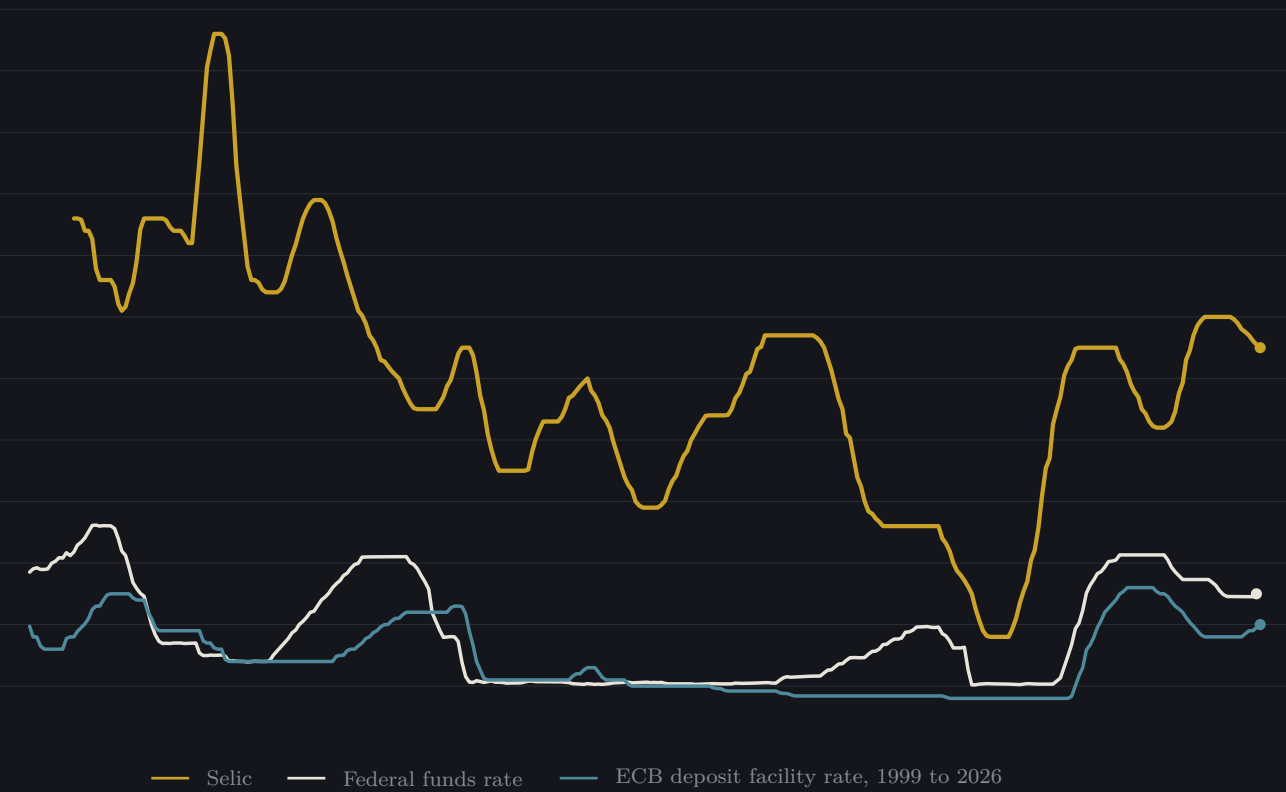
Rates, Rules and Words

Monetary Policy after Interest and Prices

Neutral rates, reaction functions, the yield curve
and central bank communication at the Fed,
the ECB and the Banco Central do Brasil

Arthur Mota



Rates, Rules and Words

Monetary Policy after Interest and Prices

Arthur Mota

Copyright © 2026 Arthur Mota

This book is free to share. It is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International licence (CC BY-NC-ND 4.0): you may copy and redistribute it in any medium or format, including posting it online, provided you credit the author, do not use it for commercial purposes and do not distribute modified versions. The licence is at creativecommons.org/licenses/by-nc-nd/4.0.

The views expressed in this book are the author's own and do not represent those of any institution with which he is or has been affiliated. Nothing in it is investment advice or a recommendation to buy or sell any asset.

Data are from public sources: the Federal Reserve Board, the Federal Reserve Banks of New York, Philadelphia, St. Louis, Chicago, Atlanta and Cleveland, the European Central Bank, Eurostat, the Banco Central do Brasil, IBGE and Tesouro Nacional. Levels are as published up to mid-October 2026. All estimates, simulations and figures are the author's, from code and data archived with the manuscript.

Typeset in Latin Modern with LaTeX.

Cataloguing-in-Publication Data

Mota, Arthur.
Rates, rules and words : monetary policy after *Interest and Prices* / Arthur Mota.
x, 263 p. : ill. ; 24 cm.
Includes bibliographical references.

1. Monetary policy. 2. Interest rates. 3. Central banks. 4. Inflation targeting. 5. Phillips curve. 6. Federal Reserve. 7. European Central Bank. 8. Banco Central do Brasil. I. Title. II. Title: Monetary policy after Interest and Prices.

CDD 332.46 CDU 336.74

First edition, 2026

Contents

I Foundations **1**

1 The Core Model **3**
1.1 The household 4
1.1.1 Preferences and the budget constraint 4
1.1.2 Optimality conditions 4
1.1.3 The log-linear Euler equation 5
1.1.4 What the Euler equation says 6
1.2 The firm 7
1.2.1 Demand for a differentiated good 7
1.2.2 Technology and flexible-price behaviour 7
1.2.3 Calvo pricing 8
1.2.4 From marginal cost to the output gap 11
1.2.5 Inflation as the present value of future gaps 12
1.3 The natural rate and the dynamic IS curve 12
1.3.1 From the Euler equation to the IS curve 12
1.3.2 The long rate is what matters 13
1.3.3 Cost-push shocks and the divine coincidence 14
1.4 Closing the model with a rule 15
1.4.1 The policy rule 15
1.4.2 Determinacy 16
1.4.3 The Taylor principle 16
1.4.4 Solving the model 18
1.5 Pegs, money and the Fisher relation 19
1.5.1 Why an interest-rate peg fails 19
1.5.2 A model without money 20
1.5.3 The Fisher relation and the neo-Fisherian question . . 21
1.5.4 Promises about the far future 22
1.6 Calibration 22

1.7 What the model says about three shocks 23
1.7.1 A monetary policy shock 23
1.7.2 A natural-rate shock 24
1.7.3 A cost-push shock 25
1.8 What the core model leaves out 25

2 The Loss Function **27**
2.1 Why inflation is costly in the model 28
2.1.1 Price dispersion 28
2.1.2 How Calvo pricing turns inflation into dispersion . . . 28
2.1.3 The welfare loss 29
2.1.4 Reading the welfare weight 30
2.2 Mandates as loss functions 31
2.2.1 The Fed: a dual mandate 31
2.2.2 The ECB: a hierarchical mandate 32
2.2.3 The BCB: a continuous target with a band and a penalty 33
2.3 Optimal policy under discretion 35
2.3.1 The bank's problem 35
2.3.2 The targeting rule 35
2.3.3 The equilibrium 35
2.3.4 The variance frontier 36
2.4 The horizon as a weight 37
2.5 Bands, asymmetries and penalties 38
2.5.1 The band 39
2.5.2 The discrete penalty 39
2.5.3 Asymmetry: the linex loss 40
2.6 Which inflation? The open economy 41
2.6.1 Domestic and imported prices 41
2.6.2 Why the BCB targets the IPCA anyway 42
2.7 Why the target is positive 42
2.7.1 Measurement bias 42
2.7.2 The lower bound 43
2.7.3 Downward wage rigidity 43
2.7.4 The cost: trend inflation and price dispersion 43
2.8 Beyond inflation and the gap 44

3 Reaction Functions and Rules **47**
3.1 Two kinds of rule 47
3.2 Commitment 48
3.2.1 The bank's problem with promises 48

3.2.2 Deriving the optimal plan . . . 49
3.2.3 What the rule says . . . 49
3.2.4 Solving for the path . . . 50
3.2.5 Why commitment is hard . . . 51
3.3 The Taylor rule and its family . . . 51
3.3.1 The original rule . . . 51
3.3.2 The rules the Fed publishes . . . 54
3.3.3 Four variants . . . 55
3.3.4 Interest-rate smoothing as history dependence . . . 56
3.4 The BCB's published reaction function . . . 56
3.4.1 How the BCB describes its policy . . . 56
3.4.2 The rule in the BCB's model . . . 56
3.4.3 Replicating the BCB's model from public information 57
3.4.4 Forecast targeting and instrument rules on forecasts . 58
3.4.5 A model with inflation inertia . . . 59
3.4.6 Comparing rules . . . 60
3.4.7 The optimal simple rule . . . 61
3.5 Estimating a reaction function . . . 63
3.5.1 The specification . . . 63
3.5.2 Estimation . . . 63
3.5.3 Real-time data . . . 63
3.5.4 Identification . . . 64

4 Policy as a Game 65
4.1 The inflation bias . . . 66
4.1.1 The model . . . 66
4.1.2 The timing . . . 66
4.1.3 Solving the game . . . 67
4.1.4 Commitment and the cost of discretion . . . 67
4.1.5 The bias in the core model . . . 68
4.2 Institutional answers . . . 68
4.2.1 The conservative central banker . . . 68
4.2.2 Goal and instrument independence . . . 70
4.2.3 The Walsh contract: a penalty that removes the bias . 71
4.2.4 Reputation . . . 72
4.2.5 Reputation under uncertainty about type . . . 73
4.3 Decisions by committee . . . 74
4.3.1 Aggregating preferences . . . 74
4.3.2 Aggregating information . . . 75
4.4 The fiscal-monetary game . . . 75

4.4.1 Two authorities, one budget constraint 75
4.4.2 Active and passive policies 76
4.4.3 The game of chicken 76
4.4.4 The arithmetic of Brazilian debt 78
4.5 Political pressure 80

5 Information **83**
5.1 Signal extraction 84
5.1.1 The formula 84
5.1.2 The Lucas islands 84
5.1.3 Learning over time and the anchoring of expectations 85
5.1.4 What the anchoring data say about q 86
5.2 The central bank's information advantage 86
5.2.1 What a decision reveals 86
5.2.2 Delphic and Odyssean 87
5.3 The social value of public information 88
5.3.1 The beauty contest 88
5.3.2 The equilibrium 88
5.3.3 Welfare 89
5.3.4 Disagreement in the data 90
5.4 Higher-order beliefs and sluggish prices 91
5.4.1 Woodford's imperfect common knowledge 91
5.4.2 Why the hierarchy adjusts slowly 92
5.5 Rational inattention 93
5.5.1 Attention as a scarce resource 93
5.5.2 Evidence and implications 93
5.5.3 Testing for rigidity in Focus forecasts 94
5.6 Choosing what to reveal 96
5.6.1 Bayesian persuasion 96
5.6.2 What this means for a central bank 96

6 The Balance Sheet in Theory **99**
6.1 Why purchases should not matter 99
6.1.1 A Modigliani-Miller theorem for open-market operations 99
6.1.2 The lower bound adds a twist 100
6.1.3 Three assumptions 100
6.2 Preferred habitat 101
6.2.1 Segmented demand and arbitrage 101
6.2.2 A one-factor model 101
6.2.3 What a purchase does 102

6.2.4 State dependence . . . 103
6.3 Credit and the capital of intermediaries . . . 104
6.4 Signalling . . . 105
6.4.1 The announcement days . . . 105
6.5 The euro area: purchases against fragmentation . . . 107
6.6 Brazil: a balance sheet in dollars . . . 108
6.7 Reserves and the operating framework . . . 110
6.7.1 Corridor and floor . . . 110
6.7.2 Where reserves become scarce . . . 111
6.8 Quantitative tightening, 2022–2026 . . . 111
6.9 What the balance sheet adds . . . 113

II The Modern Facts 115

7 The Neutral Rate 117
7.1 Measuring an unobserved rate . . . 117
7.1.1 How the model is estimated . . . 118
7.2 The United States and the euro area . . . 119
7.2.1 Replicating the estimate . . . 120
7.2.2 Why the neutral rate fell . . . 120
7.2.3 Two models from the same authors . . . 121
7.2.4 Real time against hindsight . . . 122
7.3 What the central banks estimate . . . 122
7.3.1 The Fed's longer-run dot . . . 123
7.3.2 The BCB's set of measures . . . 123
7.4 A global neutral rate? . . . 124
7.5 Brazil . . . 124
7.5.1 Where the model fails . . . 124
7.5.2 What the market expects . . . 125
7.6 Market real rates . . . 126
7.7 A suite of measures for each market . . . 127
7.8 The stance through the cycles . . . 130
7.9 Policy when r^* is uncertain . . . 131
7.9.1 Shrinking the estimate . . . 131
7.9.2 Brainard's principle . . . 132
7.9.3 Rules that do not need r^* . . . 132
7.10 What follows for policy . . . 132

8 Measuring Inflation and Expectations 135
8.1 The measures . . . 135
8.1.1 How a trimmed mean is built . . . 136
8.1.2 How expectations are measured . . . 136
8.2 Which measure predicts headline inflation . . . 137
8.3 Decomposing the measures . . . 140
8.3.1 Free and administered prices in Brazil . . . 140
8.3.2 What a breakeven contains . . . 141
8.4 Anchoring . . . 142
8.4.1 Do long-run expectations react to inflation news? . . . 143

9 Inflation after 2021 145
9.1 The flat curve before 2020 . . . 146
9.1.1 What the data showed . . . 146
9.1.2 Three explanations for flatness . . . 146
9.2 Why the curve bends . . . 147
9.2.1 State-dependent pricing . . . 147
9.2.2 The labour market and the v/u ratio . . . 148
9.2.3 Capacity and attention . . . 149
9.3 A threshold Phillips curve, 2001 to 2026 . . . 150
9.3.1 United States . . . 150
9.3.2 Could the kink have been seen in 2020? . . . 151
9.3.3 What the curve says for 2026 . . . 152
9.3.4 Brazil . . . 153
9.4 The Bernanke-Blanchard decomposition . . . 154
9.4.1 The four equations . . . 154
9.4.2 Estimates . . . 155
9.4.3 Decomposition . . . 156
9.5 The soft landing . . . 157
9.6 The euro area: an energy shock with a lagged wage response . 159
9.7 Brazil: early hikes, a second round, and record-low unemployment . . . 159
9.8 Policy with a kinked curve . . . 160
9.8.1 Precautionary policy . . . 160
9.8.2 Three lessons . . . 161

10 Estimated Reaction Functions 163
10.1 From rules to reaction functions . . . 163
10.1.1 Partial adjustment . . . 163
10.1.2 Forecasts, real-time data and identification . . . 164

10.2 The Fed . . . 165
10.2.1 Quarterly data, the lower bound and the shadow rate 165
10.2.2 What the Fed saw: real-time Greenbook data . . . 166
10.2.3 How late was the Fed in 2021? . . . 168
10.2.4 The SEP medians as a reduced-form reaction function 169
10.3 The ECB . . . 170
10.3.1 A rule written on survey expectations . . . 170
10.3.2 Late, then tight, then on the rule . . . 172
10.4 The BCB . . . 172
10.4.1 Replicating the published rule . . . 172
10.4.2 Where the Copom departed from its rule . . . 174
10.5 Estimated rules and the welfare loss . . . 175
10.6 What an estimated rule leaves out . . . 177

11 The Models Inside Central Banks 179
11.1 Two families of models . . . 179
11.1.1 What a central bank model is for . . . 179
11.1.2 Semi-structural and DSGE . . . 180
11.2 FRB/US . . . 180
11.2.1 Structure . . . 180
11.2.2 The PAC equation . . . 181
11.2.3 Prices and wages . . . 182
11.2.4 Expectations: VAR and model-consistent . . . 183
11.2.5 Running the model . . . 183
11.2.6 An oil shock and the choice of rule . . . 185
11.2.7 Rules inside FRB/US: a core inflation shock . . . 185
11.2.8 A term premium shock . . . 186
11.3 NAWM II and the euro area models . . . 187
11.3.1 Structure . . . 187
11.3.2 How a DSGE model is estimated . . . 187
11.3.3 The 2025 harmonised comparison . . . 188
11.4 SAMBA and the BCB's small models . . . 189
11.4.1 SAMBA . . . 189
11.4.2 The small semi-structural models . . . 190
11.5 The same shock in seven models . . . 190
11.6 Model uncertainty . . . 191

III Transmission **193**

12 The Yield Curve **195**
12.1 Expectations and the term premium . . . 195
12.1.1 Prices, yields and forward rates . . . 195
12.1.2 The decomposition . . . 196
12.1.3 An affine model in four steps . . . 196
12.1.4 Testing the expectations hypothesis . . . 197
12.1.5 ACM and Kim-Wright . . . 198
12.2 The shape of the curve . . . 199
12.2.1 Nelson-Siegel . . . 199
12.2.2 Can the curve be forecast? . . . 199
12.3 The slope as a recession signal . . . 201
12.4 How policy cycles reach the ten-year yield . . . 202
12.5 What the curve says in October 2026 . . . 205
12.5.1 The market-implied path . . . 205
12.5.2 The euro area . . . 205
12.5.3 Brazil . . . 206
12.6 What the curve means for policy . . . 208

13 Financial Conditions and the Balance Sheet **209**
13.1 What a financial conditions index measures . . . 209
13.1.1 From policy to asset prices . . . 209
13.1.2 Two indices built by the Fed . . . 210
13.2 Do financial conditions predict growth? . . . 212
13.2.1 The central tendency . . . 212
13.2.2 The risk to growth . . . 212
13.3 The policy rate and financial conditions in three cycles . . . 213
13.4 How a policy surprise moves financial conditions . . . 214
13.5 The balance sheet as a stock . . . 216
13.5.1 From flows to stocks . . . 216
13.5.2 Quantitative tightening and reserves . . . 217
13.6 Silicon Valley Bank, March 2023 . . . 217
13.7 The euro area and Brazil . . . 218
13.7.1 Systemic stress in the euro area . . . 218
13.7.2 An index for Brazil . . . 219
13.8 What financial conditions mean for policy . . . 220

IV Words as an Instrument 221

14 Forward Guidance 223
14.1 Two kinds of guidance . . . 223
14.2 Guidance in the core model . . . 224
14.2.1 The forward guidance puzzle . . . 224
14.2.2 Resolving it . . . 225
14.3 Measuring guidance: the US . . . 226
14.3.1 Three dimensions of a policy announcement . . . 226
14.3.2 Episodes . . . 227
14.3.3 Information and policy surprises in the data . . . 227
14.3.4 The puzzle measured . . . 228
14.4 The ECB . . . 230
14.5 The BCB . . . 231
14.5.1 Surprises against the Focus survey . . . 231
14.5.2 Guidance in two cycles . . . 232
14.6 Lessons for guidance . . . 232

15 Statements, Minutes and Press Conferences 235
15.1 Why words move prices . . . 235
15.1.1 Cheap talk and its limits . . . 235
15.1.2 Public signals . . . 236
15.2 The FOMC statement . . . 236
15.2.1 Length and change . . . 236
15.2.2 Does the redline move markets? . . . 237
15.3 The press conference . . . 238
15.3.1 The ECB: the conference is the news . . . 238
15.4 The Copom: measuring tone from the text . . . 239
15.4.1 The BCB's communication . . . 239
15.4.2 An index of tone . . . 240
15.4.3 What the tone adds . . . 241
15.4.4 Redlines and the market in Brazil . . . 242
15.4.5 Limits of a level index . . . 242
15.5 Communication as policy . . . 242

V Policy Now 245

16 Optimal Policy in the New Regime 247
16.1 The lower bound . . . 247

16.1.1 The problem 247
16.1.2 Five strategies 248
16.1.3 How often? 250
16.2 A moving neutral rate 251
16.2.1 What an error in r^* costs 251
16.2.2 The neutral rate and the bound together 251
16.3 A kinked Phillips curve 252
16.4 What each central bank would do differently 252
16.4.1 The Fed 252
16.4.2 The ECB 253
16.4.3 The BCB 253
16.5 Conclusion 254

Part I

Foundations

Chapter 1

The Core Model

On 1 September 2021 the two-year US Treasury yield stood at 0.20%. The federal funds rate was at its floor, with a target range of 0 to 0.25%, and the Federal Open Market Committee had not yet said when it would move. On 16 March 2022 it raised the range by a quarter of a point. By then the two-year yield had already reached 1.9%, and on 31 March it closed at 2.28%. Before the Fed had done anything that a reader of the overnight rate would notice, the market had priced most of a year of tightening, mortgage rates had risen by more than a point, and financial conditions had tightened accordingly.

That sequence is the subject of this chapter and, in a sense, of the whole book. A central bank sets one short interest rate for one night. Output, employment and inflation respond to that rate only because households and firms look through it to the path they expect it to follow. A model of monetary policy is therefore a model of expectations: of how today's decisions to spend and to set prices depend on what people believe about future rates, future income and future prices.

The model built here is the one Woodford (2003a) calls neo-Wicksellian and Galí (2015) the basic New Keynesian model. It has two optimising agents. A household decides how much to consume and how much to work, and in doing so produces an ***Euler equation*** linking spending today to the expected real interest rate. A firm sets the price of a differentiated good knowing it may not be able to change it for a while, and in doing so produces the ***New Keynesian Phillips curve*** linking inflation to expected future inflation and to the output gap. A third equation describes the central bank. The three together determine output, inflation and the interest rate.

Every step of the derivation is shown, because each assumption made along the way is a place where a later chapter will relax the model and where the data will push back. The chapter ends by solving the model for three shocks and reading the answers as statements about policy.

1.1 The household

1.1.1 Preferences and the budget constraint

A representative household lives forever and values consumption and leisure. Its expected lifetime utility is

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{C_t^{1-\sigma}}{1-\sigma} - \frac{N_t^{1+\varphi}}{1+\varphi} \right), \tag{1.1}$$

where C_t is a consumption basket, N_t hours worked and $\beta \in (0,1)$ the discount factor. Two parameters shape behaviour. σ governs how much the household dislikes uneven consumption over time: its inverse, $1/\sigma$, is the ***intertemporal elasticity of substitution***, the percentage change in consumption growth for a one-point change in the real interest rate. φ governs how much it dislikes variation in hours: $1/\varphi$ is the ***Frisch elasticity*** of labour supply, the percentage change in hours for a one percent change in the real wage holding the marginal utility of consumption fixed.

Each period the household earns wages $W_t N_t$, receives the profits D_t of the firms it owns, and holds one-period nominal bonds. A bond bought at price Q_t in period t pays one unit of money in $t+1$. The budget constraint, in money terms, is

$$P_t C_t + Q_t B_t \leq B_{t-1} + W_t N_t + D_t, \tag{1.2}$$

where P_t is the price of the consumption basket and B_t the number of bonds held at the end of the period. Spending on goods and new bonds cannot exceed the payoff of old bonds plus income. A no-Ponzi condition rules out borrowing without limit.

The bond is the only asset, and its price defines the interest rate. The nominal rate i_t is the yield on the bond, $Q_t = e^{-i_t}$, so $i_t = -\log Q_t$. This is the rate the central bank controls.

1.1.2 Optimality conditions

Attach a multiplier μ_t to the budget constraint of period t. The Lagrangian is

$$\mathcal{L} = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\sigma}}{1-\sigma} - \frac{N_t^{1+\varphi}}{1+\varphi} - \mu_t \left(P_t C_t + Q_t B_t - B_{t-1} - W_t N_t - D_t \right) \right]. \tag{1.3}$$

Differentiating with respect to C_t, N_t and B_t gives three conditions:

$$C_t^{-\sigma} = \mu_t P_t, \tag{1.4}$$

$$N_t^{\varphi} = \mu_t W_t, \tag{1.5}$$

$$\mu_t Q_t = \beta \, \mathbb{E}_t \, \mu_{t+1}. \tag{1.6}$$

The first says the marginal utility of a unit of money, μ_t, equals the marginal utility of consumption divided by its price. The second says the marginal disutility of an hour equals the money it earns, valued at μ_t. The third says the cost of a bond today, Q_t units of money valued at μ_t, equals the discounted expected value of the unit of money it pays tomorrow.

Dividing (1.5) by (1.4) eliminates μ_t and gives the ***labour supply*** condition

$$\frac{W_t}{P_t} = C_t^{\sigma} N_t^{\varphi}. \tag{1.7}$$

The real wage equals the marginal rate of substitution between leisure and consumption. A household that consumes more (C_t higher) values extra income less and demands a higher wage for the same hours: the wealth effect on labour supply.

Substituting (1.4) for t and $t+1$ into (1.6) gives the ***Euler equation*** in levels:

$$Q_t = \beta \, \mathbb{E}_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \frac{P_t}{P_{t+1}} \right]. \tag{1.8}$$

Read it as a pricing equation. The bond costs Q_t and pays one unit of money, worth P_t/P_{t+1} in goods of today's value. The household values goods tomorrow relative to goods today at the rate $\beta(C_{t+1}/C_t)^{-\sigma}$, the stochastic discount factor. If consumption is expected to grow fast, goods tomorrow are worth little relative to goods today, and the household will only hold the bond at a low price, that is at a high interest rate.

1.1.3 The log-linear Euler equation

Equation (1.8) is nonlinear. To combine it with the other equations of the model, approximate it around a steady state with constant consumption and zero inflation. Write lower-case letters for logs, $c_t = \log C_t$, and $\pi_{t+1} = \log(P_{t+1}/P_t)$ for inflation. Then

$$e^{-i_t} = \mathbb{E}_t \exp(-\rho - \sigma \Delta c_{t+1} - \pi_{t+1}), \qquad \rho \equiv -\log \beta. \tag{1.9}$$

In the steady state $\Delta c = \pi = 0$ and $i = \rho$: the nominal rate equals the rate of time preference. A first-order approximation of the exponential around that point, $e^x \approx 1 + x$, gives

$$i_t = \rho + \sigma\, \mathbb{E}_t \Delta c_{t+1} + \mathbb{E}_t \pi_{t+1}. \tag{1.10}$$

Rearranging, and using goods-market clearing $y_t = c_t$ (there is no investment, government or trade in this model), gives the form used throughout the book:

$$y_t = \mathbb{E}_t y_{t+1} - \frac{1}{\sigma}\left(i_t - \mathbb{E}_t \pi_{t+1} - \rho\right). \tag{1.11}$$

The approximation drops one thing worth naming. With uncertainty, the exact equation contains terms in the variance of consumption growth and inflation: a precautionary motive to save that rises with risk. These are second-order terms: a first-order log-linearisation drops them, and a second-order approximation would keep them. They matter in a crisis, which is one reason the model needs extending in chapter 13.

1.1.4 What the Euler equation says

Equation (1.11) has three readings.

Intertemporal substitution. The term $i_t - \mathbb{E}_t \pi_{t+1}$ is the expected real interest rate. When it rises above ρ, saving is more attractive and the household postpones consumption: output today falls relative to output expected tomorrow. The size of the response is $1/\sigma$. With $\sigma = 1$ (log utility), a real rate one point higher for a year raises the expected growth of consumption over that year by one point, so consumption today falls by one percent relative to the path it would otherwise follow. With $\sigma = 2$, by half a point.

Expectations of future income. Output today rises one for one with output expected tomorrow. A household that expects to be richer next year spends more now. This term carries the whole future into the present, as the next section shows.

The only channel. In this model the real rate is the only way policy reaches spending. There is no investment, so no cost-of-capital channel; no credit, so no bank-lending or balance-sheet channel; no exchange rate, so no trade channel; no wealth effect from asset prices except through future income. Chapters 12 and 13 add these; the core model gives the interest-rate channel in its purest form.

In practice: how large is $1/\sigma$

Macroeconomic estimates of the elasticity of intertemporal substitution are small. A meta-analysis of thousands of published estimates by Havránek (2015) finds that once selective reporting is corrected, the elasticity for consumption in aggregate data is around a third, which corresponds to $\sigma \approx 3$. Models used by central banks typically set σ between 1 and 2 and obtain larger interest-rate effects through investment and credit rather than through household consumption. The calibration of this chapter uses $\sigma = 1$, the textbook benchmark, and the reader should treat its consumption response as an upper bound for the household channel alone.

1.2 The firm

1.2.1 Demand for a differentiated good

The consumption basket aggregates a continuum of goods indexed by $j \in [0, 1]$:

$$C_t = \left(\int_0^1 C_t(j)^{\frac{\varepsilon-1}{\varepsilon}} dj \right)^{\frac{\varepsilon}{\varepsilon-1}}, \tag{1.12}$$

with $\varepsilon > 1$ the elasticity of substitution between goods (Dixit and Stiglitz, 1977). Given prices $P_t(j)$, the household buys the basket at minimum cost. Minimising $\int P_t(j) C_t(j)\, dj$ subject to the basket equalling C_t gives the demand for each good,

$$C_t(j) = \left(\frac{P_t(j)}{P_t} \right)^{-\varepsilon} C_t, \qquad P_t = \left(\int_0^1 P_t(j)^{1-\varepsilon} dj \right)^{\frac{1}{1-\varepsilon}}. \tag{1.13}$$

Demand for good j falls with its relative price, with elasticity ε, and rises with total spending. P_t is the price index: the minimum cost of one unit of the basket. When all prices are equal, $P_t = P_t(j)$.

1.2.2 Technology and flexible-price behaviour

Firm j produces with labour, $Y_t(j) = A_t N_t(j)^{1-\alpha}$, where A_t is aggregate productivity and $\alpha \in [0, 1)$ measures decreasing returns. Its nominal marginal cost is the wage divided by the marginal product of labour:

$$\Psi'_t(j) = \frac{W_t}{(1-\alpha) A_t N_t(j)^{-\alpha}}. \tag{1.14}$$

Because of decreasing returns, a firm that produces more has higher marginal cost.

If the firm could reset its price every period, it would choose $P_t(j)$ to maximise $P_t(j)Y_t(j) - \Psi_t(Y_t(j))$ subject to the demand curve (1.13). The first-order condition is

$$P_t(j) = \frac{\varepsilon}{\varepsilon - 1}\,\Psi_t'(j) \equiv \mathcal{M}\,\Psi_t'(j). \tag{1.15}$$

The price is a constant markup $\mathcal{M}$ over marginal cost. With $\varepsilon = 9$ the markup is 12.5%. Market power is what gives the firm a price to set. The next assumption is what makes the price sticky.

1.2.3 Calvo pricing

Each period a randomly chosen fraction $1-\theta$ of firms may reset their price. The other θ keep last period's price. The draw is independent of the firm's history and of how long its price has been in place, so the probability that a price set today is still in force k periods from now is θ^k, and the expected duration of a price is

$$\sum_{k=0}^{\infty}(1-\theta)\theta^k(k+1) = \frac{1}{1-\theta}. \tag{1.16}$$

With $\theta = 0.75$ at a quarterly frequency, prices last a year on average.

> **In practice: how sticky are prices**
>
> Bils and Klenow (2004) find, in the microdata behind the US consumer price index, that the median price lasts about four and a half months. Nakamura and Steinsson (2008) show that most of those changes are temporary sales that return to the old price; excluding them, the median duration of a regular price is eight to eleven months. The calibration $\theta = 0.75$, a year, sits at the long end of that range. It is the value that makes the model's inflation respond to monetary shocks as slowly as the data suggest, which is one sign that Calvo pricing alone understates the sources of inertia. Chapter 9 returns to the evidence that the frequency of price changes rose sharply in 2021 and 2022.

The price index

Firms that do not reset are a random sample of all firms, so their average price is last period's price index. Firms that reset all choose the same

price P_t^*, since they face the same problem. The price index (1.13) therefore satisfies

$$P_t^{1-\varepsilon} = \theta P_{t-1}^{1-\varepsilon} + (1-\theta)\,(P_t^*)^{1-\varepsilon}\,. \tag{1.17}$$

Dividing by $P_{t-1}^{1-\varepsilon}$ and writing $\Pi_t = P_t/P_{t-1}$,

$$\Pi_t^{1-\varepsilon} = \theta + (1-\theta)\left(\frac{P_t^*}{P_{t-1}}\right)^{1-\varepsilon}. \tag{1.18}$$

In a zero-inflation steady state $\Pi = 1$ and $P^* = P_{t-1}$. A first-order approximation gives

$$\pi_t = (1-\theta)\,(p_t^* - p_{t-1})\,. \tag{1.19}$$

Inflation is the fraction of firms that reset times the average size of their adjustment. It can be high because many firms change prices or because those that change, change a lot. Calvo pricing fixes the first margin, which is the assumption the 2021 to 2022 data challenge.

The reset price

A firm that resets in t knows its price may last. It chooses P_t^* to maximise the expected discounted profit over all the periods in which the price survives:

$$\max_{P_t^*} \sum_{k=0}^{\infty} \theta^k\, \mathbb{E}_t\left[\Lambda_{t,t+k}\left(P_t^* Y_{t+k|t} - \Psi_{t+k}\big(Y_{t+k|t}\big)\right)\right], \tag{1.20}$$

where $\Lambda_{t,t+k} = \beta^k (C_{t+k}/C_t)^{-\sigma}(P_t/P_{t+k})$ is the household's stochastic discount factor for money, the firm's owners' valuation, and $Y_{t+k|t} = (P_t^*/P_{t+k})^{-\varepsilon} Y_{t+k}$ the demand the firm faces at the reset price. The first-order condition is

$$\sum_{k=0}^{\infty} \theta^k\, \mathbb{E}_t\Big[\Lambda_{t,t+k}\, Y_{t+k|t}\left(P_t^* - \mathcal{M}\,\Psi'_{t+k|t}\right)\Big] = 0. \tag{1.21}$$

The firm sets its price so that, on average over the states in which the price is in force, it equals the markup over marginal cost. With flexible prices ($\theta = 0$) the sum has one term and the condition collapses to the static markup rule.

Log-linearising (1.21) around the zero-inflation steady state gives

$$p_t^* = \mu + (1-\beta\theta)\sum_{k=0}^{\infty} (\beta\theta)^k\, \mathbb{E}_t\big[\psi_{t+k|t}\big]\,, \qquad \mu = \log \mathcal{M}, \tag{1.22}$$

where $\psi_{t+k|t}$ is the log nominal marginal cost in $t+k$ of a firm that last reset in t. The reset price is the markup over a weighted average of expected future marginal costs. The weights $(1-\beta\theta)(\beta\theta)^k$ sum to one and decline at the rate $\beta\theta$: costs far in the future matter less because they are discounted and because the price is less likely to survive that long.

Two consequences follow immediately. A firm that expects costs to rise raises its price now, before the costs arrive, because the price will be stuck when they do. And a central bank that convinces firms that future costs will be lower lowers prices today. The reset price is where expectations enter the price level.

Firm-specific marginal cost

The reset firm's marginal cost differs from the average because its price differs from the average. A firm with a high relative price sells less, hires fewer hours, and, with decreasing returns, has a lower marginal cost. Using the production function and the demand curve, the log marginal cost of a firm that reset in t relates to the average real marginal cost mc_{t+k} as

$$\psi_{t+k|t} - p_{t+k} = mc_{t+k} - \frac{\alpha\varepsilon}{1-\alpha}\left(p_t^* - p_{t+k}\right). \tag{1.23}$$

This term is the source of ***strategic complementarity***: a firm's optimal price depends on other firms' prices. The larger $\alpha\varepsilon$, the more a firm that raises its price relative to others loses demand and sees its marginal cost fall, and so the smaller the adjustment it chooses. Strategic complementarity makes the aggregate price level slow to respond even when individual firms respond quickly.

The Phillips curve

Substituting the firm-specific marginal cost into (1.22) and writing the sum recursively gives

$$p_t^* - p_{t-1} = \beta\theta\,\mathbb{E}_t\left(p_{t+1}^* - p_t\right) + (1-\beta\theta)\,\Theta\,\widehat{mc}_t + \pi_t, \qquad \Theta \equiv \frac{1-\alpha}{1-\alpha+\alpha\varepsilon}, \tag{1.24}$$

where $\widehat{mc}_t = mc_t + \mu$ is the deviation of real marginal cost from its steady-state value of $-\mu$. Now use (1.19) to replace the reset price: $p_t^* - p_{t-1} = \pi_t/(1-\theta)$ and $\mathbb{E}_t(p_{t+1}^* - p_t) = \mathbb{E}_t\pi_{t+1}/(1-\theta)$. Then

$$\frac{\pi_t}{1-\theta} - \pi_t = \frac{\beta\theta}{1-\theta}\,\mathbb{E}_t\pi_{t+1} + (1-\beta\theta)\,\Theta\,\widehat{mc}_t. \tag{1.25}$$

The left side is $\theta\pi_t/(1-\theta)$. Multiplying through by $(1-\theta)/\theta$ gives the New Keynesian Phillips curve in terms of marginal cost:

$$\pi_t = \beta\,\mathbb{E}_t\pi_{t+1} + \lambda\,\widehat{mc}_t, \qquad \lambda \equiv \frac{(1-\theta)(1-\beta\theta)}{\theta}\,\Theta. \tag{1.26}$$

The coefficient λ falls as prices become stickier (θ up) and as strategic complementarity rises (Θ down). With $\beta = 0.99$, $\theta = 0.75$, $\alpha = 0.25$ and $\varepsilon = 9$, the first factor is 0.086 and $\Theta = 0.25$, so $\lambda = 0.0215$: a one percent rise in real marginal cost sustained for one quarter raises quarterly inflation by about two hundredths of a point, or 0.09 points annualised. Prices respond slowly to costs.

1.2.4 From marginal cost to the output gap

Marginal cost is not observed, but in this model it is a function of output. Real marginal cost is the real wage over the marginal product of labour:

$$mc_t = (w_t - p_t) - \left[\log(1-\alpha) + a_t - \alpha n_t\right]. \tag{1.27}$$

Substitute the labour supply condition (1.7), $w_t - p_t = \sigma c_t + \varphi n_t$; the aggregate production function, $n_t = (y_t - a_t)/(1-\alpha)$ to first order; and $c_t = y_t$. Collecting terms,

$$mc_t = \left(\sigma + \frac{\varphi+\alpha}{1-\alpha}\right) y_t - \frac{1+\varphi}{1-\alpha}\,a_t - \log(1-\alpha). \tag{1.28}$$

Marginal cost rises with output, for two reasons visible in the coefficient: workers demand a higher wage to supply more hours (φ) and to give up the higher consumption they already have (σ), and each extra hour is less productive (α). It falls with productivity.

Define ***natural output*** y_t^n as the output that would prevail if prices were flexible. With flexible prices every firm charges the markup, so real marginal cost is constant at $-\mu$. Setting $mc_t = -\mu$ in (1.28) and solving,

$$y_t^n = \psi_{ya}\,a_t + \vartheta, \qquad \psi_{ya} \equiv \frac{1+\varphi}{\sigma(1-\alpha)+\varphi+\alpha}, \tag{1.29}$$

with ϑ a constant. Subtracting the flexible-price version of (1.28) from the general one gives

$$\widehat{mc}_t = \left(\sigma + \frac{\varphi+\alpha}{1-\alpha}\right) x_t, \qquad x_t \equiv y_t - y_t^n. \tag{1.30}$$

The ***output gap*** x_t is the distance between actual and natural output, and marginal cost is proportional to it. Substituting into (1.26) gives the Phillips curve used in the rest of the book:

$$\pi_t = \beta\,\mathbb{E}_t\pi_{t+1} + \kappa\,x_t, \qquad \kappa \equiv \lambda\left(\sigma + \frac{\varphi+\alpha}{1-\alpha}\right). \tag{1.31}$$

With the baseline calibration, $\sigma + (\varphi+\alpha)/(1-\alpha) = 8.0$ and $\kappa = 0.172$.

Three features of the slope matter for the rest of the book. It is built from parameters with a micro meaning, not estimated freely, so a change in the slope must come from a change in price-setting or in labour supply. It is small, so large output gaps produce small changes in inflation. And the natural level of output is the one with flexible prices, which is not the level a statistical trend or a production-function estimate delivers. The output gap of the theory and the output gap of the data are different objects, a fact that complicates every empirical chapter.

1.2.5 Inflation as the present value of future gaps

Iterating (1.31) forward, and assuming expected inflation far in the future does not explode,

$$\pi_t = \kappa \sum_{k=0}^{\infty} \beta^k\,\mathbb{E}_t x_{t+k}. \tag{1.32}$$

Inflation today is the discounted sum of all output gaps expected from today on. It does not depend on past inflation at all.

A numerical example shows how strong this makes expectations. Suppose the public comes to expect an output gap of -1% in each of the next four quarters and zero after. By (1.32), inflation this quarter falls by $0.172\times(1+0.99+0.98+0.97) = 0.68$ points, or 2.7 points annualised, before any of the slack has materialised. In the data, inflation responds far less to announced or expected slack. The gap between the two is one of the main motives for the chapters on expectations (chapter 5) and on the Phillips curve in practice (chapter 9).

1.3 The natural rate and the dynamic IS curve

1.3.1 From the Euler equation to the IS curve

The Euler equation (1.11) relates output to its expected value. Subtracting y_t^n from both sides, and adding and subtracting $\mathbb{E}_t y_{t+1}^n$, writes it in terms

of the gap:

$$x_t = \mathbb{E}_t x_{t+1} - \frac{1}{\sigma}\left(i_t - \mathbb{E}_t \pi_{t+1} - r_t^n\right), \qquad r_t^n \equiv \rho + \sigma\, \mathbb{E}_t \Delta y_{t+1}^n. \tag{1.33}$$

This is the ***dynamic IS curve***. The new variable r_t^n is the ***natural rate of interest***: the real interest rate that would hold if prices were flexible, and therefore the real rate at which the output gap neither opens nor closes.

The natural rate moves with expected growth of natural output. From (1.29), $\mathbb{E}_t \Delta y_{t+1}^n = \psi_{ya}\, \mathbb{E}_t \Delta a_{t+1}$, so

$$r_t^n = \rho + \sigma \psi_{ya}\, \mathbb{E}_t \Delta a_{t+1}. \tag{1.34}$$

With the baseline calibration $\psi_{ya} = 1$ and $\sigma = 1$, so the natural rate moves one for one with expected productivity growth. A useful and slightly counterintuitive example: suppose productivity rises by 1% today and the gain decays at 10% a quarter, $a_t = 0.9\, a_{t-1} + \epsilon_t$. Productivity is now expected to fall back toward its old level, so $\mathbb{E}_t \Delta a_{t+1} = -0.1\%$ and the natural rate falls by 0.1 points per quarter, 0.4 annualised. A favourable but temporary supply shock lowers the natural rate, because households want to save part of the windfall and the real rate must fall to persuade them to spend it.

In richer models the natural rate also moves with fiscal policy, demographics, the demand for safe assets and the risk appetite of savers. Chapter 7 measures it. Woodford (2003a) builds his theory around one statement: policy is expansionary when the real rate is below the natural rate and contractionary when it is above, whatever the level of the nominal rate. This is Wicksell's natural rate in modern form, and it gives the model its name.

1.3.2 The long rate is what matters

Iterating (1.33) forward, with the gap expected to close eventually,

$$x_t = -\frac{1}{\sigma} \sum_{k=0}^{\infty} \mathbb{E}_t\left(r_{t+k} - r_{t+k}^n\right), \qquad r_t \equiv i_t - \mathbb{E}_t \pi_{t+1}. \tag{1.35}$$

The output gap depends on the sum of all expected future gaps between the real rate and the natural rate. That sum is, up to a scale, a long-term real interest rate. A central bank that raises the overnight rate by a point for one quarter and is expected to reverse it at once does almost nothing. A central bank that is expected to keep the real rate a point above neutral for two years does eight times as much, even if it has not yet moved.

Episode: the 2021–22 US tightening before the first hike

From September 2021 to March 2022 the federal funds rate did not move, but the two-year Treasury yield rose from 0.20% to 2.28%. Through (1.35), the relevant policy stance is the expected path, and the two-year yield is a high-frequency market proxy for it, contaminated by a term premium (chapter 12): by the end of March 2022 markets expected the funds rate to average about 2% over the following two years, against 0.2% six months earlier. A large part of the initial tightening of 2022, in its effect on spending, was priced before the first decision.

Episode: higher for longer, 2023

The Fed raised the funds rate for the last time in its cycle on 26 July 2023, to a range of 5.25 to 5.50%. Over the following three months, with no change in the policy rate, the ten-year Treasury yield rose from 3.85% on 20 July to 4.98% on 19 October. Part of the rise was a higher expected path, as markets absorbed the message that rates would stay high for longer, and part was a higher term premium. Either way, financial conditions tightened by more in those three months than in several of the hikes that preceded them. Chapter 12 separates the two components.

These episodes are the theoretical reason for the structure of the book. If the long rate is what moves spending, the yield curve is the transmission mechanism (chapter 12), and managing expectations of the future short rate (forward guidance, statements, press conferences) is an instrument in its own right (chapters 14 and 15).

1.3.3 Cost-push shocks and the divine coincidence

As written, the Phillips curve (1.31) contains no shock. That has a striking implication. If the central bank sets $i_t = r_t^n + \mathbb{E}_t \pi_{t+1}$ every period, the real rate equals the natural rate, the gap is zero by (1.35), and inflation is zero by (1.32). Stabilising output at its natural level and stabilising inflation are the same task. Blanchard and Galí (2007) call this the ***divine coincidence***. A central bank in this world faces no trade-off: it should track the natural rate.

The coincidence breaks when the natural level of output is not the efficient level and the distance between the two moves over time. Variations in the desired markup of firms, in the markup unions or workers demand over the competitive wage, or in distorting taxes all change flexible-price output without changing efficient output. They enter the Phillips curve as a ***cost-push shock*** u_t:

$$\pi_t = \beta\, \mathbb{E}_t \pi_{t+1} + \kappa\, x_t + u_t. \tag{1.36}$$

Now no path for the interest rate delivers both $\pi_t = 0$ and $x_t = 0$. The central bank must decide how to share the shock between inflation and output, and to decide it needs a loss function. Chapter 2 derives one.

Episode: the euro area energy shock, 2022

Euro area energy inflation reached 44.3% in March 2022 and headline HICP inflation peaked at 10.6% in October 2022. For a gas-importing economy, a rise in the price of energy relative to domestic output is close to a textbook cost-push shock: it raises inflation and lowers the efficient level of output at the same time. The ECB raised rates in July 2022, its first increase since 2011, and kept raising them while activity stagnated. The choice it faced is the one (1.36) describes: lean against the inflation and deepen the downturn, or accommodate it and risk the expectations that (1.32) makes decisive.

1.4 Closing the model with a rule

1.4.1 The policy rule

The IS curve (1.33) and the Phillips curve (1.36) contain three unknowns: x_t, π_t and i_t. The third equation describes how the central bank sets the interest rate. The simplest version is a Taylor-type rule,

$$i_t = \rho + \phi_\pi\, \pi_t + \phi_x\, x_t + v_t, \tag{1.37}$$

in which the bank raises the rate when inflation or the output gap rises, and v_t is a ***monetary policy shock***: the part of the decision not explained by the systematic response. Chapter 3 studies the alternatives: rules on expected inflation, inertial rules, optimal targeting rules. This chapter uses (1.37) to answer the first question any rule must answer: does it pin down a unique equilibrium?

1.4.2 Determinacy

Substitute the rule into the IS curve and write the IS and Phillips curves as a system in $\boldsymbol{z}_t = (x_t, \pi_t)'$:

$$\boldsymbol{z}_t = \boldsymbol{A}_T \, \mathbb{E}_t \boldsymbol{z}_{t+1} + \boldsymbol{B}_T \, \boldsymbol{s}_t, \qquad \boldsymbol{A}_T = \frac{1}{\sigma + \phi_x + \kappa\phi_\pi} \begin{pmatrix} \sigma & 1 - \beta\phi_\pi \\ \sigma\kappa & \kappa + \beta(\sigma + \phi_x) \end{pmatrix}, \tag{1.38}$$

where $\boldsymbol{s}_t$ collects the shocks $r_t^n - \rho$, v_t and u_t. Both x_t and π_t are ***jump variables***: nothing ties them to their own past, and they can take any value today that is consistent with expectations of tomorrow. For such a system, Blanchard and Kahn (1980) show that there is a unique bounded solution if and only if both eigenvalues of $\boldsymbol{A}_T$ lie inside the unit circle. If one or both lie outside, there is a continuum of bounded solutions, including solutions driven by beliefs with no connection to fundamentals, called ***sunspots***.

Working out the eigenvalue conditions for (1.38) with non-negative coefficients gives a single requirement (Bullard and Mitra, 2002):

$$\kappa\,(\phi_\pi - 1) + (1 - \beta)\,\phi_x > 0. \tag{1.39}$$

1.4.3 The Taylor principle

Condition (1.39) has a direct economic reading. Suppose inflation rose permanently by one point. In the long run, by the Phillips curve (1.31), the output gap would rise by $(1-\beta)/\kappa$ points, a small number because $1-\beta$ is close to zero. The rule would then raise the nominal rate by

$$\phi_\pi + \phi_x \frac{1 - \beta}{\kappa} \tag{1.40}$$

points. Condition (1.39) says this long-run response must exceed one. In other words, the real interest rate must rise when inflation rises. This is the ***Taylor principle***.

The intuition comes from a thought experiment. Suppose the public, for no fundamental reason, comes to expect higher inflation. If the central bank raises the nominal rate by less than the rise in expected inflation, the real rate falls. Through the IS curve the gap opens, and through the Phillips curve inflation rises: the expectation is validated. If the bank raises the nominal rate by more, the real rate rises, the gap closes, inflation falls, and the expectation is refuted. Only in the second case is the expectation inconsistent with equilibrium, and only then are fundamentals alone able to pin down inflation.

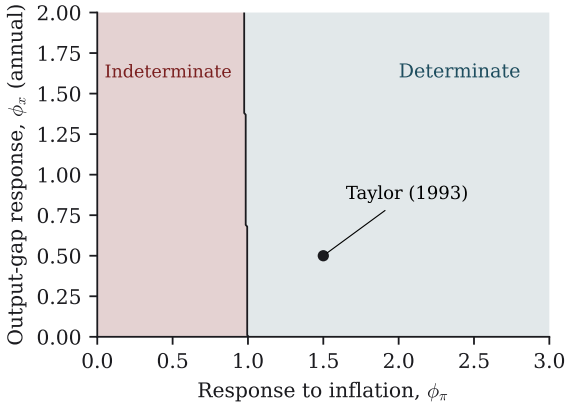


Figure 1.1. Determinacy region of the three-equation model under the rule (1.37), baseline calibration (Table 1.1). The output coefficient is shown in annualised units. The original Taylor (1993) coefficients sit well inside the determinate region.

Because $1-\beta$ is close to zero at a quarterly frequency, the output term barely relaxes the condition. Figure 1.1 shows the region. With an output coefficient of 0.5 in annualised units, the boundary sits at $\phi_\pi = 0.993$: in practice the principle reads $\phi_\pi > 1$.

Episode: the 1970s and after

Clarida et al. (2000) estimate a forward-looking version of (1.37) for the Fed. Before Paul Volcker became chairman in 1979, the estimated response to expected inflation is 0.83: when expected inflation rose, the Fed raised the funds rate by less, and the real rate fell. From 1979 to 1996 the estimate is 2.15. In the model, the earlier policy admits self-fulfilling inflation and the later one does not, which is one account of why inflation was high and volatile in the 1970s and became low and stable afterwards.

Episode: Brazil, 2021–22

In March 2021, with twelve-month IPCA inflation at 6.1% and rising, the BCB began raising the Selic from 2.00%. By August 2022 it had raised it twelve times, to 13.75%. Deflated by the Focus median of inflation expected over the next twelve months, the ex-ante real Selic

went from -1.7% in March 2021 to 8.1% in October 2022: the nominal rate rose by 11.75 points and the real rate by almost 10. Inflation peaked at 12.1% in April 2022 and fell to 3.2% by June 2023. This is the Taylor principle in operation: the nominal rate rose by more than inflation did, so the real rate rose, and inflation turned before the policy rate had come down.

1.4.4 Solving the model

With a unique equilibrium, the solution expresses x_t and π_t as functions of the shocks. Suppose each shock follows an AR(1), $s_t = \rho_s s_{t-1} + \epsilon_t^s$. Because the model is forward-looking and linear, the solution sets each endogenous variable proportional to the shock. For the policy shock, guess $x_t = \psi_x v_t$ and $\pi_t = \psi_\pi v_t$, so that $\mathbb{E}_t x_{t+1} = \rho_v \psi_x v_t$ and $\mathbb{E}_t \pi_{t+1} = \rho_v \psi_\pi v_t$.

The Phillips curve gives $\psi_\pi = \beta\rho_v\psi_\pi + \kappa\psi_x$, so

$$\psi_\pi = \frac{\kappa}{1 - \beta\rho_v}\,\psi_x. \tag{1.41}$$

The IS curve, with the rule substituted, gives

$$\psi_x = \rho_v\psi_x - \frac{1}{\sigma}\left(\phi_\pi\psi_\pi + \phi_x\psi_x + 1 - \rho_v\psi_\pi\right), \tag{1.42}$$

which rearranges to

$$\left[\sigma(1-\rho_v) + \phi_x\right]\psi_x + \left(\phi_\pi - \rho_v\right)\psi_\pi = -1. \tag{1.43}$$

Substituting the first equation into the second and solving,

$$\begin{aligned} x_t &= -(1-\beta\rho_v)\,\Lambda_v\,v_t, \qquad \pi_t = -\kappa\,\Lambda_v\,v_t, \\ \Lambda_v &\equiv \frac{1}{(1-\beta\rho_v)\left[\sigma(1-\rho_v)+\phi_x\right] + \kappa(\phi_\pi - \rho_v)}. \end{aligned} \tag{1.44}$$

Two properties are worth reading off (1.44). Λ_v falls as ϕ_π rises: a more aggressive systematic response dampens the effect of a given policy surprise, because the rule leans against the inflation the surprise causes. And Λ_v falls as the shock becomes less persistent, because a surprise expected to vanish quickly moves the long rate of (1.35) by little.

The same steps solve the model for the other two shocks, and the answers share the denominator. With the same persistence ρ for all shocks, write $\Lambda \equiv 1/\{(1-\beta\rho)[\sigma(1-\rho)+\phi_x] + \kappa(\phi_\pi - \rho)\}$. For a natural-rate shock, the

IS curve becomes $[\sigma(1-\rho)+\phi_x]\,\psi_x+(\phi_\pi-\rho)\,\psi_\pi=1$ and the Phillips curve is unchanged, so

$$x_t=(1-\beta\rho)\,\Lambda\,(r_t^n-\rho),\qquad \pi_t=\kappa\,\Lambda\,(r_t^n-\rho). \tag{1.45}$$

A natural-rate shock is a policy shock with the sign reversed: a rise in r^n with the policy rate unchanged is the same as a cut in the policy rate with r^n unchanged, because only their difference enters (1.33).

For a cost-push shock the IS curve has no shock, so $[\sigma(1-\rho)+\phi_x]\,\psi_x=-(\phi_\pi-\rho)\,\psi_\pi$, and the Phillips curve is $(1-\beta\rho)\psi_\pi=\kappa\psi_x+1$. Solving the two,

$$\pi_t=[\sigma(1-\rho)+\phi_x]\,\Lambda\,u_t,\qquad x_t=-(\phi_\pi-\rho)\,\Lambda\,u_t. \tag{1.46}$$

The coefficient on inflation falls with ϕ_π through Λ, while the coefficient on the gap, $(\phi_\pi-\rho)\Lambda$, rises with it. A more aggressive rule buys lower inflation with a deeper recession: the trade-off left by the failure of the divine coincidence, now in closed form.

1.5 Pegs, money and the Fisher relation

The determinacy result answers an old question and opens a new one. Both are worth settling here, because they recur whenever a central bank holds its rate fixed for a long time.

1.5.1 Why an interest-rate peg fails

Sargent and Wallace (1975) argued that a central bank that pegs the nominal interest rate leaves the price level undetermined: if only the rate is fixed, nothing anchors the level of prices, and any path of inflation consistent with that rate is an equilibrium. The conclusion was read for two decades as an argument for targeting money instead of interest rates.

The core model gives the modern version of the argument and its resolution. Set $\phi_\pi=\phi_x=0$ in (1.37). The eigenvalues of $\boldsymbol{A}_T$ in (1.38) have moduli 0.66 and 1.50 with the baseline calibration: one lies outside the unit circle, so the equilibrium is indeterminate. Restoring the output response to its baseline $\phi_x=0.125$, a weak inflation response, $\phi_\pi=0.8$, gives moduli of 0.72 and 1.09, still indeterminate (0.77 and 1.13 with $\phi_x=0$). With the Taylor rule, $\phi_\pi=1.5$, both are 0.85 (0.89 with $\phi_x=0$) and the equilibrium is unique.

The mechanism is the one described under the Taylor principle. Under a peg, a belief that inflation will be higher lowers the expected real rate, which

opens a positive gap, which raises inflation: the belief is self-fulfilling. What rules it out is not a quantity target but a response of the interest rate to inflation large enough to raise the real rate. Woodford (2003a) draws the central conclusion of his book from this: the problem with a peg is the absence of feedback, not the choice of an interest rate as the instrument. A central bank can determine inflation with an interest-rate rule alone, provided the rule responds strongly enough.

1.5.2 A model without money

The reader may have noticed that money does not appear in the model. There is no demand for currency, no money supply and no quantity theory. This is deliberate. Woodford (2003a) shows that in a ***cashless limit***, where money is used so little that its demand has a negligible effect on spending, the equilibrium of the three-equation model is the same as in a model with money, and the central bank determines inflation through the interest rate alone.

How does a central bank set an interest rate in a world without money? By standing ready to lend to and borrow from banks at rates it announces. Banks hold reserves at the central bank; no bank will lend overnight to another for less than the central bank pays on reserves, or borrow for more than the central bank charges. Market rates are therefore pinned between the two in the textbook corridor (in floor systems, institutions that cannot earn interest on reserves can trade slightly below the floor), and the central bank moves them by moving its own rates. The quantity of money is whatever households and banks choose to hold at those rates: it is determined by policy, not the other way round. Chapter 6 describes how the Fed, the ECB and the BCB implement this in practice, and what changes when the central bank holds a large portfolio of bonds.

> **In practice: money in the data**
>
> The decoupling of money from inflation is visible in the data. From 2008 to 2014 the monetary base in the United States grew roughly fivefold while core inflation stayed below 2%. In 2020 and 2021 broad money grew by about 40% and inflation then rose sharply. Neither episode is explained by money alone; both are explained, in the framework of this book, by the expected path of real rates relative to the

natural rate, by the fiscal response that accompanied the monetary one, and by the supply shocks of chapter 9.

1.5.3 The Fisher relation and the neo-Fisherian question

In a steady state with constant inflation π, the Euler equation (1.11) gives $i = \rho + \pi$: the nominal rate equals the real rate plus inflation. This is the ***Fisher relation***. Read from right to left, it says that a central bank that wants permanently higher inflation must eventually set a permanently higher interest rate.

Some economists have read it from left to right: if the bank raises the interest rate and keeps it there, inflation must eventually rise to match. In the perfect-foresight version of the core model, with the rate pegged at a higher level forever and expectations jumping to the new steady state, inflation rises immediately. This ***neo-Fisherian*** view implies that raising rates is inflationary, the opposite of the conventional wisdom.

The resolution again turns on determinacy. A permanent peg is indeterminate, so the perfect-foresight path is one of many, and nothing in the model says the public will coordinate on it. García-Schmidt and Woodford (2019) show that if expectations are formed by a process of reasoning about what others will do, rather than by jumping to the new steady state, a credible commitment to higher rates lowers inflation in the short run, as conventional analysis holds, and raises it only in the long run, as the Fisher relation requires. The two readings can coexist: which one governs the transition depends on how the new rate is implemented, on how expectations are formed and on fiscal backing.

Episode: cutting rates against inflation, Turkey 2021–22

From September 2021 the Central Bank of the Republic of Turkey cut its policy rate from 19% to 14% by December 2021, and to 9% by November 2022, while consumer price inflation rose to about 85% in October 2022. The cuts were justified publicly by the country's political leadership with an argument close to the neo-Fisherian one: lower rates would eventually bring lower inflation. In the model, a falling nominal rate with rising inflation is a collapsing real rate, the passive policy of the indeterminate region. The lira lost about half its value against the dollar over the period. The episode ended with a

Table 1.1. Baseline calibration, quarterly, as in Galí (2015, ch. 3).

Parameter	Value	Meaning and source
β	0.99	discount factor; steady-state real rate of 4% a year
σ	1	log utility; upper end of the household channel
φ	5	inverse Frisch elasticity of 0.2, as in micro labour studies
α	0.25	decreasing returns; labour share of three quarters
ε	9	elasticity of substitution; 12.5% markup
θ	0.75	price duration of one year
ϕ_π, ϕ_x	1.5, 0.125	Taylor (1993) rule (0.5 on the annual gap)
ρ_v, ρ_r, ρ_u	0.5	shock persistence
λ	0.0215	implied Calvo coefficient (1.26)
κ	0.172	implied Phillips curve slope (1.31)

sharp reversal from mid-2023, when the policy rate was raised in steps to 50%.

1.5.4 Promises about the far future

Equation (1.35) has one more implication that has become a research programme of its own. A promise to keep the real rate one point lower for a single quarter, ten years from now, has the same effect on today's output gap as the same cut today, because every future real-rate gap enters the sum with weight one. Through (1.32), its effect on inflation today is even larger the further away it is, because inflation accumulates the gaps in between. A central bank could, in this model, produce an enormous stimulus with a promise about 2036. No central bank believes that, and the data do not show it. McKay et al. (2016) call this the ***forward guidance puzzle*** and trace it to the assumptions that households can borrow freely against future income and that everyone hears and believes the promise. Chapter 14 measures how much distant guidance actually moves markets and spending.

1.6 Calibration

Table 1.1 lists the parameters. Each has a source outside the model.

Two choices deserve a comment. A discount factor of 0.99 implies a steady-state real rate of 4% a year, far above the estimates of chapter 7, which put the US natural rate near 1%. The responses of the model to shocks barely depend on it, but the lower bound on interest rates does,

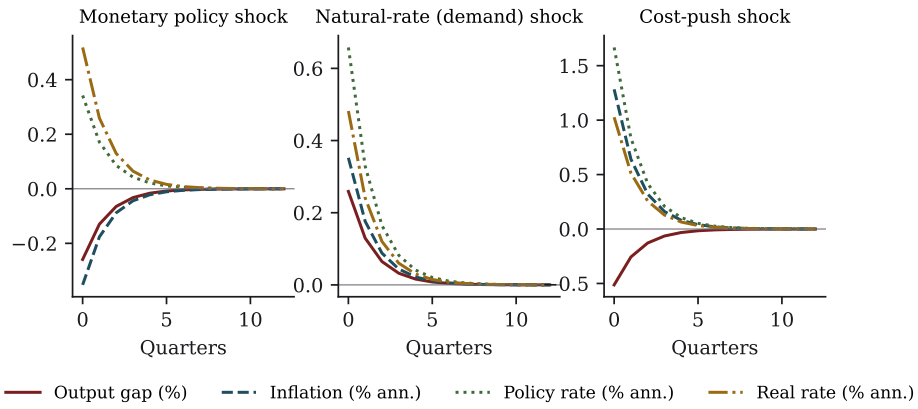


Figure 1.2. Impulse responses to a one-point (annualised) monetary, natural-rate and cost-push shock, persistence 0.5, baseline calibration and Taylor rule. Output gap in percent; inflation and rates in annualised percentage points.

which is why chapter 16 revisits it. And the price duration of a year is at the long end of the microdata, as discussed above; it stands in for sources of inertia the core model lacks.

1.7 What the model says about three shocks

Figure 1.2 feeds each shock into the calibrated model, with a size of one percentage point in annualised terms and a persistence of 0.5.

1.7.1 A monetary policy shock

From (1.44) with $\rho_v = 0.5$, $\Lambda_v = 2.05$. A shock of 0.25 in quarterly terms (one point annualised) lowers the gap by $0.505 \times 2.05 \times 0.25 = 0.26\%$ and quarterly inflation by $0.172 \times 2.05 \times 0.25 = 0.088$ points, 0.35 annualised. The policy rate is the shock plus the systematic response to these outcomes: $0.25 + 1.5 \times (-0.088) + 0.125 \times (-0.26) = 0.085$ quarterly, or 0.34 annualised. Expected inflation for next quarter is $0.5 \times (-0.088)$, so the real rate rises by $0.085 + 0.044 = 0.129$, or 0.52 annualised.

Result: the policy shock

A one-point surprise tightening lowers the output gap by 0.26% and annualised inflation by 0.35 points on impact. The nominal policy rate rises by only 0.34 points, a third of the shock, because the rule cuts the rate in response to the lower inflation and output the shock produces. The real rate rises by 0.52 points, more than the nominal rate, because expected inflation falls. Reading the size of a tightening from the change in the nominal rate understates it.

The model's response is immediate: output and inflation are at their lowest on impact and recover at the rate of the shock. In the data, the response is gradual. Christiano et al. (2005) estimate that after a policy shock output reaches its trough after about a year and a half and inflation responds even later. The core model has no source of delay: no habit in consumption, no adjustment cost in investment, no indexation of prices to past inflation. Chapter 3 adds price indexation and shows what the delay does to the choice of rule.

1.7.2 A natural-rate shock

A rise in the natural rate of one point annualised, with policy following the Taylor rule, raises the gap by 0.26% and inflation by 0.35 points. The nominal rate rises by 0.66 points and the real rate by 0.48: too little to match the one-point rise in r^n. A central bank that tracked r^n exactly would have left the gap and inflation at zero.

Result: the natural-rate shock

Under a Taylor rule, a shift in the natural rate is only partly offset, and the part that is not offset shows up as an output gap and inflation of the same sign. The same is true of an error in estimating r^n: a central bank that believes the natural rate is lower than it is sets policy too loose, and both inflation and output run above their targets. Chapter 7 shows that real-time estimates of r^* are revised by about half a point on average.

1.7.3 A cost-push shock

A cost-push shock of one point annualised raises inflation by 1.28 points and lowers the output gap by 0.51%. The policy rate rises by 1.67 points and the real rate by 1.03.

Result: the cost-push shock

Policy does not offset a cost-push shock; it splits it between inflation and output. The Taylor rule, with coefficients chosen for good behaviour after demand shocks, lets inflation absorb most of it. Whether that split is right depends on how the bank values inflation against output, which is the loss function of chapter 2.

The three results map the structure of the book. The policy shock says that the effect of policy is measured in real rates and expected paths, not in the overnight rate, which is the basis of the event studies in chapters 14 and 15. The natural-rate shock says the level of r^* is a first-order input, which is chapter 7. The cost-push shock says stabilisation involves a choice, which needs the loss function of chapter 2 and the rules and games of chapters 3 and 4.

1.8 What the core model leaves out

The model is a skeleton, and each missing piece is the subject of a later chapter.

Inertia. Both equations are purely forward-looking, so inflation has no memory and a credible change in policy moves it at once. The responses in Figure 1.2 peak on impact, against the delayed, hump-shaped responses in the data. Chapter 3 adds indexation of prices to past inflation and inertia in the policy rule.

Expectations. The model assumes everyone knows the model and forms rational expectations. Chapter 5 introduces information frictions and limited attention, and chapter 8 shows how far survey and market expectations depart from the rational benchmark, especially in Brazil.

Finance. Spending responds only to the real rate, through (1.35). There are no credit spreads, term premia, asset prices or bank balance sheets, so the model cannot say why the same policy path tightened financial conditions by different amounts in 2022 and in 2023. Chapters 12 and 13 supply the missing transmission.

The balance sheet. Money and the central bank's assets do not appear. In this model, and in Woodford (2003a), a purchase of long bonds changes nothing unless it changes expected policy rates. Chapter 6 shows which assumption breaks that result and what the purchases of 2008 to 2022 did.

The open economy. The model has no exchange rate, no imported goods and no foreign interest rate. For the Fed this is a reasonable first approximation. For the BCB it is not: the exchange rate passes through to prices within months, and the sovereign risk premium moves the real rate the economy faces. The Brazilian chapters bring these in where they decide the result.

The three equations stay throughout. Each chapter adds the piece the data say is missing, derives what it changes, and measures it.

Chapter 2

The Loss Function

Twelve-month IPCA inflation was 4.56% in January 2025, the first month of Brazil's continuous inflation target. The target was 3%, with a tolerance band reaching 4.5%. Inflation then rose to 5.06% in February, 5.48% in March and 5.53% in April, and was still 5.35% in June. Under the rules that took effect that year, six consecutive months outside the band count as a miss, and in July 2025 the Governor of the Banco Central do Brasil published a letter to the Minister of Finance explaining why inflation had left the band, what the bank was doing about it and when it expected inflation to return. It was the first letter under the new regime and the ninth since inflation targeting began in 1999, after the misses of 2001, 2002, 2003, 2015, 2017, 2021, 2022 and 2024.

The letter is a penalty. It costs the bank nothing in money and a great deal in reputation, and it is triggered by a specific event: inflation outside a band for a specific length of time. A central bank that faces such a penalty behaves differently from one that is simply asked to keep inflation close to target, and both behave differently from a bank whose mandate gives equal standing to employment. To analyse those differences the book needs a precise statement of what each bank is trying to avoid and how much it dislikes each kind of miss. That statement is the ***loss function***.

Chapter 1 ended with a cost-push shock that no interest-rate path could fully offset: policy had to split it between higher inflation and a negative output gap. A Taylor rule makes that split implicitly, through its coefficients. This chapter makes it explicit. It first derives the loss function from the welfare of the households in the model, which gives a benchmark with no free parameters. It then writes the mandates of the Fed, the ECB and the BCB as loss functions and compares them with the benchmark. It solves for the policy that minimises a loss when the bank cannot commit, which produces the trade-off frontier every later chapter uses. And it closes with losses that are not quadratic: bands, asymmetries and discrete penalties like the BCB's letter.

2.1 Why inflation is costly in the model

2.1.1 Price dispersion

In the model of chapter 1, inflation has no cost through money, because there is no money. Its cost comes from sticky prices. When the price level rises, the firms that reset charge the new, higher price, while the firms that do not reset keep their old, lower one. Goods that are equally costly to produce now sell at different prices. Households, who substitute between goods with elasticity ε, buy too much of the cheap goods and too little of the expensive ones. Firms producing the cheap goods hire more hours than efficiency requires, and firms producing the expensive goods fewer. The economy produces less useful output from the same total hours.

The loss can be measured precisely. Aggregate hours are the sum of hours across firms. Using each firm's production function and its demand curve,

$$N_t = \int_0^1 N_t(j)\, dj = \left(\frac{Y_t}{A_t}\right)^{\frac{1}{1-\alpha}} \int_0^1 \left(\frac{P_t(j)}{P_t}\right)^{-\frac{\varepsilon}{1-\alpha}} dj. \tag{2.1}$$

The integral is the ***price dispersion*** index. It equals one when all prices are equal and exceeds one otherwise. In logs, to second order,

$$\begin{aligned} n_t &= \frac{y_t - a_t + d_t}{1-\alpha}, \\ d_t &\equiv (1-\alpha)\log \int_0^1 \left(\frac{P_t(j)}{P_t}\right)^{-\frac{\varepsilon}{1-\alpha}} dj \approx \frac{\varepsilon}{2\Theta}\, \mathrm{var}_j\{p_t(j)\}, \end{aligned} \tag{2.2}$$

where $\Theta = (1-\alpha)/(1-\alpha+\alpha\varepsilon)$ is the strategic complementarity term of chapter 1. The second-order expansion uses the price-index restriction $\int (P_t(j)/P_t)^{1-\varepsilon} dj = 1$, which removes the first-order term; defining d_t with the factor $1-\alpha$ in front, as Galí (2015) does, leaves the hours equation with $d_t/(1-\alpha)$. Price dispersion raises the hours needed to produce a given output, and it does so in proportion to the cross-section variance of prices, scaled by how strongly households react to relative prices.

2.1.2 How Calvo pricing turns inflation into dispersion

Under Calvo pricing the cross-section variance of prices evolves with inflation. A fraction θ of firms keeps last period's prices, so it keeps last period's dispersion; a fraction $1-\theta$ moves to the common reset price, which sits

away from the average by an amount proportional to inflation. Working this through, the dispersion obeys

$$\operatorname{var}_j\{p_t(j)\} = \theta \operatorname{var}_j\{p_{t-1}(j)\} + \frac{\theta}{1-\theta}\,\pi_t^2. \tag{2.3}$$

Each period's inflation adds to dispersion, and dispersion decays at the rate θ as firms reset. Summing over time with discounting,

$$\sum_{t=0}^{\infty} \beta^t \operatorname{var}_j\{p_t(j)\} = \frac{\theta}{(1-\theta)(1-\beta\theta)} \sum_{t=0}^{\infty} \beta^t \pi_t^2 \tag{2.4}$$

up to terms independent of policy. The discounted sum of dispersion is proportional to the discounted sum of squared inflation. Inflation is costly in proportion to its square because dispersion is a variance.

2.1.3 The welfare loss

To turn dispersion into utility, approximate the household's period utility (1.1) to second order around the steady state. Two assumptions keep the result clean. First, the steady state is efficient: a subsidy to firms offsets the monopoly markup, so that natural output is also the efficient level of output. Second, shocks are small enough that a quadratic approximation is accurate. Then the utility of consumption contributes a term in y_t and y_t^2, the disutility of hours a term in n_t and n_t^2, and substituting (2.2) for n_t brings in the dispersion term. The linear terms cancel because the steady state is efficient, and the quadratic terms collect into two: one in the squared output gap, with coefficient $\sigma + (\varphi+\alpha)/(1-\alpha)$, the curvature of utility in output; and one in the cross-section variance of prices. Using the dispersion sum above to replace the variance with squared inflation gives the expected welfare loss, as a fraction of steady-state consumption,

$$\mathbb{W} = \frac{1}{2}\,\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\left(\sigma + \frac{\varphi+\alpha}{1-\alpha} \right) x_t^2 + \frac{\varepsilon}{\lambda}\,\pi_t^2 \right], \tag{2.5}$$

where λ is the Calvo coefficient of (1.26). The constant ε/λ combines the two steps: ε/Θ from how much dispersion costs, and $\theta/[(1-\theta)(1-\beta\theta)]$ from how much dispersion inflation creates.

Dividing by ε/λ and using the definition of κ normalises the loss to the form used throughout the book:

$$L = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\pi_t^2 + \lambda_x\, x_t^2 \right), \qquad \lambda_x = \frac{\kappa}{\varepsilon}. \tag{2.6}$$

2.1.4 Reading the welfare weight

The weight on the output gap is not a preference of the central bank. It is a property of the economy: the slope of the Phillips curve divided by the elasticity of demand. Each part has a reason.

A flatter Phillips curve lowers λ_x. When prices respond little to slack, a given output gap moves inflation little, so stabilising output does little to stabilise inflation, and the distortion that matters most, price dispersion, is the one inflation creates. A central bank in an economy with a flat Phillips curve should care relatively more about inflation, which is the opposite of the popular intuition that a flat curve makes inflation unimportant.

A higher elasticity ε also lowers λ_x. When goods are close substitutes, a given price gap moves demand a lot, so dispersion is more costly and inflation more damaging.

Result: the welfare weight

With the baseline calibration, $\lambda_x = 0.019$ with inflation measured as a quarterly rate. Mandates are written in annual rates, and with inflation in annualised points the weight becomes $16\lambda_x = 0.31$. One point of annual inflation volatility costs as much welfare as 1.8 points of output-gap volatility. In levels, a standard deviation of one point of annualised inflation costs $\frac{1}{2}(\varepsilon/\lambda)(0.0025)^2 = 0.13\%$ of steady-state consumption each quarter: a permanent loss of 0.13% of consumption, which is also 0.13% of a year's consumption.

That last number is large. It is about two and a half times the cost of all business-cycle fluctuations in consumption that Lucas famously estimated at about 0.05% of consumption. The difference comes entirely from price dispersion: Calvo pricing with an elasticity of 9 makes every point of inflation variability a misallocation of labour across firms. Whether the misallocation is that large in the data depends on whether firms' relative prices really drift apart when inflation rises, which the microdata support only partly. The book uses the welfare weight as a benchmark, not as a fact, and asks of each mandate how far it sits from it.

Three features of (2.6) matter for what follows. It is quadratic: large misses are disproportionately costly, and misses of either sign are equally bad. It is centred on zero inflation, because the model has no reason for a positive target; chapter 16 returns to the reasons central banks give for 2 or 3% (measurement bias in price indexes, room above the lower bound,

downward rigidity of nominal wages). And it is written in terms of the welfare-relevant output gap, an unobserved variable, which is why mandates refer to observable proxies such as unemployment.

2.2 Mandates as loss functions

No central bank publishes (2.6). Each has a legal mandate, written by a legislature, and a strategy, written by itself. Both can be represented by a loss function, and the representation shows where each bank departs from the benchmark. The loss functions below are reduced-form representations of mandates, not statements the banks have made about their objectives.

2.2.1 The Fed: a dual mandate

Since the 1977 amendment to the Federal Reserve Act, the Fed has had the goals of maximum employment, stable prices and moderate long-term interest rates; the third is usually treated as a consequence of the first two. In January 2012 the FOMC published its first Statement on Longer-Run Goals and Monetary Policy Strategy, defining price stability as 2% inflation in the PCE price index. It declined to fix a number for maximum employment, on the grounds that it depends on structural features of the labour market that change over time.

A natural reading of the dual mandate is

$$L^{\text{Fed}} = \mathbb{E}_t \sum_{k=0}^{\infty} \beta^k \left[(\pi_{t+k} - 2)^2 + \lambda_u \left(u_{t+k} - u^*_{t+k} \right)^2 \right], \tag{2.7}$$

with u^* the unemployment rate consistent with maximum employment. Okun's law, $u - u^* \approx -0.5\,x$, converts an unemployment weight into an output-gap weight: $\lambda_x = \lambda_u / 4$. A symmetric reading with $\lambda_u = 1$, often used in the Fed's own policy simulations, implies $\lambda_x = 0.25$ in annual units, close to the welfare benchmark of 0.31.

In August 2020 the FOMC revised its strategy in two ways that changed the shape of (2.7). It made the employment term one-sided: policy would respond to ***shortfalls*** of employment from its maximum, not to deviations in both directions. And it adopted ***flexible average inflation targeting***: after periods of inflation below 2%, it would aim for inflation moderately above 2% for some time, so that inflation averaged 2% over time. In loss-function terms, the first change sets $\lambda_u = 0$ whenever $u < u^*$, and the second replaces $\pi_t - 2$ with a moving average of past inflation minus 2. In August

2025, after the inflation of 2021 to 2023, the FOMC revised the statement again, removing both the reference to shortfalls and the averaging language and returning to a symmetric reading of the mandate.

Episode: the Fed in 2021

PCE inflation was 2.6% in March 2021 and 6.0% in December 2021. The FOMC did not raise rates until March 2022. Under the 2020 framework, inflation above 2% after a decade below it was, at first, part of the plan, and the strong labour market was not, by itself, a reason to tighten: the framework said the Committee would not tighten simply because unemployment fell below an estimate of its natural rate, although inflation and its risks still mattered. Read through the loss function, the framework had lowered the weight on overshooting at exactly the moment the economy delivered a large cost-push shock. The 2025 revision is, in part, a judgement that the asymmetry was costly.

2.2.2 The ECB: a hierarchical mandate

Article 127 of the Treaty on the Functioning of the European Union makes price stability the primary objective of the ECB. Without prejudice to it, the ECB supports the general economic policies of the Union, which include full employment and balanced growth. The mandate is ***hierarchical***: price stability first, everything else only to the extent it does not conflict.

The ECB has defined price stability three times. In 1998, as inflation in the harmonised index "below 2%". In 2003, as "below, but close to, 2% over the medium term". And in July 2021, as 2% over the medium term, with deviations in either direction "equally undesirable". The first two definitions are asymmetric: they penalise overshooting more than undershooting. Section 2.5 shows what such an asymmetry does, and the euro area data of 2013 to 2019 show it in practice.

A hierarchical mandate is lexicographic in principle: minimise any secondary loss subject to meeting the inflation objective. In practice the phrase "over the medium term" does the work of λ_x. A bank that returns inflation to target over three years rather than one chooses a path along which output deviations are smaller and inflation deviations last longer. The horizon is a weight in disguise, and section 2.4 makes the equivalence explicit.

Episode: the ECB in 2011

In April and July 2011, with headline HICP inflation at 2.7 to 2.8% driven largely by energy, the ECB raised its main refinancing rate twice, from 1.0% to 1.5%. The euro area sovereign debt crisis was then intensifying. By November and December 2011 both increases had been reversed. Under a dual mandate with a meaningful weight on activity, a cost-push shock from energy calls for looking through the first-round effect on headline inflation; under a strict reading of a hierarchical mandate with an asymmetric definition, inflation above 2% is a miss to be corrected. The 2011 hikes are usually cited as the clearest case in which the second reading prevailed and was quickly judged wrong.

2.2.3 The BCB: a continuous target with a band and a penalty

Brazil adopted inflation targeting in 1999. The National Monetary Council sets the target for the IPCA; the BCB is responsible for meeting it. Until 2024 the target applied to calendar-year inflation, with a tolerance band (1.5 points since 2017); from January 2025 it is ***continuous***: 3%, with a band of 1.5 points on either side, assessed every month on twelve-month inflation. Inflation outside the band for six consecutive months is a miss and triggers the letter. Complementary Law 179 of 2021, which made the BCB formally autonomous, gave it the objective of price stability and, without prejudice to it, of smoothing fluctuations in economic activity and fostering full employment.

Figure 2.1 shows how often the band binds. Measured month by month, as the continuous-target rule now measures it, twelve-month inflation was outside the band in 108 of the 261 months from January 2005 to September 2026, 41% of the time. The misses cluster in three episodes: above the ceiling in 2015 and 2016, during the fiscal crisis; below the floor in 2017 and 2018, after the deep recession; and above the ceiling from 2021 to 2022 and again from late 2024. Under the calendar-year regime only December counted, so many of these months carried no formal penalty. Under the continuous regime, the first six-month run, January to June 2025, produced the letter that opens this chapter.

This mandate has two parts that pull in different directions. The first is a quadratic in the distance from 3%, the loss that guides each Copom

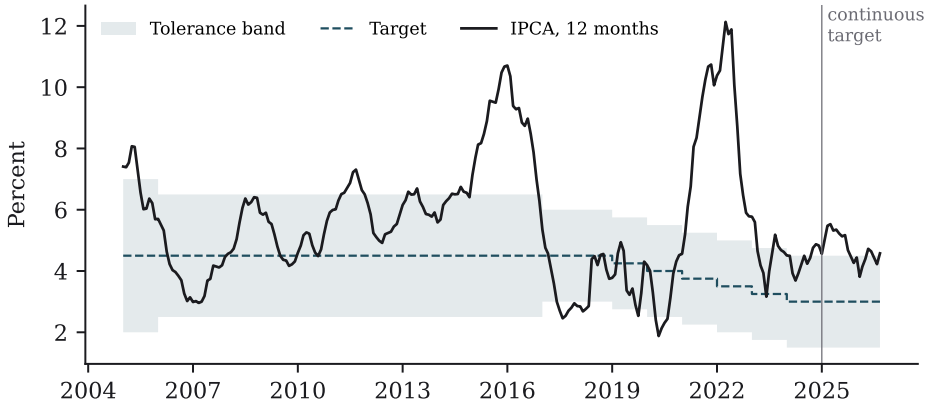


Figure 2.1. Twelve-month IPCA inflation against the target and tolerance band, January 2005 to September 2026. Targets applied to calendar-year inflation until 2024 and continuously from 2025. Source: IBGE via BCB SGS series 433; targets from the National Monetary Council.

meeting. The second is a discrete penalty, the letter, triggered only when inflation stays outside the band. Section 2.5 shows that the second part makes the bank care about the whole distribution of its inflation forecast, not just its centre.

In practice: what the mandate is about

The three banks also differ in the variable they react to. The Fed and the ECB target a measure of current inflation over a horizon and respond, in practice, to their own forecasts and to core inflation. The BCB, in its decisions and its communication, responds to inflation expectations and its own projections at the ***relevant horizon***, about six quarters ahead. Chapter 8 shows why: in Brazil, the Focus survey predicts headline inflation better than any core measure. Expected inflation enters the reaction functions of chapter 3, and the reaction functions are estimated in chapter 10.

2.3 Optimal policy under discretion

2.3.1 The bank's problem

Take the loss (2.6) and the Phillips curve with a cost-push shock (1.36). Suppose the central bank chooses x_t and π_t each period, taking the public's expectations of the future as given. This is ***discretion***: the bank re-optimises every period and cannot bind its future self. Because expectations of future inflation do not depend on what the bank does today, the problem separates into a sequence of static problems:

$$\min_{x_t,\pi_t} \pi_t^2 + \lambda_x x_t^2 \quad \text{subject to} \quad \pi_t = \kappa x_t + f_t, \qquad f_t \equiv \beta \mathbb{E}_t \pi_{t+1} + u_t. \tag{2.8}$$

The bank chooses a point on a straight line in the (x, π) plane, the Phillips curve for today, whose position f_t it cannot affect.

2.3.2 The targeting rule

Substitute the constraint into the objective and differentiate with respect to x_t: $2\kappa(\kappa x_t + f_t) + 2\lambda_x x_t = 0$, that is, $\kappa\pi_t + \lambda_x x_t = 0$. The optimal policy satisfies

$$x_t = -\frac{\kappa}{\lambda_x}\,\pi_t. \tag{2.9}$$

This is a ***targeting rule***: a condition on outcomes, not on the instrument. When inflation is above target the bank opens a negative output gap, and by more the steeper the Phillips curve (each point of gap buys more disinflation) and the less it cares about output. With the welfare weight, $\kappa/\lambda_x = \varepsilon = 9$: the bank should accept a gap of -9% for every point of quarterly inflation, or -2.25% for every point of annualised inflation. The interest rate is whatever the IS curve (1.33) requires to deliver that gap.

2.3.3 The equilibrium

Substitute (2.9) into the Phillips curve and guess that inflation is proportional to the shock, $\pi_t = \psi u_t$, with u_t following an AR(1) with persistence ρ_u:

$$\psi u_t = \beta\rho_u \psi u_t - \frac{\kappa^2}{\lambda_x}\psi u_t + u_t. \tag{2.10}$$

Solving for ψ gives the equilibrium

$$\pi_t = \frac{\lambda_x}{\kappa^2 + \lambda_x(1-\beta\rho_u)}\,u_t, \qquad x_t = -\frac{\kappa}{\kappa^2 + \lambda_x(1-\beta\rho_u)}\,u_t. \tag{2.11}$$

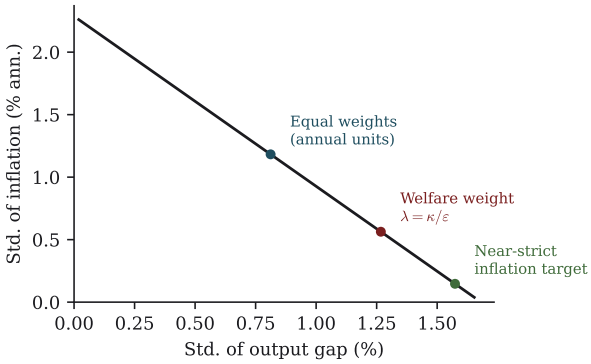


Figure 2.2. Inflation-output variance frontier under optimal discretion, cost-push shock with persistence 0.5 and innovation standard deviation of one point annualised. Points mark three weights on the output gap.

As $\lambda_x \to 0$, strict inflation targeting, inflation is fully stabilised and the gap absorbs the whole shock, $x_t = -u_t/\kappa$. As $\lambda_x \to \infty$ the gap is closed and inflation absorbs the shock, $\pi_t = u_t/(1 - \beta\rho_u)$. Every weight in between selects a point on a frontier.

The interest rate that implements this policy follows from the IS curve. Inflation and the gap are both proportional to u_t, so their expected values next period are ρ_u times their current values. Substituting into (1.33),

$$i_t = r_t^n + \rho_u \pi_t + \sigma(1-\rho_u)\frac{\kappa}{\lambda_x}\pi_t = r_t^n + \left[\rho_u + \frac{\sigma(1-\rho_u)\kappa}{\lambda_x}\right]\pi_t. \tag{2.12}$$

The optimal discretionary policy can be written as an instrument rule that tracks the natural rate and responds to inflation with a coefficient that is large when the weight on output is small. With the welfare weight and $\rho_u = 0.5$ the coefficient is $0.5 + 0.5 \times 9 = 5.0$, far above Taylor's 1.5. This is one reason estimated Taylor rules are not optimal rules: the optimal response to cost-push inflation, in this model, is much more aggressive than the response central banks show.

2.3.4 The variance frontier

Figure 2.2 traces the standard deviations of annualised inflation and the output gap implied by (2.11) as λ_x varies, for a cost-push shock with persistence 0.5 and an innovation of one point annualised.

Result: the frontier

At the welfare weight, the standard deviation of annualised inflation is 0.56 points and that of the output gap 1.27 points. Equal weights in annual units ($\lambda_x = 1/16$ quarterly) double inflation volatility to 1.18 points and cut gap volatility by a third, to 0.81. A near-strict inflation target ($\lambda_x = 0.004$) holds inflation volatility at 0.15 points at the cost of a gap volatility of 1.57. Moving from the welfare point to equal weights trades 0.62 points of inflation volatility for 0.46 points of output volatility.

The frontier is the menu; the loss function picks from it. Taylor-rule coefficients also pick a point, but not necessarily one on the frontier: a rule can be inefficient, delivering more volatility of both variables than some alternative. Chapter 10 places the estimated rules of the Fed, the ECB and the BCB relative to frontiers built from their own data.

The frontier also shows why the choice of weight is political. Moving along it does not make the economy better or worse on average; it redistributes volatility between price stability and employment. Workers whose jobs depend on the cycle and savers whose wealth depends on prices value the two differently. A legislature that writes a mandate is choosing a point on this curve, and the hierarchical and dual mandates are two different choices.

2.4 The horizon as a weight

Central banks rarely speak of weights. They speak of horizons: the ECB returns inflation to target "over the medium term", the BCB at the "relevant horizon", the Fed "over time". Under discretion with a persistent shock, inflation and the gap inherit the persistence of the shock whatever the weight, so the horizon does not appear. Under commitment, studied in chapter 3, it does. The committed bank keeps the gap open after the shock has gone, with a persistence δ that depends on the weight on output: the more the bank cares about output, the more gently it closes the gap, and the longer inflation takes to return.

Table 2.1 gives the mapping. A bank with the welfare weight closes half the gap in less than a quarter; one with equal weights in about a quarter; one with a weight thirteen times the welfare weight in two. A statement that inflation will return to target "over the medium term" rather than

Table 2.1. Weight on output and the speed of return under optimal commitment, baseline calibration. δ is the persistence of the output gap after a one-period cost-push shock; the half-life is in quarters.

Weight	λ_x (quarterly)	δ	Half-life
Near-strict inflation target	0.004	0.11	0.3
Welfare weight	0.019	0.31	0.6
Equal weights (annual units)	0.063	0.51	1.0
High weight on output	0.250	0.71	2.1

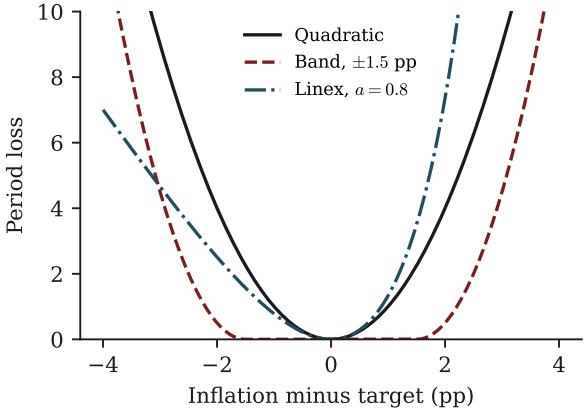


Figure 2.3. Period loss as a function of the inflation miss. Quadratic; zero inside a band of 1.5 points and quadratic outside it; and linex with asymmetry $a = 0.8$, which penalises overshooting more than undershooting.

"promptly" is, in this model, a statement that the weight on output is high. The ECB's horizon language and the Fed's dual mandate can produce the same policy.

2.5 Bands, asymmetries and penalties

The quadratic loss is a local approximation around a target. It is a poor description of a bank that faces a tolerance band, a public letter or a lower bound on the interest rate. Three alternatives cover most of what matters. Figure 2.3 draws them.

2.5.1 The band

A loss that is zero inside $\pm b$ and quadratic outside describes a bank judged only on whether it stays in the band:

$$\ell(d) = \max\{|d| - b,\, 0\}^2. \tag{2.13}$$

Taken literally, it implies inaction inside the band: any miss smaller than b costs nothing. It never applies literally, because the bank controls the expected miss, not the realised one. If the realised miss is $d \sim N(m, s^2)$, with m the bank's forecast and s the forecast uncertainty, the expected loss

$$\mathbb{E}\,\ell(d) = \int \max\{|d| - b, 0\}^2\, \phi\!\left(\frac{d-m}{s}\right) \frac{dd}{s} \tag{2.14}$$

is smooth in m even though ℓ has a flat bottom, and positive even when $m = 0$. Its slope in m is close to zero when the forecast sits in the middle of the band and its distribution is narrow, and steepest when the forecast is near an edge and the distribution wide.

A band therefore changes the timing of policy more than its average stance. The bank does little while expected inflation is comfortably inside the band and its uncertainty is low, and responds strongly when either the forecast drifts toward an edge or uncertainty rises. That is the pattern of a bank that "waits and sees" in calm periods and moves fast once expectations start to drift.

2.5.2 The discrete penalty

The BCB's letter is a fixed cost c paid when inflation leaves the band and stays out. In a one-period version, the expected penalty is

$$c\,\Pr(|d| > b) = c\left[\Phi\!\left(\frac{-b-m}{s}\right) + 1 - \Phi\!\left(\frac{b-m}{s}\right)\right], \tag{2.15}$$

which depends on the spread of the forecast distribution as much as on its centre. With $b = 1.5$, a forecast exactly at target and a forecast standard deviation of one point gives a 13% probability of a miss. Halving the uncertainty to half a point cuts the probability to 0.3%, a reduction by a factor of almost fifty, without moving the forecast at all. With the forecast one point above target, the probability is 32% at a standard deviation of one and 16% at half a point.

Result: the value of certainty

A bank facing a discrete penalty values reducing the uncertainty of its forecast for its own sake. The cheapest way to do so is not always to move the rate: anything that narrows the distribution of expected inflation, including credible communication, lowers the expected penalty. Part IV measures how far statements and press conferences narrow that distribution.

The six-month rule of the continuous target makes the penalty depend on the path, not on a single month. A shock that pushes inflation out of the band briefly costs nothing; one that keeps it out for half a year costs the letter. The rule therefore penalises persistent misses more than temporary ones, the same distinction the BCB draws when it says it looks through first-round effects of supply shocks but responds to second-round effects.

2.5.3 Asymmetry: the linex loss

A bank that fears overshooting more than undershooting is described by the ***linex*** loss (Zellner, 1986; Nobay and Peel, 2003), normalised so that it reduces to the quadratic as the asymmetry parameter a goes to zero:

$$\ell(d) = \frac{2}{a^2}\left(e^{ad} - ad - 1\right). \tag{2.16}$$

For $a > 0$ the loss rises exponentially above target and only linearly below it.

Suppose again that the miss is normal, $d \sim N(m, s^2)$, and the bank chooses m. The expected value of the exponential of a normal variable is $\mathbb{E}\, e^{ad} = e^{am + a^2 s^2/2}$, so the expected loss is

$$\mathbb{E}\,\ell = \frac{2}{a^2}\left(e^{am + a^2 s^2/2} - am - 1\right). \tag{2.17}$$

Differentiating with respect to m and setting the result to zero gives $e^{am + a^2 s^2/2} = 1$, that is,

$$m^{*} = -\frac{a\, s^2}{2}. \tag{2.18}$$

An asymmetric bank aims below target on average, and by more when uncertainty is high. With $a = 0.8$ and a forecast standard deviation of one point, the optimal expected miss is -0.4 points. Reversing the sign of a reverses the bias: a bank close to the lower bound on interest rates, which

fears undershooting because it has little room to respond to it, aims above target.

Episode: the euro area undershoot, 2013–2019

From January 2013 to December 2019 euro area HICP inflation averaged 0.96%, and it was below 1% in 40 of the 84 months, under a definition of price stability, "below, but close to, 2%", that treated overshooting and undershooting differently. The SPF expectation of inflation five years ahead drifted down to 1.67% by the end of 2019 (chapter 8). A linex loss with a positive asymmetry predicts exactly this: a systematic undershoot that grows with uncertainty. The ECB's 2021 strategy review replaced the definition with a symmetric 2% target and an explicit statement that deviations in either direction are equally undesirable, which in the language of this section sets $a = 0$.

2.6 Which inflation? The open economy

2.6.1 Domestic and imported prices

The loss (2.6) was derived for a closed economy, in which every good consumed is produced at home by firms with sticky prices. Brazil imports a substantial share of what it consumes, and the prices of imported goods and of commodities move with the exchange rate. Consumer price inflation is then a weighted average,

$$\pi_t^{CPI} = (1-\omega)\,\pi_t^{H} + \omega\,\pi_t^{F}, \tag{2.19}$$

of inflation in domestically produced goods π^H and in imported goods π^F, where ω is the import share. With complete and immediate pass-through, $\pi_t^F = \pi_t^* + \Delta e_t$, foreign inflation plus the change in the exchange rate.

Galí and Monacelli (2005) extend the core model to a small open economy and derive the welfare loss. Two results matter. Under their assumptions the loss depends on *domestic* inflation, not on consumer price inflation: price dispersion arises only among goods whose prices are sticky, and those are the domestically produced ones. And the optimal policy stabilises domestic inflation and lets the exchange rate float freely, absorbing foreign shocks through the terms of trade. A depreciation that raises import prices is, in that model, an efficient change in relative prices that policy should not resist.

2.6.2 Why the BCB targets the IPCA anyway

The BCB targets consumer price inflation, and it responds to depreciations. Three features of the Brazilian economy explain the gap between the theory and the practice.

Pass-through is not limited to imports. Many domestic prices are indexed, formally or by custom, to past consumer price inflation: rents, school fees, a large share of administered prices. A depreciation that raises IPCA today raises domestic prices tomorrow, so it becomes a shock to π^H as well. The BCB's published model makes the size explicit: a permanent 10% depreciation raises four-quarter IPCA inflation by about 0.96 points at the peak, with free prices rising 0.72 points and administered prices 1.65 points (Banco Central do Brasil, 2024a).

Expectations are formed on the IPCA. The Focus survey, wage negotiations and contracts refer to consumer price inflation. A bank that let IPCA inflation run far from target while stabilising an index of domestic prices would see expectations drift, and through the forward-looking Phillips curve the drift would feed back into domestic inflation.

Accountability is on the IPCA. The target, the band and the letter are written on consumer prices. Whatever the welfare loss in a model, the penalty the bank faces is on the IPCA.

The practical result is a compromise. The BCB looks through the first-round effect of a depreciation on import prices, which the open-economy theory says it should, and responds to the second-round effects on indexed prices and expectations, which the Brazilian structure makes inevitable. In loss-function terms its target variable is consumer price inflation, but its reaction to exchange-rate shocks is smaller than to domestic demand shocks of the same size, because part of the exchange-rate effect is expected to reverse.

2.7 Why the target is positive

The loss (2.6) is centred on zero inflation, yet every central bank in this book targets a positive rate: 2% for the Fed and the ECB, 3% for the BCB. Three arguments justify a positive target, and one counts against it.

2.7.1 Measurement bias

Price indexes overstate true inflation, because they adjust slowly for new goods, quality improvements and substitution between goods and outlets.

The Boskin Commission estimated in 1996 that the US consumer price index overstated inflation by about 1.1 points a year; methodological changes since then have reduced the bias, but most estimates put it between a half and one point. A target of zero measured inflation would be a target of true deflation.

2.7.2 The lower bound

The nominal interest rate cannot fall much below zero, so the room a central bank has to cut in a recession is roughly the neutral real rate plus the inflation target. With a neutral real rate of 1% and a target of 2%, the average policy rate is 3%, and a recession that calls for more than three points of cuts hits the floor. The data show how often that has happened. Since January 1990 the federal funds rate has been at or below 0.25% in 109 of 441 months, a quarter of the time, all of them after 2008. The ECB's deposit rate was at or below zero for 122 consecutive months, from July 2012 to August 2022, and negative for 97 of them, reaching −0.5%. A higher target buys room above the floor.

2.7.3 Downward wage rigidity

Nominal wages rarely fall. Akerlof et al. (1996) argue that at very low inflation, firms that need to cut real wages can do so only by cutting nominal wages, which they resist, so employment adjusts instead. A little inflation "greases the wheels" of the labour market by allowing real wage cuts without nominal cuts.

2.7.4 The cost: trend inflation and price dispersion

Against these benefits stands the cost the welfare derivation identified: inflation creates price dispersion. With a positive trend inflation, firms that do not reset their price fall steadily behind, and dispersion has a positive steady-state level that grows faster than linearly with the trend. Ascari (2004) shows that under Calvo pricing this cost is large and rises steeply.

Result: the cost of trend inflation

With constant returns to labour, trend inflation of 2% costs 0.16% of output each year through price dispersion, 3% costs 0.41% and 4% costs 0.82% (Table 2.2). The cost roughly quadruples each time trend

Table 2.2. Steady-state output loss from price dispersion at different annual trend inflation rates, Calvo pricing with $\theta = 0.75$ and $\varepsilon = 9$, under constant returns to labour ($\alpha = 0$) and under the chapter 1 calibration ($\alpha = 0.25$). Loss is d of (2.2) in steady state, percent of output.

Trend inflation (% a year)	0	1	2	3	4	6	8
Output loss, $\alpha = 0$ (%)	0.00	0.04	0.16	0.41	0.82	2.40	5.94
Output loss, $\alpha = 0.25$ (%)	0.00	0.15	0.68	1.75	3.61	11.71	35.47

inflation doubles, and at 8% it reaches almost 6% of output. With the decreasing returns of the chapter 1 calibration, $\alpha = 0.25$, a firm whose price is too low must hire more than proportionally to meet its demand, and the cost is four times larger: 0.68% at 2%, 1.75% at 3%. Moving the BCB's target from 3% to 2% would save between a quarter of a point and a point of output a year in this model, at the cost of less room above the lower bound and less grease for real wage adjustment. In practice the BCB's target is higher than the Fed's and the ECB's; one rationale is that Brazilian indexation makes each point of disinflation costlier, and the National Monetary Council has lowered it in steps, from 4.5% until 2018 to 3% from 2024, as indexation has weakened.

The numbers depend heavily on Calvo pricing. Under state-dependent pricing, where firms change prices more often when inflation is higher, dispersion grows more slowly and the cost of trend inflation is smaller. The microdata of chapter 1, in which the frequency of price changes rose sharply in 2021 and 2022, suggest that the truth lies between the two.

2.8 Beyond inflation and the gap

Practical loss functions carry terms that the welfare derivation does not produce. Four matter for the rest of the book.

Interest-rate smoothing. Many analyses add a term $\mu\,(\Delta i_t)^2$ to the loss, which penalises large changes in the policy rate and produces the gradual moves central banks make. Chapter 3 shows that part of the gradualism is optimal even without such a term, because inertia in the rate builds the history dependence that commitment requires.

Financial stability. A bank that cares about credit spreads, asset prices or the solvency of banks adds a term in some measure of financial conditions. Whether monetary policy should respond to financial imbalances or leave them to macroprudential policy is one of the open questions of the last fifteen years; chapter 13 measures how financial conditions transmit policy, which is the first step in answering it.

The exchange rate. For an open economy, exchange-rate volatility passes through to prices and to the balance sheets of firms with foreign-currency debt. The BCB does not target the exchange rate, but its reaction function responds to the real through expected inflation, and its foreign-exchange interventions are a second instrument with its own implicit loss.

History. The 2020 Fed framework, and price-level targeting more generally, replace the inflation term with one that depends on the past: the average of inflation over a window, or the distance of the price level from a path. Chapter 3 shows that optimal commitment in the core model is a form of price-level targeting, which gives such losses a theoretical foundation, and chapter 16 asks when they are worth their cost in credibility.

The loss function settles what the bank wants. Two questions remain: whether it can commit to the policy that minimises its loss, and whether the public believes it does. Chapter 3 answers the first with rules and chapter 4 the second with games.

Chapter 3

Reaction Functions and Rules

In the fourth quarter of 2021, core PCE inflation in the United States was 4.8% over the previous year and output stood about 2% above the Congressional Budget Office's estimate of potential. The rule John Taylor had proposed in 1993 as a description of good Fed behaviour prescribed a federal funds rate of 9.2%. The actual rate was 0.08%. Never since the rule was written had the gap between its prescription and policy been so wide. A year later the Fed had raised rates to 3.7%, and by the third quarter of 2023 to 5.3%; by then the rule, with core inflation at 3.9%, prescribed 7.6%. Through 2024 policy and the rule came back within about a point of each other.

Was the Fed wrong in 2021, or was the rule? The question has no answer without a theory of what rules are for. This chapter supplies one. It starts where chapter 2 left off, with a bank that minimises a loss function period by period, and asks what it could achieve if it could also make promises about the future. The answer, commitment, beats discretion, and the gap between them explains why central banks care so much about being predictable. It then turns to the instrument rules central banks are judged against: the Taylor rule and its family, inertial rules, and rules on expected rather than current inflation. For the last of these the chapter uses the reaction function the BCB itself publishes in its semi-structural model, and tests it alongside the textbook rules. It ends with how a reaction function is estimated from data and why the estimates are harder to interpret than they look.

3.1 Two kinds of rule

An ***instrument rule*** sets the policy rate as a function of observable variables, as the Taylor rule (1.37) does. It tells a committee what to do with given data. A ***targeting rule*** states a condition that the projected path of

the target variables must satisfy, as (2.9) does, and leaves the interest rate to be whatever achieves it. It tells a committee what the forecast should look like after the decision.

The distinction has practical consequences. An instrument rule is simple to state, simple to verify and simple to communicate: anyone can compute the Taylor rule from published data. But it embeds a model of the economy in its coefficients, and when the economy changes the coefficients become wrong. A targeting rule adapts automatically to a changing economy, because the bank uses its current model to find the rate that delivers the required forecast. But it is harder to verify, since the public must trust the bank's forecast. Inflation-targeting central banks describe themselves in the second language. The BCB's statements explain each decision by reference to its inflation projection at the relevant horizon and ask whether that projection is converging to the target.

Svensson (1999) argued that inflation targeting is best understood as a targeting rule, and that the instrument rules estimated on central bank behaviour are descriptions, not prescriptions. Woodford (2003a) agreed on the value of targeting rules and added a reason to care about the form of the rule that goes beyond its period-by-period outcomes: a rule is also a promise about the future, and promises move expectations. That is the subject of the next section.

3.2 Commitment

3.2.1 The bank's problem with promises

Return to the loss (2.6) and the Phillips curve with a cost-push shock (1.36). Under discretion, chapter 2 found that the bank treats expected future inflation as given, because it cannot influence what its future self will do. Now suppose the bank can choose, once and for all, a complete plan for $\{x_t, \pi_t\}$ in every future period and every state, and that the public believes the plan. The bank's problem is

$$\min_{\{x_t,\pi_t\}} \; \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \tfrac{1}{2} \left(\pi_t^2 + \lambda_x x_t^2\right) \quad \text{s.t.} \quad \pi_t = \beta\, \mathbb{E}_t \pi_{t+1} + \kappa x_t + u_t. \tag{3.1}$$

The difference from discretion is that $\mathbb{E}_t \pi_{t+1}$ is no longer given: it is part of the plan. A promise about inflation tomorrow changes inflation today.

3.2.2 Deriving the optimal plan

Attach a multiplier ζ_t to each period's Phillips curve. The Lagrangian is

$$\mathcal{L} = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\tfrac{1}{2} \left(\pi_t^2 + \lambda_x x_t^2 \right) + \zeta_t \left(\pi_t - \kappa x_t - \beta \pi_{t+1} - u_t \right) \right]. \tag{3.2}$$

Inflation in period t appears twice: in period t's own constraint, with coefficient ζ_t, and in period $t-1$'s constraint, as next period's inflation, with coefficient $-\beta\zeta_{t-1}$, which the extra discount factor of the earlier period turns into $-\zeta_{t-1}$. The first-order conditions are

$$\pi_t + \zeta_t - \zeta_{t-1} = 0, \tag{3.3}$$

$$\lambda_x x_t - \kappa \zeta_t = 0, \tag{3.4}$$

for $t \geq 1$, and for $t = 0$ the same with $\zeta_{-1} = 0$, because no promise was made before the plan began.

From (3.4), $\zeta_t = (\lambda_x/\kappa)\, x_t$. Substituting into (3.3) eliminates the multiplier:

$$\pi_t = -\frac{\lambda_x}{\kappa} \left(x_t - x_{t-1} \right). \tag{3.5}$$

This is the ***timeless*** targeting rule. The name, due to Woodford (2003a), reflects the fact that the rule applies in every period, including the first, if one imagines the plan was adopted long ago. At $t = 0$ the strict optimum would set $x_{-1} = 0$, treating the start of the plan as special; the timeless perspective uses the actual x_{-1}, which makes the rule the same in every period and removes the temptation to treat "now" as a fresh start.

3.2.3 What the rule says

Under discretion, inflation is linked to the level of the output gap: $\pi_t = -(\lambda_x/\kappa)\, x_t$. Under commitment it is linked to the change: $\pi_t = -(\lambda_x/\kappa)(x_t - x_{t-1})$. The committed bank promises that a negative gap opened today will be closed only gradually. Because the Phillips curve is forward-looking, the promised slack tomorrow lowers expected inflation, which lowers inflation today. The bank gets disinflation today from slack it will deliver later.

Summing (3.5) over time, and using $\pi_t = p_t - p_{t-1}$,

$$p_t = -\frac{\lambda_x}{\kappa}\, x_t + \text{constant}. \tag{3.6}$$

Optimal commitment in the core model is a target for the ***price level***, adjusted for the output gap. Once the gap closes, the price level returns to

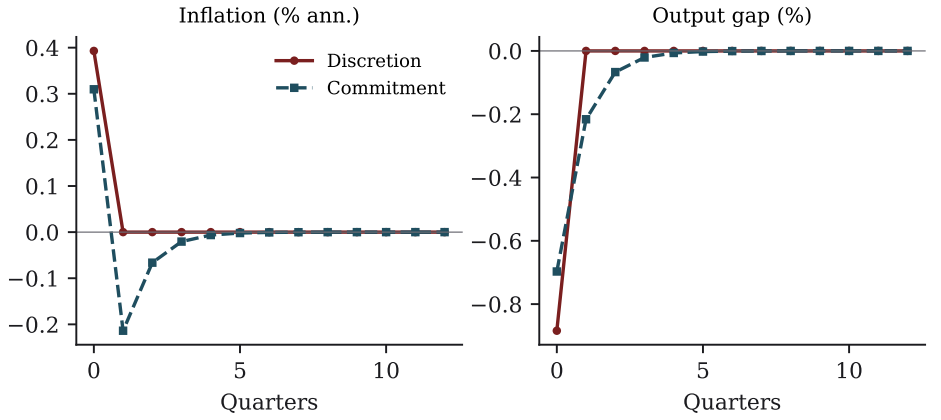


Figure 3.1. Responses to a one-period cost-push shock of one point (annualised) under optimal discretion and optimal commitment, welfare-based loss and baseline calibration.

where it would have been without the shock: inflation above target today must be followed by inflation below target later. This is history dependence in its purest form. What the bank does today depends on what happened to prices in the past, even after the shocks that moved them have vanished.

3.2.4 Solving for the path

Substituting (3.5) into the Phillips curve gives a second-order difference equation in the gap:

$$-\frac{\lambda_x}{\kappa}(x_t - x_{t-1}) = -\beta\frac{\lambda_x}{\kappa}\,\mathbb{E}_t(x_{t+1} - x_t) + \kappa x_t + u_t. \tag{3.7}$$

Its stable solution has the form $x_t = \delta\, x_{t-1} + \gamma\, u_t$. Substituting and matching the coefficients on x_{t-1} gives a quadratic in δ whose root inside the unit circle is

$$\delta = \frac{1 - \sqrt{1 - 4\beta a^2}}{2a\beta}, \qquad a \equiv \frac{\lambda_x}{\lambda_x(1+\beta) + \kappa^2}, \tag{3.8}$$

and matching the coefficients on u_t gives $\gamma = -\kappa\delta/[\lambda_x(1 - \beta\delta\rho_u)]$. The gap inherits a persistence δ even when the shock has none. That inherited persistence is the promise. Its size depends on the weight on output, as Table 2.1 showed: $\delta = 0.31$ at the welfare weight, 0.51 with equal weights.

Result: the gain from commitment

After a one-period cost-push shock of one point, the discretionary bank lets inflation rise 0.39 points and opens a gap of 0.88%, and both return to zero the next quarter. The committed bank lets inflation rise only 0.31 points with a smaller gap of 0.70%, keeps the gap open with persistence 0.31, and lets inflation fall slightly below target in the following quarters, so that the price level returns toward its path (Figure 3.1). The welfare loss is 21% lower under commitment. The difference is the ***stabilisation bias*** of discretion: unable to promise future restraint, the bank must deliver more restraint now.

3.2.5 Why commitment is hard

The committed plan contains a promise the bank will not want to keep. Once the shock has passed, inflation is at or below target and the gap is still negative. Re-optimising at that moment, a discretionary bank would close the gap at once, because nothing calls for slack any more. But closing it at once breaks the promise that made the initial disinflation cheap, and if the public anticipates that, the promise never had any effect. This is the ***time inconsistency*** of optimal policy identified by Kydland and Prescott (1977): the plan that is optimal at the start is not optimal to continue.

Chapter 4 asks how a bank that cannot sign a binding contract can still earn the credibility commitment requires. Two partial answers come from the rules themselves. One is interest-rate smoothing: Woodford (2003c) shows that inertia in the policy rate builds history dependence into policy and reproduces part of the commitment gain without an explicit promise. The other is price-level or average-inflation targeting, which writes the history dependence of (3.6) directly into the mandate. The Fed's 2020 framework was an attempt at the second; its withdrawal in 2025 is a lesson in how hard it is to make a make-up promise credible when the shock that tests it is large and the promise is new.

3.3 The Taylor rule and its family

3.3.1 The original rule

Taylor (1993) proposed that the Fed's behaviour from 1987 to 1992 was well described by

$$i_t = r^* + \pi_t + 0.5\,(\pi_t - \pi^*) + 0.5\,y_t, \tag{3.9}$$

with the natural real rate r^* and the inflation target π^* both 2%, π_t inflation over the previous four quarters and y_t the percentage output gap. The coefficient on inflation is 1.5, so the rule satisfies the Taylor principle (1.39): a one-point rise in inflation raises the real rate by half a point. At target and with a closed gap, the rule sets the nominal rate at 4%.

The rule was not derived from a model. Taylor chose round coefficients that described the data and argued that they were sensible: large enough on inflation to anchor it, moderate on output to avoid fine-tuning. That the coefficients satisfy determinacy in the New Keynesian model, and perform well across many models in simulation, is what made the rule the benchmark of the field. The Fed itself now publishes, in each semi-annual Monetary Policy Report, the prescriptions of a set of simple rules of this family next to the actual path of the funds rate.

Episode: the Fed against the rule, 2021–24

Applying (3.9) to current-vintage data on core PCE inflation and the CBO output gap gives a prescription of 3.4% in the first quarter of 2021, 9.2% in the fourth quarter and 9.8% in the first quarter of 2022, when the funds rate averaged 0.12%. The rate rose to 3.7% by the end of 2022 and 5.3% by the third quarter of 2023, while the prescription fell to 7.6% as core inflation came down. In 2024 the two were within about a point of each other: 5.3% against 6.4% in the first quarter, 4.7% against 6.1% in the fourth. The largest deviation in the rule's history was concentrated in 2021 and early 2022, the period of the 2020 framework and of the judgement that inflation was transitory.

Figure 3.2 extends the comparison over sixteen years. Three regimes stand out. From 2010 to 2015 policy sat at the lower bound while the original rule prescribed 1.5 to 3%; the balanced-approach rule, with its larger weight on a still-negative output gap, prescribed rates close to zero, and the Fed's own reading of slack was closer to it. From 2016 to 2019 the Fed stayed below the rules: in the fourth quarter of 2015, when it first raised rates, the Taylor prescription was 2.2% against an actual 0.16%, and in the fourth quarter of 2019 it was 4.0% against 1.64%, a gap the Fed attributed to a lower neutral rate than Taylor's 2%. And from 2021 the gap widened to nine points before closing in 2023 and 2024. In the second quarter of 2026 it has reopened: the funds rate averaged 3.6% against a Taylor prescription of 6.3%, with core inflation running around 3%. Whether that gap reflects a lower neutral rate, a different weight on activity or a judgement that the

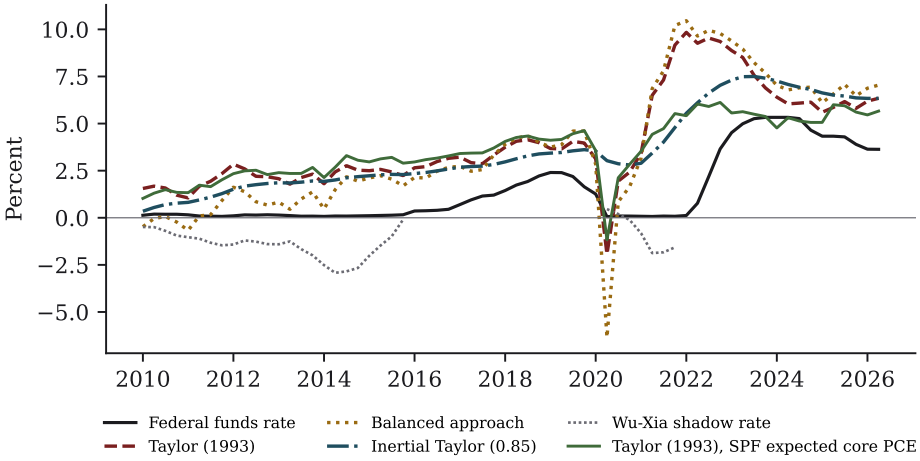


Figure 3.2. Federal funds rate, the Wu-Xia shadow rate at the lower bound, and simple-rule prescriptions, 2010Q1 to 2026Q2. The three Monetary Policy Report rules use current core PCE inflation and the CBO output gap in the current vintage; the inertial rule starts from the actual rate in 2010Q1 and closes 15% of the gap to the Taylor prescription each quarter. The solid green line applies the Taylor (1993) coefficients to the SPF median forecast of core PCE inflation over the next four quarters. Sources: BEA, CBO, Federal Reserve and Federal Reserve Bank of Philadelphia via FRED; Federal Reserve Bank of Atlanta.

inflation is transitory is exactly the question of interpretation the rest of this chapter equips the reader to ask.

Two features of the figure change the picture once the rule is read as a central bank reads it. The first is the lower bound. With the funds rate at zero, the Fed eased through asset purchases, and the shadow rate of Wu and Xia (2016), which reads the stance off the whole yield curve, fell to −2.9% in 2014 and below −1.5% in 2021: the gap between rule and stance in those years was smaller, by two to three points, than the funds rate suggests. The second is the inflation variable. Current core inflation includes the transitory part that policy is meant to look through; the Fed responds to its forecast. Applied to the SPF forecast of core PCE inflation over the next four quarters, the Taylor (1993) coefficients prescribe 5.5% in the fourth quarter of 2021 rather than 9.2%, and 5.7% in the second quarter of 2026. The rule on expectations is smoother, sits closer to what the Fed did once it moved, and still shows the Fed five to seven points behind in late 2021, measured against the funds rate and the shadow rate respectively. Chapter 10 estimates the Fed's rule on the Greenbook forecasts the committee actually saw.

Table 3.1. Federal funds rate and the prescriptions of four published simple rules, quarterly averages. Core PCE inflation over four quarters, CBO natural rate of unemployment, and the FOMC's median longer-run funds rate minus 2% as r^{LR}. Source: BEA, BLS, CBO and Federal Reserve via FRED.

Quarter	Funds rate	Core PCE	Taylor	Balanced	Shortfalls	First diff.
2021Q4	0.08	4.75	6.90	7.18	6.63	4.06
2022Q4	3.65	5.29	8.30	9.17	7.43	4.43
2023Q4	5.33	3.41	5.24	5.87	4.61	5.72
2024Q4	4.65	2.89	4.62	4.90	4.34	5.36
2025Q4	3.90	2.77	4.10	4.05	4.05	4.35
2026Q2	3.63	3.10	4.87	5.00	4.74	4.11

3.3.2 The rules the Fed publishes

Since 2017 the Fed's semi-annual Monetary Policy Report has shown the prescriptions of a set of simple rules. They are written in terms of the unemployment gap rather than the output gap, use core PCE inflation, and take the neutral real rate from the committee's own longer-run projections rather than fixing it at 2%. With r^{LR} the longer-run real rate, u^{LR} the longer-run unemployment rate and Okun's law converting the output-gap coefficient of 0.5 into an unemployment-gap coefficient of 1, the main ones are:

$$\text{Taylor (1993):} \quad i = r^{LR} + \pi + 0.5\,(\pi - 2) + (u^{LR} - u), \tag{3.10}$$

$$\text{Balanced approach:} \quad i = r^{LR} + \pi + 0.5\,(\pi - 2) + 2\,(u^{LR} - u), \tag{3.11}$$

$$\text{Balanced, shortfalls:} \quad i = r^{LR} + \pi + 0.5\,(\pi - 2) + 2\,\min\{u^{LR} - u,\, 0\}, \tag{3.12}$$

$$\text{First difference:} \quad i = i_{-1} + 0.5\,(\pi - 2) + (u^{LR} - u) - (u^{LR} - u)_{-4}. \tag{3.13}$$

The shortfalls rule responds only when unemployment is above its longer-run level, the asymmetry of the 2020 framework. The first-difference rule needs no r^{LR}.

Result: the measure of slack and the neutral rate decide the verdict

With the unemployment gap and the committee's own neutral rate, the rules put policy in 2026Q2 at 1.1 to 1.4 points below their prescriptions, against 2.7 points with the CBO output gap and a neutral

rate of 2%. In 2023Q4 and 2024Q4 the funds rate was within about half a point of every level rule and above the Taylor and shortfalls prescriptions (Table 3.1): by this measure, policy was at or slightly above the rules through 2024, which the output-gap version does not show. In 2021Q4 every level rule, including the shortfalls rule written for the 2020 framework, prescribed at least 6.6%; the first-difference rule, which moves from wherever the rate stands, prescribed 4.1%. The judgement that the Fed was far behind the curve in late 2021 does not depend on the choice of rule; the judgement about 2024 to 2026 does.

3.3.3 Four variants

The balanced-approach rule doubles the gap coefficient to 1.0. It describes a bank that weighs employment more heavily, and in 2021 and 2022 it prescribed even higher rates than the original, because the gap was positive.

The inertial rule is

$$i_t = \rho_i \, i_{t-1} + (1 - \rho_i) \, i_t^{\mathrm{T}}, \tag{3.14}$$

where i_t^{T} is the Taylor prescription. The rate closes a fraction $1 - \rho_i$ of the distance to the prescription each quarter. With $\rho_i = 0.85$ and starting from the actual rate at the end of 2020, the inertial rule prescribes 4.0% at the end of 2021 and 6.6% at the end of 2022: still far above actual policy in 2021, but much closer to it than the original rule.

The first-difference rule sets the change in the rate rather than its level:

$$\Delta i_t = \phi_\pi \, (\pi_t - \pi^*) + \phi_x \, \Delta x_t. \tag{3.15}$$

It needs no estimate of r^*, which drops out when the rule is written in differences. Chapter 7 shows that estimates of r^* are revised by about half a point on average, and that error enters every level rule one for one. It is related to the inertial rule, but not as its simple limit: writing (3.14) as $\Delta i_t = (1 - \rho_i)(i_t^{\mathrm{T}} - i_{t-1})$ shows that the change goes to zero as $\rho_i \to 1$ with the coefficients of i_t^{T} fixed. Only if the response coefficients are rescaled so that $(1 - \rho_i)\phi_\pi$ and $(1 - \rho_i)\phi_x$ stay constant does the limit become a rule in differences, and then on the level of the gap rather than its change. What the two rules share is history dependence.

The forecast-based rule replaces current inflation with its expected value at some horizon. It is the form closest to how the BCB describes its policy, and the next section studies it using the BCB's own published rule.

3.3.4 Interest-rate smoothing as history dependence

Why do central banks move gradually? One reason is uncertainty: a bank unsure of the effects of its actions moves cautiously. Another is the cost of reversals to its credibility. Woodford (2003c) adds a reason rooted in commitment. A rule of the form (3.14) makes today's rate depend on yesterday's, so a tightening today implies, through the rule, a higher rate tomorrow even if conditions return to normal. The public, understanding the rule, expects tight policy to persist, and that expectation does part of the work of the committed plan. Inertia is a promise written into the rule.

3.4 The BCB's published reaction function

3.4.1 How the BCB describes its policy

The Copom sets the Selic so that its projection of inflation at the ***relevant horizon***, about six quarters ahead, converges to the target, and it reports in each statement the projection in its reference scenario. It also watches the Focus survey of expectations as a measure of how credible that convergence is. The rationale has two parts. Policy acts with a lag of several quarters, so inflation over the next few months is already determined. And much of current inflation in Brazil comes from administered prices, food and the exchange rate, shocks that often reverse and that policy should look through unless they spread to expectations.

3.4.2 The rule in the BCB's model

The BCB publishes its small semi-structural projection models in boxes of its inflation reports, with every equation and the estimated parameters. The June 2024 update (Banco Central do Brasil, 2024a), estimated by Bayesian methods on quarterly data from 2003 to 2023, writes the reaction function as

$$i_t = \theta_1\, i_{t-1} + \theta_2\, i_{t-2} + (1-\theta_1-\theta_2)\left[\overline{rr}_t + \pi^{\text{meta}}_t + \theta_3\left(\pi^{e}_{t,t+4|t} - \pi^{\text{meta}}_t\right)\right] + \epsilon^{i}_t, \tag{3.16}$$

where $\pi^{e}_{t,t+4|t}$ is the Focus expectation of inflation over the next four quarters, π^{meta} the target and $\overline{rr}_t$ a time-varying equilibrium real rate. The posterior modes in the aggregate model are $\theta_1 = 1.48$, $\theta_2 = -0.58$ and $\theta_3 = 2.03$, with 90% credible intervals of $[1.41, 1.54]$, $[-0.63, -0.52]$ and $[1.47, 2.64]$. The rule has no output-gap term: the BCB reports that the coefficient was

not identified, and notes that this does not mean the bank ignores activity, since expectations move with the gap.

Three features of (3.16) deserve a reading.

Strong inertia. The two lags sum to 0.90, so in the long run a tenth of the target level enters the rate each quarter. Because the second lag is negative this is not a simple partial adjustment: the characteristic roots of $z^2 - 1.48z + 0.58$ are complex, $0.74 \pm 0.18i$, with modulus 0.76, so after a shock the rate approaches the path set by its target in damped waves rather than geometrically. The combined persistence is higher than the typical estimate for the Fed.

Momentum. The second lag is negative. Rewriting the inertial part, $1.48\, i_{t-1} - 0.58\, i_{t-2} = 0.90\, i_{t-1} + 0.58\, (i_{t-1} - i_{t-2})$: more than half of last quarter's change in the rate carries into this quarter's. The Selic moves in cycles of steps in the same direction, rather than jumping to a new level. This is the ***momentum*** visible in every BCB cycle, including the twelve consecutive increases of 2021 and 2022.

The Taylor principle on expectations. The response to the gap between expected inflation and the target is 2.03 in the long run. A one-point rise in expected inflation eventually raises the Selic by about two points, and the ex-ante real rate by about one. The rule satisfies the Taylor principle, written on expectations rather than on current inflation.

> **In practice: what the model says about transmission**
>
> The same box reports the model's responses. A Selic one point higher for four quarters, returning afterwards to the rule, lowers four-quarter IPCA inflation by 0.27 points at the peak, after four quarters, in the aggregate model, and by 0.24 in the disaggregated one. A permanent 10% depreciation of the real raises it by about 0.96 points at the peak, and a demand shock that raises the output gap by 1% raises it by about 0.49. A point of Selic for a year buys about a quarter of a point of inflation; a 10% move in the exchange rate costs almost a full point. These magnitudes explain why the exchange rate dominates the BCB's discussion of inflation risks.

3.4.3 Replicating the BCB's model from public information

Because the BCB publishes its equations and parameters, the model can be rebuilt by anyone with public data, which is a strong form of transparency. A replication built for this book's research workspace from the published boxes,

Table 3.2. Peak response of four-quarter IPCA inflation (percentage points) in the BCB's published small models and in a replication from public data. Source: Banco Central do Brasil (2024a); replication in the author's research workspace.

	Aggregate model		Disaggregated model	
Shock	BCB	Replication	BCB	Replication
Selic +1 point for four quarters	−0.27	−0.19	−0.24	−0.20
Output gap +1%	0.49	0.47	0.45	0.49
Exchange rate, 10% depreciation	0.96	0.53	0.87	0.35

Focus expectations, SGS series and IBGE national accounts reproduces two of the three benchmark responses closely.

The demand channel is reproduced within 0.04 points and the policy channel within 0.08 (Table 3.2). The exchange-rate channel is not: the replication delivers about half of the BCB's response. The missing half is the 24-equation model of administered prices, published separately, through which a depreciation raises fuel, electricity and regulated tariffs. In the BCB's own decomposition, administered prices contribute about 0.4 points of the 0.96. That gap is a quantitative statement of how much of Brazil's exchange-rate pass-through runs through regulated prices rather than market ones.

3.4.4 Forecast targeting and instrument rules on forecasts

There are two ways to write down a policy that responds to forecasts, and they behave differently. Svensson (1997) defines ***inflation forecast targeting*** as choosing the path of the interest rate so that the forecast of inflation satisfies a targeting rule at the horizon at which policy has its effect. The forecast is the target, not an input: the bank adjusts the rate until the forecast is where it should be. When transmission takes several quarters, this implements optimal policy in the class of models and objectives Svensson studies, and it is the policy the BCB describes in its statements.

The alternative feeds the forecast into a Taylor-type formula,

$$i_t = r^* + \pi^* + \phi_\pi \left(\mathbb{E}_t \pi_{t+h} - \pi^* \right) + \phi_x \, x_t, \tag{3.17}$$

which is what (3.16) looks like once the inertia is set aside. Bernanke and Woodford (1997) identified two weaknesses. The first is self-fulfilling equilibria: if the rate responds to forecasts and the forecasts respond to the rate, beliefs can justify themselves. In the core model with $h = 1$, checking the

eigenvalues of the system (1.38) gives the determinacy condition

$$1 < \phi_\pi < 1 + \frac{2(1+\beta)\sigma}{\kappa}. \tag{3.18}$$

The lower bound is the Taylor principle. The upper bound is new: a rule that responds too strongly to expected inflation creates oscillating self-fulfilling paths, in which expected high inflation produces a large rate increase, a deep recession and low inflation, then a large cut, and so on. With the baseline calibration the bound is 24, so it does not bind at $h = 1$, but the determinate region narrows as the horizon lengthens.

The second weakness matters more in practice. A model-consistent forecast of inflation several quarters ahead moves much less than current inflation, because shocks fade. A coefficient that is aggressive when applied to current inflation is weak when applied to such a forecast. Whether that matters for the BCB depends on what its expectations variable is, which the following experiment makes precise.

3.4.5 A model with inflation inertia

To test rules on forecasts, the core model needs inflation with a memory. The standard device, due to Christiano et al. (2005), is ***indexation***: firms that do not reset their price in a given period raise it automatically by a fraction γ of last period's inflation. Repeating the derivation of chapter 1 with this change, the reset price now anticipates the automatic increases the firm's price will receive while it is stuck, and the price index mixes reset prices with indexed old prices. Collecting terms gives the ***hybrid*** Phillips curve

$$\pi_t = \frac{\gamma}{1+\beta\gamma}\,\pi_{t-1} + \frac{\beta}{1+\beta\gamma}\,\mathbb{E}_t\pi_{t+1} + \frac{\kappa}{1+\beta\gamma}\,x_t + u_t. \tag{3.19}$$

With $\gamma = 0$ it is the curve of chapter 1; with $\gamma = 1$, inflation depends equally on its past and its expected future. With $\gamma = 0.5$, a third of last quarter's inflation carries into this quarter's directly. The BCB's own Phillips curve for free prices has a similar structure, with an inertia weight of 0.24 on lagged free-price inflation, 0.38 on lagged headline inflation and the rest on Focus expectations.

The model with (3.19) has lagged state variables and no closed-form solution. It is solved numerically by finding linear policy functions that express x_t, π_t and i_t as functions of the state (last period's inflation, the last two interest rates and the two shocks) and that satisfy the three equations for every value of the state.

Table 3.3. Rules compared in the model with inflation indexation ($\gamma = 0.5$). Standard deviations of annualised inflation and the output gap; welfare loss (3.20) relative to the Taylor rule on current inflation. Textbook rules respond to the gap with 0.5 (annual); the BCB rule has no gap term. Simulated, 200,000 quarters.

Rule	sd π	sd x	Rel. loss
Taylor, current inflation, $\phi_\pi = 1.5$	2.32	1.05	1.00
Taylor, current inflation, $\phi_\pi = 2.5$	1.54	1.43	0.61
Inertial Taylor, $\rho_i = 0.8$	1.85	0.99	0.85
Taylor, expected inflation 4 qtrs, $\phi_\pi = 1.5$	28.4	2.57	80
Taylor, expected inflation 4 qtrs, $\phi_\pi = 8$	2.84	1.15	1.28
BCB rule, expected inflation 4 qtrs	7.21	2.36	9.20
BCB rule, current inflation	1.30	1.28	0.60

3.4.6 Comparing rules

Table 3.3 compares seven rules in this model, with natural-rate and cost-push shocks of persistence 0.5 and equal size (one point annualised), on the welfare loss of chapter 2 re-derived for indexation. With indexation, firms that do not reset still raise their prices by $\gamma\pi_{t-1}$, so what creates dispersion is inflation relative to that indexed path. Repeating the steps of chapter 2 replaces π_t in (2.6) by its quasi-difference (Woodford, 2003a, ch. 6):

$$L = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[(\pi_t - \gamma\pi_{t-1})^2 + \lambda_x x_t^2 \right]. \tag{3.20}$$

A bank that lets inflation persist pays less in this loss than in (2.6), because persistent inflation is partly absorbed by indexation rather than turned into dispersion. The two BCB rows use the published coefficients of (3.16); the first applies them to model-consistent inflation four quarters ahead, the second to current inflation.

Result: rules on expected inflation

A Taylor rule that responds with the conventional 1.5 to model-consistent inflation four quarters ahead lets inflation volatility rise more than tenfold: the four-quarter forecast moves by only a fraction of current inflation once shocks fade. With the BCB's published inertia and momentum, the same weakness costs nine times the Taylor rule's loss. Applied to current inflation, the BCB's rule is the best of the

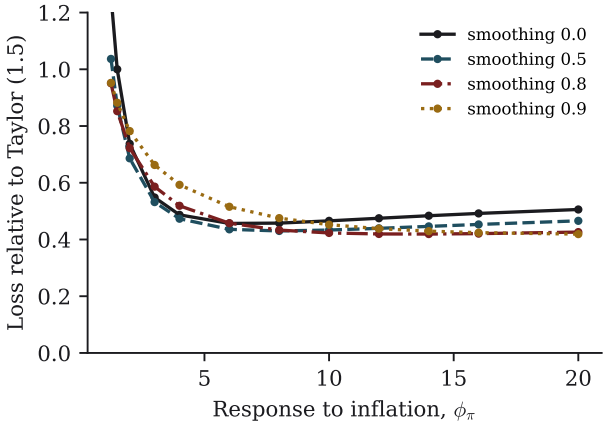


Figure 3.3. Welfare loss (3.20) relative to the Taylor rule ($\phi_\pi = 1.5$, $\phi_x = 0.125$, no smoothing) as a function of the response to inflation, for four degrees of interest-rate smoothing, gap response 0.125. Model with indexation $\gamma = 0.5$, natural-rate and cost-push shocks.

seven, with a loss 40% below the Taylor rule's, because its inertia and momentum supply the history dependence of the commitment result.

3.4.7 The optimal simple rule

The table compares a few rules. A systematic search asks which rule of the form (3.14), with a Taylor target $\phi_\pi \pi_t + \phi_x x_t$, minimises the welfare loss (3.20) in the model with indexation, over responses to inflation from 1.25 to 20, gap responses from 0 to 0.5 and smoothing from 0 to 0.9. For each combination of the three coefficients the model is solved, and the unconditional variances of inflation, the gap and the interest rate are computed exactly from the solution, by solving the discrete Lyapunov equation $\Sigma = T\Sigma T' + Q$ for the covariance of the state, rather than by simulation.

Result: the best simple rule

The loss falls steeply as the response to inflation rises from 1.25 to about 6, and is almost flat from 8 to 20 (Figure 3.3). The surface is flat near the minimum: smoothing of 0.8 with a response to inflation of 12 to 14, and smoothing of 0.9 with a response of 20, at the edge of the

grid, all give 42% of the Taylor rule's loss (0.419), and the gap response barely matters (0.421 of the Taylor loss with none, 0.419 with 0.25). Without smoothing the best attainable loss is 45% of the Taylor rule's. The larger gain from smoothing is in the policy rate itself: at the best rules its standard deviation falls from 3.9 points without smoothing to 2.4 with it, a cut of about 40%. The optimal rule responds aggressively to inflation, needs little response to the gap and moves gradually.

Three features of the result recur across the literature on simple rules. The response to inflation in the optimal rule is much larger than Taylor's, as the optimal discretionary rule of chapter 2 also implied. The gap term adds nothing once inflation is stabilised aggressively, because in this model a stable inflation rate already requires a small gap. And inertia earns its keep twice: through the history dependence of the commitment result, and through lower rate volatility, which the model does not penalise but central banks plainly do. The BCB's published rule, with inertia of 0.9, momentum, a long-run response of about 2 to expected inflation and no gap term, shares three of the four features; its response to inflation is smaller. One interpretation is that the BCB also cares about the volatility of its rate and about the exchange rate, which the loss (3.20) omits.

The two BCB rows bracket the truth, and the reason is the expectations variable. In this model, expectations are rational: the four-quarter forecast falls quickly toward the target as shocks fade, so a rule on it barely responds to a shock in progress. The Focus survey is not like that. The BCB's own expectations equation gives the lagged Focus expectation a weight of 0.75 and model-consistent expectations only 0.11: survey expectations are sticky and move with realised inflation. Chapter 8 shows that the Focus twelve-month expectation tracks headline inflation closely enough to be its best predictor a year ahead. A rule on such expectations behaves much more like the second BCB row than the first.

The lesson is about the rule's input, not its form. A forecast belongs in a targeting rule, where the bank keeps adjusting until the forecast is where it should be. As the right-hand variable of an instrument rule, model-consistent forecasts make policy too passive. Survey expectations, which carry the persistence of actual inflation and measure credibility at the same time, should make the rule behave much more like the current-inflation case, which in these simulations comes near the best simple rules (50% of the Taylor loss against 42%); the simulation uses current inflation as a stand-in

rather than an estimated law of motion for Focus. Chapter 10 estimates the BCB's rule on the data and asks which of the two it resembles.

3.5 Estimating a reaction function

3.5.1 The specification

The standard specification, from Clarida et al. (1999) and Clarida et al. (2000), combines inertia with a forward-looking target rate:

$$i_t = \rho_i \, i_{t-1} + (1-\rho_i)\left[\bar{r} + \pi^* + \phi_\pi \left(\mathbb{E}_t \pi_{t+h} - \pi^*\right) + \phi_x \, \mathbb{E}_t x_{t+q}\right] + \epsilon_t. \quad (3.21)$$

The bank has a target for the rate that depends on expected inflation h periods ahead and the expected gap q periods ahead, and it moves part of the way to the target each period. The BCB's rule (3.16) is this specification with two lags, Focus expectations in place of the econometrician's forecast and no gap term.

3.5.2 Estimation

The expectations in (3.21) are not observed. Clarida et al. (2000) replace them with realised values, π_{t+h} and x_{t+q}, and note that under rational expectations the forecast errors are uncorrelated with anything the bank knew at time t. Any variable in the bank's information set, such as lagged inflation, output and interest rates, is then a valid instrument, and the coefficients can be estimated by generalised method of moments: choose the parameters that make the residual, including the forecast errors, as close as possible to uncorrelated with the instruments. The BCB's approach is different: it uses an observed measure of expectations, the Focus survey, and estimates the whole model jointly by Bayesian methods, which avoids the rational-expectations assumption at the cost of tying the rule to one survey.

3.5.3 Real-time data

Orphanides (2001) reconstructed the data the Fed actually had at each meeting, and showed that the output gap in particular was revised heavily after the fact. In the 1970s, real-time estimates of potential output were far too optimistic, because the productivity slowdown that began in 1973 was recognised only years later. The Fed believed it faced a large negative gap when, by later estimates, the gap was small. A rule estimated on revised

data attributes to the Fed a disregard for inflation that, on the data it had, was partly a response to slack that did not exist.

The problem has not gone away. The BCB's June 2025 report (Banco Central do Brasil, 2025) lists twelve estimates of the Brazilian output gap for the first quarter of 2025, from seven univariate filters and five multivariate models. They range from 0.14% to 5.03%, with a median of 1.18%; the median estimate among market participants in the pre-Copom questionnaire was 1.2%. A reaction function with a gap term fitted on any one of them is fitted on a number its own producer regards as highly uncertain. This is one reason the BCB's estimated rule has no gap term.

3.5.4 Identification

A deeper problem was raised by Cochrane (2011). In a determinate equilibrium of the New Keynesian model, inflation responds to fundamental shocks in a way the rule itself dampens. The coefficient ϕ_π governs what the bank would do off the equilibrium path, if inflation deviated from the unique stable solution, which is never observed. On the equilibrium path, the observed co-movement of the rate and inflation reflects both the rule and the shocks, and a regression of one on the other recovers neither the rule nor anything structural unless the shocks are identified separately.

The practical consequence is modest but important. Estimated reaction functions are good descriptions of how a central bank has behaved and useful for forecasting what it will do. They are not reliable evidence of whether its behaviour satisfies the Taylor principle, which is a statement about off-equilibrium responses. Chapter 10 reads its estimates as behaviour and judges policy by its outcomes on the frontier of chapter 2.

Part I still needs two pieces. Chapter 4 asks why a bank would deviate from its own rule and how institutions can stop it; chapter 5 asks what the public learns from what the bank does and says.

Chapter 4

Policy as a Game

On 8 May 2024 the Copom cut the Selic by a quarter of a point, to 10.50%. The decision was split five to four. The five members who voted for the smaller cut had been appointed under the previous government; the four who voted for a half-point cut, including the director widely expected to become the next Governor, had been appointed by the sitting President, who had spent months criticising the level of interest rates in public. In the following days the yield curve steepened, the real weakened and inflation expectations for 2025 and 2026 in the Focus survey rose. The market had read the vote as information about the future committee, not just about the current decision. In June the committee held the rate unanimously, and from September it began raising it again, with the votes of all its members.

Nothing in chapters 1 to 3 explains that sequence. Those chapters treated the central bank as a single decision-maker with a known loss function and the public as a passive forecaster of what that decision-maker will do. In fact the public does not know the bank's preferences with certainty, the bank knows that the public is watching, and the bank is a committee whose membership changes. Policy is a game. Chapter 3 showed that commitment beats discretion by 21%; this chapter asks why, if commitment is so valuable, a bank cannot simply commit, and what institutions can do instead. It solves the classic game between a central bank and the public, then the institutional responses to it: delegation to a conservative banker, a contract with a penalty, reputation, and decision by committee. It closes with the game that matters most for an emerging market, between the central bank and the treasury.

4.1 The inflation bias

4.1.1 The model

The cleanest version of the problem is the static model of Kydland and Prescott (1977) and Barro and Gordon (1983a). It strips away the dynamics of the core model to isolate one thing: what happens when a central bank wants output above its natural level.

Output deviates from its natural level only through inflation surprises and a supply shock:

$$y = (\pi - \pi^e) + \varepsilon, \qquad \mathbb{E}\varepsilon = 0,\ \mathrm{var}(\varepsilon) = \sigma^2. \tag{4.1}$$

This is a Lucas supply curve. Only unexpected inflation raises output, because expected inflation is built into wages and prices in advance. It is the static counterpart of the Phillips curve (1.31), with expectations formed one period earlier.

The bank's loss is

$$L = \tfrac{1}{2}\left[\pi^2 + \lambda\,(y - k)^2\right], \qquad k > 0. \tag{4.2}$$

It dislikes inflation and dislikes output away from a target k above the natural level. Why would it target output above the natural level? Because the natural level is inefficiently low: monopoly power in goods and labour markets and distorting taxes keep employment below the level a planner would choose. This steady-state distortion is distinct from the cost-push shock of chapter 1: the first is a permanent wedge that makes natural output inefficiently low, the second a time-varying wedge that moves inflation and the gap in opposite directions and breaks the divine coincidence. A bank that cares about welfare wants more output than the economy produces on its own, and it has a tool, surprise inflation, that delivers more output in the short run.

4.1.2 The timing

The game has two moves. First, the public forms its expectation π^e, which is built into wage contracts and prices. Second, the bank observes the shock ε and chooses π. The public cannot react to the bank's choice until the next round.

4.1.3 Solving the game

Solve backwards. Given π^e and ε, the bank chooses π to minimise (4.2). Substituting (4.1),

$$L = \tfrac{1}{2}\left[\pi^2 + \lambda\left(\pi - \pi^e + \varepsilon - k\right)^2\right]. \tag{4.3}$$

The first-order condition is $\pi + \lambda(\pi - \pi^e + \varepsilon - k) = 0$, which gives the bank's ***best response***

$$\pi = \frac{\lambda}{1+\lambda}\left(\pi^e + k - \varepsilon\right). \tag{4.4}$$

The bank's inflation rises with expected inflation, because a higher π^e means a larger negative output surprise at any given π, which the bank partly offsets. It rises with the output ambition k. And it falls with the supply shock, because a favourable shock already pushes output toward k.

The public knows (4.4). With rational expectations it sets π^e equal to the expected value of the bank's choice. Taking expectations of (4.4), $\pi^e = \frac{\lambda}{1+\lambda}(\pi^e + k)$, which solves to $\pi^e = \lambda k$. Substituting back,

$$\pi^{\text{disc}} = \lambda k - \frac{\lambda}{1+\lambda}\,\varepsilon, \qquad y^{\text{disc}} = \frac{\varepsilon}{1+\lambda}. \tag{4.5}$$

Inflation is λk higher on average than zero, and output is no higher on average: the public anticipates the bank's temptation, so there is no surprise in expectation. The bank ends up with higher inflation and no extra output. This is the ***inflation bias*** of discretion. It does not come from bad intentions or poor information. It comes from a well-meaning bank that cannot promise not to exploit a short-run trade-off, facing a public that knows it.

4.1.4 Commitment and the cost of discretion

Suppose instead the bank could commit, before the public forms expectations, to a rule for π as a function of ε. The public would then expect whatever the rule implies, and the bank would gain nothing from average inflation, which cannot raise average output. The best rule sets average inflation to zero and keeps the stabilisation response: $\pi = -\frac{\lambda}{1+\lambda}\varepsilon$. Output is the same as under discretion; inflation is lower by λk in every state.

The expected losses can be computed directly. Under commitment, $\mathbb{E}L^{\text{comm}} = \frac{1}{2}\left[\frac{\lambda\sigma^2}{1+\lambda} + \lambda k^2\right]$. Under discretion the same terms appear plus the squared bias:

$$\mathbb{E}L^{\text{disc}} - \mathbb{E}L^{\text{comm}} = \tfrac{1}{2}\,\lambda^2 k^2. \tag{4.6}$$

The cost of discretion does not depend on the variance of shocks. It depends only on how much the bank cares about output and how ambitious its output target is. With $\lambda = 0.5$, $k = 1$ and $\sigma = 1$, the expected loss is 0.417 under commitment and 0.542 under discretion. The difference of 0.125 is pure inflation, with no gain in output.

4.1.5 The bias in the core model

The static model makes the argument transparent, but the same bias appears in the dynamic model of chapter 1. Suppose the bank's loss penalises deviations of the output gap from a target $x^* > 0$, reflecting the inefficiently low natural level of output: $\pi_t^2 + \lambda_x (x_t - x^*)^2$. Under discretion the targeting rule of chapter 2 becomes $\kappa \pi_t + \lambda_x (x_t - x^*) = 0$. In a steady state with constant inflation $\bar{\pi}$, the Phillips curve (1.31) gives $\bar{x} = (1 - \beta)\bar{\pi}/\kappa$. Combining the two,

$$\bar{\pi} = \frac{\lambda_x \, \kappa}{\lambda_x (1 - \beta) + \kappa^2} \, x^*. \tag{4.7}$$

With the welfare weight and the baseline calibration, an output ambition of 1% produces steady-state inflation of 0.11 points per quarter, 0.44 points a year, and raises output by only 0.006%: the long-run Phillips curve is almost vertical because $1 - \beta$ is close to zero. An ambition of 2% doubles the bias to 0.88 points a year. The mechanism is the same as in the static model: the public anticipates the bank's temptation to push output above its natural level, builds the anticipated inflation into prices, and the bank ends up with inflation and almost no extra output.

4.2 Institutional answers

If commitment is impossible because no one can enforce it, the solution must change the bank's incentives so that the discretionary choice is closer to the committed one. Three mechanisms have been proposed and, in some form, adopted.

4.2.1 The conservative central banker

Rogoff (1985) proposed that society appoint a central banker who cares less about output than society does, with weight $\lambda_d < \lambda$. The banker then chooses under discretion with its own weight, so the bias falls to $\lambda_d k$ and the response to supply shocks becomes $-\frac{\lambda_d}{1+\lambda_d}\varepsilon$. But society judges outcomes

with its own weight λ. Its expected loss as a function of the appointee's weight is

$$\mathbb{E}L(\lambda_d) = \tfrac{1}{2}\left[\lambda_d^2 k^2 + \left(\frac{\lambda_d}{1+\lambda_d}\right)^2 \sigma^2 + \lambda\left(\frac{\sigma^2}{(1+\lambda_d)^2} + k^2\right)\right]. \tag{4.8}$$

The first term is the bias, which falls as λ_d falls. The second and third are the volatility of inflation and output under the banker's stabilisation policy, which is distorted when $\lambda_d \neq \lambda$: a conservative banker lets output absorb too much of each supply shock. Differentiating with respect to λ_d and setting the result to zero gives a condition that balances the two effects. At $\lambda_d = \lambda$ the stabilisation terms are at their minimum, so a small reduction in λ_d costs nothing to first order in stabilisation and gains λk^2 per unit in bias: some conservatism is always optimal. At $\lambda_d = 0$ the bias is gone but the stabilisation loss is largest. The optimum lies in between. Writing out the derivative,

$$\frac{\partial \mathbb{E}L}{\partial \lambda_d} = \lambda_d k^2 + \frac{\sigma^2(\lambda_d - \lambda)}{(1+\lambda_d)^3} = 0 \quad \Longrightarrow \quad \lambda_d k^2 (1+\lambda_d)^3 = \sigma^2(\lambda - \lambda_d), \tag{4.9}$$

the left side is the marginal gain from a lower bias, the right side the marginal cost of worse stabilisation. With $\lambda = 0.5$, $k = 1$ and $\sigma = 1$ the two balance at $\lambda_d = 0.187$. The more volatile the supply shocks (σ larger), the closer the optimal banker is to society; the larger the output ambition (k larger), the more conservative.

Result: delegation

The optimal appointee has $\lambda_d = 0.19$, less than two fifths of society's weight. Delegation closes 67% of the gap between discretion and commitment (Figure 4.1). The remaining third is lost because the conservative banker stabilises supply shocks with the wrong weight. Delegation trades a bias for a distortion.

Rogoff's argument is the intellectual basis of central bank independence. It explains why legislatures choose to hand monetary policy to unelected officials with a narrow mandate, and why those officials are often chosen for a reputation for caring about inflation. Alesina and Summers (1993) found that across industrial countries, measures of legal central bank independence were negatively correlated with average inflation and uncorrelated with average growth: more independence, less inflation, no cost in output. That

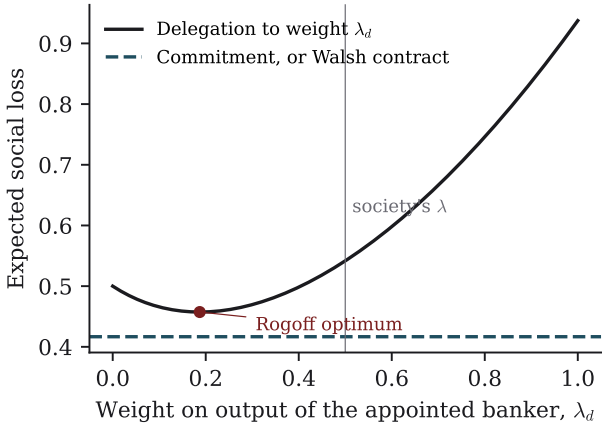


Figure 4.1. Expected social loss when policy is delegated to a banker with output weight λ_d, society's weight $\lambda = 0.5$, $k = 1$, $\sigma = 1$. The dashed line is the loss under commitment, which the Walsh contract also attains.

is the pattern (4.5) predicts if independence lowers the effective λ without changing the natural rate of output.

4.2.2 Goal and instrument independence

The literature distinguishes two kinds of independence. A bank has ***instrument independence*** if it sets the policy rate without government interference, and ***goal independence*** if it also chooses the target it pursues. The three banks of this book differ on the second. The Fed chose its own numerical target, 2% PCE inflation, in 2012, within a mandate set by Congress. The ECB defines price stability itself, within the Treaty's primary objective. The BCB has instrument independence but not goal independence: the National Monetary Council, chaired by the Minister of Finance, sets the target.

In Rogoff's terms, a government that sets the target can change the bank's effective weight on inflation without changing its law. The arrangement can strengthen credibility, by making the target a joint commitment of government and bank, or weaken it, if the public expects the government to raise the target when inflation is inconvenient.

Episode: the target debate of 2023

In early 2023, with the Selic at 13.75% and inflation falling, members of the new government argued publicly that the 3% target for 2024 and beyond was too low and should be raised. Long-run inflation expectations in the Focus survey, which had been at target through 2022, rose by more than half a point in the first months of the year (chapter 8). In June 2023 the National Monetary Council kept the target at 3% and adopted the continuous target from 2025. Expectations stopped rising but did not return to target. The episode is a direct measurement of the value of goal credibility: the mere possibility that the target would be changed moved long-run expectations as much as several years of inflation surprises had failed to.

4.2.3 The Walsh contract: a penalty that removes the bias

Walsh (1995) showed that the remaining third can be recovered by a penalty instead of a change of preferences. Suppose the principal (the government or the legislature) imposes on the banker a cost linear in inflation, so the banker minimises $L + c\,\pi$. The first-order condition becomes

$$\pi + \lambda\,(\pi - \pi^{e} + \varepsilon - k) + c = 0. \tag{4.10}$$

The constant c enters exactly where the output ambition $-\lambda k$ does, with the opposite sign. Setting $c = \lambda k$ cancels the temptation term, and the best response becomes $\pi = \frac{\lambda}{1+\lambda}(\pi^{e} - \varepsilon)$. Expectations are then $\pi^{e} = 0$ and $\pi = -\frac{\lambda}{1+\lambda}\varepsilon$: the commitment outcome, achieved under discretion.

The penalty must be linear, not quadratic, because its job is to offset a temptation that is the same at the margin in every state. A quadratic penalty would change how the bank responds to shocks and reintroduce a distortion. And the banker keeps society's weight λ, so stabilisation is undistorted.

In practice: penalties that exist

No legislature pays a central banker by the inflation rate. Accountability does the work instead. New Zealand's 1989 Reserve Bank Act came closest to Walsh's design: a contract between the Governor and the Minister of Finance specifying an inflation target, with dismissal possible for failing to meet it. The BCB's public letter, triggered by a

sustained miss (chapter 2), is a discrete penalty; the Bank of England's open letter is another. The Fed's semi-annual testimony to Congress is a continuous, reputational one. A discrete penalty is weaker than Walsh's linear one inside the band and stronger at its edge, which is the behaviour section 2.5 derived.

4.2.4 Reputation

In a repeated game, the public can punish. Barro and Gordon (1983b) consider a public that expects low inflation as long as the bank has delivered it, and reverts to the discretionary expectation λk forever after the first deviation. The bank then compares two numbers each period. To keep things simple, drop the supply shock.

Temptation. If the public expects the announced inflation π^a and the bank cheats, it chooses its best response (4.4) with $\pi^e = \pi^a$: $\pi^c = \frac{\lambda}{1+\lambda}(\pi^a + k)$. The one-period gain from cheating is the loss from keeping the promise minus the loss from breaking it.

Enforcement. After cheating, the bank is stuck with the discretionary outcome forever. The discounted cost is $\frac{\delta}{1-\delta}\left(L^{\text{disc}} - L^{\pi^a}\right)$, where δ is the bank's discount factor and L^{π^a} the per-period loss when the promise is kept.

Low inflation is sustainable when temptation does not exceed enforcement. Take the most ambitious promise, $\pi^a = 0$. Keeping it gives $L^0 = \frac{1}{2}\lambda k^2$. Cheating sets $\pi^c = \frac{\lambda k}{1+\lambda}$ and gives $L^c = \frac{1}{2}\frac{\lambda k^2}{1+\lambda}$, so the temptation is

$$L^0 - L^c = \tfrac{1}{2}\,\frac{\lambda^2 k^2}{1+\lambda}. \tag{4.11}$$

The discretionary outcome gives $L^{\text{disc}} = \frac{1}{2}(\lambda^2 k^2 + \lambda k^2)$, so the per-period cost of punishment is $\frac{1}{2}\lambda^2 k^2$ and the enforcement is $\frac{\delta}{1-\delta}\frac{1}{2}\lambda^2 k^2$. Comparing the two,

$$\frac{1}{1+\lambda} \le \frac{\delta}{1-\delta} \quad \Longleftrightarrow \quad \delta \ge \frac{1}{2+\lambda}. \tag{4.12}$$

Result: sustainable credibility

With $\lambda = 0.5$, zero inflation is sustainable if the bank's discount factor is at least 0.4. Below that threshold, the lowest sustainable inflation rises quickly toward the discretionary 0.5 (Figure 4.2). A bank with a

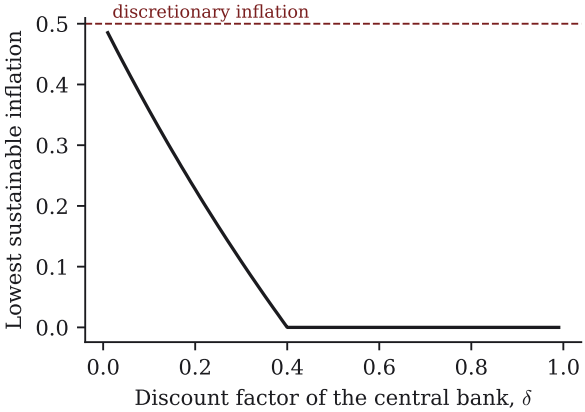


Figure 4.2. Lowest inflation sustainable under trigger strategies as a function of the bank's discount factor, $\lambda = 0.5$, $k = 1$.

long horizon can sustain the commitment outcome without any external enforcement; a bank with a short one cannot.

The discount factor here is the bank's horizon, not the household's. A bank that expects political interference, a short mandate or a change of regime discounts the future heavily, because the reputation it builds may not be its own to enjoy. Central bank independence with long, staggered terms raises δ. Brazil's Complementary Law 179 of 2021 gave the BCB's Governor and directors fixed four-year terms, staggered so that they do not coincide with the presidential term. In the model, that law raised the BCB's discount factor by making the committee's horizon independent of the electoral cycle.

4.2.5 Reputation under uncertainty about type

Trigger strategies assume the public knows the bank's preferences. Backus and Driffill (1985) let the public be uncertain whether the bank is ***committed***, a type that never inflates, or ***opportunistic***, a type with the loss (4.2). The public starts with a prior probability p that the bank is committed and updates it by Bayes' rule after each observation.

The logic of the equilibrium can be seen in two periods. In the last period an opportunistic bank always inflates, because there is no future to protect. In the first period it faces a choice: inflate now and reveal its type,

so that the public expects λk next period; or mimic the committed type, keep p high, and enjoy a lower expected inflation next period, which raises its surprise gain when it finally inflates. When p is not too small, mimicking is optimal, and the opportunistic bank behaves exactly like the committed one until near the end of its horizon. When p is small, the gain is not worth it and the opportunistic bank reveals itself early, in a mixed strategy that makes the public gradually more suspicious.

Three conclusions follow. Credibility is a belief about type, held by the public and updated by behaviour. It can be built only slowly, by a record of decisions that an opportunistic type would not have taken. And it can be lost in one meeting, by a decision that only an opportunistic type would have taken. The Copom vote of May 2024 is an example of the last: the split did not change policy, but it was evidence about the type of the committee the public would face once the new appointees became a majority, and the public updated accordingly.

4.3 Decisions by committee

The Fed, the ECB and the BCB decide by committee: twelve voters on the FOMC, the six members of the ECB's Executive Board and the governors of the euro-area national central banks on the Governing Council, nine members of the Copom. Committees change the game in two ways.

4.3.1 Aggregating preferences

Suppose members differ only in their weight λ_j and vote on a single interest rate. Each member's loss is single-peaked in the rate: it has one preferred rate and dislikes rates further from it in either direction. By the median voter theorem, the preferred rate of the median member beats any other proposal in a pairwise vote. Policy is set as if by a single banker with the median λ.

Two consequences follow. The appointment of the median member is the decisive political act, more than the appointment of the chair. And with staggered terms, the median moves slowly, which is another way of raising the committee's δ. The Copom episode that opened this chapter shows the converse: when the public expects the median to shift toward a different type at a known date, expectations move before the shift.

4.3.2 Aggregating information

If members also hold private information about the state of the economy, voting aggregates it. Suppose each of n members observes an independent signal of the right interest rate with equal precision, and the committee adopts the median of their preferred rates. For large n, the variance of the median estimate is about $\pi/(2n)$ times the variance of one member's estimate. A committee of nine reduces the error variance of a single member by about four fifths. In laboratory experiments, Blinder and Morgan (2005) find that groups make better decisions than individuals and reach them no more slowly.

The cost of committees is communication. Nine voices are harder to read than one, and a split vote is information that may be misread. The Fed publishes dissents and, with a lag, transcripts; the ECB publishes accounts without attributing views; the BCB publishes the vote and the names. Chapter 15 measures what each disclosure adds to the market's understanding of policy.

4.4 The fiscal-monetary game

4.4.1 Two authorities, one budget constraint

All of the above assumes the treasury adjusts taxes and spending to keep public debt sustainable, whatever the central bank does. That assumption fails often enough in emerging markets, and occasionally in advanced ones, to deserve its own analysis.

The government's budget constraint links the two authorities. In real terms, debt at the start of a period plus interest must be financed by the primary surplus, new debt or seigniorage. Iterating forward and imposing that debt cannot grow faster than the interest rate forever gives the intertemporal budget constraint

$$\frac{B_{t-1}}{P_t} = \mathbb{E}_t \sum_{k=0}^{\infty} \Lambda_{t,t+k}\, s_{t+k}, \tag{4.13}$$

where B_{t-1} is nominal debt outstanding, P_t the price level and s the real primary surplus (including seigniorage). The real value of the debt must equal the present value of future surpluses. Something must adjust to make this hold: future surpluses, or the price level.

Table 4.1. The fiscal-monetary game. Rows: the central bank; columns: the treasury.

	Passive treasury (surpluses respond to debt)	Active treasury (surpluses do not respond)
Active bank (Taylor principle)	Monetary dominance: the bank pins down inflation. The regime of chapters 1 to 3.	Conflict: no stable equilibrium. Debt explodes or one side yields.
Passive bank	Indeterminacy: inflation is not pinned down.	Fiscal dominance: the price level adjusts to make debt sustainable.

4.4.2 Active and passive policies

Leeper (1991) calls a treasury ***passive*** if it adjusts future surpluses to keep (4.13) satisfied at any price level, and ***active*** if it sets surpluses without regard to debt. A central bank is ***active*** if it obeys the Taylor principle and ***passive*** if it does not. Determinacy requires exactly one active authority. The four combinations have different outcomes (Table 4.1).

Under ***fiscal dominance***, equation (4.13) becomes the equation that determines the price level (Woodford, 1995). If expected surpluses fall, the price level must rise to reduce the real value of the debt. The central bank loses control of inflation, and its interest rate decisions can have perverse effects. A higher rate raises the interest bill and, if surpluses do not respond, lowers the present value of surpluses net of interest; by (4.13), the price level must rise further. Under this regime, with fiscal policy active and surpluses unresponsive, a tightening can raise the price level rather than lower it.

Sargent and Wallace (1981) gave the earlier version in terms of money. If the treasury runs deficits that it will not close, they must eventually be financed by seigniorage. A central bank that holds money growth down today forces more borrowing today, larger debt tomorrow, and therefore more seigniorage, and more inflation, later. The bank can choose when inflation comes, not whether it comes, and holding it down now may make it larger in the end. This is their ***unpleasant monetarist arithmetic***.

4.4.3 The game of chicken

When both authorities are active, the outcome is the worst cell of Table 4.1. Neither has a dominant strategy: each does better if the other yields. The

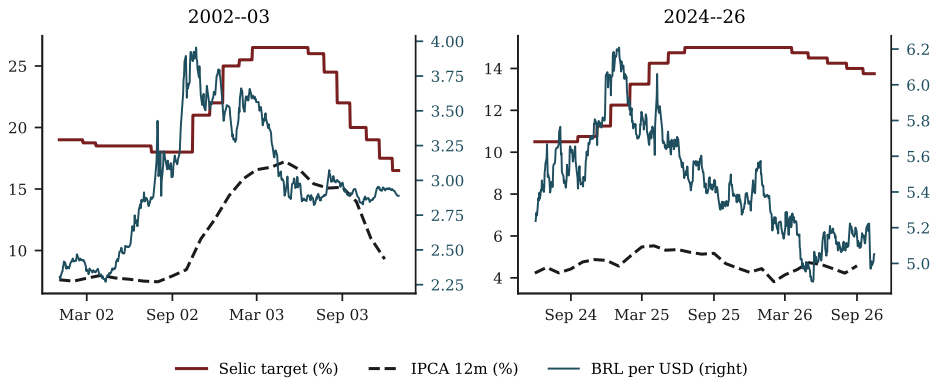


Figure 4.3. The Selic target, twelve-month IPCA inflation and the exchange rate (BRL per USD, right axis) in two fiscal episodes: 2002–03 and 2024–26. Source: BCB SGS series 432, 433 and 1.

central bank wants the treasury to adjust surpluses so that it can keep pinning down inflation; the treasury wants the central bank to accommodate so that it does not have to adjust. The game is one of chicken, and its outcome depends on which side is believed to be more committed to its stance.

Result: the signature of the worst cell

In the conflict cell, tighter monetary policy does not lower inflation expectations. Instead, sovereign spreads widen, the currency weakens and long-term inflation expectations rise, because the market prices the possibility that the central bank will eventually yield. Those three movements together, in response to tightening, are consistent with fiscal dominance. They are not unique to it: adverse supply news, simultaneous fiscal news, a loss of credibility or a global dollar shock can produce the same pattern, so the signature is a reason to look harder, not an identification.

Episode: Brazil, 2002–03

In the run-up to the October 2002 presidential election, the market priced a significant probability that the new government would not

maintain primary surpluses. The real depreciated from 2.54 per dollar in early June to 3.96 on 22 October. The BCB raised the Selic from 18% in July to 21% in an extraordinary meeting on 14 October and to 25% in December, and to 26.5% by February 2003. IPCA inflation reached 12.5% in 2002. Blanchard (2004) argues that in this period higher Selic rates raised the perceived probability of default by raising the interest bill, which weakened the real and raised inflation: the signature of the conflict cell. The episode ended when the new government raised its primary surplus target and the treasury, in the model's terms, became passive. The Selic was cut from 26.5% in June 2003 to 16.5% by December, and the real recovered.

Episode: Brazil, 2015 and 2024

In 2015, with the primary balance in deficit and IPCA inflation reaching 10.7%, the BCB raised the Selic to 14.25% in July and held it there until October 2016. Expectations began to fall in mid-2016, when a constitutional ceiling on federal spending was proposed; it was approved at the end of the year, and inflation fell to 2.95% in 2017. The ceiling was one force among several: the deep recession of 2015 and 2016, the tight Selic, the appreciation of the real and a collapse in food inflation in 2017 all pushed inflation down at the same time. In late 2024, after a fiscal package judged insufficient by the market, the real weakened from 5.81 per dollar at the start of November to 6.21 at the start of January 2025, and Focus expectations for 2025 and 2026 drifted further above the target. The BCB responded with the 100-basis-point increases and forward guidance described in chapter 3, together with heavy foreign-exchange intervention (Figure 4.3). The real recovered to about 5.05 per dollar by October 2026, inflation returned toward the band, and the Selic was cut from 15% to 13.75% between March and September 2026. Both episodes are consistent with the model's prediction that the size and duration of the tightening needed to bring inflation down depend on the fiscal response.

4.4.4 The arithmetic of Brazilian debt

Whether the treasury is passive is, in the end, a question of numbers. Let b_t be public debt as a share of GDP, i the average nominal interest rate on it, n the growth rate of nominal GDP and s_t the primary surplus as a share of

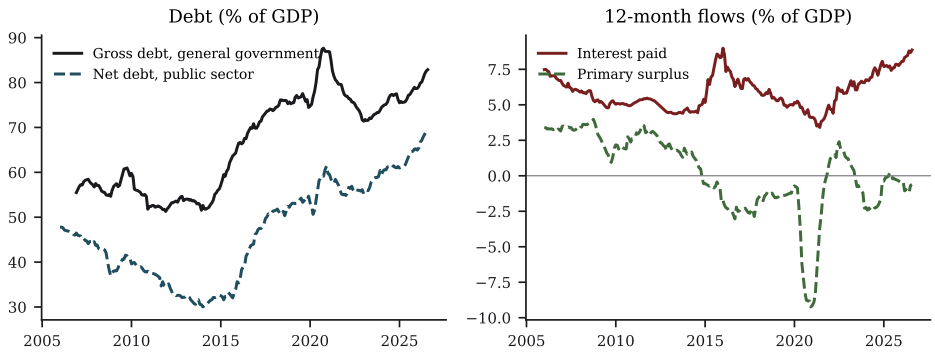


Figure 4.4. Brazil: gross general government debt and net public sector debt (left), and interest paid and the primary surplus of the consolidated public sector over twelve months (right), percent of GDP, 2006 to August 2026. Source: BCB SGS series 13762, 4513, 5727 and 5793.

GDP. The debt ratio evolves as

$$b_t = \frac{1+i}{1+n}\, b_{t-1} - s_t, \tag{4.14}$$

so the primary surplus that holds the ratio constant is

$$s^{*} = \frac{i-n}{1+n}\, b. \tag{4.15}$$

When the interest rate on debt exceeds the growth rate of the economy, a primary surplus is needed just to stand still, and it grows with the level of debt and with the gap $i - n$. The central bank sets the short end of i.

The Brazilian numbers, from the BCB's fiscal statistics, show how the game has tightened. Net public sector debt was 46.5% of GDP at the end of 2006, fell to 30.5% by the end of 2013 and has since risen to 69.3% in August 2026; gross general government debt reached 82.9%. Over the twelve months to August 2026 the consolidated public sector ran a primary deficit of 0.6% of GDP and paid interest of 8.9% of GDP, so the nominal deficit was 9.5%. The interest bill divided by net debt gives an effective net interest rate of about 12.8%: not the average coupon on a homogeneous debt, but a summary of the cost of liabilities net of the return on assets, across indexers, swaps and valuation effects. With nominal GDP growing at around 6.5% a year, equation (4.15) gives a debt-stabilising primary surplus of about $(0.128 - 0.065)/1.065 \times 0.69 \approx 4.1\%$ of GDP.

Figure 4.4 shows the two sides of the game over two decades. Until 2013 the public sector ran primary surpluses, of about 3% of GDP in 2006 and

still 1.7% in 2013, and net debt fell. From 2014 the surplus turned into a deficit, and the interest bill, which moves with the Selic, rose to a peak of 9.0% of GDP in January 2016. It fell to 3.4% in mid-2021, when the Selic was at 2%, and is back at 8.9% in 2026: over the tightening of 2021 to 2025 the share of GDP the public sector pays in interest roughly tripled, with the Selic the main driver alongside inflation-indexed debt and the growth of the debt itself. In the language of this section, the central bank's active policy raised the surplus the treasury needs, and the treasury has not supplied it.

Result: the fiscal gap

On the 2026 numbers, stabilising Brazil's net debt requires a primary surplus of about 4% of GDP; the public sector is running a deficit of 0.6%. The gap of more than 4.5 points of GDP is what the market prices when it asks whether the treasury is passive. Every point on the Selic raises the average rate on the debt with a lag, as floating-rate and short-term debt reprices, and with it the surplus needed. This is the mechanism through which, in the conflict cell of Table 4.1, tightening can raise the perceived probability that the central bank will eventually be forced to yield.

4.5 Political pressure

The institutions of section 4.2 exist because governments, which bear the electoral cost of recessions, have an incentive to press central banks for looser policy, especially before elections. The pressure has been documented in every era.

Episode: Nixon and Burns, 1971–72

Tapes and diaries released after the fact show President Nixon pressing Fed Chairman Arthur Burns for easier policy ahead of the 1972 election. Money growth was rapid in 1971 and 1972 and inflation rose sharply after the election, compounded by the oil shock of 1973. The episode is the standard example of the political business cycle and one of the reasons the Fed's independence came to be valued as it was in the following decades.

Episode: Turkey, 2019–21

Between July 2019 and March 2021 the President of Turkey dismissed three governors of the central bank, each after disagreements over interest rates; the third was removed two days after raising the policy rate. The lira fell sharply after each dismissal. In the model of section 4.2.4, a bank whose governor can be removed for raising rates has a discount factor close to zero: no reputation can be built when the person building it may not be there next month.

In 2025 the President of the United States repeatedly called in public for much lower interest rates and sought to remove a member of the Federal Reserve Board, raising for the first time in decades the question of how far the legal protections of the Fed's governors would hold. The question matters for every result in this chapter: the value of independence lies in the discount factor it gives the bank, and that value depends on the public's belief that the protection will survive a test.

The games in this chapter all turn on what the public believes about the bank and the treasury. Chapter 5 makes those beliefs explicit: what the public observes, what it infers, and when telling it more makes things better.

Chapter 5

Information

In November 2002 the Focus survey asked its panel of forecasters for their expectation of IPCA inflation over the following twelve months. Their answers differed from one another with a standard deviation of 2.6 percentage points. In March 2019 the same question produced a standard deviation of 0.23 points, a tenth as much. The forecasters had not become ten times better at forecasting in the interval, and the economy had not become ten times simpler. What had changed was how much they agreed about what the central bank would do, and therefore about where inflation was heading. Disagreement about inflation is, to a large degree, disagreement about policy.

The games of chapter 4 assumed that everyone sees the same state of the economy and differs only in what they believe the central bank will do. In fact the public sees a noisy version of the state, the bank sees a different noisy version, and each learns from the other's actions. Policy is partly an exercise in sending and reading signals. This chapter builds the information theory behind that. It starts with signal extraction, the single formula behind every inference in the chapter, and shows how it turns the Phillips curve into a relationship that changes with policy. It then extends signal extraction over time to show how expectations become anchored or drift. It turns to the central bank's information advantage and to what a policy action reveals about it. It solves the Morris-Shin model of the social value of public information, which every central bank communication department knows, and shows when a public signal can lower welfare. It introduces limits on attention, measured in bits, and closes with the bank's choice of what to reveal. Part IV takes each result to the data.

5.1 Signal extraction

5.1.1 The formula

Suppose an agent wants to know a variable θ, about which it has a prior belief that is normal with mean μ_0 and variance σ_0^2. It observes a signal $s = \theta + \eta$ with independent normal noise of variance σ_η^2. Bayes' rule gives a posterior that is normal with mean

$$\mathbb{E}[\theta \mid s] = \mu_0 + g\,(s - \mu_0), \qquad g = \frac{\sigma_0^2}{\sigma_0^2 + \sigma_\eta^2}, \tag{5.1}$$

and variance $(1 - g)\sigma_0^2$. The agent moves its belief toward the signal by a fraction g, the ***weight*** or ***gain***, which is the share of the signal's variance that comes from the thing the agent cares about. A precise signal (small σ_η^2) gets a weight close to one; a noisy signal gets a weight close to zero. In terms of precisions, the reciprocals of variances, $g = \tau_\eta/(\tau_0 + \tau_\eta)$: each source of information is weighted by its precision relative to the total.

Every inference in this chapter is an application of (5.1). What changes from section to section is what the signal is, who observes it, and whether others' beliefs also matter.

5.1.2 The Lucas islands

Lucas (1972) gave the first equilibrium model in which money matters only because people confuse it with something else. A producer sees the price of its own good, $p_i = p + z_i$, which mixes the aggregate price level p with a relative-price shock z_i specific to its market. It wants to produce more only if its relative price has risen, but it cannot observe z_i directly. With normal priors for p (variance σ_p^2 around its expected value $\mathbb{E}p$) and for z_i (variance σ_z^2 around zero), the signal-extraction formula gives its best estimate of the relative price:

$$\mathbb{E}\,[z_i \mid p_i] = \theta\,(p_i - \mathbb{E}p), \qquad \theta = \frac{\sigma_z^2}{\sigma_z^2 + \sigma_p^2}. \tag{5.2}$$

The producer raises output in proportion to this estimate. Averaging over producers, the relative shocks cancel and aggregate output responds to the aggregate price surprise with slope θ: a Phillips curve.

The slope is not a structural parameter. Written as the response of output to surprise inflation, it falls when the variance of aggregate prices

rises, because producers learn to attribute more of any price change to the aggregate. A central bank that tries to exploit the trade-off by making inflation more variable shrinks it, until surprise inflation buys almost no output. That is the Lucas critique in one equation. In the orientation used in the rest of this book, inflation on slack, the same result reads the other way: the slope of inflation on output is $1/\theta$, which rises with σ_p^2. It reappears in chapter 9: a long period of stable inflation, with σ_p^2 low, makes inflation respond little to output, a flat curve in the data, and the curve can steepen abruptly when inflation becomes volatile again.

5.1.3 Learning over time and the anchoring of expectations

Signal extraction repeated over time is ***learning***. Suppose the public is trying to learn the central bank's true long-run inflation target π_t^*, which may drift as a random walk with innovation variance σ_ω^2, and each period observes inflation $\pi_t = \pi_t^* + e_t$, with transitory noise of variance σ_e^2. This is a local-level model, and the Kalman filter is the optimal learning rule. In its steady state the long-run expectation evolves as

$$\pi_t^{LR} = \pi_{t-1}^{LR} + g\,(\pi_t - \pi_{t-1}^{LR}), \qquad g = \frac{-q + \sqrt{q^2 + 4q}}{2}, \quad q = \frac{\sigma_\omega^2}{\sigma_e^2}. \tag{5.3}$$

The gain g depends only on the signal-to-noise ratio q: how much the target is believed to move relative to how much inflation fluctuates around it.

Deriving the gain

The gain in (5.3) comes from two steps repeated each period. Let P_{t-1} be the variance of the public's error about the target after observing inflation in $t-1$. Before seeing π_t, the target may have drifted, so the variance of the error rises to $P_{t-1} + \sigma_\omega^2$: the ***prediction*** step. On seeing π_t, the public applies the signal-extraction formula (5.1) with prior variance $P_{t-1}+\sigma_\omega^2$ and noise variance σ_e^2, so the gain is

$$g_t = \frac{P_{t-1} + \sigma_\omega^2}{P_{t-1} + \sigma_\omega^2 + \sigma_e^2}, \tag{5.4}$$

and the variance falls to $P_t = (1 - g_t)(P_{t-1} + \sigma_\omega^2)$: the ***update*** step. In the steady state $P_t = P_{t-1} = P$. Substituting the gain into the update and writing $p = P/\sigma_e^2$ and $q = \sigma_\omega^2/\sigma_e^2$ gives $p = (p+q)/(p+q+1)$, which rearranges to the quadratic $p^2 + q\,p - q = 0$ with positive root $p = [-q +$

$\sqrt{q^2+4q}]/2$. The same equation says that the steady-state gain, $(p+q)/(p+q+1)$, equals p itself, so the gain is the expression in (5.3). A small q means the public believes the target barely drifts; its uncertainty about the target becomes small, and it puts little weight on each new observation.

This is the precise meaning of ***anchored expectations***. If the public believes the target never moves ($q = 0$), the gain is zero and long-run expectations ignore current inflation entirely. If it believes the target can drift, any surprise in inflation moves long-run expectations by g times the surprise. A credible central bank is one whose public has a small q. A bank that tolerates persistent misses teaches its public that the target moves, raises q, and makes long-run expectations follow current inflation. Chapter 8 measures this sensitivity for the three banks; the Brazilian long-run expectations that jumped in 2023 and did not come back are an increase in q.

5.1.4 What the anchoring data say about q

Equation (5.3) can be inverted. Solving the steady-state gain formula for the signal-to-noise ratio gives $q = g^2/(1-g)$, so an estimate of how much long-run expectations move with inflation surprises recovers how much the public believes the target drifts. Chapter 8 estimates that a one-point change in twelve-month CPI inflation moves the US 5-year 5-year breakeven by about 0.07 points, and that the euro area SPF expectation five years ahead moves by about 0.02. Read as Kalman gains, which is an illustrative mapping rather than a structural estimate (the regressions use different horizons and instruments, and breakevens contain inflation risk and liquidity premia), these imply q of about 0.005 for the United States and 0.0004 for the euro area: the public behaves as if the standard deviation of drift in the target were 7% of the standard deviation of transitory inflation in the United States and 2% in the euro area. For Brazil the estimated response of the Focus expectation three years ahead is zero, so q is zero by this measure, which is why the jump in Brazilian long-run expectations in 2023 must be read as a discrete change in beliefs about the regime rather than as learning from inflation data.

5.2 The central bank's information advantage

5.2.1 What a decision reveals

Suppose the bank and the market share a prior about the state of the economy s, but the bank observes an additional signal before it decides. Its

rate decision then carries two pieces of news. One is the systematic response to what the bank knows; the other is the deviation from its usual rule. A market that understands the rule but not the bank's signal cannot fully separate them. Write the market's update of its expectation of output after a surprise Δi as

$$\Delta \mathbb{E} y = \underbrace{-b\,\Delta i}_{\text{policy news}} + \underbrace{c\,\Delta i}_{\text{information about } s}. \tag{5.5}$$

The first term is the effect of tighter policy: higher rates lower output. The second is the inference the market draws from the decision: if the bank tightened more than expected, perhaps it knows that demand is stronger than the market thought. The coefficient c is the signal-extraction weight the market puts on the decision as a signal of s, and it is larger the more the market believes the bank has better information.

If $c > b$, a surprise tightening raises expected output. Nakamura and Steinsson (2018) find evidence of exactly that in high-frequency data. In a thirty-minute window around FOMC announcements, they measure the surprise in interest rates and then look at how private forecasts of growth change over the following weeks. Growth forecasts move in the direction of the rate surprise, not against it. In their estimates, the information the surprise reveals offsets much of what the tightening itself does to expected growth.

5.2.2 Delphic and Odyssean

Campbell et al. (2012) give the vocabulary. ***Delphic*** communication forecasts: it tells the market what the bank expects to happen, and therefore what it expects to do, much as the oracle at Delphi foretold events without causing them. ***Odyssean*** communication commits: it binds the bank, as Odysseus had himself bound to the mast, to a path it would not otherwise choose.

The distinction matters because the two have opposite effects. The commitment gain of chapter 3 requires Odyssean guidance: a promise to keep rates low after the economy recovers, which stimulates today. The information effect of (5.5) is Delphic: a statement that rates will stay low may be read as a statement that the economy will stay weak. A market that reads Odyssean guidance as Delphic hears "the economy is weak" when the bank means "we will stay low for longer", and tightens its own conditions instead of easing them. The same words can stimulate or depress, depending on

what the public believes the bank knows. Chapter 14 separates the two in the data.

5.3 The social value of public information

5.3.1 The beauty contest

Is more transparency always better? Morris and Shin (2002) answer no, with a model built on Keynes's image of the stock market as a beauty contest, in which each judge tries to pick not the face he finds prettiest but the one he expects the others to pick.

A continuum of agents indexed by i choose actions a_i. Each wants its action to be close to a fundamental θ and also close to the average action $\bar{a}$ of the others. The utility of agent i is

$$u_i = -(1-r)\,(a_i-\theta)^2 - r\,(L_i - \bar{L}), \qquad L_i = \int (a_j - a_i)^2\, dj, \tag{5.6}$$

where $r \in (0,1)$ measures the coordination motive. The term L_i is the distance of agent i's action from everyone else's, and $\bar{L}$ its average; together they make the coordination motive a zero-sum game among agents, so that society as a whole cares only about the distance of actions from the fundamental. Think of firms setting prices that depend on the state of the economy and on other firms' prices, or of investors who care about fundamentals and about what other investors will do.

Each agent sees a private signal $x_i = \theta + \epsilon_i$ with precision β, independent across agents, and a public signal from the central bank, $y = \theta + \eta$, with precision α, the same for everyone.

5.3.2 The equilibrium

The first-order condition for agent i is

$$a_i = (1-r)\,\mathbb{E}_i\theta + r\,\mathbb{E}_i\bar{a}. \tag{5.7}$$

Guess that every agent uses a linear rule, $a_j = \kappa\, x_j + (1-\kappa)\, y$. Then the average action is $\bar{a} = \kappa\theta + (1-\kappa)y$, because the private noise averages out. Agent i's expectation of the fundamental, from signal extraction with the two signals, is $\mathbb{E}_i\theta = (\alpha y + \beta x_i)/(\alpha+\beta)$. Its expectation of the average action is therefore $\mathbb{E}_i\bar{a} = \kappa\,\mathbb{E}_i\theta + (1-\kappa)\, y$. Substituting both into the first-order

condition and matching the coefficient on x_i with the guess gives

$$\kappa = \frac{\beta}{\alpha+\beta}\left[(1-r)+r\kappa\right] \quad\Longrightarrow\quad \kappa = \frac{\beta(1-r)}{\alpha+\beta(1-r)}, \tag{5.8}$$

so the equilibrium action is

$$a_i = \frac{\alpha\, y + \beta(1-r)\, x_i}{\alpha+\beta(1-r)}. \tag{5.9}$$

Compare this with the weight an agent would put on the public signal if it cared only about the fundamental, $\alpha/(\alpha+\beta)$. In equilibrium the public signal gets more weight, because private precision is discounted by $1-r$. The public signal is valuable twice: it is information about the fundamental, and it is information about what everyone else will do, since everyone sees it. Agents overweight it relative to its accuracy.

5.3.3 Welfare

Welfare depends only on the distance of actions from the fundamental. Substituting (5.9) and taking expectations,

$$\mathbb{E}(a_i-\theta)^2 = \frac{\alpha+\beta(1-r)^2}{\left(\alpha+\beta(1-r)\right)^2}. \tag{5.10}$$

Differentiating with respect to α: write $D=\alpha+\beta(1-r)$ and $N=\alpha+\beta(1-r)^2$; the derivative of N/D^2 is $(D-2N)/D^3$, and $D-2N=-\alpha+\beta(1-r)(2r-1)$. The loss therefore rises with the precision of the public signal whenever

$$\alpha < (2r-1)(1-r)\,\beta. \tag{5.11}$$

More precise public information can make society worse off. Because agents overweight the public signal, its errors become common errors, shared by everyone, which do not average out. If the public signal is imprecise relative to private information, a small improvement in it increases the herding on it more than it improves its accuracy.

Result: how much transparency can hurt

The product $(2r-1)(1-r)$ is at most $1/8$, reached at $r=3/4$. Condition (5.11) therefore requires both a coordination motive above one half and a public signal less than one eighth as precise as private infor-

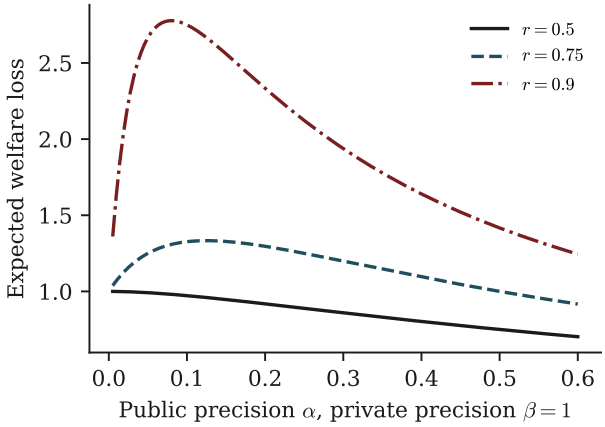


Figure 5.1. Expected welfare loss (5.10) as a function of the precision of the public signal, private precision normalised to one, for three degrees of coordination.

mation; with $r = 0.9$ the threshold is 8% (Figure 5.1). Central bank forecasts of inflation and activity are typically at least as accurate as private ones, so for them the condition fails and more precision helps. The danger is real only for statements about things the bank knows little better than the market, such as the neutral rate, the length of a supply shock or the effect of a new policy. Svensson (2006) makes this point for the original calibration. The practical rule is to be precise about what you know and explicit about what you do not.

5.3.4 Disagreement in the data

The model has a direct counterpart in surveys: the dispersion of private forecasts around the consensus is a measure of how much weight agents put on private rather than common information. Figure 5.2 shows the cross-sectional standard deviation of Focus forecasts of inflation over the next twelve months.

Result: disagreement and credibility

Disagreement peaked at 2.6 points in November 2002, as the market priced the fiscal-monetary game of chapter 4, and was still 1.9 points in January 2003. From 2010 to 2019 it averaged 0.38 points, with a

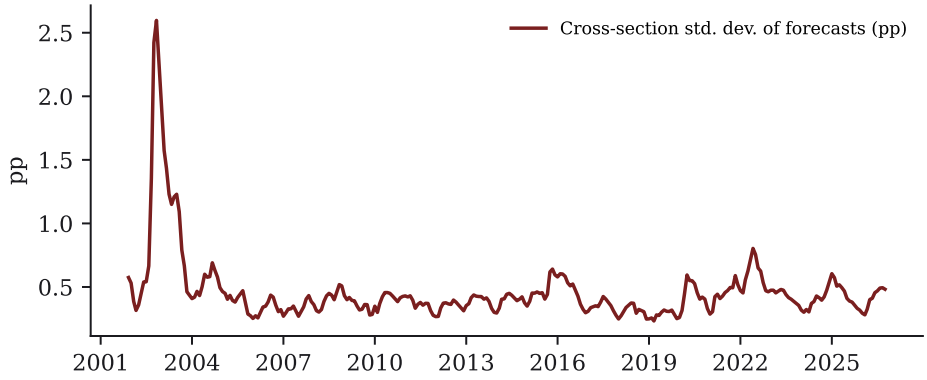


Figure 5.2. Disagreement among Focus forecasters: cross-sectional standard deviation of twelve-month IPCA expectations, monthly average, 2001 to October 2026. Source: BCB, Focus survey.

low of 0.23 in March 2019. It rose to 0.59 at the end of 2015 and to about 0.8 in mid-2022, in the two inflation surges, and stands at 0.48 in October 2026. Disagreement rises with inflation and with doubts about the regime, and falls when the target is credible: the common signal, the target itself, then carries most of the weight.

5.4 Higher-order beliefs and sluggish prices

5.4.1 Woodford's imperfect common knowledge

Signal extraction explains why each firm learns slowly. It does not yet explain why the price level, which aggregates many firms that may each learn reasonably fast, moves so slowly. Woodford (2003b), published the same year as *Interest and Prices*, supplies the missing step: when firms' prices depend on other firms' prices, what matters is not only what each firm believes about the economy, but what it believes others believe, and so on.

Let nominal aggregate demand be q, and let each firm i set its price as a weighted average of what it expects nominal demand and the average price to be:

$$p_i = \mathbb{E}_i \left[\xi \, q + (1 - \xi) \, p \right], \tag{5.12}$$

where $\xi \in (0, 1]$ measures how much a firm's desired price responds to demand rather than to competitors' prices. A small ξ means strong strategic

complementarity, the same property that lowered Θ in chapter 1. Averaging over firms, $p = \xi\,\bar{\mathbb{E}}q + (1-\xi)\,\bar{\mathbb{E}}p$, where $\bar{\mathbb{E}}$ is the average expectation. Substituting the same equation for p inside the expectation, repeatedly,

$$p = \xi \sum_{k=1}^{\infty} (1-\xi)^{k-1}\,\bar{\mathbb{E}}^{(k)} q, \tag{5.13}$$

where $\bar{\mathbb{E}}^{(k)}$ is the k-th order average expectation: the average expectation of the average expectation of … of q. The price level depends on the whole hierarchy of beliefs, with weights that decline slowly when ξ is small.

5.4.2 Why the hierarchy adjusts slowly

Suppose each firm observes q with private noise and puts weight g on its signal, as in (5.1), starting from a prior of zero. The first-order average expectation is $\bar{\mathbb{E}}^{(1)}q = g\,q$. A firm forming a second-order expectation must forecast what others expect; since others put weight g on their signals, and the firm itself only sees its own noisy signal, $\bar{\mathbb{E}}^{(2)}q = g \cdot g\,q = g^2 q$. In general $\bar{\mathbb{E}}^{(k)}q = g^k q$. Substituting,

$$p = \xi\, g \sum_{k=0}^{\infty} \left[(1-\xi)\, g\right]^k q = \frac{\xi\, g}{1-(1-\xi)\, g}\; q. \tag{5.14}$$

Real output, $y = q - p$, responds to nominal demand by the remainder.

Result: strategic complementarity amplifies inattention

With $g = 0.5$, firms individually absorb half of a change in nominal demand. Without strategic complementarity ($\xi = 1$) the price level also absorbs half, and output the other half. With strong complementarity ($\xi = 0.15$) the price level absorbs only 13%, and output 87%. Even with $g = 0.8$, prices absorb only 38%. Higher-order expectations adjust more slowly than first-order ones, because each layer compounds the incompleteness of the one below, and with strong complementarity the price level is driven mostly by the slow upper layers.

In a dynamic version, firms keep learning over time and the higher-order expectations catch up gradually, so a monetary shock has real effects that build and then fade over several quarters, even though no price is sticky. The same mechanism explains why a central bank announcement can be

more powerful than its content: a public statement is common knowledge, so it moves not only each firm's first-order expectation but every layer of the hierarchy at once. That is the formal counterpart of the coordination role of public information in the Morris-Shin model.

5.5 Rational inattention

5.5.1 Attention as a scarce resource

Most households and firms do not read monetary policy statements. Sims (2003) models the limit on attention with Shannon's measure of information. Reducing the uncertainty of a belief from variance σ_0^2 to σ_1^2 requires processing

$$I = \tfrac{1}{2}\log_2 \frac{\sigma_0^2}{\sigma_1^2} \quad \text{bits} \tag{5.15}$$

of information. An agent with a capacity of $\mathcal{K}$ bits per period can reduce the variance of its belief about a variable by a factor of $2^{-2\mathcal{K}}$ each period. Inverting the formula, the posterior variance is $\sigma_1^2 = 2^{-2\mathcal{K}}\sigma_0^2$, and by the signal-extraction formula the belief moves toward the truth with weight

$$g = 1 - 2^{-2\mathcal{K}}. \tag{5.16}$$

Expectations become a moving average of past news, with a sluggishness determined by attention rather than by an assumed rule of thumb.

Result: attention and expectations

With 0.1 bits per quarter, expectations absorb 13% of a change each quarter and take five quarters to absorb half of a permanent change in the inflation target. With 0.25 bits the half-life is two quarters; with one bit, half a quarter (Figure 5.3). When inflation is low and stable, paying attention to it is not worth much, capacity devoted to it is low and expectations adjust slowly: the Phillips curve looks flat. When inflation becomes salient, agents reallocate their limited attention towards it, the capacity devoted to inflation rises, g rises, and expectations, and with them the Phillips curve, start to respond.

5.5.2 Evidence and implications

Coibion et al. (2018) surveyed firm managers in New Zealand, the country that invented inflation targeting, and found that most did not know

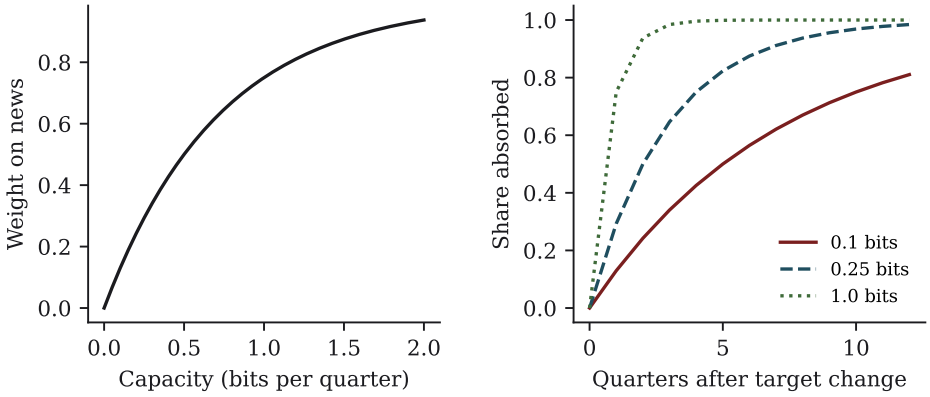


Figure 5.3. Left: weight on new information as a function of attention capacity. Right: share of a permanent change in the inflation target absorbed by expectations, for three capacities.

the central bank's target or the current inflation rate, and that their inflation expectations were dispersed and far from the professional consensus. Maćkowiak and Wiederholt (2009) show why this can be optimal: a firm's profits depend much more on conditions in its own market than on aggregate inflation, so it rationally allocates its attention to the former. The result is slow price responses to monetary shocks even when prices can change frequently, a second source of the inertia the core model lacks.

Two implications run through the rest of the book. Inattention links the information of this chapter to the nonlinear Phillips curve of chapter 9: the curve steepens when inflation is high enough to command attention. And it says what communication is for. A central bank cannot make people pay attention, but it can lower the cost of the information it wants absorbed. A clear number in a statement, a simple rule or a visible commitment substitutes for capacity the public does not have.

5.5.3 Testing for rigidity in Focus forecasts

Inattention and sticky information leave a footprint in forecasts. Coibion and Gorodnichenko (2015) show that if forecasters update their information only occasionally, or weight new information too little, the average forecast adjusts too slowly, and its errors are predictable from its own revisions. Regress the error of the consensus forecast on the latest revision:

$$\pi_Y - F_t\pi_Y = \alpha + \beta\,(F_t\pi_Y - F_{t-1}\pi_Y) + e_t, \tag{5.17}$$

Table 5.1. Coibion-Gorodnichenko regression (5.17) on Focus median forecasts of calendar-year IPCA inflation, quarterly revisions. Newey-West standard errors in parentheses. Source: BCB, Focus survey; IBGE.

Forecast of	2002–2025	2005–2019	2020–2025
Current year	0.20 (0.14)	0.30 (0.25)	0.12 (0.29)
Next year	−0.26 (0.26)	1.12 (1.06)	−2.50 (1.05)

where $F_t\pi_Y$ is the forecast made at t of inflation in calendar year Y. Under full information and rational expectations $\beta = 0$: nothing known at t, including the revision, predicts the error. A positive β means that an upward revision is followed by further upward surprises, the signature of underreaction; in a sticky-information model the fraction of forecasters who do not update in a period is $\beta/(1+\beta)$. A negative β means revisions are followed by surprises in the opposite direction: overreaction.

Result: how Focus forecasters update

For the current year, β is small and positive, 0.20 over the whole sample: a mild underreaction that implies about one forecaster in six not updating in a given quarter, and that is not statistically significant. For the following year the picture changes with the period. From 2005 to 2019 the estimate is positive and imprecise. From 2020 to 2025 it is −2.5 with a standard error of 1.05: revisions of next year's forecast were followed by errors of the opposite sign, the pattern of overreaction that Bordalo et al. (2020) document for individual forecasters in the United States. In the surge and its aftermath, Focus forecasters moved their forecasts for the following year too far in response to news, and then had to move them back.

Calendar-year forecasts have a horizon that shortens through the year, which can create mechanical patterns in revisions and errors. Two checks address it. With quarter-of-year controls the estimates barely change: 0.24 (0.14) for the current year and −2.64 (1.08) for the following year in 2020 to 2025. With a fixed horizon, the Focus forecast of inflation over the next twelve months against realised IPCA inflation over the same months, β is 0.20 (0.31) over the whole sample and 0.92 (1.21) in 2020 to 2025. The mild underreaction in the full sample survives every version; the overreaction of 2020 to 2025 appears only in calendar-year forecasts of the following year,

so it should be read as a feature of that horizon, not of Focus forecasters in general.

The contrast matters for policy. A central bank facing forecasters who underreact can expect its credible actions to be absorbed slowly, which argues for clarity and repetition. One facing forecasters who overreact to news about the following year can expect expectations to swing with each surprise, which argues for anchoring communication on the target and the horizon rather than on the latest data.

5.6 Choosing what to reveal

5.6.1 Bayesian persuasion

Transparency is a design choice. In the framework of Kamenica and Gentzkow (2011), a sender commits, before the state is known, to an information structure: a rule for which signal it will send in each state. The receiver sees the signal, updates its belief by Bayes' rule and acts. The sender cannot lie, because the structure is public, but it can choose how much the signal reveals.

A simple example shows the logic. The state is either good or bad, with prior probability 0.3 of good. The receiver takes a costly action only if it believes the state is good with probability at least one half; the sender always wants the action taken. With full disclosure, the action is taken only in the good state, 30% of the time. With no disclosure, the posterior equals the prior of 0.3 and the action is never taken. The optimal structure is in between: in the good state always send the signal "act"; in the bad state send it with a probability chosen so that, after seeing "act", the receiver's posterior is exactly one half. That probability is $0.3/0.7 \approx 0.43$, and the action is then taken 60% of the time, twice as often as with full disclosure.

The general result works through ***concavification***. Because the receiver's posteriors must average to the prior, the sender can achieve any average of payoffs at posteriors whose mean is the prior. Its value is the smallest concave function lying above its payoff as a function of the receiver's posterior, evaluated at the prior.

5.6.2 What this means for a central bank

Two conclusions carry over. A bank gains from revealing information where its payoff is convex in the public's belief, which is where the public's behaviour changes sharply with what it believes: near a policy turn, near the

edge of a tolerance band (chapter 2), near a run on a bank or a currency. And, in this setup, a bank can gain from coarse communication, revealing a range or a direction rather than a point, where its payoff is concave, which is the formal case for language like "data dependent" or "the committee will remain vigilant" in calm periods. Precision is not always the goal; matching precision to the convexity of the problem is.

Forward guidance, statements, minutes and press conferences are choices on this menu. Chapter 14 measures what guidance does to expected rates and chapter 15 what statements and press conferences add.

Part I now has every element except one. Chapters 1 to 5 treat the short rate as the only instrument. Chapter 6 adds the balance sheet.

Chapter 6

The Balance Sheet in Theory

On 18 March 2009, with the federal funds rate already at zero, the FOMC announced that it would buy up to 300 billion dollars of longer-term Treasury securities over the following six months, on top of its purchases of mortgage-backed and agency debt. The ten-year Treasury yield fell from 3.02% to 2.51% that day, by 51 basis points. Over the next thirteen years the Fed's balance sheet grew from 6% of GDP at the end of 2007 to 35% in early 2022, peaking at 8.97 trillion dollars in April of that year. The ECB's grew from 1.3 trillion euros in January 2008 to 8.8 trillion in June 2022. Both have since shrunk, to 6.8 trillion dollars and 5.9 trillion euros in October 2026.

Interest and Prices gave little reason to expect any of this to work. In the core model of chapter 1, and in the result Eggertsson and Woodford (2003) proved the same year, swapping central bank reserves for government bonds changes nothing unless it changes the expected path of the policy rate. Yet the yield fell 51 basis points in a day. This chapter reconciles the two. It starts from the irrelevance result and shows exactly which assumption each channel of quantitative easing (QE) breaks. It then builds the preferred-habitat model of the term premium, the framework chapters 12 and 13 take to the data, and works out how large an effect a purchase should have. It closes with the plumbing: how the quantity of reserves changes the way a central bank sets the policy rate itself.

6.1 Why purchases should not matter

6.1.1 A Modigliani-Miller theorem for open-market operations

Wallace (1981) proved that, under certain conditions, the composition of the public sector's liabilities is irrelevant, in the same way that the Modigliani-

Miller theorem of corporate finance says the mix of debt and equity does not affect the value of a firm. The argument runs as follows.

Suppose every asset is priced by the stochastic discount factor of a representative household: the price of any asset paying X_{t+1} is $\mathbb{E}_t[\Lambda_{t,t+1}X_{t+1}]$. The central bank buys a long-term bond from the household and pays with newly created reserves. The household now holds reserves instead of the bond. But the bond's risk has not disappeared. If the bond's price falls, the central bank takes the loss, its remittances to the treasury fall, and the treasury must raise taxes or cut spending to make up the difference. The taxes fall on the household. The household therefore bears exactly the same risk as before, in a different form, and values the bond exactly as before. Since its consumption plan is unchanged in every state, its discount factor is unchanged, and so is every asset price.

6.1.2 The lower bound adds a twist

Eggertsson and Woodford (2003) apply the argument at the lower bound on interest rates. When the short rate is zero, reserves and short-term bills are perfect substitutes: both pay zero and both are safe. Expanding reserves by buying bills is then irrelevant for the simplest of reasons. And buying long bonds is irrelevant for Wallace's reason. The only way the central bank can lower long rates is to change expectations of future short rates: to promise, credibly, that policy will stay low after the lower bound stops binding.

Under this view every effect of QE runs through forward guidance. The 51 basis points of March 2009 would be the market's inference that a central bank buying hundreds of billions of long bonds intended to keep short rates at zero for a long time.

6.1.3 Three assumptions

The irrelevance result rests on three assumptions.

1. All investors can trade all assets at all maturities, without limits or costs, so that one pricing kernel prices everything.
2. The risk the central bank absorbs returns to the household unchanged, through taxes and transfers the household fully anticipates.
3. The purchase carries no information about future policy.

Each channel of QE is the failure of one of them. The ***duration*** or ***portfolio-balance*** channel breaks the first, the ***credit*** channel the second and the ***signalling*** channel the third. The rest of the chapter takes them in turn.

6.2 Preferred habitat

6.2.1 Segmented demand and arbitrage

Vayanos and Vila (2021) break the first assumption with an old idea in a modern form. Some investors have a ***preferred habitat***: they want bonds of particular maturities for reasons unrelated to expected returns. Pension funds and life insurers have long liabilities and want long bonds to match them; banks and money market funds want short ones. Their demand for each maturity responds only weakly to its yield.

Between the habitats sit ***arbitrageurs***, such as hedge funds and dealers, who trade across maturities to exploit differences in expected returns. Arbitrageurs are risk averse and have limited capital, so they bear interest-rate risk only if they are paid for it. Bond prices are set by their willingness to hold whatever duration risk habitat investors do not absorb.

6.2.2 A one-factor model

The mechanism can be seen in a model with a single risk factor. Let the short rate follow an AR(1),

$$r_{t+1} = \bar{r} + \phi\,(r_t - \bar{r}) + \sigma\,\epsilon_{t+1}, \tag{6.1}$$

with persistence ϕ and shock variance σ^2. In an affine term-structure model the log price of a zero-coupon bond with τ periods to maturity is linear in the short rate, $\log P_t^{(\tau)} = -A(\tau) - B(\tau)\,r_t$, with

$$B(\tau) = \frac{1 - \phi^{\tau}}{1 - \phi}. \tag{6.2}$$

$B(\tau)$ measures how much a bond's price falls when the short rate rises by one point: the bond's duration measured in units of short-rate risk. It rises with maturity, but less than proportionally, because shocks to the short rate are expected to fade.

An arbitrageur with mean-variance preferences and risk aversion a chooses holdings of bonds of each maturity to maximise expected excess return minus $a/2$ times its variance. Because all bonds are exposed to the same shock, the variance of the portfolio's return is σ^2 times the square of its total exposure,

$$D_t = \sum_{\tau} X_t^{(\tau)}\,B(\tau - 1), \tag{6.3}$$

where $X^{(\tau)}$ is the holding of maturity τ. The first-order condition for each maturity sets its expected excess return equal to its marginal contribution to the portfolio's risk:

$$\mathbb{E}_t \, rx_{t+1}^{(\tau)} = a \, \sigma^2 \, B(\tau - 1) \, D_t. \tag{6.4}$$

Risk aversion times the quantity of risk arbitrageurs hold, times the bond's own exposure to that risk. In equilibrium arbitrageurs hold the net supply left over after habitat demand, so D_t is the duration risk the market as a whole must be paid to carry.

The ***term premium*** on a τ-period yield is the difference between that yield and the average short rate expected over its life. It equals the average of the bond's expected excess returns until maturity:

$$tp_t^{(\tau)} = \frac{1}{\tau} \sum_{k=1}^{\tau} a \, \sigma^2 \, B(k - 1) \, D_t. \tag{6.5}$$

6.2.3 What a purchase does

A central bank that buys long bonds removes them from the market, reducing the duration risk D_t arbitrageurs must hold. By (6.5), every term premium falls in proportion, and by more at longer maturities, because B rises with maturity. Habitat investors push back partially: as long yields fall, some of them hold fewer long bonds, which returns part of the duration to arbitrageurs.

Result: the duration channel

With the ten-year term premium at 1.0% before the purchase, removing 15% of the duration risk held by arbitrageurs lowers it to 0.85%, by 15 basis points. If habitat investors give back 30% of the reduction, the fall is 10.5 basis points. The two-year premium falls by 2.5 basis points and the thirty-year by 19 (Figure 6.1). In the one-factor model the effect is proportional to the duration removed, so a purchase should be measured in ten-year-equivalent duration, not in dollars: a billion dollars of thirty-year bonds removes several times the risk of a billion of five-year bonds.

The proportionality is the model's testable content. Greenwood and Vayanos (2014) show that the maturity-weighted supply of Treasuries predicts excess bond returns, as (6.4) requires: when the government issues

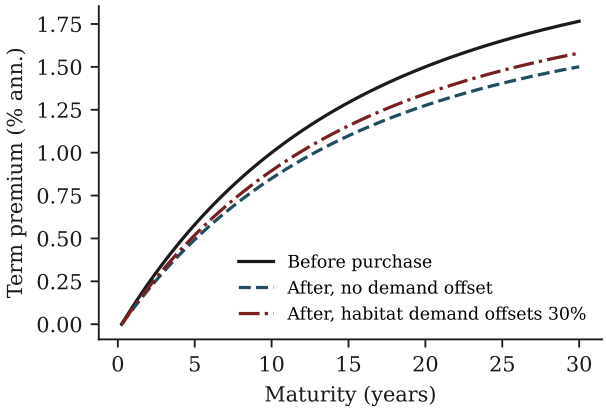


Figure 6.1. Term premium by maturity in a one-factor preferred-habitat model before and after a purchase that removes 15% of the duration risk held by arbitrageurs. Short-rate persistence 0.97 per quarter, innovation standard deviation 30 basis points, risk aversion set so that the ten-year premium starts at 1.0%.

more long debt, long bonds subsequently earn more. Central bank purchases are the same mechanism in reverse.

6.2.4 State dependence

Equation (6.5) contains a second lesson. The effect of a purchase is proportional to a, the arbitrageurs' risk aversion. In a crisis, when dealers' and hedge funds' capital is impaired and a is high, a given purchase has a much larger effect than in calm markets. This is why the first purchase programmes, announced when markets were dysfunctional, had large announcement effects, and later programmes of similar size had smaller ones. It also means the effect need not be symmetric: quantitative tightening (QT), which returns duration to arbitrageurs, raises term premia by the same amount per unit only if their risk appetite is the same as when the bonds were bought.

Figure 6.2 shows the US experience. The Fed's assets rose from 6.0% of GDP at the end of 2007 to 25.0% at the end of 2014, fell to 17.5% by the third quarter of 2019, jumped to 35.2% in the first quarter of 2022 and have since fallen to 20.6%. The Kim-Wright estimate of the ten-year term premium fell from about 1% in late 2008 to a low of −0.61% in July 2020, in the second purchase wave, and has since risen to 1.08% in October 2026. The broad co-movement is what the habitat model predicts. But the term

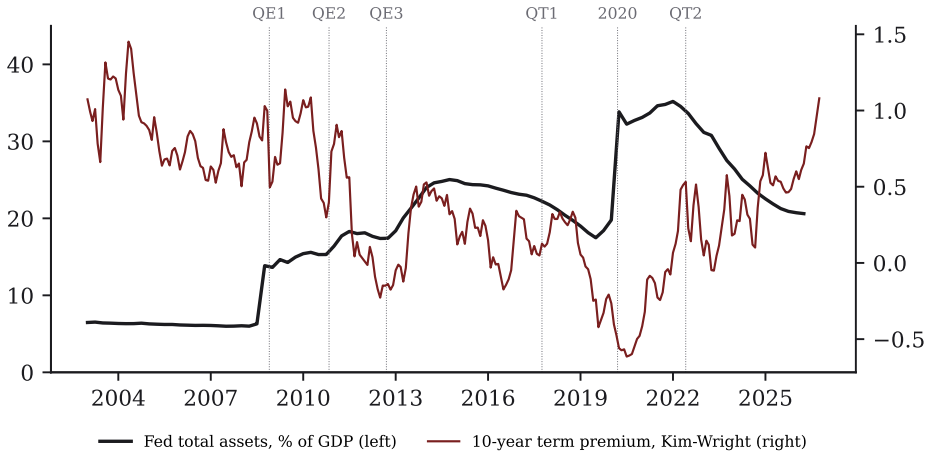


Figure 6.2. Federal Reserve total assets as a share of nominal GDP (left) and the Kim-Wright estimate of the ten-year term premium (right), 2003 to 2026. Vertical lines mark the start of the main purchase and run-off programmes. Source: Federal Reserve and BEA via FRED.

premium also moves with inflation uncertainty, with the supply of Treasuries from fiscal deficits and with global demand for safe assets, so the figure alone cannot isolate the effect of the purchases. Chapter 13 does that with event studies and the term-structure model of chapter 12.

6.3 Credit and the capital of intermediaries

The second assumption fails when the risk the central bank absorbs does not return to the household unchanged. The clearest case is when financial intermediaries are constrained by their own capital.

Gertler and Karadi (2011) build a model in which banks finance loans with deposits and their own net worth, and depositors lend to a bank only if its net worth gives it a large enough stake in repaying them. When a crisis erodes bank capital, banks must shrink their lending, and the spread between the return on loans and the deposit rate widens far beyond what fundamental credit risk justifies. In that state, a central bank that buys private assets, or lends directly to the private sector, substitutes for the intermediation banks can no longer provide. Because the central bank is not capital-constrained in the same way, its purchases narrow credit spreads and

support lending. The effect is largest exactly when private intermediaries are weakest.

This is the theory behind the Fed's purchases of mortgage-backed securities in 2008 and 2009, its corporate bond facilities in 2020, and the ECB's targeted longer-term refinancing operations, which lent to banks at favourable rates conditional on their lending to firms and households. It is also why the effect of such programmes is measured in spreads, not in Treasury yields.

6.4 Signalling

The third assumption fails when the purchase changes expectations of the policy rate. A central bank holding a large portfolio of long bonds takes a capital loss if it raises short rates earlier than the market expects, because long bond prices fall. If the bank dislikes such losses, perhaps because they interrupt its remittances to the treasury and invite political scrutiny, the portfolio is a costly signal that rates will stay low. The mechanism is contested: a central bank does not face the capital constraint of a private intermediary, so the cost it bears is political and reputational rather than financial. To the extent the cost is real, the portfolio works as a commitment (chapter 4) backed by the bank's own balance sheet.

How much of the fall in long yields on QE announcement days came from lower expected short rates rather than lower term premia is the central disagreement of the empirical literature. The answer depends on the term-structure model used to split the yield, because neither component is observed, and estimates of the share due to signalling range from small to more than half. Chapter 13 decomposes each Fed and ECB announcement with the model of chapter 12 and reports the range.

6.4.1 The announcement days

The cleanest evidence comes from the days on which purchases were announced. Some of those days also carried rate decisions or guidance, so the changes mix the purchase with conventional and information effects; they are the best available measure, not a pure one. Table 6.1 reports the one-day change in the two- and ten-year Treasury yields on twelve Fed announcements, and splits the ten-year change into the Kim-Wright estimate of the term premium and the remainder, the change in expected short rates over ten years.

Table 6.1. One-day changes in Treasury yields on Fed balance-sheet announcements, basis points. The term premium is the Kim-Wright estimate; the expected-path component is the ten-year change minus the term premium change (Path). TP: term premium. Source: Federal Reserve via FRED.

Date	Announcement	2y	10y	TP	Path
25 Nov 2008	QE1: MBS and agency debt	−16	−24	−7	−17
1 Dec 2008	QE1: Treasuries mentioned	−10	−21	−10	−11
16 Dec 2008	QE1: rate to 0–0.25%	−10	−16	−8	−8
28 Jan 2009	QE1: FOMC	+2	+12	+3	+9
18 Mar 2009	QE1: Treasuries, expansion	−23	−51	−19	−32
10 Aug 2010	MBS reinvestment	−2	−7	−3	−4
3 Nov 2010	QE2: 600 bn Treasuries	0	+4	−1	+5
21 Sep 2011	Maturity extension	+3	−7	−2	−5
13 Sep 2012	QE3: open-ended MBS	−1	−2	0	−2
12 Dec 2012	QE3: Treasuries added	+1	+6	+2	+4
16 Mar 2020	Pandemic: 700 bn	−13	−21	−7	−14
23 Mar 2020	Pandemic: open-ended	−9	−16	−9	−7

Result: what the announcements did

The five QE1 announcement days lowered the ten-year yield by a cumulative 100 basis points. On the Kim-Wright split, 41 of them came from the term premium and 59 from lower expected short rates. The later programmes, QE2, the maturity extension and QE3, moved ten-year yields by a few basis points in either direction on announcement, because they were largely anticipated and because markets were functioning normally. The two pandemic announcements of March 2020 lowered the ten-year yield by 37 basis points, again split roughly evenly. The pattern matches the state dependence of the habitat model: large effects when markets are stressed and arbitrage capital is scarce, small effects otherwise.

Two cautions apply. The split depends on the term-structure model: another estimate of the term premium would move basis points between the two columns, which is exactly the disagreement about signalling described above. And announcement-day changes measure surprises, not total effects: a programme the market expected moved yields before the announcement, and the small reactions to QE2 and QE3 partly reflect that anticipation. Chapter 13 extends the analysis to the ECB and to quantitative tightening.

Episode: the backstop: ECB 2020–22 and the Bank of England 2022

Two programmes show that a balance sheet can also act as insurance. In March 2020 the ECB announced the Pandemic Emergency Purchase Programme, initially of 750 billion euros, with flexibility to buy more of the bonds of countries whose spreads were widening. In July 2022 it created the Transmission Protection Instrument, which allows purchases of a country's bonds, without a limit set in advance, if its spread rises for reasons the ECB judges unwarranted. In September 2022, after a fiscal announcement in the United Kingdom, long gilt yields rose so fast that pension funds using liability-driven investment strategies faced collateral calls they could not meet, and the Bank of England bought long gilts for about two weeks to stop the spiral. In all three cases the announcement mattered more than the purchases: by promising to act, the bank removed the tail of the distribution the market feared.

6.5 The euro area: purchases against fragmentation

In a monetary union the balance sheet has a fourth use that a single country does not need. The ECB sets one policy rate for twenty-one countries since Bulgaria joined in January 2026, but each country issues its own debt, and a spread between, say, Italian and German yields means that the same policy rate translates into different borrowing costs. When spreads widen for reasons unrelated to each country's fundamentals, the single monetary policy is no longer single. The ECB calls this ***fragmentation***, and from 2012 it built a sequence of instruments against it.

Result: spreads and the ECB's instruments

Italian ten-year spreads over Germany peaked at 519 basis points in November 2011 and Spanish spreads at 555 in July 2012. After the ECB President's statement of 26 July 2012 that the bank would do "whatever it takes" to preserve the euro, and the announcement of Outright Monetary Transactions that followed, the Italian spread fell from 476 basis points that month to 231 by the end of 2013, without a single purchase under the programme. In March 2020 it rose to 209

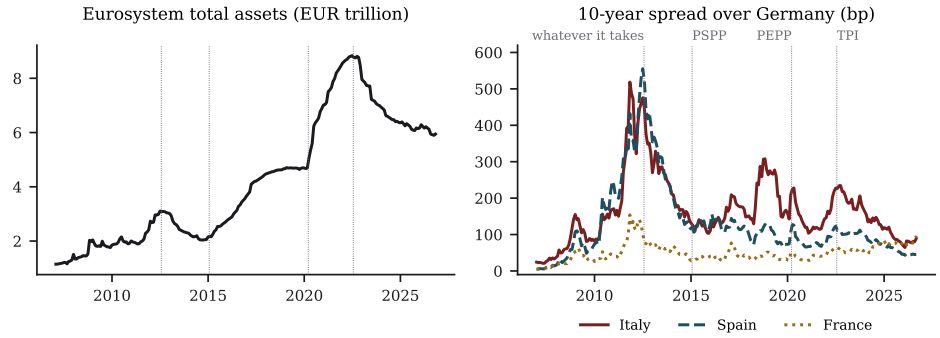


Figure 6.3. Left: Eurosystem total assets, trillions of euros. Right: ten-year government bond yield spreads over Germany, basis points, monthly, with the dates of the main ECB announcements. Source: ECB Data Portal; ECB via FRED.

and fell to 120 by December, as the pandemic purchase programme bought Italian bonds well beyond Italy's capital key. In June 2022, as the ECB prepared to raise rates, it widened to 219; the Transmission Protection Instrument was announced in July. In September 2026 the Italian spread is 89 basis points, the Spanish 45 and the French 97: France now pays more than Italy (Figure 6.3).

The euro area experience is the clearest evidence for the insurance role of the balance sheet. The largest single fall in Italian spreads came from a sentence and a programme that was never used. In the language of chapter 5, the statement changed the distribution of outcomes the market had to price, by removing the tail in which a country is forced out of the union; in the language of chapter 4, it was a commitment by the central bank to act as lender of last resort to sovereigns, conditional on their fiscal behaviour, which changed the equilibrium of a game with multiple equilibria. The cost of that commitment is the moral hazard it creates, which the conditionality attached to OMT and the TPI is designed to contain.

6.6 Brazil: a balance sheet in dollars

The BCB has never bought government bonds on the scale of the Fed or the ECB to lower long rates. Its balance-sheet policy has another target: the exchange rate and the country's exposure to a sudden stop in foreign financing. The two instruments are international reserves and foreign-exchange swaps.

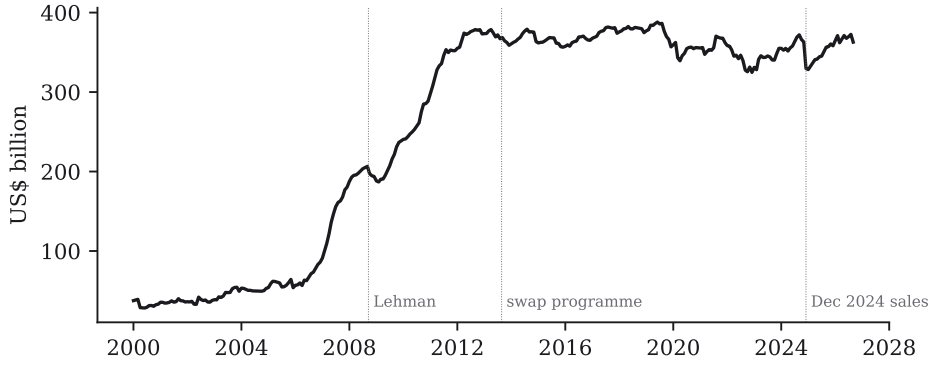


Figure 6.4. Brazil's international reserves, monthly, US$ billion, 2000 to September 2026. Source: BCB SGS series 3546.

Reserves as self-insurance. Reserves rose from 38 billion dollars in January 2000 and 54 billion at the end of 2005 to 194 billion at the end of 2008 and 373 billion at the end of 2012, peaking at 388 billion in June 2019. In September 2026 they stand at 363 billion (Figure 6.4). The accumulation was a deliberate response to the crises of 1999 and 2002, when Brazil had too few dollars to meet its obligations: with reserves larger than its external public debt, a sudden stop is far less likely to become a default: the reserves greatly reduce external-liquidity and foreign-currency funding risk, although they cannot make default impossible. The insurance has a cost. Reserves are financed, in effect, by the repurchase operations through which the BCB absorbs the liquidity their purchase created, which pay roughly the Selic, while reserves earn foreign rates. With reserves of about 1.8 trillion reais at the September 2026 exchange rate and a gap of about ten points between the Selic and dollar yields, the ex-ante carry is on the order of 180 billion reais a year. That is an opportunity cost, not the realised fiscal cost, which also depends on the return on the reserve assets and on the exchange rate: a depreciation produces valuation gains that offset it.

Swaps as a dollar balance sheet without dollars. In August 2013, as the prospect of Fed tapering hit emerging-market currencies, the BCB began daily auctions of ***foreign-exchange swaps***: contracts settled in reais that pay the holder the change in the exchange rate plus a dollar interest rate, in exchange for the domestic interest rate. A swap gives the market the hedge it would have obtained by buying dollars, without the BCB selling any reserves. The stock of swaps grew to around 100 billion dollars by 2015. It is the Brazilian counterpart of the duration channel: instead of absorbing

duration risk from arbitrageurs, the central bank absorbs currency risk from the private sector.

Episode: December 2024

In December 2024, with the real at record lows after a fiscal package the market judged insufficient (chapter 4), the BCB sold dollars in the spot market and through repurchase lines on an unusual scale. Reserves fell from 363 billion dollars at the end of November to 330 billion at the end of December, a fall of 33 billion dollars in one month, the largest monthly fall in the series. Together with the 100-basis-point increases in the Selic, the sales were followed by a stabilisation of the real; without a counterfactual, the separate contribution of each is not identified. It is the clearest recent case in which the BCB used its balance sheet and its policy rate together against the same shock.

6.7 Reserves and the operating framework

6.7.1 Corridor and floor

Purchases are paid for with reserves, and the quantity of reserves changes how the central bank sets the policy rate. Banks hold reserves at the central bank to settle payments and to meet liquidity requirements. Their demand for reserves falls as the overnight interest rate rises, because holding reserves has an opportunity cost. The demand curve is steep when reserves are scarce and flat when they are abundant.

With scarce reserves, the central bank operates a ***corridor***: it sets a lending rate above and a deposit rate below its target, and steers the market rate between them by small daily changes in the supply of reserves along the steep part of the demand curve. With abundant reserves, demand is flat at the rate the central bank pays on reserves, which then sets the overnight rate directly: a ***floor*** system. The bank sets policy by setting the rate it pays.

The Fed has operated a floor since the large purchases. In October 2026 it paid 3.90% on reserve balances, and the effective funds rate traded just below that. The ECB moved to a floor for the same reason and in 2024 described its framework as a demand-driven floor, in which it supplies reserves through lending operations at banks' request. The BCB has never accumulated reserves on the same scale through bond purchases. Brazil's

banking system holds a large structural surplus of liquidity, much of it created by the accumulation of foreign-exchange reserves, and the BCB absorbs it daily with repurchase operations backed by government bonds, keeping the overnight Selic at its target.

6.7.2 Where reserves become scarce

The demand curve for reserves is nonlinear. Above some level it is flat and quantitative tightening has no effect on money-market rates. Below it, the curve steepens and small reductions in reserves move short rates sharply.

Episode: September 2019

On 16 September 2019 the Secured Overnight Financing Rate, the benchmark repo rate in the United States, was 2.43%. On 17 September it was 5.25%, more than twice the top of the Fed's target range, before falling back to 2.55% the next day after the Fed injected reserves through repo operations. A corporate tax date and a large Treasury settlement had drained reserves at the same moment, and the system turned out to be on the steep part of the demand curve. The Fed had already ended its balance-sheet run-off in August; after the September stress it supplied reserves through repo operations and, from October, began buying Treasury bills to restore an ample buffer of reserves. The episode showed that the point where reserves become scarce was higher than expected, and that it moves with regulation and with banks' own liquidity preferences.

The practical question for quantitative tightening is where that point lies. Chapter 13 estimates it from the relationship between reserves and the spread of money-market rates over the rate paid on reserves.

6.8 Quantitative tightening, 2022–2026

From June 2022 the Fed let its holdings run off, at monthly caps that it later slowed. The theory of this chapter makes two predictions about such a run-off: term premia should rise as duration returns to private hands, and money markets should be unaffected until reserves approach the steep part of the demand curve. The second can be checked directly against the liabilities side of the Fed's balance sheet.

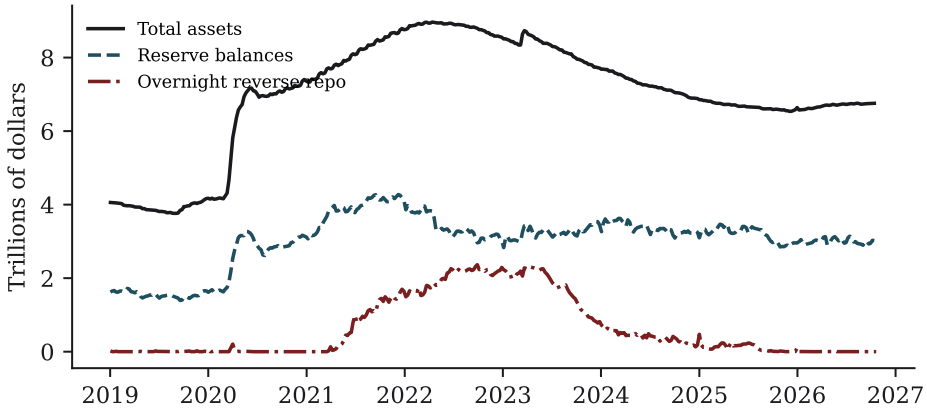


Figure 6.5. Federal Reserve total assets, reserve balances and the overnight reverse repo facility, weekly, 2019 to October 2026, trillions of dollars. Source: Federal Reserve H.4.1 via FRED.

Result: where the run-off came from

Total assets fell from their peak of 8.97 trillion dollars in April 2022 to 6.76 trillion in October 2026, a reduction of 2.2 trillion: 1.2 trillion of Treasuries and 0.8 trillion of mortgage-backed securities. On the liabilities side, balances in the overnight reverse repo facility, where money market funds park cash at the Fed, fell from more than 2.3 trillion dollars in late 2022 to almost zero. Reserve balances fell much less, from 3.3 trillion in June 2022 to 3.0 trillion, and remain more than twice their level of September 2019, when repo rates spiked (Figure 6.5). Most of the run-off drained the reverse repo facility, not bank reserves.

The sequence is the one the demand curve for reserves predicts. Money market funds held cash at the reverse repo facility because it paid more than private alternatives; as the Treasury issued more bills and the Fed's holdings shrank, the funds moved their cash into bills and the facility emptied. Only once it was close to empty did further run-off fall on reserves. With reserves still at about 3 trillion dollars, the system remains on the flat part of the curve, and the policy rate is still set by the rate paid on reserves.

6.9 What the balance sheet adds

The balance sheet gives a central bank three channels for which the short rate is an imperfect substitute: the term premium, through duration; credit spreads, through the capital of intermediaries; and the credibility of a path, through signalling. It also gives the bank exposures it did not have. A bank holding long bonds financed by reserves that pay the policy rate loses money when it raises rates, as the Fed and the ECB did from 2022, and its losses interrupt the remittances to the treasury on which the fiscal-monetary game of chapter 4 depends.

With the balance sheet, Part I is complete. The core model gives the transmission of the short rate; the loss function gives the objective; rules and games give the behaviour of a bank that wants to be believed; information theory gives what the public learns and how expectations become anchored; and the balance sheet adds the instruments that act on the long end and on credit. Part II takes these to the data, starting with the variable every rule depends on: the neutral rate.

Part II

The Modern Facts

Chapter 7

The Neutral Rate

Every rule in chapter 3 starts from r^*, and the natural-rate result of chapter 1 says an error in it becomes an inflation and output error of the same sign. This chapter measures r^* for the three economies of the book: with the model the Fed staff and most of the profession use for the United States and the euro area, and with survey data for Brazil, where that model fails. The answer has three parts. The level fell for two decades and has partly recovered. The uncertainty around it is as large as the level. And real-time estimates are revised by about half a percentage point on average, which is the size of a policy move.

7.1 Measuring an unobserved rate

Woodford (2003a) defines the natural rate inside the model; it cannot be observed. The standard empirical approach, due to Holston, Laubach and Williams (HLW), writes the core model of chapter 1 in a form the Kalman filter can estimate. Output has a trend and a gap; the gap responds to the distance between the real rate and r^* through an IS curve; inflation responds to the gap through a Phillips curve with backward-looking expectations; and the natural rate moves with trend growth g and with a residual z that captures everything else that shifts saving and investment:

$$\tilde{y}_t = a_{y,1}\tilde{y}_{t-1} + a_{y,2}\tilde{y}_{t-2} + \tfrac{a_r}{2}\sum_{j=1}^{2}\left(r_{t-j} - r^*_{t-j}\right) + \epsilon_{\tilde{y},t}, \tag{7.1}$$

$$\pi_t = b_\pi\,\pi_{t-1} + (1-b_\pi)\,\pi_{t-2,4} + b_y\,\tilde{y}_{t-1} + \epsilon_{\pi,t}, \tag{7.2}$$

$$r^*_t = c\,g_t + z_t. \tag{7.3}$$

Here $\tilde{y}_t$ is the output gap, $\pi_{t-2,4}$ the average of inflation two to four quarters back, and r_t the real interest rate, measured as the policy rate minus a four-quarter average of past core inflation. Trend growth g_t and the residual z_t follow random walks, as does the level of potential output, whose growth is g_t.

7.1.1 How the model is estimated

All of $\tilde{y}_t$, g_t, z_t and potential output are unobserved. The model is a ***state-space system***: observed output, inflation and interest rates are linear functions of the unobserved states, and the states evolve as random walks or autoregressions. The ***Kalman filter*** computes, for each quarter, the best estimate of the states given the data up to that quarter. It is signal extraction (chapter 5) applied recursively: each new observation of output and inflation is split into what was expected and a surprise, and the surprise is attributed to the different states in proportion to how much each could have caused it.

One problem dominates. The innovations to trend growth and to z are small relative to the other shocks, and maximum likelihood tends to estimate their variances at zero, which would make the natural rate a constant. This is the ***pile-up problem***. HLW solve it in three stages. First they estimate a version of the model without the interest rate and use a ***median-unbiased estimator***, due to Stock and Watson, to measure how much potential growth varies, by testing for a structural break in its mean and converting the test statistic into a ratio of variances, λ_g. Second they do the same for z, obtaining λ_z. Third they impose these two ratios and estimate the remaining parameters by maximum likelihood. Since 2023 the model also lets the variance of shocks rise in 2020 to 2023 and includes an index of pandemic restrictions, so that the collapse and rebound of output in 2020 are not read as movements in potential output or in the natural rate.

The 2026 parameter estimates for the United States illustrate what the model sees. The slope of the IS curve, a_r, is -0.071: a real rate one point above r^* for a year lowers the output gap by about 0.07 points per quarter. The slope of the Phillips curve, b_y, is 0.075. Both are small, and r^* is identified only through them. The signal-to-noise ratios are $\lambda_g = 0.069$ and $\lambda_z = 0.018$: the model lets trend growth and the residual drift, but slowly.

Two features matter for what follows. The estimate is identified only through the IS curve: if the real rate has no measurable effect on the output gap (a_r close to zero), any path of r^* fits the data equally well. And the one-sided (filtered) estimate, which uses only data available at each date, is what a policymaker could have known; the two-sided (smoothed) estimate uses data from later quarters and is not. The New York Fed publishes the one-sided estimate, and this chapter uses it throughout. An earlier version of the model, by Laubach and Williams, is still published alongside; it differs mainly in its treatment of the pandemic.

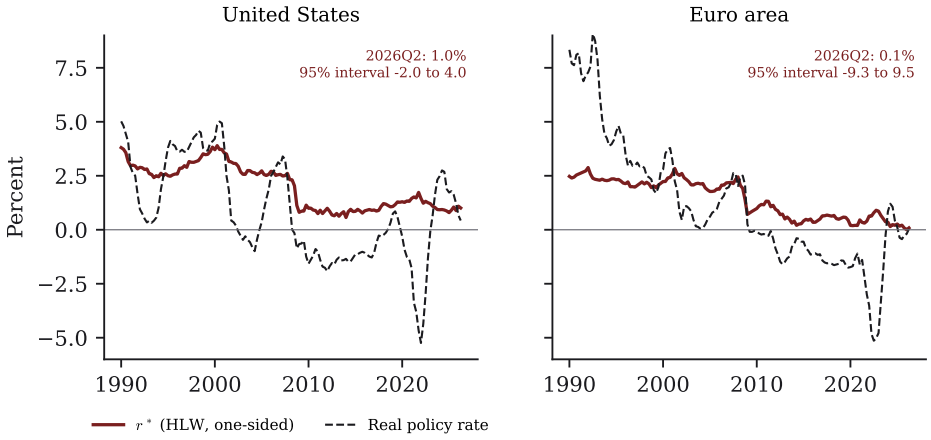


Figure 7.1. HLW one-sided estimates of r^* and the real policy rate (policy rate minus HLW's measure of expected inflation), 1990Q1 to 2026Q2. The note in each panel gives the last estimate and the 95% interval implied by its standard error.

Calibration and data: sources and vintages

United States and euro area: HLW one-sided estimates, parameters and input data from the Federal Reserve Bank of New York, vintage of 2026Q2, and its real-time files for vintages 2015Q4 to 2020Q2 and 2022Q4 to 2026Q2. Brazil: IBGE quarterly GDP (SIDRA table 1621), BCB SGS series 4466 (IPCA-MS core) and 4189 (Selic), the Focus survey (annual and twelve-month medians) through 9 October 2026, and the Oxford stringency index. Public providers only; no terminal data.

7.2 The United States and the euro area

Result: the level and its uncertainty

In 2026Q2 the HLW estimate of r^* is 1.0% for the United States and 0.1% for the euro area. Both fell sharply before the pandemic (from 2.6% and 2.3% in 2007Q4 to 1.3% and 0.4% in 2019Q4) and have not recovered. The standard errors of the final estimates are 1.5 points for the United States and 4.8 for the euro area, so the 95% intervals run

from −2.0% to 4.0% and from −9.3% to 9.5% (Figure 7.1). These bands are filter uncertainty conditional on the estimated parameters; they leave out parameter and model uncertainty, so the total uncertainty is larger still. For the euro area the model says almost nothing about the current level.

The real policy rate in the model is the policy rate minus a four-quarter average of past core inflation. In 2026Q2 it stands at 0.4% for the United States and 0.0% for the euro area. Measured against HLW's own r^*, the Fed is 0.6 points below neutral and the ECB at neutral. Both readings sit well inside the interval around r^*. The model cannot say with any confidence whether either policy is restrictive.

7.2.1 Replicating the estimate

The New York Fed publishes its code and its input data, so the estimate can be reproduced exactly. Running the official code on the published US input file, as this book's scripts do, gives a one-sided r^* of 1.00865% for 2026Q2, identical to the published figure to six decimal places; the largest difference over the whole sample is three ten-millionths of a point. The replication matters less as a check on the New York Fed than as a foundation for the experiments in this chapter: the Brazilian run below uses the same code with Brazilian data.

7.2.2 Why the neutral rate fell

The model splits r^* into a part due to trend growth and a part due to everything else: $r^*_t = c\, g_t + z_t$. The estimated c is 1.11 for the United States and 0.96 for the euro area, close to one, as the Euler equation (1.11) predicts when σ is close to one.

Result: growth and everything else

In the United States, trend growth fell from 3.6% in 2000 to 2.4% in 2026, which lowered r^* by about 1.4 points; the residual z fell from −0.3 to −1.7, another 1.3 points. In the euro area, trend growth fell from 2.3% to 1.3%, and z from zero to −1.1 (Figure 7.2). In both economies, half of the decline in the neutral rate since 2000 is slower growth and half is something else: higher desired saving as populations age, a higher global demand for safe assets, lower investment demand.

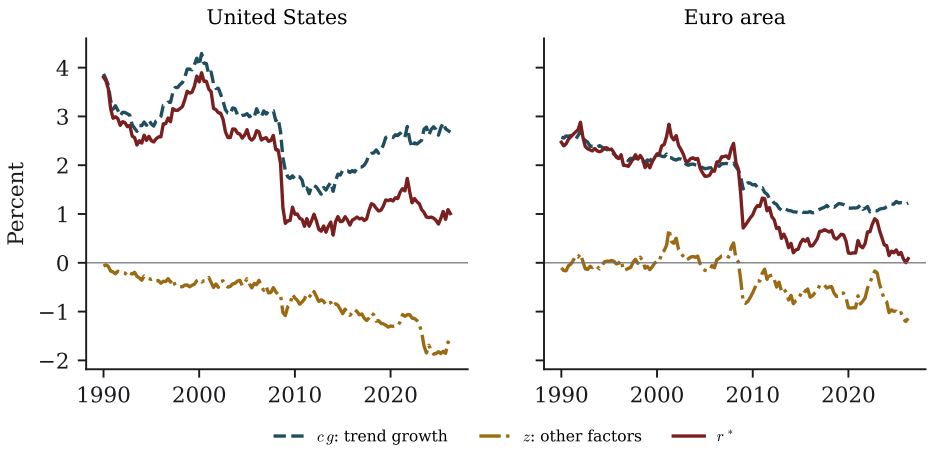


Figure 7.2. Decomposition of the HLW natural rate into trend growth (cg) and other determinants (z), one-sided estimates, 1990Q1 to 2026Q2. Source: Federal Reserve Bank of New York.

Since 2019 trend growth has stabilised in both, while z has kept falling, so the model attributes the absence of a recovery in r^* to forces other than growth.

7.2.3 Two models from the same authors

The New York Fed also publishes the older Laubach-Williams model, which covers only the United States, uses a simpler treatment of the pandemic and differs in a few details of specification. The two models agree closely before 2008: both put the US natural rate at about 2.6% at the end of 2007. They then diverge. At the end of 2019, Laubach-Williams put r^* at 1.74% against 1.27% for HLW; at the end of 2022, 1.72% against 1.25%; and in 2026Q2, 1.65% against 1.01%, the widest gap since 2000. Two closely related models, estimated by the same authors on the same data, disagree by more than half a point on the current neutral rate. The choice of model is itself a source of uncertainty of the same size as the revisions below.

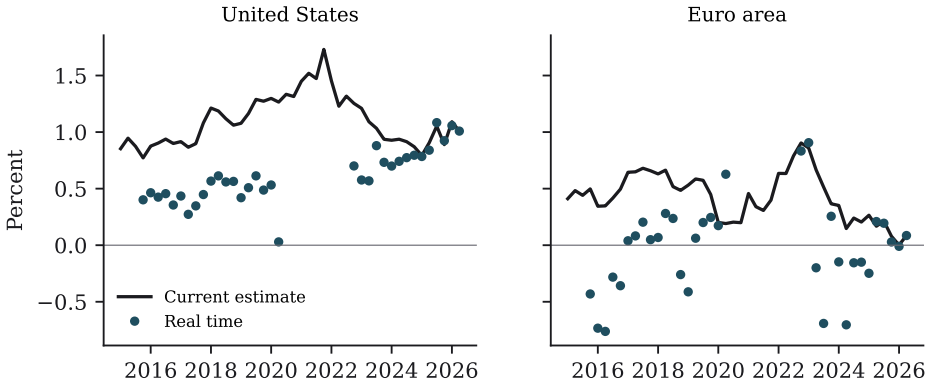


Figure 7.3. r^* as estimated in real time (each vintage's last point) against the current one-sided estimate for the same quarter. No real-time estimates were published between 2020Q3 and 2022Q3.

7.2.4 Real time against hindsight

Result: revisions

The real-time estimate differs from the current estimate of the same quarter by 0.43 points on average for the United States and 0.48 for the euro area. At the end of 2022, in the middle of the fastest tightening in four decades, real-time HLW put the US natural rate at 0.70%; today's estimate for that quarter is 1.25%. A Fed that set policy against the real-time number judged its stance tighter by 0.55 points than it now appears to have been (Figure 7.3).

A revision of half a point is two standard policy moves. It enters every rule of chapter 3 one for one. That is not the case for the first-difference rule, which needs no r^*, and for policy that responds to inflation and activity rather than to an estimated gap between the real rate and an unobserved neutral level.

7.3 What the central banks estimate

The HLW model is one estimate among many, and the central banks use several. Their own numbers are a useful check on the model, and a measure of how uncertain the people setting policy believe the neutral rate to be.

Table 7.1. Estimates of the Brazilian neutral real interest rate by family of methods, second quarter of 2024, percent a year. Group medians. Source: Banco Central do Brasil, Relatório de Inflação, June 2024.

Method	2024Q2
Focus survey: real rate expected four years ahead, and one year ahead filtered	5.4
Output-gap-based, high and low frequency	4.9
NTN-B market real rates net of the term premium (ACM model), 5 to 20 years	5.6
Uncovered interest parity: US rates plus country and currency risk	4.3
BCB models: SAMBA forward rates and the small semi-structural models	4.4
Pre-Copom questionnaire of market participants	5.0
Median of all estimates (25th to 75th percentile)	5.0 (4.7 to 5.5)

7.3.1 The Fed's longer-run dot

Each quarter, every FOMC participant reports the federal funds rate they expect to prevail in the longer run, once the economy has settled at full employment and 2% inflation. Subtracting 2% gives each participant's estimate of the real neutral rate. The median of these projections was 4.3% when they were first published in January 2012, a real neutral rate above 2%. It fell steadily, to 2.5% in 2019 and a low of 2.4% in March 2022, a real rate of about half a point. It has since risen, to 3.0% at the end of 2024 and 3.2% in September 2026: a real neutral rate of 1.2%. The committee's view moved in the same direction as HLW over the whole period and is now about 0.2 points above HLW's one-sided estimate of 1.0%.

7.3.2 The BCB's set of measures

The BCB publishes, in its June inflation reports, a box comparing estimates of the Brazilian neutral real rate from six families of methods (Banco Central do Brasil, 2024b). Table 7.1 reproduces the June 2024 summary.

The spread across families is itself a result. The BCB's own DSGE model, SAMBA, put the neutral rate two and five years ahead at 3.6% and 3.5%; its small semi-structural models at 5.5% and 5.2%; market NTN-B rates, cleaned of the term premium, at 5.4% to 5.7%. The interquartile range of all estimates, 4.7% to 5.5%, is narrower than the US interval from HLW, but the gap between the lowest and the highest family, more than two points,

is as wide. The dispersion is the reason the BCB publishes the whole set rather than a single number: no one estimate is reliable enough to carry a decision alone.

7.4 A global neutral rate?

The HLW estimates for the United States, Canada and the euro area move together. Since 1972 the correlation of their levels is between 0.78 and 0.89, and since 1990 the correlation of the residual z is between 0.74 and 0.83. All three fell from between 2.5% and 4% in 1990 to between 0.1% and 1.8% in 2026. That co-movement is the basis for the view that the neutral rate is largely global, set by world saving and investment in a world of open capital markets.

The co-movement is mostly in the trend. Over five-year windows the correlation of changes is 0.71 between the United States and Canada but only 0.26 between the United States and the euro area. The euro area's natural rate fell faster after 2008 and has not recovered at all since 2019, while the US and Canadian rates stabilised. In 2026 the gap between them is 0.9 points for the US and the euro area, and 1.7 points between Canada and the euro area. A global component sets the level around which national neutral rates move; national growth, fiscal positions and risk premia set how far each one sits from it.

For an emerging market, the global component arrives through the exchange rate and the country risk premium. The BCB's uncovered-interest-parity measures of the neutral rate (Table 7.1) are built on that logic: a US real rate, plus Brazil's sovereign risk premium, plus a currency risk premium. They put the Brazilian neutral rate at about 4.3% in 2024, below the domestic measures, which is consistent with a premium for Brazilian domestic risk that the UIP construction does not fully capture.

7.5 Brazil

7.5.1 Where the model fails

Running the HLW code on Brazilian data (2000Q1 to 2026Q2, IPCA-MS core inflation, Selic, the same pandemic adjustments) does not produce an estimate. The IS-curve slope a_r and the Phillips-curve slope b_y both settle on the constraints the code imposes to keep them away from zero. With no measurable effect of the real rate on the gap, the filter attributes every

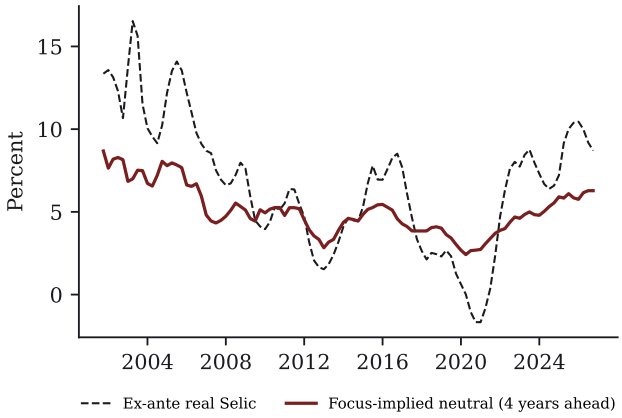


Figure 7.4. Brazil: real rate implied by Focus medians for the Selic and IPCA four calendar years ahead, and the ex-ante real Selic (Selic deflated by Focus twelve-month expected inflation). Quarterly averages, 2001 to October 2026.

movement in the real rate to r^* itself, and the estimate swings from 30% at the start of the sample to -9% at the end. These numbers are not an estimate of anything.

The failure has a reason. In Brazil the real rate the economy faces moves with sovereign risk, the exchange rate and administered prices, none of which is in the model, and the policy rate responds to all three. A closed-economy IS curve with only the domestic real rate cannot separate a tight policy from a high risk premium. This is the fiscal-monetary signature of chapter 4 showing up as a failure of identification.

7.5.2 What the market expects

A measure that does not depend on the model is the real rate the market expects once the cycle is over. The Focus survey asks for the Selic and for IPCA inflation in each of the next four calendar years. The rate expected four years ahead, deflated by inflation expected for the same year, is the survey's answer to where the real rate settles, and the measure the BCB itself uses in its neutral-rate box:

$$r_t^{\text{Focus}} = \frac{1 + i^e_{t,\,t+4}}{1 + \pi^e_{t,\,t+4}} - 1. \tag{7.4}$$

Result: the Brazilian neutral rate

The Focus-implied neutral real rate fell from 7.7% at the end of 2005 to 3.0% at the end of 2019 and a low of 2.4% in the second quarter of 2020. It has since risen to 6.3% in October 2026. The ex-ante real Selic is 8.7%, so policy is 2.4 points above the neutral rate the survey's respondents expect (Figure 7.4). Against the median of all the Brazilian measures, 5.0% (Table 7.2), the gap is 3.7 points: large enough to call policy restrictive with confidence.

The survey measure has its own weakness. A survey expectation is not a bond yield and carries no term premium, but respondents' views of where the Selic settles absorb their views of fiscal risk and of the credibility of the regime, and the two moved together in every fiscal scare since 2015. Its rise from 3% to over 6% is consistent with a higher expected equilibrium rate driven partly by fiscal and credibility concerns, which the survey alone cannot separate. Chapter 12 separates the two with the yield curve.

7.6 Market real rates

The models estimate r^* from the behaviour of output and inflation. Financial markets offer a second view. The yield on ten-year inflation-protected Treasury securities (TIPS) is a real interest rate the market pays today for ten years. It is not an estimate of r^*: it contains a term premium, and it reflects expected real policy rates over the next decade, which need not equal the neutral rate. But over long periods it moves with the market's view of where real rates will settle.

Result: models, committee and markets

The ten-year TIPS yield was 2.44% in the second quarter of 2007, fell to −0.76% at the end of 2012 and to a low of −1.02% in the third quarter of 2021, and rose to 2.15% by the end of 2023 and 2.46% in the third quarter of 2026 (Figure 7.5). From 2007 to 2021 it fell by almost three and a half points, against 1.3 points for HLW from 2007 to 2019 and 1.9 points for the FOMC's longer-run real rate from 2012 to 2022. Since 2022 the three have diverged: the market real rate is back near its 2007 level, the committee's estimate has risen by 0.8 points from its low, and HLW has not moved. Either markets demand

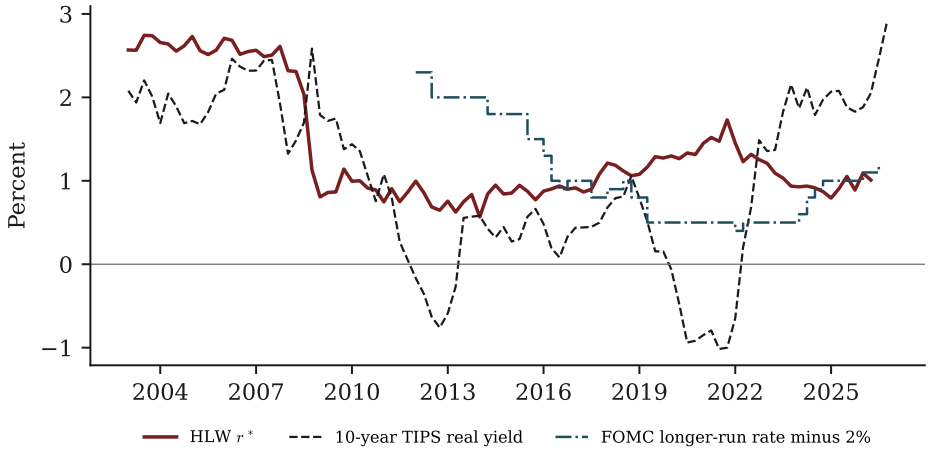


Figure 7.5. HLW one-sided r^*, the ten-year TIPS real yield and the FOMC's median longer-run funds rate minus 2%, 2003 to 2026. Source: Federal Reserve Bank of New York; Federal Reserve via FRED.

> a large real term premium, or they expect real policy rates over the next decade well above what the models call neutral.

7.7 A suite of measures for each market

No single estimate of r^* is reliable enough to carry a decision, and the BCB's practice of publishing the whole set of families, with a median per family and the median across families, is the right discipline. It can be carried to the other two economies, with the families adapted to what each market offers. The uncovered-interest-parity family, which for Brazil adds sovereign and currency risk premia to a US rate, has no counterpart for the United States or the euro area and is dropped there. The families are:

Surveys. For the United States, the SPF's ten-year expectation of the three-month Treasury bill rate minus its ten-year expectation of CPI inflation. For Brazil, the Focus real rate four years ahead. The euro area has no survey that asks for a long-run rate in this form.

Gap-based models. HLW and Laubach-Williams for the United States, HLW for the euro area, the BCB's high- and low-frequency gap measures for Brazil, which are not public series and enter at the median the BCB last published, 4.9% in its June 2024 box.

Table 7.2. Neutral real rate by family, latest quarter (2026Q3 for the United States and the euro area; for Brazil 2026Q4 to date, through mid-October), percent. Market measures are four-quarter averages; models use the last published quarter (2026Q2).

Family	United States	Euro area	Brazil
Survey	0.7		6.3
Gap-based models	1.3	0.1	4.9 (box)
Market real rate, 5y5y	1.8	1.4 (gross)	7.3 (gross)
Central bank	1.2	0.0	5.0
Central bank models			4.4 (box)
Uncovered parity			4.3 (box)
Questionnaire	1.1	0.0	5.5 (2025)
Median	1.2	0.1	5.0

Market real rates, net of the term premium where available. For the United States, the five-year rate five years ahead from the ACM model's risk-neutral yields, which removes the term premium as the BCB does for NTN-Bs, minus the SPF's ten-year inflation expectation. For the euro area, the five-year rate five years ahead on the ECB's own AAA yield curve minus the SPF's longer-term inflation expectation; for Brazil, the same forward on Tesouro IPCA+ (NTN-B principal) real yields. These two are gross of the term premium, which was negative in the euro area in the years of asset purchases and is positive in Brazil. All market measures are four-quarter averages.

Central bank estimates. The FOMC's longer-run median minus 2%, the ECB staff range published in 2025 (1.75% to 2.25% nominal, about zero real), and the BCB's reference value of 5%. For Brazil two further families from the BCB's box, the neutral rates implied by its own models (SAMBA and the small semi-structural models, 4.4%) and by uncovered interest parity with country and currency risk (4.3%), enter at their last published values; they are held constant after June 2024, and should be updated when the BCB publishes a new box.

Questionnaires of market participants. The New York Fed's Survey of Market Expectations (longer-run funds rate median 3.13% in March, June and September 2026), the ECB Survey of Monetary Analysts (long-run deposit rate 2.00%), and the BCB's pre-Copom questionnaire.

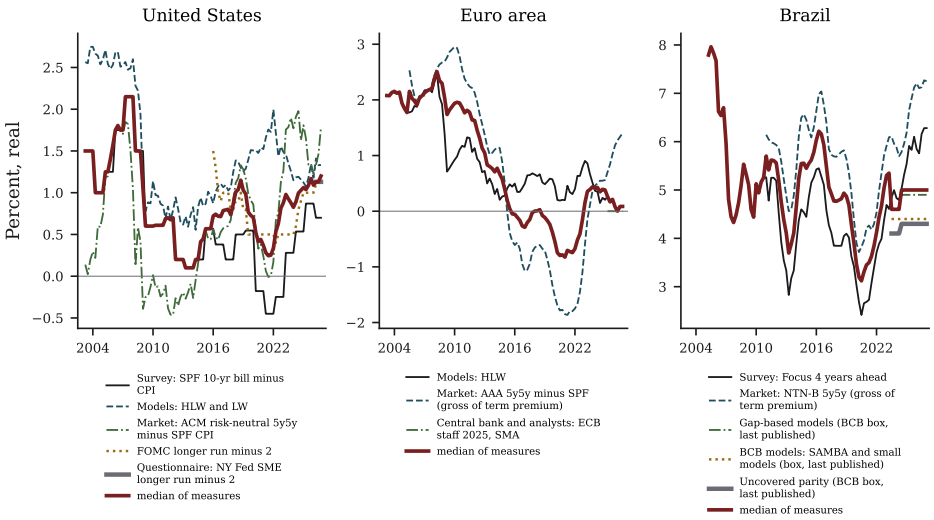


Figure 7.6. Neutral real rate measures by family and their median, 2003 to 2026, quarterly. Market measures are four-quarter averages. Families from the BCB's box that are not public series are shown at their last published values. Sources: Federal Reserve Bank of New York (HLW, LW, ACM, SME), Federal Reserve Bank of Philadelphia (SPF), Federal Reserve, ECB (yield curve, SPF, SMA, Economic Bulletin), Tesouro Nacional, BCB.

Result: the neutral rate as a set

In the third quarter of 2026 the US measures run from 0.7% (survey) to 1.8% (market, net of the term premium), with a median of 1.2%; HLW alone, at 1.0%, sits near the bottom (Table 7.2, Figure 7.6). In the euro area the model, the ECB's staff and the analysts' survey all point to a real neutral rate of about zero, and the market forward of 1.4% is the outlier, consistent with a term premium that has turned positive. In Brazil the survey and the market place the neutral rate at 6.3% and 7.3%, above the BCB's reference value of 5%, the pre-Copom questionnaire's 5.5% and the model, gap and parity families of the BCB's box at 4.3% to 4.9%; the median of the seven families is 5.0%. Excluding the three families held at their 2024 values, the median would be 5.9%. The gap between the market measure and the central bank's has widened since the BCB's 2024 box, when its own NTN-B measure net of the term premium was 5.6%; ours is gross of the term premium.

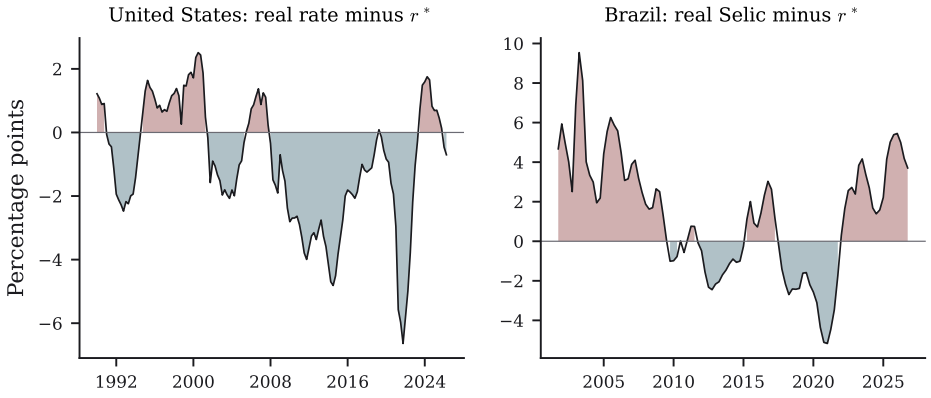


Figure 7.7. Policy stance: real policy rate minus the median of the neutral-rate suite. United States, 1990 to 2026, with the Wu-Xia shadow rate in 2009 to 2015 and 2020 to 2021; Brazil, ex-ante real Selic, 2001 to 2026. Shaded red when restrictive, blue when accommodative.

The median across families is the neutral rate used in the rest of this book wherever a rule or a stance needs one: in the stance measures below and in the estimated reaction functions of chapter 10.

7.8 The stance through the cycles

Subtracting r^* from the real policy rate gives the stance of policy: positive when policy is restrictive, negative when it is accommodative. Figure 7.7 shows the measure for the United States and for Brazil, using the median of the suite in each case. For the United States the real rate uses the Wu-Xia shadow rate in the lower-bound years, so that the easing done through asset purchases counts.

Result: the stance in each cycle

US policy was accommodative in all but one quarter from 2009 to 2021. The stance reached −6.6 points at the end of 2021, the most accommodative reading in the series, as inflation surged with the policy rate at zero and asset purchases still running. It turned restrictive in the third quarter of 2023, peaked at +1.8 points in the second quarter of 2024, and turned mildly accommodative again in 2026 (−0.7 in

the second quarter). In Brazil the stance was −5.2 points in the first quarter of 2021, at the bottom of the 2% Selic, swung to +2.7 by the end of 2022 and +4.2 in the third quarter of 2023, eased to +1.4 in mid-2024, and rose again with the 2024 to 2025 tightening to +5.5 in the first quarter of 2026; it remains restrictive, at +3.7, in the fourth quarter of 2026. The Brazilian swing of almost eight points in two years is the largest of the three economies in this book. Both stance measures use an ex-ante real rate: the policy rate less expected inflation (the HLW expectations measure for the United States, Focus twelve months ahead for Brazil).

The figure also shows why the stance is a fragile guide. In the United States the measure says policy was loose for thirteen years; inflation was below target for most of them. Either the neutral rate was lower than even the median of the suite suggests, or the stance mattered less for inflation than the models imply, or both. The measure is informative about the direction of policy within a cycle and much less about its level across cycles.

7.9 Policy when r^* is uncertain

7.9.1 Shrinking the estimate

A policymaker who knows that r^* is estimated with a standard error of one and a half points should not use the point estimate as if it were known. The signal-extraction formula of chapter 5 says how to combine it with a prior. Suppose the policymaker's prior for the US neutral rate, from long experience and from other models, is centred at μ_0 with standard deviation $\sigma_0 = 1$, an illustrative value rather than an estimated one, and the HLW estimate $\hat{r}$ has a standard error of $\sigma_{\hat{r}} = 1.5$. The best estimate is

$$\mathbb{E}[r^* \mid \hat{r}] = \mu_0 + \frac{\sigma_0^2}{\sigma_0^2 + \sigma_{\hat{r}}^2}\,(\hat{r} - \mu_0) = \mu_0 + 0.31\,(\hat{r} - \mu_0). \tag{7.5}$$

The model's estimate receives less than a third of the weight. For the euro area, with a standard error of 4.8, it would receive about 4%: the HLW estimate is, for practical purposes, uninformative there, and the prior does all the work.

7.9.2 Brainard's principle

Uncertainty about the effect of policy, as opposed to uncertainty about the level of r^*, calls for a different correction. Brainard (1967) considers a policymaker who wants an outcome y close to a target y^*, where $y = b\,i + u$ and the effect b of the instrument is uncertain, with mean μ_b and variance σ_b^2. Minimising $\mathbb{E}(y - y^*)^2$ gives

$$i^* = \frac{\mu_b}{\mu_b^2 + \sigma_b^2}\,(y^* - \mathbb{E}u), \tag{7.6}$$

which is the certainty-equivalent policy $(y^* - \mathbb{E}u)/\mu_b$ multiplied by $\mu_b^2/(\mu_b^2 + \sigma_b^2) < 1$. When the effect of policy is uncertain, the optimal policy does less than it would if the effect were known: moving the instrument adds variance to the outcome. This is the formal case for gradualism, and it is distinct from the commitment argument of chapter 3. Uncertainty about the level of r^*, by contrast, is additive: it shifts where the rate should be, not how strongly it should respond.

7.9.3 Rules that do not need r^*

Orphanides and Williams (2002) take the next step. They compare rules across a range of assumptions about how much r^* and the natural rate of unemployment move over time and how well they can be estimated. Rules that set the level of the rate using estimates of r^* perform well when the estimates are accurate and badly when they are not. Rules written in first differences, which adjust the rate in response to inflation and to the change in activity and never use r^*, perform almost as well when r^* is known and much better when it is not. With real-time revisions of half a point and model disagreements of more than half a point, the case for difference rules, or for heavy inertia in level rules, is strong for the Fed and the ECB. For the BCB, whose reaction function has inertia of 0.9 and momentum, the published rule is already close to a difference rule.

7.10 What follows for policy

The neutral rate is estimated with an error of the same order as its level for the United States, larger than its level for the euro area, and not at all by the standard model for Brazil. Three consequences run through the rest of the book. Rules that depend on r^* should be scored with the revision error built in, which chapter 10 does. Policy that responds to the observed

behaviour of inflation and activity is less sensitive to r^* error than policy set against an estimated gap. And for an economy with a sovereign risk premium, the neutral rate and the premium must be estimated together, which takes the yield curve.

Chapter 8

Measuring Inflation and Expectations

A central bank targets one price index and watches a dozen. The Fed targets headline PCE inflation and talks about core; the ECB targets headline HICP and watches core and services; the BCB targets headline IPCA and publishes five core measures, but sets policy against expectations. This chapter asks a narrow question of each of them: which measure, read today, best predicts headline inflation a year from now. It then measures how well long-run expectations have stayed anchored, which is the credibility the games of chapter 4 are about.

8.1 The measures

Every core measure removes something from headline inflation. ***Exclusion*** measures drop fixed categories: food and energy (core PCE, core CPI, core HICP), administered prices and food at home (the BCB's IPCA-EX0). ***Trimmed means*** drop the largest price changes each month, whichever items they are: the Dallas Fed trimmed-mean PCE, the Cleveland Fed median CPI, and the BCB's IPCA-MA, with IPCA-MS as a smoothed version that spreads large infrequent changes over twelve months. ***Reweighting*** measures, like the BCB's double-weighting IPCA-DP, down-weight volatile items instead of dropping them. Each answers a different question: what is persistent, what is broad, what is not driven by a few items.

Expectations come from three sources. ***Surveys of households*** (the Michigan survey) measure what people who set wages and spending believe. ***Surveys of professionals*** (the Philadelphia Fed and ECB Surveys of Professional Forecasters, the BCB's Focus) measure what forecasters believe, usually closer to the target. ***Market measures*** (breakeven inflation from inflation-linked bonds) measure what investors will pay for, and include risk and liquidity premia.

8.1.1 How a trimmed mean is built

Each month the statistical agency publishes price changes for many items, $\pi_{j,t}$, with expenditure weights w_j summing to one. Headline inflation is the weighted average $\sum_j w_j \pi_{j,t}$. A trimmed mean sorts the items by their price change, removes those whose cumulative weight falls in the bottom α_L and top α_U of the distribution, and averages the rest with their weights rescaled:

$$\pi_t^{TM} = \frac{1}{1-\alpha_L-\alpha_U} \sum_{j \in \mathcal{J}_t} w_j \, \pi_{j,t}, \qquad \mathcal{J}_t = \{ j : \alpha_L < W_{j,t} \leq 1-\alpha_U \}, \tag{8.1}$$

where $W_{j,t}$ is the cumulative weight of items with price changes up to item j's. The median is the limiting case in which only the item at the 50th percentile of the weight distribution survives. The Dallas Fed's trimmed-mean PCE trims 24% of weight from the bottom and 31% from the top, an asymmetric trim chosen because the distribution of price changes is skewed to the right. The BCB's IPCA-MA trims 20% from each tail. Its smoothed version, IPCA-MS, first spreads the changes of items that are adjusted infrequently, such as some administered prices and school fees, evenly over twelve months, and then trims, so that a single large annual adjustment does not appear as an outlier in the month it happens.

The logic of a trimmed mean is that large relative price changes in a given month are more likely to be temporary and specific to a sector than broad and persistent. The logic of an exclusion measure is that some categories are volatile by nature. The two are different bets, and the test below shows which paid off in each economy.

8.1.2 How expectations are measured

The Focus survey, run by the BCB, collects forecasts from about a hundred and fifty banks, asset managers and consultancies, which can update their forecasts on any business day. Each Monday the BCB publishes the median of the forecasts submitted over the previous thirty days, for inflation, the Selic, the exchange rate, GDP and other variables, for each calendar year up to four years ahead and for the next twelve months. Because respondents can update daily, the median responds quickly to news, and because their accuracy is ranked and published, they have an incentive to be right rather than to agree.

The Philadelphia Fed's Survey of Professional Forecasters has been run quarterly since 1968; the ECB's, also quarterly, since 1999, and it asks for inflation one, two and five years ahead. The University of Michigan asks

a sample of about five hundred households each month what they expect prices to do over the next year and over the next five to ten.

Market measures come from comparing nominal government bonds with inflation-linked ones. The ***breakeven*** inflation rate at maturity n is the difference between the nominal and the real yield,

$$BEI_t^{(n)} = y_t^{(n),\text{nom}} - y_t^{(n),\text{real}} = \mathbb{E}_t \bar{\pi}_t^{(n)} + IRP_t^{(n)} - LP_t^{(n)}, \tag{8.2}$$

which equals expected average inflation over the bond's life plus an ***inflation risk premium***, which investors demand for bearing the risk that inflation turns out higher than expected, minus a ***liquidity premium*** that inflation-linked bonds pay for being harder to trade. The five-year, five-year-forward rate, $(10\,BEI_t^{(10)} - 5\,BEI_t^{(5)})/5$, isolates expected inflation from five to ten years ahead, the horizon at which the current cycle should no longer matter, and is the standard market measure of anchoring.

Calibration and data: sources and vintages

FRED for US price indexes, trimmed-mean PCE, median CPI, Michigan expectations and the 5-year 5-year forward breakeven, through September 2026. ECB Data Portal for HICP headline, core and services (through December 2025) and the SPF five-year expectation (through 2026Q3). BCB SGS for IPCA (433) and cores IPCA-MS (4466), IPCA-MA (11426), IPCA-EX0 (11427), IPCA-EX2 (16121) and IPCA-DP (16122), and the Focus survey, through September and October 2026.

8.2 Which measure predicts headline inflation

The test is simple. Take each measure's twelve-month rate in month t as the forecast of headline twelve-month inflation in month $t+12$, and compute the root mean squared error over all forecast origins from January 2005. Add a benchmark that always forecasts the target. A second statistic, the ***attractor coefficient*** b from

$$\pi_{t+12}^{h} - \pi_t^{h} = a + b\left(m_t - \pi_t^{h}\right) + e_t, \tag{8.3}$$

says whether headline inflation moves toward the measure m: $b = 1$ means the gap between the two closes in a year.

Table 8.1. Forecasting headline inflation twelve months ahead with each measure's current twelve-month rate. RMSE in percentage points, forecast origins from January 2005; relative RMSE against headline itself; attractor coefficient b.

Measure	RMSE	Pre-2020	2020 on	Relative
United States, target: headline PCE				
Core PCE ($b = 1.43$)	1.45	1.05	2.18	0.83
Trimmed-mean PCE ($b = 0.77$)	1.57	1.15	2.35	0.90
Constant at 2%	1.57	1.09	2.42	0.90
Core CPI ($b = 0.86$)	1.59	1.19	2.35	0.91
Headline PCE	1.75	1.48	2.33	1.00
Michigan one-year expectations	1.86	1.93	1.69	1.06
Median CPI	1.90	1.45	2.76	1.08
Euro area, target: headline HICP				
Constant at 2%	2.03	1.19	3.51	0.90
Core HICP ($b = 0.75$)	2.13	1.02	3.87	0.94
Services HICP	2.16	1.07	3.90	0.96
Headline HICP	2.25	1.28	3.93	1.00
Brazil, target: headline IPCA				
Focus 12-month expectations ($b = 1.04$)	2.25	1.60	3.40	0.83
Constant at target	2.58	2.06	3.60	0.95
IPCA-MA, trimmed mean	2.70	2.01	3.96	0.99
Headline IPCA	2.71	2.19	3.74	1.00
IPCA-DP, double weighting	2.74	2.09	3.98	1.01
IPCA-EX2	2.82	2.01	4.27	1.04
IPCA-EX0	2.83	1.85	4.47	1.05
IPCA-MS, smoothed trimmed mean	2.93	2.14	4.36	1.08

Result: the United States

Core PCE is the best single predictor of headline PCE a year ahead, with an error 17% smaller than headline itself. Its attractor coefficient of 1.43 means headline does more than close the gap: when headline runs above core, it falls below core the following year, the signature of energy shocks that reverse. Trimmed-mean PCE and core CPI tie with the naive forecast of 2%. Household expectations were the worst predictor before 2020 and the best from 2020 on, with an error of 1.69 points against 2.18 for core PCE: households saw the surge earlier than the core measures did.

Episode: households in 2022

The Michigan measure of household expectations for the year ahead rose to 5.4% in March 2022, when headline CPI inflation was 8.5%, while the Fed's own projections still showed inflation returning toward 2% within two years. Households were the least accurate forecasters before 2020 and the most accurate after it, because the surge was driven by food, energy and the prices of goods they buy often, the prices that dominate household perceptions. In the language of chapter 5, the shock made inflation salient and households paid attention.

Result: the euro area

No measure beats the target. A forecast of 2% every month has an error of 2.03 points, against 2.13 for core HICP and 2.25 for headline. Before 2020, core was the best predictor (1.02); after 2020 every measure failed by about the same margin, because the surge came from energy and food prices that core measures exclude by construction and that then fed into everything else.

Result: Brazil

None of the BCB's five core measures predicts headline IPCA better than headline itself does; the smoothed trimmed mean IPCA-MS is the worst, 8% less accurate than headline. The Focus twelve-month expectation is the best predictor by a wide margin, 17% more accurate than headline, and its attractor coefficient is 1.04: headline inflation closes the full gap to the market's expectation within a year. The BCB's practice of setting policy against expectations rather than against core inflation is what the data support.

Figure 8.1 shows why. The Focus expectation tracks realised inflation closely in normal times: from 2005 to 2019 realised inflation exceeded the expectation by 0.38 points on average. The errors are concentrated in shocks the survey could not foresee, such as the 2002 crisis, when inflation realised over the following year exceeded the May 2002 expectation by almost 13 points, and the 2021 surge. Over the whole sample realised inflation exceeded the expectation by one point on average, which means the survey has been biased toward the target in a country where inflation was usually above it.

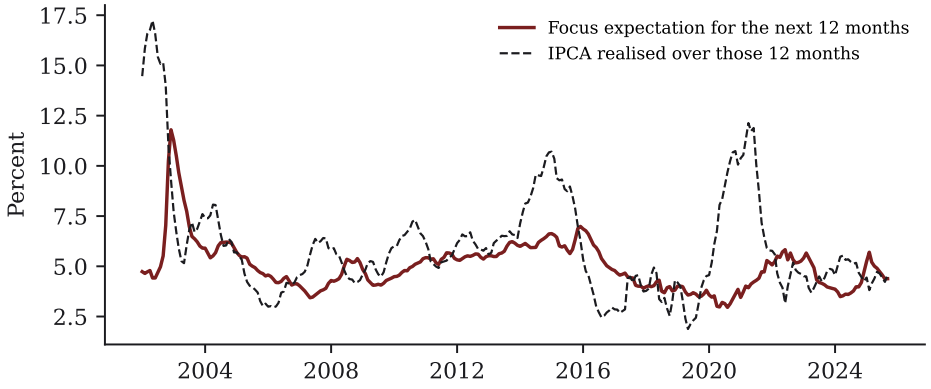


Figure 8.1. Focus median expectation of IPCA inflation over the next twelve months, against the inflation actually realised over those twelve months, by month of the forecast, 2002 to 2025. Source: BCB (Focus), IBGE.

The Brazilian result has a mechanical part. Administered prices and food are a large share of IPCA and move with the exchange rate and the weather; cores that exclude them miss shocks that are not transitory in Brazil, and trimmed means drop exactly the large changes that carry the information. Focus respondents see the same data plus the announced schedule of administered price adjustments and the exchange rate. Chapter 10 uses Focus expectations as the inflation term in the BCB reaction function for this reason.

8.3 Decomposing the measures

8.3.1 Free and administered prices in Brazil

About a quarter of the IPCA consists of ***administered prices***: electricity, fuels, public transport fares, health plans, regulated medicines and other items whose prices are set or approved by government bodies or follow contractual indexation rules. They respond little to demand and a lot to policy decisions, the exchange rate and past inflation. The rest are ***free prices***, set in markets.

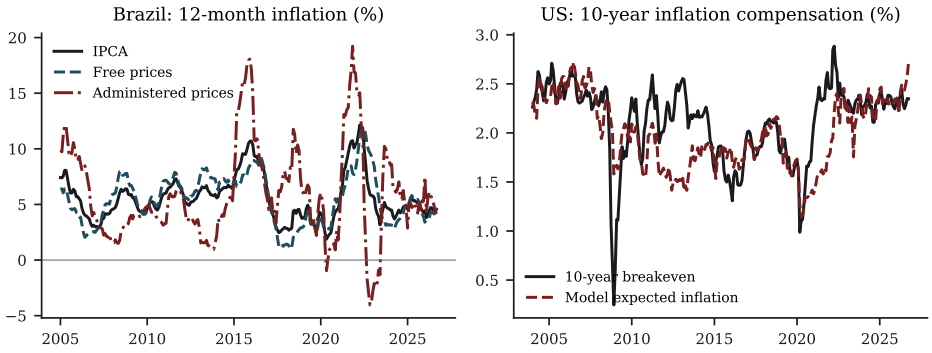


Figure 8.2. Left: twelve-month IPCA inflation, free prices and administered prices, 2005 to 2026 (BCB SGS 433, 11428 and 4449). Right: US ten-year breakeven inflation and the Cleveland Fed's model estimate of expected inflation over ten years, monthly averages (FRED).

Result: two Brazilian surges, two different stories

In 2015 administered prices rose 18.1% as the government corrected electricity and fuel prices that had been held down in previous years, while free prices rose 8.5%: the surge was largely a relative-price correction. In 2022 the order reversed. Free prices rose 9.4% in the year, peaking at 11.9% in June, while administered prices fell 3.8%, reaching −4.0% in November, after tax cuts on fuel and electricity (Figure 8.2). The tax cuts lowered headline inflation mechanically without changing the underlying pressure, which is why the Copom discounted them in its assessment. A bank that looked only at headline would have been misled in both years, in opposite directions.

8.3.2 What a breakeven contains

The breakeven decomposition above says that market inflation compensation equals expected inflation plus an inflation risk premium minus a liquidity premium. The Cleveland Fed estimates expected inflation with a model that combines Treasury yields, inflation data, swaps and surveys, so the difference between the breakeven and the model estimate is a measure of the two premia together.

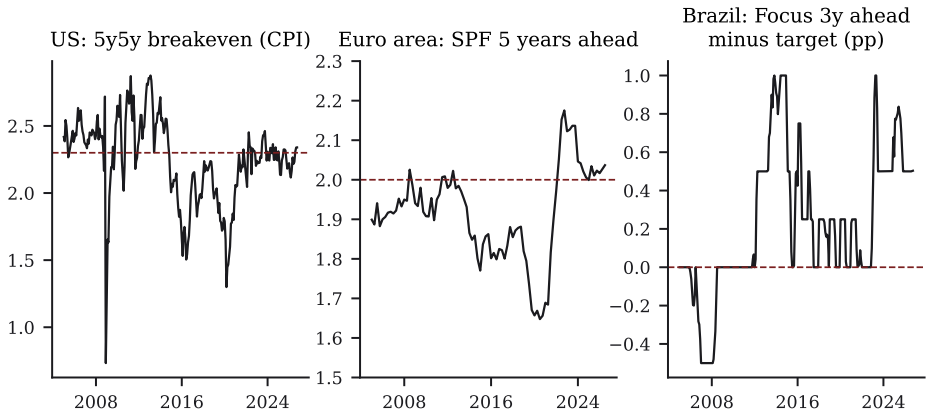


Figure 8.3. Long-run inflation expectations. United States: 5-year 5-year forward breakeven from CPI-linked Treasuries (dashed line at 2.3%, the CPI rate consistent with 2% PCE). Euro area: SPF expectation five years ahead (dashed line at the 2% target). Brazil: Focus median for IPCA three calendar years ahead minus the target for that year.

Result: premia in US inflation compensation

From 2010 to 2019 the ten-year breakeven exceeded the Cleveland Fed's expected inflation by 0.21 points on average. In December 2008 it was 1.33 points below, as investors fled inflation-protected securities for the liquidity of nominal Treasuries, and in March 2020 0.28 points below for the same reason. In April 2022, at the height of the surge, it was 0.67 points above: investors paid to insure against inflation. In October 2026 the breakeven, at 2.35%, is 0.38 points below the model's 2.73% (Figure 8.2). A reader of the breakeven alone would conclude that long-run expectations are well anchored near 2.3%; the model says expected inflation is higher and that the market is charging little for the risk, or paying for liquidity.

8.4 Anchoring

Expectations far ahead should not move with current inflation if the target is credible. Figure 8.3 shows one long-horizon measure for each economy.

Table 8.2. Response of long-run expectations to changes in twelve-month inflation, coefficient β with standard error. United States: 5-year 5-year breakeven on CPI, monthly. Euro area: SPF five years ahead on HICP, quarterly. Brazil: Focus three calendar years ahead on IPCA, monthly.

	2005–2019	2020–2022	2023–2026
United States	0.068 (0.030)	0.065 (0.036)	0.036 (0.026)
Euro area	0.020 (0.010)	0.039 (0.013)	−0.001 (0.012)
Brazil	−0.002 (0.017)	−0.008 (0.013)	−0.054 (0.042)

Result: anchoring through the surge

In the United States the 5-year 5-year breakeven averaged 2.28% from 2005 to 2019 with a standard deviation of 0.35, and 2.17% from 2020 on with a smaller standard deviation of 0.24; its highest monthly average since 2020 is 2.46%. In the euro area the SPF five-year expectation fell to 1.67% at the end of 2019, the low point of a slow drift below target, and is 2.04% in 2026Q3. In Brazil the Focus expectation three years ahead deviated from the target by 0.28 points on average from 2005 to 2019 and 0.37 since 2020; it is 0.50 points above the target in October 2026, the same order as in the fiscal episodes of 2013 to 2015, when it peaked at 1.0.

8.4.1 Do long-run expectations react to inflation news?

A level that stays near target is one test of anchoring. A sharper one asks whether long-run expectations move when current inflation moves. If they are anchored, news about inflation today should say nothing about inflation five years from now. The regression

$$\Delta \pi_t^{LR} = a + \beta \, \Delta \pi_t + e_t \tag{8.4}$$

relates the change in a long-run expectation to the change in twelve-month inflation over the same month (quarter for the euro area). Under perfect anchoring $\beta = 0$.

Result: what moves long-run expectations

In the United States a one-point rise in twelve-month CPI inflation raised the 5-year 5-year breakeven by about 0.07 points before 2020

and during the surge, and by 0.04 since: a small, measurable sensitivity, partly from the inflation risk premium in the breakeven. In the euro area the SPF expectation moved by 0.02 points per point of inflation before 2020 and by 0.04 in 2020 to 2022, then not at all. In Brazil the Focus expectation three years ahead does not respond to monthly inflation news in any period (Table 8.2). Yet it jumped by more than half a point in early 2023 and stayed there. Brazil's de-anchoring was not inflation news passing through to the long run; the timing is consistent with a discrete revision of beliefs about the regime, through fiscal and central bank credibility channels, the shocks of chapters 4 and 7.

The euro area entered the pandemic with expectations drifting below target, the asymmetric-loss bias of chapter 2 in the data, and left it anchored at target. The United States kept long-run expectations anchored through the largest inflation surge in four decades, and disinflated without a recession. Among the measures used here, which mix a breakeven, a professional survey and a market survey at different horizons, Brazil's lies furthest from target: at target through 2022, they jumped to 0.6 to 0.9 points above it in the first half of 2023 and have not come back. The drift coincides with the rise in the neutral rate measured in chapter 7. Both are consistent with a common fiscal risk, though without a formal decomposition the two also carry distinct components: real term and inflation risk premia, liquidity, and the expected policy path.

Chapter 9

Inflation after 2021

In June 2022 US consumer price inflation reached 9.1% over the previous year, the highest since 1981. Euro area inflation peaked at 10.6% in October 2022, the highest since the euro was created. Brazil's IPCA had peaked earlier, at 12.1% in April 2022. Eighteen months before, all three central banks and almost every private forecaster had expected inflation in 2022 to be close to target. The surge was the largest forecasting failure in monetary policy since the 1970s, and it was followed by a second surprise: in the United States inflation came down by more than two thirds while unemployment rose by less than one point.

Both surprises point to the same object, the Phillips curve. Chapter 1 derived it as a straight line with a slope κ fixed by the frequency of price changes and the structure of the labour market. Estimated on the two decades before 2020, that line was flat: slack seemed to have almost no effect on inflation, which is why few expected the tight labour markets of 2021 to matter. This chapter argues that the curve is not a line. It is flat when the economy has slack and steep when the labour market is tight, and supply shocks move it. It derives three mechanisms that bend it, estimates a threshold Phillips curve for the United States and Brazil on data through 2026, finds that the two curves bend in opposite places, re-estimates the wage-price model that Bernanke and Blanchard (2023) built to decompose the surge, and works out what a kinked curve does to optimal policy. The questions are not historical. In 2026 oil prices are above 100 dollars a barrel again, euro area energy inflation is close to 19%, and the US labour market sits on the kink.

9.1 The flat curve before 2020

9.1.1 What the data showed

From 2001 to 2019 US core PCE inflation moved in a band of roughly 1 to 2.5% while unemployment ranged from 3.5% to 10%. Regress core inflation, measured relative to the Cleveland Fed's estimate of one-year expected inflation, on the ratio of job openings to unemployed workers, and add the New York Fed's Global Supply Chain Pressure Index (GSCPI) as a control. On data through 2019 the slope on tightness is -0.17, with a standard error of 0.17, and the regression explains 3% of the variation in inflation. On this evidence there was no Phillips curve at all.

9.1.2 Three explanations for flatness

The literature offered three readings, and they have different implications for what would happen in a very tight labour market.

Anchoring. With expectations firmly at target, firms set prices on the basis of the target rather than of current conditions. The reduced-form curve flattens even if the structural slope is unchanged, and it can steepen again if expectations come loose.

Structural change. Globalisation, the decline of unions and online price comparison made prices less responsive to domestic slack. On this reading the slope really had fallen, and a tight labour market would raise inflation only a little.

Identification. A central bank that offsets demand shocks leaves mostly supply shocks in the data, and supply shocks move inflation and slack in opposite directions. The estimated slope is then a mix of the structural curve and the central bank's reaction function. The point can be made exactly in the model of chapter 2. Under optimal discretion with cost-push shocks, the first-order condition is

$$x_t = -\frac{\kappa}{\lambda_x}\,\pi_t, \tag{9.1}$$

and the equilibrium is (2.11). If demand shocks are fully offset, every observation in the data lies on this targeting rule, not on the Phillips curve. A regression of inflation on the gap recovers

$$\hat{b} = \frac{\operatorname{cov}(\pi, x)}{\operatorname{var}(x)} = -\frac{\lambda_x}{\kappa}, \tag{9.2}$$

the inverse slope of the targeting rule, with a negative sign. With the welfare weight $\lambda_x = 0.019$ and $\kappa = 0.17$, $\hat{b} = -0.11$ on a quarterly gap, the same sign

and order of magnitude as the -0.17 the US data produce before 2020. A central bank that does its job well hides its own Phillips curve. The estimate before 2020 says less about the slope than about the success of policy, and nothing about the slope in a state the economy had not visited.

9.2 Why the curve bends

9.2.1 State-dependent pricing

In Calvo pricing the probability of a price change, $1-\theta$, is constant. In the data, firms change prices more often when the gap between their current and desired price is large, and that gap widens faster when inflation is high. Models in which firms pay a fixed cost to change a price, such as Golosov and Lucas (2007), produce this ***state dependence***: a firm leaves its price unchanged while the gap stays inside a band and resets when it leaves.

A reduced form captures the effect on the slope. Let the quarterly frequency of price changes be the sum of a part driven by idiosyncratic shocks, f_0, and a part driven by aggregate inflation eroding relative prices,

$$1-\theta(\bar{\pi}) = f_0 + \frac{\bar{\pi}/4}{b}, \tag{9.3}$$

where $\bar{\pi}$ is annual inflation and b the width of the band of inaction, so that $b/(\bar{\pi}/4)$ is the number of quarters inflation alone needs to push a price out of the band. Substituting into the slope of the Phillips curve from (1.26) and (1.31),

$$\kappa(\bar{\pi}) = \frac{\big(1-\theta(\bar{\pi})\big)\big(1-\beta\theta(\bar{\pi})\big)}{\theta(\bar{\pi})}\,\Theta\left(\sigma + \frac{\varphi+\alpha}{1-\alpha}\right), \tag{9.4}$$

which rises with inflation through two channels: more firms reset in each quarter, and each resetting firm, expecting to reset again sooner, puts less weight on the future.

Result: the slope rises with inflation

Halving the average duration of prices, from a year to two quarters, raises the slope of the Phillips curve sixfold, from 0.17 to 1.01 (Table 9.1). With state-dependent frequency, moving from 2% to 8% inflation nearly triples the slope, from 0.26 to 0.75, and at 12% it exceeds 1.3. The curve steepens exactly when inflation is already high, so a

Table 9.1. Slope of the Phillips curve under Calvo pricing as the duration of prices falls (top), and under the state-dependent frequency (9.3) with $f_0 = 0.25$ and a band of 10% (bottom). Other parameters as in Table 1.1.

Average price duration (quarters)	4.0	3.0	2.5	2.0
θ	0.75	0.67	0.60	0.50
κ	0.17	0.33	0.54	1.01
Annual inflation (%)	2	4	8	12
Quarterly frequency $1-\theta$	0.30	0.35	0.45	0.55
κ	0.26	0.38	0.75	1.36

shock that starts inflation rising also makes the next unit of demand more inflationary.

Consumer price microdata from 2021 and 2022 show the frequency of price changes rising sharply in the United States and the euro area, with most of the increase coming from price rises rather than cuts. The mechanism was visible in the prices themselves.

9.2.2 The labour market and the v/u ratio

Benigno and Eggertsson (2023) locate a second nonlinearity in the labour market. The model has three steps.

Step 1: matching. Firms post vacancies v_t and unemployed workers u_t search. Hires follow a matching function $m_t = \mu v_t^{\eta} u_t^{1-\eta}$, so the probability that a worker finds a job is $m_t/u_t = \mu\, \vartheta_t^{\eta}$, with $\vartheta_t \equiv v_t/u_t$ labour-market tightness. Tightness is the sufficient statistic for how easy it is to change jobs.

Step 2: two wage regimes. Incumbent workers are paid a wage norm that moves with expected inflation and is rigid, $W_t^n = W_{t-1}^n(1+\pi_t^e)$. A worker can also go to the market, where the wage is set by Nash bargaining and rises with the worker's outside option, which is the job-finding rate. The paid wage is the larger of the two,

$$W_t = \max\{W_t^n,\ W_t^b(\vartheta_t)\}, \qquad \frac{\partial W^b}{\partial \vartheta} > 0. \tag{9.5}$$

When tightness is low the norm binds and wages move only with expected inflation. When tightness rises above the point $\bar{\vartheta}$ where $W^b(\bar{\vartheta}) = W^n$, the bargained wage binds and wages rise with tightness.

Step 3: prices. Firms set prices as a markup over unit labour cost, as in chapter 1, so price inflation follows wage inflation net of productivity. Log-linearising the two regimes and combining them gives the kinked curve

$$\pi_t = \pi_t^e + \gamma_1\,\vartheta_t + \gamma_2 \max\{\vartheta_t - \bar{\vartheta},\, 0\} + \delta\, s_t + \epsilon_t, \tag{9.6}$$

where s_t is a supply-shock variable and γ_1 is small because the norm binds below the threshold. Below $\bar{\vartheta}$ the slope is γ_1; above it, $\gamma_1 + \gamma_2$. Benigno and Eggertsson (2023) put the threshold for the United States near $v/u = 1$, the point at which there is one vacancy for each unemployed worker.

The ratio v/u has an advantage over the unemployment rate or the output gap as a measure of slack: it is observed, it needs no estimate of a natural rate, and it captures movements in vacancies that leave unemployment unchanged, which turn out to be most of the story after 2021. The price is comparability: a slope per unit of v/u is not a slope per point of output gap, and mapping it into the κ of chapter 1 needs a Beveridge and Okun relation that itself shifted after 2020. The estimates below are read as shapes and relative magnitudes, not as values of κ.

9.2.3 Capacity and attention

Two further mechanisms add convexity. The first is physical capacity. If marginal cost rises as output approaches a capacity limit $\bar{y}$,

$$mc_t = -\eta_c \ln\!\left(1 - \frac{y_t}{\bar{y}}\right), \qquad \frac{\partial\, mc_t}{\partial y_t} = \frac{\eta_c}{\bar{y} - y_t}, \tag{9.7}$$

the slope of marginal cost in output, and with it the slope of the Phillips curve, grows without bound as output nears capacity. In 2021 spending shifted from services to goods while ports, shipping and semiconductor plants were near their limits. The GSCPI, which aggregates delivery times, backlogs and shipping costs, and is a proxy for supply-chain pressure rather than a pure supply shock, since congestion also reflects the exceptional demand for goods, reached 4.4 standard deviations above its average in December 2021, against a previous maximum of 1.6 in 2011.

The second is attention. Chapter 5 showed that agents choose how much attention to pay to inflation, and pay more when it is costly to ignore. When inflation is low, expectations barely react to it and the reduced-form curve is flat even if the structural one is not. When inflation becomes salient, the Kalman gain on actual inflation rises, expectations move, and wages indexed to expectations follow. In the US data the Cleveland Fed's one-year expected

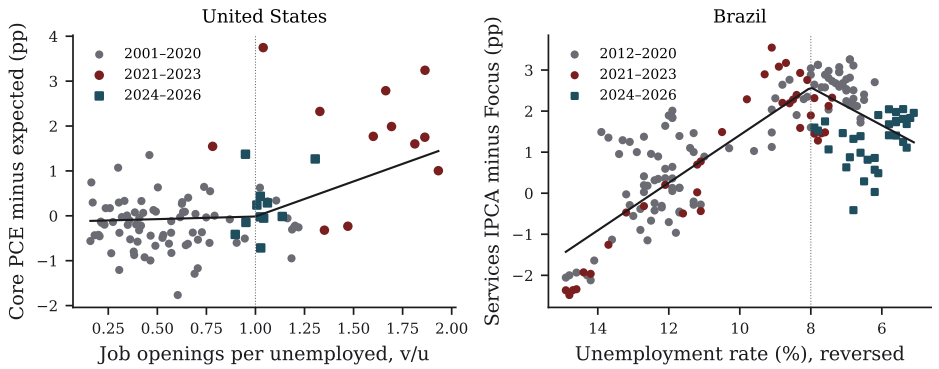


Figure 9.1. Left: US core PCE inflation minus one-year expected inflation against the ratio of job openings to unemployed, quarterly, 2001 to 2026, with the fitted threshold curve at the mean supply-pressure index. Right: Brazilian services inflation minus the Focus twelve-month expectation against the unemployment rate (axis reversed), monthly, 2012 to 2026, with the fitted threshold curve. Sources: BEA, BLS, Cleveland Fed and New York Fed via FRED; IBGE and BCB.

inflation rose from 1.7% at the end of 2019 to 3.8% in mid-2022, while ten-year expectations moved only from 1.7% to 2.3%: short-run attention rose, long-run anchoring held.

9.3 A threshold Phillips curve, 2001 to 2026

9.3.1 United States

Equation (9.6) is estimated on quarterly US data from 2001 to the third quarter of 2026: core PCE inflation (annualised quarterly rate) minus the Cleveland Fed's one-year expected inflation on the left, the ratio of job openings to unemployed with a threshold at one and the GSCPI on the right. Standard errors are Newey-West with four lags.

Result: the US curve has a kink

With a threshold at $v/u = 1$ the slope below the threshold is 0.12 (standard error 0.30), indistinguishable from zero, and above it rises by 1.45 (0.75) to 1.56 (0.54) per unit of v/u. A one-standard-deviation rise in supply-chain pressure adds 0.37 points (0.15) to core inflation relative to expectations (Table 9.2, Figure 9.1). In the first quarter of

Table 9.2. Phillips curve estimates. United States: quarterly, 2001Q1 to 2026Q3, dependent variable core PCE inflation minus one-year expected inflation. Brazil: monthly, March 2012 to August 2026, dependent variable twelve-month services IPCA minus the Focus twelve-month expectation. Newey-West standard errors in parentheses.

	United States			Brazil
	Linear	Threshold	Pre-2020	Threshold
Slope, slack region	0.64 (0.25)	0.12 (0.30)	−0.17 (0.17)	−0.58 (0.09)
Slope, tight region		1.56 (0.54)		0.46 (0.15)
Supply pressure (GSCPI)	yes	0.37 (0.15)	yes	
Threshold		$v/u = 1$		$u = 8\%$
R^2	0.43	0.46	0.03	0.67
Observations	102	102	76	174

2022, when v/u averaged 1.86 and core inflation ran 3.2 points above expectations, the model attributes 1.4 points to tightness beyond the threshold and 1.2 to supply pressure, for a fitted gap of 2.5.

The Cleveland Fed's expectation refers to CPI, while the dependent variable uses core PCE, so the gap carries a level wedge between the two indices. Re-estimating with core CPI inflation, which matches the expectation's concept, leaves the shape intact and sharpens it: the slope below the threshold is 0.40 (0.43), the extra slope above it 1.88 (0.88), so 2.27 in the tight region, with supply pressure at 0.27 (0.19) and an R^2 of 0.43. The pre-2020 slope is 0.37 (0.31) with an R^2 of 0.03, flat again. The average gap is 0.33 points with core CPI against 0.07 with core PCE, the composition wedge the mismatch introduces.

The threshold is imposed, not estimated, so its robustness matters. Moving it to 0.8, 0.9, 1.1 or 1.2 gives an extra slope above the threshold of 1.42, 1.48, 1.58 and 1.82, with the slope below it between −0.12 and 0.24 and never significant. The kink is a feature of the data, not of where it is placed.

9.3.2 Could the kink have been seen in 2020?

A curve fitted after the event can always be bent to pass through it. The fair test is out of sample: estimate each specification on 2001 to 2020 and ask how well it predicts core inflation relative to expectations in 2021 to

Table 9.3. Out-of-sample test. Each specification estimated on 2001Q1 to 2020Q4 and used to predict US core PCE inflation minus one-year expected inflation in 2021Q1 to 2023Q4, given the actual paths of slack and supply pressure. Actual mean gap over the test period: 1.77 points.

Specification	Slopes to 2020	RMSE	Mean error
Linear in v/u	−0.09	2.05	1.78
Threshold at $v/u = 1$	−0.03, extra −0.46	2.21	1.96
Unemployment gap (CBO)	0.04	2.10	1.82
Supply pressure only		1.98	1.70

2023. The training sample is not empty of tight labour markets: v/u was above one in nine quarters of 2018 and 2019.

Result: the kink was not visible before it was used

Every specification estimated before 2021 under-predicts the surge by about 1.8 points a quarter on average, and none beats a model with no slack at all (Table 9.3). The threshold model does worst: the nine tight quarters of 2018 and 2019 came with low inflation, so the pre-2021 data put a negative slope above the kink. The kink is identified by the surge itself.

This is the identification problem of the previous section in its sharpest form. The economy of 2018 and 2019 was barely above the threshold, with v/u at most 1.22 on a quarterly average, and expectations were anchored after a decade of below-target inflation; the steep region had not been visited. The lesson is not that the kink is an artefact. Over the full sample the threshold model fits better than any linear alternative, and a linear curve in the unemployment gap has a slope of −0.06 (0.05) and an R^2 of 0.37, against 0.46 for the threshold in v/u. The lesson is that theory, not the pre-2020 data, was the only available guide to how the economy would behave far above the threshold, and the theory of the previous section was available in 2021.

9.3.3 What the curve says for 2026

With the full-sample threshold estimate and supply pressure at its September 2026 reading of 1.3, the fitted gap of core inflation over expectations is 0.44 points at $v/u = 1.0$, 0.42 at 0.8 and 0.76 at 1.2. At the kink, a loos-

ening of the labour market buys almost nothing; a tightening of the same size costs 0.3 points. Expected inflation, the other term, has not yet taken in the oil shock: the Cleveland one-year measure averaged 2.5% in the third quarter of 2026, down from 2.8% a year earlier. If it follows oil, as it did in 2021, the gap term matters less than the level it is added to.

9.3.4 Brazil

Brazil's labour-market data begin in 2012 with the continuous household survey, PNAD Contínua, and there is no long series for vacancies. The estimate uses the unemployment rate, monthly from March 2012 to August 2026, with a threshold at 8%, near the middle of the sample range. The dependent variable is twelve-month services inflation minus the Focus twelve-month expectation. Services are the component most driven by domestic wages; administered and food prices are dominated by policy decisions, the weather and the exchange rate (chapter 8).

Result: the Brazilian curve bends the other way

A linear fit gives a slope of −0.36 per point of unemployment (0.09). Split at 8%, the slope is −0.58 (0.09) when unemployment is above the threshold and +0.46 (0.15) below it, with an R^2 of 0.67 (Table 9.2). Between 2012 and 2021 a fall in unemployment raised services inflation relative to expectations by half a point per point. From 2022 the relation disappears: on 2022 to 2026 alone the slope is 0.08 (0.12) and the R^2 0.03. Unemployment fell from 14.9% in September 2020 to a record low of 5.1% in December 2025 and was 5.3% in August 2026, while services inflation relative to expectations fell from about 3 points in 2022 to between 0.3 and 2 points.

Two readings fit this, and they have opposite policy implications. On the first, the natural rate of unemployment fell: the 2017 labour reform, the spread of platform work and a shift in the composition of employment mean that 6% unemployment in 2025 is no tighter than 8 or 9% was in 2014, and measured unemployment overstates tightness. On the second, the pressure went somewhere the dependent variable does not see. The gap is services inflation minus expected inflation, and expected inflation itself rose: the Focus twelve-month median averaged 4.2% from 2024 to 2026, well above the 3% target, while services inflation held at 5.5%. If tightness raised expectations, and expectations raised services inflation one for one, the gap

would show nothing and the level would show everything. The evidence is consistent with the second reading, though with an unemployment-only threshold, shifts in u^*, in expectations and in the composition of employment remain hard to separate. Services inflation never returned near target in the period of record-low unemployment, and the Copom's statements through 2025 cited the strength of the labour market and the de-anchoring of expectations together as reasons to tighten.

The US and Brazilian estimates therefore make different points. In the United States, with long-run expectations anchored, tightness shows up in the gap and the curve is kinked upward at $v/u = 1$. In Brazil, with expectations moving, tightness shows up in expectations, and a Phillips curve written relative to expected inflation flattens precisely when the labour market is tightest. This is the anchoring reading of the flat curve from the opening section, observed in the data rather than assumed.

9.4 The Bernanke-Blanchard decomposition

A single-equation Phillips curve cannot say how a shock to energy prices becomes a rise in wages, or how long it lasts. Bernanke and Blanchard (2023) built a small dynamic model for that purpose and published its code; in 2024 the authors extended it to eleven economies with research teams from their central banks. This section re-estimates the US version on data through the second quarter of 2026 and uses it to split inflation into its sources.

9.4.1 The four equations

All variables are annualised quarterly rates. Wage growth g_t^w (the Employment Cost Index) depends on its own lags, on lagged short-run expected inflation π^{e1}, on tightness $\vartheta = v/u$, and on a catch-up term cu_t, the gap between price inflation over the past year and what was expected a year earlier:

$$g_t^w = \sum_{i=1}^{4} a_i \, g_{t-i}^w + \Big(1 - \sum_i a_i\Big)\pi_{t-1}^{e1} + \sum_{i=1}^{4} b_i \, \vartheta_{t-i} + \sum_{i=1}^{4} c_i \, cu_{t-i} + \varepsilon_t^w. \quad (9.8)$$

Price inflation π_t (CPI) depends on its lags, on wage growth net of trend productivity g_t^{wp} (an eight-quarter average of productivity growth), and on three supply variables: energy and food price growth relative to wages (*grpe*,

Table 9.4. Re-estimated Bernanke-Blanchard model, US, quarterly. Wage and price equations 2002Q2 to 2026Q2; expectations equations 1991Q2 to 2026Q2. Sums of lag coefficients. Short-run expectations from the Cleveland Fed one-year series, long-run from its ten-year series.

Equation	Coefficient sum	Estimate	R^2
Wages	own lags	0.73	0.50
	short-run expectations	0.27	
	v/u	0.31	
	catch-up	−0.02	
Prices	own lags	0.87	0.96
	wages net of productivity	0.13	
	relative energy price	0.018	
	relative food price	0.082	
	shortages (GSCPI)	0.13	
Short-run expectations	own lags	−0.14	0.42
	long-run expectations	1.04	
	past inflation	0.10	
Long-run expectations	own lags	0.99	0.99
	past inflation	0.01	

$grpf$) and a shortage index sh, here the GSCPI:

$$\pi_t = \sum_{i=1}^{4} d_i\,\pi_{t-i} + \sum_{i=0}^{4} e_i\,g^{wp}_{t-i} + \sum_{i=0}^{4}\big(f_i\,grpe_{t-i} + h_i\,grpf_{t-i} + k_i\,sh_{t-i}\big) + \varepsilon^p_t, \quad (9.9)$$

with $\sum d_i + \sum e_i = 1$. Short-run expectations depend on their own lags, on long-run expectations and on past inflation; long-run expectations depend on their own lags and on past inflation. Each expectations equation has coefficients summing to one, so that in a steady state expectations equal inflation.

The homogeneity constraints make the model a Phillips curve in the accelerationist sense: a permanent rise in tightness raises wage growth permanently and, through the price equation, inflation. The constraints are imposed by estimating each equation in deviations, for example $g^w_t - \pi^{e1}_{t-1}$ on $g^w_{t-i} - \pi^{e1}_{t-1}$, and recovering the omitted coefficient as one minus the sum of the others.

9.4.2 Estimates

Three features of Table 9.4 drive the decomposition. Price inflation is persistent: lagged inflation carries a weight of 0.87, so a shock to prices de-

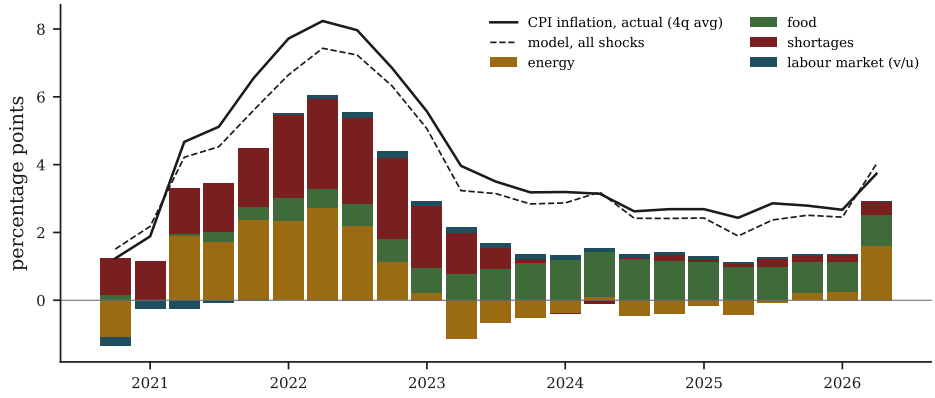


Figure 9.2. Contributions to US CPI inflation from the re-estimated Bernanke-Blanchard model, four-quarter averages, 2020Q4 to 2026Q2. The solid line is actual CPI inflation, the dashed line the model simulated with all exogenous variables at their observed values and residuals at zero. Sources: BLS, BEA, Cleveland Fed and New York Fed via FRED; author's estimates.

cays slowly. Wages respond to tightness, with a long-run effect of $0.31/(1 - 0.73) = 1.15$ points of wage growth per unit of v/u, but pass into prices with a small weight. And long-run expectations barely respond to inflation: a coefficient of 0.01 on past inflation means that the largest surge in four decades moved ten-year expectations by a fraction of a point. The catch-up term, which in the original estimate captured workers recovering lost real wages, is close to zero on the longer sample. Short-run expectations in the Cleveland series follow long-run expectations almost one for one, with past inflation adding 0.10 and their own lags, which the restriction makes the residual, entering at -0.14, so the series reverts quickly to the long-run anchor; this is a property of a model-based expectations series, which leans on long-run anchors.

9.4.3 Decomposition

The decomposition runs the model forward from the first quarter of 2020 with all residuals set to zero, so that inflation is driven only by the observed paths of the exogenous variables. Each source is then removed in turn, by holding energy and food relative prices and v/u at their 2018 to 2019 averages and the shortage index at zero, and the difference between the two simulations is that source's contribution.

Result: the US surge was mostly supply, made persistent by the price equation

In the second quarter of 2022, when four-quarter CPI inflation averaged 8.2%, the model attributes 2.7 points to energy, 2.7 to shortages, 0.6 to food and 0.1 to the labour market (Figure 9.2). The simulation with all shocks gives 7.4% against an actual 8.2%. By the end of 2024 energy had turned negative (−0.4), shortages had faded to 0.2, and the largest remaining source was the slow decay of the 2022 food shock (1.2). In the second quarter of 2026 energy is again the largest source, at 1.6 points, as the oil shock of 2026 enters CPI.

The labour-market contribution is smaller here than in the original estimate, which found it growing into the main source of persistence in 2023. Two differences explain most of the gap. The sample now runs three years longer, through a period in which v/u fell sharply and wage growth fell from 5.5% (four-quarter average, mid-2022) to 3.1% (mid-2026), which weakens the estimated link from wages to prices. And the shortage measure here is the GSCPI, which loads more of 2021 and 2022 onto shortages than the index of internet searches for "shortage" used in the original. The two versions agree on the headline: the surge began as a supply shock, and the policy question was whether a tight labour market would turn it into a wage-price spiral. In the United States, with long-run expectations anchored, it did not.

The threshold curve and the dynamic model are complements. The threshold curve explains core inflation, from which energy and food are removed, and finds a large role for tightness once the curve is allowed to bend. The dynamic model explains headline inflation, imposes a linear effect of v/u on wages, and finds most of the headline surge in energy, food and shortages. Tightness mattered most where energy did not reach, in core services, and most when it was far above the kink.

9.5 The soft landing

The kinked curve also explains the second surprise. If the curve is steep above the threshold, reducing tightness from far above it to near it lowers inflation a lot while doing little to unemployment, because what falls is vacancies, not employment. The Beveridge curve, the relation between the vacancy rate and the unemployment rate, shows the path (Figure 9.3, left).

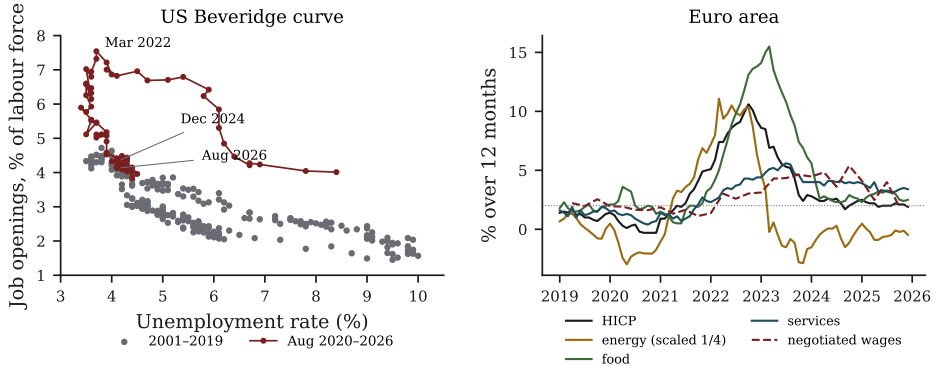


Figure 9.3. Left: US Beveridge curve, job openings as a share of the labour force against the unemployment rate, monthly, 2001 to 2019 (dots) and April 2020 to August 2026 (line). Right: euro area HICP inflation and its components, and the ECB indicator of negotiated wages, twelve-month rates, 2019 to 2025; energy scaled by one quarter. Sources: BLS via FRED; ECB Data Portal.

Job openings rose from 4.1% of the labour force at the end of 2019 to 7.5% in March 2022, the highest in the series, and fell back to 4.1% by August 2026: from 12.3 million to 7.1 million openings. Over the same period the unemployment rate rose from 3.4% at its low in April 2023 to 4.1% at the end of 2024 and 4.2% in September 2026. The ratio of openings to unemployed fell from 2.04 to 1.01.

Result: disinflation along the Beveridge curve

Headline disinflation in 2022 to 2024 was driven mostly by the reversal of the supply shocks, as the decomposition above shows. The labour market's part was in core: moving down the steep part of the Phillips curve, by reducing vacancies rather than raising unemployment, removed the domestic persistence that would otherwise have kept core inflation high once energy and shortages faded. With the pre-2020 slope the same fall in core inflation would have required an increase in unemployment of several points. The soft landing was in large part a property of where the economy started.

The same logic carries a warning. With v/u at 1.01 in August 2026, the labour market sits on the kink. Further cooling would move it onto the flat part of the curve, where each point of disinflation costs much more

unemployment. The cheap part of the disinflation is over, and it ended as a new energy shock arrived.

9.6 The euro area: an energy shock with a lagged wage response

The euro area surge had a different composition (Figure 9.3, right). Energy inflation reached 44.3% in March 2022, driven by the gas price after the invasion of Ukraine; the euro area imports most of its gas, and much of its electricity price is set at the margin by gas-fired plants. Headline inflation followed to 10.6% in October 2022. Food inflation, which embeds energy and fertiliser costs with a lag, peaked later, at 15.5% in March 2023. Core inflation peaked at 5.7% in March 2023 and services at 5.6% in July 2023, more than a year after energy.

The labour market tightened less than in the United States: unemployment fell from 7.5% at the end of 2019 to a low of 6.2% in November 2024. But wages are set in multi-year collective agreements, and they responded late. The ECB's indicator of negotiated wages rose from 2.0% in the last quarter of 2019 to 5.4% in the third quarter of 2024, two years after the energy peak, before falling to 1.9% a year later. This is the catch-up mechanism of (9.8), slower and larger than in the United States because contracts are longer. The ECB raised its deposit rate from -0.50% in July 2022 to 4.00% in September 2023, cut it to 2.00% by June 2025, and raised it to 2.25% in June 2026 and 2.50% in September 2026.

The 2026 episode repeats the first stage on a smaller scale. Eurostat's flash estimate for September 2026 put headline inflation at 3.8%, energy at 18.8%, services at 3.2% and core at 2.5%. With unemployment at 6.5% and negotiated wage growth back near 2%, the euro area is on the flat part of its curve, and the question for the ECB is the one it faced in 2022: whether energy reaches wages through catch-up.

9.7 Brazil: early hikes, a second round, and record-low unemployment

Brazil entered the surge with a Selic of 2%, its lowest ever, set in August 2020. Twelve-month IPCA inflation was already 4.5% in December 2020, driven by food and the exchange rate, while services inflation was 1.5%, held down by the pandemic. The Copom began raising the Selic in March 2021, a year

before the Fed, and took it to 13.75% by August 2022. Inflation peaked at 12.1% in April 2022, pushed by fuel, by electricity during a drought that emptied hydroelectric reservoirs, and by food. Services inflation, the part the Phillips curve speaks to, rose from 1.5% in January 2021 to 8.9% in July 2022 as unemployment fell from 14.4% to 9.1%.

The early start paid off in the first round. Headline inflation fell to 3.2% by June 2023, helped by tax cuts on fuel and electricity in 2022, and the Copom cut from August 2023. The second round came through expectations. As unemployment fell below 7% in mid-2024, services inflation stopped falling at around 5% and the Focus twelve-month median rose from 3.6% in May 2024 to 5.7% in February 2025. The Copom resumed hiking in September 2024, from 10.50%, raised the rate by 100 basis points at each of three consecutive meetings from December 2024, and reached 15% in June 2025, the highest since 2006. It held there until March 2026, with unemployment at a record 5.1% in December 2025, and then cut in steps to 13.75% by September 2026. Twelve-month IPCA inflation was 4.6% in September, just above the 4.5% upper limit of the tolerance band.

Result: in Brazil the labour market works through expectations

The second Brazilian tightening responded to expectations, not to a measured gap: the BCB's published reaction function (chapter 3) has no gap term and a coefficient of 2.03 on the distance of Focus expectations from the target. The Phillips curve estimates explain why that design suits Brazil. With expectations that move with the labour market, the expectations term carries the information a gap would carry, and it is observed every week. Holding expectations down at record-low unemployment has required an ex-ante real rate, the Selic less the Focus twelve-month expectation, of about 11% in 2025 and 9% in late 2026.

9.8 Policy with a kinked curve

9.8.1 Precautionary policy

A linear Phillips curve and a quadratic loss give certainty equivalence: the central bank sets its expected gap as if it knew the state of the economy. A kink breaks this. Let the central bank choose expected tightness $\bar{\vartheta}$, with realised tightness $\vartheta = \bar{\vartheta} + \zeta$ and $\zeta \sim N(0, \sigma^2)$ reflecting shocks and measurement error. Inflation relative to target is $g(\vartheta) = a(\vartheta - 1) + b \max\{\vartheta - 1, 0\}$,

Table 9.5. Optimal expected tightness $\bar{\vartheta}$ under a linear curve ($a = 0.64$, $b = 0$) and the estimated kinked curve ($a = 0.12$, $b = 1.45$), with $\lambda = 0.25$ and $\vartheta^* = 1.2$, as uncertainty about realised tightness rises. Computed by simulation with 200,000 draws.

Standard deviation of ζ	0	0.15	0.30
Linear curve	1.08	1.08	1.08
Kinked curve	1.02	0.92	0.79

and the loss adds a penalty for tightness away from a preferred level ϑ^*:

$$\mathcal{L}(\bar{\vartheta}) = E\big[g(\bar{\vartheta} + \zeta)^2\big] + \lambda\, E\big[(\bar{\vartheta} + \zeta - \vartheta^*)^2\big]. \tag{9.10}$$

The first-order condition is

$$E\big[g(\bar{\vartheta} + \zeta)\, g'(\bar{\vartheta} + \zeta)\big] + \lambda(\bar{\vartheta} - \vartheta^*) = 0. \tag{9.11}$$

With $b = 0$, $g\,g' = a^2(\vartheta - 1)$ is linear, the expectation depends only on $\bar{\vartheta}$, and σ drops out: $\bar{\vartheta} = (a^2 + \lambda\vartheta^*)/(a^2 + \lambda)$. With $b > 0$, $g\,g'$ is convex in ϑ, so by Jensen's inequality its expectation rises with σ: a draw above the kink costs more than an equal draw below it saves. The optimal $\bar{\vartheta}$ falls as uncertainty grows.

Result: uncertainty plus a kink calls for caution on the tight side

With the US estimates, a central bank uncertain about tightness by 0.3 units of v/u should aim for a ratio near 0.8 rather than 1.0 (Table 9.5). Under a linear curve the same uncertainty changes nothing.

9.8.2 Three lessons

Asymmetric risk. The loss functions of chapter 2 are symmetric in the gap, but the inflation a gap produces is not. The Fed's 2020 framework made the employment term one-sided, responding to shortfalls from maximum employment rather than to deviations in both directions. With a linear curve that is a matter of preference. With a kinked curve it reduced the role of preemptive tightening, based on uncertain estimates of maximum employment, in exactly the region where tightness is most inflationary, though it left the response to inflation itself in place. The 2025 revision dropped the shortfalls language and restored symmetry; the argument above suggests the response should if anything be stronger above the kink than below it.

The cost of waiting. When the labour market is tight and the curve steep, a delay lets inflation rise quickly, and expectations, more attentive when inflation is high, begin to follow. The Fed's first hike came in March 2022, when v/u had been above one for most of a year and core PCE inflation was above 5%. The Copom started a year earlier, and its first-round disinflation was faster.

The cheap disinflation and its end. The same steepness makes disinflation cheap at first: reducing demand from far above the kink lowers inflation with little cost in unemployment. Near the kink the trade-off worsens sharply. Policy that tightens early and stops near the kink exploits the shape of the curve; policy that tightens late and overshoots into the flat region pays twice. In late 2026 the United States is at the kink, the euro area below it with slack, and Brazil at record-low unemployment with expectations above target, and a common energy shock is arriving in all three.

Chapter 10 asks how the Fed, the ECB and the BCB responded to these conditions, by estimating their reaction functions and evaluating them on the welfare loss.

Chapter 10

Estimated Reaction Functions

Chapter 3 asked which rules a central bank should follow. This chapter asks which rules the Fed, the ECB and the BCB have followed, by estimating their reaction functions on their own decisions, and then evaluates the estimates on the welfare loss of the model, so that the question of what they did can be compared with the question of what they should have done.

An estimated reaction function is a description, not a commitment. No committee votes on a formula, and every one of the three banks says it does not follow one. But decisions that respond systematically to inflation and slack leave a systematic pattern in the data, and markets price that pattern: a bond trader who knows the reaction function can forecast the policy rate from the inflation and labour market data alone. The estimates below show how stable that pattern has been, how far policy departed from it after 2020, and one case of replication: re-estimated from public data at the BCB's own quarterly frequency, the BCB's rule returns inertia coefficients within a standard error of the published ones and a response to expected inflation that brackets the published 2.03.

10.1 From rules to reaction functions

10.1.1 Partial adjustment

Taylor's (1993) rule sets the rate directly at its target level. Central banks move in steps instead, and the standard way to capture that is ***partial adjustment***. Let the target level be

$$i_t^* = c + \phi_\pi \, \pi_t + \phi_x \, x_t, \tag{10.1}$$

and let the bank close a fraction $1-\rho$ of the distance between the current rate and the target in each period:

$$i_t = \rho\, i_{t-1} + (1-\rho)\, i_t^{*} + \varepsilon_t. \tag{10.2}$$

The coefficients ϕ_π and ϕ_x are long-run responses: a permanent one-point rise in inflation raises the rate by $(1-\rho)\phi_\pi$ on impact and by ϕ_π eventually. With quarterly data, estimates of ρ are typically between 0.75 and 0.95, so the impact response is a fifth or less of the long-run response. The Taylor principle concerns the long-run coefficient: $\phi_\pi > 1$.

The intercept collects the neutral rate and the target, $c = r^{*} + (1-\phi_\pi)\pi^{*}$, so a rule with a constant intercept assumes a constant neutral rate. Chapter 7 showed that r^{*} fell by more than a point in the United States and by about two in the euro area between 2007 and 2019. A rule estimated before the fall prescribes rates that are too high after it, and the error looks like a policy deviation. The alternative is to write the target as

$$i_t^{*} = r_t^{*} + \pi_t + a\,(\pi_t - \pi^{*}) + b\, x_t, \tag{10.3}$$

with r_t^{*} an external estimate and $\phi_\pi = 1 + a$. This is the form of the rules the Fed reports in its Monetary Policy Report (chapter 3), with r_t^{*} fixed at the median of the FOMC's longer-run projections.

10.1.2 Forecasts, real-time data and identification

Three problems complicate estimation.

Forecasts. Central banks respond to expected inflation, not current inflation, because policy acts with a lag. Clarida et al. (2000) replace π_t in (10.2) with $\mathbb{E}_t \pi_{t+h}$ and estimate by generalised method of moments, using lagged variables as instruments for the unobserved forecast. When a survey or the bank's own forecast is observed, it can enter directly. The BCB's published rule uses Focus expectations; the ECB estimates below use the ECB Survey of Professional Forecasters; the Fed's Summary of Economic Projections (SEP) provides a third kind of evidence, the committee's own forecasts of inflation, unemployment and the policy rate.

Real-time data. Orphanides (2001) showed that the Fed's decisions in the 1970s look much less inflationary when the rule is estimated with the output gap the Fed saw at the time, which was large and negative, rather than the revised gap, which was close to zero. A rule fitted to revised data attributes to the policymaker information it did not have.

Identification. Cochrane (2011) made a deeper point. In the determinate equilibrium of the New Keynesian model, the response to inflation is

not identified from the equilibrium data. The argument is short. With the rule $i_t = \phi_\pi \pi_t + v_t$ and an AR(1) monetary shock v_t, the model of chapter 1 has the solution $\pi_t = \psi_v v_t$ for some ψ_v that depends on ϕ_π. Then $i_t = (\phi_\pi \psi_v + 1) v_t$, and both i and π are proportional to v. A regression of i on π recovers $(\phi_\pi \psi_v + 1)/\psi_v$, which is not ϕ_π. The off-equilibrium threat that delivers determinacy, a large response to inflation that would occur if inflation left its path, never shows up in the data, because in a determinate equilibrium it is never carried out. Estimated rules escape the problem only to the extent that inflation moves for reasons other than the policy shock itself, such as supply shocks that the bank chooses to accommodate partly. The practical consequence, which the estimates below confirm, is that inflation coefficients are least precise when inflation has been most stable.

10.2 The Fed

10.2.1 Quarterly data, the lower bound and the shadow rate

Two choices determine whether an estimated Fed rule means anything. The first is the inflation variable. The Fed sets policy on its forecast of inflation, and a rule written on current inflation attributes to the committee a response to numbers it was looking through; the estimates below use expected core CPI inflation over the next four quarters, from the Greenbook to 2006 and from the Survey of Professional Forecasters (SPF) median from 2007, against an objective of 2.3%, the core CPI equivalent of 2% PCE inflation. The second is the lower bound. From December 2008 to December 2015, and from March 2020 to March 2022, the funds rate was pinned near zero while the Fed eased through asset purchases and guidance (chapters 6 and 14). A rule that sees a rate of zero sees no easing. The ***shadow rate*** of Wu and Xia (2016), the short rate implied by the whole yield curve under a term-structure model in which the observed short rate is the maximum of the shadow rate and the lower bound, measures the stance in those years. It reached -2.9% in the second quarter of 2014 and -1.9% in mid-2021. The estimates use it in place of the funds rate during both lower-bound episodes, with the target written around the time-varying r^* as in (10.3).

Result: on quarterly data the inflation response is not identified

With expectations, a moving r^* and the shadow rate, the quarterly estimates pin down the response to the output gap, about 1.5 with

Table 10.1. Fed reaction function, quarterly: expected core CPI inflation (Greenbook to 2006, SPF from 2007), CBO output gap, median of the US neutral-rate suite (chapter 7), Wu-Xia shadow rate at the lower bound. Long-run response to inflation $\phi_\pi = 1 + a$. Standard errors in parentheses.

Sample	ρ	ϕ_π	ϕ_x	R^2	n
1990Q1 to 2008Q3	0.85 (0.04)	0.15 (0.43)	1.42 (0.39)	0.95	75
1990Q1 to 2026Q2, shadow rate	0.90 (0.02)	0.41 (0.56)	1.51 (0.28)	0.98	146

a standard error of 0.3 on the full sample, and leave the response to expected inflation undetermined: point estimates below one with standard errors as large as the coefficients (Table 10.1). Expected core inflation in the SPF and the Greenbook stayed between 1.5% and 3% for most of three decades, and the identification problem of the previous section takes the form it predicts.

One way forward is better information per observation: the decision taken at each FOMC meeting, set against what the staff told the committee before that meeting. More frequent data help only if they carry information the quarterly averages lose; they do not by themselves solve the identification problem.

10.2.2 What the Fed saw: real-time Greenbook data

The estimates above use revised data: today's vintage of core PCE inflation and today's CBO estimate of the output gap. The Fed set rates on what its staff knew and forecast at the time, which the Federal Reserve Bank of Philadelphia publishes, with a five-year lag, as the Greenbook (now Tealbook) data set. For each FOMC meeting from August 1987 to October 2008, 167 meetings, it gives the staff's real-time estimate of the output gap for the current quarter and its forecast of core CPI inflation for the next four quarters. The rule is re-estimated at meeting frequency, with the target rate decided at each meeting on the left, once with the Greenbook variables and once with ex-post data for the same dates (realised core CPI inflation over the next four quarters and the current CBO gap). The ex-post version differs in two ways at once: the gap is a revised vintage, and the inflation variable is the realisation rather than the forecast, so it adds the forecast error. Two further versions change one at a time.

Table 10.2. Fed reaction function at FOMC meeting frequency, August 1987 to October 2008: real-time Greenbook data, one-at-a-time changes, and ex-post data. Inflation: core CPI over the next four quarters, Greenbook forecast or realised. Gap: Greenbook real-time estimate or current CBO estimate. Smoothing per meeting (about six weeks). Standard errors in parentheses. Sources: Federal Reserve Bank of Philadelphia Tealbook data set; FRED; author's estimates.

Data	ρ	ϕ_π	ϕ_x	R^2	Meetings
Real-time, 1987–2008	0.92 (0.02)	1.57 (0.19)	0.86 (0.13)	0.991	167
Forecast, revised gap	0.91 (0.02)	1.40 (0.16)	1.34 (0.16)	0.992	167
Realised π, real-time gap	0.90 (0.02)	1.69 (0.16)	0.85 (0.10)	0.992	167
Ex post, 1987–2008	0.93 (0.01)	1.36 (0.21)	1.44 (0.22)	0.992	167
Real-time, 1987–1999	0.88 (0.02)	1.32 (0.14)	0.63 (0.08)	0.991	96
Real-time, 2000–2008	0.90 (0.04)	−0.05 (1.40)	1.48 (0.40)	0.983	71

Result: on the data it had, the Fed followed the Taylor principle

With the Greenbook forecasts the long-run response to expected inflation is 1.57 with a standard error of 0.19, clearly above one, and the response to the gap is 0.86 (Table 10.2). With ex-post data for the same meetings, the inflation response falls to 1.36 and the gap response rises to 1.44. Changing one input at a time shows where the difference comes from: replacing the real-time gap with the revised one gives 1.40 and 1.34, almost the whole change, while replacing the inflation forecast with realised inflation gives 1.69 and 0.85, no change in the gap response. Revisions to the output gap make the Fed look less concerned with inflation and more with output than it was. The real-time gap was revised up by 0.54 points on average, with a standard deviation of 0.97, and the correlation between the two measures is 0.84.

The split sample repeats the identification problem: from 2000 to 2008, with forecast core inflation between 1.5% and 2.5% throughout, the inflation coefficient cannot be estimated, while the gap coefficient is precise.

The real-time data settle one argument. Taylor (2007) argued that the Fed held rates "too low for too long" in 2003 and 2004, contributing to the housing boom: with revised data, the Taylor (1993) rule called for 3.0% on average in 2003 and 4.2% in 2004, when the funds rate averaged 1.1% and 1.5%. The Greenbook tells a different story. In 2003 the staff estimated the output gap at −2.4%, a point more slack than the CBO estimate today, and

forecast core inflation of 1.65%, with deflation discussed as a risk. On those numbers the real-time rule called for 0.8% in 2003 and 1.7% in 2004, within 0.3 points of what the Fed did. The deviation is a property of the revised data, not of the policy the Fed thought it was following. Whether the Fed should have seen the economy more accurately is a different question from whether it departed from its rule.

10.2.3 How late was the Fed in 2021?

The real-time rule gives the benchmark for 2021, once it is written around the current r^* and fed with the expectations available at the time: $i_t^* = r_t^* + \pi_t^e + 0.57(\pi_t^e - 2.3) + 0.86\,x_t$, with π^e the SPF median for core CPI over the next four quarters and r^* the median of the neutral-rate suite of chapter 7, which fell from 2.2% in 2007 to 0.7% at the end of 2019 and 0.3% at the end of 2021, and has since recovered to 1.1%. Its value, the ***notional*** rate the Fed would have set had it closed the whole gap at once, was 1.8% in the first quarter of 2021, with expected core inflation at 2.1% and an output gap of −0.5%. By the fourth quarter, with expected inflation at 2.6% and a gap of +2.1%, it was 4.9%. The funds rate was 0.08%, and the shadow rate −1.6%, because asset purchases were still running (Figure 10.1, left). The Taylor (1993) coefficients give 4.2% for the same quarter. On either rule the Fed was four to six and a half points looser than its own pre-2008 behaviour implied, measured against the funds rate and the shadow rate respectively.

The Fed began raising rates in March 2022. Under the framework adopted in August 2020 (chapter 2), the delay was deliberate: the FOMC had said it would not raise rates until inflation had reached 2% and was on track to exceed it moderately, and until labour market conditions had reached its assessment of maximum employment. By the time both conditions were judged met, the economy was far above the threshold of the nonlinear Phillips curve of chapter 9. The framework did what it said: it removed the pre-emptive response to a tight labour market that the 1987 to 2008 rule embodied. The estimated rule makes the cost of that choice visible. The funds rate reached the rule's notional level in the third quarter of 2023, at 5.3% against 5.2%, and stayed close to it through 2024. In 2025 and 2026, as the oil shock lifted expected core inflation to between 2.6% and 3.2% and the CBO gap stayed above one, the rule moved back above the funds rate: 5.2% against 3.6% in the second quarter of 2026, a gap the September 2026 increase began to close.

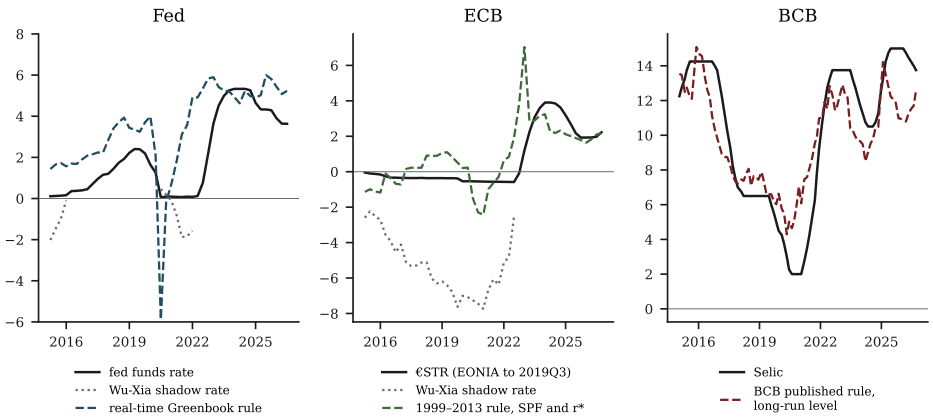


Figure 10.1. Policy rates, shadow rates and rule-implied notional rates, 2015 to 2026. Left: fed funds rate, Wu-Xia shadow rate at the lower bound, and the target level of the real-time Greenbook rule on SPF expected core CPI inflation around the median of the neutral-rate suite. Centre: euro area overnight rate, Wu-Xia shadow rate, and the target level of the ECB rule estimated on 1999 to 2013 with SPF expectations and the median of the euro area suite. Right: Selic and the long-run level of the BCB's published rule, at Copom meetings. Sources: FRED, Federal Reserve Bank of Atlanta, J. C. Wu, Federal Reserve Bank of Philadelphia, ECB Data Portal, New York Fed, BCB; author's estimates.

10.2.4 The SEP medians as a reduced-form reaction function

Since 2012 every FOMC participant has published, four times a year, projections of the fed funds rate, core PCE inflation and unemployment for the current year and the following two or three, together with longer-run values. The medians of the three projections are separate marginal medians: they need not belong to the same participant, and individual rate projections are not published alongside the same participant's macroeconomic projections. A regression of median on median is therefore a reduced-form summary of how the central tendency of projected rates moves with the central tendency of projected inflation and unemployment, not the reaction function of any one participant. From the real-time vintages in ALFRED, 43 SEP rounds from December 2015 to September 2026 give 129 projection points for the current year and the next two. Regressing the projected rate minus the longer-run rate on projected core PCE inflation minus 2 and projected

unemployment minus its longer-run level,

$$i_{s,h}^{SEP} - i_s^{LR} = -0.54 + 1.18\,(\pi_{s,h}^{SEP} - 2) - 0.25\,(u_{s,h}^{SEP} - u_s^{LR}), \qquad R^2 = 0.32, \tag{10.4}$$

with standard errors of 0.12, 0.16 and 0.14. The constant is negative: in the median projection, the rate sits half a point below its longer-run level when inflation and unemployment are at their longer-run values, which reflects the long stretch at the zero bound in the early part of the sample.

Result: the SEP reaction function changed in 2022

Split at 2022, the two periods look different. The comparison is descriptive: horizon effects, a common trend, the change of framework and different shocks all differ between the periods, and the samples are small. In the 72 projections from December 2015 to December 2021, the coefficient on projected inflation is −0.26: higher projected inflation came with lower projected rates, which is what a framework that promised to let inflation overshoot looks like in the data. In the 57 projections from March 2022 to September 2026, the coefficient is +0.68, and the constant turns from −1.26 to +0.53.

The September 2026 SEP is the latest application. The median projections were core PCE inflation of 3.4% for 2026, 2.5% for 2027 and 2.2% for 2028, unemployment of 4.1% against a longer-run 4.2%, and a fed funds rate of 4.1% at the end of 2026 and 2027, against a longer-run rate of 3.2%. The committee had just raised the target range to 3.75 to 4.00% on 16 September, its first increase since July 2023, in response to the oil shock described in chapter 9. The projection implies one further quarter-point increase, a rate held almost a point above neutral while inflation falls back, and it is consistent with the post-2022 coefficients.

10.3 The ECB

10.3.1 A rule written on survey expectations

The ECB's mandate is price stability, defined since 2021 as 2% inflation over the medium term, symmetric. It has no employment mandate, but it responds to economic activity because activity forecasts inflation. The natural estimation uses the ECB Survey of Professional Forecasters: inflation and unemployment expected for the next calendar year, and the longer-term

Table 10.3. Estimated ECB reaction functions, quarterly, nonlinear least squares. Inflation and unemployment: SPF expectations for the next calendar year; unemployment gap relative to the SPF longer-term expectation. The r^* version uses the median of the euro area neutral-rate suite of chapter 7. Standard errors in parentheses.

Specification and sample	ρ	ϕ_π	ϕ_u	R^2	n
Constant intercept, 1999 to 2013	0.77 (0.04)	0.96 (0.91)	−2.10 (0.35)	0.96	54
With r_t^*, 1999 to 2013	0.73 (0.05)	1.48 (0.75)	−1.33 (0.17)	0.96	54

expectation of unemployment as the benchmark for slack. The policy rate is measured by the overnight market rate (EONIA until September 2019, €STR from October 2019), quarterly averages, as a proxy for the realised stance; the administered rates move it, and with excess liquidity from 2008 it traded close to the deposit facility rate. Re-estimating the r^* version on the deposit facility rate itself gives a response to expected inflation of 1.73 (standard error 1.58) and to the unemployment gap of −1.69 (0.34), the same picture with less precision.

Result: the ECB's inflation response depends on the neutral rate it is measured against

Before the rate reached its floor, the ECB's long-run response to expected inflation was 0.96 with a constant intercept and 1.48 with the intercept written around the time-varying neutral rate (Table 10.3). The response to the expected unemployment gap is precise and large: −1.33 to −2.10 per point, equivalent with an Okun coefficient of two (a point of unemployment gap per two points of output gap, both quarterly averages) to 0.7 to 1.0 per point of output gap. Extending the sample beyond 2013 does not help: from 2014 to 2022 the overnight rate sat at its floor, and with the Wu-Xia shadow rate in its place, which fell to −7.7% at the end of 2020 while expected inflation barely moved, the inflation coefficient becomes implausibly large.

The comparison carries a lesson about the intercept. With a constant intercept the fall in the neutral rate after 2008 is read as a softer response to inflation: the rate came down while expected inflation did not, and a regression with a fixed r^* can explain that only with a smaller ϕ_π. Letting the neutral rate move, the estimate rises to 1.48 and the rule satisfies the Taylor principle, $\kappa(\phi_\pi - 1) + (1 - \beta)\phi_x = 0.172 \times 0.48 + 0.01 \times 0.17 > 0$.

The standard error, 0.75, is still large: expected inflation in the SPF varied between 1.1 and 2.6% from 1999 to 2013, and a survey expectation is itself an outcome of policy, so a credible bank sees little movement in it.

10.3.2 Late, then tight, then on the rule

Figure 10.1 (centre) compares the overnight rate with the notional rate of the r^* version estimated on 1999 to 2013. The rule called for a rate of 0.8% in the first quarter of 2022, 3.9% in the third and 7.0% in the fourth, when SPF expected inflation for 2023 jumped to 5.8% and the overnight rate was −0.58%, −0.08% and 1.2%: the ECB, like the Fed, was late by four to six points at the peak. The overnight rate understates the easing before 2022. With the Wu-Xia shadow rate, which reads the stance off the whole yield curve and so includes the asset purchases, the stance from 2019 to 2021 was between −6% and −7.7%, against rule levels between −2.4% and 0.5%: the balance sheet delivered an easing several points larger than the rule called for, and it was still running when inflation took off. It then overshot. In the third quarter of 2023 the rate was 3.6% against a rule level of 3.2%, and in early 2024, with expected inflation back near 2%, 3.9% against 2.3%. From mid-2025 the two coincide: 1.92% against 1.79% in the third quarter of 2025 and 1.93% against 1.82% in the first quarter of 2026. The increases of June and September 2026, to a deposit rate of 2.50%, put the rate slightly above the rule as SPF expected inflation for 2027 rose to 2.2%.

The sequence has a simple reading in terms of chapter 9. The ECB treated the 2021 to 2022 energy shock as a supply shock to look through, as theory recommends when expectations are anchored, until the risk that it would reach wages through catch-up became the dominant concern. It then held rates above the rule for a year while negotiated wage growth peaked at 5.4%. The rule estimated on the calm period could not have prescribed either phase, because the calm period contained neither a large energy shock nor a lagged wage response.

10.4 The BCB

10.4.1 Replicating the published rule

The BCB is unusual in publishing an estimated reaction function as part of its model suite. The June 2024 Inflation Report gives

$$i_t = 1.48\, i_{t-1} - 0.58\, i_{t-2} + 0.10\left[\bar{r}_t + \pi^*_t + 2.03\,(\mathbb{E}_t\pi_{t+h} - \pi^*_t)\right] + \varepsilon_t, \quad (10.5)$$

Table 10.4. The BCB reaction function: published coefficients (Banco Central do Brasil, 2024a), re-estimates on quarterly data from public sources, and a separate reaction function at Copom meeting frequency. Nonlinear least squares. The shift is an additive constant on the target level, zero in the published rule. Standard errors in parentheses. Meeting-frequency lags are per meeting and not comparable with quarterly lags.

	i_{t-1}	i_{t-2}	Shift	$\mathbb{E}\pi - \pi^*$	R^2	n
Published (Q), 2003–23	1.48	−0.58		2.03		
Quarterly, 2003–23	1.40 (0.08)	−0.53 (0.07)	0.12 (0.78)	2.94 (0.80)	0.977	81
Quarterly, 2005–23	1.56 (0.09)	−0.69 (0.08)	0.45 (0.67)	1.77 (0.71)	0.980	76
Quarterly, 2003–26	1.39 (0.07)	−0.52 (0.06)	0.20 (0.73)	3.10 (0.73)	0.977	92
Meetings, 2006–26	1.78 (0.05)	−0.82 (0.04)	0.63 (0.63)	1.92 (0.61)	0.996	165
Meetings, 2006–19	1.78 (0.06)	−0.81 (0.05)	0.52 (0.82)	1.36 (0.82)	0.994	111

estimated on quarterly data from 2003Q4 to 2023Q4, where $\bar{r}_t$ is the neutral real rate, π_t^* the target and $\mathbb{E}_t\pi_{t+h}$ the Focus expectation (chapter 3). The two lags give the rule momentum: the impact of last quarter's move carries into this quarter's, which reproduces the Copom's habit of moving in sequences of equal steps.

The replication uses only public data and the BCB's own frequency. The policy rate is the Selic target at the end of each quarter (SGS 432). Expected inflation is the quarterly average of the Focus median for the next twelve months. The target is the weighted average of the targets for the current and the following calendar year over the twelve-month horizon. The neutral rate is the Focus-implied real rate four years ahead, the Focus median for the Selic deflated by the Focus median for inflation, which is the survey measure in the BCB's own neutral-rate box (chapter 7).

A second exercise asks a different question: what reaction function describes the Copom decision by decision? It uses the 165 Copom meetings from January 2006 to September 2026, with the Focus median from the last survey before each meeting. Its lag coefficients are not comparable with the published ones, because a lag of one meeting, about six weeks, is not a lag of one quarter; its long-run response to the expectations gap is comparable, because a long-run response does not depend on the frequency at which it is measured.

Result: the published rule replicates at its own frequency

On quarterly data over the BCB's sample, the inertia coefficients are 1.40 and −0.53, within one standard error of the published 1.48 and −0.58 (Table 10.4). The response to the expectations gap is 2.94 (standard error 0.80) when the sample starts in 2003Q4 and 1.77 (0.71) when it starts in 2005, after the disinflation of 2003 and 2004; the published 2.03 lies between the two and within one standard error of each. The shift on the target level is not significantly different from zero, which confirms the published intercept $\bar{r}_t + \pi_t^*$. The published rule, applied with its own coefficients to the public quarterly data, tracks the Selic with a root mean squared error of 0.63 points over 2003Q4 to 2023Q4.

The meeting-level reaction function gives the same long-run picture. Its response to the expectations gap is 1.92 (0.61) over 2006 to 2026, close to the quarterly estimates, and it predicts each decision with a root mean squared error of 0.22 points, less than one quarter-point step. Estimated on 2006 to 2019 alone, the response is 1.34 (0.82). Because the 2020 to 2026 window is short and dominated by two large cycles, it is not estimated separately; the change between the two periods is better read as a descriptive comparison of regimes than as a structural break. Part of it is the inflation-targeting regime working as designed: after the de-anchoring of 2015, when Focus expectations for the next twelve months reached 7.1%, the Copom under new leadership from mid-2016 rebuilt credibility by responding more strongly. Part is the period itself: the 2021 to 2022 and 2024 to 2025 tightenings, the two largest moves in the sample, both came with expectations far above target, which identifies the coefficient well.

10.4.2 Where the Copom departed from its rule

The meeting-level rule's residuals show the episodes in which the Copom did something its own systematic behaviour did not predict. Thirteen decisions since 2006 differ from it by 0.4 points or more.

2008 to 2011. The Copom raised rates in April 2008 by more than the rule predicted as commodity prices rose, and then, after the Lehman failure, cut more slowly than the rule in January and March 2009 before catching up. The rule responds to expectations, and in early 2009 Focus expectations had not yet fallen with the collapse in activity. In August 2011 it cut by

50 basis points when the rule, and the Focus median, expected a hold: the largest surprise in the sample (chapter 14).

2019 and 2021. In July 2019 the Copom began an easing cycle with a 50 basis point cut, half a point more than the rule. In March and October 2021 it moved by more than the rule in the other direction, with a 75 basis point increase from 2% and a 150 basis point increase, as expectations began to drift.

2024. In December 2024 the Copom raised the Selic by 100 basis points, with guidance of two more of the same size, half a point more than the rule. In each of the tightening episodes expectations had begun to drift and the Copom chose to front-load.

The long-run level of the published rule, $\bar{r}_t + \pi_t^* + 2.03(\mathbb{E}_t\pi_{t+h} - \pi_t^*)$, which does not depend on the frequency of the lags, is plotted with the Selic in Figure 10.1 (right). It rose from 4.6% in August 2020 to 12.8% in June 2022 and to 14.2% in January 2025, when Focus expected inflation for the next twelve months reached 5.7%. The Selic then went above it: 15% against 12.2% in June 2025 and 11.0% in November 2025, as expectations fell back towards 4% while the Copom held. At the September 2026 meeting the long-run level was 12.6% (Focus 4.6%, target 3%, Focus-implied neutral rate 6.3%) and the Selic 13.75%. The meeting-level rule predicted 13.76% for that meeting, a cut of a quarter point; the Copom cut by a quarter point.

In practice: reading a Copom meeting with a reaction function

Three inputs give the meeting-level rule's prediction for the next meeting: the last two decisions, the latest Focus median for the relevant horizon and the Focus-implied neutral rate. The difference between the prediction and market pricing for the meeting measures how much the market expects the Copom to deviate from its own rule. When the two diverge by more than a quarter point, the statement and minutes usually explain why.

10.5 Estimated rules and the welfare loss

The estimated rules can be evaluated in the model of chapter 3, the core model with price indexation of $\gamma = 0.5$ and natural-rate and cost-push shocks of persistence 0.5, by solving the model under each rule and computing the welfare loss (3.20), which uses the quasi-difference of inflation appropriate to the indexed model. The coefficients are converted to the

Table 10.5. Estimated reaction functions evaluated in the model with indexation ($\gamma = 0.5$). Standard deviations of annualised inflation, of its quasi-difference $q_t = \pi_t - \gamma\pi_{t-1}$ (annualised), which is what the loss penalises, and of the output gap; welfare loss (3.20) relative to the Taylor (1993) rule. Simulated, 100,000 quarters.

Rule	sd π	sd q	sd x	Rel. loss
Taylor (1993)	2.32	1.68	1.05	1.00
Fed, real-time 1987–2008, current π	1.97	1.64	0.92	0.94
Fed, real-time 1987–2008, expected π, 4 qtrs	5.84	4.29	1.60	6.09
ECB, 1999–2013, current π	1.97	1.60	0.96	0.90
ECB, 1999–2013, expected π, 4 qtrs	6.76	4.85	1.87	7.82
BCB, quarterly 2003–2026, current π	1.06	0.98	1.43	0.50
BCB, quarterly 2003–2026, expected π, 4 qtrs	6.11	4.45	1.80	6.61
Best simple rule in the chapter 3 search	0.65	0.63	1.75	0.42

model's quarterly units; unemployment gaps become output gaps with an Okun coefficient of two. Rules estimated on expected inflation are evaluated twice, once on model-consistent inflation four quarters ahead and once on current inflation, because the model's rational forecast moves much less than a survey forecast does.

Result: the estimated rules lose more than the best simple rules

None of the estimated rules reaches the best simple rules' loss (Table 10.5, Figure 10.2). The Fed's real-time rule, with smoothing of 0.85 a quarter, a response of 1.57 to inflation and 0.86 to the gap, reaches 94% of the Taylor loss when applied to current inflation, mostly by stabilising the gap; the ECB rule, estimated around the moving neutral rate, reaches 90%. The BCB's rule, estimated at quarterly frequency and applied to current inflation, does best of the estimated rules, at 50%, because its momentum supplies the history dependence of the commitment solution. The best simple rules of chapter 3 reach 42% of the Taylor loss, by responding to inflation four or more times as strongly as any of the banks does.

The rows on expected inflation repeat the warning of chapter 3. A rule that responds about one for one, or twice, to a model-consistent forecast four quarters ahead stabilises very little, because that forecast moves by a fraction of current inflation once shocks fade. The ECB and BCB rules are written on survey expectations, which in the data move far more than a

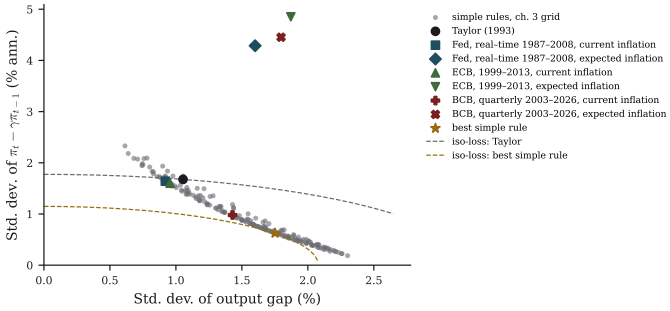


Figure 10.2. Estimated reaction functions and the simple rules of the chapter 3 search, in the space of output-gap volatility and the volatility of $\pi_t - \gamma\pi_{t-1}$, the inflation term the indexed loss penalises; model with $\gamma = 0.5$. Dashed curves are iso-loss contours at the model's weight on the gap, through the Taylor rule and the best simple rule: a point below a contour has lower loss. Ranking by distance from the origin is valid only when one point is lower on both axes.

rational forecast would: Focus twelve-month expectations moved from 2.9% to 5.9% between 2020 and 2022, and the BCB's coefficient of two applied to that movement produced the 11.75 point tightening of 2021 to 2022. The welfare cost of a rule on expectations depends on how much the expectations move, which is an empirical property of the survey, not of the rule.

10.6 What an estimated rule leaves out

Four things make the residual of an estimated reaction function larger than its standard error suggests.

Regime changes. The Fed's SEP coefficients changed sign in 2022, and the BCB's long-run response to expectations rises from 1.36 on 2006 to 2019 meetings to 1.92 when 2020 to 2026 is added; the samples are nested, so the difference understates the later period's own coefficient. A rule estimated across a regime change averages two policies that never coexisted.

The lower bound and the balance sheet. From 2009 to 2015 and in 2020 to 2021 the Fed's rate was at zero, and the ECB's overnight rate was at or below zero from 2014 to 2022. Policy in those years ran through asset purchases and forward guidance (chapters 6 and 14), which a rule for the short rate does not see.

Risk management. The Fed's three cuts of 2019, with unemployment at 3.6% and core inflation near 1.7%, were described as insurance against downside risks from trade policy and global growth. A rule that responds

to point forecasts cannot represent a response to the distribution around them.

Uncertainty about the model. Brainard's (1967) result, that uncertainty about the effect of policy calls for smaller moves, shows up in estimated rules as smoothing. The estimates of ρ between 0.73 and 0.93 (per quarter or per meeting, as estimated), and the BCB's lag sum of 0.96 at meeting frequency (0.87 on quarterly data), are larger than any loss function without uncertainty would justify. Part of what an estimated rule calls inertia is caution.

The models central banks use to produce the forecasts these rules respond to are the subject of the next chapter.

Chapter 11

The Models Inside Central Banks

Every forecast a central bank publishes, every scenario in a press conference and every staff memo on the cost of waiting comes out of a model. The three-equation model of chapter 1 is the skeleton of those models, but none of the three banks in this book uses it directly. The Federal Reserve Board's main model, FRB/US, has 367 variables. The ECB's New Area-Wide Model II (NAWM II) is an estimated dynamic stochastic general equilibrium (DSGE) model with a banking sector. The BCB runs a family of small semi-structural models next to a DSGE model called SAMBA. They are built differently, and they answer the same question differently: how much does a percentage point on the policy rate lower inflation, how fast, and at what cost in output?

This chapter opens the models. It runs FRB/US itself, using the Fed's published Python implementation and the July 2026 database, reads the ECB's 2025 harmonised comparison of the models used across the euro area's central banks, and decodes the published results for SAMBA and the BCB's small models. The answer to the question above differs across them by a factor of more than five, and the difference is mostly in one parameter: the slope of the Phillips curve.

11.1 Two families of models

11.1.1 What a central bank model is for

Central bank models have three jobs. They produce the ***baseline*** forecast, conditional on assumptions for the policy rate, fiscal policy, oil and the exchange rate. They produce ***scenarios***: what happens to the forecast if oil rises 30%, if the exchange rate depreciates 10%, if the policy rate follows market pricing rather than the staff assumption. And they evaluate ***policy***: which path of the rate brings inflation back to target at least cost under the

bank's loss function. The first job rewards fit to the data. The third rewards structure, because a policy evaluation asks what happens under a rule that has not been followed before, and a model fitted to past correlations cannot answer that without a theory of why the correlations hold.

11.1.2 Semi-structural and DSGE

The two families resolve the trade-off differently.

A ***semi-structural*** model, of which FRB/US is the largest example, writes each behavioural equation in two parts: a long-run target derived from theory (desired consumption from permanent income, desired capital from the cost of capital, the desired price from a markup over cost) and a short-run adjustment equation estimated freely on the data. The long-run targets give the model an economic interpretation; the free dynamics let it fit the data equation by equation. The BCB's small models and the ECB's ECB-BASE are of this type.

A ***DSGE*** model derives every equation from the optimisation problem of a household, a firm or a bank, so that all equations share a single set of deep parameters, and estimates them jointly, usually by Bayesian methods. Smets and Wouters (2007) showed that a model of this kind, with enough frictions (habit in consumption, investment adjustment costs, sticky wages and prices with indexation), forecasts about as well as a vector autoregression. NAWM II and SAMBA belong to this family.

> **In practice: how the families are used**
>
> The ECB's staff projections are produced with a suite in which semi-structural models carry most of the forecast and DSGE models interpret it. The Fed's staff forecast is judgmental, with FRB/US used for scenarios and policy alternatives. The BCB's published forecasts in its Monetary Policy Report come from the small semi-structural models, with SAMBA and other models used for cross-checking. In all three, the model the public sees is not the only model the staff uses.

11.2 FRB/US

11.2.1 Structure

FRB/US went into use at the Board in 1996, and the Board has published its equations, data and code since 2014 (Brayton et al., 2014). The cur-

rent model file has 367 variables: 120 behavioural equations, 164 identities and 83 exogenous variables. It covers household spending in four categories, business investment in three, housing, inventories, trade, the labour market, wages and prices, the federal and state government sectors, and a financial block of interest rates, equity prices and the dollar.

Two design choices make it distinctive. First, almost every behavioural equation is a ***polynomial adjustment cost*** (PAC) equation. Second, expectations can be formed in two ways, which the user chooses.

11.2.2 The PAC equation

The derivation is short. A firm or household wants a variable y to equal a target y^*, which theory supplies, but changing y is costly. Let it minimise

$$\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \Big[(y_{t+j} - y^*_{t+j})^2 + b_1 (\Delta y_{t+j})^2 + b_2 (\Delta^2 y_{t+j})^2 \Big], \tag{11.1}$$

where the second term penalises changing y and the third penalises changing the rate of change. The first-order condition is a fourth-order difference equation in y. Factoring it, with two stable roots, gives the decision rule

$$\Delta y_t = a_0 \, (y^*_{t-1} - y_{t-1}) + a_1 \Delta y_{t-1} + \mathbb{E}_t \sum_{j=0}^{\infty} d_j \, \Delta y^*_{t+j}, \tag{11.2}$$

in which the coefficients a_0, a_1 and d_j are all functions of b_1, b_2 and β. The equation has an error-correction term, which pulls y toward its target; a momentum term, which carries last period's change forward; and a forward-looking term, a weighted sum of expected future changes in the target. That last term is where expectations of policy enter: a household that expects higher rates for longer expects a lower path of desired consumption and starts adjusting now.

The one-cost case shows the mechanics in closed form. With $b_2 = 0$ the first-order condition for y_t is

$$(y_t - y^*_t) + b_1 (y_t - y_{t-1}) - \beta b_1 \, \mathbb{E}_t (y_{t+1} - y_t) = 0, \tag{11.3}$$

or $\beta b_1 \mathbb{E}_t y_{t+1} - (1 + b_1 + \beta b_1) \, y_t + b_1 y_{t-1} = -y^*_t$. Its characteristic equation $\beta b_1 \lambda^2 - (1 + b_1 + \beta b_1)\lambda + b_1 = 0$ has one root inside the unit circle, λ, and one outside, $1/(\beta\lambda)$. Solving the unstable root forward gives

$$y_t = \lambda \, y_{t-1} + (1 - \lambda)(1 - \beta\lambda) \sum_{j=0}^{\infty} (\beta\lambda)^j \, \mathbb{E}_t y^*_{t+j}. \tag{11.4}$$

The agent closes a fraction $1-\lambda$ of the gap each period, toward a discounted average of future targets rather than today's. With $b_1 = 4$ and $\beta = 0.98$, $\lambda = 0.61$: 39% of the gap is closed in a quarter, and the target expected four quarters ahead carries a weight of $(0.98 \times 0.61)^4 = 0.12$ relative to today's. The second adjustment cost, b_2, adds the momentum term of (11.2), which is what lets FRB/US reproduce the hump-shaped responses of spending seen in the data.

11.2.3 Prices and wages

The core inflation equation is a hybrid New Keynesian Phillips curve in the form derived by Cogley and Sbordone (2008), with prices not reset optimally indexed to a mix of last quarter's inflation and trend inflation. In the model file it reads

$$\pi_t = \frac{\gamma\,\pi_{t-1} + \beta\,\mathbb{E}_{t-1}\pi_{t+1} + (1-\beta)(1-\gamma)\,\pi^*_{t-1} + \kappa_p\,\mu_{t-1}}{1+\gamma\beta}, \tag{11.5}$$

with $\gamma = 0.67$, $\beta = 0.98$, π^* the ten-year inflation expectation from surveys, and μ the price markup gap, which enters with $\kappa_p = 0.0019 \times 400$ per unit of log markup. The wage equation has the same form, with a four-quarter average of past wage inflation, trend wage growth equal to trend inflation plus trend productivity, and the unemployment gap with a coefficient of -0.0167.

> **Result: FRB/US has a very flat Phillips curve**
>
> In FRB/US, a one-point unemployment gap held for a year lowers annualised wage inflation by about 0.06 points in that year, and the markup gap moves price inflation by less than a hundredth of a point per percentage point. Unemployment does not enter the price equation at all; it reaches prices only through wages and expected inflation. The estimated slopes are an order of magnitude smaller than the calibrated $\kappa = 0.17$ of chapter 1, and close to the pre-2020 reduced-form estimate of chapter 9.

The model's inflation is anchored mainly by the trend term π^*, survey expectations ten years ahead, which in the July 2026 database stand near 2.2%. Monetary policy moves inflation in FRB/US chiefly by moving the output gap a great deal for a long time.

Table 11.1. FRB/US: peak responses to a funds rate 100 basis points higher for four quarters (quarter of peak in brackets).

	VAR expectations	Model-consistent
10-year Treasury yield (pp)	+0.40 [3]	+0.22 [1]
Real dollar (%)	+1.75 [3]	+0.80 [1]
Output gap (%)	−0.74 [8]	−0.29 [6]
Unemployment rate (pp)	+0.41 [9]	+0.17 [6]
Core PCE inflation, four-quarter (pp)	−0.05 [11]	−0.02 [7]

11.2.4 Expectations: VAR and model-consistent

The expectation terms in (11.2) and in the price and wage equations have to come from somewhere. FRB/US offers two options. Under ***VAR expectations***, agents forecast with a small vector autoregression in inflation, the output gap and the funds rate, estimated on history, plus agents' knowledge of long-run anchors. They know how the economy has usually behaved, not the current policy. Under ***model-consistent expectations*** (MCE), expectations are the model's own forecasts: agents know the model and the policy being followed. The Board's staff typically use MCE in financial markets, where asset prices react to news about policy, and VAR expectations elsewhere. PyFRB/US allows MCE in asset pricing alone, in wage and price setting alone, or everywhere.

11.2.5 Running the model

All the simulations below use PyFRB/US with the July 2026 database, whose baseline follows the median SEP projections where they are available and a model-guided extrapolation beyond them. The simulation starts in the first quarter of 2027. The experiment holds the funds rate 100 basis points above its baseline path for four quarters and then returns policy to the model's inertial Taylor rule. This is the experiment the BCB reports for its own models, which allows a direct comparison later in the chapter.

Result: a year of tighter policy in FRB/US

With VAR expectations, a year of rates 100 basis points higher raises unemployment by 0.41 points after two years and lowers the output gap by 0.74%, but reduces core PCE inflation by only 0.05 points at the peak, almost three years out (Table 11.1, Figure 11.1). The

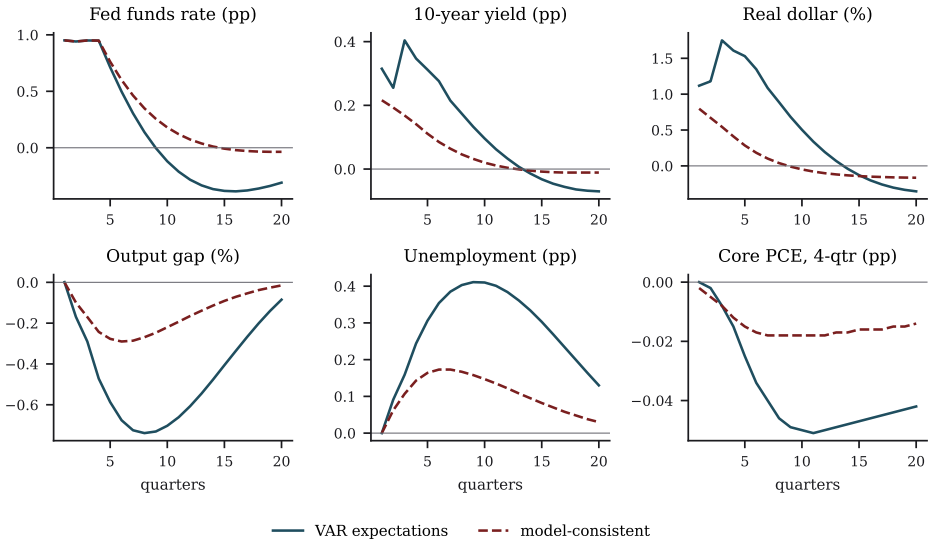


Figure 11.1. FRB/US responses to a fed funds rate 100 basis points above baseline for four quarters, then the inertial Taylor rule, from 2027Q1. Deviations from baseline. Blue: VAR expectations. Red dashed: model-consistent expectations in asset pricing and wage and price setting. Source: PyFRB/US, July 2026 database; author's simulations.

ratio of output lost to inflation reduced is about fifteen to one at the peak. With model-consistent expectations, the effects are less than half as large, because markets see from the start that the tightening is temporary: the 10-year yield rises by 22 basis points on impact rather than climbing to 40, and the dollar appreciates less.

The MCE result runs against a common intuition, that rational expectations make policy more powerful. They do when policy is persistent or announced in advance, because the forward-looking terms in (11.2) then bring the effect forward. A tightening known to last only four quarters is, under MCE, a small change in the expected path of rates beyond a year, and long-term rates, which drive spending on housing, durables and investment, move little. Under VAR expectations agents read the higher rate through their historical forecasting equation, which predicts that rates stay high for longer than they do. This is the forward guidance puzzle of chapter 14 seen from the other side: the power of a rate move in a model depends on what agents are assumed to know about its duration.

11.2.6 An oil shock and the choice of rule

The second experiment raises the relative price of imported oil by 50% in the first quarter of 2027, somewhat smaller than the rise in WTI from 58 dollars in December 2025 to 97 in September 2026, and lets the model's own oil equation govern its persistence. It is run under four of the policy rules FRB/US carries: the inertial Taylor rule, the Taylor (1993) rule, a Taylor rule on the unemployment gap, and an estimated rule.

Result: FRB/US looks through oil shocks under any rule

The oil price rises 47% in the second quarter and falls back to baseline within two years. Headline PCE inflation rises by 0.47 points in the second quarter and then falls 0.39 below baseline in the sixth as oil reverses. Core PCE inflation does not move by more than 0.01 points under any of the four rules, and the funds rate moves by at most 12 basis points, slightly down, as the oil shock lowers output. In FRB/US the choice of rule does not matter for an oil shock, because nothing in the model transmits oil to wages or core prices.

This is the property that made FRB/US, and models like it, recommend looking through the energy shock of 2021 and 2022. Chapter 9 found the opposite in the data: in the re-estimated Bernanke-Blanchard model, energy contributed 2.7 points to US CPI inflation at the peak and its effect persisted through the price equation. The difference is the transmission channel. FRB/US has trend inflation fixed by long-run survey expectations, which barely moved, and no catch-up of wages to past price inflation; the data in 2022 showed short-run expectations, wages and core prices all responding to energy.

11.2.7 Rules inside FRB/US: a core inflation shock

The oil experiment cannot distinguish rules because it does not reach core inflation. A shock that does is a direct shock to the core inflation equation: core PCE inflation one point higher, annualised, for four quarters before any feedback, which in FRB/US becomes a rise of 1.7 points in four-quarter core inflation through the indexation term. Table 11.2 reports the outcome under the four rules, with discounted sums of squared deviations of quarterly core inflation and the output gap over ten years as the two parts of the loss.

Table 11.2. FRB/US, VAR expectations: core inflation shock under four policy rules. Four-quarter core PCE inflation and unemployment, deviations from baseline; loss components are discounted sums over 40 quarters ($\beta = 0.99$) of squared deviations of annualised quarterly core inflation and of the output gap. Source: PyFRB/US, July 2026 database; author's simulations.

	Core PCE, 4-qtr		Unemp.	Peak	Loss components	
Rule	qtr 4	qtr 8	peak	funds	infl.	gap
Inertial Taylor	1.71	0.76	0.36	1.08	15.3	6.4
Taylor (1993)	1.70	0.72	0.63	1.92	14.9	8.9
Taylor, unemp. gap	1.70	0.71	0.74	2.18	14.9	13.8
Estimated rule	1.71	0.75	0.53	1.58	15.3	8.4

Result: in FRB/US, tighter rules buy almost no inflation

Raising the funds rate twice as much, under the Taylor rule on the unemployment gap rather than the inertial rule, doubles the rise in unemployment, from 0.36 to 0.74 points, and lowers four-quarter core inflation after two years by 0.05 points, from 0.76 to 0.71. The inflation part of the loss falls by 2.5% and the output part more than doubles. Inflation returns to baseline within three years under every rule, because the shock and the indexation term die out on their own.

With a flat Phillips curve, the optimal response to a transitory cost-push shock in the model of chapter 2 is close to zero: $x_t = -(\kappa/\lambda_x)\pi_t$, so for a given weight λ_x a small κ means a small gap. If λ_x is itself derived from the micro foundations, $\lambda_x = \kappa/\varepsilon$ and the ratio is ε whatever the slope; the case for looking through then rests on the shock being transitory rather than on the slope. FRB/US reaches the same conclusion through simulation. The recommendation is only as good as the assumption that the shock is transitory and expectations stay anchored, which is what the episode of 2021 to 2022 tested.

11.2.8 A term premium shock

The model also gives a quantitative answer to a question that returns in chapters 12 and 13: what a rise in term premia does to the economy. A 50 basis point rise in the 10-year term premium, decaying through its own equation to 30 basis points after two years and 10 after five, lowers the output gap by 0.20% and raises unemployment by 0.10 points at the peak,

and the inertial rule cuts the funds rate by 16 basis points in response. In FRB/US a rise of 50 basis points in the term premium is worth about a fifth of the output effect of a year of funds rate 100 basis points higher.

11.3 NAWM II and the euro area models

11.3.1 Structure

NAWM II (Coenen et al., 2018) is the ECB's main DSGE model for the euro area. It extends the original NAWM, an open-economy model with the frictions of Smets and Wouters (2007), in three directions: a banking sector in which banks finance firms and face a leverage constraint, in the spirit of Gertler and Karadi (2011); long-term government bonds that households hold imperfectly, so that bond purchases can lower term premia, which gives the model a channel for the asset purchases of chapter 6; and a gradual pass-through of the policy rate to bank lending rates. It is estimated by Bayesian methods, and was re-estimated in 2022 with an adjustment to its steady-state real interest rate.

The model's price Phillips curve, written on real marginal cost rather than the output gap, has an estimated slope of 0.0072, and its wage Phillips curve 0.0065. The comparable slope in this book's calibration, on marginal cost rather than on the gap, is $\kappa/(\sigma + (\varphi + \alpha)/(1 - \alpha)) = 0.172/8 = 0.022$, three times NAWM II's. Like FRB/US, the euro area's main DSGE model has a much flatter Phillips curve than the textbook calibration.

11.3.2 How a DSGE model is estimated

NAWM II, SAMBA and most central bank DSGE models are estimated by the same Bayesian procedure, which has four steps.

Step 1: solve. For given parameters θ, the model's equilibrium conditions are log-linearised around the steady state and solved, as in chapter 1, for a state-space form

$$s_t = T(\theta)\, s_{t-1} + R(\theta)\, \varepsilon_t, \qquad y_t = Z(\theta)\, s_t + \eta_t, \tag{11.6}$$

where s_t is the vector of model states, ε_t the structural shocks, y_t the observed data (GDP growth, inflation, the policy rate, wages and so on) and η_t measurement error.

Step 2: likelihood. The Kalman filter computes, recursively, the forecast of y_t given data up to $t-1$ and its variance, and from them the likelihood $p(y_{1:T} \mid \theta)$, the probability of the observed data under the parameters.

Step 3: prior. A prior $p(\theta)$ summarises what is known before seeing the data: that the discount factor is close to 0.99, that prices change on average every few quarters, that the response to inflation in the rule is above one.

Step 4: posterior. Bayes's rule gives the posterior

$$p(\theta \mid y_{1:T}) \propto p(y_{1:T} \mid \theta)\, p(\theta), \tag{11.7}$$

which is explored by a Markov chain Monte Carlo algorithm, usually random-walk Metropolis-Hastings: propose a new θ near the current one, accept it with probability equal to the ratio of posteriors (capped at one), and repeat several hundred thousand times. The accepted draws describe the posterior distribution of each parameter.

The prior matters. With thirty years of quarterly data and forty or more parameters, some parameters, such as the slope of the Phillips curve, are weakly identified by the likelihood, and their posterior stays close to the prior. A slope estimated at 0.0072 is partly a statement about the data and partly a statement about what the modellers believed. SAMBA's estimates of the rule, reported below, show the reverse case: a prior on the response to expected inflation centred on 2 with a standard deviation of 0.35, and a posterior mean of 2.43 that the data pulled away from it.

11.3.3 The 2025 harmonised comparison

In 2025 the ECB published the results of a harmonised exercise across the macroeconomic models used by the central banks of the European System of Central Banks, DSGE and semi-structural, together with a meta-analysis of several hundred empirical studies (Bobasu et al., 2025). Each model was hit with the same unexpected 25 basis point increase in the policy rate, under its own rule and under a common rule, $R_t = 0.87 R_{t-1} + 0.13[\bar{R} + 1.7(\pi_t - \bar{\pi})] + 0.1\,\Delta Y_t$.

Result: euro area models and evidence on a 25 basis point shock

In the meta-analysis of empirical studies, the median response of output to an unexpected 25 basis point increase peaks at -0.15% after two years and the median response of the price level at -0.11% after four years; the means are -0.25% and -0.18%. Correcting for publication bias roughly halves these numbers, to medians of -0.08% and -0.05%. The models produce responses of broadly the same size. DSGE models front-load the effect and produce larger inflation responses; semi-structural models, such as ECB-BASE, produce slower and more per-

sistent responses, closer to the empirical benchmarks. Harmonising the policy rule narrows the range across models, which shows that much of the disagreement among models is about the assumed reaction function rather than about the economy.

Scaled linearly to 100 basis points, the median empirical estimate is a fall of 0.6% in output after two years and 0.44% in the price level after four. If the price-level effect accumulates evenly, an assumption the profile does not guarantee, that is roughly 0.1 to 0.15 points a year off inflation. That is two to three times FRB/US's response with VAR expectations and five or more times its response with MCE.

The same paper reports three findings about the 2022 to 2023 tightening that bear on chapters 9 and 10. The inflation response to policy was stronger than in earlier cycles, with a broadly unchanged activity response, so the sacrifice ratio was lower, consistent with the state-dependent slope of chapter 9; transmission through bank credit was stronger than in earlier cycles; and policy reacted more strongly to supply and energy shocks than in 2005 to 2008, that is, the ECB looked through the energy shock less than its models would have recommended.

11.4 SAMBA and the BCB's small models

11.4.1 SAMBA

SAMBA (de Castro et al., 2015) is the BCB's DSGE model, built by the Research Department and first published in 2011. It is a small open economy model with the standard frictions plus four features chosen for Brazil: imported goods enter as inputs into domestic production rather than as final goods, so the exchange rate reaches prices through costs; a large share of consumer prices is administered, set by contract or regulation and indexed to past inflation; a fraction of households cannot borrow, which breaks Ricardian equivalence; and fiscal policy follows a rule for the primary surplus, as Brazil's fiscal framework required from 1999. It was estimated on data from 1999, when inflation targeting began.

The estimated policy rule in SAMBA responds to expected inflation with a long-run coefficient of 2.43 (90% interval 1.98 to 2.89), with smoothing of 0.79 and a small response to the output gap. The long-run inflation response is close to the 2.03 the BCB publishes for its own rule and to the 1.8 to 2.9 re-estimated on quarterly data in chapter 10.

Result: a policy shock in SAMBA

An unexpected increase of 100 basis points in the annualised Selic, 25 in quarterly terms, under the estimated rule, lowers four-quarter IPCA inflation by about 0.25 points at the trough, in the fourth quarter. The response of freely set prices is larger and earlier than that of administered prices, whose trough comes in the sixth quarter, because administered prices are indexed to past inflation. Because the rule responds to the fall in inflation, the Selic is back near baseline by the fourth quarter; holding it fixed for a year would produce a larger fall in inflation. The exchange rate appreciates on impact, and both appreciation and lower activity lower marginal cost.

11.4.2 The small semi-structural models

The models behind the BCB's published forecasts are the small semi-structural models described in chapter 3: an aggregate model and a disaggregated one that separates free and administered prices, each with an IS curve, a Phillips curve with Focus expectations, an uncovered interest parity condition and the BCB's published rule. Their benchmark experiment is the one run in FRB/US above: the Selic one point higher for four quarters, then the rule. In the aggregate model, four-quarter IPCA inflation falls by 0.27 points at the peak, after four quarters; in the disaggregated model, 0.24. The replication from public data in chapter 3 gives 0.19 and 0.20.

11.5 The same shock in seven models

Table 11.3 puts the results side by side. Only rows with the same experiment compare exactly. On the matched experiment, a year of +100 basis points with four-quarter inflation, the BCB small model's response is four to five times FRB/US's under VAR expectations (-0.27 and -0.19 against -0.05), and it peaks within a year rather than three. SAMBA and the ESCB median use one-off shocks and, for the ESCB, the price level, so they place the euro area evidence in between only as an order of magnitude.

Three differences account for most of the spread.

The exchange rate. In Brazil a higher Selic appreciates the real, and the real's pass-through to prices is large: the BCB's aggregate model puts the effect of a 10% depreciation at 0.96 points of IPCA. In FRB/US the dollar

Table 11.3. Effect of a 100 basis point tightening on inflation across models. Peak response of inflation, percentage points; ESCB numbers are scaled linearly from a 25 basis point shock and refer to the price level. Experiments differ in shock duration and inflation measure, so rows are comparable as orders of magnitude, and exactly only where the experiment column matches. Sources: chapter 1; PyFRB/US simulations; Bobasu et al. (2025); de Castro et al. (2015); Banco Central do Brasil (2024a); chapter 3 replication.

Model	Experiment	Inflation measure	Peak effect
Core NK model, ch. 1	one-off, $\rho_v = 0.5$	qtr, annualised	−0.35, impact
FRB/US, VAR expectations	+100 bp for 4 qtrs	core PCE, 4-qtr	−0.05, qtr 11
FRB/US, MCE	+100 bp for 4 qtrs	core PCE, 4-qtr	−0.02, qtr 7
ESCB evidence, median	one-off	price level	−0.44, year 4
SAMBA (BCB DSGE)	one-off	IPCA, 4-qtr	−0.25, qtr 4
BCB small, aggregate	+100 bp for 4 qtrs	IPCA, 4-qtr	−0.27, qtr 4
BCB small, replication	+100 bp for 4 qtrs	IPCA, 4-qtr	−0.19, qtr 4

appreciates 1.75% under VAR expectations, and import prices are a small share of US consumption. A large part of the BCB's policy transmission runs through the exchange rate, and that channel is fast.

The Phillips curve. FRB/US and NAWM II have estimated slopes far below the textbook calibration, which forces the transmission through a long and large fall in output. The BCB's Phillips curves, estimated on a sample with more inflation variability, are steeper.

Expectations. The more policy shifts expected inflation directly, the less it needs to move output. In the BCB's models, Focus expectations respond to policy; in FRB/US, long-run expectations are pinned by surveys that barely move.

11.6 Model uncertainty

A policymaker who knew which model was right would choose policy optimally in it. A policymaker who does not, which is every policymaker, needs a rule that performs acceptably in all the plausible models. The Macroeconomic Model Data Base of Wieland et al. (2012), which puts more than a hundred published models, including versions of central bank models, on a common footing, was built for that purpose. Its consistent finding is that the rules that are optimal in one model can perform very badly in another, and that simple rules with interest rate smoothing and a strong response to inflation perform best across models.

The comparison in this chapter adds one result specific to 2021 to 2026. The main central bank models entered the surge with flat Phillips curves and long-run expectations that moved little: pinned to surveys in FRB/US, tied to the target under the rule in NAWM II and SAMBA, partly adaptive in the BCB small model. That made them recommend looking through the energy shock, and it made them predict a costly disinflation, which held where the economy started near the threshold of chapter 9 and was too pessimistic for the US and euro area headline, where the reversal of supply shocks did much of the work. What they could not represent was the steep part of the curve, because their parameters were estimated on two decades in which the economy had not been there. The models were right about the region they had been estimated on and silent about the one the economy entered. The ECB's finding of a lower sacrifice ratio in 2022 to 2023, and the threshold estimates of chapter 9, are the same fact seen in the evidence rather than in the models.

The next part of the book turns from the models of the economy to the markets that transmit policy to it, starting with the yield curve.

Part III

Transmission

Chapter 12

The Yield Curve

A central bank sets an overnight rate. Households and firms borrow at rates of two, five, ten and thirty years. Everything in the transmission of policy that runs through financial markets runs through the link between the two, and that link is the yield curve. It has two parts. One is the market's expectation of the path of the policy rate over the life of the bond; the other is the ***term premium***, the extra return investors demand for holding duration rather than rolling over short-term debt. A central bank moves the first directly, through its decisions and its guidance. It moves the second only indirectly, through its balance sheet and through the uncertainty its own policy creates.

This chapter derives the decomposition, estimates it with the models central banks themselves publish, the New York Fed's ACM model and the Board's Kim-Wright model, and measures how each US policy cycle since 1994 passed through to the ten-year yield. The answer varies from cycle to cycle by a factor of six, and in the easing of 2024 and 2025 it had the wrong sign: the Fed cut by 1.44 points and the ten-year yield rose by half a point. The chapter then sets the Treasury curve beside the euro area curve, using the ECB's own estimated AAA curve, and the Brazilian curve, using Tesouro Nacional prices, and ends with what the curve says about policy in 2026.

12.1 Expectations and the term premium

12.1.1 Prices, yields and forward rates

Let $P_t^{(n)}$ be the price at t of a zero-coupon bond paying one dollar in n periods, and $p_t^{(n)} = \ln P_t^{(n)}$. The continuously compounded yield is $y_t^{(n)} = -p_t^{(n)}/n$, and the one-period rate is $i_t = y_t^{(1)}$. The forward rate for borrowing

between $t+n-1$ and $t+n$, contracted today, is

$$f_t^{(n)} = p_t^{(n-1)} - p_t^{(n)} = n\, y_t^{(n)} - (n-1)\, y_t^{(n-1)}. \tag{12.1}$$

A yield is the average of the forward rates up to its maturity: $y_t^{(n)} = \frac{1}{n}\sum_{k=1}^{n} f_t^{(k)}$.

12.1.2 The decomposition

The ***expectations hypothesis*** says that each forward rate equals the expected short rate for that period. Then

$$y_t^{(n)} = \frac{1}{n}\sum_{k=0}^{n-1} \mathbb{E}_t\, i_{t+k}. \tag{12.2}$$

The hypothesis fails in the data: forward rates predict future short rates with a bias, and the bias varies over time (Campbell and Shiller, 1991). The term premium is defined as the residual,

$$y_t^{(n)} = \underbrace{\frac{1}{n}\sum_{k=0}^{n-1} \mathbb{E}_t\, i_{t+k}}_{\text{expectations component}} + \underbrace{tp_t^{(n)}}_{\text{term premium}}, \tag{12.3}$$

and the expectations component is called the ***risk-neutral yield***: the yield that would prevail if investors were indifferent to risk. Equation (12.3) is an identity. Everything depends on how $\mathbb{E}_t i_{t+k}$ is measured, and that requires a model.

12.1.3 An affine model in four steps

The simplest model that delivers a time-varying term premium is a discrete-time version of Vašíček (1977).

Step 1: the short rate. Let the short rate follow an AR(1),

$$i_{t+1} = \mu + \phi\,(i_t - \mu) + \sigma\,\varepsilon_{t+1}, \qquad \varepsilon \sim N(0,1). \tag{12.4}$$

Step 2: the stochastic discount factor. Bond prices satisfy the pricing equation $P_t^{(n)} = \mathbb{E}_t[M_{t+1} P_{t+1}^{(n-1)}]$, with the log discount factor

$$m_{t+1} = -i_t - \tfrac{1}{2}\lambda^2 + \lambda\,\varepsilon_{t+1}, \tag{12.5}$$

where λ is the market price of interest-rate risk. The sign convention makes marginal utility high when the short rate rises, which is when bond prices fall, so $\lambda > 0$ means bonds pay off badly in bad states and carry a positive premium. With $\lambda = 0$ investors are risk-neutral.

Step 3: guess and verify. Guess $p_t^{(n)} = -A_n - B_n i_t$. Substituting into the pricing equation, using the log-normal formula $\ln \mathbb{E}e^X = \mathbb{E}X + \frac{1}{2}\mathrm{Var}\, X$, and matching the constant and the coefficient on i_t gives the recursions

$$B_n = 1 + \phi B_{n-1}, \qquad A_n = A_{n-1} + B_{n-1}(1-\phi)\mu + B_{n-1}\lambda\sigma - \tfrac{1}{2}B_{n-1}^2\sigma^2, \tag{12.6}$$

with $A_0 = B_0 = 0$, so $B_n = (1-\phi^n)/(1-\phi)$.

Step 4: the term premium. The yield is $y_t^{(n)} = (A_n + B_n i_t)/n$. Setting $\lambda = 0$ gives the risk-neutral yield; the difference is

$$tp^{(n)} = \frac{\sigma\lambda}{n}\sum_{k=1}^{n-1} B_k. \tag{12.7}$$

The premium rises with the price of risk λ, with the volatility σ, and with the persistence ϕ, which makes long bonds more exposed to short-rate shocks. In this one-factor model the premium is constant; making λ a linear function of the state, $\lambda_t = \lambda_0 + \lambda_1 i_t$, keeps prices affine and makes the premium move over time.

12.1.4 Testing the expectations hypothesis

If the expectations hypothesis held, the slope of the curve would predict how long yields move. With n-year zero-coupon yields, a bond held for a year becomes an $(n-1)$-year bond, and the hypothesis implies

$$y_{t+1}^{(n-1)} - y_t^{(n)} = a + b\,\frac{y_t^{(n)} - y_t^{(1)}}{n-1} + e_{t+1}, \qquad b = 1. \tag{12.8}$$

A steep curve should be followed by rising long yields, by exactly enough to offset the higher yield on the long bond. Campbell and Shiller (1991) found the opposite. On the ACM model's fitted zero-coupon yields, monthly from 1961 to 2026 with a twelve-month holding period, the estimated b is -0.46 (standard error 0.46) at two years, -0.72 (0.51) at three, -1.18 (0.58) at five and -2.20 (0.69) at ten. A steep curve is followed by falling long yields, not rising ones: the extra yield on long bonds is a premium that investors earn, not compensation for expected capital losses. The term premium of (12.3) is not a modelling device; it is what the data require.

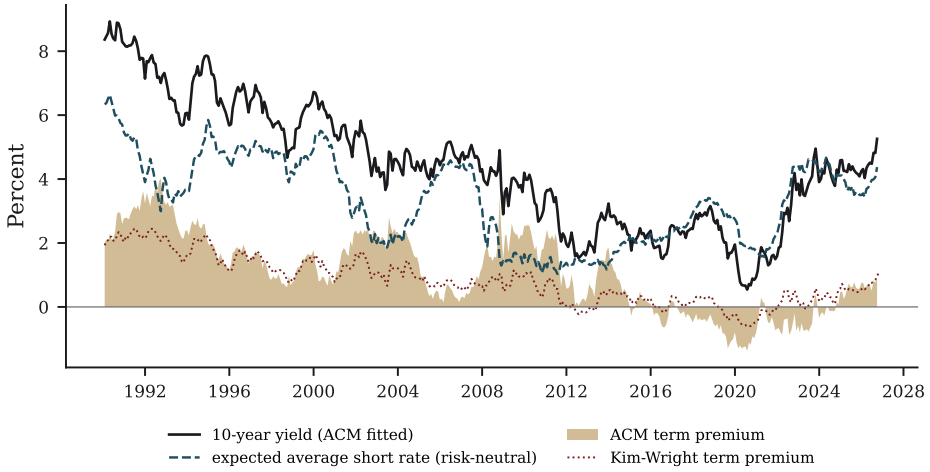


Figure 12.1. The ten-year Treasury yield split into the expected average short rate and the term premium, ACM model, monthly, 1990 to September 2026, with the Kim-Wright term premium for comparison. Sources: Federal Reserve Bank of New York; Federal Reserve Board via FRED.

12.1.5 ACM and Kim-Wright

The models the Fed staff use generalise this in two ways: several factors instead of one, and a time-varying price of risk. The Adrian et al. (2013) model, published daily by the New York Fed, takes as factors the first five principal components of Treasury yields and estimates the model in three ordinary least squares steps. First, a VAR for the factors gives their dynamics and innovations. Second, a regression of one-month excess returns on bonds of each maturity on lagged factors and factor innovations gives each bond's exposure to the risks. Third, a cross-sectional regression of the return coefficients on the exposures recovers the prices of risk λ_0 and λ_1. Risk-neutral yields follow by pricing the bonds with the prices of risk set to zero. The model of Kim and Wright (2005), published by the Board, uses three latent factors estimated by maximum likelihood and adds survey forecasts of short rates to pin down expectations.

Result: the ten-year term premium

The ACM ten-year term premium was 1.8 points in January 2000, 0.6 in mid-2007 and 0.9 in mid-2014. It fell below zero during the asset

purchases and reached its minimum, −1.66 points, on 9 March 2020. It was −1.4 in July 2020, rose to 0.45 in October 2023, fell back to −0.28 on the eve of the first cut in September 2024, and rose to 0.75 at the end of 2025 and 0.81 in October 2026 (Figure 12.1). The ACM and Kim-Wright premia have a correlation of 0.86 on daily data since 1990; they differ more in level than in direction.

12.2 The shape of the curve

12.2.1 Nelson-Siegel

The decomposition explains what a yield is. To describe the whole curve in a few numbers, Nelson and Siegel (1987) write the instantaneous forward rate at maturity τ as

$$f(\tau) = \beta_1 + \beta_2\, e^{-\lambda\tau} + \beta_3\, \lambda\tau\, e^{-\lambda\tau}. \tag{12.9}$$

Averaging the forward curve from zero to τ gives the yield curve,

$$y(\tau) = \beta_1 + \beta_2 \left(\frac{1 - e^{-\lambda\tau}}{\lambda\tau} \right) + \beta_3 \left(\frac{1 - e^{-\lambda\tau}}{\lambda\tau} - e^{-\lambda\tau} \right). \tag{12.10}$$

The loading on β_1 is one at every maturity, so β_1 is the ***level***. The loading on β_2 starts at one and decays to zero, so $-\beta_2$ is the ***slope***, long minus short. The loading on β_3 starts and ends at zero and peaks in between, so β_3 is the ***curvature***. Diebold and Li (2006) fix $\lambda = 0.0609$ with maturities in months, which puts the peak of the curvature loading at thirty months, and estimate the three factors each month by ordinary least squares. Svensson (1994) adds a second hump; the ECB's published euro area curve and the Fed's own published curve are Svensson curves.

Fitted to ten maturities from three months to thirty years, monthly from 1990 to September 2026, the model misses observed yields by 8.8 basis points on average (Figure 12.2). The level factor tracks the ten-year yield with a correlation of 0.97. The slope factor, written as short minus long, has a correlation of 0.52 with the change in the funds rate over the previous year: tightening flattens the curve and easing steepens it, because the short end moves with policy and the long end moves less.

12.2.2 Can the curve be forecast?

A forecaster of the curve has three natural candidates: the random walk, which says today's curve is the best forecast of next year's; the forward

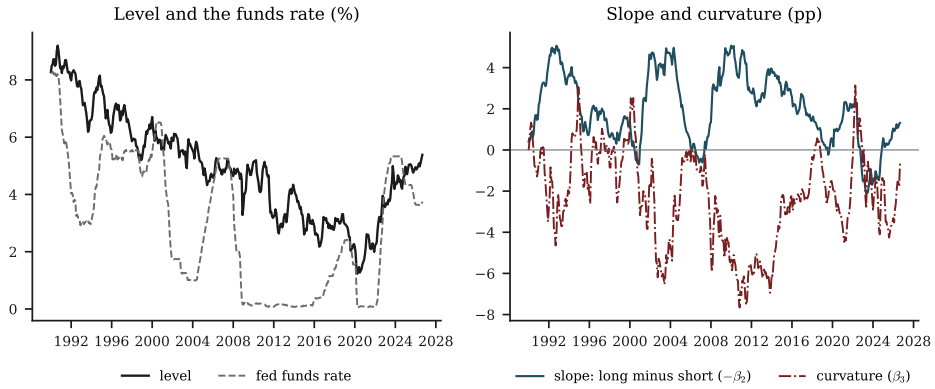


Figure 12.2. Nelson-Siegel factors of the Treasury curve, monthly, 1990 to September 2026 (ten maturities from three months to thirty years, $\lambda = 0.0609$), and the effective federal funds rate. Source: FRED; author's estimates.

Table 12.1. Twelve-month-ahead forecasts of Treasury yields, monthly origins January 2000 to September 2025. Root mean squared error and mean error (actual minus forecast), percentage points.

	RMSE			Mean error		
Forecast	2y	5y	10y	2y	5y	10y
Random walk	1.24	0.97	0.80	−0.09	−0.08	−0.06
Diebold-Li, AR(1) factors	1.25	1.01	0.83	−0.08	−0.23	−0.22
Forward rate	1.29	1.10	0.88	−0.47	−0.53	−0.36

curve, which is the forecast under the expectations hypothesis; and the Diebold-Li model, which forecasts each Nelson-Siegel factor with an AR(1) and rebuilds the curve. All three are run here out of sample, re-estimated every month from January 2000, for yields twelve months ahead.

Result: the random walk is hard to beat, and the forward curve is biased

Over 2000 to 2025, neither the Diebold-Li model nor the forward curve beats the random walk twelve months ahead (Table 12.1). The forward curve does worst and is biased: on average it predicted the ten-year yield 36 basis points higher than it turned out and the two-year yield 47 basis points higher. The average bias estimates the average term premium of (12.3): forward rates contain a premium over expected

short rates, so on average they over-predict. Any single forecast error adds the realisation error to the premium, so only the average is informative.

Diebold and Li (2006) found that their model beat the random walk at the twelve-month horizon on data to 2000. The result does not survive the quarter century since, which contained a lower bound, two rounds of asset purchases and an inflation surge, none of which an AR(1) on three factors anticipates. The practical models used on trading desks combine the structural decomposition with statistical learning. The ensemble built in this book's research workspace for the Treasury curve, for example, weights three components: a gradient-boosted model on the curve's own shape and momentum, a second on the policy cycle and mean reversion, and a structural anchor that bootstraps the market-implied path of the funds rate from bills and adds historical term premia. The structural anchor is equation (12.3) in operational form.

12.3 The slope as a recession signal

The slope of the curve has a second use, older than term-premium models: it predicts recessions. Estrella and Mishkin (1998) showed that the spread between the ten-year Treasury yield and the three-month bill rate, entered in a probit model, forecasts a recession twelve months ahead better than any other financial or macroeconomic indicator. The New York Fed publishes the model's probability each month:

$$\Pr(\text{recession}_{t+12}) = \Phi\big(\alpha + \beta\,(y_t^{10} - y_t^{3m})\big). \tag{12.11}$$

The logic follows from (12.3). An inverted curve says that the market expects short rates to fall, which it does when it expects the central bank to cut in a downturn, and that the term premium is low, which it is when investors seek the safety of duration.

Result: the longest inversion without a recession

Re-estimated on monthly data from 1959 to 2025, the probit gives $\alpha = -0.67$ and $\beta = -0.56$ (standard errors 0.08 and 0.06), close to the published coefficients. Every recession since 1960 was preceded by an inversion; the 2020 recession followed the short inversion of 2019, though outside the twelve-month horizon of the probit and from a

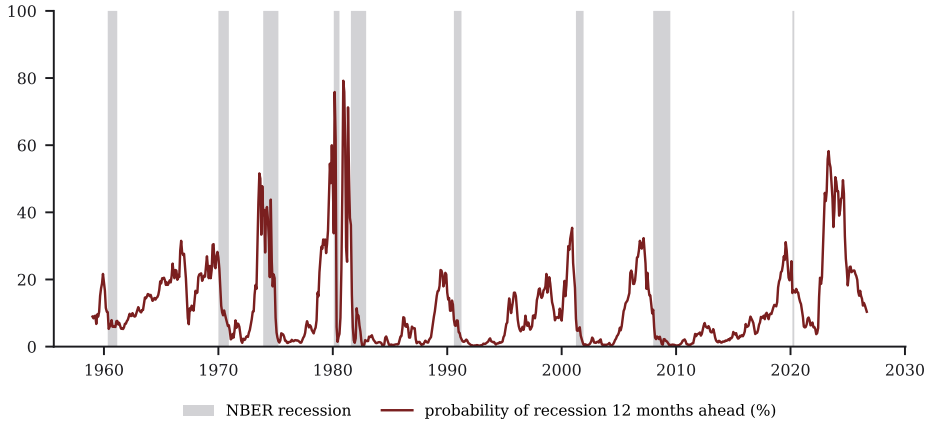


Figure 12.3. Probability of a US recession twelve months ahead from a probit on the ten-year minus three-month Treasury spread, estimated on 1959 to 2025, and NBER recessions. Sources: FRED; author's estimates.

pandemic the curve could not have priced. The inversion that began in November 2022 lasted 25 months, to November 2024, the longest in the sample, and raised the model's probability to 58% in May 2023, the highest since 1981 (Figure 12.3). No recession followed. In September 2026, with the spread at 1.05 points, the probability is 10%.

The failure of 2023 is instructive, and the decomposition explains it. The inversion was deep because the expected path of rates fell from a high level: the market expected the Fed to reverse the tightening of 2022 once inflation fell, which is what an inverted curve priced and what happened in 2024. But the expected cuts were a return from restrictive to neutral, driven by disinflation along the steep part of the Phillips curve of chapter 9, not a response to a recession. The model reads every expected cut as recession; it cannot tell a normalisation from an easing in response to a downturn. When the policy rate starts far above neutral, an inverted curve is a forecast of disinflation, not necessarily of contraction.

12.4 How policy cycles reach the ten-year yield

The decomposition answers the question every central bank asks at the start of a cycle: how much of a change in the policy rate will reach long-term rates.

Table 12.2. US policy cycles: changes in the effective funds rate and Treasury yields from the day before the first move to the last move of the cycle, and the ACM decomposition of the ten-year change. Residual: observed change minus the two ACM components, the model's fitting error, so the components do not sum exactly to the observed change. Percentage points. Pass-through: change in ten-year yield over change in the funds rate. The 2026 cycle is measured to 14 October 2026. Sources: FRED; Federal Reserve Bank of New York; author's calculations.

Cycle	Funds	2y	10y	Expect.	Premium	Resid.	Pass-through
Tightening 1994–95	+2.13	+3.05	+1.85	+1.95	−0.29	+0.19	0.87
Tightening 1999–2000	+1.22	+1.21	+0.50	+0.87	−0.49	+0.12	0.41
Easing 2001–03	−5.49	−3.57	−1.54	−2.59	+1.09	−0.04	0.28
Tightening 2004–06	+3.95	+2.38	+0.52	+1.82	−1.55	+0.25	0.13
Easing 2007–08	−5.16	−3.43	−2.11	−2.75	+1.24	−0.60	0.41
Tightening 2015–18	+2.05	+1.65	+0.49	+1.13	−0.76	+0.12	0.24
Easing 2019	−0.57	−0.24	−0.28	−0.10	−0.18	0.00	0.49
Tightening 2022–23	+5.00	+2.97	+1.71	+1.66	+0.05	0.00	0.34
Easing 2024–25	−1.44	−0.05	+0.48	−0.37	+0.97	−0.12	−0.33
Tightening 2026	+0.25	+0.08	+0.18	+0.04	+0.14	0.00	

Table 12.2 and Figure 12.4 show, for ten US cycles since 1994, the change in the funds rate, the two-year and ten-year yields, and the split of the ten-year change into its expectations component and its term premium.

Result: the term premium decides the pass-through

The expectations component of the ten-year yield moves with policy in every cycle, by between 0.16 and 0.92 of the change in the funds rate. The term premium moves against policy in seven of the ten cycles: it fell in tightenings and rose in easings. In 2004 to 2006 it fell by 1.55 points while the funds rate rose by four, which is the "conundrum" of long rates that did not rise; in 2001 to 2003 and 2007 to 2008 it rose by more than a point while the Fed cut by five. The pass-through to the ten-year yield ranged from 0.13 in 2004 to 2006 to 0.87 in 1994 to 1995.

The easing of 2024 to 2025 is the exception that defines the present. The Fed cut by 1.44 points between September 2024 and December 2025. The expectations component of the ten-year yield fell by 0.37, as it should, but the term premium rose by 0.97 and the yield rose by 0.48. Mortgage rates and corporate borrowing costs, which price off the long end, rose through an easing cycle. The usual explanation, supported by the timing, is a combi-

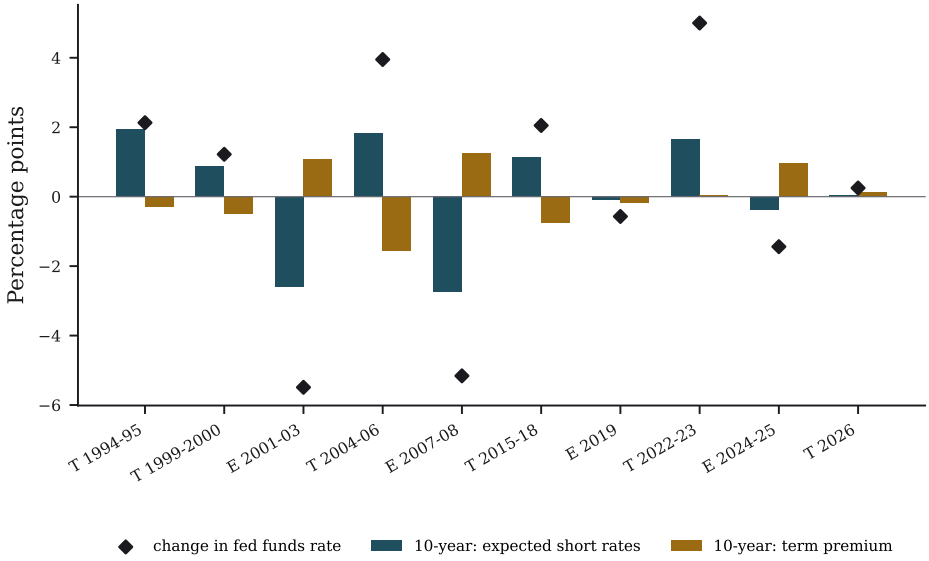


Figure 12.4. Change in the ten-year Treasury yield over each policy cycle, split into expected short rates and the term premium (ACM), and the change in the funds rate. Sources as in Table 12.2.

nation of fiscal supply, with deficits near 6% of GDP and rising issuance at long maturities, and the uncertainty that the inflation of 2021 to 2023 left about the inflation regime. Both raise the compensation investors demand for duration, and neither responds to the funds rate directly, though both can move with it through inflation expectations and the fiscal cost of debt.

The policy lesson follows from (12.3). A central bank controls the expectations component through the funds rate and guidance. It influences the term premium through its balance sheet, which is the subject of chapter 13, and through the credibility of its inflation target. When the term premium moves against policy, as it did in 2004 and 2024, a central bank that wants a given change in long rates has to move the short rate further, or use the balance sheet.

12.5 What the curve says in October 2026

12.5.1 The market-implied path

Applying (12.3) to today's curve gives the path of the funds rate the market expects. ACM's risk-neutral yields are the expected average short rate at each horizon; one-year forward rates computed from them are the expected short rate year by year. On 14 October 2026, with the target range at 3.75% to 4.00%, the risk-neutral one-year yield was 4.31%, and the expected short rate was 4.56% in the second year, 4.53% in the third and 4.41% in the fifth. The market expects the September increase to be followed by about two more, and rates to stay near 4.5% for several years.

Result: the market and the committee

The FOMC's September 2026 projections put the funds rate at 4.1% at the end of 2026 and 2027 and 3.9% at the end of 2028, against a longer-run 3.2%. The market's expected path, net of the term premium, is about half a point higher in 2027 and 2028. Forward rates without the correction would overstate the gap: the raw one-year forwards are 5.0% to 5.1%, and the difference, 0.4 to 0.7 points, is term premium, not expected policy. Reading the gap between the dots and the forward curve without a term-premium model would overstate the market's disagreement with the committee by half.

The ten-year yield averaged 4.99% in September 2026, up from 4.12% a year earlier, and reached 5.17% on 14 October. In ACM's September averages, 0.71 points of the yield was term premium and 4.26% the expected average short rate over ten years. Of the rise of about 0.8 points over the year, 0.73 was in expected short rates and less than 0.1 in the premium: unlike 2024 to 2025, the 2026 rise in long rates is a change in the market's view of where policy goes, driven by the oil shock, not a change in the price of duration.

12.5.2 The euro area

The ECB publishes a daily AAA government bond curve, a Svensson curve fitted by its staff. The same exercise on that curve shows the same pattern in two cycles (Figure 12.5, left).

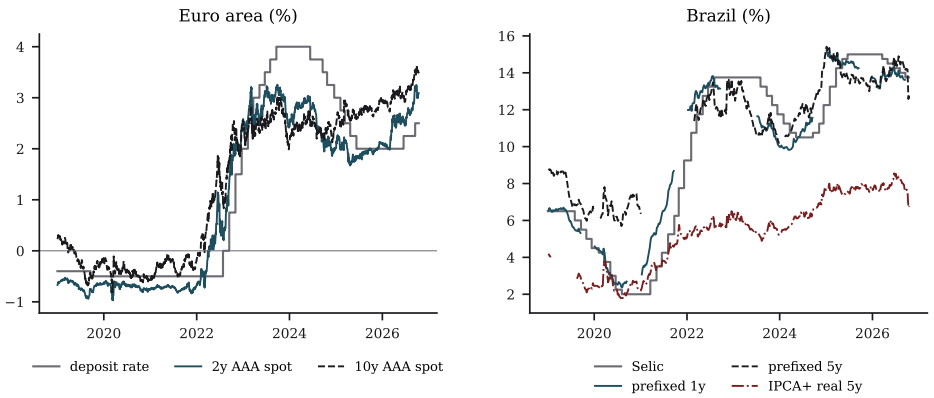


Figure 12.5. Left: euro area deposit facility rate and two- and ten-year spot rates on the ECB's AAA Svensson curve, 2019 to October 2026. Right: Selic target, Tesouro Prefixado (LTN) yields interpolated at one and five years, and the Tesouro IPCA+ (NTN-B principal) real yield at five years, weekly. Sources: ECB Data Portal; FRED; Tesouro Nacional (Tesouro Direto prices); BCB.

Result: euro area pass-through

From July 2022 to September 2023 the ECB raised its deposit rate by 4.5 points; the AAA two-year rate rose by 2.66, the ten-year by 1.46 and the thirty-year by 1.23. From June 2024 to June 2025 it cut by 2 points; the two-year fell by 1.13, the ten-year rose by 0.05 and the thirty-year by 0.31. The euro area curve repeated the US pattern of 2024 to 2025: easing did not reach the long end. With the increases of June and September 2026 the curve rose almost in parallel, by 0.57 at two years and 0.52 at ten. On 14 October 2026 the AAA curve ran from 2.92% at one year to 3.09% at two, 3.51% at ten and 3.75% at thirty, a year after 1.88%, 1.92%, 2.68% and 3.15%.

With the deposit rate at 2.50%, a two-year rate of 3.09% prices further increases, as in the United States. The long end, which rose by 0.83 points in a year, carries the fiscal expansion in Germany and the larger supply of euro area government debt as well as the energy shock.

12.5.3 Brazil

Brazil has no published term-structure model of the ACM kind at daily frequency, but its curve can be read from the prices at which the Treasury

Table 12.3. Brazil: changes in the Selic target and in Tesouro Direto yields over Copom cycles. Prefixed (nominal) rates at one, three and five years; IPCA+ real rates at five and ten years. Percentage points. The 2026 easing is measured to 14 October 2026. Sources: Tesouro Nacional; BCB; author's calculations.

Cycle	Selic	Pre 1y	Pre 3y	Pre 5y	IPCA+ 5y	IPCA+ 10y
Tightening 2021–22	+11.25	+8.81	+5.47	+6.51	+3.00	+2.71
Easing 2023–24	−3.00	−1.07	+0.65	+0.76	+1.13	+1.05
Tightening 2024–25	+4.25	+2.99	+1.74	+1.69	+1.17	+0.92
Easing 2026	−1.25	−0.18	−0.97	−1.05	−1.06	−0.68

sells its bonds to the public through Tesouro Direto: prefixed bonds (LTN) give the nominal curve and IPCA-linked zero-coupon bonds (NTN-B principal) the real curve. These are retail purchase rates, a few basis points above wholesale levels, which does not affect changes over a cycle (Figure 12.5, right).

Result: Brazil: the long end moves on fiscal news

In the 2021 to 2022 tightening the prefixed one-year rate rose by 8.8 points and the five-year by 6.5, a pass-through to the long end of over half (Table 12.3). In the easing of August 2023 to May 2024 the Selic fell by 3 points and the one-year prefixed rate by only 1.1, while the five-year rose by 0.8 and the five-year real rate by 1.1: the long end rose through the easing, as it did in the United States a year later, on doubts about the fiscal framework. The 2026 easing is the reverse: the Selic has fallen by 1.25, the one-year rate barely moved, and the five-year real rate fell by 1.06. On 14 October 2026 the prefixed curve ran from 13.6% at one year to 12.8% at five, the five-year real rate was 6.8%, and five-year breakeven inflation 5.6%.

A real rate of 6.8% at five and ten years is above every Brazilian measure of the neutral rate in chapter 7 except the market's own, and the breakeven of 5.6% is above the 4.5% ceiling of the tolerance band. The real yield contains real term and liquidity premia, and the breakeven an inflation risk premium, so neither is a pure expectation. With that caveat the pricing is consistent with the Focus expectations of chapter 10: a higher expected real-rate path and larger real and inflation risk premia than the BCB's reference values assume. The term premium that separates the two is the

BCB's quantity of interest in its neutral-rate box, where it is removed with an ACM model fitted to NTN-Bs.

12.6 What the curve means for policy

Three points carry over to the chapters that follow.

Expectations are the lever. Over ten US cycles, the expectations component of the ten-year yield moved with the funds rate every time. A central bank moves long rates mainly by moving the expected path, which is why guidance, the subject of chapter 14, is a policy instrument and not a courtesy.

The term premium is the wild card. It moved against policy in seven cycles out of ten, by up to 1.55 points. It is driven by the supply of duration, by inflation uncertainty and by the balance sheet, and its swings are as large as the policy-driven part of the yield.

Reading the curve requires a model. Forward rates are biased forecasts of policy, by about half a point at two years since 2000. A committee that compares its own projections with the forward curve, or a market economist who reads the dots against forwards, needs the term-premium correction of (12.3) before concluding that the market and the central bank disagree.

Chapter 13

Financial Conditions and the Balance Sheet

Between March 2022 and July 2023 the Fed raised its policy rate by five points, the fastest tightening in four decades. Over the same period the Chicago Fed's index of national financial conditions tightened by 0.13 of a standard deviation, the stock market rose by 7%, and the dollar ended where it started. The thirty-year mortgage rate rose by almost three points, and that was most of what the tightening did to the prices that households and firms face. Between September 2024 and December 2025 the Fed cut by 1.44 points; the mortgage rate did not move and the stock market rose by 22%.

The policy rate reaches spending through a set of asset prices, and those prices respond to much besides the policy rate. This chapter measures that set. It defines financial conditions indices and asks whether they predict growth, finds that they predict the risk to growth far better than its central tendency, measures what the Fed's balance sheet did to the term premium as a stock rather than as a sequence of announcements, and works through the one episode since 2020 in which financial conditions threatened to take policy out of the central bank's hands: the failure of Silicon Valley Bank in March 2023. The euro area and Brazil close the chapter.

13.1 What a financial conditions index measures

13.1.1 From policy to asset prices

In the model of chapter 1 the only financial price is the real short rate. A minimal extension adds an index of financial conditions fci_t, positive when conditions are tight, to the IS curve:

$$x_t = \mathbb{E}_t x_{t+1} - \frac{1}{\sigma}\left(i_t - \mathbb{E}_t \pi_{t+1} - r_t^n\right) - \psi\, fci_t, \tag{13.1}$$

where ψ measures how much tighter conditions lower demand beyond the effect of the policy rate. Most indices contain rates themselves (the FCI-G includes the funds rate and the ten-year yield), so fci_t in (13.1) has to be read as the part of conditions not already in the real policy rate, the component orthogonal to it, or the policy rate is counted twice. Equation (13.1) has a direct implication for policy. A shock to financial conditions that is unrelated to the policy rate, a fall in equity prices or a widening of credit spreads, enters exactly like a fall in the natural rate: it is a demand shock. Under the optimal policy of chapter 2, which offsets demand shocks fully, the policy rate should move by

$$\frac{\partial i_t}{\partial fci_t} = -\sigma\psi \tag{13.2}$$

for each unit of exogenous tightening. A central bank that does not offset the shock lets it pass to output and inflation. The question is then empirical: how large is ψ, how fast do conditions move, and can they be measured well enough to respond to them.

13.1.2 Two indices built by the Fed

The Federal Reserve Bank of Chicago's National Financial Conditions Index (NFCI) is a dynamic factor model of 105 weekly, monthly and quarterly indicators of risk, credit and leverage in money markets, debt and equity markets and the banking system (Brave and Butters, 2011). It is normalised to have mean zero and standard deviation one since 1971; positive values are tighter than average. The adjusted version, ANFCI, removes the part explained by economic activity and inflation, leaving the financial component.

The Board's FCI-G (Ajello et al., 2023) is built differently. It takes seven variables, the funds rate, the ten-year Treasury yield, the mortgage rate, the BBB corporate yield, equity prices, house prices and the dollar, and weights their changes by the effect each has on GDP growth over the following year in the Board's models, FRB/US among them:

$$FCIG_t = \sum_{j} \sum_{k=0}^{K} w_{j,k}\, \Delta z_{j,t-k}, \tag{13.3}$$

where $w_{j,k}$ is the impulse response of growth to a change in variable j k quarters earlier. The index is measured in percentage points of growth: a value of $+1$ means financial conditions will subtract one point from GDP growth

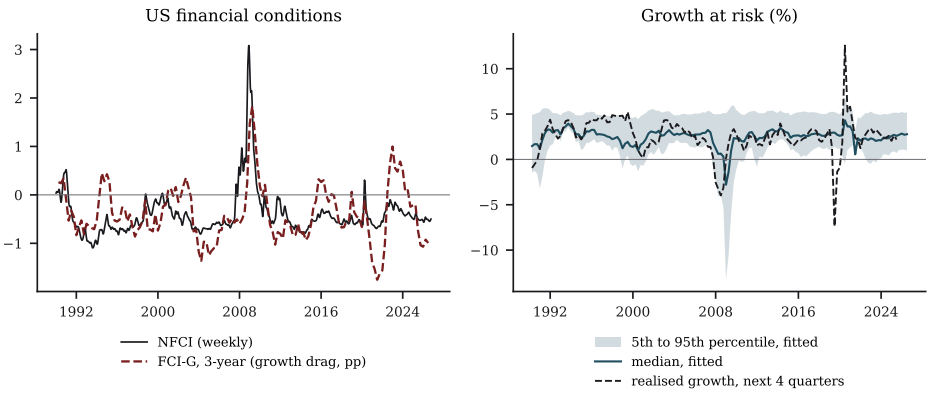


Figure 13.1. Left: Chicago Fed NFCI (weekly) and the Board's FCI-G with a three-year lookback (quarterly, percentage points of GDP growth subtracted over the next year), 1990 to 2026. Right: GDP growth over the next four quarters, realised, and its 5th, 50th and 95th percentiles fitted by quantile regression on the NFCI and current growth, 1990 to 2026. Sources: Federal Reserve Bank of Chicago, Federal Reserve Board, BEA via FRED; author's estimates.

over the next year. Unlike the NFCI, it has a structural interpretation, and its components can be read separately.

Result: financial conditions in 2026

The FCI-G with a three-year lookback stood at −0.99 in the second quarter of 2026: financial conditions were adding about one point to US growth over the following year, almost three-quarters of it (−0.71) from equity prices and −0.15 from the dollar, against contributions near zero from the funds rate and Treasury yields. The NFCI was −0.50 in the week of 9 October 2026, looser than its historical average, after the September rate increase. Its peaks were 3.08 in the 2008 crisis and 0.30 in March 2020; the FCI-G's was 1.83 in the first quarter of 2009 (Figure 13.1).

The contrast between the two halves of that reading is the point of the chapter's opening. A central bank that has raised rates by five points and then held them at 3.6% to 4% has not, by the FCI-G's measure, tightened financial conditions in 2026: the stock market has offset the policy rate.

Table 13.1. US GDP growth over the next four quarters: in-sample fit, 1990 to 2025, and out-of-sample root mean squared error, forecasts from 2000Q1 with expanding estimation windows. The benchmark uses past growth only. Newey-West standard errors in parentheses.

Index	Coefficient	R^2	RMSE, out of sample
None (past growth only)	—	0.00	2.41
NFCI	−2.08 (0.63)	0.19	2.40
ANFCI	−1.54 (0.50)	0.19	2.42
FCI-G (one-year lookback)	−0.70 (0.39)	0.04	2.47

13.2 Do financial conditions predict growth?

13.2.1 The central tendency

A natural test regresses GDP growth over the next four quarters on current growth and an index:

$$g_{t\to t+4} = a + b\, g_{t-4\to t} + c\, FCI_t + e_t. \tag{13.4}$$

Over 1990 to 2025, in sample, the NFCI coefficient is −2.08 (standard error 0.63): a one-standard-deviation tightening lowers growth over the next year by two points, and the regression explains 19% of the variation. The ANFCI gives −1.54 (0.50) and the same fit; the FCI-G, −0.70 (0.39), explains 4%.

Result: no gain for the central forecast

Out of sample, forecasting from 2000 with expanding windows on current-vintage data (GDP and the indices as revised today, so not a real-time test), none of the indices improves on past growth: the root mean squared errors are 2.40 to 2.47 points against 2.41 (Table 13.1). The in-sample fit comes almost entirely from the 2008 crisis, when financial conditions and growth collapsed together, and from 2020; outside those episodes the indices add little to a forecast of average growth.

13.2.2 The risk to growth

Adrian et al. (2019) showed that the result reverses for the tails. Replacing the least-squares regression with a ***quantile regression***, which estimates

the q-th percentile of growth conditional on the regressors by minimising

$$\min_{\beta} \sum_{t} \rho_q\big(g_{t\to t+4} - X_t'\beta\big), \qquad \rho_q(u) = u\,\big(q - \mathbf{1}\{u < 0\}\big), \tag{13.5}$$

allows the effect of financial conditions to differ across the distribution of outcomes.

Result: growth at risk

The NFCI coefficient is -4.41 (standard error 0.38) at the 5th percentile of next-year growth, -2.06 (0.32) at the median and -0.32 (0.63) at the 95th. Tight financial conditions shift the lower tail of growth down by more than twice as much as the centre and leave the upper tail unchanged. In the third quarter of 2008 the fitted 5th percentile of growth over the next year was -5.3%, with the median at 0.3%; realised growth was -3.1%. In the second quarter of 2026, with the NFCI at -0.50, the fitted 5th, 50th and 95th percentiles were 1.0%, 2.8% and 5.2% (Figure 13.1, right).

Financial conditions are information about risk, not about the central forecast. For a central bank with the loss function of chapter 2, which penalises deviations quadratically, a fatter lower tail matters even when the median forecast is unchanged: the expected loss rises with the variance of outcomes and, under asymmetric objectives, with downside tail risk. That is the formal case for the "risk management" language central banks use when financial conditions tighten, and for the Fed's insurance cuts of 2019, taken with the median forecast close to target.

13.3 The policy rate and financial conditions in the last three cycles

Table 13.2 shows the link between the policy rate and financial conditions over the last three cycles. In the 2022 to 2023 tightening the policy rate moved by five points and the broad index barely moved; the mortgage rate carried most of the tightening, and the equity market gained. In the 2024 to 2025 easing the mortgage rate was flat because the term premium of chapter 12 rose by a point, and equities rose by a fifth. The policy rate is one input to financial conditions among several, and in the last three years it has rarely been the dominant one.

Table 13.2. Changes in the funds rate and financial conditions over the last three Fed cycles. Sources: FRED (Chicago Fed, Freddie Mac, S&P Dow Jones Indices, Federal Reserve Board); author's calculations.

Cycle	Funds	NFCI	Mortgage	S&P 500, %	Dollar, %
Tightening 2022–23	+5.00	+0.13	+2.93	+7.1	+0.4
Easing 2024–25	−1.44	−0.10	−0.01	+22.2	−0.8
Tightening 2026	+0.25	+0.04	+0.64	+2.2	+2.3

13.4 How a policy surprise moves financial conditions

The cycle-by-cycle comparison mixes the effect of policy with everything else that happened in each cycle. A policy component can be separated with the high-frequency surprises of chapter 14: the change in short-rate futures in the thirty minutes around each FOMC announcement, purged of the part predictable from public data before the announcement (Bauer and Swanson, 2023), which is designed to isolate an exogenous policy component under the identification assumptions of high-frequency studies, information effects included. Aggregated to months, these surprises are a monetary policy shock that can be fed into a local projection (Jordà, 2005),

$$y_{t+h} - y_{t-1} = a_h + b_h\, \varepsilon_t^{mp} + \sum_{k=1}^{12} c_{h,k}\, \Delta y_{t-k} + \sum_{k=1}^{12} d_{h,k}\, \varepsilon_{t-k}^{mp} + u_{t+h}, \quad (13.6)$$

estimated separately for each horizon h. The coefficient b_h is the response of y after h months to a surprise. The sample is 1988 to 2023, 432 months; the standard deviation of the monthly surprise is 6 basis points, and the responses are scaled to a 10 basis point tightening surprise. Among the outcomes is the ***excess bond premium*** of Gilchrist and Zakrajšek (2012), published by the Board: the part of corporate bond spreads not explained by expected default, a measure of the risk appetite of lenders.

Result: a policy surprise tightens financial conditions first, and activity later

A 10 basis point tightening surprise raises the two-year yield by 8 basis points on impact. Financial conditions tighten within three months: the NFCI rises by 0.08 standard deviations (standard error 0.04) and the excess bond premium by 8 basis points (0.05), both fading within



Figure 13.2. Responses to a 10 basis point tightening surprise, monthly local projections, 1988 to 2023, with 90% confidence bands. Shock: Bauer-Swanson orthogonalised monetary policy surprise. Sources: Bauer and Swanson (2023); Federal Reserve Board (excess bond premium); FRED; author's estimates.

a year. Industrial production falls steadily, by 0.9% after a year and 1.4% after two (standard error 0.46); the CPI falls by 0.4% after two years (0.22); the unemployment rate rises by about a quarter of a point after eighteen months, imprecisely (Figure 13.2).

The ordering is the transmission of (13.1) in the data: the surprise moves risk premia and credit spreads quickly, the tighter conditions then lower activity, and prices follow last. The size is large relative to the shock, as Bauer and Swanson (2023) also find: a surprise is the part of policy the market did not expect, and the response to it includes the revision of the expected path that the surprise signals. A 10 basis point surprise is a statement about the next year of policy, not a 10 basis point change in the rate.

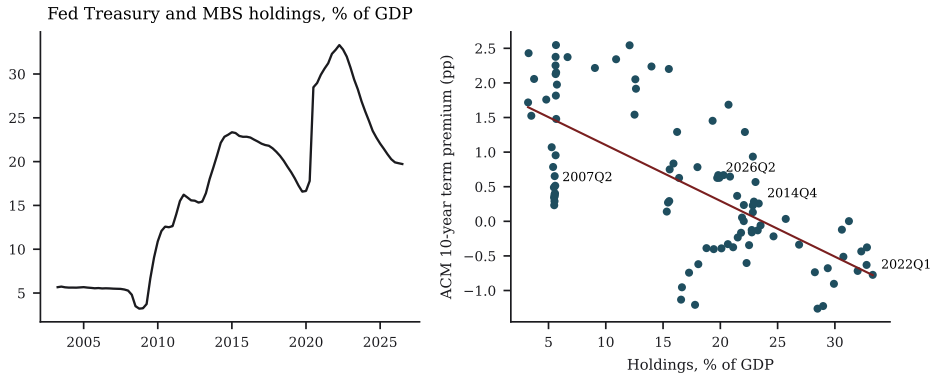


Figure 13.3. Left: Federal Reserve holdings of Treasury securities and mortgage-backed securities, percent of nominal GDP, 2003 to 2026. Right: the ACM ten-year term premium against holdings, quarterly. Sources: Federal Reserve, BEA via FRED; Federal Reserve Bank of New York.

13.5 The balance sheet as a stock

13.5.1 From flows to stocks

Chapter 6 measured the effect of asset purchases on the days they were announced: the first round of purchases lowered the ten-year yield by about a point, of which four-tenths was term premium. The preferred-habitat theory of that chapter implies that what matters for the term premium is the stock of duration the central bank holds relative to what the market has to absorb, not the announcement. A direct test regresses the ACM ten-year term premium on the Fed's holdings of Treasuries and mortgage-backed securities as a share of GDP, controlling for the stock of federal debt and for long-run inflation expectations:

$$tp_t^{(10)} = a + b\,\frac{\text{Fed holdings}_t}{\text{GDP}_t} + c\,\frac{\text{Federal debt}_t}{\text{GDP}_t} + d\,\pi_t^{e,10} + e_t. \tag{13.7}$$

Result: a stock effect of about 9 basis points per point of GDP

On quarterly data from 2003 to 2026, each percentage point of GDP in Fed holdings lowers the ten-year term premium by 8.8 basis points (standard error 3.2); the federal debt ratio has no measurable effect once holdings are controlled for, and long-run inflation expectations

enter with a coefficient of 0.93 (Figure 13.3). Holdings rose from 5.5% of GDP in mid-2007 to 23.4% at the end of 2014 and 33.3% in early 2022, and fell under quantitative tightening to 19.7% in mid-2026.

A point of GDP is about 300 billion dollars, so the estimate implies a fall of about 15 basis points in the term premium per 500 billion dollars of purchases, in the range of published estimates for the US purchase programmes. The regression is an association, not a causal estimate, so it is used here as a calibration. Applied to the whole path, it associates quantitative tightening from 2022 to 2026, a fall of 13.6 points of GDP in holdings, with about 1.2 points on the ten-year term premium, roughly the rise the ACM model shows from early 2022. That is an upper bound: term premia fell after 2008 and rose after 2021 for reasons the regression attributes partly to the balance sheet, the decline and recovery of inflation uncertainty above all. The evidence is consistent with the stock of duration the Fed holds being a first-order determinant of long rates, and that a reduction in it tightens financial conditions through the channel the policy rate does not reach.

13.5.2 Quantitative tightening and reserves

Chapter 6 described the mechanics of the run-off: the overnight reverse repo facility fell from 2.3 trillion dollars to almost nothing, absorbing most of the early run-off and cushioning the decline in reserves, which let the balance sheet shrink for two years without strain in money markets. The stock estimate adds the other side. The run-off withdrew duration as well as reserves, and the term premium rose as it did. A central bank that wants to keep the policy stance unchanged while shrinking its balance sheet has to set the policy rate lower than it otherwise would. FRB/US gives the exchange rate between the two (chapter 11): a 50 basis point rise in the ten-year term premium lowers the output gap by about a fifth as much as a funds rate 100 basis points higher for a year, so on the calibration above a rise of 1.2 points is worth roughly half of that year of tighter policy.

13.6 Silicon Valley Bank, March 2023

On 8 March 2023 Silicon Valley Bank announced losses on its securities portfolio and a capital raise; on 10 March it was closed. On 12 March the Treasury, on the recommendation of the FDIC and Fed boards, invoked the systemic risk exception, which let the FDIC make all depositors of SVB and

Signature Bank whole; the Fed opened the Bank Term Funding Program, with a Treasury backstop, which lent against Treasury and agency securities at par rather than at market value.

Result: the financial shock and the backstop

The two-year Treasury yield fell from 5.05% on 8 March to 4.03% on 13 March and 3.76% on 24 March, 129 basis points in twelve trading days, the largest decline over twelve trading days since at least 1990, larger than in March 2020 or September 2001. The Fed's balance sheet, which had been shrinking under quantitative tightening, rose from 8,342 billion dollars on 8 March to 8,734 billion on 22 March, as banks borrowed at the discount window and the new facility. The NFCI moved only from -0.21 to -0.15. On 22 March the FOMC raised the funds rate by 25 basis points.

The episode separated two instruments. The balance sheet, through lending at par, dealt with the run; the policy rate continued to deal with inflation. The NFCI's small move is consistent with the backstop working as much as with a small shock, though a before and after comparison cannot separate the two: the market priced a large probability of rate cuts, visible in the two-year yield, and then withdrew it as the backstop held. The lesson for the IS curve (13.1) is that fci_t can be moved by tools other than the policy rate, and that separating the two lets a central bank keep its policy rate on its inflation objective through a financial shock.

13.7 The euro area and Brazil

13.7.1 Systemic stress in the euro area

The ECB's Composite Indicator of Systemic Stress (CISS) aggregates stress in money, bond, equity, foreign exchange and financial intermediary markets, weighting each by its correlation with the others so that the index rises most when stress is widespread (Holló et al., 2012). It runs from zero to one. Its peaks were 0.94 in November 2008, 0.72 in October 2011 during the sovereign crisis, 0.69 in April 2020 and 0.74 in October 2022, at the height of the energy crisis and the gilt market turmoil (Figure 13.4, left). It stood at 0.01 on 14 October 2026, against a median of 0.09 since 2000: euro area markets are calm by this measure despite the energy shock, the September rate increase and a long end that rose by 0.8 points in a year.

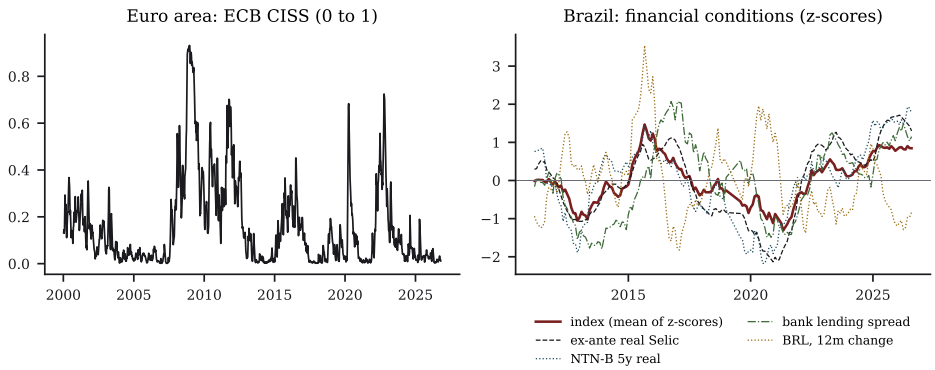


Figure 13.4. Left: ECB Composite Indicator of Systemic Stress, weekly averages, 2000 to October 2026. Right: a Brazilian financial conditions index, the mean of standardised ex-ante real Selic, five-year NTN-B real yield, bank lending spread and twelve-month change in BRL per dollar, and its components, monthly, 2011 to 2026. Sources: ECB Data Portal; BCB; Tesouro Nacional; author's calculations.

The ECB's instrument against the fragmentation that drove the 2011 peak, the Transmission Protection Instrument, was described in chapter 6.

13.7.2 An index for Brazil

Brazil has no public financial conditions index of the NFCI kind at monthly frequency. A simple one averages the standardised values of four variables that the BCB's own analysis of transmission emphasises: the ex-ante real Selic, the five-year real yield on NTN-Bs, the average bank lending spread and the twelve-month change in the exchange rate (Figure 13.4, right).

Result: Brazilian financial conditions

The index was tightest in September 2015, at 1.47 standard deviations above its mean, during the fiscal crisis and the loss of the investment grade, and loosest in May 2021, at −1.31, with the Selic at its low. It tightened to 0.87 in mid-2025 and stood at 0.85 in August 2026. As a predictor of growth in the IBC-Br over the next twelve months it adds nothing: the coefficient is −0.98 with a standard error of 0.88 and an R^2 of 0.03.

The weak result has a reason. Two of the four components, the long real rate and the exchange rate, tighten in Brazil mainly on fiscal and external

news, which also shifts the expected path of the Selic and of fiscal policy. An index that mixes the policy stance with the risk premium measures the joint outcome, and the growth that follows depends on which of the two moved. It is the same identification problem as the Focus-implied neutral rate of chapter 7, which mixes the neutral rate with survey views on fiscal and inflation risk; the survey carries no duration premium of its own, but its forecasters price the same risks.

13.8 What financial conditions mean for policy

Respond to exogenous shocks to conditions, not to conditions. By (13.1), a tightening of financial conditions caused by the policy rate is the transmission working; one caused by a shock to risk appetite is a demand shock to offset. A rule that responds to the level of an FCI confuses the two. The useful information in an index is its change not explained by policy.

Use the tail. Financial conditions say little about median growth and a great deal about the probability of a bad outcome. Their place in the reaction function is in the assessment of risks around the forecast.

Keep instruments separate. The balance sheet moves the term premium and can stop a run; the policy rate moves expected short rates. In 2023 the Fed used the first to contain a bank failure and the second to keep tightening, and in 2022 to 2026 the run-off of its holdings is associated, on the stock estimate above, with about a point on the term premium. The next part of the book turns to the third instrument, words.

Part IV

Words as an Instrument

Chapter 14

Forward Guidance

On 9 August 2011 the FOMC replaced its statement that rates would stay low "for an extended period" with a date: at least through mid-2013. The funds rate did not change, and there was no new asset purchase. In the thirty minutes after the release the two-year Treasury yield fell by 9 basis points, the ten-year by 6, and the stock market fell by 1%. On 11 December 2024 the Copom raised the Selic by 100 basis points, 25 more than the Focus median expected, and said it anticipated two further increases of the same size; the three- and five-year prefixed rates fell the next morning. On 4 July 2013 the ECB, which until then had said it never pre-committed, announced that it expected rates to remain at present or lower levels for an extended period, and Italian ten-year yields fell by 13 basis points during the press conference.

All three are forward guidance: statements about the future path of policy, used as an instrument in their own right. Chapter 12 showed that the expectations component of long rates is what moves with policy; guidance acts on it directly. This chapter sets out the theory, including a result that makes guidance look implausibly powerful in the standard model, measures its effects with the high-frequency data the Fed, the ECB and academic researchers publish, separates guidance that commits from guidance that reveals, and reads the BCB's use of it against the Focus survey.

14.1 Two kinds of guidance

Campbell et al. (2012) distinguished two things a statement about the future can do. ***Odyssean*** guidance commits: like Odysseus tied to the mast, the central bank binds itself to a path it might later want to leave, and the commitment is valuable for the reason chapter 2 gave, that history dependence improves the trade-off between inflation and output. ***Delphic*** guidance forecasts: the central bank tells the public what it expects to do because of what it expects the economy to do. A Delphic statement that rates will stay low

for long is also a statement that the economy will be weak for long, and can depress spending rather than raise it.

The distinction matters because the two have opposite effects on private expectations of the economy, and so on the asset prices that depend on them. An Odyssean easing lowers rates and raises equity prices: the same economy with easier money. A Delphic easing lowers rates and lowers equity prices: the central bank has revealed bad news. That difference in the sign of the comovement between rates and stocks is how the two are separated in the data.

14.2 Guidance in the core model

14.2.1 The forward guidance puzzle

Iterating the Euler equation of chapter 1 forward from (1.33),

$$x_t = -\frac{1}{\sigma}\sum_{j=0}^{\infty}\mathbb{E}_t\big(i_{t+j} - \mathbb{E}_{t+j}\pi_{t+j+1} - r^n_{t+j}\big). \tag{14.1}$$

Today's output gap depends on the whole future path of real rates, with equal weight on every period. Suppose the central bank announces credibly that the real rate will be one point (annualised) lower in quarter k only, and unchanged in every other quarter. By (14.1), the output gap rises by $(1/\sigma)\times 0.25 = 0.25\%$ with $\sigma = 1$ in every quarter from today to k, whatever k is. Inflation, from the Phillips curve, is the discounted sum of future gaps,

$$\pi_t = \kappa\sum_{j=0}^{\infty}\beta^j\mathbb{E}_t x_{t+j}, \tag{14.2}$$

and so rises with the number of quarters for which the gap is higher: by $\kappa\,0.25\,(1-\beta^{k+1})/(1-\beta)$ in quarterly terms.

Result: the forward guidance puzzle

A one-quarter cut in the real rate raises inflation today by 0.17 points if it happens now, 0.84 if it is promised a year ahead, 2.1 if three years ahead and 5.8 if ten years ahead (Figure 14.1). The further away the promise, the larger its effect today. The output gap rises by 0.25% today regardless of the horizon. No central bank has seen anything like this; Del Negro et al. (2023) call it the forward guidance puzzle.

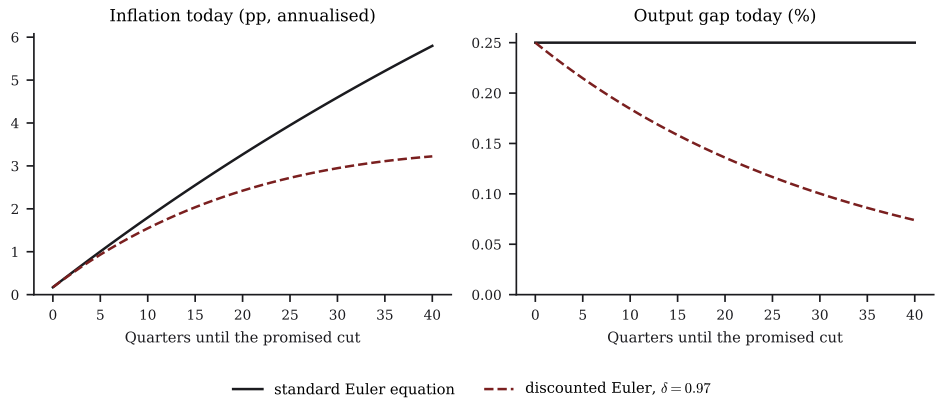


Figure 14.1. Effect today of a credible promise to lower the real rate by one point (annualised) for one quarter, k quarters ahead, in the core model (calibration of Table 1.1), with the standard Euler equation and with a discounted Euler equation. Author's calculations.

14.2.2 Resolving it

Three resolutions have been proposed, and each has a policy implication.

Discounting. McKay et al. (2016) show that if households face uninsurable income risk and borrowing limits, the aggregate Euler equation becomes

$$x_t = \delta\, \mathbb{E}_t x_{t+1} - \frac{1}{\sigma}\big(i_t - \mathbb{E}_t \pi_{t+1} - r_t^n\big), \qquad \delta < 1, \tag{14.3}$$

so that news about distant real rates is discounted. With $\delta = 0.97$ a quarter, the output effect of the promise falls from 0.25 to 0.07 at ten years and the inflation effect from 5.8 to 3.2 points: still large, but no longer explosive in the horizon.

Finite planning horizons and attention. If households plan only a few quarters ahead, or attend imperfectly to statements about the far future (chapter 5), distant promises reach their decisions weakly. The result is a discount of the same form.

Credibility. The derivation assumes the promise is believed. A promise to keep rates low after the economy has recovered is exactly the kind a future committee will want to break, and the public knows it. The time-consistency problem of chapter 4 applies to guidance: the more valuable the commitment, the less credible it is.

The implication for practice is that guidance works best at short and medium horizons, where it is both believed and not discounted away, and

Table 14.1. Response of asset prices in thirty-minute windows around scheduled FOMC announcements to the three Swanson factors, July 1991 to June 2019. Basis points for yields, percent for the S&P 500, per one-standard-deviation surprise. Robust standard errors in parentheses. Sources: Swanson (2021) factors; Bauer and Swanson (2023) announcement data; author's regressions.

	2-year	5-year	10-year	S&P 500
Funds rate factor	3.15 (0.41)	2.03 (0.28)	1.16 (0.18)	−0.24 (0.08)
Forward guidance factor	4.29 (0.18)	4.53 (0.16)	3.42 (0.11)	−0.19 (0.04)
Asset purchase factor	0.09 (0.25)	−2.06 (0.33)	−3.20 (0.37)	0.08 (0.09)
R^2	0.89	0.91	0.92	0.21

when it is tied to conditions the public can check, which reduces the temptation to renege. That is the history of guidance at the three banks.

14.3 Measuring guidance: the US

14.3.1 Three dimensions of a policy announcement

Gürkaynak et al. (2005) showed that the asset-price responses to FOMC announcements in narrow windows are driven by two factors: a surprise in the current funds rate and a surprise in its expected path. Swanson (2021) added a third, for asset purchases, identified as the part of the response of long yields not explained by the other two. Each factor is normalised to a standard deviation of one. On 226 scheduled FOMC announcements from July 1991 to June 2019, regressing thirty-minute changes in Treasury yields and the stock market on the three factors gives Table 14.1.

Result: guidance moves the curve further out than the rate does

A one-standard-deviation surprise in the funds rate moves the two-year yield by 3.2 basis points and the ten-year by 1.2. A guidance surprise of the same size moves them by 4.3 and 3.4: almost all of the effect reaches the ten-year, against little more than a third for the funds rate (Table 14.1). Asset purchases move the ten-year by 3.2 basis points and leave the two-year untouched. The three factors explain about 90% of the variance of Treasury yields in announcement windows.

14.3.2 Episodes

The factors identify the content of individual announcements. On 9 August 2011, the introduction of calendar guidance, the guidance factor was −1.52 standard deviations: a large easing of the expected path, with no change in the rate. The two-year yield fell by 9.4 basis points and the ten-year by 5.7. On 30 January 2019, when the FOMC said it would be "patient" and ended the run of increases, the guidance factor was −0.67 and the two-year fell by 4.3 basis points.

The taper tantrum of 2013 is the counter-example, guidance about the balance sheet that markets read as guidance about rates. On 22 May Chairman Bernanke said in testimony that purchases could be reduced "in the next few meetings". Between 21 May and 5 September the ten-year yield rose by 1.17 points. In the ACM decomposition of chapter 12, 0.93 of that was term premium and only 0.24 an increase in expected short rates. The FOMC announcement of 19 June 2013, at which the taper plan was described, had a guidance factor of +1.28 and an asset-purchase factor of −1.96: the market read a tightening of both.

14.3.3 Information and policy surprises in the data

Jarociński and Karadi (2020) separate policy surprises from information surprises by the sign of the comovement between rates and stocks in the announcement window. Applying the simplest version of their rule to the Bauer and Swanson (2023) surprise series, an announcement is an ***information*** surprise, Delphic in spirit, when the policy surprise and the stock market move in the same direction, and a ***pure policy*** surprise otherwise (Figure 14.2). The split identifies whether news about the economy dominated the window, not whether the central bank committed: commitment, the Odyssean case, is a property of the guidance, which a sign restriction cannot observe.

Result: three in ten FOMC surprises are information surprises

Of 230 scheduled FOMC announcements with a surprise of at least half a basis point, 30% moved rates and stocks in the same direction. The share has been stable: 30% in 1988 to 1999, 28% in 2000 to 2007, 30% in 2008 to 2015 and 35% in 2016 to 2023. In pure policy announcements, a 10 basis point tightening surprise lowers the S&P 500 by 0.64% (standard error 0.07); in information ones it raises it by 0.80%

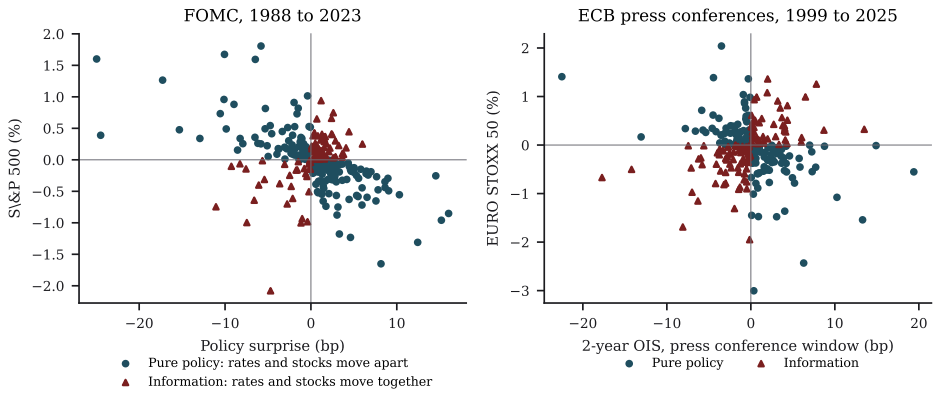


Figure 14.2. Policy surprises and stock returns in announcement windows. Left: Bauer-Swanson monetary policy surprise and S&P 500 return, scheduled FOMC announcements, 1988 to 2023. Right: two-year OIS change and EURO STOXX 50 return in the ECB press conference window, 1999 to 2025. Sources: Bauer and Swanson (2023); Euro Area Monetary Policy event study Database (Altavilla et al., 2019).

(0.15). Both raise the ten-year yield by about half a basis point per basis point of surprise.

The calendar guidance of August 2011 falls in the information group: rates and stocks fell together, consistent with a promise to hold rates at zero for two years telling the market how weak the Fed thought the economy was, though markets were still volatile after the US sovereign downgrade four days earlier, so the sign is a weak guide to the message. The "patient" statement of January 2019 falls in the pure policy group: rates fell and stocks rose by 0.9%. A central bank that wants guidance to stimulate has to make it Odyssean, by tying it to its own reaction function rather than to its forecast. The Fed's December 2012 thresholds, which promised no increase while unemployment stayed above 6.5% and projected inflation below 2.5%, were designed for that purpose: they described a rule, not a forecast.

14.3.4 The puzzle measured

The forward guidance puzzle is a statement about the effect of guidance on output and inflation, not on yields. It can be confronted with data by aggregating the Swanson factors to months and estimating local projections

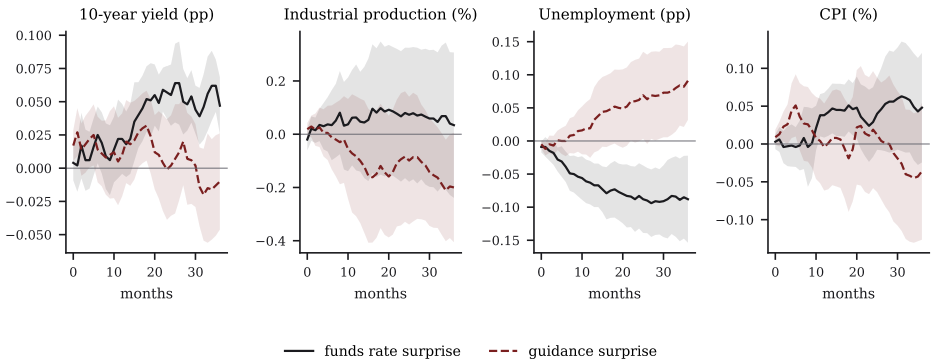


Figure 14.3. Responses to funds rate and forward guidance surprises scaled to raise the two-year Treasury yield by 10 basis points, monthly local projections, July 1991 to December 2019, with 90% confidence bands. Sources: Swanson (2021) factors; FRED; author's estimates.

of the economy on them, as chapter 13 does for the overall policy surprise. To compare like with like, each factor is scaled so that the surprise raises the two-year yield by 10 basis points in the announcement window: about 3.2 standard deviations of the funds rate factor and 2.3 of the guidance factor. The sample is July 1991 to December 2019.

Result: guidance works, modestly

A guidance tightening that raises the two-year yield by 10 basis points raises the unemployment rate by 0.06 points after two years and 0.09 after three (standard error 0.04), and lowers industrial production by 0.2% after three years, imprecisely; its effect on the CPI is not distinguishable from zero. The surprise in the current funds rate gives responses with the wrong sign: unemployment falls and prices rise after a tightening (Figure 14.3).

Two lessons follow. The empirical effects of guidance are an order of magnitude smaller than the puzzle's: in the model, a promise that raised the expected real rate path by 10 basis points for two years would lower output by about 0.2% immediately, and the effect of a promise grows without limit in the horizon. In the data, guidance moves the economy in the right direction, slowly and by little. And the wrong-signed response to funds rate surprises is the information effect of the previous section in a different form:

surprises in the current rate are disproportionately information surprises, a tightening that tells the market the Fed sees a stronger economy than it thought. The orthogonalised surprises of chapter 13, which remove the part of each surprise predictable from public news, give the expected signs; the raw factors do not. Identification is the whole difficulty.

14.4 The ECB

The ECB began giving guidance in July 2013, made it state-contingent in 2016 and 2019, tied it to inflation outcomes in July 2021 after its strategy review, and abandoned it in July 2022 for decisions taken "meeting by meeting" as inflation rose. The Euro Area Monetary Policy event study Database of Altavilla et al. (2019), maintained by the ECB, measures asset-price changes separately in the window around the press release, which contains the rate decision, and the window around the press conference, which contains most of the guidance.

Result: the press conference is where guidance happens

On 263 press conferences from 1999 to 2025, a one basis point change in the two-year OIS rate in the press conference window moves the German two-year yield by 1.03 basis points, the German ten-year by 0.54, the Italian ten-year by 0.48 and the euro by 0.044% against the dollar. Measured on the same instrument in both windows, the decision moves the short end and the words the rest of the curve: the standard deviation of the one-month OIS change is 3.2 basis points in the release window and 1.6 in the conference window, while that of the two-year OIS is 2.5 in the release window and 4.2 in the conference window, and of the German ten-year 1.9 and 3.1. In 45% of press conferences with a two-year OIS move of more than half a basis point, rates and stocks moved together, a higher information share than at the Fed.

The introduction of guidance on 4 July 2013 lowered the two-year OIS rate by 4.5 basis points and the Italian ten-year yield by 13, and raised stocks by 1.4%: a pure policy easing by the sign rule, aimed at separating euro area rates from the US taper tantrum. On 21 July 2022 the ECB raised rates by 50 basis points against an expected 25, and the one-month OIS rose by 2.6 basis points in the release window; in the press conference, where the Transmission Protection Instrument was presented, the two-year

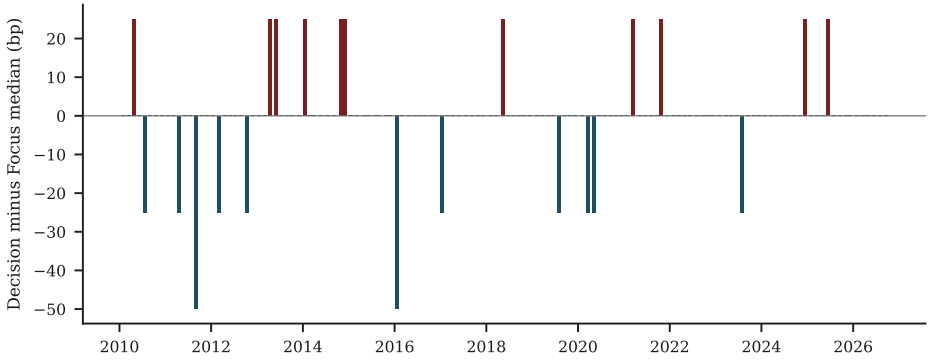


Figure 14.4. Copom decisions minus the Focus median expectation for the meeting, basis points, January 2010 to September 2026. Sources: BCB (Focus, SGS); author's calculations.

OIS fell by 6.2 and the German ten-year by 9.4. On the narrative reading, the decision tightened and the words eased, two instruments pointing in different directions on the same day, though the conference window held other news as well.

14.5 The BCB

14.5.1 Surprises against the Focus survey

The BCB's guidance is read against the Focus survey, which since 2010 has asked for the expected Selic at each Copom meeting. The surprise at a meeting is the decision minus the Focus median in the last survey before it. Over 134 meetings from January 2010 to September 2026, the Copom matched the median in 84% of them, and the average absolute surprise was 4.5 basis points (Figure 14.4).

Result: the Copom has become more predictable

The average absolute surprise fell from 6.2 basis points in 2010 to 2015 to 3.9 in 2016 to 2019 and 3.2 in 2020 to 2026. The two largest surprises, both of 50 basis points, are the cut of August 2011, when the market expected a hold, and the hold of January 2016, when it expected an increase; both preceded periods in which the Copom's credibility was

questioned. Since February 2016 no surprise has exceeded 25 basis points.

14.5.2 Guidance in two cycles

The Copom used explicit guidance at both ends of the 2020 to 2025 cycle. In August 2020, with the Selic at 2%, it adopted conditional guidance: it did not intend to reduce the degree of stimulus as long as expectations and projections stayed below the target and fiscal policy was unchanged. The guidance was conditional, tied to observable variables rather than to a forecast, which made it accountable: when expectations crossed the target the guidance was dropped and the Copom raised rates in March 2021 by 75 basis points, 25 more than the Focus median, a surprise that lifted the one-year prefixed rate by about 20 basis points the next morning.

In December 2024, with Focus expectations rising towards 5%, the Copom raised the Selic by 100 basis points against an expected 75 and announced two more increases of the same size. The one-year rate did not trade cleanly at that point of the curve, but the three-year prefixed rate fell by 10 basis points and the five-year by 17 the next morning: a tightening surprise that lowered longer rates, consistent with guidance that buys credibility, though a daily retail quote also takes in the day's other news. The two promised increases were delivered in January and March 2025.

Across all 134 meetings the response of Tesouro Direto prefixed rates to surprises is small and imprecise, 0.27 basis points per basis point of surprise at one year with a standard error of 0.19. The data are the cause as much as the economics: retail purchase rates measured once a day include all the other news of the day, and the market's view of fiscal policy, which drives the long end in Brazil (chapter 12), moves more on most days than any Copom surprise. A high-frequency study of Brazilian guidance needs intraday DI futures, which are not public.

14.6 Lessons for guidance

Describe the reaction function. Guidance that describes the reaction function moves rates and preserves confidence; guidance that reveals a forecast can move rates and depress confidence. Thresholds and conditions, as in the Fed's 2012 statement and the BCB's 2020 guidance, are the practical device; they become Odyssean only when the commitment behind them is credible.

Keep it near. The forward guidance puzzle says distant promises are powerful in models; the evidence and the time-consistency problem say they are weak in practice. Guidance is most reliable over the horizon of the next few meetings.

Know when to drop it. The ECB in 2022 and the Fed in 2021 held guidance written for a world of low inflation into one of high inflation, and paid for it with the delays measured in chapter 10. The value of guidance is the value of commitment, and the cost is the same: a commitment that turns out wrong has to be broken, and every broken commitment makes the next one less credible. The next chapter turns from what central banks say about the future to how they say everything else.

Chapter 15

Statements, Minutes and Press Conferences

A central bank speaks through four channels. The statement, released with the decision, is the committee's collective voice, negotiated word by word. The minutes, released three weeks later at the Fed and six days later at the BCB, explain the decision and record the disagreements. The press conference puts one person in front of the market for an hour of unscripted questions. And the speeches of individual members, between meetings, add or subtract emphasis. Chapter 14 measured the content of words about the future path of rates. This chapter measures the words themselves: how much statements change from meeting to meeting, how markets respond to the change, what the press conference adds to the statement, and how the BCB's tone can be read from its own texts.

15.1 Why words move prices

15.1.1 Cheap talk and its limits

A central bank's statement costs nothing to issue. The theory of Crawford and Sobel (1982) says that costless messages from a sender whose interests differ from the receiver's can convey only coarse information: if the receiver knows the sender wants to shade the message, the message can carry only a few distinct meanings. The bias in monetary policy is the inflation bias of chapter 4: a central bank that would like the public to expect lower inflation than it intends to deliver cannot persuade it by saying so. Two features of central bank communication overcome the limit. The first is reputation: a bank that is held to account for its statements over many meetings has an incentive to make them accurate, and the public learns that they are. The second is commitment: statements that describe the

bank's reaction function, as the conditional guidance of chapter 14 does, are checkable against later decisions, which makes them costly to falsify.

15.1.2 Public signals

Chapter 5 showed, following Morris and Shin (2002), that a precise public signal can lower welfare when agents also care about coordinating with each other, because they overweight it relative to their private information. The result gives one theoretical reason to keep statements spare, and it is model-dependent: Svensson (2006) showed that for plausible parameters public information raises welfare in the same model. If every word is read as a signal and the market coordinates on it, a word inserted for a minor reason can move prices by more than the information it carries justifies. The opposite risk is that a statement says too little to anchor expectations, which is the case for the press conference: it gives the committee a channel to explain without changing the formal text.

15.2 The FOMC statement

15.2.1 Length and change

From January 2011 to September 2026 the FOMC released 128 statements on policy decisions, scheduled and unscheduled. Their length followed the committee's tools (Figure 15.1). It rose from about 480 words in 2011 to 820 in 2014, when the statement described the asset purchase programme, the thresholds for the first increase and the reinvestment policy; fell to about 350 words by 2018; rose again to 460 in 2021 with the guidance of the 2020 framework; and fell below 300 words from 2022, as the committee returned to setting a single rate. In June 2026 the FOMC adopted a new and much shorter format, of about 130 words, beginning with the vote and the decision and dropping most of the standing language on the outlook and the strategy.

The share of words changed from one statement to the next, measured as one minus the similarity of the two word sequences, averages 18%. It is close to zero in many statements, 1% in January 2020 and March 2024, and near 90% only when the committee rewrote the text: the emergency statements of March 2020 and the new format of June 2026. Between those extremes the redline is the market's object of study. Analysts compare each statement with the last one word by word, because the committee changes words deliberately and rarely.

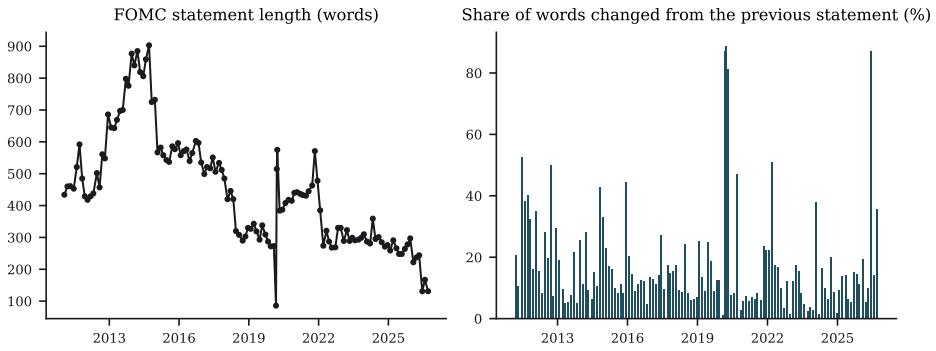


Figure 15.1. FOMC policy statements, January 2011 to September 2026: length in words (excluding the vote and the implementation note), and the share of words changed from the previous statement, measured as one minus the word-sequence similarity ratio. Source: Federal Reserve Board; author's calculations.

15.2.2 Does the redline move markets?

Result: the redline and the long end

On 105 FOMC announcements from 2011 to 2023 with high-frequency data, the absolute change in the ten-year Treasury yield in the thirty-minute window rises by 3.7 basis points (standard error 1.7) when the share of words changed rises from zero to 100%, controlling for the surprise in the current funds rate; the effect on the two-year yield is smaller and imprecise. Announcements in the top quarter of the redline distribution, with 41% of words changed on average, moved the two-year yield by 4.0 basis points; those in the bottom quarter, with 5% changed, by 2.7.

The effect is on the ten-year yield because a change in words is mostly a change in guidance and in the outlook, which the analysis of chapter 14 showed reaches the long end, while a surprise in the current rate moves the short end. The size of the redline is a weak measure of content: a single changed word, "patient" in January 2019 or "some additional policy firming" in 2023, can carry more than a rewritten paragraph describing the economy. A count of changes is the starting point of a reading, not a substitute for it.

15.3 The press conference

The Fed began holding press conferences in April 2011, four times a year, at the meetings that also published the Summary of Economic Projections. From January 2019 the Chair has held one after every meeting. The schedule from 2012 to 2018 provides a natural comparison: meetings with and without a press conference, under the same Chair, statement format and committee.

Result: press-conference and SEP meetings generate larger moves

From 2012 to 2018 the absolute daily change in the two-year Treasury yield averaged 4.6 basis points on the 28 FOMC days with a press conference and 2.2 on the 28 days without one, against 1.9 on all trading days. From 2019 to 2025, with a press conference at every meeting, it averaged 6.2 basis points on FOMC days against 4.0 on all days.

The comparison overstates the effect of the press conference alone, since the press conference days also carried the committee's projections, including the dots. But it shows that the meetings at which the committee explained itself at length were the ones that moved the market, and that the Fed concentrated its new information on them; in 2012 to 2018 almost every change of policy, the first increase of December 2015 included, came at a meeting with a press conference.

15.3.1 The ECB: the conference is the news

The ECB has held a press conference after every monetary policy meeting since 1999, with the decision released shortly before it. The event database of chapter 14 measures market moves separately in the two windows, which gives a direct measure of the share of a meeting's news carried by the press conference.

Result: words dominated, until the decisions became the news

From 1999 to 2013 three-quarters or more of the variance of the change in the two-year OIS rate over a meeting occurred in the press conference, and the share reached 89% in the crisis years (Table 15.1). From 2020 to 2025 it fell to 43%. The decisions of the 2022 to 2023 tightening, several of which surprised the market in size, and the abandonment of

Table 15.1. ECB monetary policy meetings: share of the variance of the change in the two-year OIS rate and the German ten-year yield over the meeting that occurs in the press conference window, as opposed to the press release window. Source: Euro Area Monetary Policy event study Database (Altavilla et al., 2019); author's calculations.

Period	Meetings	Two-year OIS	German ten-year
1999 to 2007	143	0.76	0.86
2008 to 2013	73	0.89	0.81
2014 to 2019	52	0.66	0.84
2020 to 2025	47	0.43	0.54

guidance for meeting-by-meeting decisions moved the news from the words to the decision.

The shift coincides with changes in instruments, in the rate regime and in macro volatility, so the variance shares describe where the news arrived rather than measuring the effect of communication alone.

The ECB's experience gives a general result. Guidance shifts the news to the press conference, because the press conference is where the future is discussed. When a central bank stops guiding, the decision itself carries the information, and the press conference becomes an explanation of what has already been priced.

15.4 The Copom: measuring tone from the text

15.4.1 The BCB's communication

The Copom does not hold a press conference after its meetings. It communicates through a statement released after the meeting, minutes released six days later on Tuesday morning, and the quarterly Monetary Policy Report, which the Governor presents to the press. The statement changed form in October 2016: before it, statements were two sentences, a decision, a bias and a vote; since then they run to several paragraphs on the outlook, the balance of risks and the Committee's guidance. The minutes have been detailed throughout.

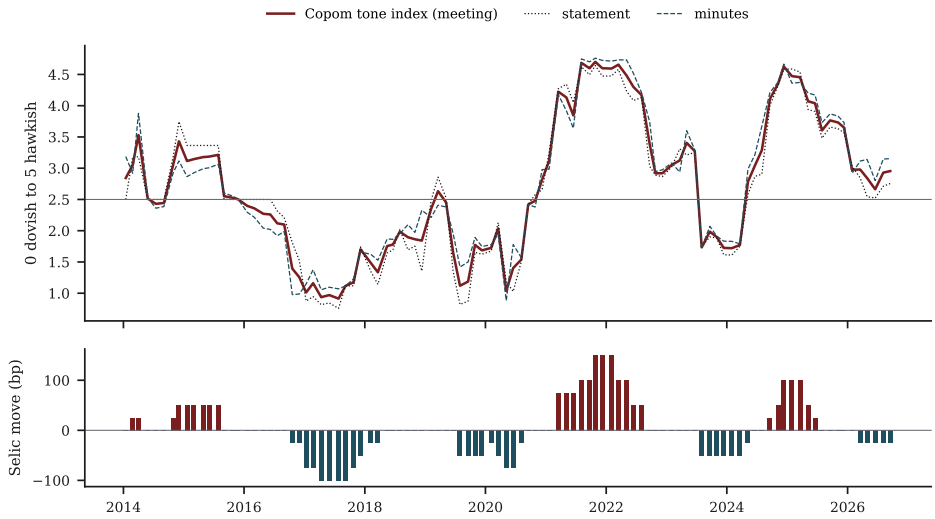


Figure 15.2. Copom tone index, meeting level, with the statement and minutes scores, and the Selic decision at each meeting, January 2014 to September 2026. Sources: BCB statements and minutes; author's research workspace.

15.4.2 An index of tone

The Copom tone index built in this book's research workspace scores every statement and set of minutes from January 2014 to September 2026, 102 meetings, on a scale from 0 (most dovish) to 5 (most hawkish), from the BCB's own texts alone. It has four components, scored separately and weighted: the policy action (25%), the guidance, read from an ordered ladder of the Committee's recurring formulas (30%), the balance of risks, from the stated asymmetry or the count of enumerated upside against downside risks (15%), and the balance of hawkish and dovish vocabulary, after removing the decision sentence, the vote and the symmetric list of risks (30%). The vocabulary score is standardised within each type of document and mapped through

$$s = 2.5 + 2.5 \tanh\!\left(\frac{z}{1.5}\right), \tag{15.1}$$

which keeps the ordering of extreme meetings without letting them saturate. The meeting index is the average of the statement and minutes scores. The parsed decisions are checked against the official Selic series (SGS 432) and reconcile in every meeting checked.

Result: the tone of the Copom, 2014 to 2026

The index was most hawkish, at 4.7, in October 2021, when the Copom raised the Selic by 150 basis points, and most dovish, at 0.9, in July 2017, in the middle of the 2016 to 2018 easing (Figure 15.2). In the last cycle it rose from 3.0 in June 2024 to 4.1 at the first increase of September 2024 and 4.6 in December 2024, when the Committee guided two further increases of 100 basis points; it fell below 3.0 in January 2026, two months before the first cut, and stood at 2.95 at the meeting of September 2026.

15.4.3 What the tone adds

A tone index is useful if it says something the decision does not. Two tests answer that.

The next decision. Regressing the Selic change at the next meeting on the current change and the current tone, over 101 meetings, the tone coefficient is 5.6 basis points per point of the index (standard error 4.4), and the fit rises only from 84% to 85%. The index includes the current action, which takes about 30% of the effective weight once a missing component is dropped, so part of what it measures is the decision itself. A text-only index, which drops the action and reweights guidance, risks and vocabulary, correlates 0.96 with the full one and gives 4.1 basis points per point (standard error 3.1) with the same fit. The Copom's moves come in sequences of equal steps, so the current move already predicts most of the next one; on average the tone, with or without the action, adds little. The minutes score and the statement score carry about the same information.

Turning points. Where the tone matters is in the meetings before a change of direction. On the text-only index, which cannot move with the decision, the tone fell by 0.9 points between December 2025 and January 2026 with the Selic unchanged at 15%, and the first cut came in March; it rose by 0.4 points between June and July 2024, two holds, before the first increase in September. In the two turning points highlighted here the text turned before the rate, and a reader who watched only the decisions would have seen the turn a meeting later. Two episodes are not a general result, and the regression above says the average signal is weak.

15.4.4 Redlines and the market in Brazil

The redline can be measured for the Copom as for the FOMC. The 102 statements from January 2014 to September 2026 averaged 122 words before October 2016 and 772 since; the narrative statements change by 24% of their words from one meeting to the next on average, more than the FOMC's 18%. The largest changes came at the start of the 2016 to 2018 easing (59% in November 2016, 57% in February 2017), at the end of the 2021 to 2022 tightening (48% in September 2022, when the Committee stopped raising rates at 13.75% and replaced its guidance on the cycle with a commitment to vigilance), and in May 2020, at the depth of the pandemic.

Tesouro Direto's morning rates give a rough measure of the market's response. From October 2016 to September 2026 the three-year prefixed rate moved by 12.5 basis points on average, in absolute value, between the morning before the statement and the morning after it, against 10.1 on the morning of the minutes and 9.7 on all days. Regressing the move after a statement on the redline, the absolute surprise against Focus and the absolute change in the tone index, the redline has the largest coefficient, 20 basis points for a fully rewritten statement, but with a standard error of 12; the data, one retail quote a day, are too coarse to separate the channels with confidence. Descriptively, the statement moves the curve more than the minutes; the minutes, whose content is mostly anticipated by then, move it little more than an ordinary day.

15.4.5 Limits of a level index

The index scores what a document says, not what changed from the previous one. When the Committee deletes a paragraph that justified its path, the index sees only the absence of the words, which understates the signal. It is relative to Copom language from 2014 to 2026, not to an absolute scale, and the lexicon on which 30% of the weight rests is hand-built. The tone index is therefore a complement to the redline, not a replacement: the redline says what changed, the index says in which direction the document as a whole leans.

15.5 Communication as policy

Three conclusions follow for central bank practice.

Change words rarely and deliberately. The market reads every change. The FOMC's practice of leaving most of the statement unchanged

from meeting to meeting, which leaves 18% of words changed on average, makes each change informative; the cost is that the statement becomes formulaic and the press conference has to carry the explanation.

Put the news where the explanation is. The Fed's 2012 to 2018 practice of moving only at meetings with press conferences, and the ECB's press conference as the main source of news before 2020, both reflect the same principle: a decision that changes expectations needs an hour of explanation, not two paragraphs.

A short statement transfers weight to other channels. The FOMC's new format of June 2026, at about 130 words, says less in the statement and so places more weight on the press conference, the minutes and the projections. The BCB, without a press conference after its meetings, relies on the statement and the minutes, and the statement carries most of the market's reaction; in the two recent turning points the tone of its texts moved a meeting before its decisions.

Part V

Policy Now

Chapter 16

Optimal Policy in the New Regime

The theory of Part I was built for a world with a stable neutral rate well above zero, a Phillips curve whose slope did not depend on the state of the economy, and a policy rate that could always be cut. The evidence of Parts II to IV says that none of the three holds. The neutral rate fell after 2008 and is uncertain within a range of more than a point in every economy (chapter 7); at its euro area level of about zero, the lower bound is a frequent constraint. The Phillips curve is kinked, flat with slack and steep when the labour market is tight (chapter 9). And the policy rate is one of three instruments, with the balance sheet and words (chapters 13 to 15).

This chapter puts the three facts back into the model of Part I and asks what optimal policy looks like with them. It solves the core model at the lower bound under five strategies, measures how often each economy should expect to be there, derives what an error in r^* costs under different rules, collects the implications of the kinked Phillips curve, and closes with what the Fed, the ECB and the BCB would do differently if they followed the results.

16.1 The lower bound

16.1.1 The problem

At the lower bound the nominal rate cannot fall below $\underline{i}$, which in deviations from steady state is minus the nominal neutral rate, $-(r^* + \pi^*)$. A shock that pushes the natural real rate far enough below r^* makes the rate the bank would like to set infeasible. In the IS curve,

$$x_t = \mathbb{E}_t x_{t+1} - \frac{1}{\sigma}\big(\underline{i} - \mathbb{E}_t \pi_{t+1} - r_t^n\big), \tag{16.1}$$

the real rate is then too high, the gap falls, inflation falls with it, and lower expected inflation raises the real rate further. The bank cannot cut; it can only promise. Every remedy for the lower bound works through $\mathbb{E}_t x_{t+1}$ and $\mathbb{E}_t \pi_{t+1}$: a credible promise of higher inflation and output after the shock lowers the real rate during it. This is Eggertsson and Woodford (2003)'s result, and it is the commitment result of chapter 3 in its sharpest form.

16.1.2 Five strategies

The experiment is a fall in the natural real rate of four points, annualised, for eight quarters, known when it begins: a recession of the size of 2008 to 2009 in real-rate terms. The model is the core model of chapter 1 with the welfare loss of chapter 2. Five strategies are compared:

- a Taylor rule, $i_t = \max\{\underline{i},\, 1.5\pi_t + 0.125 x_t\}$;
- optimal discretion, which sets the rate to keep $\kappa\pi_t + \lambda_x x_t = 0$ when it can and is at the bound otherwise;
- average inflation targeting, the targeting criterion $\sum_{s=t-7}^{t} \pi_s + (\lambda_x/\kappa)\, x_t = 0$, a makeup over two years;
- price-level targeting, $p_t + (\lambda_x/\kappa)\, x_t = 0$ with p_t the price level relative to its target path, which is the timeless commitment rule of chapter 3, equation (3.5);
- optimal commitment, the path of the rate that minimises the loss subject to the bound.

Each rule-based strategy is solved exactly: the bound binds from the start for some number of quarters L, after which the rule holds, and L is found as the shortest duration consistent with the rule's prescription staying above the bound afterwards. The commitment path is found by numerical minimisation over the whole path of the rate subject to the bound.

The slope of the Phillips curve matters enormously at the bound, because the deflationary spiral works through it. With the textbook calibration, $\kappa = 0.17$, a Taylor rule facing this shock with $r^* = 1.2\%$ produces annualised deflation of 25% and an output gap of -13%; with the slope closer to those estimated in the central bank models of chapter 11, $\kappa = 0.03$, the same rule produces deflation of 2.7% and a gap of -4.9%. The second is the plausible case and the one reported in detail.

Table 16.1. The lower bound in the core model ($\kappa = 0.03$). Natural real rate 4 points below r^* for 8 quarters; inflation target 2%. Loss: $1{,}000 \times \sum_{t=0}^{59} \beta^t(\pi_t^2 + \lambda_x x_t^2)$, the discounted welfare loss of chapter 2 over the simulation horizon, with quarterly inflation and the gap in percent; trough of the output gap; trough and peak of annualised inflation; quarters at the bound. Strategies: Taylor rule, optimal discretion, average inflation targeting, price-level targeting, optimal commitment. Author's calculations.

	$r^* = 1.2\%$					$r^* = 0$				
Strategy	Loss	x min	π min	π max	Bound	Loss	x min	π min	π max	Bound
Taylor rule	1312	−4.9	−2.7	0.0	4	1742	−6.4	−3.2	0.0	5
Discretion	171	−2.1	−1.0	0.0	8	1066	−5.4	−2.5	0.0	8
Average inflation	25	−1.2	−0.4	0.1	8	235	−3.4	−1.2	0.1	10
Price level	26	−1.2	−0.4	0.1	8	162	−3.0	−0.9	0.3	8
Commitment	14	−0.9	−0.1	0.1	8	96	−2.3	−0.4	0.4	10

Result: makeup strategies do most of what commitment does

With $r^* = 2.5\%$ the shock is absorbed without reaching the bound under every strategy except the Taylor rule. With $r^* = 1.2\%$, the median US estimate of chapter 7, optimal discretion has a loss twelve times that of commitment; average inflation targeting and price-level targeting have losses within a factor of two of commitment, and cut the trough of the output gap from −2.1% to −1.2% and the trough of inflation from −1.0 to −0.4 points (Table 16.1, Figure 16.1). With $r^* = 0$, the euro area's median, every loss is larger and the ordering is the same; commitment keeps the rate at the bound two quarters after the shock has ended.

The mechanism is visible in the figure. Under discretion the bank raises the rate to neutral the moment the shock ends, so the public expects nothing better after it, and the real rate during the shock stays high. Under the makeup strategies the bank promises to make up the shortfall of inflation, which means keeping rates below neutral after the shock and letting inflation run slightly above target; the promise lowers expected real rates during the shock and cushions it. The Taylor rule does worst of all because it responds only to the inflation and output it observes, not to the natural rate; it lifts off early, in the middle of the shock, because inflation is recovering, and the early tightening deepens the recession.

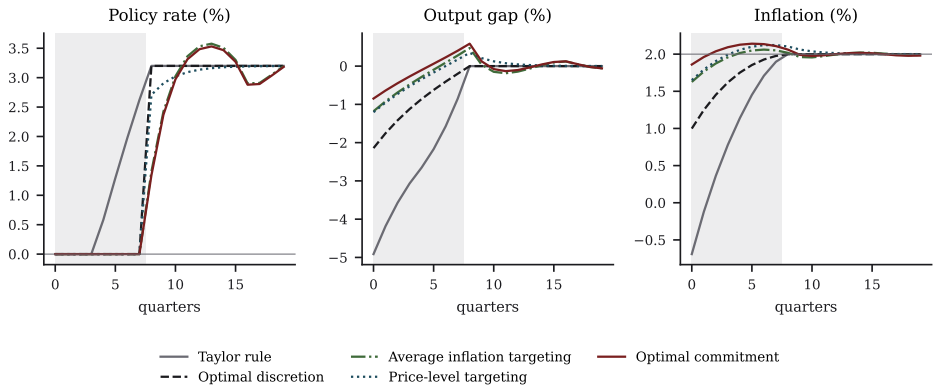


Figure 16.1. Policy rate, output gap and inflation under five strategies after a fall of 4 points in the natural real rate for 8 quarters (shaded), $r^* = 1.2\%$, inflation target 2%, $\kappa = 0.03$. Author's calculations.

16.1.3 How often?

How much weight the lower bound deserves in a strategy depends on how often it binds. The chapter 3 model with a Taylor rule produces a standard deviation of the policy rate of 3.0 points under the shocks used there. Treating the rate as normally distributed around the nominal neutral rate, the probability of a prescription below zero is $\Phi\big(-(r^* + \pi^*)/3.0\big)$.

Result: the bound is a euro area problem

With the median neutral rates of chapter 7, the probability that the policy rate prescription falls below zero is 15% in the United States (r^* 1.2%, target 2%), 25% in the euro area (r^* 0.1%) and 0.4% in Brazil (r^* 5.0%, target 3%). Each point of neutral rate matters: with a 2% target, the probability is 5% at $r^* = 3\%$, 9% at 2%, 16% at 1%, 25% at zero and 31% at -0.5%.

These are crude numbers, from a model that is not estimated, but the ranking does not depend on the details. For the ECB, a strategy that ignores the bound ignores a quarter of the distribution of outcomes; for the BCB the bound is not a practical concern, and the strategic questions lie elsewhere.

16.2 A moving neutral rate

16.2.1 What an error in r^* costs

Every rule written around r^* inherits its uncertainty. Take a Taylor rule with a perceived neutral rate $\hat{r}^*$,

$$i_t = \hat{r}^* + \pi_t + (\phi_\pi - 1)(\pi_t - \pi^*) + \phi_x x_t. \tag{16.2}$$

In a steady state the Fisher relation requires $i = r^* + \bar{\pi}$ and the gap is zero. Substituting,

$$r^* + \bar{\pi} = \hat{r}^* + \bar{\pi} + (\phi_\pi - 1)(\bar{\pi} - \pi^*) \quad \Longrightarrow \quad \bar{\pi} - \pi^* = \frac{r^* - \hat{r}^*}{\phi_\pi - 1}. \tag{16.3}$$

Result: the cost of a wrong neutral rate

Under a Taylor rule with $\phi_\pi = 1.5$, overestimating the neutral rate by one point produces an inflation rate two points below target, permanently. With $\phi_\pi = 3$, the error shrinks to half a point. A rule that responds strongly to inflation is also a rule that is insensitive to errors in r^*.

The dispersion of the neutral-rate measures in chapter 7, 0.7% to 1.8% for the United States and 5% to 7.3% for Brazil, is large enough to make this practical. Three responses are available. A strong response to inflation, as in the optimal simple rules of chapter 3, makes the error small. A first-difference rule, $i_t = i_{t-1} + a(\pi_t - \pi^*) + b\,\Delta x_t$, contains no r^* at all and cannot be biased by it, at the cost of responding slowly to large shocks (Orphanides and Williams, 2002). And an estimate of r^* built from many measures, as the BCB does and chapter 7 does for all three economies, reduces the error itself.

16.2.2 The neutral rate and the bound together

The two problems compound. A lower neutral rate raises the probability of the bound; an uncertain neutral rate raises the probability of a policy that is too tight in a downturn, because the bank does not know how far below neutral it needs to go. Both argue for the same thing: a strategy with a makeup element, which commits to compensate for shortfalls, and a rule that does not lean on a precise estimate of r^*. The Fed's 2020 framework adopted

the first, and its 2025 revision withdrew it. On the reading of this book, the inflation of 2021 to 2023 had shown the cost of applying it asymmetrically.

16.3 A kinked Phillips curve

Chapter 9 estimated a US Phillips curve with a slope of about 0.1 when the ratio of vacancies to unemployed is below one and 1.6 above it, and derived two results for policy. First, with a kink, certainty equivalence fails: a bank uncertain about where the economy sits relative to the kink should aim lower, at a ratio of about 0.8 rather than 1.0 when its uncertainty is 0.3 units. Second, the cost of disinflation depends on where it starts: steep and cheap above the kink, flat and expensive below it.

Combined with the lower bound, the kink gives policy an asymmetry in both directions. Overheating is costly because the curve is steep; deep recessions are costly because the bound limits the response. The optimal strategy leans against both tails: it tightens early when the labour market is tight, which the precautionary result recommends, and it commits to make up shortfalls when the economy falls into a slump, which the lower-bound result recommends. Within the model built here, a symmetric makeup strategy, which promises to compensate for both undershoots and overshoots, satisfies both. The 2020 framework failed the first test because its makeup applied only to undershoots and its employment objective only to shortfalls.

16.4 What each central bank would do differently

16.4.1 The Fed

Respond to the kink. The Fed's estimated rule responds to the output gap with a coefficient of about 0.9 on real-time data (chapter 10). With a kinked curve the response should be larger above full employment than below it. In October 2026 the US labour market sits at the kink, with the ratio of vacancies to unemployed at 1.0 and an oil shock under way; the precautionary result says that uncertainty about where the economy is relative to the kink argues for leaning tight.

Keep a symmetric makeup in reserve. With a 15% probability of the bound, the Fed will face it again. The lower-bound results argue for a makeup strategy at the bound, announced in advance and applied symmetrically, rather than the discretion that the 2025 revision restored.

Read the curve with a term-premium model. In 2026 the market's expected path of rates is about half a point above the committee's projections once the term premium is removed (chapter 12). The gap is a fact about expectations the committee can address with guidance.

16.4.2 The ECB

Plan for the bound. With a neutral rate near zero by most of the model and staff measures used here, the ECB faces the bound in a quarter of outcomes. The evidence here supports keeping the full set of balance-sheet tools, which the stock estimate of chapter 13 shows to be powerful, and a makeup commitment at the bound.

Respond to inflation more strongly. The ECB's estimated rule, written around the moving neutral rate, has a response to expected inflation of about 1.5 and a large response to slack (chapter 10). The robustness argument above, and the optimal simple rules of chapter 3, both favour a stronger response to inflation, which also protects against error in a neutral rate whose market measure disagrees with its model and staff measures by more than a point.

16.4.3 The BCB

Keep the rule on expectations, with care. The BCB's published rule, replicated in chapter 10, responds to Focus expectations with a coefficient of 2. Chapter 3 showed that the same coefficients applied to a model-consistent forecast four quarters ahead would perform badly; they work in practice because survey expectations move far more than a model forecast does. The design depends on that property of the survey, and should be revisited if expectations become better anchored and move less.

Take the neutral rate seriously, and its premium. The Brazilian measures of r^* run from 4.3% to 7.3%, and the gap between the BCB's reference value and the market's is consistent with a sizable fiscal and risk premium component, alongside the expected real policy path, term and liquidity premia (chapters 7 and 12). The cost of an error in r^* derived above applies directly: with a coefficient of 2 on expectations, a one-point underestimate of the neutral rate produces an inflation bias of a point. A neutral rate estimated from the full set of measures, published with its range, is the protection.

Watch the fiscal boundary. The interaction of fiscal and monetary policy of chapter 4 is the BCB's binding constraint, as the lower bound is

the ECB's. The real rate above 6% that the market prices is consistent with the same constraint, seen in the yield curve.

16.5 Conclusion

Woodford (2003a)'s *Interest and Prices* showed that a central bank that sets an interest rate can control inflation without controlling money, provided it manages expectations. Two decades of practice have confirmed the central claim and complicated every assumption around it. Expectations are managed with words as much as with rates, and guidance meant to stimulate through the expected path needs a credible commitment, while its information content has to be recognised. The neutral rate around which the policy rate is set moves, and cannot be estimated with confidence by any one method. The Phillips curve that links the policy rate to inflation is flat most of the time and steep when it matters. The lower bound that was a curiosity in 2003 bound the Fed for seven years and the ECB for eight. And the balance sheet, which the theory said should not matter, moves the term premium by an amount comparable in its effect on output to a year of tighter policy rates.

The policy that follows from all of this is the one the theory always recommended, made resistant to error: respond strongly to expected inflation, commit to make up for misses, measure the neutral rate as a set rather than a number, and say clearly what will be done, conditional on what can be checked. The Fed, the ECB and the BCB each do some of this. None does all of it.

Bibliography

Tobias Adrian, Richard K. Crump, and Emanuel Moench. Pricing the term structure with linear regressions. *Journal of Financial Economics*, 110 (1):110–138, 2013.

Tobias Adrian, Nina Boyarchenko, and Domenico Giannone. Vulnerable growth. *American Economic Review*, 109(4):1263–1289, 2019.

Andrea Ajello, Michele Cavallo, Giovanni Favara, William B. Peterman, John W. Schindler, and Nitish R. Sinha. A new index to measure U.S. financial conditions. FEDS Notes, Board of Governors of the Federal Reserve System, June 2023.

George A. Akerlof, William T. Dickens, and George L. Perry. The macroeconomics of low inflation. *Brookings Papers on Economic Activity*, 1996 (1):1–76, 1996.

Alberto Alesina and Lawrence H. Summers. Central bank independence and macroeconomic performance: Some comparative evidence. *Journal of Money, Credit and Banking*, 25(2):151–162, 1993.

Carlo Altavilla, Luca Brugnolini, Refet S. Gürkaynak, Roberto Motto, and Giuseppe Ragusa. Measuring Euro Area monetary policy. *Journal of Monetary Economics*, 108:162–179, 2019.

Guido Ascari. Staggered prices and trend inflation: Some nuisances. *Review of Economic Dynamics*, 7(3):642–667, 2004.

David Backus and John Driffill. Inflation and reputation. *American Economic Review*, 75(3):530–538, 1985.

Banco Central do Brasil. Atualização dos modelos semiestruturais de pequeno porte. Box, Relatório de Inflação, June 2024, pp. 97–106, 2024a.

Banco Central do Brasil. Atualização de medidas de taxa de juros real neutra no Brasil. Box, Relatório de Inflação, June 2024, pp. 94–96, 2024b.

Banco Central do Brasil. Atualização das medidas de hiato do produto no Brasil. Box, Relatório de Política Monetária, June 2025, pp. 75–78, 2025.

Robert J. Barro and David B. Gordon. A positive theory of monetary policy in a natural rate model. *Journal of Political Economy*, 91(4):589–610, 1983a.

Robert J. Barro and David B. Gordon. Rules, discretion and reputation in a model of monetary policy. *Journal of Monetary Economics*, 12(1): 101–121, 1983b.

Michael D. Bauer and Eric T. Swanson. A reassessment of monetary policy surprises and high-frequency identification. *NBER Macroeconomics Annual*, 37:87–155, 2023.

Pierpaolo Benigno and Gauti B. Eggertsson. It's baaack: The surge in inflation in the 2020s and the return of the non-linear Phillips curve. Working Paper 31197, National Bureau of Economic Research, 2023.

Ben Bernanke and Olivier Blanchard. What caused the U.S. pandemic-era inflation? Working Paper 31417, National Bureau of Economic Research, 2023.

Ben S. Bernanke and Michael Woodford. Inflation forecasts and monetary policy. *Journal of Money, Credit and Banking*, 29(4):653–684, 1997.

Mark Bils and Peter J. Klenow. Some evidence on the importance of sticky prices. *Journal of Political Economy*, 112(5):947–985, 2004.

Olivier Blanchard. Fiscal dominance and inflation targeting: Lessons from Brazil. Working Paper 10389, National Bureau of Economic Research, 2004.

Olivier Blanchard and Jordi Galí. Real wage rigidities and the New Keynesian model. *Journal of Money, Credit and Banking*, 39(s1):35–65, 2007.

Olivier Jean Blanchard and Charles M. Kahn. The solution of linear difference models under rational expectations. *Econometrica*, 48(5):1305–1311, 1980.

Alan S. Blinder and John Morgan. Are two heads better than one? Monetary policy by committee. *Journal of Money, Credit and Banking*, 37(5):789–811, 2005.

Alina Bobasu, Matteo Ciccarelli, and Alessandro Notarpietro. Monetary policy transmission: A reference guide through ESCB models and empirical benchmarks. Occasional Paper 377, European Central Bank, 2025.

Pedro Bordalo, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer. Overreaction in macroeconomic expectations. *American Economic Review*, 110 (9):2748–2782, 2020.

William C. Brainard. Uncertainty and the effectiveness of policy. *American Economic Review*, 57(2):411–425, 1967.

Scott Brave and R. Andrew Butters. Monitoring financial stability: A financial conditions index approach. *Economic Perspectives, Federal Reserve Bank of Chicago*, 35(1):22–43, 2011.

Flint Brayton, Thomas Laubach, and David Reifschneider. The FRB/US model: A tool for macroeconomic policy analysis. FEDS Notes, Board of Governors of the Federal Reserve System, April 2014.

James Bullard and Kaushik Mitra. Learning about monetary policy rules. *Journal of Monetary Economics*, 49(6):1105–1129, 2002.

Jeffrey R. Campbell, Charles L. Evans, Jonas D. M. Fisher, and Alejandro Justiniano. Macroeconomic effects of Federal Reserve forward guidance. *Brookings Papers on Economic Activity*, 2012(1):1–80, 2012.

John Y. Campbell and Robert J. Shiller. Yield spreads and interest rate movements: A bird's eye view. *Review of Economic Studies*, 58(3):495–514, 1991.

Lawrence J. Christiano, Martin Eichenbaum, and Charles L. Evans. Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, 113(1):1–45, 2005.

Richard Clarida, Jordi Galí, and Mark Gertler. The science of monetary policy: A New Keynesian perspective. *Journal of Economic Literature*, 37(4):1661–1707, 1999.

Richard Clarida, Jordi Galí, and Mark Gertler. Monetary policy rules and macroeconomic stability: Evidence and some theory. *Quarterly Journal of Economics*, 115(1):147–180, 2000.

John H. Cochrane. Determinacy and identification with Taylor rules. *Journal of Political Economy*, 119(3):565–615, 2011.

Günter Coenen, Peter Karadi, Sebastian Schmidt, and Anders Warne. The New Area-Wide Model II: An extended version of the ECB's microfounded model for forecasting and policy analysis with a financial sector. Working Paper 2200, European Central Bank, 2018.

Timothy Cogley and Argia M. Sbordone. Trend inflation, indexation, and inflation persistence in the New Keynesian Phillips Curve. *American Economic Review*, 98(5):2101–2126, 2008.

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

Olivier Coibion, Yuriy Gorodnichenko, and Saten Kumar. How do firms form their expectations? New survey evidence. *American Economic Review*, 108(9):2671–2713, 2018.

Vincent P. Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, 50(6):1431–1451, 1982.

Marcos R. de Castro, Solange N. Gouvea, André Minella, Rafael C. dos Santos, and Nelson F. Souza-Sobrinho. SAMBA: Stochastic analytical model with a Bayesian approach. *Brazilian Review of Econometrics*, 35 (2):103–170, 2015.

Marco Del Negro, Marc P. Giannoni, and Christina Patterson. The forward guidance puzzle. *Journal of Political Economy Macroeconomics*, 1(1): 43–79, 2023.

Francis X. Diebold and Canlin Li. Forecasting the term structure of government bond yields. *Journal of Econometrics*, 130(2):337–364, 2006.

Avinash K. Dixit and Joseph E. Stiglitz. Monopolistic competition and optimum product diversity. *American Economic Review*, 67(3):297–308, 1977.

Gauti B. Eggertsson and Michael Woodford. The zero bound on interest rates and optimal monetary policy. *Brookings Papers on Economic Activity*, 2003(1):139–233, 2003.

Arturo Estrella and Frederic S. Mishkin. Predicting U.S. recessions: Financial variables as leading indicators. *Review of Economics and Statistics*, 80(1):45–61, 1998.

Jordi Galí. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*. Princeton University Press, Princeton, 2nd edition, 2015.

Jordi Galí and Tommaso Monacelli. Monetary policy and exchange rate volatility in a small open economy. *Review of Economic Studies*, 72(3): 707–734, 2005.

Mariana García-Schmidt and Michael Woodford. Are low interest rates deflationary? A paradox of perfect-foresight analysis. *American Economic Review*, 109(1):86–120, 2019.

Mark Gertler and Peter Karadi. A model of unconventional monetary policy. *Journal of Monetary Economics*, 58(1):17–34, 2011.

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

Mikhail Golosov and Robert E. Lucas, Jr. Menu costs and Phillips curves. *Journal of Political Economy*, 115(2):171–199, 2007.

Robin Greenwood and Dimitri Vayanos. Bond supply and excess bond returns. *Review of Financial Studies*, 27(3):663–713, 2014.

Refet S. Gürkaynak, Brian Sack, and Eric T. Swanson. Do actions speak louder than words? The response of asset prices to monetary policy actions and statements. *International Journal of Central Banking*, 1(1): 55–93, 2005.

Tomáš Havránek. Measuring intertemporal substitution: The importance of method choices and selective reporting. *Journal of the European Economic Association*, 13(6):1180–1204, 2015.

Dániel Holló, Manfred Kremer, and Marco Lo Duca. CISS: A composite indicator of systemic stress in the financial system. Working Paper 1426, European Central Bank, 2012.

Marek Jarociński and Peter Karadi. Deconstructing monetary policy surprises: The role of information shocks. *American Economic Journal: Macroeconomics*, 12(2):1–43, 2020.

Òscar Jordà. Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182, 2005.

Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

Don H. Kim and Jonathan H. Wright. An arbitrage-free three-factor term structure model and the recent behavior of long-term yields and distant-horizon forward rates. Finance and Economics Discussion Series 2005-33, Board of Governors of the Federal Reserve System, 2005.

Finn E. Kydland and Edward C. Prescott. Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy*, 85(3): 473–491, 1977.

Eric M. Leeper. Equilibria under "active" and "passive" monetary and fiscal policies. *Journal of Monetary Economics*, 27(1):129–147, 1991.

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

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

Alisdair McKay, Emi Nakamura, and Jón Steinsson. The power of forward guidance revisited. *American Economic Review*, 106(10):3133–3158, 2016.

Stephen Morris and Hyun Song Shin. Social value of public information. *American Economic Review*, 92(5):1521–1534, 2002.

Emi Nakamura and Jón Steinsson. Five facts about prices: A reevaluation of menu cost models. *Quarterly Journal of Economics*, 123(4):1415–1464, 2008.

Emi Nakamura and Jón Steinsson. High-frequency identification of monetary non-neutrality: The information effect. *Quarterly Journal of Economics*, 133(3):1283–1330, 2018.

Charles R. Nelson and Andrew F. Siegel. Parsimonious modeling of yield curves. *Journal of Business*, 60(4):473–489, 1987.

A. Robert Nobay and David A. Peel. Optimal discretionary monetary policy in a model of asymmetric central bank preferences. *Economic Journal*, 113(489):657–665, 2003.

Athanasios Orphanides. Monetary policy rules based on real-time data. *American Economic Review*, 91(4):964–985, 2001.

Athanasios Orphanides and John C. Williams. Robust monetary policy rules with unknown natural rates. *Brookings Papers on Economic Activity*, 2002 (2):63–145, 2002.

Kenneth Rogoff. The optimal degree of commitment to an intermediate monetary target. *Quarterly Journal of Economics*, 100(4):1169–1189, 1985.

Thomas J. Sargent and Neil Wallace. "Rational" expectations, the optimal monetary instrument, and the optimal money supply rule. *Journal of Political Economy*, 83(2):241–254, 1975.

Thomas J. Sargent and Neil Wallace. Some unpleasant monetarist arithmetic. *Federal Reserve Bank of Minneapolis Quarterly Review*, 5(3):1–17, 1981.

Christopher A. Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690, 2003.

Frank Smets and Rafael Wouters. Shocks and frictions in US business cycles: A Bayesian DSGE approach. *American Economic Review*, 97(3):586–606, 2007.

Lars E. O. Svensson. Estimating and interpreting forward interest rates: Sweden 1992–1994. Working Paper 4871, National Bureau of Economic Research, 1994.

Lars E. O. Svensson. Inflation forecast targeting: Implementing and monitoring inflation targets. *European Economic Review*, 41(6):1111–1146, 1997.

Lars E. O. Svensson. Inflation targeting as a monetary policy rule. *Journal of Monetary Economics*, 43(3):607–654, 1999.

Lars E. O. Svensson. Social value of public information: Comment: Morris and Shin (2002) is actually pro-transparency, not con. *American Economic Review*, 96(1):448–452, 2006.

Eric T. Swanson. Measuring the effects of Federal Reserve forward guidance and asset purchases on financial markets. *Journal of Monetary Economics*, 118:32–53, 2021.

John B. Taylor. Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy*, 39:195–214, 1993.

John B. Taylor. Housing and monetary policy. In *Housing, Housing Finance, and Monetary Policy*, pages 463–476. Federal Reserve Bank of Kansas City, 2007.

Oldřich Vašíček. An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5(2):177–188, 1977.

Dimitri Vayanos and Jean-Luc Vila. A preferred-habitat model of the term structure of interest rates. *Econometrica*, 89(1):77–112, 2021.

Neil Wallace. A Modigliani–Miller theorem for open-market operations. *American Economic Review*, 71(3):267–274, 1981.

Carl E. Walsh. Optimal contracts for central bankers. *American Economic Review*, 85(1):150–167, 1995.

Volker Wieland, Tobias Cwik, Gernot J. Müller, Sebastian Schmidt, and Maik Wolters. A new comparative approach to macroeconomic modeling and policy analysis. *Journal of Economic Behavior and Organization*, 83 (3):523–541, 2012.

Michael Woodford. Price-level determinacy without control of a monetary aggregate. *Carnegie-Rochester Conference Series on Public Policy*, 43: 1–46, 1995.

Michael Woodford. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, Princeton, 2003a.

Michael Woodford. Imperfect common knowledge and the effects of monetary policy. In Philippe Aghion, Roman Frydman, Joseph Stiglitz, and Michael Woodford, editors, *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, pages 25–58. Princeton University Press, Princeton, 2003b.

Michael Woodford. Optimal interest-rate smoothing. *Review of Economic Studies*, 70(4):861–886, 2003c.

Jing Cynthia Wu and Fan Dora Xia. Measuring the macroeconomic impact of monetary policy at the zero lower bound. *Journal of Money, Credit and Banking*, 48(2–3):253–291, 2016.

Arnold Zellner. Bayesian estimation and prediction using asymmetric loss functions. *Journal of the American Statistical Association*, 81(394):446–451, 1986.