gamma}{1-\gamma} (1-\nu) \frac{\alpha\delta}{\delta+\psi} \frac{\psi+\delta-1}{1-\delta}, \\ \lambda_{18} &= q_{wI}, \quad \lambda_{19} = 1. \end{aligned} \quad (\text{D.51})$$

Here the premium split $\hat{R}_t^P = \rho_1 \hat{e}_t + \hat{r}_t^P$ has already been imposed, so ρ_1 appears in λ_{16} and $\lambda_{19} = 1$.

D.4 Reduced Equilibrium System and Steady State

Collecting the derived conditions gives a system for domestic inflation, the rental return, the durable price, the wage, the durable stock, and employment:

$$\begin{aligned}
\hat{\pi}_{H,t} &= \beta E_t \hat{\pi}_{H,t+1} + \lambda_1 \hat{r}_t^X + \lambda_2 \hat{X}_t^* + \lambda_3 \hat{e}_t, \\
\hat{r}_t^X &= \lambda_4 \hat{w}_t + \lambda_5 \hat{e}_t + \lambda_6 \hat{p}_t^X, \\
\hat{p}_t^X &= \lambda_7 \hat{X}_t^* + \lambda_8 \hat{e}_t + \lambda_9 \hat{e}_{t-1} + \lambda_{10} E_t \hat{r}_{t+1}^X + \lambda_{11} E_t \hat{r}_{t+1}^P \\
&\quad + \lambda_{12} E_t \hat{p}_{t+1}^X + \lambda_{13} E_t \hat{e}_{t+1} + \lambda_{14} \hat{\mu}_t, \\
\hat{w}_t &= \lambda_{15} \hat{X}_t^* + \lambda_{16} \hat{e}_t + \lambda_{17} \hat{e}_{t-1} + \lambda_{18} \hat{p}_t^X + \lambda_{19} \hat{r}_t^P, \\
\hat{X}_{t+1}^* &= (1 - \delta) \hat{X}_t^* + \psi \hat{e}_t, \\
\hat{e}_t &= \lambda_{20} \hat{X}_t^* + \lambda_{21} \hat{e}_{t-1} + \lambda_{22} E_t \hat{e}_{t+1} + \lambda_{23} \hat{r}_t^P + \lambda_{24} \hat{\mu}_t.
\end{aligned} \tag{D.52}$$

Finally, the steady-state block follows from the bank zero-profit condition, the consumption Euler equation with monetary policy, durable accumulation, goods production, the durable Euler equation, the wage equation, and the rental-rate condition:

$$\begin{aligned}
\bar{R}^P &= \frac{1 + (1 - \bar{e}) \phi \Phi}{\bar{e} + (1 - \bar{e}) \phi}, \\
1 &= \beta \Theta \bar{R}^P [\bar{e} + (1 - \bar{e}) \phi] \bar{e}^{\varphi_e}, \\
\bar{I}^X &= \psi \bar{e}^\alpha, \quad \bar{X} = \frac{\psi \bar{e}^\alpha}{\delta}, \quad \bar{Y}_H = \bar{e}^\alpha + \bar{X}, \quad \bar{C} = \bar{Y}_H^{1-\nu}, \\
1 &= \beta \left(\frac{\bar{r}^X}{\bar{p}^X \bar{R}^P [\bar{e} + (1 - \bar{e}) \phi]} + 1 - \delta \right), \\
\bar{w} (1 - \bar{e}) [(1 - \gamma) \bar{C}]^{-\omega} &= \Omega_w \bar{R}^P, \\
\bar{r}^X &= \frac{\bar{w}}{\alpha \bar{e}^{\alpha-1}} - \psi \bar{p}^X, \\
\bar{Y}_H^{-\nu} &= \frac{\epsilon}{\epsilon - 1} \bar{r}^X.
\end{aligned} \tag{D.53}$$

The last condition imposes firms' steady-state optimal pricing without a subsidy. With foreign output normalized to one, market clearing implies $\bar{p}_H = \bar{Y}_H^{-\nu}$. We retain the calibrated employment target $\bar{e} = 1 - 0.058$ and determine the leisure weight η_L and policy intercept Θ jointly with the remaining steady-state objects. In particular, optimal pricing determines $\bar{r}^X$, the durable Euler condition determines $\bar{p}^X$, and the rental-rate condition determines $\bar{w}$. The wage and consumption Euler conditions then give

$$\eta_L = \frac{(1 - \bar{e})[(1 - \gamma)\bar{C}]^{-\omega}}{\bar{e} + (1 - \bar{e})\phi} \left[\frac{\bar{w}}{\bar{R}^P} - (1 - \phi)(\bar{C} + \bar{p}^X \bar{I}^X) \right],$$

$$\Theta = \frac{1}{\beta \bar{R}^P [\bar{e} + (1 - \bar{e})\phi] \bar{e}^{\varphi_e}}.$$

The calibration requires $\eta_L > 0$. The calibrated values of η_L and Θ are held fixed when deriving the dynamics.

D.5 Coefficient Reference for Numerical Implementation

For implementation, the following list collects the auxiliary objects and coefficients derived above. The dynamic system and steady-state conditions are repeated in compact form so that the required inputs and equations can be read together. Inputs fixed before steady-state calibration: $\alpha, \beta, \gamma, \delta, \omega, \nu, \theta, \epsilon, \phi, \Phi, \psi, \varphi_e, \rho_1, \rho_2, \rho_3$, together with the target $\bar{e}$.

The parameters η_L and Θ are calibrated from the steady-state conditions.

Steady-state objects: $\bar{e}, \bar{C}, \bar{X}, \bar{I}^X, \bar{Y}_H, \bar{p}^X, \bar{R}^P, \bar{r}^X, \bar{w}$.

Auxiliary objects

$$\kappa \equiv \frac{(1 - \theta)(1 - \beta\theta)}{\theta} \frac{\alpha}{\alpha + (1 - \alpha)\epsilon'},$$

$$\Gamma_1 \equiv \frac{\bar{r}^X}{\bar{p}^X \bar{R}^P [\bar{e} + (1 - \bar{e})\phi]},$$

$$\Omega_w \equiv \eta_L [\bar{e} + (1 - \bar{e})\phi] + (1 - \phi)(1 - \bar{e})[(1 - \gamma)\bar{C}]^{-\omega} (\bar{C} + \bar{p}^X \bar{I}^X),$$

$$q_{w0} \equiv \frac{\eta_L [\bar{e} + (1 - \bar{e})\phi]}{\Omega_w},$$

$$q_{wC} \equiv \frac{(1 - \phi)(1 - \bar{e})[(1 - \gamma)\bar{C}]^{-\omega} \bar{C}}{\Omega_w}, \quad q_{wI} \equiv \frac{(1 - \phi)(1 - \bar{e})[(1 - \gamma)\bar{C}]^{-\omega} \bar{p}^X \bar{I}^X}{\Omega_w},$$

$$\hat{R}_t^P = \rho_1 \hat{e}_t + \hat{r}_t^P, \quad \hat{r}_t^P = \rho_2 \hat{e}_t^2 + \rho_3 \hat{e}_t^3.$$

Employment-equation coefficients

$$m_1 = (1 - \nu) \frac{\alpha\delta}{\delta + \psi} (1 + \gamma - \psi) + \frac{1 - \gamma}{\omega} \frac{(1 - \phi)\bar{e}}{\bar{e} + (1 - \bar{e})\phi},$$

$$m_2 = \gamma(1 - \nu) \frac{\alpha\delta}{\delta + \psi} \left(1 - \frac{\psi}{1 - \delta} \right),$$

$$m_3 = -(1 - \nu) \frac{\alpha\delta}{\delta + \psi} \left(\delta + \gamma - \frac{\gamma}{1 - \delta} \right),$$

$$\begin{aligned}
m_4 &= (1 - \nu) \frac{\alpha \delta}{\delta + \psi} - \frac{1 - \gamma}{\omega} \varphi_e, \\
m_5 &= m_6 = -\frac{1 - \gamma}{\omega}, \\
\alpha_{10} &= \frac{m_3}{m_1}, \quad \alpha_{20} = \frac{m_2}{m_1}, \quad \alpha_{30} = \frac{m_4}{m_1}, \\
\alpha_{40} &= \frac{m_5}{m_1}, \quad \alpha_{50} = \frac{m_6}{m_1}, \\
\alpha_1 &= \frac{\alpha_{10}}{1 - \alpha_{40} \rho_1}, \quad \alpha_2 = \frac{\alpha_{20}}{1 - \alpha_{40} \rho_1}, \\
\alpha_3 &= \frac{\alpha_{30}}{1 - \alpha_{40} \rho_1}, \quad \alpha_4 = \frac{\alpha_{40}}{1 - \alpha_{40} \rho_1}, \\
\alpha_5 &= \frac{\alpha_{50}}{1 - \alpha_{40} \rho_1}.
\end{aligned}$$

State-system coefficients

$$\begin{aligned}
\lambda_1 &= \kappa, \\
\lambda_2 &= \kappa \nu \frac{\alpha \delta}{\delta + \psi}, \\
\lambda_3 &= \kappa \nu \frac{\alpha \delta}{\delta + \psi}, \\
\lambda_4 &= \frac{\bar{w}}{\alpha \bar{e}^{\alpha-1} \bar{r}^X}, \\
\lambda_5 &= \lambda_4 (1 - \alpha), \\
\lambda_6 &= -\frac{\psi \bar{p}^X}{\bar{r}^X}, \\
\lambda_7 &= \frac{\omega}{1 - \gamma} \frac{(1 - \nu) \alpha \delta}{(\delta + \psi) \beta (\Gamma_1 + 1 - \delta)} \left(\delta + \gamma - \frac{\gamma}{1 - \delta} \right), \\
\lambda_8 &= \frac{\omega}{1 - \gamma} \frac{(1 - \nu) \alpha \delta}{(\delta + \psi) \beta (\Gamma_1 + 1 - \delta)} (1 + \gamma - \psi), \\
\lambda_9 &= \frac{\omega}{1 - \gamma} \frac{(1 - \nu) \alpha \delta}{(\delta + \psi) \beta (\Gamma_1 + 1 - \delta)} \frac{\gamma (\psi + \delta - 1)}{1 - \delta}, \\
\lambda_{10} &= \frac{\Gamma_1}{\Gamma_1 + 1 - \delta}, \quad \lambda_{11} = -\frac{\Gamma_1}{\Gamma_1 + 1 - \delta}, \quad \lambda_{12} = \frac{1 - \delta}{\Gamma_1 + 1 - \delta}, \\
\lambda_{13} &= -\frac{\omega}{1 - \gamma} \frac{(1 - \nu) \alpha \delta}{(\delta + \psi) \beta (\Gamma_1 + 1 - \delta)} - \frac{\Gamma_1}{\Gamma_1 + 1 - \delta} \left[\frac{(1 - \phi) \bar{e}}{\bar{e} + (1 - \bar{e}) \phi} + \rho_1 \right], \\
\lambda_{14} &= \frac{1}{\beta (\Gamma_1 + 1 - \delta)}, \\
\lambda_{15} &= q_{w0} \left(\frac{\omega}{1 - \gamma} - \frac{\omega \gamma}{(1 - \gamma) (1 - \delta)} \right) (1 - \nu) \frac{\alpha \delta}{\delta + \psi} + q_{wC} (1 - \nu) \frac{\alpha \delta}{\delta + \psi},
\end{aligned}$$

$$\begin{aligned}
\lambda_{16} &= q_{w0} \frac{\bar{e}}{1 - \bar{e}} + \frac{\eta_L (1 - \phi) \bar{e}}{\Omega_w} + q_{w0} \frac{\omega}{1 - \gamma} (1 - \nu) \frac{\alpha \delta}{\delta + \psi} + q_{wC} (1 - \nu) \frac{\alpha \delta}{\delta + \psi} + q_{wI} \alpha + \rho_1, \\
\lambda_{17} &= q_{w0} \frac{\omega \gamma}{1 - \gamma} (1 - \nu) \frac{\alpha \delta}{\delta + \psi} \frac{\psi + \delta - 1}{1 - \delta}, \\
\lambda_{18} &= q_{wI}, \quad \lambda_{19} = 1, \\
\lambda_{20} &= \alpha_1, \quad \lambda_{21} = \alpha_2, \quad \lambda_{22} = \alpha_3, \\
\lambda_{23} &= \alpha_4, \quad \lambda_{24} = \alpha_5.
\end{aligned}$$

Dynamic equilibrium system

$$\begin{aligned}
\hat{\pi}_{H,t} &= \beta E_t \hat{\pi}_{H,t+1} + \lambda_1 \hat{r}_t^X + \lambda_2 \hat{X}_t^* + \lambda_3 \hat{e}_t, \\
\hat{r}_t^X &= \lambda_4 \hat{w}_t + \lambda_5 \hat{e}_t + \lambda_6 \hat{p}_t^X, \\
\hat{p}_t^X &= \lambda_7 \hat{X}_t^* + \lambda_8 \hat{e}_t + \lambda_9 \hat{e}_{t-1} + \lambda_{10} E_t \hat{r}_{t+1}^X + \lambda_{11} E_t \hat{r}_{t+1}^P \\
&\quad + \lambda_{12} E_t \hat{p}_{t+1}^X + \lambda_{13} E_t \hat{e}_{t+1} + \lambda_{14} \hat{\mu}_t, \\
\hat{w}_t &= \lambda_{15} \hat{X}_t^* + \lambda_{16} \hat{e}_t + \lambda_{17} \hat{e}_{t-1} + \lambda_{18} \hat{p}_t^X + \lambda_{19} \hat{r}_t^P, \\
\hat{X}_{t+1}^* &= (1 - \delta) \hat{X}_t^* + \psi \hat{e}_t, \\
\hat{e}_t &= \lambda_{20} \hat{X}_t^* + \lambda_{21} \hat{e}_{t-1} + \lambda_{22} E_t \hat{e}_{t+1} + \lambda_{23} \hat{r}_t^P + \lambda_{24} \hat{\mu}_t.
\end{aligned}$$

Steady-state block

$$\begin{aligned}
\bar{R}^P &= \frac{1 + (1 - \bar{e}) \phi \Phi}{\bar{e} + (1 - \bar{e}) \phi}, \\
1 &= \beta \Theta \bar{R}^P [\bar{e} + (1 - \bar{e}) \phi] \bar{e}^{\varphi_e}, \\
\bar{I}^X &= \psi \bar{e}^\alpha, \quad \bar{X} = \frac{\bar{I}^X}{\delta}, \\
\bar{Y}_H &= \bar{e}^\alpha + \bar{X}, \quad \bar{C} = \bar{Y}_H^{1-\nu}, \\
1 &= \beta \left(\frac{\bar{r}^X}{\bar{p}^X \bar{R}^P [\bar{e} + (1 - \bar{e}) \phi]} + 1 - \delta \right), \\
\bar{w} (1 - \bar{e}) [(1 - \gamma) \bar{C}]^{-\omega} &= \Omega_w \bar{R}^P, \\
\bar{r}^X &= \frac{\bar{w}}{\alpha \bar{e}^{\alpha-1}} - \psi \bar{p}^X, \\
\bar{Y}_H^{-\nu} &= \frac{\epsilon}{\epsilon - 1} \bar{r}^X.
\end{aligned}$$

Appendix E Wild-Bootstrap Inference

This appendix describes the wild-bootstrap construction used for the recurrence test in Table 1 and forecast inference in Table 4, Appendix C, and Section 5.3. All these exercises generate exchange-rate panels under a driftless random-walk null. Appendix A.1 explains how the simulated panels are filtered to evaluate alignment. Here, we set out the common construction and its forecasting applications. Currency-specific tests compare each observed Clark–West statistic with its own simulated distribution. The joint tests evaluate the number of MSFE improvements and the largest Clark–West statistic across the baseline panel.

E.1 Null Hypothesis and Maintained Forecasting Design

Let $s_{i,t} = \log \mathcal{E}_{i,t}$ denote the log exchange rate of currency i . The null hypothesis imposes unpredictable changes with zero drift:

$$s_{i,t} = s_{i,t-1} + u_{i,t}, \quad \mathbb{E}[u_{i,t} \mid \mathcal{F}_{t-1}] = 0, \quad (\text{E.1})$$

where $\mathcal{F}_{t-1}$ is the information available before date t . The corresponding forecast of the cumulative exchange-rate change is zero at every horizon. In the no-intercept limit-cycle regression, this forecast is nested at a recurrence coefficient of zero.

The baseline forecasting design uses $B = 10,000$ replications and the quarterly sample ending in 2020Q1. Every replication maintains the twenty currency series, observed availability pattern, cycle length $p = 32$, forecast horizons $h \in \{12, 16, 20\}$, and 64-observation rolling estimation window. The forecast dates and completed-outcome restriction also match the baseline. The evaluation contains 70, 62, and 54 forecasts per currency at the three horizons, respectively. Thus, inference evaluates the baseline procedure with its cycle length, window, horizons, and sample endpoint held fixed.

E.2 Generating the Wild-Bootstrap Panels

We first calculate observed quarterly changes and remove each currency’s sample mean. Let $\mathcal{D}_i$ be the dates with available quarterly innovations and define

$$\hat{u}_{i,t} = \Delta s_{i,t} - \overline{\Delta s}_i, \quad \overline{\Delta s}_i = \frac{1}{|\mathcal{D}_i|} \sum_{t \in \mathcal{D}_i} \Delta s_{i,t}. \quad (\text{E.2})$$

The vector $\hat{\mathbf{u}}_t$ contains these demeaned innovations for the currencies observed at date t .

For replication b , draw one Rademacher multiplier per date:

$$\Pr(\eta_t^{(b)} = 1) = \Pr(\eta_t^{(b)} = -1) = \frac{1}{2}, \quad \mathbf{u}_t^{*(b)} = \eta_t^{(b)} \hat{\mathbf{u}}_t. \quad (\text{E.3})$$

The signs are independent across dates and replications. Within a date, the same sign multiplies every available currency innovation. Cumulating the innovations gives the simulated log exchange rates:

$$s_{i,t}^{*(b)} = s_{i,t-1}^{*(b)} + u_{i,t}^{*(b)}. \quad (\text{E.4})$$

An additive constant in the initial log level has no effect on the exercise, because the predictors, targets, and forecasts are expressed as changes. The observed availability pattern determines the usable regression and evaluation observations in each replication.

The common sign preserves the observed contemporaneous relationships. For currencies i and j observed at the same date,

$$|u_{i,t}^{*(b)}| = |\hat{u}_{i,t}|, \quad u_{i,t}^{*(b)} u_{j,t}^{*(b)} = \hat{u}_{i,t} \hat{u}_{j,t}. \quad (\text{E.5})$$

Each replication therefore retains the innovation magnitudes and products at each date, including the common movements across dollar exchange rates and euro-area series. Large and small innovations remain at their original dates. Meanwhile, independence of the signs across dates implies $\mathbb{E}^*[u_{i,t}^{*(b)} \mid \mathcal{F}_{t-1}^*] = 0$, where $\mathbb{E}^*$ conditions on the observed innovation vectors. The reference distribution therefore conditions on those vectors and uses a symmetric sign-randomization design to impose the null.

E.3 Repeating the Rolling Forecast Exercise

The complete forecast is reconstructed from every simulated history. For each currency, horizon, and empirical forecast date t , define

$$y_{i,t,h}^{*(b)} = s_{i,t+h}^{*(b)} - s_{i,t}^{*(b)}, \quad (\text{E.6})$$

$$x_{i,t,h}^{*(b)} = s_{i,t-p+h}^{*(b)} - s_{i,t-p}^{*(b)}. \quad (\text{E.7})$$

Since $p > h$ at all three horizons, the predictor is observable at the forecast date. Let $\mathcal{W}_{i,t,h}$ contain the 64 most recent available regression rows whose outcomes have been completed by t . Each included date τ satisfies $\tau + h \leq t$. The no-intercept rolling estimate

is

$$\hat{\beta}_{i,t,h}^{*(b)} = \frac{\sum_{\tau \in \mathcal{W}_{i,t,h}} x_{i,\tau,h}^{*(b)} y_{i,\tau,h}^{*(b)}}{\sum_{\tau \in \mathcal{W}_{i,t,h}} (x_{i,\tau,h}^{*(b)})^2}. \quad (\text{E.8})$$

The limit-cycle and random-walk forecasts are, respectively,

$$\mathbb{E}_t^{LC} [y_{i,t,h}^{*(b)}] = \hat{\beta}_{i,t,h}^{*(b)} x_{i,t,h}^{*(b)}, \quad \mathbb{E}_t^{RW} [y_{i,t,h}^{*(b)}] = 0. \quad (\text{E.9})$$

Here $\mathbb{E}_t^m[\cdot]$ denotes model m 's forecast computed using the simulated exchange-rate history available at date t in replication b , following the notation in Section 4.1. Re-estimating the coefficient at every forecast date incorporates the sampling variation from the rolling regressions. Constructing all horizons from the same simulated quarterly path also reproduces overlap among forecast targets and estimation windows.

E.4 Forecast Losses and Currency-Specific Inference

Let $\mathcal{T}_{i,h}$ denote the empirical set of forecast dates and $n_{i,h} = |\mathcal{T}_{i,h}|$. For $m \in \{LC, RW\}$, the simulated forecast error is $e_{i,t,h}^{m,*(b)} = y_{i,t,h}^{*(b)} - \mathbb{E}_t^m [y_{i,t,h}^{*(b)}]$. Relative MSFE is

$$R_{i,h}^{*(b)} = \frac{\sum_{t \in \mathcal{T}_{i,h}} (e_{i,t,h}^{LC,*(b)})^2}{\sum_{t \in \mathcal{T}_{i,h}} (e_{i,t,h}^{RW,*(b)})^2}. \quad (\text{E.10})$$

A value below one records an improvement in realized forecast accuracy.

For inference using the Clark–West statistic, define the adjusted loss differential following [Clark and West \(2007\)](#):

$$d_{i,t,h}^{CW,*(b)} = (e_{i,t,h}^{RW,*(b)})^2 - (e_{i,t,h}^{LC,*(b)})^2 + \left(\mathbb{E}_t^{LC} [y_{i,t,h}^{*(b)}] \right)^2. \quad (\text{E.11})$$

The final term adjusts for estimation noise in the larger forecasting model; it is the squared difference between the two forecasts because the random-walk forecast is zero. Averaging over forecast dates gives $\bar{d}_{i,h}^{CW,*(b)}$. Studentization uses a standard error robust to heteroskedasticity and autocorrelation (HAC), with the same Bartlett weights and 32 lags as in the empirical calculation:

$$T_{i,h}^{CW,*(b)} = \frac{\bar{d}_{i,h}^{CW,*(b)}}{\widehat{\text{se}}_{HAC}(\bar{d}_{i,h}^{CW,*(b)})}. \quad (\text{E.12})$$

Larger values favor the limit-cycle benchmark. The standard error is recalculated from each simulated loss sequence.

The two inference columns in Table 4 use the same observed statistic $T_{i,h}^{CW,obs}$. The conventional column uses $1 - \Phi(T_{i,h}^{CW,obs})$, where Φ is the standard normal distribution function. The wild-bootstrap column ranks it against the simulated statistics for the same currency and horizon:

$$\hat{p}_{i,h}^{wild} = \frac{1 + \sum_{b=1}^B \mathbf{1}\{T_{i,h}^{CW,*^{(b)}} \geq T_{i,h}^{CW,obs}\}}{B + 1}. \quad (\text{E.13})$$

Here $\mathbf{1}\{\cdot\}$ is the indicator function. The plus-one correction prevents a reported simulation probability of zero. These currency-specific probabilities evaluate each comparison separately.

E.5 Forecasting Extensions

Tables C.1 and C.2 use the same wild-bootstrap procedure with 10,000 complete panels for each design. The shorter-horizon exercise retains the quarterly sample ending in 2020Q1, the 32-quarter cycle length, and the 64-quarter rolling window, and evaluates $h = 1, 4, 8$ quarters. The monthly exercise uses data from 1973M1–2020M1, a cycle length of 96 months, a rolling window of 192 observations, and horizons of 36, 48, and 60 months.

For each design, innovations are calculated and demeaned using its own frequency and sample. Each replication applies common signs across currencies, reconstructs the predictors, and re-estimates the rolling regression at every forecast date, using only outcomes then available. Forecast dates and evaluation samples match the corresponding empirical exercise. The same Clark–West and HAC calculations are applied to the observed and simulated loss sequences within each design. Each observed statistic is compared with its own design-specific simulated distribution using equation (E.13); the baseline simulated critical values are not reused for the extensions.

E.6 Joint Forecast Statistics

Let $\mathcal{C} = \{(i, h) : i = 1, \dots, 20, h \in \{12, 16, 20\}\}$ denote the set of 60 currency–horizon comparisons. All statistics below are computed from the same complete panel within a replication. This preserves their dependence across currencies and horizons when constructing the joint reference distributions.

Number of MSFE improvements across all currencies. The first statistic is the number of comparisons with lower realized MSFE:

$$N_{\text{all}}^{*(b)} = \sum_{(i,h) \in \mathcal{C}} \mathbf{1}\{R_{i,h}^{*(b)} < 1\}, \quad N_{\text{all}}^{\text{obs}} = 43. \quad (\text{E.14})$$

Each favorable comparison contributes one, regardless of the size of the MSFE reduction. Its null tail counts panels with $N_{\text{all}}^{*(b)} \geq 43$. The count uses realized MSFE reductions; it does not require each comparison to be individually statistically significant.

Number of improvements with one euro-area representative. Define $\mathcal{I}_{\text{one euro}}$ to contain Germany and the ten non-euro currencies: Australia, Canada, Denmark, Japan, Korea, New Zealand, Norway, Sweden, Switzerland, and the United Kingdom. The restricted set $\mathcal{C}_{\text{one euro}} = \mathcal{I}_{\text{one euro}} \times \{12, 16, 20\}$ contains 33 comparisons. The second statistic is

$$N_{\text{one euro}}^{*(b)} = \sum_{(i,h) \in \mathcal{C}_{\text{one euro}}} \mathbf{1}\{R_{i,h}^{*(b)} < 1\}, \quad N_{\text{one euro}}^{\text{obs}} = 22. \quad (\text{E.15})$$

The same restriction is applied to the observed results and every simulated panel. The null tail counts panels with at least 22 gains among these 33 comparisons. The restriction reduces repeated euro-area representation while retaining the simulated dependence among the remaining series.

Maximum Clark–West statistic. The third statistic selects the strongest result in each panel:

$$M^{*(b)} = \max_{(i,h) \in \mathcal{C}} T_{i,h}^{\text{CW},*(b)}, \quad M^{\text{obs}} = \max_{(i,h) \in \mathcal{C}} T_{i,h}^{\text{CW},\text{obs}}. \quad (\text{E.16})$$

The observed maximum is attained by Denmark at $h = 16$ and is displayed as 11.38 in Table 5. Comparisons use the underlying statistic before rounding. Every simulated panel selects its own maximizing currency and horizon, which can differ from Denmark at $h = 16$. The null tail counts panels with $M^{*(b)} \geq M^{\text{obs}}$, accounting for the search over the 60 specified comparisons.

E.7 Tail Frequencies and Interpretation

For any of the three joint statistics $S \in \{N_{\text{all}}, N_{\text{one euro}}, M\}$, define

$$K_S = \sum_{b=1}^B \mathbf{1}\{S^{*(b)} \geq S^{\text{obs}}\}, \quad \hat{q}_S = \frac{K_S}{B}. \quad (\text{E.17})$$

Table 5 expresses the simulated tail frequencies as percentages, $100\hat{q}_S$. The corresponding plus-one Monte Carlo convention is

$$\hat{p}_S = \frac{1 + K_S}{B + 1}. \quad (\text{E.18})$$

At $B = 10,000$, the difference between these conventions is at most $1/(B + 1)$ in probability units and does not change the significance conclusions for the reported joint results. Both use an upper-tail comparison because larger counts and a larger maximum favor the limit-cycle benchmark.

The reported tail frequencies are 2.89 percent for the full-panel count, 1.25 percent for the count with one euro-area representative, and 0.49 percent for the maximum Clark–West statistic. These correspond to probabilities of 0.0289, 0.0125, and 0.0049. The first two statistics reject the specified random-walk null at the five-percent level, and the third rejects at the one-percent level. Their interpretation is the probability of obtaining an equally or more favorable statistic under the simulated null. The count tests ask how often a random-walk panel produces at least as many MSFE improvements as observed. The maximum test asks how often its largest Clark–West statistic reaches the observed maximum, allowing the same search in every simulated panel.