INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 2 Practical PID Control

CHAPTER 11

Inflation Targeting as Sampled, Quantized Control of a Lagged Economy

Pages 272–296

Edited by John D. Hedengren

How to cite this chapter. "Inflation Targeting as Sampled, Quantized Control of a Lagged Economy," in *Practical PID Control*, John D. Hedengren, Ed., vol. 2 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 11, pp. 272–296.

The page numbers in this file are those of the printed volume.

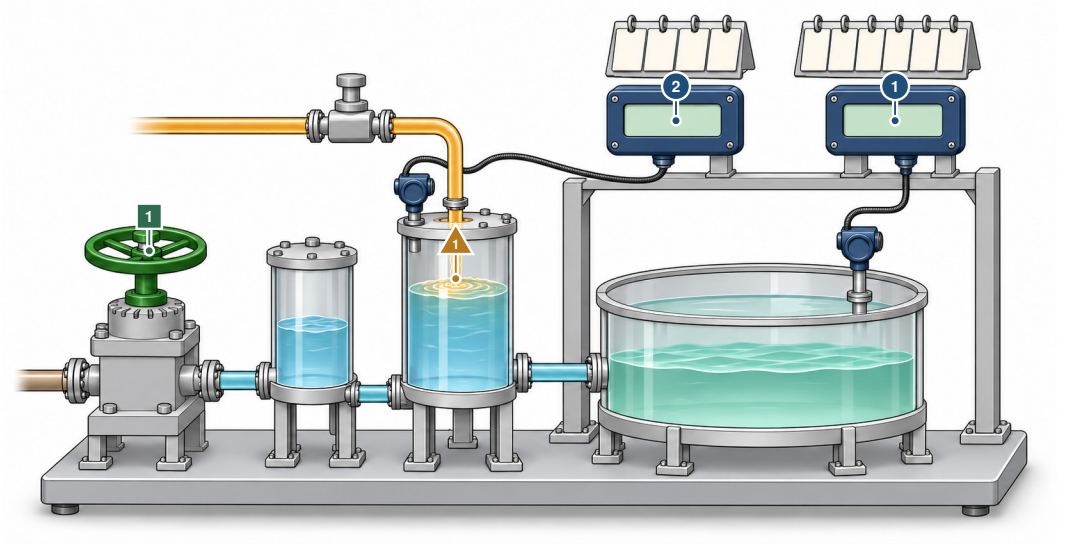
The app for this chapter
https://apmonitor.com/apps/inflation-targeting/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v02-practical-pid/#inflation-targeting

Copyright © 2026 Advanced Process Solutions, LLC. All rights reserved. Produced by Advanced Process Solutions, LLC.

CHAPTER 11

Inflation Targeting as Sampled, Quantized Control of a Lagged Economy



The controller moves 1 the policy rate, moved at scheduled meetings in quarter-point steps

The controller reads 1 the price index, published weeks after each month

2 the reading of output, published after each quarter and revised

What disturbs it 1 a shock to spending or to prices

Notice the two narrow vessels between the wheel and the wide vessel: a move of the rate passes through both before prices answer, and each reading reaches its display only on its own calendar. Illustration, not to scale.

IN PLAIN TERMS Steering Prices With One Interest Rate

The main idea. A central bank tries to keep prices rising slowly and steadily by moving one interest rate. A higher rate slows borrowing and spending, and slower spending later slows the rise in prices.

The challenge. The rate reaches spending after about a year and prices after about two, and the figures that show prices come out weeks late. The bank meets only eight times a year and moves the rate in small fixed steps.

Why it matters. A bank that keeps pushing on the newest figure can turn a small boom into a deep slump.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the history of the first inflation targets (Section 11.2), the fact on the target range (Section 11.5), the try-it segments of cases 2 and 4 and the fact on the floor (Section 11.6), and the key points
Study	45 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the results (Section 11.6), Technology in Practice (Section 11.7), Limitations (Section 11.8) and the open challenge

The control question A central bank moves one rate eight times a year in quarter-point steps, on statistics published late, and the economy answers in output a year later and in prices two years later. Which way of deciding holds inflation near its target through those lags, and what do the calendar, the quantum and the lower bound cost it?

Learning objectives

- Derive the Taylor-principle boundary on the inflation response from the model's steady state, with and without an anchored public.
- Determine the action of the inflation PID from the sign of the process gain in the convention of Chapter 1.
- Explain why the transmission is three lags in series and not a dead time, and estimate the effective dead time a step test returns.
- Predict what the decision calendar, the rate quantum and the floor do to a demand boom and a financial crisis.
- Compare a proportional rule, the PID and forecast targeting on an unannounced fall of the neutral rate and on an oil shock.

Prerequisites Sections 1.4.2 (proportional offset), 1.4.3 (controller action), 1.4.4, 1.6.2, 1.4.9 and 1.4.11 (sample time and a quantized output). Chapter 9 and Chapter 10 decide on the same kind of calendar. Section 1.6.3 is needed only for forecast targeting.

The app https://apmonitor.com/apps/inflation-targeting/

See also Volume 5, Chapter 3 and Volume 5, Chapter 4 for control through a late measurement, and Volume 5, Chapter 9 for a slow stock refilled by a sampled, capped decision. Volume 9, Chapter 4 for an estimator that carries an unmeasured bias.

11.1 Introduction

A policy rate is changed at a meeting, and the economy it is meant to steer answers over the following two years. In the simulator of this chapter, a rise of 100 bp held for one quarter moves the output gap most 1.04 yr after the rise and inflation most after 2.40 yr, and the bank sees both only through a consumer price index (CPI) published two weeks after each month and a gross domestic product (GDP) estimate published a month after each quarter and revised for a year. A chair who raises the rate 25 bp at every meeting for as long as the inflation print exceeds 2.25 % takes it to 8.875 %, a year and a half after inflation peaked, and drives the output gap to −4.93 % and inflation to −1.82 %. The same demand boom under the Taylor rule peaks the rate at 5.625 % and the gap stays between +1.61 % and −0.90 %.

Economists have described this problem in the language of feedback control for seventy years. Stabilization policy was written in 1954 as proportional, integral and derivative correction of a production gap acting through lags [1], and in 1972 short-term stabilization was posed as a discrete-time linear quadratic tracking problem on a 28-state quarterly econometric model [2]. The rule proposed in 1993 sets the rate from the inflation gap and the output gap with fixed coefficients [3], which is a proportional controller on two measured errors. Inflation-forecast targeting chooses the rate so that the bank's own forecast of inflation meets the target over a policy horizon [4], which is receding-horizon optimal control. A control paper is the right frame because the domain already uses it, with one standing objection. Under rational expectations a discretionary optimum does not maximize the social objective, and there is then no way control theory can be applied to economic planning [5]. The simulator takes the other branch, with backward-looking expectations that can be integrated forward in time, and states that choice as a limitation.

The structure is general. A slow plant reached through a chain of lags, a measurement that is a late and revised average of the past, an actuator that moves only at scheduled instants and in fixed steps, a hard lower limit, and an operating point that drifts without being announced appear together in process plants whose quality is known only when a laboratory result returns, whose analyzers

report on a cycle, and whose operators change setpoints at shift meetings. The column behind a slow analyzer of Volume 5, Chapter 3 and the digester graded at its blowline in Volume 5, Chapter 4 are two of them. Each of those features has a textbook answer on its own, such as dead-time compensation, sampled-data design, anti-windup and integral action [6]. What a transfer function from the policy rate to inflation does not show is which vintage of which statistic the controller held when it decided. The belief this plant overturns is that a controller responding harder or more often to the newest measurement must control better. In this economy a stronger inflation response raises the integrated loss of a demand boom from 4.24 to 4.85, and deciding every week instead of eight times a year lowers it only from 3.64 to 3.46.

However, the policy-rule literature evaluates rules on estimated models, where the information set and the calendar are assumptions of the analysis [7, 8], so a student rarely has a plant on which the chair, the rule, a proportional-integral-derivative (PID) controller and an optimizer act on the same published data through the same quantized calendar. This chapter works with such a plant and measures on it what the lags, the data and the calendar do to each policy. The structure, the Taylor rule, forecast targeting and the real-time data problem are all taken from the literature cited below. The simulator joins a continuous-time economy to a statistics agency that keeps every vintage and to the framework's shared decision calendar, with each equation, constant, delay and limit written out. Nine cases and a challenge each isolate one mechanism, and each was measured on the running engine.

The chapter is organized as follows. Section 11.2 traces the problem's history, Section 11.3 writes out the economy and Section 11.4 its dynamic character. Section 11.5 gives the bank's four ways of deciding, and Section 11.6 reports the measured cases and the challenge. Section 11.7 sets the real statistics, rate and committee beside the model, Section 11.8 states the limitations, Section 11.9 summarizes the results, and the exercises follow in Section 11.10.

11.2 Stabilization Policy as Feedback, 1954 to Forecast Targeting

The control description of macroeconomic policy is older than the policy rules in use today. In 1954 a model of a closed economy was given a government that corrects the deviation of production from its desired level in proportion to the deviation, to its time integral and to its rate of change, with the correction reaching demand through lags [1]. Chapter 10 rebuilds that economy and its PID policy as a plant of its own. Four years later the same author fitted the relation between unemployment and the rate of change of money wage rates in the United Kingdom from 1861 to 1957 [9], the curve that later carried his name and that the Phillips-curve view of this application draws. How long monetary policy takes to act was by 1961 a question argued in its own right [10].

The 1970s brought optimal control to the problem and then the objection to it. A quadratic tracking formulation was solved analytically and applied to a 28-state quarterly econometric model, and the experiments were presented as a tool both for policy planning and for analyzing the dynamics of econometric models [2]. Within five years two arguments limited what such a design could claim. A fitted model's coefficients need not survive a change in the policy that produced the data [11], and discretionary policy, the best decision given the current situation, does not maximize an agreed social objective when the economy's agents are rational [5]. Both arguments are about the plant changing when the controller does, and both return in the limitations of this chapter.

Inflation targeting turned the objective into a published number, and the history box after this paragraph follows the first regimes and what they shared [7]. Forecast targeting made the bank's forecast the intermediate target [4]. A rate rule with fixed responses to inflation and the output gap was proposed in 1993 as a description of practice [3], and the efficient rules of a simple backward-looking model, those minimizing a weighted sum of output variance and inflation variance, turned out to be Taylor rules and to be equivalent to inflation targets pursued at different speeds [8]. The policy-rule literature of the same decade set the rules inside a common framework of optimal policy, credibility and imperfect information [12].

Table 11.1 Signals of the Economy and the Bank by Role, With Their Defaults

Symbol	Role	Default	Range
π_{pub}	CV, latest monthly CPI print	SP π^* 2 %	0–6 %
π, y	objectives, true inflation and gap, hidden from the bank	2 %, 0	—
y_{pub}	measured, latest quarterly gap print	0	vintages 1–3
i	MV, applied policy rate	2.875 %	0.125–20 %
i^w	request, from the chair's slider or a controller	2.875 %	0.125–19.875
d_D, d_S	DVs, demand and cost-push shocks	0, 0	events, ±6, ±3
r^*, g^*	DVs, true neutral rate and potential growth	0.875 %, 2	−2–4, 0–5
r_1, r_2	states, transmission stages	0.875 %	—
π^a, λ	states, adaptive expectations, anchored share	2 %, 0.7	0–1 for λ
m, P	states, potential misperception, log potential	0, 0	—
I_y, I_{yp}, I_π	states, integrals for period averages	0	—
u, g	outputs, unemployment and real GDP growth	4 %, 2	—
$\hat{x}$	bank's estimate of seven model states	truth at start	—

Shocks: d_D in percent of potential, d_S and g^* in percent a year.

HISTORY Inflation Targets as Published Numbers, 1998 to 2012

In early 1998 explicit inflation targeting was pursued in New Zealand, Canada, the United Kingdom, Sweden and Australia, and arguably also in Finland and Spain. The regimes shared three marks: a publicly announced numerical target, given as a range, a point or a point with a tolerance interval, a framework for decisions that compared an inflation forecast with the target, so that the forecast served as an intermediate target, and a higher than average degree of transparency and accountability [7]. In the United States the Federal Open Market Committee adopted, effective January 24, 2012, a statement judging that inflation of 2 percent, measured by the annual change in the price index for personal consumption expenditures, is most consistent over the longer run with its statutory mandate [13].

Two problems kept the rule from being a formula. A rule evaluated with the information available in real time prescribes rates that differ considerably from the same rule evaluated with revised data [14], and ex post revisions of the estimated output gap are of the same order of magnitude as the gap itself and highly persistent, mostly because the trend of output is unreliable at the end of the sample [15]. The second problem was the floor. By 2003 the overnight rate in Japan had been within 50 basis points of zero since October 1995, so that the usual remedy of a lower short-term rate was unavailable [16]. A dynamic model with moderate nominal rigidities was by then able to account for the observed inertia of inflation and persistence of output after a monetary shock [17], and the lags of the present model give the same qualitative picture, with output moving well before inflation.

11.3 An Economy of Stocks and Lags

The simulated economy is a continuous-time teaching analogue of the quarterly backward-looking model of [7], in years, drawn in Figure 11.1 as the stocks and flows it is. The policy rate i is the one in force, held between decisions by the bank of Section 11.5. Inflation π, the output gap y and unemployment u are in percent, rates of change in percent a year, and time in years. Table 11.1 classifies the signals.

11.3.1 Transmission and Prices

The nominal rate less expected inflation is the real rate, and it reaches spending through two first-order stages, as shown in Equation (11.1). These are lag states, not market rates.

$$\tau_1 \dot{r}_1 = i - \pi^e - r_1, \qquad \tau_2 \dot{r}_2 = r_1 - r_2 \tag{11.1}$$

where $\tau_1 = 0.25\,\text{yr}$ stands for credit and asset prices and $\tau_2 = 0.5\,\text{yr}$ for spending decisions. The output gap answers the second stage against the true neutral real rate, and the gap moves inflation, as

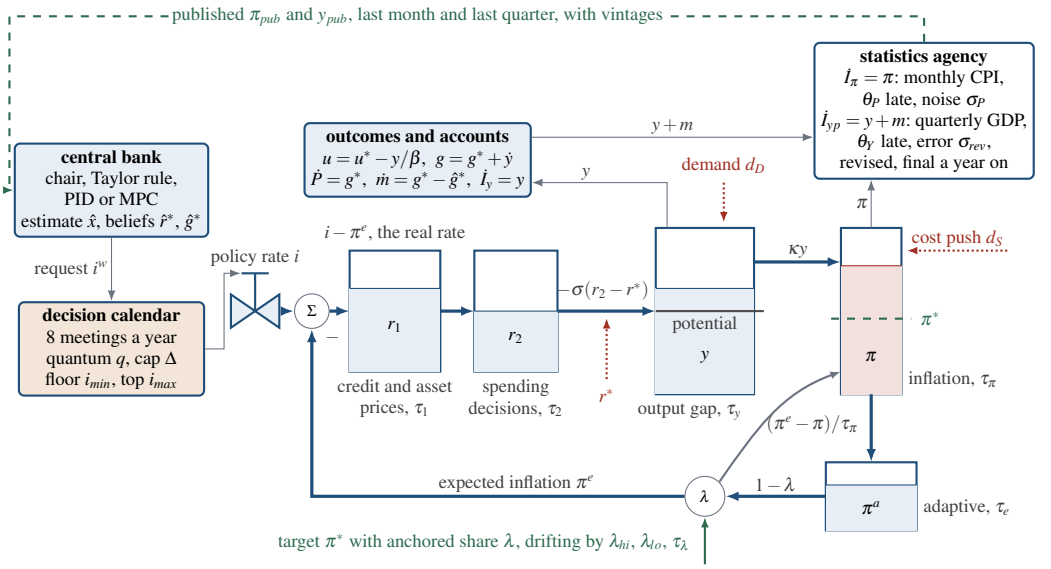


Figure 11.1 The economy as stocks and lags, the statistics agency, and the bank with its calendar. The real rate $i - \pi^e$ fills the credit stage r_1, which fills the spending stage r_2, and the second stage drains the output gap y against the true neutral rate r^*. The gap raises inflation π at the rate κy while inflation relaxes toward expected inflation π^e, a mix with share λ of the target π^* and the adaptive stock π^a. Densely dotted red arrows are disturbances, thick blue arrows are flows between stocks and thin gray arrows are signals. The agency integrates π and the gap the bank can measure, $y + m$, and returns monthly and quarterly averages late along the dashed green path, where the bank's request i^w passes the calendar's quantum q, cap Δ and bounds before it reaches the valve.

shown in Equation (11.2).

$$\begin{aligned} \tau_y \dot{y} &= -y - \sigma(r_2 - r^*) + d_D \\ \dot{\pi} &= \kappa y + (\pi^e - \pi)/\tau_\pi + d_S \end{aligned} \tag{11.2}$$

where $\tau_y = 0.75\,\text{yr}$, $\sigma = 1$ point of gap per point of real rate, $r^* = 0.875\,\%$, $\kappa = 0.5$ percent a year per point of gap and $\tau_\pi = 2\,\text{yr}$. The demand disturbance d_D is in percent of potential and the cost push d_S in percent a year. At a steady gap inflation settles at $\pi = \pi^e + \kappa\tau_\pi y$, a short-run Phillips curve whose slope against unemployment is $-\kappa\tau_\pi\beta = -2$ with the Okun coefficient below.

11.3.2 Expectations and Credibility

Expected inflation is a share λ anchored on the announced target and an adaptive remainder, as shown in Equation (11.3).

$$\begin{gathered} \pi^e = \lambda\pi^* + (1-\lambda)\pi^a, \qquad \tau_e \dot{\pi}^a = \pi - \pi^a \\ \tau_\lambda \dot{\lambda} = \begin{cases} \lambda_{hi} - \lambda & |\pi - \pi^*| \le 1 \\ 0 & 1 < |\pi - \pi^*| < 3 \\ \lambda_{lo} - \lambda & |\pi - \pi^*| \ge 3 \end{cases} \end{gathered} \tag{11.3}$$

where $\tau_e = 1.5\,\text{yr}$, and the drift of the second line acts only when the learner switches it on, with $\lambda_{hi} = 0.9$, $\lambda_{lo} = 0.3$ and $\tau_\lambda = 6\,\text{yr}$. The share is clipped to $[0,1]$ inside the derivative and after every step. With $\lambda = 0$ a gap held at one point raises inflation at the constant rate $\kappa\tau_\pi/(\tau_\pi + \tau_e) = 0.286$ points a year, the accelerationist curve of the 1970s, which the application's check confirms to 5×10^{-9} after fifteen years. With $\lambda = 1$ inflation returns to target after an oil shock with the rate held at neutral, which the check confirms within 0.01 points after fifteen years.

11.3.3 Outcomes and Accounts

Unemployment, growth and the accounts that the statistics need are given by Equation (11.4).

$$\begin{aligned} u &= u^* - y/\beta, \qquad g = g^* + \dot{y} \\ \dot{P} &= g^*, \qquad \dot{m} = g^* - \hat{g}^* \\ \dot{I}_y &= y, \qquad \dot{I}_{yp} = y + m, \qquad \dot{I}_\pi = \pi \end{aligned} \tag{11.4}$$

where $u^* = 4\,\%$, $\beta = 2$ points of gap per point of unemployment, $g^* = 2$ percent a year, P is the log of potential output in percent with the GDP index $100\exp[(P+y)/100]$, and m is how far the bank's potential has drifted from the truth when its belief $\hat{g}^*$ differs from g^*, so that the gap it can measure is $y+m$, the Orphanides problem [15]. Both identities hold in the exported file, Okun's law to 5×10^{-9} on every row and the growth identity to 0.0015 points against a central difference of the GDP index, away from the instants where a shock starts or ends. The integrals give exact period averages, interpolated to the period boundary inside the step. Quarterly growth, annualized, is four times the sum of the quarter's change in log potential and in average gap, two negative quarters in a row raise RECESSION, and calendar-year growth is the change of log potential over the year plus the change of the year's average gap.

11.3.4 What the Agency Publishes

The bank never sees the economy. Each month's average inflation, from I_π, is published as the CPI print of Equation (11.5) two weeks after the month ends and is never revised.

$$\pi_{\text{pub},j} = 12\int_{(j-1)/12}^{j/12} \pi\,\mathrm{d}t + \sigma_P\,\nu_j \tag{11.5}$$

where $\sigma_P = 0.2$ points and ν_j is a standard normal deviate. Each quarter's average of $y+m$, from I_{yp}, receives a first print one month after the quarter, modeled on the United States advance estimate, with an error $\varepsilon \sim N(0, 0.5^2)$, a second print two months later that keeps 0.4ε, and a final print a year after the first with no error. The published gap is the newest quarter's latest vintage. A misleading first print can be scheduled, two points too low and revised later. Every deviate comes from the engine's seeded generator in a fixed order, one per CPI print and one per first GDP print, so a reset replays every release. In 15 years the agency publishes 179 CPI prints and 173 GDP prints. At $t = 0$ the bank holds a print of last month and last quarter at the starting point, as though the economy had been there before.

11.3.5 Shocks, Start and Integration

Three event shocks have declared shapes in the time t' since they started. The oil shock adds $d_S = 3$ percent a year for a year and then decays with a 0.5 yr time constant, and the challenge fires it at 6. The demand boom adds $d_D = 2$ for two years and decays with the same time constant. The financial crisis starts at $d_D = -6$ and recovers in a straight line over three years. Two sliders add a persistent d_D of up to ±6 and d_S of up to ±3, a de-anchoring event sets $\lambda = 0.4$, and the true r^* and g^* are sliders hidden from the bank. A case starts in equilibrium at its target unless it declares otherwise, with $r_1 = r_2 = r^* - y_0/\sigma + \bar{d}_D/\sigma$ for a starting gap y_0 and a persistent demand $\bar{d}_D$, $\pi = \pi^a = \pi^*$, and the first rate on the grid at $r_2 + \pi^e = 2.875\,\%$. Classical fourth-order Runge-Kutta integrates the eleven states on a fixed step of a week, $\Delta t = 1/52$ year, the controller executes every week before the step, and the historian logs every week. The equilibrium holds with zero derivatives and no drift at all over ten noise-free years, and halving the step leaves the challenge's verdicts unchanged and its peak unemployment within 10^{-4} points. The default seed is 20261004, and the engine of Section 1.6.1 runs the economy at up to two simulated years a second. Table 11.2 gathers the constants.

Table 11.2 Model Constants, Marked M When the Show Model Dialog Sets Them, A When an Advanced Control Does and C When Only the Source Does

Symbol	Meaning	Value	Range	Class	Symbol	Meaning	Value	Range	Class
τ_1	credit stage lag	0.25 yr	0.05–1	M	τ_2	spending stage lag	0.5 yr	0.05–1.5	M
τ_y	output-gap lag	0.75 yr	0.1–2	M	σ	interest sensitivity of spending	1.0	0.2–2	M
κ	Phillips slope, %/yr per point of gap	0.5	0.1–1.5	M	τ_π	pull of inflation toward π^e	2 yr	0.5–5	M
τ_e	adaptive expectations lag	1.5 yr	0.25–5	M	λ_{hi}, λ_{lo}	anchored share near and far from target	0.9, 0.3	0.5–1, 0–0.6	M
τ_λ	credibility drift lag	6 yr	1–20	M	u^*, β	natural unemployment, Okun coefficient	4 %, 2	2–8, 1–3.5	M
θ_Y	first GDP print after the quarter	1/12 yr	0–26 wk	M	σ_{rev}	first-print error	0.5	0–2	M
θ_P	CPI print after the month	2 wk	0–13 wk	M	σ_P	CPI noise	0.2	0–1.5	M
r^*	true neutral real rate	0.875 %	−2–4	A	g^*	true potential growth	2 %/yr	0–5	A
—	second print, share of error kept	1/6 yr, 0.4	—	C	—	final print after the first	1 yr	—	C
—	oil shock, hold, decay	3 %/yr, 1, 0.5 yr	6 in challenge	C	—	boom, hold, decay	2, 2, 0.5 yr	—	C
—	crisis and linear recovery	−6, 3 yr	—	C	—	drift bands on $\lvert\pi - \pi^*\rvert$	1, 3 points	—	C
—	STAGFLATION	$\pi > \pi^* + 2$, $y < -1$	—	C	—	OVERHEATING, DE-ANCHORED	$y > 2$, $\lambda < 0.5$	—	C
Δt	integration step, control period	1/52 yr	—	C	—	log rate, seed	52 per yr, 20261004	—	C

11.3.6 Assumptions

Expectations look backward. The anchored share and the adaptive stock of Equation (11.3) make the model causal, so an engine can step it forward. A public that learns the bank's rule would change these relationships, which is the critique of [11], and the model therefore shows what a rule does to this economy, not what it would do to one whose public anticipates it.

Teaching values. Every coefficient, lag, shock size and threshold is a declared teaching value. The lags were chosen so that a pulse lands its troughs in design windows (Section 11.4), and estimated coefficients come from models such as [7] and [17], not from this one.

A local, linear economy. The investment-saving (IS) curve of Equation (11.2), the Phillips curve and Okun's law are linear near the operating point, so a gap near −7 %, which a chair moving 50 bp a meeting reaches in case 1 (−6.93 %), is outside the range they describe.

One instrument, one closed economy. There is no exchange rate, no financial sector beyond the two stages, no fiscal policy, no forward guidance and no asset purchases.

Constructed statistics. The CPI and GDP prints are period averages of the model's own states with a simplified revision schedule, not reconstructions of any agency's procedure. The anchored share is read by the bank as an exact survey.

Deterministic shocks. Only the releases are random. The crisis is a demand shock, not a credit crunch with dynamics of its own.

11.4 How the Economy Answers a Move

11.4.1 Three Lags in Series, Not a Dead Time

Nothing in the economy forgets a move, and nothing waits. A point held from $t = 0$ on the noise-free reference has moved the first stage 0.63 points after a quarter but the second only 0.15, and the gap is −0.09 % after half a year, −0.35 % after a year and −0.76 % after two, while inflation has moved only −0.27 points by then. The response starts at once and arrives late, the signature of three lags in series rather than a transport delay. A first-order-plus-dead-time fit of such a response returns an effective dead time made of lags (Section 1.4.4). Released after a quarter, the same rise moves the gap most at 1.04 yr by −0.13 % and inflation most at 2.40 yr by −0.07 points, quarters 4.2 and 9.6. Held a year, the troughs come at quarters 6.1 and 11.3, by −0.50 % and −0.26 points, as shown in Figure 11.2(a). Each rise is set before Run, which puts it in force from $t = 0$, so every time is measured from the rise.

🕮 **Pause.** Try Exercise 11.2 before reading on.

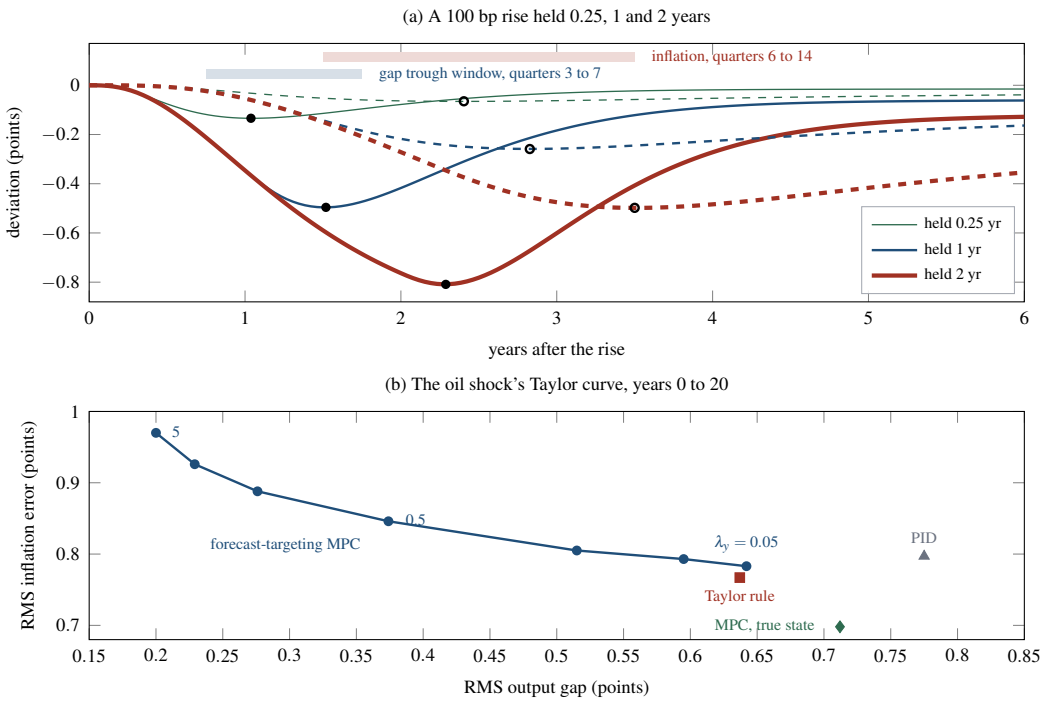


Figure 11.2 Two measurements from the application's engine. (a) The output gap (solid) and inflation (dashed) after a 100 bp rise held 0.25 years (thin lines), 1 year (medium) and 2 years (thick) on the noise-free reference with weekly decisions and no quantum, in force from $t = 0$, with each gap trough marked by a filled dot, each inflation trough by an open circle and the design windows shaded above. (b) The oil shock of case 5 over 20 years, root-mean-square (RMS) inflation error against RMS output gap for the model predictive controller (MPC) at output weights from 0.05 to 5 (line with dots), the Taylor rule (square), the PID (triangle) and the MPC given the true state (diamond).

The application's plan placed the troughs of a rise held two years inside quarters 3 to 7 for the gap and 6 to 14 for inflation. The model cannot do that, and the check was restated rather than the model bent. The gap keeps falling for as long as the real rate is held up, so a rise held two years has its gap trough after the release, at quarter 9.2, with inflation's at quarter 14.0. The windows are therefore asserted on rises held a quarter and a year, which land inside both, and the two-year rise is checked for a gap trough between quarters 8 and 10. The windows themselves are design targets of the application and not estimates.

11.4.2 The Taylor Principle With an Anchor

The Taylor rule of Section 11.5 is stable only if the real rate rises when inflation does. Setting the derivatives of Equations (11.1) to (11.3) to zero with the rule acting on true values and $\hat{r}^* = r^*$ gives $\pi^a = \pi$, $y = \lambda x/(\kappa\tau_\pi)$ and $y = -\sigma[(a_\pi - 1 + \lambda)x + a_y y]$ for an inflation deviation $x = \pi - \pi^*$, so a deviation can persist without a shock only when the bracket of Equation (11.6) vanishes.

$$\left[a_\pi - 1 + \lambda\left(1 + \frac{1 + \sigma a_y}{\kappa\sigma\tau_\pi}\right)\right]x = 0 \tag{11.6}$$

where a_π is the rule's total inflation response and a_y its gap response. The bracket is proportional to the determinant of the closed loop's state matrix, so a real root crosses zero where it vanishes. For the linearized loop without publication delays or quantum the slowest root is $+0.0148$ per year at $a_\pi = 0.9$ with no anchor and -0.0148 at 1.1, so a deviation feeds itself below the boundary,

Table 11.3 Bank Settings, Defaults and Ranges. A Marks an Advanced Control

Symbol	Setting	Default	Range
Policy and calendar			
π^*	inflation target	2 %	0–6
—	controller in AUTO	Taylor rule	3 policies
—	calendar	8 meetings a year	4 calendars
Δ, q	largest move, quantum (A)	0.5, 0.25	0.25–2, 0–1
i_{min}, i_{max}	lower bound (A), top	0.125 %, 20 %	−2–2, fixed
$\hat{r}^*$, $\hat{g}^*$	bank's neutral rate and growth (A)	0.875 %, 2	−2–4, 0–5
Taylor rule			
a_π, a_y	inflation and gap responses	1.5, 0.5	0–3, 0–2
ρ	partial adjustment (A)	0	0–0.9
PID			
K_c, τ_I	gain, integral time	1.5, 8 yr	0.1–5, 0–20
τ_D	derivative time (A)	0	0–2 yr
K_y	gain on the published gap	0.5	0–2
—	floor tracking (A)	on	on, off
Forecast-targeting MPC			
λ_y, ν	weight on the gap, move penalty	0.5, 0.3	0.05–5, 0.05–3
H	horizon (A)	3 yr	1–6
—	information (A)	estimator	estimator, truth
Public expectations			
λ_0	anchored share at the start	0.7	0–1
—	credibility drift	off	off, on

slowly. The boundary is $a_\pi = 1$ with no anchor, the classic Taylor principle, and 0.25 with 30 percent anchored at $a_y = 0.5$, because an anchored public pulls inflation back by itself. Between the two, Equation (11.6) puts it at 0.75 for an anchored share of 0.1 and 0.5 for 0.2. A forty-year sweep on case 3's boom confirms both. With no anchor, $a_\pi = 0.9$ leaves inflation 1.00, 1.06 and 1.43 points over target at years 10, 20 and 40, while 1.1 leaves 0.79, 0.59 and 0.47. With 30 percent anchored, 0.2 diverges, from 1.02 to 1.30 points between years 10 and 40, while 0.3 converges, from 0.90 to 0.69. The plan expected 0.8 to diverge at an anchored share of 0.3. It converges there, to 0.04 points at year 40, so the case runs with no anchor.

11.4.3 A Grid That Holds the Neutral Rate

The rate lives on the grid of target-range midpoints, $0.125 + 0.25n$ percent. The neutral real rate is 0.875 % rather than a round 1 % so that the neutral nominal rate, 2.875 %, lies on that grid and the equilibrium holds exactly. With a 50 bp quantum the nearest rates are 2.625 % and 3.125 %, neither of which holds the economy still, and the Taylor rule settles into a relay cycle whose size the quantum sets, not the gain. From year 20 to 30 of case 6 the rate alternates between them, rising every 3.6 years, with the gap within ±0.15 % and inflation between 1.96 % and 2.04 %. An inflation response of 3 instead of 1.5 leaves the gap within ±0.16 %. Chapter 1 treats the period and the step of a sampled output in Section 1.4.11. The starting rate is rounded to the quantum in force at start-up, so a one-point quantum starts the rate at 3.125 % and leaves a gap cycle from –0.24 % to +0.27 %.

🕮 **Pause.** Try Exercise 11.3 before reading on.

11.5 How the Bank Decides

The controller executes every week. Each execution reads the releases that came out since the last one and corrects the bank's estimate, asks the selected policy for a request i^w, and asks the calendar whether this execution is a decision. The chair and the three automatic policies share that last step, so they differ only in how they form the request. Table 11.3 lists the settings.

11.5.1 The Calendar and the Shape of a Decision

The calendar and the shaping of each decision come from the framework's shared decision-calendar module, `core/cadence.js`, which also moves the contribution of Volume 5, Chapter 9 on a calendar with a quantum and a cap. Eight meetings fall at the fractions 0.075, 0.205, 0.340, 0.455, 0.575, 0.715, 0.840 and 0.955 of each year, spaced as the Federal Open Market Committee's schedules usually are but a teaching calendar, not any year's dates. Monthly and quarterly calendars are regular intervals, and the continuous one makes every weekly execution a decision. An execution is a decision when a calendar instant falls at or before it and after the previous decision, so a meeting between two weekly executions is taken at the first one after it. The eight meetings therefore fall 6 to 8 weeks apart, monthly decisions 4 or 5 weeks apart and quarterly ones every 13 weeks. The module also keeps a log of the decisions, which the meeting strip of the Economy view draws as one block for each quantum step, above its line for a rise and below it for a cut. At a decision the request is shaped as shown in Equation (11.7).

$$\begin{aligned} \tilde{\imath} &= b + q\,\mathrm{round}\big((i^{w} - b)/q\big), \qquad \Delta_q = q\max\big(1, \lfloor \Delta/q \rfloor\big) \\ i_k &= \mathrm{clip}_{[i_{min},\, i_{max}]}\Big(i_{k-1} + \mathrm{sat}_{\Delta_q}\big(\tilde{\imath} - i_{k-1}\big)\Big) \end{aligned} \tag{11.7}$$

where $b = 0.125$ is the grid's base, rounding takes ties away from zero, the bounds are first pulled in to the grid, sat limits a move to $\pm\Delta_q$, and a request that is not finite holds the rate. The default $q = 0.25$, $\Delta = 0.5$, $i_{min} = 0.125\,\%$ and $i_{max} = 20\,\%$ give the familiar 25 bp steps, at most 50 bp a meeting. The rate is then held to the next decision. The application's check finds every decision of four calendars on the grid, inside the cap and the bounds and on its calendar instant, and the floor uncrossed in every case under every policy.

> ✦ DID YOU KNOW? **A Target Range of 0 to 1/4 Percent, 2008**
>
> On December 16, 2008 the Federal Open Market Committee set a target range for the federal funds rate of 0 to 1/4 percent [18]. The model's floor of 0.125 % is the midpoint of that range, and its grid of midpoints keeps ranges a quarter point wide.

11.5.2 The Chair

In MANUAL the request is the chair's slider, from 0.125 % to 19.875 % in 25 bp steps, and the slider is a request, which the next meeting shapes exactly as it shapes a controller's. In AUTO the slider follows the controller's request, not the rate after the quantum, the cap and the floor, so a switch to MANUAL leaves the chair asking for what the controller asked for. A switch to AUTO, or to another controller, holds the rate in force at the first decision after it, and every policy transfers without a jump (Section 1.4.10).

11.5.3 The Taylor Rule

The rule is a proportional controller on the latest prints, as shown in Equation (11.8).

$$\begin{aligned} i_T &= \hat{r}^{*} + \pi^{*} + a_{\pi}(\pi_{\mathrm{pub}} - \pi^{*}) + a_{y} y_{\mathrm{pub}} \\ i^{w} &= \rho\, i + (1 - \rho)\, i_T \end{aligned} \tag{11.8}$$

where $a_{\pi} = 1.5$ is the total response to a point of inflation, Taylor's coefficient of one plus one half [3], $a_y = 0.5$, and the partial adjustment ρ is zero by default. The intercept is the bank's belief $\hat{r}^{*} = 0.875\,\%$, which an unannounced move of r^{*} makes wrong.

11.5.4 The PID

The PID is the framework's velocity-form PID of Section 1.6.2. It runs every week on the CPI print, and the published gap is added outside it, as shown in Equation (11.9).

$$\begin{aligned} e_k &= \pi_{\text{pub},k} - \pi^*, \qquad P_k = K_c e_k \\ D_k &= K_c \tau_D (\pi_{\text{pub},k} - \pi_{\text{pub},k-1})/\Delta t \\ S_k &= S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\ \tilde{u}_k &= \hat{r}^* + \pi^* + S_k + I_{k-1}, \quad u_k = \text{clip}_{[\underline{u},\bar{u}]}(\tilde{u}_k) \\ I_k &= I_{k-1} + \left[K_c e_k/\tau_I + (u_k - \tilde{u}_k)/\tau_I\right]\Delta t \\ i^w &= u_k + K_y y_{\text{pub}} \end{aligned} \tag{11.9}$$

where $K_c = 1.5$, $\tau_I = 8\,\text{yr}$, $\tau_D = 0$ and $K_y = 0.5$. These are Equations (1.23) to (1.26) with the bias $u_b = \hat{r}^* + \pi^*$, $\tilde{u}_k = v_k$ and a tracking time equal to τ_I, so the running sum is never clipped and the back-calculation acts on the integral only.

The error $e_k = \pi_{\text{pub},k} - \pi^*$ is the framework's $a(\text{SP} - \text{PV})$ with $a = -1$. A higher rate lowers inflation, so the process gain is negative, and the controller must raise its output as the PV rises. Inflation above target raises the rate. In the convention of Section 1.4.3 this is a direct-acting controller, and the loop gain aK_cK_p is positive. The app's guide to the PID also calls it direct acting. Named by its push against inflation, the same loop would be called reverse acting, so the error equation, not the name, fixes the sign.

🕮 **Pause.** Try Exercise 11.1 before reading on.

With floor tracking on, $\underline{u} = i_{min} - K_y y_{\text{pub}}$ and $\bar{u} = i_{max} - K_y y_{\text{pub}}$, so that the request is held inside the rate's own bounds and the integral stops at the floor. With it off, $\underline{u} = -70\,\%$, which removes the lower clip while the upper one still back-calculates. The derivative is taken on the measurement with no filter, so a nonzero τ_D differentiates the CPI's noise, and $\tau_I = 0$ removes the integral. The meeting's quantum and cap are not fed back to the integral, which keeps integrating in the weeks between meetings while the rate is held. Whenever the PID is not the policy in force, in MANUAL or with another controller selected, it tracks the rate in force less $K_y y_{\text{pub}}$. With its integral off and $K_c = 1.5$, $K_y = 0.5$, the PID is the Taylor rule, and the application's check finds the two runs of case 4 identical to 10^{-9} at every step.

11.5.5 The Bank's Estimate

The model predictive controller (MPC) reads a state estimate $\hat{x} = [r_1, r_2, y, \pi, \pi^a, b_D, b_S]$ of the bank's own linear model, which is Equations (11.1) to (11.3) with $\hat{r}^*$ in place of r^*, the bias b_D in place of d_D and b_S in place of d_S, $\dot{b}_D = 0$, $\dot{b}_S = -b_S/\tau_b$ with $\tau_b = 1\,\text{yr}$, and the anchored share read each week as an exact survey and frozen over a forecast. The demand bias absorbs whatever the bank has wrong about the neutral rate, so its implied estimate is $\hat{r}^* + b_D/\sigma$. The offset-free controller of Volume 9, Chapter 4 carries an unmeasured fuel bias in its estimator in the same way. The estimate is propagated every week by fourth-order Runge-Kutta with the rate in force, and the bank keeps its own monthly and quarterly averages of $\hat{\pi}$ and $\hat{y}$. Each release corrects it at the middle of the period it describes and carries the correction to now, as shown in Equation (11.10).

$$\begin{aligned} &\hat{x} \leftarrow \hat{x} + e^{A(t - t_m)} L \nu \\ &L = \begin{cases} (0,0,0,0.5,0.25,0,0.4)^\top & \text{CPI print} \\ (0,0,0.8,0,0,1,0)^\top & \text{GDP print} \end{cases} \end{aligned} \tag{11.10}$$

where A is the bank model's state matrix, t_m is the middle of the month or quarter, and the innovation ν is the print less the bank's own average for a CPI print or a first GDP print, and the change from the previous print for a revision. The gains are teaching values, not a Kalman design. A perfect-state setting gives the MPC the true states and the true shocks as its biases, as a labeled comparison.

11.5.6 Forecast Targeting as Model Predictive Control

At every decision, and whenever a setting or a switch has made its last plan stale, the MPC discretizes the bank model exactly for a zero-order hold over one decision interval, $T_s = 1/8$ year on the default calendar and never less than a month, and solves the quadratic program (QP) of Equation (11.11) with the framework's linear MPC, `core/linmpc.js` (Section 1.6.3).

$$\begin{aligned} \min_{u_0,\ldots,u_9} \; & \textstyle\sum_{k=1}^{N} T_s\left[(\hat{\pi}_k - \pi^*)^2 + \lambda_y \hat{y}_k^2\right] \\ & + \nu \textstyle\sum_{j=0}^{9} (u_j - u_{j-1})^2 \\ \text{s.t.} \; & i_{min} \le u_j \le i_{max}, \qquad |u_0 - i| \le \Delta \\ & |u_j - u_{j-1}| \le \Delta\, p\, n_{j-1}, \qquad j = 1,\ldots,9 \end{aligned} \tag{11.11}$$

where $N = H/T_s = 24$ steps, $u_{-1} = i$, the ten blocks hold $n_j = 1,1,1,1,2,2,3,4,6,3$ steps, p is the number of decisions a step spans, one on the default calendar and $52/12$ on the continuous one, $\lambda_y = 0.5$ and $\nu = 0.3$. The program has 10 variables and 40 rows, and a ridge of $10^{-9}(1 + d_{max})$, with d_{max} the largest diagonal entry of the Hessian, keeps it strictly convex. The dual active-set method of [19] solves it, the first move becomes the request and the calendar shapes it, and a failed solve holds the rate and is counted. Over the 161 solves of case 8 it needs 1.18 active-set iterations on average and 12 at most, with none failing. The forecast freezes the anchored share and knows nothing of shocks to come.

With a one-step horizon the program has a closed form. If b_π is the inflation entry of the discrete input matrix and $\hat{\pi}_1$ the one-step forecast with the rate held, the move is

$$\Delta i = \frac{T_s\,|b_\pi|}{\nu + T_s b_\pi^2}\,(\hat{\pi}_1 - \pi^*) \tag{11.12}$$

a velocity-form integral move on the forecast error. The application checks that the shared PID run as a pure integral controller with that gain gives the same move to 10^{-12} and that the QP agrees to the ridge, on a horizon below the slider's one year that only the check can set. Over one meeting the rate barely reaches inflation, $b_\pi = -4.47 \times 10^{-5}$, so with the check's $\nu = 1$ and a forecast of 2.96 % the move is 5.35×10^{-6} points and both policies hold the rate. Forecast targeting earns its lead only over a horizon long enough for the lags to act.

11.5.7 Economy, Belief and Benchmark

Equations (11.1) to (11.5), with the shocks, the drift of the anchored share and the release noise, are the simulated economy. The bank's model is their linear copy with the believed neutral rate, two biases, a frozen anchored share and no foresight, corrected only by the releases. The model both stand for is an estimated quarterly economy such as [7] or a simple backward-looking model such as [8], whose coefficients were estimated or derived and not copied here. No model is identified from data inside the application, and the PID and the Taylor rule use no model at all.

11.6 Measured Runs of the Economy

Table 11.4 lists the cases. Each starts from the defaults of Tables 11.2 and 11.3, selected from the dropdown, with the change named in its row, and runs for its own length at the default seed (Section 1.6.4). A variant is the case with one setting changed before Run, and the figures were produced by the engine alone. Table 11.5 collects the outcomes. Table 11.7 lists what each case requires before the application credits it. A case is credited once its listed actions have been taken and the economy has reached the listed state, and the challenge counts only on its own model and disturbances, so changing either voids an attempt. The integrated loss is $\int[(\pi - \pi^*)^2 + \lambda_y y^2]\,\mathrm{d}t$ on the true states with the MPC's weight, so it is compared only within one case and one run length.

Table 11.4 The Nine Cases, the Challenge and the Reference

#	Change from the reference, run length	Teaching point
1	chair in MANUAL, boom at 0.25 yr, 8 yr	the overshoot the lags make
2	Taylor rule, the same boom, 10 yr	a proportional rule
3	$a_\pi = 0.8$, $\lambda_0 = 0$, boom, 40 yr	the Taylor principle
4	PID, r^* falls a point at 1 yr, 20 yr	integral action
5	MPC, oil shock at 1 yr, 20 yr	an infeasible frontier
6	Taylor rule, boom at 1 yr, noise-free data, 10 yr	cadence and quantum
7	PID with $\tau_I = 2$ yr, crisis at 1 yr, 12 yr	the floor and windup
8	MPC, boom, oil and crisis at 1, 5 and 9 yr, 20 yr	optimizer against PID
9	$\pi = \pi^a = 10\,\%$, $y = 2.4$, $\lambda = 0.3$ drifting, 20 yr	the sacrifice ratio
C	oil at 6 %/yr at 1 and 5 yr, $\lambda_0 = 0.4$ drifting, 15 yr	a 1970s decade
R	noise-free data, no shock, 5 yr	the equilibrium

Table 11.5 Final and Peak Inflation, Lowest Gap, Peak Unemployment, Worst Calendar-Year Growth and Loss Over Each Case's Own Run

Case and configuration	π_{end}	π_{max}	y_{min}	u_{max}	g_{min}	Loss
1, rate held at 2.875 %	2.67	3.24	0.00	4.00	1.19	9.84
1, ±25 bp while $\lvert\pi_{\text{pub}} - 2\rvert > 0.25$	−1.82	3.05	−4.93	6.47	−0.20	63.67
2, Taylor rule	2.06	2.95	−0.90	4.45	0.44	4.24
2, $a_\pi = 2.5$	1.95	2.90	−1.33	4.66	0.23	4.85
2, $a_y = 1$	2.03	2.88	−1.10	4.55	0.39	3.93
2, MPC	2.06	2.85	−1.66	4.83	0.13	5.19
3, $a_\pi = 0.8$	4.77	4.77	−0.23	4.12	0.78	118.95
3, $a_\pi = 1.5$	2.08	3.10	−0.72	4.36	0.55	7.90
4, Taylor rule	1.52	2.00	−0.61	4.31	1.64	4.60
4, PID	1.86	2.00	−0.61	4.31	1.64	1.78
4, PID, $\tau_I = 4$ yr	1.98	2.00	−0.61	4.31	1.64	1.05
4, MPC	2.01	2.15	−0.58	4.29	1.68	0.68
5, MPC, $\lambda_y = 0.05$	1.97	4.75	−2.20	5.10	0.61	12.68
5, MPC, $\lambda_y = 0.5$	1.99	4.77	−1.26	4.63	1.21	15.73
5, MPC, $\lambda_y = 5$	1.99	4.81	−0.49	4.24	1.61	22.82
5, Taylor rule	1.96	4.72	−2.23	5.12	0.65	15.84
6, Taylor, weekly	2.02	2.85	−0.80	4.40	0.83	3.46
6, Taylor, 8 a year	2.02	2.87	−0.86	4.43	0.77	3.64
6, Taylor, quarterly	2.02	2.90	−0.97	4.48	0.65	4.03
7, PID, floor tracking	2.04	2.74	−3.25	5.63	−0.24	17.43
7, floor tracking off	2.08	3.23	−3.25	5.63	−0.24	22.84
7, $i_{min} = -2\,\%$	1.99	3.20	−3.25	5.63	−0.24	22.58
7, $\pi^* = 1\,\%$	1.12	1.52	−3.25	5.63	−0.24	17.34
8, MPC	2.02	4.30	−3.49	5.75	0.43	28.93
8, PID	1.91	4.43	−3.86	5.93	0.70	41.56
8, Taylor rule	1.96	4.58	−3.32	5.66	0.80	32.57
8, MPC, true state	1.91	4.38	−3.48	5.74	0.10	24.26
9, Taylor rule	1.99	10.00	−3.75	5.88	−0.83	239.23
9, $a_\pi = 2$, $a_y = 0$	1.96	10.00	−5.79	6.90	−1.02	235.94
9, MPC, $\lambda_y = 2$	2.10	10.00	−1.98	4.99	−0.20	295.82

Inflation, gap and unemployment in %, growth in %/yr, all on the true states.

11.6.1 The Chair and the Rule

In case 1 a fiscal stimulus adds two points of demand for two years from year 0.25. Left at 2.875 %, the boom takes the gap to +1.89 % at year 2.3 and inflation to 3.24 % at year 3.3. The first CPI print above 2.25 % is published at year 0.38, and it is noise on a month whose true average was 2.00 %, while the true gap is already +0.33 %. A chair who raises 25 bp at every meeting while the print is above 2.25 %, and cuts while it is below 1.75 %, keeps raising for a year and a half after inflation peaked at 3.05 % at year 2.8, reaches 8.875 % at the meeting of year 4.3, and takes the gap to −4.93 % at year 6.1, unemployment to 6.47 % and inflation to −1.82 % at year 8, with a recession on the way. Steps of 50 bp make it worse, a rate of 11.875 % and a gap of −6.93 %. That chair is an integral

controller on a plant with two years of lag.

The Taylor rule on the same boom in case 2 raises the rate to 5.625 % at year 2.5, holds inflation to 2.95 % and the gap between +1.61 % and −0.90 %, and scores a ten-year loss of 4.24. A stronger inflation response mostly makes a bigger move at the wrong time. At $a_\pi = 2.5$ the inflation peak falls only to 2.90 % while the gap falls to −1.33 % and the loss rises to 4.85. A stronger gap response helps, 3.93 at $a_y = 1$ with inflation peaking at 2.88 %, because on a demand shock the gap moves a year before inflation and is the earlier warning of the same disturbance. Dropping the gap term costs 5.49, a partial adjustment of 0.7 smooths the rate to 41 moves instead of 66 for a loss of 5.17, and the MPC at its shipped weight scores 5.19 by pushing the gap to −1.66 %.

> **TRY IT IN THE APP 2 · The Taylor Rule on the Same Boom**
>
> Load the case and press Run. Watch the applied policy rate on the trend. The rule raises it in quarter-point steps to a peak of 5.625 % at year 2.5, inflation peaks at 2.95 %, and after ten years the scorecard gives an integrated loss of 4.24.

11.6.2 The Principle, the Bias and the Frontier

In case 3 no one is anchored and the rule answers each point of inflation with 0.8. The real rate then sits on average 0.19 points below the true neutral rate from year 10 on, the gap stays near +0.19 %, and inflation climbs without end, 2.92 % at year 5, 3.17 % at 10, 3.48 % at 20 and 4.77 % at 40. At 1.5 the same boom dies out, 2.47 % at year 10 and 2.08 % at 40. The climb is slow because the gap that drives it is small, and its direction is the lesson.

In case 4 the true neutral rate falls from 0.875 % to −0.125 % at year 1 and the bank is not told. The Taylor rule, a proportional controller with the wrong intercept, can run a lower rate only while inflation stays below target, and it settles at 1.52 % at year 20. The PID's integral moves the intercept down by 0.84 points by year 20 and ends at 1.86 %, and an integral time of four years moves it 1.05 points and ends at 1.98 %. The integral time that brings inflation within 0.05 points of the target by year 20 therefore lies between those two, since eight years leave 0.14 points and four leave 0.02. The MPC, whose estimator carries a demand bias, ends at 2.01 % with a loss of 0.68 against the rule's 4.60. The integral finds the neutral rate that real-time data misstate [14].

> **TRY IT IN THE APP 4 · The PID Finds a Neutral Rate That Moved**
>
> Load the case and run twenty years: under the PID inflation ends at 1.86 %. Then set the controller to the Taylor rule, press Reset and run again. Without an integral the rule settles at 1.52 %, below the target for good, because a proportional controller with the wrong intercept needs an error to hold a lower rate.

In case 5 an oil shock adds 3 percent a year to inflation for a year. At the shipped weight the MPC lets inflation reach 4.77 % at year 2.4 and the gap fall to −1.26 % at year 3.6. Across the slider the peak inflation error stays between 2.75 and 2.81 points while the worst gap ranges from −2.20 % to −0.49 %, and even the MPC given the true state, at a weight of 0.05, peaks 2.59 points over target. No output weight, rule or information keeps inflation and the gap both within a point, which the application's check confirms over seven weights, the true state, the Taylor rule and the PID. The weight trades the RMS gap, from 0.64 to 0.20 points, against the RMS inflation error, from 0.78 to 0.97, as shown in Figure 11.2(b). The Taylor rule sits at 0.64 and 0.77, a little inside the MPC's low-weight end, because the MPC's curve is the frontier of its own family on an estimated state with a frozen anchor, not the efficient frontier of the economy [8]. The true state moves it to 0.71 and 0.70.

11.6.3 Cadence and the Floor

Case 6 repeats the boom with noise-free releases, so that runs differ only in their calendar. Over ten years the Taylor rule's loss is 3.46 deciding every week, with 22 moves in 520 decisions, 3.53 monthly, 3.64 at eight meetings a year, with 19 moves in 80, and 4.03 quarterly, and inflation peaks

at 2.85 %, 2.86 %, 2.87 % and 2.90 %. The lag dominates the sampling. The MPC moves harder and gains more from frequent decisions, 3.40 weekly against 4.23 at eight meetings and 5.44 quarterly. The relay cycle of a coarse quantum is the subject of Section 11.4.

In case 7 a financial crisis takes six points off demand at year 1. The PID, given a two-year integral time so that windup shows, reaches the floor at year 2.1 and spends 2.25 years there. The gap falls to −3.25 % at year 2.0 and inflation to 0.34 % at year 3.1. The trough is the lag's, not the floor's. With the floor removed, $i_{min} = -2$ %, it is the same −3.25 %, because the crisis does its damage before any cut reaches spending. With floor tracking off the integral winds up, the PID asks for −2.79 % at year 3.5, holds the floor until year 5.1 instead of 4.5, and the recovery overshoots into OVERHEATING, a gap of +2.19 % at year 5.75 and inflation of 3.23 % at year 7.6, raising the loss from 17.43 to 22.84. A target of 1 % lines the economy out at a neutral rate a point nearer the floor, so the rate spends 3.71 years at the floor and inflation is negative for 2.71 years, down to −0.74 %, while the gap's trough stays −3.25 %.

> **DID YOU KNOW? Removing the Floor Does Not Lower the Loss**
>
> From the simulator: with the lower bound moved from 0.125 % to −2 % in case 7, the gap's trough stays at −3.25 %. The PID with floor tracking at the real floor scores a loss of 17.43, the run with the floor at −2 % scores 22.58 (Table 11.5).

11.6.4 The Optimizer and Credibility

Case 8 fires a boom at year 1, an oil shock at 5 and a crisis at 9 with the same prints in every run. The MPC's twenty-year loss is 28.93 against the PID's 41.56 and the Taylor rule's 32.57, with lower peaks of unemployment (5.75 % against 5.93 %) and inflation (4.30 % against 4.43 %) than the PID. Its worst calendar year is worse, 0.43 percent growth in the year ending at 10 against the PID's 0.70 in the year ending at 5, and it has two recessions, at years 3.75 and 9.50, to the PID's none, because it brings the gap down faster and growth is the gap's rate of change. Given the true state its loss falls to 24.26, a larger gain than its whole lead over the Taylor rule, so in this economy information is worth more than the optimizer.

Case 9 starts at 10 % inflation held steady by a gap of +2.4 % against an anchor of 0.3 that drifts with performance. The Taylor rule raises the rate to 14.125 % at year 2.0, the gap falls to −3.75 % at year 3.6, and inflation is within half a point of the target for good from year 12.9, by which time half the public is anchored again, 0.78 by year 20. The disinflation costs 13.8 gap point-years, a sacrifice ratio of 1.73 per point. A faster rule, $a_\pi = 2$ and $a_y = 0$, gets there at year 5.4 for a ratio of 2.43 and unemployment of 6.90 %, and the MPC at $\lambda_y = 2$ takes 15.7 years at a ratio of 1.35. A faster disinflation costs more output in this model, and whether speed buys credibility faster in an economy that learns is the question the backward model cannot answer [11].

11.6.5 A 1970s Decade

The challenge fires two oil shocks of 6 percent a year for a year, at years 1 and 5, with 40 percent of the public anchored and the share drifting, eight meetings a year and 25 bp steps. It asks for inflation within half a point of 2 % throughout years 12 to 15, no calendar year of growth below −2 % and unemployment never above 8 %, judged on the true states. Table 11.6 compares six configurations, and Figure 11.4 charts two of its columns. As shipped, the Taylor rule lets inflation reach 8.26 % at year 6.4 and expected inflation 4.94 %, and STAGFLATION shows for 5.4 years. In the slump after the second shock its gap term eases while expectations are still high. The real rate is zero at year 10 with inflation at 3.23 %, inflation is up to 1.10 points off target in the last three years, and the anchored share, 0.48 at year 2, has fallen to 0.38 by year 8 and is still 0.38 at year 12. The growth and unemployment limits are met, with a worst year of 0.10 % and unemployment of 5.85 %, so the rule fails on the late inflation alone.



Figure 11.3 The two-oil-shock challenge at $t = 15.00$ yr under forecast-targeting MPC with an output weight of 0.2. The figure is a composite of the Economy view, eight scorecard tiles set at its right and the Trends panel with a 20-year window. The header and the control panel are cut. The badges read inflation 2.05 %, gap −0.15 %, rate 3.375 %, unemployment 4.1 %, growth 1.7 %, real rate 1.33 %, anchored share 0.70 and a last move 1.1 quarters ago. (1) The valve holds 3.375 % while the MPC asks 3.35 %, with the next decision in four weeks. (2) The two pipes, tinted against the hidden 0.88 % neutral rate, carry 1.33 % into the first stage, 1.38 % out of it and 1.28 % out of the second. (3) The published gap reads 0.56 % against a true −0.15 %, and (4) the thermometer's CPI print 1.93 % against a true 2.05 % and an expected 2.0. (5) The meeting strip counts seven moves in the last two years, four up and three down. In the trends the published CPI and the rate rise twice after the shocks at years 1 and 5 (dotted marks), and the dashed marks are, by the run's record, DE-ANCHORED at the start and STAGFLATION at years 2.77 and 5.10. The scorecard gives a loss of 152.82, peak inflation 8.18 %, peak unemployment 6.14 %, a worst calendar year of 0.25 % growth, no time at the floor, no recession, 121 solves with no failure and the challenge met.

Table 11.6 The Challenge of 15 Years in Six Configurations

Configuration	$\lvert\pi - \pi^*\rvert_{12-15}$	g_{min}	u_{max}	π_{max}	Target
as shipped, Taylor rule	1.10	0.10	5.85	8.26	not met
MPC, $\lambda_y = 0.2$, suggested	0.16	0.25	6.14	8.18	met
MPC, $\lambda_y = 0.5$	0.92	0.46	5.77	8.70	not met
Taylor rule, $a_y = 0$	0.17	0.06	6.06	7.54	met
Taylor rule, $a_\pi = 2.5$	0.46	−0.11	6.48	7.57	met
PID	1.99	−0.01	6.75	7.55	not met

Limits 0.5 points, −2 %/yr and 8 %. Deviation in points, growth in %/yr.

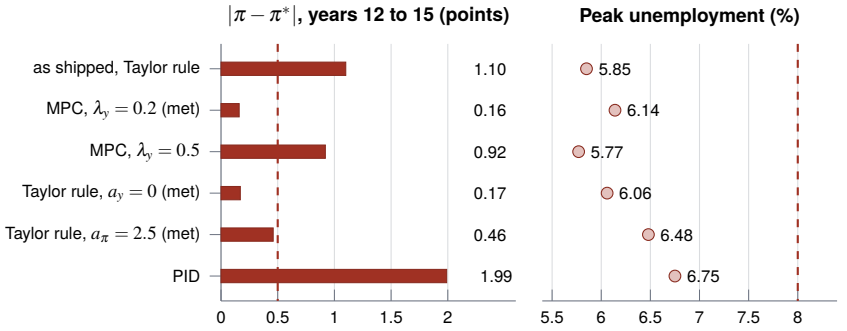


Figure 11.4 Two columns of Table 11.6, the 1970s challenge over 15 years: the largest inflation error over years 12 to 15 (bars, values at the right) against the limit of 0.5 points, and the peak unemployment (dots) against the limit of 8 % (dashed lines). Notice that every configuration that meets the target peaks higher in unemployment than the shipped rule and the MPC at its shipped weight, and that no configuration comes near the unemployment limit.

The suggested MPC, at an output weight of 0.2, keeps the real rate up until its forecast is at target, 0.63 % at year 10 against the rule's zero, and holds inflation within 0.16 points over years 12 to 15 (Table 11.6), with unemployment peaking at 6.14 % and a worst year of 0.25 %, as the scorecard of Figure 11.3 shows. Inflation returns within a point of the target early enough for the anchor to drift back, from 0.38 at year 8 to 0.56 at year 12 and 0.70 at year 15, against the rule's 0.55 at the end. A Taylor rule with no gap response also passes (0.17), as does an inflation response of 2.5 (0.46, at unemployment of 6.48 % and a year of –0.11 %), while the MPC at its shipped weight of 0.5 does not (0.92) and neither does the PID (1.99). The verdict is not a property of one seed. On the default seed and seeds 1 to 4 the shipped rule misses by 1.06 to 1.13 points and the suggested MPC passes with 0.10 to 0.23, and over seeds 1 to 20 the rule fails every time (1.06 to 1.28) and the MPC passes every time (0.10 to 0.37). No configuration of Table 11.6 comes near the growth or unemployment limit, the worst year being –0.11 % and the highest unemployment 6.75 %. Judged against a stricter unemployment limit of 6 %, none of the six configurations meets the challenge. Both lessons of the decade are about the weight on output, which is a tuning parameter here.

11.7 Technology in Practice

The bank of Section 11.5 reads two published statistics and moves one administered rate on a calendar. This section sets the real statistics, the rate and the committee beside the model.

What Measures It

Inflation and output are compiled from surveys and records, not read by a device. In the United States the consumer price index of a month is released in the middle of the next month, and the first estimate of a quarter's output about one month after the quarter, followed by a second and a third estimate and later comprehensive revisions [20]. The model's CPI print two weeks after the month, with noise σ_P and no revision, and its three GDP vintages of Section 11.3 stand for these releases. The published gap also needs potential output, whose end-of-sample estimate is the larger source of error in real time (Section 11.2). The model gives the print noise but no bias, so the fault described below is one the simulated bank never meets.

Table 11.7 Scenarios and Completion Requirements

#	Scenario	Kind	Required to complete
1	The chair	Experiment	In MANUAL, request a rate above 3.125 %, keep chairing to year 8, and confirm that the first rise crossed both transmission stages before the gap turned.
2	Taylor rule	Experiment	Ten years with the rule as shipped ($a_\pi = 1.5$), then ten years with $a_\pi = 2.5$ after Reset.
3	Taylor principle	Demonstration	With no anchor, forty years at $a_\pi < 1$ that end with inflation more than 2 points over target, then forty years at $a_\pi = 1.5$ after Reset that end within 0.3 points of it.
4	Neutral rate	Experiment	Twenty years of the PID with integral action after the hidden fall in r^*, then twenty years of the Taylor rule after Reset.
5	Oil shock	Demonstration	Twenty years of the MPC, then twenty more after Reset with $\lambda_y \leq 0.1$ or $\lambda_y \geq 2$, and confirm that no run keeps inflation and the gap both within a point.
6	Cadence	Experiment	Ten years at eight meetings a year with the 25 bp quantum, ten years on the weekly or the quarterly calendar after Reset, then thirty years with a quantum of at least 50 bp.
7	The floor	Experiment	With floor tracking on, run until the rate reaches the floor and on to year 12, then twelve years with floor tracking off and twelve with a target of at most 1 %, each after Reset.
8	MPC and PID	Experiment	Twenty years of the MPC through the three shocks, then twenty years of the PID after Reset.
9	Disinflation	Experiment	After year 1, hold inflation within 0.5 points of the target for three years with the shipped rule, then again with $a_\pi = 2$ and $a_y = 0$ after Reset.
C	1970s challenge	Challenge	Over the full 15 years, inflation within 0.5 points of the target from year 12 to 15, no calendar year of growth below −2 % and unemployment never above 8 %, on the case as given.[a]
R	Reference	Exploration	Three years of the equilibrium in AUTO, then in MANUAL a request of at least 3.375 % held until the gap falls below −0.2 %.

[a] As shipped, the Taylor rule misses by 1.10 points (worst year +0.10 %, peak unemployment 5.85 %). The suggested MPC at $\lambda_y = 0.2$ passes with 0.16 points, a margin of 0.34 on the band, a worst year of +0.25 % and unemployment of 6.14 %, revealed before Run or carried through Reset alike.

> **TECHNOLOGY** **Sensor: The Consumer Price Index**
>
> No instrument reads inflation. A statistical agency collects prices and weights them with expenditure estimates from a consumer expenditure survey [21]. In the United States the index of a month is released in the middle of the next month, 18 days after the month ends [20]. Its typical fault is a bias, not noise. A review in 2001 estimated that the index overstated the change in the cost of living by about 0.6 percentage point a year, with a confidence interval from 0.1 to 1.2, roughly half of it from quality change and new items, against an estimate of about 1.1 points made in 1996 [21]. On a trend a bias does not show: the print looks steady while the truth sits below it.

What Moves It

The final element is an administered short rate. A committee sets a target range for an overnight rate, and administered rates, among them the rate paid on reserve balances, support that range [22]. The quantum q of 25 bp and the grid of midpoints of Equation (11.7) copy a range a quarter point wide, the cap Δ limits a meeting's move, and the floor i_{min} stands for the effective lower bound. Figure 11.5 draws the loop with these institutions in place of the schematic's symbols.

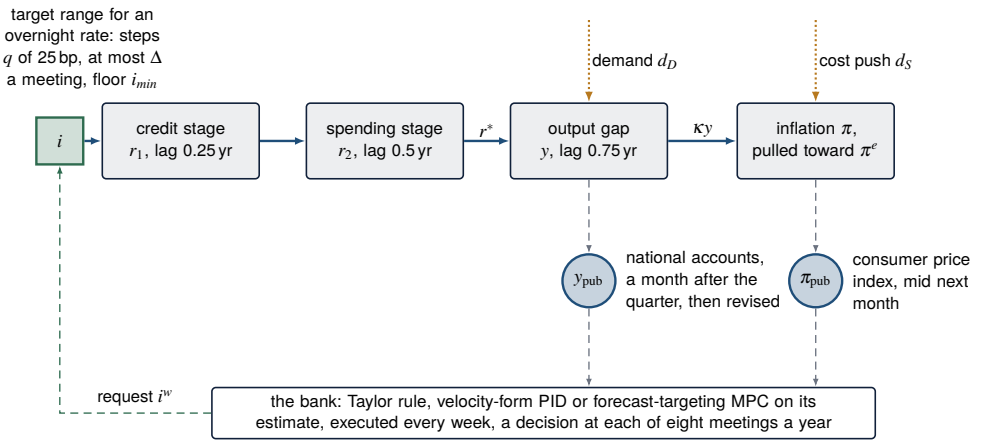


Figure 11.5 The policy loop with the institutions of a real central bank in place of the schematic's symbols. A price index and the national accounts give the published inflation π_{pub} and gap y_{pub}, the bank's policy forms a request each week, and only a meeting turns it into a step of the target range for the rate i. Notice the three lags in series between the rate and the gap, and that both readings leave the economy after its last stages and return weeks late.

> **TECHNOLOGY** **Final Element: A Target Range for an Overnight Rate**
>
> A central bank moves a rate it administers. The Federal Open Market Committee's primary means of adjusting the stance of policy is a change in the target range for the federal funds rate [13]. The Board sets the rate paid on reserve balances, an important tool, and the implementation note issued with each decision gives the settings of the tools that support the target range. The posted rate is updated each business day with the next business day's rate [22]. The element saturates: the same statement names a funds rate constrained by its effective lower bound [13], and central banks have judged that bound to be close to zero [23].

Where the Controller Runs

No computer closes this loop. The Federal Open Market Committee holds eight regularly scheduled meetings a year and other meetings as needed [24]. The simulated bank executes every week, $\Delta t = 1/52$ year, and its default calendar makes one execution in six to eight weeks a decision, so the meeting, not the weekly execution, is the period that counts. In the application the Taylor rule, the velocity-form PID of Equation (11.9) and the MPC fill the same slot ahead of the framework's shared calendar module, `core/cadence.js`, which rounds, caps and holds each request.

> **IN PRACTICE** **Rules as Benchmarks, Not Autopilots**
>
> The Federal Open Market Committee regularly consults the prescriptions of several rules with other information and does not follow them mechanically. Its published set holds the Taylor rule, a balanced-approach rule with twice the gap coefficient, a rule adjusted for the effective lower bound, an inertial rule that keeps 0.85 of the previous rate, and a first-difference rule that sets only the change of the rate and needs no estimate of the neutral rate or of potential output. Following a rule strictly would first require a choice of inflation measure, activity measure and neutral rate [23]. The partial adjustment ρ of Equation (11.8) has the inertial form, and a first-difference rule is an integral controller in velocity form.

What Is Modernizing It

Three developments sharpen what the bank reads. The first is nowcasting. A dynamic factor model in state-space form reads many series as they are released, a Kalman filter updates its estimate with each

release, and one central bank's staff nowcast is updated daily when new data arrive and published weekly [20]. The second is nowcasting of prices alone: a model built on relatively few variables nowcasts headline and core inflation in the consumer price index and in the price index for personal consumption expenditures, with accuracy that improves as information accumulates over a month or a quarter [25]. The third is a filter for the unobserved intercept. The natural rate of interest, potential output and trend growth can be estimated jointly by a Kalman filter [26]. The bank's estimator of Equation (11.10) is a filter of the same kind with fixed teaching gains, whose demand bias stands in for an error in the neutral rate, and forecast targeting by MPC is the chapter's own model-based plan on that estimate.

11.8 Limitations

The economy is a teaching model. Its expectations look backward, so the Lucas critique applies with full force [11], and the objection that control theory does not apply to planning under rational expectations [5] is avoided by assumption, not answered. The lags, the slope of the Phillips curve and every threshold are teaching values, the design windows for the troughs are the application's own, and no number here is an estimate for any country. The output gap and Okun's law are local and linear.

The bank's estimator has fixed teaching gains rather than a Kalman design, reads the anchored share as an exact survey, and is corrected only by the two releases. The MPC freezes the anchored share over its horizon, knows nothing of shocks to come, and uses one decision interval as its step with blocked moves later in the horizon. The policy rate is the only instrument, so a crisis that needs a rate below the floor has no other answer. The statistics are constructed averages of the model's own states.

The inflation target, the quantum and the floor also set the starting state. A change made before Run restarts the case with it, as Reset does (Section 1.6.1), so the application gives one run whichever way a setting is chosen.

> **OPEN CHALLENGE** **Where Is the Neutral Rate Now?**
>
> Every rule with an intercept needs the neutral real rate, which no one observes. A study that estimated it with potential output and trend growth by a Kalman filter found substantial variation over four decades in the United States, an average standard error of 1 to 2-1/2 percentage points, and a significant deterioration of stabilization when policymakers mismeasure the rate, which timely updating of the estimates can mitigate [26]. The fall of one point in case 4 lies inside that error. Progress would be estimates precise enough in real time to set the intercept, or rules that, like the first-difference rule, do not need one [23].

11.9 Conclusion

This chapter's simulator joins a continuous economy of three lags to a statistics agency that publishes late averages with revisions and to a bank that may move one rate on a calendar, in quarter-point steps, above a floor.

A rise held a quarter peaks in output after 1.04 years and in inflation after 2.40, and a chair who chases the print takes the rate to 8.875 % and the gap to −4.93 %. The Taylor principle's boundary falls from 1.0 with no anchor to 0.25 with 30 percent anchored, the PID's integral finds an unannounced fall in the neutral rate where the rule settles at 1.52 %, and no policy holds both inflation and the gap within a point through an oil shock. Deciding every week lowers a boom's loss only from 3.64 to 3.46, the floor's trough is the lag's at −3.25 %, and in the 1970s challenge the shipped rule misses by 1.10 points where forecast targeting at an output weight of 0.2 holds 0.16.

The expectations are backward-looking and every coefficient is a teaching value. A fuller model would add a learning public, a second instrument such as asset purchases or forward guidance, and an estimator designed from the release statistics. The simulator makes the lag, the vintage and the calendar of monetary policy measurable in the same loop.

KEY POINTS

- The economy answers a rate move in output after about a year and in prices after about two, so a chair who keeps raising while the print is high takes the rate to 8.875 % and the gap to −4.93 %.
- The Taylor principle's boundary falls from 1.0 with no anchor to 0.25 with 30 percent of the public anchored.
- The PID's integral finds an unannounced fall in the neutral rate where the proportional rule settles at 1.52 %.
- Deciding every week lowers a boom's loss only from 3.64 to 3.46, because the lag dominates the sampling.

Where else this applies. Volume 5, Chapter 9 moves a contribution through the same decision calendar, quantum and cap, and Volume 5, Chapter 3 controls a column through a late measurement.

References

[1] A. W. Phillips, "Stabilisation policy in a closed economy," *The Economic Journal*, vol. 64, no. 254, pp. 290–323, 1954.

[2] R. Pindyck, "An application of the linear quadratic tracking problem to economic stabilization policy," *IEEE Transactions on Automatic Control*, vol. 17, no. 3, pp. 287–300, 1972.

[3] J. B. Taylor, "Discretion versus policy rules in practice," *Carnegie-Rochester Conference Series on Public Policy*, vol. 39, pp. 195–214, 1993.

[4] L. E. O. Svensson, "Inflation forecast targeting: Implementing and monitoring inflation targets," *European Economic Review*, vol. 41, no. 6, pp. 1111–1146, 1997.

[5] F. E. Kydland and E. C. Prescott, "Rules rather than discretion: The inconsistency of optimal plans," *Journal of Political Economy*, vol. 85, no. 3, pp. 473–491, 1977.

[6] D. E. Seborg, T. F. Edgar, D. A. Mellichamp, and F. J. Doyle, III, *Process Dynamics and Control*, 4th ed. Hoboken, NJ: John Wiley & Sons, 2016.

[7] G. Rudebusch and L. E. Svensson, "Policy rules for inflation targeting," National Bureau of Economic Research, Cambridge, MA, NBER Working Paper w6512, Apr. 1998.

[8] L. Ball, "Efficient rules for monetary policy," *International Finance*, vol. 2, no. 1, pp. 63–83, 1999.

[9] A. W. Phillips, "The relation between unemployment and the rate of change of money wage rates in the United Kingdom, 1861–1957," *Economica*, vol. 25, no. 100, pp. 283–299, 1958.

[10] M. Friedman, "The lag in effect of monetary policy," *Journal of Political Economy*, vol. 69, no. 5, pp. 447–466, 1961.

[11] R. E. Lucas, "Econometric policy evaluation: A critique," *Carnegie-Rochester Conference Series on Public Policy*, vol. 1, pp. 19–46, 1976.

[12] R. Clarida, J. Galí, and M. Gertler, "The science of monetary policy: A new Keynesian perspective," *Journal of Economic Literature*, vol. 37, no. 4, pp. 1661–1707, 1999.

[13] "Statement on longer-run goals and monetary policy strategy," Board of Governors of the Federal Reserve System, Federal Open Market Committee. Adopted effective January 24, 2012; as reaffirmed effective January 27, 2026, 2026, accessed 2026-11-20. [Online]. Available: https://www.federalreserve.gov/monetarypolicy/files/FOMC_LongerRunGoals.pdf

[14] A. Orphanides, "Monetary policy rules based on real-time data," *American Economic Review*, vol. 91, no. 4, pp. 964–985, 2001.

[15] A. Orphanides and S. van Norden, "The unreliability of output-gap estimates in real time," *Review of Economics and Statistics*, vol. 84, no. 4, pp. 569–583, 2002.

[16] G. B. Eggertsson and M. Woodford, "Zero bound on interest rates and optimal monetary policy," *Brookings Papers on Economic Activity*, vol. 2003, no. 1, pp. 139–233, 2003.

[17] L. J. Christiano, M. Eichenbaum, and C. L. Evans, "Nominal rigidities and the dynamic effects of a shock to monetary policy," *Journal of Political Economy*, vol. 113, no. 1, pp. 1–45, 2005.

[18] "FOMC statement, december 16, 2008," Board of Governors of the Federal Reserve System, press release, 2008, accessed 2026-11-20. [Online]. Available: https://www.federalreserve.gov/newsevents/pressreleases/monetary20081216b.htm

[19] D. Goldfarb and A. Idnani, "A numerically stable dual method for solving strictly convex quadratic programs," *Mathematical Programming*, vol. 27, no. 1, pp. 1–33, 1983.

[20] B. Bok, D. Caratelli, D. Giannone, A. Sbordone, and A. Tambalotti, "Macroeconomic nowcasting and forecasting with big data," Federal Reserve Bank of New York, Staff Report 830, 2017. [Online]. Available: https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr830.pdf

[21] D. E. Lebow and J. B. Rudd, "Measurement error in the consumer price index: Where do we stand?" Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series 2001-61, 2001. [Online]. Available: https://www.federalreserve.gov/pubs/feds/2001/200161/200161pap.pdf

[22] "Interest on reserve balances," Board of Governors of the Federal Reserve System, 2026, accessed 2026-11-20. [Online]. Available: https://www.federalreserve.gov/monetarypolicy/reserve-balances.htm

[23] "Policy rules and how policymakers use them," Board of Governors of the Federal Reserve System, Monetary Policy Principles and Practice, 2026, accessed 2026-11-20. [Online]. Available: https://www.federalreserve.gov/monetarypolicy/policy-rules-and-how-policymakers-use-them.htm

[24] "Meeting calendars, statements, and minutes (2021–2027)," Board of Governors of the Federal Reserve System, Federal Open Market Committee, 2026, accessed 2026-11-20. [Online]. Available: https://www.federalreserve.gov/monetarypolicy/fomccalendars.htm

[25] E. S. Knotek, II and S. Zaman, "Nowcasting U.S. headline and core inflation," Federal Reserve Bank of Cleveland, Working Paper 14-03, 2014.

[26] T. Laubach and J. C. Williams, "Measuring the natural rate of interest," Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series 2001-56, 2001. [Online]. Available: https://www.federalreserve.gov/pubs/feds/2001/200156/200156pap.pdf

11.10 Exercises

11.1 Concept check · 2 min

The PID of Equation (11.9) forms its error as $e_k = \pi_{\text{pub},k} - \pi^*$. In the terms of Section 1.4.3, which statement is right?

(a) Direct acting with $a = -1$, because a higher rate lowers inflation and the rate must rise with the PV

(b) Reverse acting with $a = +1$, because the bank pushes against inflation that is above the target

(c) Reverse acting with $a = -1$, because a higher rate raises inflation through the output gap

(d) Direct acting with $a = +1$, because the published CPI and the rate are both in percent

11.2 Concept check · 2 min

On the noise-free reference, a rate held one point up from $t = 0$ has moved the output gap by –0.09 % after half a year and by –0.35 % after a year. What does the reading at half a year show?

(a) The rate reaches the gap after a transport delay of about a year, and the early reading is noise

(b) The transmission is a chain of lags in series, so the response starts at once and arrives late

(c) The two-week delay of the CPI print holds back the bank, so the gap answers only to the old rate

(d) The rise waits for the next of eight meetings, so half a year covers only three or four decisions

11.3 Concept check · 2 min

In case 6, with a rate quantum of 0.5 points and the Taylor rule on eight meetings a year, the rate alternates between 2.625 and 3.125 % from year 20 to 30. An inflation response of 3 instead of 1.5 leaves the gap cycle almost unchanged. What sets the size of the cycle?

(a) The inflation response, because a proportional rule with a high gain is close to instability

(b) The noise of the CPI print, which pushes the request across the rounding boundary at random

(c) The quantum, because neither grid rate next to the neutral 2.875 % holds the economy still

(d) The integral of the rule, which winds up between meetings while the rate is held

11.4 Theory and calculation · 7 min

At a steady state with the Taylor rule acting on true values and $\hat{r}^* = r^*$, the model gives $\pi^a = \pi$, $y = \lambda x/(\kappa\tau_\pi)$ and $y = -\sigma[(a_\pi - 1 + \lambda)x + a_y y]$ for an inflation deviation $x = \pi - \pi^*$. Use $\sigma = 1$, $\kappa = 0.5$ and $\tau_\pi = 2\,\text{yr}$.

(a) Eliminate y to recover Equation (11.6), and evaluate the factor $1 + (1 + \sigma a_y)/(\kappa\sigma\tau_\pi)$ at $a_y = 0.5$. *Checkpoint:* 2.5, so the boundary is $a_\pi = 1 - 2.5\lambda$.

(b) Give the boundary on a_π at anchored shares of 0, 0.1, 0.2 and 0.3. *Checkpoint:* 1.0, 0.75, 0.5 and 0.25.

(c) At $\lambda = 0.3$, find the boundary for $a_y = 0$ and for $a_y = 1$, and explain why a stronger gap response lowers it. *Checkpoint:* 0.4 and 0.1.

11.5 Theory and calculation · 6 min

Use Equations (11.2) and (11.3) with $\kappa = 0.5$, $\tau_\pi = 2\,\text{yr}$, $\tau_e = 1.5\,\text{yr}$, the credibility drift off and no shocks.

(a) With $\lambda = 0$ and the gap held at one point, let $c = \pi - \pi^a$. Show that c settles and find the constant rate at which inflation then rises. *Checkpoint:* $c = 0.43$ points, a rise of 0.286 points a year.

(b) Case 9 starts with $\pi = \pi^a = 10\,\%$, $\lambda = 0.3$ and $\pi^* = 2\,\%$. Find the expected inflation and the gap that hold inflation still. *Checkpoint:* $\pi^e = 7.6\,\%$ and $y = 2.4\,\%$.

(c) With $u^* = 4\,\%$, $\beta = 2$ and $r^* = 0.875\,\%$, find the starting unemployment and the real rate r_2 that holds that gap, and the nominal rate $r_2 + \pi^e$ before it is rounded to the grid. *Checkpoint:* 2.8 %, −1.525 % and 6.075 %.

11.6 Theory and calculation · 5 min

Use the shaping of Equation (11.7) with $b = 0.125$, $i_{min} = 0.125\,\%$ and $i_{max} = 20\,\%$.

(a) The rate in force is 2.875 %, $q = 0.25$, $\Delta = 0.5$, and the controller asks for 3.70 % at three meetings in a row. Give the rate after each meeting. *Checkpoint:* 3.375, 3.625 and 3.625 %.

(b) With $q = 0.5$ and the largest move set to $\Delta = 0.25$, what is the largest move a meeting can make? *Checkpoint:* 0.5 points.

(c) In the application's one-step check, $T_s = 1/8$ year, $\nu = 1$, $b_\pi = -4.47 \times 10^{-5}$ and the forecast is 2.958 %. Find the move of Equation (11.12), and the forecast error that would be needed for the 0.25 grid to move the rate at all. *Checkpoint:* 5.35×10^{-6} points, and an error of about 22 000 points.

11.7 App challenge · 8 min

Load 2 · The Taylor rule on the same boom and run the full ten years at the default seed.

(a) Read the Integrated policy loss and Peak inflation tiles. *Checkpoint:* 4.24 and 2.95 %, within 0.01.

(b) Set Total inflation response to 2.5, press Reset and run ten years again. Read the same tiles and the lowest true gap, column y of the exported file. *Checkpoint:* 4.85, 2.90 % and −1.33 %, within 0.01.

(c) Set Total inflation response back to 1.5 and Output-gap response to 1, press Reset and run again. *Checkpoint:* 3.93 and 2.88 %, within 0.01.

Explain why the stronger gap response helps on this boom and the stronger inflation response does not.

11.8 App challenge · 8 min

Load 4 · The PID finds a neutral rate that moved. The true neutral rate falls a point at year 1 and the bank is not told.

(a) Run twenty years at the default seed. Read the last value of `pi` in the exported file and how far the PID card says the integral has moved the intercept. *Checkpoint:* 1.86 % and −0.84 points, within 0.01.

(b) Set Integral time (0 disables) to 0, press Reset and run twenty years. Then set Controller to Taylor rule, press Reset and run again. Read inflation at year 20 and the Integrated policy loss for both runs. *Checkpoint:* 1.52 % and 4.60 in both runs, within 0.01.

(c) Set Controller to PID + output gap and Integral time to 4, press Reset and run twenty years. *Checkpoint:* 1.98 %, −1.05 points and a loss of 1.05, within 0.01.

Explain the result of part (b) with Section 1.4.2.

11.9 App challenge · 9 min

Load 7 · The floor: a financial crisis, in which the PID has a two-year integral time.

(a) Run twelve years at the default seed. Read Time at rate floor, Peak inflation and Integrated policy loss. *Checkpoint:* 2.25 yr, 2.74 % and 17.43, within 0.01.

(b) Switch to Advanced, set PID floor tracking to Off, press Reset and run twelve years. Read the same tiles, and find the lowest request in column `want` of the exported file and its time. *Checkpoint:* 3.00 yr, 3.23 % and 22.84, –2.79 % at year 3.5, within 0.02.

(c) The lowest true gap is –3.25 % at year 2.0 in both runs. Explain why switching off floor tracking leaves the trough alone and still costs 5.4 in loss. *Checkpoint:* The trough comes before any cut reaches spending.

11.10 Application area · 4 min

The true potential growth falls from 2 to 1.5 % a year at $t = 0$, and the bank keeps believing $\hat{g}^* = 2\,\%$ a year.

(a) Use Equation (11.4) to find m after four years, and say by how much the Taylor rule with $a_y = 0.5$ then sets the rate differently from the same rule fed the true gap. *Checkpoint:* $m = -2$ points, a rate one point lower.

(b) Name a process measurement that fails in the same way, and the feedback in this chapter that removes the resulting bias. *Checkpoint:* A drifting analyzer or inferential estimate, and integral action on an unbiased measurement.

11.11 Application area · 3 min

In case 7 the rate cannot go below 0.125 %.

(a) With a target of 1 % instead of 2, the rate spends 3.71 years at the floor instead of 2.25. Explain why from the neutral nominal rate. *Checkpoint:* 1.875 % instead of 2.875 %, a point less room to cut.

(b) Name the two instruments the chapter says the model lacks, and the plant practice that the higher target corresponds to. *Checkpoint:* Asset purchases and forward guidance. Running farther from an actuator limit.

11.12 Review the agent · 5 min

An AI assistant proposes a model for an inner loop from the policy rate to the output gap: “The gap answers a rise held a quarter most after 1.04 years, so the path is a first-order-plus-dead-time process with $K_p = -1$ point of gap per point of rate, $\theta_p = 1.04$ yr and $\tau_p = \tau_y = 0.75$ yr. SIMC with $\tau_c = \theta_p$

then gives $|K_c| = 0.36$ and $\tau_I = 0.75\,\text{yr}$." Find the flaw, give a better estimate of the dead time and time constant, and say how you would check it in the app.