Diagnostic Expectations and the Global Financial Cycle*

Do Lee†

November 2026

Abstract

Equity prices co-move strongly across countries, while cross-border equity flows co-move much less. This paper studies the role of expectations in explaining this feature of the global financial cycle. Using analysts' earnings forecasts for firms in 97 countries over 1995 to 2024, I extract a real-time global belief factor from forecast revisions. Long-horizon forecasts over-react to the global belief factor, which also predicts a boom and bust in the global financial cycle. In a multi-country model where foreign and home investors overweight recent news about global earnings growth, calibrated beliefs reproduce the boom and bust and the gap between price and flow comovement.

JEL Classification: D84, F32, F36, G15, G17, G41

Keywords: Global Financial Cycle, Diagnostic Expectations, Analyst Forecasts, Capital Flows, Equity Comovement

*I am grateful to Julien Acalin, Jaroslav Borovička, Daniel Greenwald, Sydney Ludvigson, Albert Marcet, and participants at seminars and conferences for their valuable comments. All errors remain my own. The views expressed are those of the author and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

†International Monetary Fund, 700 19th Street NW, Washington, DC 20431. Email: dlee2@imf.org.

1 Introduction

Risky-asset prices move together across countries (Rey, 2015; Miranda-Agrippino and Rey, 2020). In my sample the average pairwise correlation of quarterly country equity returns is 0.63, while that of gross portfolio-debt inflows is 0.07 (Cerutti et al., 2019; Avdjiev et al., 2022). If each market's return were a constant multiple of the flow into it plus idiosyncratic noise, no return correlation could exceed the corresponding flow correlation in absolute value, however large the price impact of flows. On the countries and quarters that both series cover, average absolute correlations are 0.65 for returns and 0.13 for flows. A constant price impact of flows cannot generate that gap, which points to a revaluation of existing holdings (Appendix C). This paper asks how much of the revaluation comes from the expectations of the investors who hold the assets.

Most work identifies the global financial cycle from prices and quantities and traces its common component to U.S. monetary policy, to the balance sheets of global intermediaries, and to shifts in risk appetite (Miranda-Agrippino and Rey, 2020; He and Krishnamurthy, 2013; Bruno and Shin, 2015). Cash-flow expectations in these accounts are rational, so common movement in prices that fundamentals do not explain must come from discount rates. A separate literature shows that expectations depart from rational benchmarks in ways that matter for asset prices and credit, and that long-horizon forecasts over-extrapolate recent news (Bordalo et al., 2018, 2019, 2020). That evidence is almost entirely domestic. Whether the common component of the global cycle contains a measurable belief distortion, and whether that distortion can account for the gap between price and flow comovement, are open questions.

I answer them with consensus analyst forecasts of earnings per share from I/B/E/S for firms in 97 countries between 1995 and 2024. Forecasts have two advantages over prices for this purpose. They measure beliefs about cash flows directly, without a model of discount rates. And realized earnings provide a benchmark, so I can test the rationality of the common component rather than assume it. For each country and horizon I average firms' revisions of the forecast for a fixed future date, and I extract a global belief factor from these series with a dynamic factor model in the spirit of Miranda-Agrippino and Rey (2020). I re-estimate the model each year on past

data and use its one-sided filtered estimate, so the factor at each date uses no later information. Local beliefs are each country's revision net of the global factor and of what countries share in the quarter.

The first result is that long-horizon forecasts over-react to the global belief factor. In the test of Coibion and Gorodnichenko (2015), the slope of the forecast error on the global factor is -0.47, -0.65, and -0.68 at three, four, and five years: after a one-unit upward global revision, the five-year realization falls about two thirds of a unit below the forecast. Local news shows the opposite pattern. Analysts under-react to it at one year (slope 0.23), and the local slope is close to zero at four and five years. A single slope on the total revision, the usual object in this literature, is close to zero at one year and averages the two patterns at long horizons. Separating the source of news reveals both.

The second result is that the global belief factor predicts a boom and bust in the global financial cycle. A one-standard-deviation global revision comes with a rise of 0.61 in the global factor of Miranda-Agrippino and Rey (2020), about two thirds of that factor's standard deviation. Half of the rise is gone a year later, and after two years the coefficient is 0.06. The total country revision, which mixes global and local news, also predicts a rise but no complete reversal within three years. The reversal belongs to the global component.

Firm-level data support a belief interpretation of these patterns. Across 30,573 firms, realized three-year earnings growth loads 0.20 on the analysts' long-term growth forecast, far below the approximate benchmark of one for an unbiased forecast, and the most optimistic firms fall furthest short of their forecasts. Within a country and quarter, firms with more optimistic forecasts see lower subsequent price changes. These firm-level patterns are associations that complement the country forecast-error tests, whose predictors use only information available at the forecast date.

I then ask whether the measured belief distortion can account for the price-flow gap. I build a model in which foreign investors, who hold every country's equity, and home investors, who hold only their own, set prices. Investors and analysts filter the trend in global earnings growth, which they believe drifts, and over-weight the latest update to their estimate. The underlying estimate follows a steady-state Kalman filter, which is constant-gain learning, and the over-weighting

extends the diagnostic Kalman filter of Bordalo et al. (2018) and Bordalo et al. (2020). Home investors see world news with noise and absorb more local news than foreign investors, and foreign investors rebalance gradually. The model has a simple separation property. Prices average the valuations of foreign and home investors, while flows follow the difference between them, so a belief that foreign and home investors share moves prices in every market and moves no flows. World news is nearly public, so revisions of global beliefs are largely shared and move prices across countries with little trade. Flows come from gaps between foreign and home valuations, which differ across countries. Investors recognize an earnings cycle gradually, so valuations first follow analysts' optimism and fall a few quarters later.

Chosen to fit forecast, price, and flow moments, the model reproduces the boom and bust, the gap between price and flow comovement, and the pattern of forecast-error slopes across horizons and sources of news, though its five-year global slope is less negative than in the data. Return comovement is 0.57 in the model and 0.63 in the data, and flow comovement is 0.060 and 0.078. When I give investors rational beliefs on the same shocks, flows disappear, return comovement falls to 0.30, and the global factor no longer reverses: two years after a global revision its response is 1.06 times its impact value.

The paper makes three contributions to the literatures discussed in Section 2. Relative to work on the global financial cycle, it measures a belief component of the cycle from forecasts instead of inferring it from prices, and shows that this component predicts the cycle's reversal. Relative to work on expectations, it shows that over-extrapolation attaches to the global component of earnings news while local news is under-reacted to, a pattern that a single slope on total revisions hides. Relative to models of international portfolios, it shows how shared belief revisions separate the comovement of prices from that of flows, so that a large common revaluation coexists with weakly correlated flows.

Section 2 discusses related work and Section 3 sets out what the forecast-error tests identify. Section 4 describes the data and Section 5 constructs global and local beliefs. Section 6 reports the forecast-error tests and Section 7 relates global beliefs to the cycle and to firm outcomes. Section 8 presents the model, and Section 9 concludes.

2 Related Literature

The global financial cycle. Rey (2015) documents that gross capital flows, asset prices, and leverage move together worldwide, and Miranda-Agrippino and Rey (2015, 2020) extract a global factor from risky asset prices and relate it to U.S. monetary policy. Later work links the common variation to the risk-bearing capacity of global intermediaries (He and Krishnamurthy, 2013; Bruno and Shin, 2015), to monetary policy and risk premia (Kekre and Lenel, 2022), and to surges and retrenchments in flows (Forbes and Warnock, 2012; Davis and van Wincoop, 2024), and documents its footprint across capital-account regimes (Acalin and Rebucci, 2020). Other measures of risk appetite come from option prices and credit spreads (Bekaert et al., 2022; Gilchrist and Zakrajšek, 2012). These accounts identify the cycle from prices and quantities, and they attribute its common component to discount rates and the capacity to bear risk. I add a component measured from forecasts, which realized earnings can check. Because the common move is a revaluation, the mechanism sits with revaluation accounts such as Kekre and Lenel (2024) more than with transmission through gross positions. Cerutti et al. (2019) find that push factors explain a small share of the variation in gross flows, and Camous et al. (2025) document and model a related disconnect between the synchronization of asset prices and of output in a rational economy. My object is the gap between returns and gross flows, and my mechanism is a measured belief distortion.

Diagnostic expectations and extrapolation. Bordalo et al. (2018) formalize diagnostic expectations and use them to explain credit cycles. Bordalo et al. (2019) show that analysts' long-term growth forecasts over-react and that stocks with optimistic forecasts earn low returns, and Bordalo et al. (2020) show that individual forecasters over-react to their own news across macroeconomic series while consensus forecasts can under-react. De la O and Myers (2021) find that subjective cash-flow expectations explain much of the variation in price-dividend ratios usually attributed to discount rates, and Bordalo et al. (2024) link belief overreaction to stock market puzzles. In related models, agents who fit simple models to persistent processes extrapolate too far (Fuster et al., 2010), extrapolation from past prices generates momentum and reversal in

house prices (Glaeser and Nathanson, 2017), constant-gain learning about capital gains explains stock price booms (Adam et al., 2017), diagnostic beliefs about fundamentals generate bubbles (Bordalo et al., 2021), and sentiment interacts with intermediaries (Maxted, 2024). Most of this work is domestic and studies a single aggregate or the cross-section of firms within a country. I take its tests to an international panel and separate the global from the local component of news. The separation matters: the two components carry opposite departures from rationality, and only the global component predicts the reversal of the cycle. Beliefs in my model follow a diagnostic Kalman filter, which nests the constant-gain learning of Adam et al. (2017) and reduces to it without the diagnostic distortion. The model adds foreign and home investors who see different news, which generates flows.

Reaction to aggregate news. Agents often under-react to aggregate relative to idiosyncratic information, because information is noisy or attention limited (Lucas, 1972; Maćkowiak and Wiederholt, 2009; Coibion and Gorodnichenko, 2015; Angeletos et al., 2021; Kohlhas and Walther, 2021; Bouchaud et al., 2019). The international panel shows that the direction depends on the source of news and the horizon: analysts under-react to local news at short horizons and over-react to global news at long horizons. The model reconciles the two with one information structure. Analysts absorb announced local news only in part, which produces under-reaction, while global news enters a learned trend that analysts extrapolate into every future year, which produces over-reaction at long horizons.

Open economies with non-rational expectations. Na and Xie (2025) and Na and Yoo (2025) build departures from rational expectations into otherwise standard open economies to explain exchange-rate puzzles and external-balance dynamics. Demand-based models of international capital markets (Gabaix and Maggiori, 2015; Itskhoki and Mukhin, 2021) supply the portfolio structure my model shares, and Gabaix and Koijen (2021) show how inelastic demand amplifies the price impact of flows. I measure the distortion from forecasts, focus on its cross-country component, and study the joint behavior of prices and flows across many countries. The price-flow gap I document is a fact about comovement, and a constant price impact of flows,

however large, cannot generate it.

3 What the Forecast-Error Tests Identify

The test of Coibion and Gorodnichenko (2015) regresses the forecast error on the forecast revision. A rational forecast has an error that is uncorrelated with any information available when it was made, including its own revision, so the slope is zero. Departures from rationality produce slopes of either sign. Three simple cases fix ideas and connect the tests to the model of Section 8. Throughout, $F_t x_{t+h}$ is the forecast at t of a variable x at $t+h$, the error is $x_{t+h} - F_t x_{t+h}$, and the revision is $F_t x_{t+h} - F_{t-1} x_{t+h}$.

Partial absorption of announced news. Suppose $x_{t+h} = n_t + \epsilon_{t+h}$, where n_t is news announced at t and ϵ_{t+h} is unforecastable, and suppose forecasters absorb a share K of the news, $F_t x_{t+h} = K n_t$. The revision is $K n_t$ and the error is $(1-K)n_t + \epsilon_{t+h}$, so the slope is

$$\beta = \frac{1-K}{K} > 0 \quad \text{for } K < 1. \tag{1}$$

Partial absorption produces under-reaction: errors follow revisions in the same direction. Noisy information and inattention generate the same sign (Coibion and Gorodnichenko, 2015).

Diagnostic expectations about a persistent state. Suppose forecasters observe $x_t = \rho x_{t-1} + u_t$ and hold the diagnostic expectation $F_t x_{t+h} = \rho^h (x_t + \theta u_t)$, which over-weights the latest news by $\theta > 0$ (Bordalo et al., 2018). The revision is $\rho^h[(1+\theta)u_t - \rho\theta u_{t-1}]$ and the error is $\sum_{j=1}^{h} \rho^{h-j} u_{t+j} - \rho^h \theta u_t$. With shocks of constant variance,

$$\beta = -\frac{\theta(1+\theta)}{(1+\theta)^2 + \rho^2\theta^2} < 0 \tag{2}$$

at every horizon. Over-weighting produces over-reaction: errors follow revisions in the opposite direction.

Learning a trend that does not exist. Suppose growth x_t is serially uncorrelated with mean zero and variance σ^2, but forecasters learn its mean with a constant gain, $b_t = b_{t-1} + \kappa(x_t -$

b_{t-1}), the steady-state Kalman filter for a mean believed to follow a random walk, and forecast cumulative growth from t through $t+H$. The forecast made at $t-1$ is $(H+1)b_{t-1}$ and the forecast made at t is $x_t + Hb_t$, so the revision of this fixed target is $(1+H\kappa)(x_t - b_{t-1})$ and the error is $\sum_{j=1}^{H} x_{t+j} - Hb_t$. Using $\text{Var}(b) = \kappa\sigma^2/(2-\kappa)$,

$$\beta = -\frac{H\kappa}{2(1+H\kappa)}. \tag{3}$$

Constant-gain learning reads transitory news as a change in trend and extrapolates it into every future period, so it produces over-reaction even without a diagnostic distortion. The over-reaction grows with the horizon, because the extrapolated trend accounts for a larger share of a longer forecast, and approaches $-1/2$ as $H\kappa$ grows.

These cases make two points for what follows. First, the sign of the slope distinguishes under-reaction from over-reaction, and different information structures for different kinds of news can produce opposite signs in the same forecasts. Second, over-reaction that comes from a perceived trend grows with the horizon, as equation (3) shows. Earnings already reported and announced news anchor a short-horizon forecast, so trend extrapolation is a small part of it; a long-horizon forecast is mostly trend. The model of Section 8 combines partial absorption of local news with a diagnostic filter of a global trend, and so predicts under-reaction to local news at short horizons and over-reaction to global news at long horizons.

4 Data

Forecasts. I use I/B/E/S median consensus forecasts of earnings per share at fiscal-year horizons one to four, the quarterly forecasts, and the long-term growth forecast, for firms in 97 countries from 1995 to 2024. For each firm and quarter I take the last consensus of the quarter and build log forecast earnings for future dates by interpolating between fiscal-year forecasts. Beyond the last fiscal-year forecast I extend the path with the analysts' own long-term growth forecast. Writing $F_{i,t}(\tau)$ for the forecast at t of log earnings at date τ, the revision compares two forecasts of the same date made four quarters apart,

$$\Delta F_{i,t}(t+h) = F_{i,t}(t+h) - F_{i,t-4}(t+h), \qquad h = 1, \ldots, 5 \text{ years}.$$

Comparing forecasts of the same date removes the mechanical change in a rolling-horizon forecast as its base moves. I keep a revision only if analysts made the forecasts at both dates, either between fiscal-year forecasts or with their long-term growth forecast, and only if the firm had announced the base earnings, those of its last fiscal year, at both dates. The second condition ensures that I never score a forecast relative to earnings that were not yet public. Revisions enter after clipping at percentiles computed from past data only. Appendix A gives the construction in detail.

Realized earnings. Realized log earnings at τ come from I/B/E/S actual earnings per share, which share the street definition of the forecasts, interpolated between the annual actuals that bracket τ. The forecast error $e_{i,t}(\tau)$ is realized minus forecast log earnings for date τ. Because forecast and realization refer to the same date and the same accounting definition, the error does not mix realized growth into the forecast or compare different earnings concepts.

Coverage. The panel has 35,879 firms with at least one revision, from 95 countries at the one-year horizon and 73 at the five-year horizon after 1995. Coverage falls with the horizon because fewer firms have forecasts that reach five years. The median country-quarter averages over 21 to 32 firms, depending on the horizon, while 14 percent of country-horizon-quarter cells contain one or two firms. Forecast errors are negative on average and grow with the horizon, from -14 log points at one year to -36 at five: analysts are optimistic on average, and more so at long horizons. The tests below include country fixed effects, so they ask whether revisions predict errors relative to each country's average bias.

Prices and flows. The global financial cycle factor is that of Miranda-Agrippino and Rey (2020), 2024 vintage, the quarterly mean of their monthly factor. Country equity returns are U.S. dollar total returns of iShares MSCI country funds for 42 countries from 1995. The motivating price-flow comparison uses gross portfolio-debt inflows from Avdjiev et al. (2022), the quarterly gross flow series with the widest coverage. The model comparison uses portfolio-equity inflows from IMF balance-of-payments statistics, as a percent of GDP, for 43 countries. I weight countries

by their previous-year share of equity market capitalization; the United States carries 41 percent of the weight on average.

Firm prices. The firm-level tests link I/B/E/S to split-adjusted monthly prices from Compustat Global and, for U.S. and Canadian firms, Compustat North America. Compustat Global prices start in 2007, so before 2007 the firm sample is North American.

5 Global and Local Beliefs

5.1 A real-time factor model

Let $x_{c,h,t}$ be the mean revision $\Delta F_{i,t}(t+h)$ of the firms in country c, and $y_{c,h,t} = (x_{c,h,t} - \mu_{c,h})/s_{c,h}$ its standardized value, using the mean and standard deviation through the previous year. Following Miranda-Agrippino and Rey (2020), each series loads on a global and a regional factor,

$$y_{c,h,t} = \lambda_{c,h}\,\xi_t + \lambda^r_{c,h}\,\xi^{r(c)}_t + u_{c,h,t}, \qquad \xi_t = \Phi\,\xi_{t-1} + \eta_t, \qquad \mathrm{Var}(u_{c,h,t}) = \frac{\sigma^2_{c,h}}{n_{c,h,t}}, \tag{4}$$

where $\xi^{r(c)}$ is the factor of country c's region (six regions) and the factors follow independent autoregressions. The panel has 437 country-horizon series. The idiosyncratic variance falls with the number of firms $n_{c,h,t}$ in the cell: a country mean over many firms measures the country's revision precisely, while a mean over one or two firms is mostly firm noise. This weighting lets the panel include small markets without letting them drive the factor. I estimate the model (Stock and Watson, 2016) by quasi-maximum likelihood with the EM algorithm and the Kalman filter and smoother, allowing for missing observations, and re-estimate it from scratch each year on data through the previous year (Appendix B).

The *global belief factor* is the filtered estimate $G_t = \xi_{t|t}$. It uses revisions through quarter t and parameters estimated on earlier years only. The tests ask whether beliefs at t predict later outcomes, and a smoothed factor or full-sample parameters would use information from after t. The real-time factor and the full-sample smoothed factor have a correlation of 0.84 from 1995, so the real-time constraint costs some precision but leaves the factor recognizable.

Local beliefs are the part of each country's five-year revision that is neither global nor common to countries that quarter,

$$L_{c,t} = x_{c,t} - \mu_c - s_c \lambda_c \, \xi_{t|t} - \bar{u}_t, \qquad \bar{u}_t = \frac{\sum_c w_{c,t} \left(x_{c,t} - \mu_c - s_c \lambda_c \, \xi_{t|t} \right)}{\sum_c w_{c,t}}, \tag{5}$$

with $w_{c,t}$ the market-capitalization weights, so $\sum_c w_{c,t} L_{c,t} = 0$ each quarter. Removing $\bar{u}_t$ keeps common news that the single factor misses out of local beliefs. The regional factor stays in local beliefs, so local beliefs include news shared by a region but not by the world.

5.2 Properties of the factor

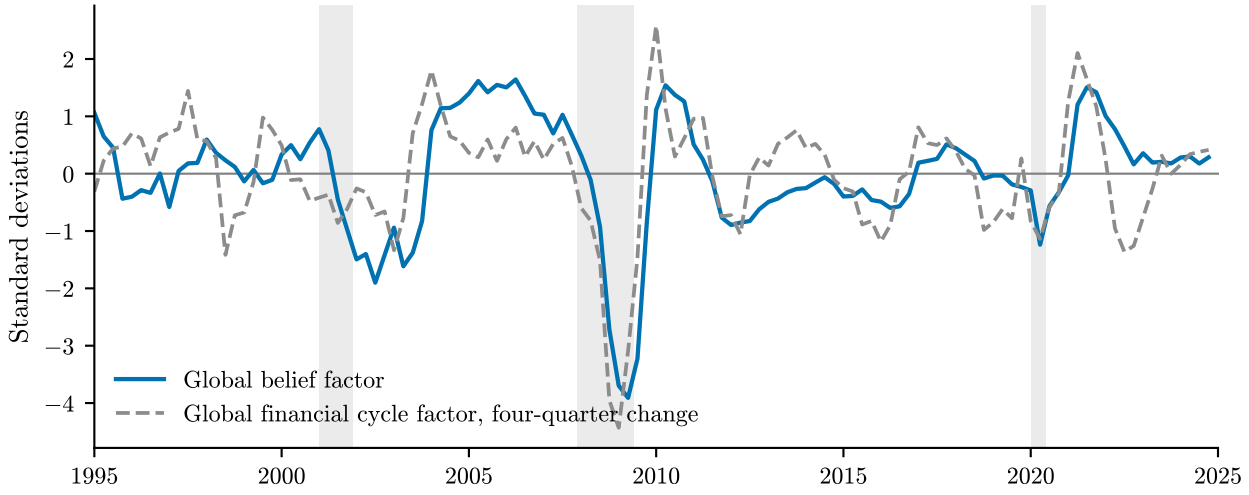


Figure 1: The global belief factor and the global financial cycle

Notes: Real-time filtered global belief factor and four-quarter change in the global financial cycle factor of Miranda-Agrippino and Rey (2020), each standardized. Shaded: 2001, 2007Q4–2009Q2, and 2020Q1–Q2.

Figure 1 plots the global belief factor against the four-quarter change in the global financial cycle factor. Beliefs fall sharply in the three global downturns of the sample and recover after them, and the two series move together, with a correlation of 0.64. The correlation of the belief factor with the VIX is −0.51. The factor is persistent, with a first-order autocorrelation of 0.87. Over the full sample the global factor accounts for a median of 29 percent of the variance of a country-horizon series and the regional factors for 12 percent. The global share is largest where the markets are largest: it is 94 percent for the North American series, 40 percent in Europe, 28 percent in Asia-Pacific, 26 percent in Latin America, 22 percent in the Middle East and Africa, and 10 percent for offshore centres. Asia-Pacific is the one region where the regional factor

matters as much as the global one, with a median share of 28 percent. Analysts' revisions are thus far from purely local even in markets that are not closely integrated, and they are dominated by a common component in the largest markets.

Both series respond to the same news, so their comovement is no evidence of a distortion. Two features of the factor are more telling. The global component of revisions has a correlation of -0.26 with its own value two years later, so global revisions tend to be partly undone. And as the next section shows, the factor predicts forecast errors, which a rational forecast rules out.

6 Forecasts Over-React to the Global Belief Factor

6.1 Specification

I regress the country-mean forecast error on global and local beliefs,

$$e_{c,t}(t+h) = \alpha_c + \beta_h^G \, G_t + \beta_h^L \, L_{c,t} + \varepsilon_{c,t}, \tag{6}$$

with G_t in revision units ($\xi_{t|t}$ times the slope of the weighted world mean revision on it), so that β_h^G and β_h^L are comparable with the slopes in Section 3. Under rational expectations both slopes are zero at every horizon, whatever the information structure, because G_t and $L_{c,t}$ are built from information available at t. I include country fixed effects, weight by market capitalization, and use Driscoll and Kraay (1998) standard errors with $4h$ lags. These allow for the overlap of forecast horizons, since forecasts made in consecutive quarters for horizons of h years share most of their realization, and for correlation of errors across countries in the same quarter, which is large because realized global growth is common. The sample starts in 1995Q1.

6.2 Results

Figure 2 and Table 1 report the slopes. The global slope is insignificant at one and two years (-0.04 and -0.32) and negative from three years on: -0.47, -0.65, and -0.68 at three to five years. After a one-unit upward global revision of the five-year forecast, the five-year realization falls about two thirds of a unit below it. The global slope becomes more negative with the horizon, as trend extrapolation predicts in equation (3).

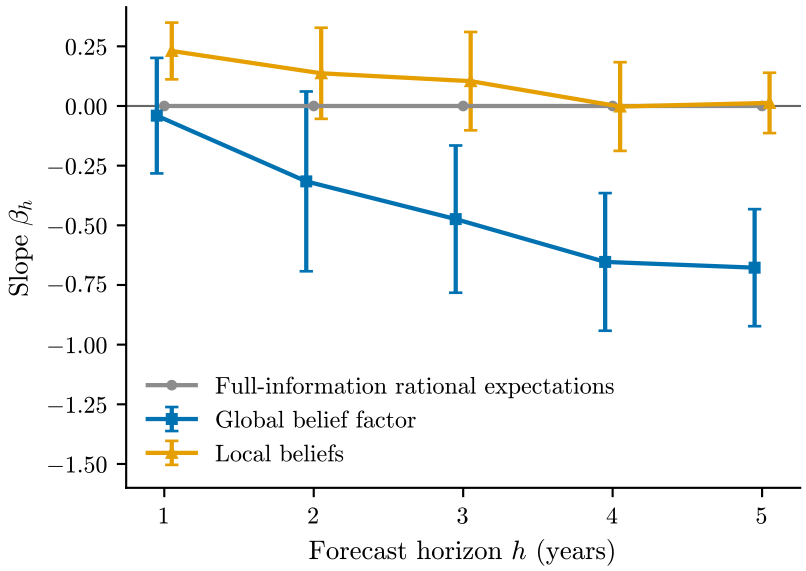


Figure 2: Forecast errors on global and local beliefs, by horizon

Notes: Slopes of equation (6) with 95 percent confidence intervals. Country fixed effects; market-capitalization weights; Driscoll-Kraay standard errors with $4h$ lags; 1995Q1–2024Q4. Full-information rational expectations implies a zero slope at every horizon.

Table 1: Forecast errors on global and local beliefs

Horizon h (years)	1	2	3	4	5
Global belief factor β_h^G	−0.04	−0.32	−0.47	−0.65	−0.68
	(−0.3)	(−1.6)	(−3.0)	(−4.4)	(−5.4)
Local beliefs β_h^L	0.23	0.14	0.10	0.00	0.01
	(3.8)	(1.4)	(1.0)	(0.0)	(0.2)
Total country revision	0.03	−0.18	−0.22	−0.31	−0.34
	(0.3)	(−1.5)	(−2.4)	(−4.2)	(−5.6)
Country-quarters	7,094	6,036	5,177	4,823	4,573

Notes: Equation (6); the last block uses the total country revision alone. Country fixed effects; market-capitalization weights; Driscoll-Kraay t-statistics with $4h$ lags in parentheses; 1995Q1–2024Q4.

Local news shows the opposite pattern. The local slope is 0.23 at one year: after a local revision, realized earnings exceed the revised forecast, so the revision was too small. Read through equation (1), a slope of 0.23 corresponds to analysts absorbing about four fifths of local news at the one-year horizon. The local slope falls to 0.14 and 0.10 at two and three years and to 0.00 and 0.01 at four and five years, as local news that has already entered realized earnings stops being forecastable.

The two components behave differently enough that a single slope on the total revision is misleading. That slope is close to zero at one year, where the positive local slope and the insignificant global slope offset each other, and −0.22 to −0.34 from three years on, about half

the global slope. Pooling hides the short-horizon under-reaction and understates the long-horizon over-reaction. The horizon pattern is the one Section 3 predicts for trend extrapolation of global news combined with partial absorption of local news. Reported earnings and announced news anchor short-horizon forecasts, while long-horizon forecasts are mostly a perceived trend, and a trend revision that the data do not bear out produces a long-horizon error of the opposite sign.

The over-reaction does not come from the United States alone, although the United States carries 41 percent of the weight. Excluding it, the global slopes at three to five years are -0.37, -0.53, and -0.55. It is, however, a property of the forecasts for the large markets that account for most of world equity. With equal country weights, which give a small economy covered by a few firms the weight of a large one, the global slopes are -0.33 to -0.37 and not significant.

7 The Global Belief Factor and the Global Financial Cycle

7.1 Boom and bust

I regress the subsequent path of the global financial cycle factor m_t on global and local beliefs,

$$m_{t+k} - m_{t-4} = \alpha_c + \beta_k^G G_t + \beta_k^L L_{c,t} + \varepsilon_{c,t}, \tag{7}$$

with standardized five-year beliefs, the fixed effects and weights of equation (6), and Driscoll-Kraay standard errors with $k+4$ lags. I measure the outcome from the start of the four-quarter revision window, so the coefficient at $k = 0$ includes the move of the factor during the revision, and coefficients for $k > 0$ trace its path afterwards. The outcome is common to all countries, so variation over time identifies the global coefficient, and the panel structure weights the global factor by the countries whose revisions define it.

A one-standard-deviation upward global revision comes with a rise of 0.61 in the factor, about two thirds of the factor's standard deviation of 0.98 (Figure 3 and Table 2). The rise starts before the revision quarter, 0.42 two quarters earlier, because the four-quarter revision window and the outcome window overlap: beliefs and prices move together during the revision. The coefficient then declines gradually. It is 0.48 two quarters later, 0.32 after a year, 0.14 after six quarters, and 0.06 after two years, and it stays near zero through three years. The point estimate returns

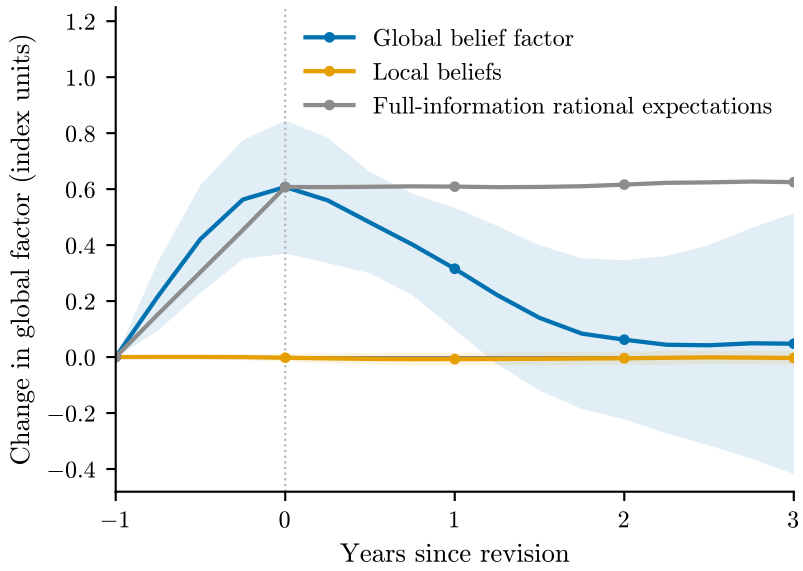


Figure 3: The global financial cycle factor after global and local revisions

Notes: Coefficients of equation (7), the change in the global financial cycle factor (2024 vintage, index units) from $t-4$ to $t+k$, with 95 percent confidence intervals. Country fixed effects; market-capitalization weights; Driscoll-Kraay standard errors with $k+4$ lags; 1995–2024. Full-information rational expectations: a full-information version of the model of Section 8, with an announced share of local news of 0.8 and stochastic volatility, scaled to the data at $k=0$.

Table 2: The global financial cycle factor after global and local revisions

Quarters k	-2	0	2	4	6	8	10	12
Global belief factor	0.42	0.61	0.48	0.32	0.14	0.06	0.04	0.05
	(4.3)	(5.0)	(5.2)	(2.8)	(1.1)	(0.4)	(0.2)	(0.2)
Local beliefs	0.00	0.00	−0.01	−0.01	−0.01	0.00	0.00	0.00
	(0.1)	(−0.3)	(−0.6)	(−0.6)	(−0.5)	(−0.3)	(−0.1)	(−0.3)
Total country revision	0.66	1.04	0.86	0.62	0.37	0.27	0.32	0.48
	(4.6)	(5.1)	(5.3)	(3.4)	(2.1)	(1.5)	(1.5)	(2.0)
Country-quarters	5,724	5,724	5,627	5,533	5,437	5,341	5,244	5,150

Notes: Equation (7); the last block uses the standardized total country revision alone. Country fixed effects; market-capitalization weights; Driscoll-Kraay t-statistics with $k+4$ lags in parentheses.

close to zero within two years, though the two-year coefficient is imprecise, with a 95 percent interval of −0.22 to 0.35.

Local beliefs have coefficients near zero at every horizon. This is largely mechanical: local beliefs sum to zero across countries each quarter while the outcome is common, so they cannot predict it. The informative comparison is with the total country revision, which mixes global and local news. Its coefficient peaks at 1.04, falls to 0.27 after two years, and is 0.48 after three, so it predicts a rise but no complete reversal. Only the global component of revisions predicts the bust.

Country equity returns tell the same story with less precision. The cumulative world return from $t-3$ to $t+k$ on the four-quarter change of the common five-year forecast, normalized to one at $k=0$, is 0.42 after one year and 0.37 after two, with standard errors of 0.22 and 0.26. The return reversal is smaller and less precisely estimated than the reversal of the price factor, and it depends on the regressor: with the change in the rolling five-year forecast instead, the values are 0.26 and 0.16. The price factor, which extracts the common component of many risky asset prices, measures the cycle with less noise than an average of country equity returns.

The boom and bust has two readings. A common discount rate that falls with good news and then mean-reverts would raise prices and later lower them, with forecasts of cash flows unbiased. A belief distortion would raise prices because investors expect more growth than the data deliver, and prices would fall as the shortfall becomes evident. The forecast errors of Section 6 separate the two: after upward global revisions realized earnings fall short of forecasts, which a rational forecast rules out. The evidence here is predictive. It shows that the measured belief component forecasts the cycle's reversal, not that beliefs cause it; the model of Section 8 provides a structure in which they do.

7.2 Realized growth falls short of optimistic forecasts

The analysts' long-term growth forecast is their forecast of annual earnings growth over the next three to five years, and it is the forecast in which Bordalo et al. (2019) document over-reaction for U.S. firms. Across 30,573 firms from 1981 to 2024, realized three-year growth has a slope of 0.20 on the forecast and realized five-year growth a slope of 0.16, against one for an unbiased forecast. One is an approximate benchmark here, because realized growth is log growth while the forecast is in percent, and forecasts are matched to fiscal years by calendar year rather than by exact date, without the announced-base requirement of the country tests. The mean forecast is 15.0 percent a year and mean realized growth about 5 percent. Firms in the most optimistic quintile fall 24 percentage points a year short of their forecast, while those in the least optimistic quintile slightly exceed theirs.

Splitting the forecast into a global component (the mean across all firms with a forecast in

the year), a country component, and a firm component, with standard errors clustered by firm and year, the firm component predicts realized three-year growth with a loading of 0.20 and the country component with a loading of 0.27. Firm-specific and country-specific optimism carry information about realized growth, though far less than one for one.

7.3 Optimistic firms have lower subsequent price changes

If prices reflect these forecasts, firms with optimistic forecasts should see lower subsequent price changes as reality disappoints the optimism. Within a country and quarter, a one-standard-deviation more optimistic long-term growth forecast predicts a log price change lower by 3.4, 6.3, and 8.5 percent over one, two, and three years (standard errors clustered by firm and quarter). Country-by-quarter fixed effects absorb all common shocks, including the global cycle, so the estimate comes from differences across firms in the same market at the same time. Controlling for size moves the two-year estimate from -6.2 to -5.8 percent on the firms with market capitalization. Outside the United States the estimate is -1.3 percent, smaller but of the same sign, in a sample that starts only in 2007.

The global part of firms' optimism, identified over time without country-quarter fixed effects, predicts two-year price changes lower by 11.9 percent per standard deviation, and the country part predicts nothing. This is the firm-level counterpart of the boom and bust: the component of optimism that is common across countries predicts the largest subsequent fall. Differences in risk across firms could explain these price results. The evidence that the optimism is a distortion comes from the country forecast-error tests of Section 6, and the firm shortfall of realized earnings is consistent with it.

8 A Model of the Global Financial Cycle under Diagnostic Expectations

The evidence shows a belief distortion in the global component of forecasts and a reversal of the global cycle that the distortion predicts. This section asks whether the distortion can account quantitatively for the reversal and for the gap between price and flow comovement. The model has to produce three things: over-reaction to global news and under-reaction to local news in

forecasts, a boom and bust in prices after global revisions, and strong return comovement with weak, persistent flows that move with returns.

8.1 Environment

Earnings. Time is quarterly and there are $N = 42$ countries. Log earnings growth in country c is $x_{c,t} = g_t + \ell_{c,t}$. Global growth has a persistent part, an earnings cycle, and transitory noise,

$$g_t = z_t + z_t^C + \varepsilon_t^g, \qquad z_t = \rho z_{t-1} + u_t, \qquad z_t^C = \rho_C z_{t-1}^C - e_C\, C_{t-1} + u_t^C, \qquad C_t = C_{t-1} + z_t^C. \tag{8}$$

Cycle growth z^C error-corrects toward the level C, so earnings return to trend after a cycle: a period of high growth is followed by a period of low growth. Local growth has a persistent part, transitory noise, and news $n_{c,t}$ that firms announce four quarters before it enters growth. The true trend of growth is zero. Appendix D gives the full specification.

Investors and information. Foreign investors, with wealth share ω, hold every country's equity; home investors hold their own country's. Foreign investors observe every country's growth, hence world growth $\bar{x}_t = N^{-1}\sum_c x_{c,t}$. Home investors observe their own country's growth and a noisy signal of world growth,

$$s_{c,t} = \bar{x}_t + n_H \sigma_g \big(\sqrt{1 - c_H}\, \nu_{c,t} + \sqrt{c_H}\, \nu_t\big), \tag{9}$$

and split their own growth into global and local parts by signal extraction (Lucas, 1972). A share c_H of the noise variance is common to all countries, and its common part ν_t is persistent (Lorenzoni, 2009). Home investors absorb a share K_N of announced local news and foreign investors a smaller share K_N^G, a local information advantage in line with Bae et al. (2008) and Van Nieuwerburgh and Veldkamp (2009). Analysts publish a consensus that mixes home and foreign analysts with the measured local share of 0.88 and includes stale forecasts. The consensus does not enter prices; it is the model's counterpart of the forecasts in the data.

Diagnostic beliefs about the global trend. Investors and analysts believe that the trend in global growth follows a random walk, although it is truly zero, and filter it from growth. In

steady state the Kalman filter updates the estimate with a constant gain κ, which rises with the perceived variance of trend shocks relative to transitory growth. Foreign investors filter world growth, $b^G_t = b^G_{t-1} + \kappa(\bar{x}_t - b^G_{t-1})$, and home investors apply the same gain to their estimate of global growth, a simplification. Each estimate over-weights its departure from a slowly moving reference $\underline{b}_t$, a memory of past estimates,

$$\tilde{b}^j_t = b^j_t + \theta\left(b^j_t - \underline{b}^j_t\right), \qquad \underline{b}^j_t = \underline{b}^j_{t-1} + \underline{\kappa}\left(b^j_{t-1} - \underline{b}^j_{t-1}\right), \qquad j = G, H, \tag{10}$$

which extends the diagnostic Kalman filter of Bordalo et al. (2018) and Bordalo et al. (2020). With $\underline{\kappa} = 1$ the reference is the previous estimate, as in their filter; a slower reference lets the distortion build and fade over several quarters. With $\theta = 0$ beliefs reduce to constant-gain learning. Learning updates b, not $\tilde{b}$, so the distortion does not compound over time, and it fades as the reference catches up with the belief. Investors, but not analysts, also recognize the earnings cycle gradually, $\hat{C}_t = \hat{C}_{t-1} + \kappa_C(C_t - \hat{C}_{t-1})$, and expect earnings to return to trend.

Valuations, prices, and flows. Foreign ($j = G$) and home ($j = H$) investors value country c's equity, relative to its earnings, at

$$V^j_{c,t} = \chi\, \tilde{b}^j_{c,t} + K^j_N\, N_{c,t} - d_I\, \hat{C}_t, \tag{11}$$

the perceived trend times a loading χ, plus absorbed pending news $N_{c,t} = \sum_{j=0}^{3} n_{c,t-j}$, minus the recognized cycle level. These are reduced-form valuation loadings: a higher perceived trend raises expected growth in every future period, pending news raises it over the next year, and a recognized high cycle level implies lower growth ahead. Home investors hold $q^H_{c,t} = V^H_{c,t} - p_{c,t}$. Foreign investors move a share ψ of the way to their target each quarter, $q^G_{c,t} = (1-\psi)q^G_{c,t-1} + \psi(V^G_{c,t} - p_{c,t})$, which captures the slow rebalancing of international portfolios. Fixed supply clears each market, $\omega q^G_{c,t} + (1-\omega)q^H_{c,t} = \Pi$, which gives the price and gross foreign inflows,

$$p_{c,t} = V^H_{c,t} + \omega\, \mathcal{D}_{c,t} - \Pi, \qquad f_{c,t} = (1-\omega)\,\Delta\mathcal{D}_{c,t}, \qquad \mathcal{D}_{c,t} = \delta\, \mathcal{D}_{c,t-1} + (1-\delta)\left(V^G_{c,t} - V^H_{c,t}\right), \tag{12}$$

with $\delta = (1-\psi)(1-\omega)/(1-\omega+\omega\psi)$. Country returns are $r_{c,t} = p_{c,t} - p_{c,t-1} + x_{c,t}$, and the model's global financial cycle factor is the cumulative world-mean return, as the published factor cumulates a return factor.

Equilibrium. Beliefs at t do not depend on $p_{c,t}$, so each market clears at a unique price and I solve the model forward. Beliefs need not be consistent with the equilibrium price process, so the equilibrium is a temporary equilibrium, defined in Appendix D. I compare it with two benchmarks on the same shocks. Under full-information rational expectations (FIRE), investors know the state and value each market at expected future growth. With rational investor beliefs, investors hold the FIRE earnings beliefs and the rest of the model stays the same.

8.2 How beliefs move prices and flows

Equation (12) implies a separation between prices and flows that drives the model's comovement results.

Proposition 1. *Write each valuation as $V^j_{c,t} = A_t + \Delta^j_{c,t}$, where A_t is a component common to foreign and home investors and to all countries. Then A_t moves every price one for one, $\partial p_{c,t}/\partial A_t = 1$, and does not enter flows: $f_{c,t}$ depends only on the history of $V^G_{c,t} - V^H_{c,t}$. If foreign and home investors hold the same valuations at all dates and $\mathcal{D}_{c,0} = 0$, flows are zero.*

Appendix D gives the proof. The proposition explains why shared belief revisions produce strong price comovement with weak flow comovement. To the extent that foreign and home investors revise their global beliefs together, the revision raises prices in every market and leaves flows unchanged. Flows come only from the valuation gap,

$$V^G_{c,t} - V^H_{c,t} = \underbrace{\chi\big[(1+\theta)(b^G_t - b^H_{c,t}) - \theta(\underline{b}^G_t - \underline{b}^H_{c,t})\big]}_{\text{world-news gap}} + \underbrace{(K^G_N - K_N)\, N_{c,t}}_{\text{local-news gap}}. \tag{13}$$

The world-news gap arises because home investors see world news with noise. After good world news, foreign valuations rise more than home valuations, and foreign investors buy in every market as prices rise, so flows move with returns. The diagnostic term amplifies the difference in current beliefs by $1+\theta$. The local-news gap arises because foreign investors absorb less local news, and it pushes foreign and home valuations apart country by country, which keeps flow comovement low. The common part of home investors' noise pushes valuations apart in the same direction everywhere and is a source of common flows. Gradual rebalancing spreads each

change in the gap over several quarters, and the persistence of the common noise makes the common part of the gap move slowly, so flows are persistent.

The forecast patterns follow Section 3. Constant-gain learning reads part of each transitory global surprise as a change in trend and extrapolates it into every future year, and the diagnostic term adds to the revision, so long-horizon forecasts over-react to global news. Announced local news enters forecasts only in part, so further revisions in the same direction follow local revisions. The boom and bust comes from the interaction of the trend belief and the cycle. After good global news, valuations follow the optimistic trend belief. As investors recognize the cycle and expect earnings to return to trend, and as the diagnostic distortion fades with the reference, valuations fall.

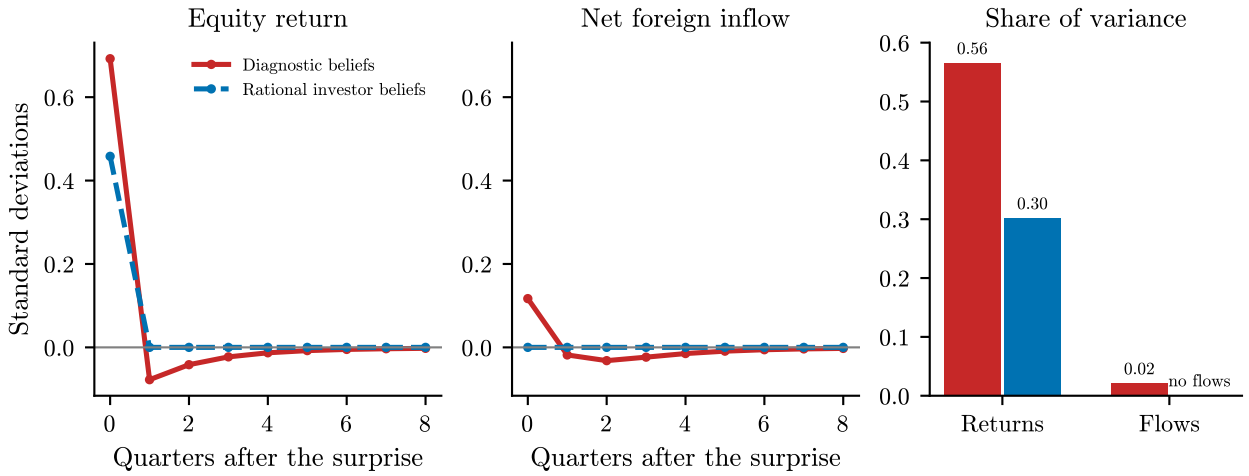


Figure 4: Model responses to a transitory global earnings surprise

Notes: Exact impulse responses: the model is linear in the shocks, and responses are differences between paired simulations that differ only by the impulse. Left and middle panels: response to a one-standard-deviation transitory surprise in global earnings growth, which moves every country alike: quarterly equity return and net foreign inflow, each in units of the country's unconditional standard deviation and averaged across countries. Right panel: share of the variance of a country's quarterly return and net foreign inflow due to global earnings shocks (transitory, persistent, and cycle), averaged across countries; with rational investor beliefs there are no flows. Under rational investor beliefs the cycle shock does not move prices, because investors lower their valuation by exactly the deviation of earnings from trend that the cycle creates.

Figure 4 shows why prices comove while flows do not. A one-standard-deviation transitory global earnings surprise raises returns by an average of 0.69 standard deviations across countries on impact, against 0.46 with rational investor beliefs, and about a fifth of the impact is given back within a year as the trend belief that the surprise created fades. Foreign investors see world news without noise and home investors with noise, so the surprise opens a small valuation gap: net foreign inflows rise by 0.12 standard deviations and turn slightly negative as the gap closes.

Global earnings shocks account for 56 percent of the variance of a country's return and 2 percent of the variance of its flows. With rational investor beliefs they account for 30 percent of return variance, and there are no flows. Flows come mostly from local news, which is independent across countries.

8.3 Parameters

Table 3: Parameters

Symbol	Meaning	Value
Fixed		
N, σ_g^2	countries; global variance parameter	42; 0.355
ρ, s_z	persistence and variance share of persistent news	0.8; 0.035
χ, ω, Π	valuation loading; foreign investors' share; premium	1.0; 0.3; 0.1
s_N	announced share of local variance	1.0
–	analysts: local share; stale share; delayed share	0.88; 0.26; 0.5
Chosen to fit the targets		
κ; θ, $\underline{\kappa}$	global trend gain; diagnostic strength, reference gain	0.20; 0.87, 0.50
s_C, ρ_C, e_C	earnings cycle: variance share, persistence, error correction	0.25; 0.92; 0.09
κ_C, d_I	cycle-recognition gain; weight on the recognized cycle	0.42; 1.69
n_H, c_H, ρ_ν	home noise on world news; common share; its persistence	0.26; 0.12; 0.50
K_N, K_N^G	absorbed announced news: home investors and analysts; foreign investors	0.30; 0.23
ψ	foreign investors' rebalancing speed	0.65

I fix the parameters in the upper block of Table 3, including the analysts' local and stale shares, which I measure, and their delayed share, which I set to one half. I choose the others by a search that minimizes the sum of squared standardized deviations from 25 targets (the ten forecast-error slopes of Figure 2, the response of the global financial cycle factor in Figure 3, the stress ratio of flow comovement, cross-country earnings comovement, and two earnings autocorrelations), plus terms for the response of cumulative world returns, flow comovement, the flow-return correlations, and the factor's path in the third year. The search keeps only versions with the targets, the return response, and seven moments of flow dynamics inside their bands: 95 percent confidence bands for estimated moments and fixed tolerances for the earnings moments. I simulate 42 countries for 4,000 quarters and average moments over three seeds. I chose the parameters in an earlier search, on a version of the forecast data that did not require announced base earnings and with a coding error in the third-year term, which used the forecast-error regressions in place of the boom-bust regressions. On the corrected data and objective, a search

over 800 nearby parameter vectors finds no version with all 25 targets inside their bands: in every version the largest standardized deviation is at least 2.03. I keep the original parameters, and the comparisons below use the corrected data.

Different moments pin down different parameters. The long-horizon global slopes and the size of the boom pin down the impact response of beliefs to global news, $(1+\theta)\kappa \approx 0.37$, with $\theta = 0.87$ and $\kappa = 0.20$. The gain corresponds to a perceived ratio of trend-shock to transitory variance of $\kappa^2/(1-\kappa) \approx 0.05$. Learning and the diagnostic distortion both contribute to this response, and the moments pin down mainly their combination. A version with learning alone, $\theta = 0$ and $\kappa = 0.37$, fits the 25 targets about as well, and across three sets of simulation seeds its objective is 4 to 7 percent higher than the baseline's, mainly from the flow moments, and none of 800 versions with $\theta = 0$ and the other parameters re-chosen does better than the baseline on the baseline seeds, and the six best of them also do worse on both other seed sets (Table 5). Given θ, I use the path of the price factor and the earnings autocorrelations to choose the reference gain and the cycle, which govern how fast the boom reverses. The local slopes, which fade as announced news enters growth, pin down $K_N = 0.30$. Flow comovement and the world flow-return correlation pin down $K_N^G = 0.23$ and the common noise share $c_H = 0.12$, and flow persistence pins down $\psi = 0.65$ and the persistence of the common noise. These last four parameters rest on flow moments, so the model's flow fit is not an independent test.

8.4 Results

Beliefs. The model over-reacts to global news at long horizons and under-reacts to local news at short horizons (Figure 5, left). Its five-year global slope is -0.40 against -0.68 in the data, outside the data band ($z = 2.2$), and its one-year local slope is 0.32 against 0.23. The other nine slopes lie inside their bands. The model's global slopes flatten beyond three years while the data's keep falling, so the model understates the over-reaction at the longest horizon.

Boom and bust. The model's global factor is at 0.46 of its impact value after one year and 0.14 after two, against 0.52 and 0.10 in the data (Figure 5, right), and every coefficient through two years lies inside the data band. With rational investor beliefs the response after two years

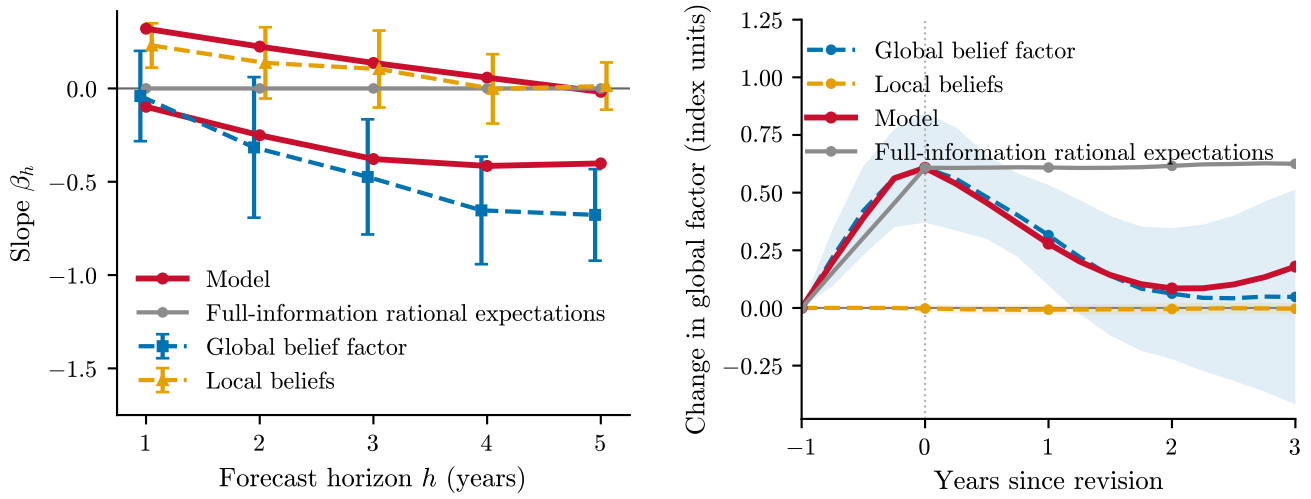


Figure 5: Model and data: forecast errors (left) and the global financial cycle factor (right)

Notes: Dashed: data (Figures 2 and 3) with 95 percent confidence intervals. Solid: model, mean of three simulation seeds with 42 countries. Right panel: model and FIRE responses scaled to the data at $k = 0$; FIRE is a full-information version with an announced share of local news of 0.8 and stochastic volatility, and differs from the rational-investor benchmark discussed in the text.

is 1.06 times its impact value: prices move with the persistent part of the news and stay there. Because the model's factor is the cumulative world return, the same path also compares with the data's cumulative world return, 0.42 and 0.37, and lies inside that band. In the third year the model's response rises again, which the data do not show.

Table 4: Return and flow comovement

	Data	Model	Rational investor beliefs
Return comovement	0.635	0.57	0.30
Flow comovement	0.078	0.060	no flows
Flows with returns, within country	0.22	0.23	
Flows with returns, world	0.52	0.47	
Flow persistence, country	0.20	0.22	
Flow persistence, world	0.35	0.25	
Return stress ratio	5.1	1.6	

Notes: Comovement: average pairwise cross-country correlation of quarterly returns (42 countries) and of portfolio-equity inflows as a percent of GDP (43 countries, 1995–2024). Flows with returns: same-quarter correlation within country and of world aggregates. Persistence: first-order autocorrelation. Stress ratio: mean comovement contribution in stress quarters relative to other quarters, with stress quarters the top fifth by the VIX in the data and by the rolling volatility of world earnings growth in the model. I used the flow moments to choose K_N^G/K_N, c_H, ψ, and ρ_ν, so the flow fit is not an out-of-sample test.

Prices and flows. Return comovement is 0.57 in the model and 0.635 in the data, and flow comovement 0.060 and 0.078 (Table 4). The model reproduces the gap: returns comove about ten times as much as flows in the model and about eight times as much in the data. Flows move with returns, with correlations of 0.23 within country and 0.47 for world aggregates (data 0.22

and 0.52), and they persist, with autocorrelations of 0.22 and 0.25 (data 0.20 and 0.35). Flow comovement is below the data ($z = -1.7$), and the return stress ratio is 1.6 against 5.1, since the model has no stochastic volatility and so no episodes in which comovement rises sharply. In the data, the comovement of returns is concentrated in high-volatility quarters such as 2008 and 2020, when risk-bearing capacity also falls. Matching that concentration would require a version of the model with time-varying risk, and the belief channel and a risk channel would then interact in the episodes that matter most for financial stability.

The role of beliefs. With rational investor beliefs on the same shocks, foreign and home investors agree on every valuation, so by Proposition 1 there are no flows; return comovement falls from 0.57 to 0.30; and the global factor does not reverse. Shared fundamentals and the earnings cycle thus account for return comovement of about 0.30 in the model, and learning, the diagnostic distortion, and the information gap between foreign and home investors account for the rest. The decomposition holds within the model. The data show a belief distortion in forecasts and a predictable reversal in prices, but they do not by themselves measure what share of observed comovement comes from beliefs.

The forecast-error tests establish that the global component of analysts' long-horizon forecasts is not a rational expectation: its revisions predict errors of the opposite sign, at horizons and in samples where the test has power. The boom-bust regressions establish that this component predicts a reversal of the global financial cycle that the total revision does not predict. The firm-level evidence shows that realized earnings growth falls far short of optimistic long-term forecasts, consistent with the forecast-error tests. None of these results identifies a causal effect of beliefs on prices, and none measures what share of the cycle's variation comes from beliefs.

The model adds a mechanism and a quantitative check. It shows that a distortion close to the size measured in forecasts, shared by foreign and home investors, produces a boom and bust of the size observed in the price factor and a price-flow gap of the size observed in the data, while rational investor beliefs on the same shocks produce neither. Learning alone with the same impact response produces similar results (Table 5), so these moments do not separate

the diagnostic distortion from learning. The model has misses that matter for this conclusion. Its five-year global slope is less negative than in the data, so it uses less over-reaction than the forecasts show; its return stress ratio is far below the data, so it does not capture the concentration of comovement in crises; and its flow parameters are chosen with flow moments. The model therefore shows that the calibrated beliefs, whose forecast-error slopes lie inside their 95 percent bands for nine of the ten horizon and news combinations, produce approximately the observed reversal and gap; it does not show that beliefs are the only force behind them.

9 Conclusion

Equity prices co-move strongly across countries, while cross-border equity flows co-move much less. This paper studies the role of expectations in explaining this gap, using a real-time global belief factor extracted from revisions of analysts' earnings forecasts for firms in 97 countries. Long-horizon forecasts over-react to the global belief factor, which also predicts a boom and bust in the global financial cycle. In a multi-country model in which foreign and home investors over-weight recent news about global earnings growth, calibrated beliefs reproduce this boom and bust and the gap between price and flow comovement. Shared revisions of global beliefs reprice every market at once, while flows follow the smaller differences between foreign and home valuations. This belief channel complements the monetary, intermediary, and risk-appetite channels studied in the literature, and forecasts measure it in real time.

References

Acalin, J. and Rebucci, A. (2020). Global business and financial cycles: A tale of two capital account regimes. *NBER Working Paper*, (27739).

Adam, K., Marcet, A., and Beutel, J. (2017). Stock price booms and expected capital gains. *American Economic Review*, 107(8):2352–2408.

Angeletos, G.-M., Huo, Z., and Sastry, K. A. (2021). Imperfect macroeconomic expectations: Evidence and theory. *NBER Macroeconomics Annual*, 35:1–86.

Avdjiev, S., Hardy, B., Kalemli-Özcan, Ş., and Servén, L. (2022). Gross capital flows by banks, corporates, and sovereigns. *Journal of the European Economic Association*, 20(5):2098–2135.

Bae, K.-H., Stulz, R. M., and Tan, H. (2008). Do local analysts know more? A cross-country study of the performance of local analysts and foreign analysts. *Journal of Financial Economics*, 88(3):581–606.

Bekaert, G., Engstrom, E. C., and Xu, N. R. (2022). The time variation in risk appetite and uncertainty. *Management Science*, 68(6):3975–4004.

Bordalo, P., Gennaioli, N., La Porta, R., and Shleifer, A. (2019). Diagnostic expectations and stock returns. *Journal of Finance*, 74(6):2839–2874.

Bordalo, P., Gennaioli, N., La Porta, R., and Shleifer, A. (2021). Diagnostic bubbles. *Journal of Financial Economics*, 141(3):1060–1077.

Bordalo, P., Gennaioli, N., La Porta, R., and Shleifer, A. (2024). Belief overreaction and stock market puzzles. *Journal of Political Economy*, 132(5):1450–1484.

Bordalo, P., Gennaioli, N., Ma, Y., and Shleifer, A. (2020). Overreaction in macroeconomic expectations. *American Economic Review*, 110(9):2748–2782.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2018). Diagnostic expectations and credit cycles. *Journal of Finance*, 73(1):199–227.

Bouchaud, J.-P., Krueger, P., Landier, A., and Thesmar, D. (2019). Sticky expectations and the profitability anomaly. *Journal of Finance*, 74(2):639–674.

Bruno, V. and Shin, H. S. (2015). Capital flows and the risk-taking channel of monetary policy. *Journal of Monetary Economics*, 71:119–132.

Camous, A., Monnet, E., and Puy, D. (2025). World cycles revisited: Diverging trends in prices and quantities. *American Economic Journal: Macroeconomics*. Forthcoming.

Cerutti, E., Claessens, S., and Puy, D. (2019). Push factors and capital flows to emerging markets: Why knowing your lender matters more than fundamentals. *Journal of International Economics*, 119:133–149.

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

Davis, J. S. and van Wincoop, E. (2024). A theory of capital flow retrenchment. *Journal of International Economics*, 150:103952.

De la O, R. and Myers, S. (2021). Subjective cash flow and discount rate expectations. *Journal of Finance*, 76(3):1339–1387.

Driscoll, J. C. and Kraay, A. C. (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics*, 80(4):549–560.

Forbes, K. J. and Warnock, F. E. (2012). Capital flow waves: Surges, stops, flight, and retrenchment. *Journal of International Economics*, 88(2):235–251.

Fuster, A., Laibson, D., and Mendel, B. (2010). Natural expectations and macroeconomic fluctuations. *Journal of Economic Perspectives*, 24(4):67–84.

Gabaix, X. and Koijen, R. S. J. (2021). In search of the origins of financial fluctuations: The inelastic markets hypothesis. Working Paper 28967, National Bureau of Economic Research.

Gabaix, X. and Maggiori, M. (2015). International liquidity and exchange rate dynamics. *Quarterly Journal of Economics*, 130(3):1369–1420.

Gilchrist, S. and Zakrajšek, E. (2012). Credit spreads and business cycle fluctuations. *American Economic Review*, 102(4):1692–1720.

Glaeser, E. L. and Nathanson, C. G. (2017). An extrapolative model of house price dynamics. *Journal of Financial Economics*, 126(1):147–170.

He, Z. and Krishnamurthy, A. (2013). Intermediary asset pricing. *American Economic Review*, 103(2):732–770.

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

Kekre, R. and Lenel, M. (2022). Monetary policy, redistribution, and risk premia. *Econometrica*, 90(5):2249–2282.

Kekre, R. and Lenel, M. (2024). The flight to safety and international risk sharing. *American Economic Review*, 114(6):1650–1691.

Kohlhas, A. N. and Walther, A. (2021). Asymmetric attention. *American Economic Review*, 111(9):2879–2925.

Lorenzoni, G. (2009). A theory of demand shocks. *American Economic Review*, 99(5):2050–2084.

Lucas, R. E. (1972). Expectations and the neutrality of money. *Journal of Economic Theory*, 4(2):103–124.

Maćkowiak, B. and Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review*, 99(3):769–803.

Maxted, P. (2024). A macro-finance model with sentiment. *Review of Economic Studies*, 91(1):438–475.

Miranda-Agrippino, S. and Rey, H. (2015). World asset markets and the global financial cycle. *NBER Working Paper*, (21722).

Miranda-Agrippino, S. and Rey, H. (2020). U.s. monetary policy and the global financial cycle. *Review of Economic Studies*, 87(6):2754–2776.

Na, S. and Xie, Y. (2025). Expectations and the uip puzzles when foresight is limited. *Journal of International Economics*, 158.

Na, S. and Yoo, D. (2025). Overreaction and macroeconomic fluctuation of the external balance. *Journal of Monetary Economics*, 151.

Rey, H. (2015). Dilemma not trilemma: The global financial cycle and monetary policy independence. *NBER Working Paper*, (21162).

Stock, J. H. and Watson, M. W. (2016). Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics. *Handbook of Macroeconomics*, 2:415–525.

Van Nieuwerburgh, S. and Veldkamp, L. (2009). Information immobility and the home bias puzzle. *Journal of Finance*, 64(3):1187–1215.

Appendix

A Data Construction

A.1 Forecasts and revisions

Source and basis. Forecasts and realized earnings come from the I/B/E/S unadjusted summary (consensus) and actuals files, U.S. and international, and the issuer identification file. I put each firm's forecasts and actuals in one currency and on one share basis, so that a forecast and the realization it is scored against are comparable. I assign a firm to the country of its I/B/E/S issuer record. Each country code is mapped to a country by the modal reporting currency of its firms (national currencies before the euro for euro-area firms) and, where several countries share a currency, by firm names.

Forecast path. For each firm I keep the last consensus of each calendar quarter. Let t be its statistical date and let $E^p_{i,t}$ be the firm's last annual actual for a fiscal year that had ended by t, the base of the forecast. The anchors of the forecast path are log base earnings $\log E^p_{i,t}$ at t, the log median consensus forecasts of annual earnings per share at the ends of the next one to four fiscal years, and the log of the sum of the next four quarterly forecasts at the end of the last of those quarters. I interpolate linearly in months between anchors to obtain $F_{i,t}(\tau)$, the forecast of log earnings at date τ. Beyond the last anchor I extend the path with the analysts' long-term growth forecast $g^{LT}_{i,t}$, adding $\log(1+g^{LT}_{i,t}/100)$ every twelve months. A forecast for date τ is *analyst-made* if it is interpolated between fiscal-year anchors or extended with the long-term growth forecast. I drop forecasts that would require extrapolating beyond the analysts' own forecasts in any other way, except that a target up to two months beyond the last anchor counts as interpolated.

Revisions for the same date. The revision at horizon h is $\Delta F_{i,t}(t+h) = F_{i,t}(t+h) - F_{i,t-4}(t+h)$. The two statistical dates are not exactly four quarters apart, so I approximate the earlier forecast at the target date by moving along its forecast path by $\delta = (t - t_{-4}) - 1$ years, using the slope between its yearly horizon values; when no slope is available and the dates differ by at

most 14 days, I use the earlier forecast unadjusted. For the five-year horizon the earlier forecast needs a six-year forecast, which I build as the earlier five-year forecast plus one year of growth at the earlier long-term growth forecast. I keep a revision only if the forecasts at both dates are analyst-made and the base earnings E^p had been announced by both statistical dates.

Realized earnings and errors. I use positive annual actuals, one per firm and fiscal year. Realized log earnings at date τ are interpolated linearly in time between the two annual actuals whose fiscal-year ends bracket τ, if the two are at most 400 days apart. The forecast error is $e_{i,t}(t+h) = a_{i,t+h} - F_{i,t}(t+h)$, with a realized log earnings. Because the forecast and the realization are levels of log earnings at the same date, the base of the forecast cancels.

Outliers. In the country tests I clip revisions and long-term growth forecasts at the first and ninety-ninth percentiles of observations from earlier years, so that the regressors use only information available at the forecast date, and winsorize forecast errors, the outcome, at full-sample percentiles. The firm-level tests of Section 7 winsorize long-term growth forecasts, realized earnings growth, and subsequent price changes at full-sample first and ninety-ninth percentiles.

Country series. The country revision $x_{c,h,t}$ is the mean of $\Delta F_{i,t}(t+h)$ over the country's firms in the quarter, and $n_{c,h,t}$ is their number. The country forecast error is the mean of $e_{i,t}(t+h)$ over the country's firms with an error.

A.2 Weights, prices, and flows

Country weights. The weight $w_{c,t}$ is the country's share of equity market capitalization across the panel countries in the previous year, from the World Bank, with OECD data for Estonia and Lithuania, Compustat for Taiwan, and an imputation from the ratio of market capitalization to GDP for Iceland. A country without a value in a year receives the smallest share observed among panel countries that year. I carry values forward, never backward.

Global financial cycle factor. I use the 2024 vintage of the factor of Miranda-Agrippino and Rey (2020), estimated on the updated list of series over 1980–2024, and average the monthly

values within each quarter.

Equity returns and flows. Country equity returns are quarterly U.S. dollar total returns of the iShares MSCI country funds, with the S&P 500 fund for the United States. Gross portfolio-debt inflows are from Avdjiev et al. (2022). Portfolio-equity inflows are the net incurrence of portfolio-equity liabilities in IMF balance-of-payments statistics, as a percent of GDP.

Firm prices. I link I/B/E/S firms to Compustat Global securities through the Compustat cross-reference to I/B/E/S tickers, and U.S. and Canadian firms to Compustat North America through the CRSP link, within the link's valid dates. Prices are monthly closes divided by the cumulative adjustment factor, for the security linked to the I/B/E/S ticker. The quarter-end price is the last observation in the quarter, and the h-quarter log price change is the difference of log adjusted prices. I drop price changes whose two endpoints are quoted in different currencies. When several I/B/E/S tickers map to one Compustat firm, I keep one firm-quarter and average the forecast.

B The Factor Model

Series and standardization. Each series is a country-horizon mean revision $x_{c,h,t}$, $h = 1, \ldots, 5$. In the estimation for year y, I standardize each series with its mean and standard deviation over quarters before year y, and a series enters once it has eight such observations. The panel starts in 1990, so the first estimation, for 1995, uses five years of data.

State space. Equation (4) defines the measurement equation. Each series loads on the global factor and on the factor of its region: North America, Latin America, Europe, Asia-Pacific (including Australia and New Zealand), Middle East and Africa, and offshore centres. A region with fewer than four eligible series has no regional factor, and its series load on the global factor only. The factors follow independent first-order autoregressions with unit innovation variances, which normalizes their scale. The idiosyncratic variance of series i in quarter t is $\psi_i / n_{i,t}$, so the information a series carries rises with the number of firms behind it.

Estimation. For each year y from 1995, I estimate the loadings, autoregressive coefficients, and idiosyncratic variances by quasi-maximum likelihood on data through the end of year $y-1$, using the EM algorithm with the Kalman smoother, allowing for missing observations, and starting from principal components. I then run the Kalman filter with these parameters through the quarters of year y and record the filtered global factor $\xi_{t|t}$. Each year's estimation starts afresh, so no parameter uses information from after the year before the one it is applied to.

Global and local beliefs. For the forecast-error tests, G_t enters in revision units: $\xi_{t|t}$ times the slope of the weighted cross-country mean of $x_{c,t} - \mu_c$ on $\xi_{t|t}$. The country's global part is $s_c\lambda_c\xi_{t|t}$ in revision units, and local beliefs follow equation (5) under the weights of each specification. For the boom-bust regressions, both regressors are standardized.

Fit. Across the 52 countries with at least 20 quarters of cells with three or more firms since 1995, the median correlation between a country's five-year revision and its real-time global part, over those quarters, is 0.58.

C Prices Comove, Flows Do Not

Across 42 countries, the average pairwise correlation of quarterly country equity returns is 0.63, and the first principal component explains 64 percent of their variance. For gross portfolio-debt inflows as a percent of GDP (Avdjiev et al., 2022), the same statistics are 0.07 and 13 percent. Portfolio-equity inflows comove at 0.078 across 43 countries. Regressing the common return factor on the common flow factor gives $R^2 \approx 0.05$.

A bound for flow-driven transmission. Suppose each market's return is a constant multiple of the net flow into it plus idiosyncratic noise,

$$r_c = M\, f_c + \varepsilon_c, \tag{14}$$

with ε_c uncorrelated across countries and with flows. Then

$$\operatorname{corr}(r_c, r_{c'}) = \operatorname{corr}(f_c, f_{c'})\sqrt{\lambda_c\lambda_{c'}}, \qquad \lambda_c = \frac{M^2 \operatorname{Var}(f_c)}{M^2 \operatorname{Var}(f_c) + \operatorname{Var}(\varepsilon_c)} \leq 1, \tag{15}$$

so $|\text{corr}(r_c, r_{c'})| \leq |\text{corr}(f_c, f_{c'})|$ for every pair, whatever the size of M. On the 39 countries and the quarters that both samples cover, the average absolute return correlation, 0.65, is about five times the average absolute flow correlation, 0.13, which no constant multiplier can produce (Gabaix and Koijen, 2021).

What the bound does not establish. The bound assumes a constant multiplier on each market's own, correctly measured net demand, and it applies only insofar as gross portfolio-debt inflows track the flow that moves prices. It does not rule out intermediary and portfolio-adjustment models (Bruno and Shin, 2015; Davis and van Wincoop, 2024; He and Krishnamurthy, 2013), in which required returns move in equilibrium and reprice existing holdings without a matching flow. A time-varying or cross-market price impact, or large measurement error in the relevant flow, could also produce the gap; with equal reliability across countries and classical measurement errors uncorrelated across countries, measurement error would have to account for about 89 percent of measured flow variance to turn a true flow correlation of 0.63 into the measured 0.07. Nor does the gap distinguish rational from belief-driven repricing, since a common discount-rate factor looks the same in prices either way. The forecast errors of Section 6 make that distinction, and the firm earnings shortfall of Section 7 is consistent with it.

D The Model

D.1 Environment

Timing and shocks. Time is quarterly, $t = 0, 1, \ldots$, and there are N countries. Each quarter, earnings growth is realized, news is announced, investors and analysts update their beliefs, investors trade, and markets clear. All primitive shocks are independent standard normal draws, scaled as below.

Earnings. Log earnings growth is $x_{c,t} = g_t + \ell_{c,t}$, with variance parameters $\sigma_g^2 = 0.355$ and $\sigma_\ell^2 = 1 - \sigma_g^2$. Global growth is

$$g_t = z_t + z_t^C + \sigma_g \sqrt{(1 - s_z)(1 - s_C)}\, \varepsilon_t, \qquad z_t = \rho z_{t-1} + u_t, \qquad \text{var}(z) = s_z \sigma_g^2,$$

and the earnings cycle is

$$z_t^C = \rho_C z_{t-1}^C - e_C\, C_{t-1} + u_t^C, \qquad C_t = C_{t-1} + z_t^C,$$

with both states scaled so that the cycle has variance share s_C of the non-persistent part of global growth. The scale is computed once from the simulated path. Because cycle growth error-corrects toward the cycle level, the earnings level returns to trend after a cycle. Local growth is

$$\ell_{c,t} = z_{c,t}^L + \sigma_\ell \sqrt{(1-s_z)\max\Big(0,\, 1 - \tfrac{s_N}{1-s_z}\Big)}\;\epsilon_{c,t} + n_{c,t-4},$$
$$z_{c,t}^L = \rho z_{c,t-1}^L + u_{c,t}^L, \qquad n_{c,t} \sim N(0,\, s_N \sigma_\ell^2),$$

where $n_{c,t}$ is local news announced at t and realized in growth at $t+4$. With $s_N = 1$ firms announce all transitory local variation four quarters ahead. The true trend of growth is zero.

Information. Foreign investors observe $x_{c,t}$ for every c, hence world growth $\bar{x}_t = N^{-1}\sum_c x_{c,t}$. Home investors in country c observe $x_{c,t}$ and the signal $s_{c,t}$ of equation (9), with $\nu_{c,t}$ independent across countries and the common part $\nu_t = \rho_\nu \nu_{t-1} + \sqrt{1-\rho_\nu^2}\,\varepsilon_t^\nu$, which has unit variance. They estimate global growth by static signal extraction with perceived variances $V_g = \sigma_g^2$, $V_\ell = 1 - \sigma_g^2$, and $V_n = n_H^2 \sigma_g^2$,

$$\hat{g}_{c,t} = c_x\, x_{c,t} + c_s\, s_{c,t}, \qquad c_x = \frac{V_g V_n}{\Delta}, \quad c_s = \frac{V_g V_\ell}{\Delta}, \quad \Delta = (V_g + V_\ell)(V_g + V_n) - V_g^2.$$

Announced news is public. Home investors and analysts absorb a share K_N of it, and foreign investors a share K_N^G. Pending news is $N_{c,t} = \sum_{j=0}^{3} n_{c,t-j}$.

Beliefs. Foreign investors learn the trend from world growth, $b_t^G = b_{t-1}^G + \kappa(\bar{x}_t - b_{t-1}^G)$, and home investors from their estimate, $b_{c,t}^H = b_{c,t-1}^H + \kappa(\hat{g}_{c,t} - b_{c,t-1}^H)$. Each belief has a reference $\underline{b}_t^j = \underline{b}_{t-1}^j + \underline{\kappa}(b_{t-1}^j - \underline{b}_{t-1}^j)$, and the distorted belief is $\tilde{b}_t^j = b_t^j + \theta(b_t^j - \underline{b}_t^j)$. Investors recognize the cycle level, $\hat{C}_t = \hat{C}_{t-1} + \kappa_C(C_t - \hat{C}_{t-1})$.

Valuations and demands. Investor j values country c's equity, relative to its earnings, at $V_{c,t}^j$ in equation (11). Let $p_{c,t}$ be the log price relative to earnings. Investors have mean-variance

demands with constant perceived risk, normalized so that the target holding of type j is $V^j_{c,t} - p_{c,t}$, the expected excess value per unit of risk. Home investors hold their target, $q^H_{c,t} = V^H_{c,t} - p_{c,t}$. Foreign investors adjust a share ψ of the way to theirs each quarter,

$$q^G_{c,t} = (1-\psi)\, q^G_{c,t-1} + \psi\,(V^G_{c,t} - p_{c,t}).$$

Equity supply is fixed, and clearing requires $\omega q^G_{c,t} + (1-\omega) q^H_{c,t} = \Pi$, where Π is supply times the price of risk.

Analysts. Analysts' belief about country c mixes home and foreign analysts with the measured local share, $b^A_{c,t} = 0.88\,\tilde{b}^H_{c,t} + 0.12\,\tilde{b}^G_t$. A forecast made at τ of growth over the next J quarters is

$$\Phi_{c,\tau}(J) = J\, b^A_{c,\tau} + S(J)\,(z_\tau + z^C_\tau + z^L_{c,\tau}) + K_N \sum_{j=1}^{\min(J,4)} n_{c,\tau+j-4}, \qquad S(J) = \rho\,\frac{1-\rho^J}{1-\rho},$$

so analysts know the persistent news, treat cycle growth as autoregressive news without the error correction, and absorb K_N of announced news. The published consensus at t for growth over the next J quarters averages vintages $\tau = t-j$, $j = 0, \ldots, 15$, each evaluated at horizon $J+j$ and net of growth realized since the vintage,

$$F_t\big(Y_{c,t+J} - Y_{c,t}\big) = \sum_{j=0}^{15} w_j \Big[\big(Y_{c,t-j} - Y_{c,t}\big) + \Phi_{c,t-j}(J+j)\Big], \qquad Y_{c,t} = \textstyle\sum_{s\le t} x_{c,s},$$

with weights proportional to 0.26^j, half of each shifted one quarter back because forecasts enter the consensus with a delay. The consensus does not enter prices.

D.2 Equilibrium

Definition. Given initial conditions and the shocks, a *temporary equilibrium* is a sequence of beliefs $\{b^j_t, \underline{b}^j_t, \hat{C}_t\}$, holdings $\{q^G_{c,t}, q^H_{c,t}\}$, and prices $\{p_{c,t}\}$ such that, each quarter t: (i) beliefs follow the updating rules above, using earnings, signals, and news dated t and earlier and taking the cycle scale as given; (ii) holdings are the demands above at $p_{c,t}$; and (iii) every market clears. Beliefs are not required to be consistent with the equilibrium price process, so this is not a rational-expectations equilibrium.

Existence and uniqueness. Beliefs at t do not depend on $p_{c,t}$, and each demand is linear in $p_{c,t}$ with a negative slope, so each clearing condition has a unique solution given the predetermined states. Markets are coupled only through world growth and the common home noise, which enter beliefs, so given beliefs each market clears separately and the model can be solved forward one quarter at a time.

Prices and flows. Define $\mathcal{D}_{c,t} = (q^G_{c,t} - \Pi)/(1-\omega)$. The clearing condition gives $p_{c,t} = V^H_{c,t} - \Pi/(1-\omega) + \omega q^G_{c,t}/(1-\omega) = V^H_{c,t} + \omega \mathcal{D}_{c,t} - \Pi$. Substituting this price into the foreign demand,

$$\Pi + (1-\omega)\mathcal{D}_{c,t} = (1-\psi)\big[\Pi + (1-\omega)\mathcal{D}_{c,t-1}\big] + \psi\big(V^G_{c,t} - V^H_{c,t} + \Pi - \omega\mathcal{D}_{c,t}\big),$$

so that

$$\mathcal{D}_{c,t} = \delta\, \mathcal{D}_{c,t-1} + (1-\delta)\big(V^G_{c,t} - V^H_{c,t}\big), \qquad \delta = \frac{(1-\psi)(1-\omega)}{1-\omega+\omega\psi}.$$

Gross foreign inflows are net purchases by foreign investors, $f_{c,t} = q^G_{c,t} - q^G_{c,t-1} = (1-\omega)\Delta\mathcal{D}_{c,t}$. With $\psi = 1$, $\delta = 0$, the price is $\omega V^G + (1-\omega)V^H - \Pi$, and inflows are $(1-\omega)\Delta(V^G - V^H)$. At $\omega = 0.3$ and the chosen $\psi = 0.646$, $\delta = 0.28$.

Proof of Proposition 1. The recursion for $\mathcal{D}_{c,t}$ involves the valuations only through $V^G_{c,t} - V^H_{c,t} = \Delta^G_{c,t} - \Delta^H_{c,t}$, in which A_t cancels. Hence $\mathcal{D}_{c,t}$ and $f_{c,t} = (1-\omega)\Delta\mathcal{D}_{c,t}$ do not depend on A_t, and $p_{c,t} = A_t + \Delta^H_{c,t} + \omega\mathcal{D}_{c,t} - \Pi$ moves one for one with A_t. If $V^G_{c,t} = V^H_{c,t}$ at all dates and $\mathcal{D}_{c,0} = 0$, the recursion gives $\mathcal{D}_{c,t} = 0$ and $f_{c,t} = 0$ for all t. □

Benchmarks. Under full-information rational expectations, investors know z_t, $z^L_{c,t}$, the cycle state, and all announced news, and value each market at expected future growth, $V_{c,t} = \frac{\rho}{1-\rho}(z_t + z^L_{c,t}) + N_{c,t} - C_t$. Prices equal $V_{c,t}$ up to the constant premium, and there are no flows. With rational investor beliefs, foreign and home investors both hold this valuation, and the analysts' consensus and everything else in the model are unchanged.

D.3 Moments and simulation

Targets. The 25 targets are the ten forecast-error slopes of equation (6) (global and local, $h = 1, \ldots, 5$), the response of the global financial cycle factor in equation (7) relative to its value at $k = 0$, the stress ratio of flow comovement, the average pairwise correlation of country earnings growth, and two autocorrelations of earnings growth. In the model, forecast errors and revisions come from the simulated consensus, the global financial cycle factor is the cumulative world-mean return, and flows are foreign inflows. The seven moments of flow dynamics are the first-order autocorrelations of country and world flows, the within-country correlation of flows with returns, the correlations of world flows with the world return in the previous, same, and next quarter, and the first-order autocorrelation of country returns.

Flow dynamics. Without gradual rebalancing and with serially independent common noise, the first-order autocorrelation of flows is -0.15 for world aggregates and -0.06 within country. Gradual rebalancing alone raises them to 0.09 and 0.20, and adding persistence of the common noise raises them to 0.25 and 0.22.

Simulation. I simulate $T = 4000$ quarters for 42 countries, discard the first 400, and average moments over three seeds. Two further sets of three seeds give the same conclusions.

Learning alone. Table 5 compares the baseline with versions without the diagnostic distortion, $\theta = 0$. Keeping $\kappa = 0.20$ roughly halves the impact response of beliefs and roughly triples the objective. Raising the gain to $\kappa = 0.37$, which keeps the impact response $(1 + \theta)\kappa$ at its baseline value, restores the fit to the 25 targets; the objective remains 4 to 7 percent higher than the baseline's on each set of seeds, mainly from the autocorrelations of flows and the other flow moments. A search over 800 versions with $\theta = 0$, κ drawn over $[0.12, 0.60]$, and the other parameters drawn around the baseline finds none with a lower objective than the baseline on the baseline seeds; the best has $\kappa = 0.36$, and the six best also have higher objectives than the baseline on both other seed sets.

Table 5: The baseline and versions with learning alone

	Baseline	$\theta = 0$ $\kappa = 0.20$	$\theta = 0$ $\kappa = 0.37$	$\theta = 0$, best of 800
θ	0.87	0	0	0
κ	0.20	0.20	0.37	0.36
Objective, baseline seeds	31.8	92.7	33.7	34.4
Objective, seed set 2	31.4	93.8	32.6	33.1
Objective, seed set 3	33.1	87.6	35.5	35.5
25-target criterion, baseline seeds	19.1	37.9	18.5	21.6
Largest $\lvert z \rvert$ among the 25 targets	2.20	2.37	2.20	2.16
Flow autocorrelation term	2.7	16.7	4.3	3.2

Notes: The objective is the sum of squared standardized deviations from the 25 targets plus the terms for the return response, flow moments, third-year factor path, and flow autocorrelations described in Section 8. Each seed set averages three simulations of 4,000 quarters for 42 countries. The best of 800 is the version with the lowest objective on the baseline seeds among draws with $\theta = 0$; its other parameters differ from the baseline.