Optimal Monetary Policy When the Economy Must Be Learned

Frank Schmid
frank.schmid@frankschmid.com

This version: 15 November 2026

First version: 6 September 2026

Abstract. The volume on inflation-forecast targeting edited by Adrian, Laxton, and Obstfeld sets beside a simple rule a policy that minimizes an explicit loss function against a model. This paper asks what such a policy is when the model must be learned from data, and it measures learning at four levels on U.S. data that are real-time but for the neutral-rate path, the price index, and the term premium. Three alternatives to model-consistent expectations are each estimated at the model-consistent value. The likelihood determines the scale of the estimated policy equation only weakly. Against a null distribution simulated with a period at the effective lower bound, nonparametric corrections to the residuals show evidence of signal in the investment–saving equation and none in the Phillips equation. In the paper's reproduction of an experiment of the volume, a rule trained by reinforcement learning inside the estimated model holds the policy rate at the effective lower bound beyond the liftoff of the volume's inflation-forecast-based rule. At the decision points of the Federal Open Market Committee, it prescribes rates above the realized federal funds rate on average; the ratio of that mean difference to its standard error exceeds its critical value over the whole sample and not in every subperiod. A rule estimated and trained on data through 2015 and held fixed prescribes rates nearer the realized rate thereafter, in mean and in root-mean-square difference, than the rule trained on the whole sample. The ratio of the mean difference between the two rules' prescriptions to its standard error does not exceed its critical value over 2016 through 2026 or over 2016 through 2020 and exceeds it over 2021 through 2026, when no observation of the output gap enters either estimate's filter. The same holds for the two parts into which those years are divided at any quarter from the third quarter of 2019 through the second quarter of 2025, so that a change cannot be dated at the end of 2020, and the paper reports no test of a change. In an accounting that exchanges three things between the two rules and their estimates (the trained coefficients, the parameters of the investment–saving and Phillips equations, and the output gaps the estimates infer), the difference between the two rules' prescriptions over 2021 through 2026 is accounted for by the gaps and not by the coefficients or the parameters. A rule re-estimated and retrained each quarter prescribes rates above the realized rate on average from 2016.

JEL classification: E52, E58, E37, C45, D83.

This paper was developed with assistance from generative AI tools across the design, analysis, coding, and drafting stages. The author supervised all stages, verified the results, and takes full responsibility for the content.

Contents

Abbreviations and Acronyms ii

1 Introduction 1

2 The Real-Time Data 6

3 Optimal Policy, Rules, and Learners 9
3.1 The volume's experiment on the learned rule 10

4 The World Model 12
4.1 Blocks and expectations 12
4.2 Estimator and Monte Carlo study 13
4.3 The multi-start search and the baseline estimate 14
4.4 The model's forecasts against the staff's 17

5 Model-Consistent and Alternative Expectations 18

6 Residual Components and a Null Distribution with an ELB Episode 21

7 The Policy Agent 22
7.1 Loss and training environment 23
7.2 Training with and without the residual components 23
7.3 Training runs and the selected checkpoint 24

8 The Learned Rule and the FOMC 27
8.1 The baseline agent and the realized federal funds rate 27
8.2 The agent frozen at the end of the first ELB episode 31
8.3 The agent re-estimated and retrained each quarter 37
8.4 Counterfactual outcomes 39

9 Sensitivity and Limitations 41

10 Conclusion 46

A Conditions of the Monte Carlo Study 48

B Five Questions on the Estimates and the Agents 49

C Mean Differences by Treatment, Paired Differences, and Losses under the CBO Measure 50

Bibliography 53

Abbreviations and Acronyms

ACM	Adrian–Crump–Moench term-premium model
ALFRED	Archival Federal Reserve Economic Data
bp	basis points (one percentage point is 100 basis points)
CBO	Congressional Budget Office
CPI	consumer price index
DM1	dual-mandate loss function of Adrian, Laxton, and Obstfeld's volume, chapter 12, at its baseline weights, and the policy that minimizes it (DM2 halves the weight on the output gap)
ELB	effective lower bound
FOMC	Federal Open Market Committee
FRB/US	the Federal Reserve Board's large-scale model of the U.S. economy
GDP	gross domestic product
GP	Gaussian process
HLW	Holston–Laubach–Williams neutral-rate model
IFB	inflation-forecast-based (rule)
IS	investment–saving (equation)
LOO	leave-one-out
MAP	maximum a posteriori
PCE	personal consumption expenditures
pp	percentage points
RMS	root mean square
RMSE	root-mean-square error
RTDSM	Real-Time Data Set for Macroeconomists
SD	standard deviation
SE	standard error
SPF	Survey of Professional Forecasters

Greek letters and other symbols

α (alpha)	weight on squared inflation deviations in the loss
β (beta)	weight on the squared output gap in the loss
γ (gamma)	weight on the squared policy-rate change in the loss
ε (epsilon)	structural shock of the IS (ε^y) and Phillips (ε^π) equations
η (eta)	innovation to the shock of the estimated policy equation
κ (kappa)	slope of the Phillips curve
λ_1 (lambda)	weight on expected four-quarter inflation in the Phillips curve
π (pi)	inflation
ρ (rho)	interest-rate smoothing
σ (sigma)	standard deviation
ϕ_u (phi)	persistence of the shock of the estimated policy equation
c_π	response to the inflation forecast's deviation from target
c_y	response to the output gap
c_0	intercept, in percentage points

ρ, c_π, c_y, and c_0 are the coefficients of the estimated policy equation (4) and, in Section 7, of the agent's rule of the same form.

1 Introduction

The volume on inflation-forecast targeting edited by Adrian, Laxton, and Obstfeld sets, in its chapter on the United States, a policy that minimizes an explicit loss over the forecast beside a policy that follows a simple rule.[1] The loss is quadratic in the deviation of inflation from target, the output gap, and the change in the policy rate. Because the loss is quadratic, a large deviation weighs more than in proportion to its size, and the volume argues that this gives policy the risk avoidance a central bank needs near what Blanchard calls a "dark corner," such as the effective lower bound (ELB) on the policy rate. The volume credits the phrase to Blanchard where it introduces the loss, and the working paper from which the volume's model comes is titled for it.[2] In the volume's experiment from the state of 2014:Q4, the loss-minimizing path holds the policy rate at the ELB until the middle of 2016, while an inflation-forecast-based (IFB) rule raises the rate above the ELB, or lifts off, at once. That optimal policy is optimal against a calibrated model, and the volume's confidence bands come from the history of forecast errors around it.

This paper asks what optimal policy is when the model of the economy must be learned from data. A policy that minimizes a loss inside a model depends on the model: a policy optimized against a model is optimized against the model's errors as well. Learning enters at four levels, and the paper measures each. The *economy's agents*, households and firms, may or may not learn: the expectations block is where forward guidance at the ELB has its effects, and it is where the behavioral and learning literatures have proposed departures from model-consistent expectations.[3] The *econometrician* learns the model, with estimators whose behavior on simulated economies with an ELB episode is measured in a Monte Carlo study, and with a search from many starting values, because the likelihood has more than one local optimum. A *correction* learns what the model's structure misses: Gaussian process (GP) components fitted to the residuals of the model's investment–saving (IS) and Phillips equations, whose evidence of signal is judged against a null distribution simulated with an ELB episode. And the *policymaker* learns the rule. A policy agent is trained by reinforcement learning inside the estimated model (Section 3 states the method) on draws of four of the model's coefficients around the estimate (Section 7.3). The training proceeds on episodes, an episode being a simulated path of the estimated economy from a real-time starting state, and the rule reported is selected among eight training runs, each with its own random-number stream, by its loss on a validation set, a fixed set of episodes not used in training. Figure 1 shows the four learners and the checks applied to them; the economy's agents and the correction are both part of the world model, the estimated model of the economy inside which the agent is trained.

The model is a reduced form of the semi-structural economy of Adrian, Laxton, and Obstfeld's volume: an IS equation, a hybrid Phillips curve, an expectations block specified in this paper, and a policy equation with an ELB. It is estimated on 122 quarters of real-time data from 1996:Q1 through 2026:Q2, which contain 36 quarters at the ELB in two episodes and the inflation of 2021–2023;

[1] Rania Al-Mashat et al., "United States: Federal Reserve's Dual Mandate," in *Advancing the Frontiers of Monetary Policy*, ed. Tobias Adrian, Douglas Laxton, and Maurice Obstfeld (Washington, DC: International Monetary Fund, 2018), 209–237.

[2] Olivier Blanchard, "Where Danger Lurks," *Finance & Development* 51, no. 3 (2014); Ali Alichi, Kevin Clinton, Charles Freedman, Ondra Kamenik, Michel Juillard, Douglas Laxton, Jarkko Turunen, and Hou Wang, "Avoiding Dark Corners: A Robust Monetary Policy Framework for the United States," IMF Working Paper 15/134 (Washington, DC: International Monetary Fund, 2015).

[3] Xavier Gabaix, "A Behavioral New Keynesian Model," *American Economic Review* 110, no. 8 (2020): 2271–2327; George W. Evans and Seppo Honkapohja, *Learning and Expectations in Macroeconomics* (Princeton, NJ: Princeton University Press, 2001).

the neutral-rate path, the price index, and the term premium are the inputs that are not real-time (Section 2). The policy agent is an IFB rule with four learned coefficients. It is trained on episodes that start from the real-time states of the 118 decision points from 1997:Q1 through 2026:Q2, each the final scheduled meeting of the Federal Open Market Committee (FOMC) in its quarter. At each decision point the rate the agent prescribes for the decision quarter is compared with the realized federal funds rate, the quarter-to-date average of the effective federal funds rate published before the meeting (Section 2 gives the full definition). Differences of interest rates are in percentage points (pp), and the smallest in basis points (bp), one pp being 100 bp. A difference is the agent's prescribed rate minus the realized federal funds rate, so that a positive difference means the agent prescribes a higher rate.

The first finding concerns the economy's agents. A constant-gain mean of realized inflation, which the paper calls the learned level, has an estimated gain of 0.030 per quarter (a half-life of about 23 quarters) when the long-run survey expectation of inflation is taken as an observation of it. The survey's measurement error about the learned level then has an estimated standard deviation of 0.30 pp, larger than the survey's own standard deviation over the sample, 0.21 pp, while the learned level's is 0.32 pp; the survey lies more than 0.21 pp below the learned level in 43 of the 122 quarters, all of them before 2010:Q1 or from 2023:Q1 on. The expectations that enter the IS and Phillips equations place no weight on the learned level. Three specifications of the expectations block, an anchor, cognitive discounting, and an anchor learned from experience, each nest model-consistent expectations at one value of a parameter, and each estimates that parameter at the model-consistent value on the sample. Holding the anchored specification's weight at 0.8, with the anchor's level as measured from the survey, lowers the log-likelihood by 166.1, and holding the cognitive discount at 0.8 lowers it by 224.6. In the Monte Carlo study, on simulated economies with an ELB episode and at one true value away from model-consistent expectations for each specification, the estimates of the anchored specification's weight and of the gain at which the anchor is learned lie on both sides of their true values, and four of the five estimates of the cognitive discount lie above its true value (Section 5).

The second finding concerns the econometrician. The likelihood determines the scale of the estimated policy equation only weakly: along its inflation response, offsetting moves of the intercept and of the policy shock's standard deviation keep much of the fit to the federal funds rate, and the paper normalizes that inflation response to 3.0. Under the normalization the likelihood has more than one local optimum. The highest, the baseline estimate, is reached from two of four non-random starting values, which are optima of the likelihood found outside the random search, one of them the optimum of the likelihood profile at the normalized value, and from none of 40 random starting values. Its log-likelihood exceeds by about 120 that of the alternative estimate, the optimum reached from the data-driven starting value, which the estimator selects by a grid search over the IS, Phillips, and policy coefficients. The baseline estimate has a nearly forward-looking Phillips curve, with a weight of 0.946 on expected inflation, a Phillips slope of 0.0596 (standard error 0.0166), and a persistent shock to the policy equation, with an autocorrelation of 0.959. Scored at decision points between 2000:Q1 and 2020:Q4, the baseline estimate's forecasts of inflation and of the output gap have a higher root-mean-square error than the Federal Reserve Board staff's Tealbook projections on every line of Table 2; a two-sided test at the 5 percent level of the mean difference in squared error distinguishes it from zero on one line alone, that of four-quarter inflation four quarters ahead over its 75 consecutive quarters through 2018:Q3 (Section 4.4).

The third finding concerns the correction. An ELB episode puts structure into the filtered shocks even at the true parameters, so the null distribution of the components' test statistic, the

leave-one-out (LOO) log-score excess of the GP over a Gaussian with no structure, must be simulated with one. The 95th percentile of the Phillips component's null distribution is 0.079 on simulated economies without an ELB episode and 0.385 on economies with one. Against the null distribution with an ELB episode, the Phillips component's statistic, 0.215, lies at the 48th percentile, which is no evidence of signal. The IS component's statistic, 0.239, exceeds all 25 simulated statistics of its null distribution, whose 95th percentile is 0.145; this is evidence of signal the IS equation does not capture, and the paper states it as a limitation of that equation. The statistic is computed without the cap that each component carries where it enters a simulation; the capped components give the same two results against a null capped in the same way (Section 6). An agent trained with the components in its environment is evaluated in the estimated model without them (Section 7.2): from the 2014:Q4 state, its loss is higher than that of the volume's IFB rule at the baseline estimate and on each of 100 draws around the estimate, and the baseline agent's is lower than the IFB rule's on each of those draws. With the components in the model, at the one state and estimate evaluated, the agent trained with them has the lower loss of the two.

The fourth finding concerns the policymaker. Under the volume's loss, the baseline agent, the rule selected from the training runs on the baseline estimate, is an inertial rule: the agent's smoothing coefficient is 0.736, its response to the inflation gap 3.353 (the long-run response to inflation being one plus this coefficient; equation (4)), and its response to the output gap 3.159. The levels of its prescriptions correlate at 0.98 with those of the first-difference rule reported in the Federal Reserve Board's *Monetary Policy Report*, which shares the agent's lagged rate, and at 0.72 with those of the Taylor (1993) rule. In the volume's experiment from the 2014:Q4 state, run in the paper's reproduction of the volume's model, it holds the rate at the ELB for five quarters and lifts off in 2016:Q2, five quarters after the IFB rule and one quarter after the path of the policy that minimizes the volume's dual-mandate loss, DM1, computed in that reproduction, whose rate in 2016:Q1 stands 0.5 bp above the ELB. In the estimated model, from the same state, its loss is lower than that of the volume's IFB rule on 100 of 100 draws around the estimate. Over the 118 decision points from 1997:Q1 through 2026:Q2, its mean difference from the realized federal funds rate is +0.723 pp, with a standard error of 0.213 pp (standard errors of mean differences are by the equal-weighted cosine estimator of the long-run variance, Section 8.1), and the realized rate lies below the 10th–90th percentile band of the agent's rate across 20 draws around the estimate in 83 of the 118 quarters. The agent prescribes higher rates than the realized federal funds rate on average in every period of Table 4. The ratio of the mean difference to its standard error, in absolute value, is set against a critical value, that of a two-sided test at the 5 percent level under Student's t distribution with as many degrees of freedom as the estimator's rule gives terms for the period's length (Section 8.1). At the state of Section 2, which sets the decision quarter's IS and Phillips shocks to zero, the ratio exceeds its critical value over all quarters, over 1997:Q1 through 2008:Q4, over 2016:Q1 through 2020:Q1, over 2022:Q2 through 2026:Q2, and over 2021:Q1 through 2026:Q2; over 2021:Q1 through 2023:Q4, a period of 12 quarters, and over 2016:Q1 through 2020:Q4, which adds the first three quarters of the second ELB episode to 2016:Q1 through 2020:Q1, it does not. Section 8.1 sets the same ratios under three other treatments of the state: with those shocks at the filter's estimates, with the decision quarter's inflation withheld from the filter, and with the preceding quarters' inflation read, in addition, from the vintage of the price index available at the decision point. The ratios over all quarters, over 2022:Q2 through 2026:Q2, and over 2021:Q1 through 2026:Q2 exceed their critical values under all four treatments; that over 1997:Q1 through 2008:Q4 exceeds it under three and that over 2016:Q1 through 2020:Q1 under two; that over 2016:Q1 through 2020:Q4 under one; and that over 2021:Q1 through 2023:Q4 exceeds it under none (Table C.1 in Appendix C). Over the first ELB episode, 2009:Q1 through 2015:Q4, the mean

difference is +1.3 bp, and over the second, 2020:Q2 through 2022:Q1, it is +0.924 pp; no standard error is reported for these two means. In the quarters of the two ELB episodes the realized rate is set to the model's lower bound, below which the agent's rate does not fall, so that the difference there is zero or positive.

Two further agents condition on the data available at the time. The frozen agent, estimated and trained on data through 2015:Q4 and held fixed, has a mean difference of +0.212 pp over the 42 decision points from 2016:Q1 through 2026:Q2, with a standard error of 0.199 pp and a ratio to it below the critical value, against a mean difference of +1.269 pp for the baseline agent over the same quarters, the baseline agent's ratio exceeding that critical value under each of the four treatments. The frozen agent's prescription less the baseline agent's at the same decision point has a mean of −1.057 pp over those quarters and a standard error of 0.449 pp, a ratio of 2.35 against a critical value of 2.78, which it does not exceed. Through 2020:Q4 the mean is −0.195 pp and the ratio 1.28 against 4.30; from 2021:Q1 the mean is −1.840 pp and the ratio 7.03 against 3.18, which it exceeds. At the state of Section 2 the same pair of results holds at each division of those quarters from 2019:Q3 through 2025:Q2 and at none before, the later period resting on one term of the estimator from 2023:Q3, so that a change cannot be dated at 2020:Q4 from these ratios; the paper reports no test of the change between the two periods' means (Section 8.2). From 2021:Q1 the difference between the two agents' prescriptions is accounted for by the output gap each estimate's filter infers. No observation of the gap enters either estimate's filter after 2020:Q4, and from 2021:Q3 the baseline estimate's filtered gap exceeds the frozen estimate's by 2.76 to 4.60 percentage points of potential output. In an exchange that replaces one agent's rule, IS and Phillips parameters, and filtered gap by the other's, one at a time and in each order, replacing the baseline estimate's filtered gap by the frozen estimate's lowers the baseline agent's mean difference from 2021:Q1 by more than the two agents' mean differences stand apart, while, averaged over the orders, replacing the rule or the parameters raises it (Section 8.2). The frozen estimate's Phillips slope is 0.0172 (standard error 0.0037), against the baseline estimate's 0.0596. The two estimates also differ in the weight on expected inflation in the Phillips curve (0.693 against 0.946), in the IS weights on the lagged and the expected gap (0.883 and 0.087 against 0.771 and 0.199), in the gap response of the policy equation (1.31 against 0.84), in its smoothing and intercept, in the persistence and standard deviation of the policy shock, and in the residual components and error moments that accompany them. Exchanged between the two estimates, both in the rule's training and in the model with which its rate is computed, the Phillips slope does not account for the difference between the two agents' prescriptions: a rule trained on the baseline estimate with the frozen estimate's slope has a mean difference of +1.421 pp over the same quarters, and a rule trained on the frozen estimate with the baseline estimate's slope one of +0.117 pp, so that neither exchange moves the mean difference toward the other agent's; the first of the two rules stands at a checkpoint within the final 200 steps of its search, and whether a longer search would select a rule with another mean difference is not established (Section 8.2). The exchange of the slope leaves each estimate's filtered gap unchanged, and the paper does not separate which of the two estimates' differences, in the parameters (the slope among them) and in the error moments and the survey's wedge that their filters read, produce the divergence of the two filtered gaps. The rolling agent, re-estimated and retrained each quarter from 2006:Q1 on the data available then, has mean differences of +1.7 bp over 2006:Q1 through 2015:Q4, +0.569 pp over 2016:Q1 through 2020:Q1, and +2.105 pp over 2021:Q1 through 2023:Q4. Its mean difference from 2016:Q1 is +1.366 pp, and the ratio of that mean to its standard error exceeds its critical value; the mean difference is larger over 2021:Q1 through 2023:Q4 than over 2016:Q1 through 2020:Q1, and neither of those two periods' ratios exceeds its critical value; the paper reports no test of the change between those two periods' means. Its ratios over 2022:Q2

through 2026:Q2 and over 2021:Q1 through 2026:Q2 exceed their critical values (Section 8.3). In a counterfactual simulation with the residuals recovered at each estimate, the baseline rule's loss, with the gap measured by the baseline estimate's filter, lies below that of the path simulated under the estimated policy equation over the quarters from 2016:Q1 and not over those before. It lies below that loss in both periods under a measure of the gap from outside the model and in the model re-estimated with the Phillips equation's weight on expected inflation held at the volume's value. Against the path simulated under the frozen estimate's own policy equation, the frozen rule's loss is the lower over the quarters on which it was trained, and after them the two losses are alike at the precision of Table 9; under the outside measure of the gap the frozen rule's loss is the higher in both periods (Section 8.4).

The paper relates to four literatures. In real-time policy evaluation, Orphanides shows that a rule evaluated on revised data misstates what a policymaker could have done with the data available at the time; for this reason every observation in the agent's state is a real-time vintage, with the three exceptions that Section 2 states: the neutral-rate path, the price index, which is read from a later vintage and through the decision quarter, and the term premium.[4] In optimal control at the Federal Reserve, the volume's loss-minimizing paths are of the kind the Federal Reserve Board's staff compute in FRB/US, the Board's large-scale model of the U.S. economy, and the simple rules reported in the Board's *Monetary Policy Report*, the Taylor (1993) rule among them, are the comparators of Section 8.1.[5] On policy under model uncertainty, Brainard's attenuation, the robust-control program of Hansen and Sargent, and the parameter-uncertainty results of Wieland and of Edge, Laubach, and Williams describe how a policymaker acts who knows that the model is uncertain, and Sargent and Primiceri model the policymaker's own learning in the history of U.S. inflation; the agent here is trained on draws of four coefficients around the estimate, a training that varies those coefficients and no others.[6] In monetary economics, Hinterlang and Tänzer train a policy by reinforcement learning against a neural-network transition model estimated on U.S. data, and Chen and coauthors use deep reinforcement learning as a model of bounded rationality inside a general-equilibrium model; the method here, a policy trained inside an estimated model of its environment, follows the Dyna architecture of Sutton and the world-model program of Ha and Schmidhuber.[7] The paper differs from this work in four respects: the data are real-time except for

[4] Athanasios Orphanides, "Monetary Policy Rules Based on Real-Time Data," *American Economic Review* 91, no. 4 (2001): 964–985.

[5] Flint Brayton, Thomas Laubach, and David Reifschneider, "Optimal-Control Monetary Policy in the FRB/US Model," FEDS Notes (Washington, DC: Board of Governors of the Federal Reserve System, November 21, 2014); John B. Taylor, "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy* 39 (1993): 195–214; John B. Taylor and John C. Williams, "Simple and Robust Rules for Monetary Policy," in *Handbook of Monetary Economics*, vol. 3B, ed. Benjamin M. Friedman and Michael Woodford (Amsterdam: Elsevier, 2010), 829–859.

[6] William C. Brainard, "Uncertainty and the Effectiveness of Policy," *American Economic Review* 57, no. 2 (1967): 411–425; Lars Peter Hansen and Thomas J. Sargent, *Robustness* (Princeton, NJ: Princeton University Press, 2008); Volker Wieland, "Monetary Policy, Parameter Uncertainty and Optimal Learning," *Journal of Monetary Economics* 46, no. 1 (2000): 199–228; Rochelle M. Edge, Thomas Laubach, and John C. Williams, "Welfare-Maximizing Monetary Policy under Parameter Uncertainty," *Journal of Applied Econometrics* 25, no. 1 (2010): 129–143; Thomas J. Sargent, *The Conquest of American Inflation* (Princeton, NJ: Princeton University Press, 1999); Giorgio E. Primiceri, "Why Inflation Rose and Fell: Policy-Makers' Beliefs and U.S. Postwar Stabilization Policy," *Quarterly Journal of Economics* 121, no. 3 (2006): 867–901.

[7] Natascha Hinterlang and Alina Tänzer, "Optimal Monetary Policy Using Reinforcement Learning," Discussion Paper 51/2021 (Frankfurt am Main: Deutsche Bundesbank, 2021); Mingli Chen, Rama Cont, Andreas Joseph, Michael Kumhof, Xinlei Pan, Wei Xiong, and Xuan Zhou, "Deep Reinforcement Learning in a Monetary Model," Staff Working Paper 1,142 (London: Bank of England, 2025); Richard S. Sutton, "Dyna, an Integrated Architecture for Learning, Planning, and Reacting," *ACM SIGART Bulletin* 2, no. 4 (1991): 160–163; David Ha and Jürgen Schmidhuber, "Recurrent World Models Facilitate Policy Evolution," in *Advances in Neural Information Processing Systems* 31 (2018): 2450–2462.

three inputs (Section 2); the estimators' behavior on simulated economies with an ELB episode is measured in a Monte Carlo study; the expectations block is tested against learning and behavioral alternatives; and the rule is trained across draws around the estimate. The paper also measures how a rule trained with a learned correction of the model's residuals performs in the estimated model without the correction, across draws around the estimate (Section 7.2).

The paper proceeds as follows. Section 2 describes the real-time data and defines the realized federal funds rate. Section 3 states the volume's concept of optimal policy, the distinction between an optimal path and an optimal rule, the four learners, and the training method, and Section 3.1 runs the volume's experiment on the baseline agent. Section 4 states the model, its estimator, the Monte Carlo study of the estimators, the multi-start search that determines the baseline estimate, and the model's forecasts against the Federal Reserve Board staff's Tealbook projections. Section 5 tests the expectations block against the three alternatives. Section 6 fits the residual components and constructs their null distribution with an ELB episode. Section 7 trains the agent, compares it with an agent trained with the components, and selects the baseline agent among training runs. Section 8 compares the baseline agent with the realized federal funds rate, repeats the comparison for the frozen and rolling agents, and simulates the outcomes each rule would have produced with the shocks recovered from the data. Section 9 reports the sensitivity to the loss weights and to the training run and states the limitations, and Section 10 concludes. Appendix A lists the conditions of the Monte Carlo study and the parameter bounds, Appendix B takes up five questions on the estimates and the agents, and Appendix C reports, by period, the baseline agent's mean difference under four treatments of the state, the differences between two agents' prescriptions at the same decision points, and the losses under a measure of the output gap from outside the model.

2 The Real-Time Data

With three exceptions described below, the neutral-rate path, core personal consumption expenditures (PCE) inflation, and the term premium, every observation the model reads at a decision point is one the FOMC could have read before its meeting. A decision point is the final scheduled FOMC meeting of a quarter. The data hold 146 quarterly decision points from 1990:Q1 through 2026:Q2, each assembled from the latest vintage of each series released strictly before the meeting. They come from 30 series, each held with every release as a separate vintage: the Archival Federal Reserve Economic Data (ALFRED) vintages of the core PCE price index, the unemployment rate, the effective federal funds rate, and the two-year Treasury yield, among others; the Real-Time Data Set for Macroeconomists (RTDSM) of the Federal Reserve Bank of Philadelphia and, from the same source, the Federal Reserve Board staff's Tealbook projections of the output gap and core inflation; the Survey of Professional Forecasters (SPF) at the four-quarter and ten-year horizons; the Holston–Laubach–Williams (HLW) neutral-rate estimates;[8] and the Adrian–Crump–Moench (ACM) term premiums.[9] The vintages are those available on 5 September 2026.

The realized federal funds rate at a decision point is the quarter-to-date average of the effective federal funds rate published before the meeting. It is the outcome of the decisions taken before

[8] Kathryn Holston, Thomas Laubach, and John C. Williams, "Measuring the Natural Rate of Interest: International Trends and Determinants," *Journal of International Economics* 108, supplement 1 (2017): S59–S75; the estimates in use are those of the model of Kathryn Holston, Thomas Laubach, and John C. Williams, "Measuring the Natural Rate of Interest after COVID-19," Federal Reserve Bank of New York Staff Reports, no. 1063 (June 2023).

[9] Tobias Adrian, Richard K. Crump, and Emanuel Moench, "Pricing the Term Structure with Linear Regressions," *Journal of Financial Economics* 110, no. 1 (2013): 110–138.

the decision point and does not include the decision taken at that meeting, so a change of the target range at a quarter's final meeting first appears in the realized rate of the following quarter. A quarter is at the ELB where the lower end of the target range for the federal funds rate is zero at the decision point. This gives 36 quarters at the ELB in two ELB episodes, 2009:Q1 through 2015:Q4 and 2020:Q2 through 2022:Q1. In quarters at the ELB the realized federal funds rate is set to the model's lower bound of 0.125 percent, and the observation of the policy rate is censored at that bound. At the decision point of 2020:Q1, the meeting of 29 January 2020 (the quarter's later scheduled meeting was not held), the latest release of the effective federal funds rate before the meeting, of 2 January 2020, holds no observation of the quarter, and the realized federal funds rate there is the midpoint of the target range of 1.50 to 1.75 percent in force at the meeting. The agent's prescribed rate at a decision point is the rate the agent prescribes for the decision quarter at that point's real-time state: the first-quarter rate of the model's solution from that state, in which the lagged output gap is the filter's estimate of the previous quarter's gap given the observations through the decision quarter, and the IS and Phillips shocks of the decision quarter are zero. Section 8.1 also reports the comparison with those two shocks at the filter's estimates and the comparison with the decision quarter's inflation withheld from the filter.

The estimation window is 1996:Q1 through 2026:Q2, 122 quarters with 836 observation rows, an observation row being one observed quantity in one quarter: the staff's output-gap series contributes three rows a quarter, its estimate of the previous quarter, its nowcast, and its projection of the next quarter, and the row of the policy rate is absent in the quarters at the ELB. The agents are trained from, and compared at, the 118 decision points from 1997:Q1 through 2026:Q2. The staff projections are released with a five-year lag, and at a decision point they are admitted only from a vintage dated within that point's own quarter; the final 22 quarters therefore carry no staff observations, and the sample there consists of the observations of inflation, the federal funds rate, the two-year rate, and the survey. The neutral-rate path is the ex post one-sided HLW estimate over the window from its latest vintage. It is not real-time: the neutral rate on which expectations are formed at a decision point is an estimate made with later data, and the model's solution from a decision point reads the path's later values, except for the rolling agent of Section 8.3, which reads the real-time HLW vintages from 2016:Q1 on; Section 9 states the limitation. Core PCE inflation is not real-time either. The inflation observations of the likelihood, the four lags of inflation in every state, and the inflation the published rules of Section 8.1 read are computed from the price index's vintage of 26 August 2026, and the filter at a decision point reads the decision quarter's own inflation, a quarter whose three months of the price index had been published at none of the 118 meetings; the lagged gap of the state is inferred with that observation. At the 103 decision points whose real-time vintages cover the four preceding quarters, the state's four-quarter inflation differs from the value those vintages give by 0.28 pp root mean square, with a mean of +0.13 pp (the later vintage's value less the earlier), and most at 2009:Q2, by 0.89 pp; Section 8.1 reports the prescriptions at states that carry the vintages' inflation. The third exception, the term premium, is read from one release and is described with the filter's moments below.

The estimator's filter (Section 4.2) also reads moments of the data, each computed for one date over the quarters known then, with the realized inflation rate and the term premium read from later releases as stated below: 2026:Q2 for the baseline estimate and the agents trained on it, 2015:Q4 for the frozen agent of Section 8.2, and each decision point for the rolling agent. The first set comprises the means, standard deviations, and first-order autocorrelations of the revision, nowcast, and projection errors of the staff's output gap and the correlations among them. Each error is measured against the staff's estimate of the same quarter at the decision point five quarters after that of the nowcast or the projection, or, for the revision, five quarters after the quarter estimated,

and the set is computed over the quarters from 1996:Q1 that lie five or more quarters before the date for which it is computed and for which the data hold both estimates (112 quarters at 2026:Q2 for the revision and the nowcast error, and 111 for the projection error). The second is the mean, standard deviation, and first-order autocorrelation of the error of the staff's projection of core PCE inflation over the following four quarters against the four-quarter rate then realized, the realized rate read from the price index's vintage of 26 August 2026 and not, as in Section 4.4, from the vintage five quarters on. It is computed over the quarters from 2000:Q1 that lie five or more quarters before that date and for which the data hold the projection and the vintage five quarters on already holds the realized rate, so that a quarter enters the moments known at a date only once its realized rate has been published (87 quarters at 2026:Q2; the vintage five quarters on lacks the realized rate of 2018:Q4, as in Section 4.4). The third is not an error: it is the mean, standard deviation, and first-order autocorrelation of the term premium through which the two-year rate is measured, computed from the quarterly means of the ACM two-year premium over the quarters completed before that date (121 quarters at 2026:Q2, 1996:Q1 through 2026:Q1, from the release retrieved on 5 September 2026). In the model the two-year rate is observed as the mean expected policy rate over the next eight quarters plus the term premium, a first-order autoregressive error whose mean, standard deviation, and first-order autocorrelation are this third set, and the model has no separate convenience component, so a movement of the two-year Treasury's convenience yield is absorbed in that error. Krishnamurthy and Chumbo find that by 2025 the convenience yield on short-term dollar safe assets was about where it had stood in 2019, while that on ten-year Treasurys had fallen over the same years; this paper does not measure the convenience yield at the two-year maturity.[10] One further set, the error of the real-time HLW estimate against its ex post value (32 vintages at 2026:Q2, the first from 2016), enters only the simulated economies of the Monte Carlo study of Section 4.2 and no estimate on the real-time data.

Because the staff's projections are released with a five-year lag, the data hold no path made after the decision point of 2020:Q4, and in the data every later decision point carries the paths of that one, the gap's through 2023:Q4 and the inflation projection's through 2022:Q4. A carried path enters no observation row, but the moments are computed without regard to whether a quarter's path is a carried one, so the first two sets of moments known at 2026:Q2 include quarters whose errors are computed from a carried path. Of the 112 quarters of the gap's revision, in 4 (2019:Q4 through 2020:Q3) the estimate five quarters later is read from a carried path, and in 13 (2020:Q4 through 2023:Q4) the first estimate, which is made a quarter later, is read from a carried path as well; the revision is zero in those 13 quarters and in 2020:Q3, since in each of these 14 quarters both estimates are read from the path of 2020:Q4. Each of the four decision points of 2021 reads the inflation projection of 2020:Q4. Over the quarters whose errors are computed from no carried path, 95 for each of the gap's three errors and 83 for the inflation projection's, the standard deviation of the gap's revision is 0.64 percentage points of potential output, against the 0.59 the filter reads, and that of the gap's projection error is 0.88 against 1.09 percentage points of potential output, with a first-order autocorrelation of 0.66 against 0.28. Over those quarters the standard deviation of the inflation projection's error is 0.68 pp against 0.97 pp, with a first-order autocorrelation of 0.70 against 0.85. The moments known at 2015:Q4 include no such quarter; those of the rolling windows include them in the 22 windows from 2021:Q1 for the gap's errors and in the 17 from 2022:Q2 for the inflation projection's. The simulated economies of the Monte Carlo study of Section 4.2 and of the null distribution of Section 6 draw their measurement errors at the moments known at 2026:Q2. Section 9 states the limitation.

[10] Arvind Krishnamurthy and Miguel Chumbo, "Disappearing Dollar Convenience?" NBER Working Paper 35742 (Cambridge, MA: National Bureau of Economic Research, 2026).

3 Optimal Policy, Rules, and Learners

The optimal policy of Adrian, Laxton, and Obstfeld's volume minimizes, at each decision, the loss

$$L_t = \sum_{i \geq 1} \left[\alpha \left(\pi^4_{t+i} - \pi^* \right)^2 + \beta \, y^2_{t+i} + \gamma \left(i_{t+i} - i_{t+i-1} \right)^2 \right], \tag{1}$$

with π^4 four-quarter core PCE inflation, π^* the inflation target of 2 percent, y the output gap in percent of potential output, i the policy rate, and weights α and β of one and γ of one half. The sum runs from one quarter ahead, the lower limit the volume's equation (12.3) prints; in training and in simulation the sum is truncated at a finite horizon, which Sections 3.1 and 7.1 state. In the volume's argument, the quadratic form makes a large deviation weigh more than proportionally, so that the loss-minimizing response to a shock is stronger the larger the shock and the nearer the economy stands to the ELB, from which recovery is slow; the volume's simulations show the loss-minimizing policy holding the rate at the ELB for an extended period after a negative shock where a linear rule lifts off early. The volume contrasts this policy with two rules, the Taylor rule and the IFB rule, which respond linearly to a small set of variables and carry no such risk avoidance.

Two distinctions organize what follows. The first is between an optimal *path* and an optimal *rule*. The volume's DM1 policy, the label the volume gives the policy that minimizes its dual-mandate loss at these weights, is a path: the sequence of rates that minimizes equation (1) from the current state under the model, re-optimized as the state moves. The baseline agent of this paper is a rule: an IFB rule with four coefficients chosen to minimize the same loss on average over the real-time states of the sample and over draws around the estimate, inside the estimated model. A rule can be announced and committed to, and its prescriptions can be compared with realized decisions state by state; the path is the benchmark against which Section 3.1 measures the rule, in the volume's own experiment.

The second distinction is among four senses in which something is learned, drawn in Figure 1. The economy's agents learn in the sense of the expectations literature: their expectation of inflation may be the model's forecast, a weighted average of that forecast and an anchor, a discounted forecast, or a level they update from experience (Section 5). The econometrician learns the model's equations from the real-time data in the statistical sense; the paper studies the estimator in a Monte Carlo study on simulated economies (Section 4.2) and characterizes the likelihood its search explores (Section 4.3). The residual correction learns nonparametrically what the equations miss, and the paper tests whether what it learns is signal (Section 6). The policymaker learns the rule by optimization against the estimated model, and because that model is itself estimated, the rule can exploit the errors of the first three learners, in the sense of fitting itself to the model's errors; Section 7.2 measures, for the residual correction, the loss of a rule trained with the correction in the estimated model without it and with it. The paper's conclusions about optimal policy are conditional on all four learners, and Section 9 and the Conclusion state what would change them.

The policymaker's learning is what this paper calls reinforcement learning, and the method is the following. The policy, or agent, is parametric: the IFB rule of equation (4) below with the shock u_t set to zero, so that four coefficients, (ρ, c_π, c_y, c_0), describe it. An episode is one simulation of the estimated model under the rule, from a real-time state of the sample, at one draw of the IS and Phillips parameters (a_{lag}, a_{lead}, λ_1, and κ) around the estimate, and the episode's loss is the loss (1) along the simulated path. Training is gradient descent on the expected loss over such episodes: at each step the coefficients move against the gradient of the mean loss over a batch of episodes, and the gradient is taken through the model's equilibrium, which is solved exactly at every quarter.

The solution is piecewise smooth in the coefficients, and the gradient is that of the solution with the set of quarters at the ELB held fixed. Methods of this kind are policy-gradient methods with a known, differentiable transition, sometimes called model-based or analytic policy gradients; the method is reinforcement learning in that sense, a parametric policy improved by the loss its own prescriptions produce in simulated episodes. The data enter training only through the estimated model and the starting states: the outcomes realized after a starting state do not enter any episode's loss. A training run is one sequence of such steps, each run with its own random-number stream for the draws and the starting states; a checkpoint is the rule's coefficients at one step of a run; the validation set is a fixed set of starting states and of draws of the IS and Phillips parameters, the draws entering no training episode, and checkpoints are compared on it to select one; and the test set is a second such set, which the selection of the reported checkpoint does not use. Section 7 states the horizon, the starting states, and the selection.

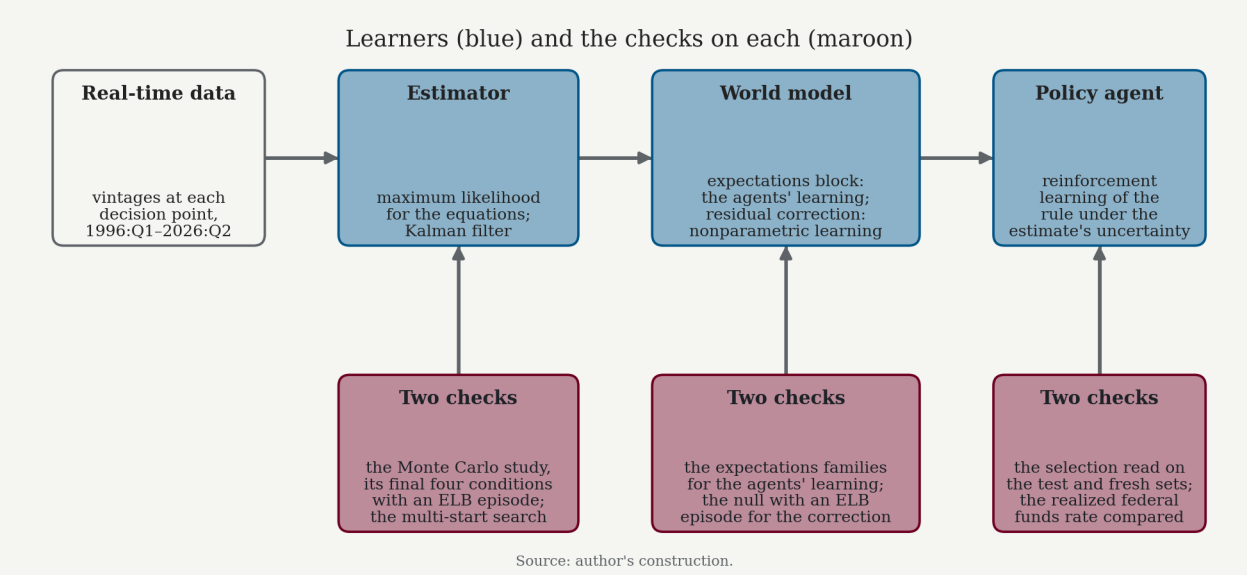


Figure 1. The four learners of optimal monetary policy when the economy must be learned. The top row runs from the real-time data through the estimator, the econometrician's learner, which learns the model's equations, and the world model, which holds two learners, the economy's agents, whose learning enters through its expectations block, and the residual correction, which learns what the equations miss, to the policy agent, which learns the rule by reinforcement learning on draws around the estimate. The bottom row holds the checks on each box: for the estimator, the Monte Carlo study, whose final four conditions carry an ELB episode, and the multi-start search; for the world model, the expectations families of Section 5, which test the economy's agents' learning, and the null distribution of the residual correction's test statistic on simulated economies with an ELB episode; for the agent, the test and fresh sets on which the selection of a checkpoint across training runs is read (Section 7.3), and the comparison with the realized federal funds rate. Blue boxes hold the learners and maroon boxes the checks. Source: author's construction.

3.1 The volume's experiment on the learned rule

The volume's Figure 12.5 starts the economy at its 2014:Q4 state, with a gap of -2 percent of potential output, four-quarter inflation of 1.5 percent, and the policy rate at the ELB, and runs the IFB rule and the loss-minimizing DM1 path forward with no shocks.

Figure 2 shows the same experiment run in the paper's reproduction of the volume's model, in the volume's calibration with this paper's IS coefficients (Table 1, final column), on the baseline agent, the learned rule selected in Section 7.3, with DM1 re-optimized each quarter. Each policy's loss is the loss (1) over the sixteen quarters 2015:Q1 through 2018:Q4 plus the loss along the policy's perfect-foresight path from the state of 2018:Q4. The IFB rule lifts off in 2015:Q1; the DM1 path holds the rate at the ELB through 2015:Q4, and its rate in 2016:Q1 stands 0.5 bp above the ELB, where the volume has its DM1 path at the ELB until the middle of 2016; the learned rule holds the rate at the ELB for five quarters and lifts off in 2016:Q2. The losses are 14.44 for the IFB rule, 9.95 for DM1, and 10.83 for the learned rule. A rule trained inside a model estimated on the real-time data, and evaluated here in the reproduction of the volume's model, lifts off one quarter after the DM1 path computed there, and it holds the rate at the ELB where the IFB rule lifts off at once. The learned rule's path overshoots: four-quarter inflation ends at 2.22 percent, above the target, and the output gap is positive from 2016:Q4 on. The rule's loss lies between the path's and the IFB rule's.

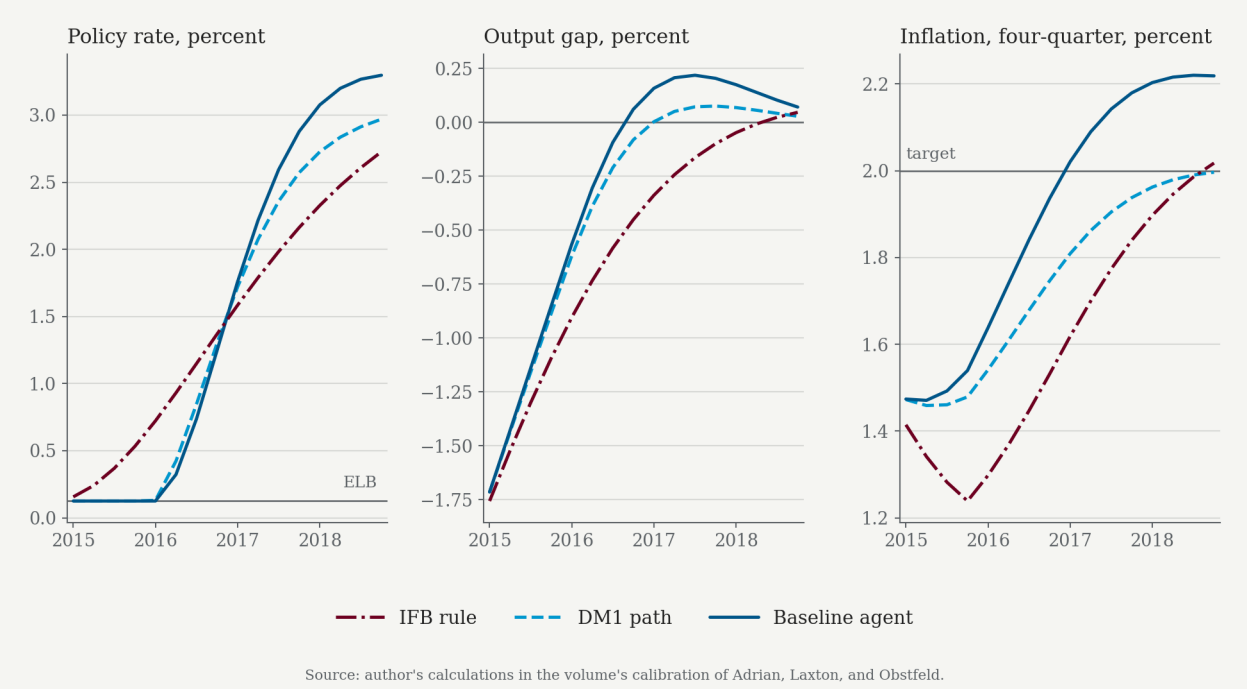


Figure 2. The volume's experiment from the 2014:Q4 state run on the learned rule. Each panel carries one variable over 2015:Q1 through 2018:Q4: the policy rate in percent, the output gap in percent of potential output, and four-quarter core PCE inflation in percent, under the IFB rule at the volume's coefficients (maroon, dash-dot), the loss-minimizing DM1 path re-optimized each quarter (bright blue, dashed), and the baseline agent (deep blue, solid), in the volume's calibration with no shocks; gray lines mark the ELB, zero, and the 2 percent target. Liftoff: the IFB rule in 2015:Q1, DM1 in 2016:Q1, the learned rule in 2016:Q2. Losses: 14.44, 9.95, and 10.83, each the loss (1) over the sixteen quarters shown plus the loss along the policy's perfect-foresight path from the final state. Source: author's calculations in the volume's calibration (Table 1, final column).

4 The World Model

The world model is the estimated model of the economy inside which the agent is trained: the semi-structural economy of the volume's chapter on the United States, in a reduced form, estimated on the real-time data of Section 2 by maximum likelihood. This section states its equations, its estimator and the Monte Carlo study of the estimator, the search that locates the baseline estimate, and the estimate's forecasts against the Federal Reserve Board staff's projections.

4.1 Blocks and expectations

The economy is this paper's reduced form of the semi-structural model of the volume's chapter 12, the chapter on the United States, written here as it is solved. The chapter prints the model's Phillips curve and its policy rules and describes its output-gap equation in words alone; that equation, its calibration, and the rest of the model are in the working paper the chapter builds on, which sets out a stripped-down Global Projection Model.[11] The IS block is

$$y_t = a_{\text{lag}}\, y_{t-1} + a_{\text{lead}}\, E_t y_{t+1} - a_r \left(i_t - E_t \pi^4_{t+4} - \bar{r}_t\right) + \varepsilon^y_t, \tag{2}$$

with $\bar{r}_t$ the neutral real rate of Section 2; the Phillips block is

$$\pi_t = \lambda_1\, E_t \pi^4_{t+4} + (1-\lambda_1)\, \pi^4_{t-1} + \kappa\, y_{t-1} + \varepsilon^\pi_t, \tag{3}$$

with π_t quarterly core PCE inflation at an annual rate and π^4_t its four-quarter average, so that λ_1 is the weight on the four-quarter-ahead expectation and $1-\lambda_1$ the weight on lagged four-quarter inflation, and the block is purely forward-looking at $\lambda_1 = 1$. The policy equation is the IFB rule

$$\begin{aligned} i_t &= \max\left\{\underline{i},\ \rho\, i_{t-1} + (1-\rho)\left[\bar{r}_t + E_t \pi^4_{t+3} + c_\pi (E_t \pi^4_{t+3} - \pi^*) + c_y\, y_t + c_0\right] + u_t\right\}, \\ u_t &= \phi_u u_{t-1} + \eta_t, \end{aligned} \tag{4}$$

with smoothing ρ, responses c_π and c_y to the model's own forecast of four-quarter inflation three quarters ahead, $E_t \pi^4_{t+3}$, and to the contemporaneous gap, an intercept c_0 in percentage points, a persistent shock u_t with autocorrelation ϕ_u and innovation standard deviation σ_η, and the lower bound $\underline{i}$ at 0.125 percent, the model's ELB. The coefficient c_π multiplies the forecast's deviation from target on top of the one-for-one term, so that the rule's long-run response to the inflation forecast is $1 + c_\pi$. Expectations are model-consistent, as in the model of the volume's chapter 12. The volume's chapter 9, on India, sets the inflation expectation from the model-consistent solution and the inflation target, as the anchored family of Section 5 does; the cognitive-discounting and learned-anchor families are this paper's, and each of the three families nests the model-consistent case at one parameter value. At each quarter the perfect-foresight equilibrium over a forward horizon of 48 quarters is solved with the ELB as a complementarity condition, the structural shocks at zero over the horizon except where a section feeds them, the gap at zero and inflation at the target beyond it, and the neutral-rate path over the horizon read from the path of Section 2 from that

[11] Alichi et al., "Avoiding Dark Corners," to which the chapter's own footnote refers for the model's equations. Its output-gap equation responds to a long real rate, a weighted average of the current short real rate and its expected averages over the next four, twelve, and twenty quarters, carries a zero coefficient on the current short real rate, and carries a bank-lending-tightening term. The IS block here responds to the current short real rate alone and carries no credit term; both are reductions made for estimation on the real-time data, and the coefficients reported for the IS block in the final column of Table 1 are this paper's calibration of that reduced block, not values the chapter prints.

quarter forward and held at its final value, 1.01 percent, beyond the sample. The agent's rule of Section 7 takes the form of equation (4) with u_t set to zero, and there ρ, c_π, c_y, and c_0 denote the agent's coefficients.

4.2 Estimator and Monte Carlo study

The parameters are estimated by maximum likelihood over the 836 observation rows of Section 2, drawn from six series: the staff's gap series, three rows a quarter (300 rows); the staff's projection of inflation (84); realized inflation (122); the federal funds rate, absent at the ELB (86); the two-year Treasury yield (122); and the SPF four-quarter inflation expectation (122). The likelihood is the Gaussian prediction-error decomposition of a Kalman filter that is linear conditional on whether the ELB binds at each quarter of the equilibrium; the equilibrium is re-solved at every quarter, and a penalty is added for any quarter at which no equilibrium exists at the trial parameters. The objective is the negative log-likelihood per observation row, so that a lower objective is a better fit and a difference in the objective times 836 is a difference in log-likelihood. Standard errors come from the outer product of the finite-difference Jacobian of the standardized innovations. The standard deviations of the IS and Phillips shocks are not estimated. They are held at 0.40 percent of potential output for the IS shock and 0.25 pp of annualized quarterly inflation for the Phillips shock, values this paper assumes, the volume stating none; they are free only in condition O12 of the Monte Carlo study. Every difference in log-likelihood and every standard error in the paper is conditional on them. The filter's standardized innovations have a standard deviation of 1.16 in the quarters off the ELB and 1.47 at it, where a likelihood whose scale matched the data would give one; a likelihood that understates the innovations' variance overstates differences in log-likelihood and understates standard errors, in proportion to the understatement where it is uniform. Over 1997:Q1 through 2026:Q2 the filtered Phillips shocks have a standard deviation of 0.34 pp at the baseline estimate and 0.75 pp at the estimate on the sample through 2015:Q4 (Section 8.2), and the filtered IS shocks 0.46 and 0.50 percent of potential output. The measurement error of the survey row, the Survey of Professional Forecasters' four-quarter inflation expectation in core PCE units, is estimated with a standard deviation of 0.798 pp (Table 1), against a standard deviation of 0.42 pp of the 122 survey observations themselves. The model has not been estimated with the two standard deviations free on the sample or at other values (Section 9). The estimator's behavior is measured in a Monte Carlo study on economies simulated at known parameters: 18 conditions and 75 estimation runs. The final four conditions carry an ELB episode, the final three of them with the sample's pattern of missing observation rows as well; condition O7 carries a scripted ELB episode under which the simulated economy has no equilibrium, and it is not estimated. Under an ELB episode with the staff's gap rows present in every quarter (condition O15), the estimator of the anchored-expectations specification of Section 5 returns anchoring weights from 0.769 to 0.809 on five simulated samples against a true value of 0.80, with standard errors of 0.037 to 0.046; with the sample's missing-row pattern as well (O16), it returns weights from 0.737 to 0.805 with standard errors of 0.039 to 0.048. In the 20 samples of the four conditions with an ELB episode, the estimates of λ_1 range from 0.631 to 0.760 against a true value of 0.70, and those of κ from 0.078 to 0.133 against 0.10. The sum of the two IS weights lies at the estimator's cap of 0.97 (Section 4.3) in 18 of those samples, against a true sum of 0.95, and the inflation response of the policy equation, whose true value is 0.91, has estimates between -0.46 and 2.64, with standard errors of 0.45 to 1.87, and its gap response, whose true value is 0.21, has estimates between -0.50 and 0.51, with standard errors of 0.27 to 0.68: the study does not show the recovery of these three quantities. Appendix A lists the Monte Carlo conditions.

4.3 The multi-start search and the baseline estimate

The likelihood has more than one local optimum, and along one direction, the scale of the policy equation, it varies less than along the others. With the inflation response c_π held at each point of the grid 1, 2, 3, 4, 6, 8, 12, and 16, and every other free parameter re-optimized from two starting values, the objective is 0.1899, 0.1549, 0.1445, 0.1511, 0.1416, 0.1349, 0.1450, and 0.1543, respectively, a range of 46 log-likelihood units (Figure 3, left panel). The profile has local minima at 3.0 and 8, and the log-likelihood at 8 is higher than at 3.0 by about 8 log-likelihood units. The objective falls from $c_\pi = 4$ to 8, with the policy shock's persistence ϕ_u at its upper bound of 0.97 from $c_\pi = 6$ on, and it rises beyond 8, with the intercept c_0 at its upper bound of 8 pp from $c_\pi = 12$ on. Along the profile the intercept and the policy shock's standard deviation move with c_π (at 8 they stand at 6.94 pp and 1.60 pp), so that the likelihood lies on a ridge on which c_π, c_0, and σ_η move together, and it determines the scale of the policy equation only weakly. The paper normalizes the inflation response of the estimated policy equation to $c_\pi = 3.0$, the only local minimum of the profile at which no parameter lies at a bound: at 6, 8, 12, and 16, ϕ_u lies at its bound. The volume's value of 0.91 lies below the grid's lowest point, at which the objective is 0.1899. Under the normalization, the intercept sets the policy equation's steady state: with u_t at zero, the gap at zero, and inflation at the target, equation (4) gives $i = \bar{r} + \pi^* + c_0$, so that the intercept is the distance between the equation's steady-state rate and $\bar{r} + \pi^*$, 3.00 pp at the baseline estimate. The persistent policy disturbance offsets this: its mean over the 118 decision points from 1997:Q1 is -2.06 pp, against the intercept's term in equation (4), $(1-\rho)c_0$, of 1.03 pp. Every estimate of the world model in what follows, on the full sample, on the sample through 2015:Q4, and in each rolling window, holds c_π at 3.0 and carries no standard error for it, and Section 9 returns to the normalization among the limitations.

With c_π normalized, the search runs from 40 random starts drawn uniformly within the parameter bounds (Table A.2) and from four non-random starting values, optima of the objective found outside the random search, each with c_π set to 3.0; two of them are the profile's optimum at $c_\pi = 3.0$ and the optimum with the lowest objective of a search in which c_π is estimated. The deepest optimum, at an objective of 0.1445, is reached from those two starting values and from none of the random starts, the deepest of which stops at 0.1583. A group of 20 starts (18 random starts and the two other non-random starting values) stops at 0.1583 to 0.1612, where the log-likelihood is lower than at the deepest optimum by 11.5 to 13.9 log-likelihood units, with a Phillips block of the same kind, λ_1 from 0.91 to 0.95 and a slope from 0.053 to 0.062. The remaining 22 random starts stop at 0.1993 to 0.3169, with λ_1 from 0.62 to 0.95 and, in nine of them, a Phillips slope below 0.001, which in seven of these is zero. Of the group of 20 starts, six were optimized again under the wider bounds of Table A.2, with the upper bound on λ_1 at 0.995 in place of the table's upper bound of one: four random starts, each from its optimum under the search's bounds, and two non-random starting values, each from the starting value itself. The four random starts keep their objectives to within 0.0001. Of the two non-random starting values, one stops at 0.1541, below its objective in the search, and one at 0.1670, above it, with λ_1 within 0.0001 of the upper bound of 0.995; optimized from the starting value and not from the search's optimum, it stops at another point, with the higher objective. The random starts alone do not locate the deepest optimum; this, with the profile, is the evidence the paper offers on the likelihood's shape. The alternative estimate is not one of the search's starting values; Figure 3 draws it as a reference.

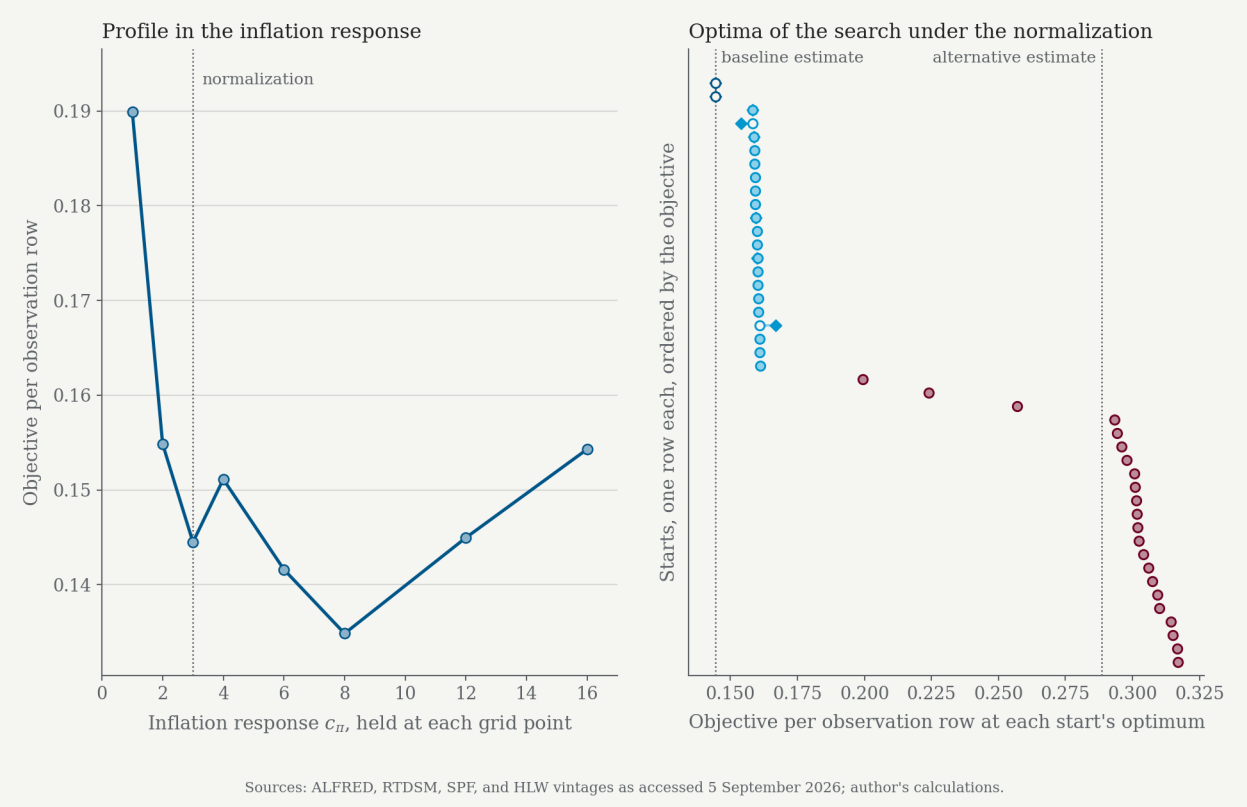


Figure 3. The profile in the inflation response and the optima of the search, 1996:Q1–2026:Q2. Left panel: the objective, the negative log-likelihood per observation row (lower is better), with c_π held at each point of the grid and every other free parameter re-optimized, the better of two starting values at each point; the dotted line marks the normalized value, 3.0. Right panel: the 44 starts of the search with c_π normalized, ordered by the objective at each start's optimum, one row each, the vertical position the start's rank and not a value; deep blue marks the deepest optimum, blue the group of 20 starts, and maroon the remaining starts, with open markers for the four non-random starting values and filled markers for the 40 random starts; where a start was also optimized again under bounds on λ_1 and c_0 wider than the search's, a second marker, joined to the first, gives that objective. Dotted lines mark the alternative estimate (0.2885) and the baseline estimate (0.1445). Sources: ALFRED, RTDSM, SPF, and HLW vintages as accessed 5 September 2026; author's calculations.

The baseline estimate is the deepest optimum, re-optimized with the bounds on λ_1 and c_0 set wider than in the search (Table A.2); its objective, 0.1445, equals the search's deepest at the precision shown. In the search the deepest optimum's intercept stands at the search's upper bound of 3 pp; under the wider bounds the baseline estimate's lies 0.08 bp above it. No parameter lies at a bound, and λ_1 stands at 0.946 (standard error 0.041), short of its structural limit of one. The estimate lies on one joint constraint: the two weights of the IS block lie on the cap $a_{\text{lag}} + a_{\text{lead}} \leq 0.97$ that the estimator imposes for a stationary solution, as do the weights of the frozen estimate of Section 8.2 and of each rolling window's estimate of Section 8.3. The weights' standard errors are computed at that point as if the two weights were free, so that the two standard errors differ although the constraint makes the weights' moves equal and opposite. Holding the cap at each point of a grid, with every other free parameter re-optimized from the baseline and the alternative estimates as starting values and c_π at its normalized value, gives the objective's profile in the cap. The log-likelihood rises at every point of the grid as the cap rises: at a cap of 0.995 it lies 18.9

log-likelihood units above its value at 0.97, and at each point the two weights stand on that point's cap. The sum of the two weights is therefore not an estimate: the cap of 0.97 is part of the model's specification, and the profile is reported in its place. Along the profile the Phillips equation moves within two of the baseline estimate's standard errors, its slope from 0.0596 to 0.0372 and λ_1 from 0.946 to 0.970 at a cap of 0.995. Figure 3 shows the 44 optima, and Table 1 reports the baseline estimate beside the alternative estimate, which is the optimum reached from the data-driven starting value (the starting value the estimator selects by a grid search over the IS, Phillips, and policy coefficients, with the remaining coefficients from least-squares regressions on the observed series), and beside the volume's calibration.

The baseline and alternative estimates describe the federal funds rate and inflation differently. The alternative estimate carries a Phillips block with a lower weight on expected inflation (λ_1 of 0.662) with a slope of 0.0314, and a rule with smoothing of 0.839 and a policy shock with persistence ϕ_u of 0.790 and innovation standard deviation σ_η of 0.496. The baseline estimate carries a nearly forward-looking Phillips block (λ_1 of 0.946) with a slope of 0.0596 (standard error 0.0166), a gap response of 0.84, and a policy shock with ϕ_u of 0.959 and σ_η of 0.947. Its log-likelihood exceeds the alternative estimate's by about 120. Under the baseline estimate the equilibrium exists at every quarter of the sample, both ELB episodes included, and the standardized innovations over the quarters at the ELB have 1.28 times the standard deviation of those elsewhere.

Table 1. The baseline estimate of the world model, 1996:Q1–2026:Q2. Alternative is the alternative estimate, the optimum reached from the data-driven starting value; Baseline is the baseline estimate, the deepest optimum of the search of Section 4.3; in both, the inflation response c_π is normalized to 3.0 and carries no standard error. SE is the standard error of the baseline estimate from the finite-difference Jacobian of the standardized innovations; SD is a standard deviation. The standard deviations of the IS and Phillips shocks are held at 0.40 percent of potential output and 0.25 pp in every column and are not estimated. The intercept c_0 and the standard deviations σ_η and σ_{spf} are in percentage points. The final column is the volume's calibration, used for this paper's experiments in that calibration: λ_1 and κ as chapter 12 prints them, ρ, c_π, and c_y those of the volume's IFB rule, and the IS coefficients this paper's, the chapter printing no IS equation (Section 4.1). The real-rate coefficient a_r is calibrated. The two IS weights lie on the cap of 0.97 on their sum, and their standard errors are computed as if the weights were free (Section 4.3). Sources: ALFRED, RTDSM, SPF, and HLW vintages as accessed 5 September 2026; author's calculations.

Parameter	Symbol	Alternative	Baseline	SE	Volume's calibration
IS: lagged gap	a_{lag}	0.684	0.771	0.030	0.900
IS: expected gap	a_{lead}	0.286	0.199	0.034	0.050
IS: real rate (calibrated)	a_r	0.100	0.100	—	0.100
Phillips: expectation weight	λ_1	0.662	0.946	0.041	0.700
Phillips: slope	κ	0.0314	0.0596	0.0166	0.1000
Rule: smoothing	ρ	0.839	0.656	0.027	0.710
Rule: inflation response	c_π	3.00	3.00	—	0.91
Rule: gap response	c_y	1.58	0.84	0.28	0.21
Rule: intercept (pp)	c_0	0.48	3.00	0.80	—
Rule shock: persistence	ϕ_u	0.790	0.959	0.016	—
Rule shock: SD (pp)	σ_η	0.496	0.947	0.103	—
Survey noise: SD (pp)	σ_{spf}	0.473	0.798	0.077	—

4.4 The model's forecasts against the staff's

The agent is trained inside the estimated model, and Table 2 sets that model's forecast errors, at the baseline estimate, beside those of the Federal Reserve Board staff's Tealbook projections at the decision points from 2000:Q1, the first at which the data hold the staff's projection of inflation, through 2020:Q4, the final one that carries a Tealbook vintage dated within the decision point's quarter. At each, the model's one-quarter-ahead gap and inflation and its four-quarter inflation four quarters ahead are computed from the filtered state under the estimated policy equation, from the IS and Phillips blocks alone and with the residual components of Section 6 added; the staff's projections are read from the same vintage. The realized values are read from the vintage five quarters after the decision point: the gap as the staff's real-time estimate in that vintage, and inflation as quarterly core PCE inflation over that vintage's complete quarters. The gap lines therefore score the 79 decision points from 2000:Q1 through 2019:Q3: the realized gap needs a Tealbook vintage five quarters on, and the data hold Tealbook vintages through 2020:Q4. The inflation lines score 84 decision points, and the four-quarter line 83, 2018:Q4 lacking its realized value: the vintage five quarters on, that of the meeting of 29 January 2020, does not yet hold core PCE inflation for the final quarter of 2019.

Table 2. The baseline model against the Federal Reserve Board staff's Tealbook projections at the decision points from 2000:Q1 through 2020:Q4 with a Tealbook vintage dated within the decision point's quarter. One ahead and four ahead are the horizons in quarters; blocks is the model's forecast from the IS and Phillips blocks alone, and components the forecast with the residual components added. Root-mean-square error (RMSE) and mean error are of the forecast against the realized value, the error being the forecast minus the realized value, in pp (for the gap, in percentage points of potential output); the loss differential is the mean of the model's squared error less the staff's, in squared pp (for the gap, squared percentage points of potential output), with its standard error by the equal-weighted cosine estimator of Section 8.1 in parentheses (SE); for the four-quarter line, which lacks 2018:Q4, no standard error is given over its 83 quarters, and the row below it scores the 75 consecutive quarters through 2018:Q3. Realized values are read from the vintage five quarters after the decision point: the staff's real-time gap in that vintage, and quarterly core PCE inflation from that vintage's complete quarters; the gap lines end at 2019:Q3. Sources: Tealbook and RTDSM vintages of the Federal Reserve Bank of Philadelphia and ALFRED vintages as accessed 5 September 2026; author's calculations.

Forecast	Quarters	Model		Staff		Loss differential
		RMSE	Mean err.	RMSE	Mean err.	Mean (SE)
Gap, one ahead, blocks	79	1.276	+0.495	0.925	−0.041	+0.772 (0.781)
Gap, one ahead, components	79	1.227	+0.263	0.925	−0.041	+0.648 (0.627)
Inflation, one ahead, blocks	84	0.772	+0.071	0.640	+0.061	+0.186 (0.095)
Inflation, one ahead, components	84	0.731	+0.045	0.640	+0.061	+0.124 (0.067)
Four-quarter inflation, four ahead	83	0.745	+0.100	0.562	−0.013	+0.238
The same through 2018:Q3	75	0.559	+0.193	0.380	+0.035	+0.168 (0.062)

The staff's RMSE is the lower on every line. For one-quarter-ahead inflation it is 0.640 pp, against the model's 0.772 pp from the blocks alone and 0.731 pp with the components; for four-quarter inflation four quarters ahead it is 0.562 pp against 0.745 pp, a ratio of the model's RMSE to the staff's of 1.32; and for the one-quarter-ahead gap it is 0.925 against 1.276 and 1.227, in percentage points of potential output. The components lower the model's RMSE on both lines they enter. The mean loss differential, the model's squared error less the staff's, is +0.186 (standard error 0.095) for one-quarter inflation from the blocks alone and +0.124 (0.067) with the components,

+0.238 for four-quarter inflation over its 83 quarters, and +0.772 (0.781) and +0.648 (0.627) for the gap, each in squared pp (for the gap, squared percentage points of potential output). The ratio of the differential to its standard error is 1.95 for one-quarter inflation from the blocks alone and 1.86 with the components, and 0.99 and 1.03 for the gap; none of these four exceeds the critical value of 2.36 (Section 8.1 states the estimator). The four-quarter line lacks 2018:Q4, and the estimator is defined for consecutive quarters. Over the 75 quarters through 2018:Q3 the differential is +0.168 (standard error 0.062), a ratio of 2.71, which exceeds 2.36: over those quarters the model's four-quarter inflation forecast has the higher squared error. Over the eight quarters after the missing one, 2019:Q1 through 2020:Q4, the differential is +0.894 (standard error 1.089), a ratio of 0.82 against 12.71 on one term; no standard error is computed across the missing quarter. The comparison measures the accuracy of the two forecasts and does not separate a difference in information from a difference in model: the staff's projections draw on data that the model's observation rows do not include.

5 Model-Consistent and Alternative Expectations

In the model, forward guidance at the ELB acts through model-consistent expectations of the future policy rate, inflation, and the output gap, and the behavioral and learning literatures cited in Section 1 propose departures from such expectations. This section estimates three families of departures. Each family nests the baseline specification at one value of its parameter, at which expectations are model-consistent. For each family, the Monte Carlo study of Section 4.2 simulates economies with an ELB episode at one value of the family's parameter away from the model-consistent one; the estimates it returns are reported there, with the cognitive-discounting and learned-anchor families below, and in Appendix A. Each family is estimated on the full sample, with c_π at its normalized value, from two starting values, the baseline estimate and the alternative estimate, and at each value of the family's parameter the deeper of the two optima is the family's estimate; the likelihood is profiled in the parameter, and Table 3 collects the estimates.

With its parameter free, every family's deeper optimum lies at the baseline estimate's parameters, at the model-consistent value, and from the alternative estimate as the starting value every family stops at the alternative estimate, again at the model-consistent value. Held at 0.9 (the anchored family) or 0.95 (cognitive discounting), the deeper optimum lies near the baseline estimate with λ_1 at its structural limit of one; held lower, it lies in the region of the alternative estimate, with λ_1 from 0.53 to 0.64.

The anchored family places the four-quarter expectation at a weight w on the model-consistent expectation and $1 - w$ on a fixed level. The level is either the inflation target or the level measured from the Survey of Professional Forecasters' ten-year expectation of consumer price index (CPI) inflation, converted to core PCE units by subtracting the sample's mean wedge between headline CPI and core PCE inflation. The converted survey level averages 1.91 percent across the 122 quarterly surveys (standard deviation 0.21 pp), and the anchor is fixed at that mean. With either level, w lies at its upper bound of one, where the usual asymptotic standard error does not apply, so the likelihood profile is the measure of the evidence (Figure 4, left panel): holding w at 0.9 and 0.8 lowers the log-likelihood by 38.6 and 166.1 with the level as measured, and by 42.4 and 173.4 with the level at the target.

The cognitive-discounting family attenuates every forward expectation k quarters ahead by m^k toward its steady-state value, the target for inflation and zero for the gap. On simulated economies with a true discount of 0.85 and an ELB episode (condition O17 of Appendix A), the estimator

returns discounts from 0.828 to 0.931; four of the five estimates lie above the true value, and the highest lies 3.6 of its standard errors from it. On the sample, m lies at one, and holding it at 0.95, 0.9, and 0.8 lowers the log-likelihood by 73.5, 187.3, and 224.6. The fit over the quarters at the ELB does not improve with discounting: the standard deviation of the standardized innovations over those quarters is 1.28 times that over the other quarters at m of one, and 1.30, 1.65, and 1.70 times at 0.95, 0.9, and 0.8. Discounting lowers the log-likelihood of the full sample by the amounts above and raises the dispersion of the innovations at the ELB relative to the other quarters.

Table 3. Estimates of the three expectations families on the full sample, 1996:Q1–2026:Q2. Each family's parameter indexes its departure from model-consistent expectations, which the family nests at a value of one: w is the weight on the model-consistent expectation against a level (the target, the level measured from the survey, or the learned level), and m is the cognitive discount per quarter of horizon. Family's optimum: the deeper of the optima reached from the baseline and the alternative estimates as starting values; from the alternative estimate: the optimum reached from the alternative estimate as the starting value. Every estimate lies at its upper bound of one. The penalty is the fall in the log-likelihood over the 836 observation rows when the parameter is held at 0.9 or 0.8 with every other free parameter re-optimized, c_π at its normalized value; the learned anchor is profiled in its gain (Figure 4). Sources: as Table 1; author's calculations.

		Family's optimum			From the alternative estimate	
Family		Estimate	Penalty, 0.9	Penalty, 0.8	Estimate	Penalty, 0.8
Anchor, target	w	1.000	42.4	173.4	1.000	53.0
Anchor, measured	w	1.000	38.6	166.1	1.000	45.7
Cognitive	m	1.000	187.3	224.6	1.000	104.2
Learned anchor	w	1.000	—	—	1.000	—

The learned-anchor family makes the anchor's level a constant-gain mean of past realized inflation and takes the ten-year survey as an observation of that level, so that it is estimated on 958 observation rows, the 836 of Section 2 and the 122 of the ten-year survey. On simulated economies with a gain of 0.05 and an ELB episode (condition O18), the estimator returns gains from 0.0483 to 0.0534 with standard errors below 0.0028. On the sample the gain is 0.030 (standard error 0.0024), a half-life of about 23 quarters, and the learned level stands at 2.85 percent at the start of the sample and at 2.61 percent at its end. The weight on the model-consistent expectation lies at its upper bound of one, where the family's free optimum lies at the baseline estimate's parameters, and on the likelihood profile in the gain it is no lower than 0.992 at any gain. The estimated gain is therefore a property of the ten-year survey taken as an observation of the learned level; the survey's measurement error about that level has an estimated standard deviation of 0.30 pp (standard error 0.028 pp), larger than the survey's standard deviation over the sample, 0.21 pp. The gain is not a property of the expectations that enter the IS and Phillips equations, which place no weight on the learned level and remain model-consistent.

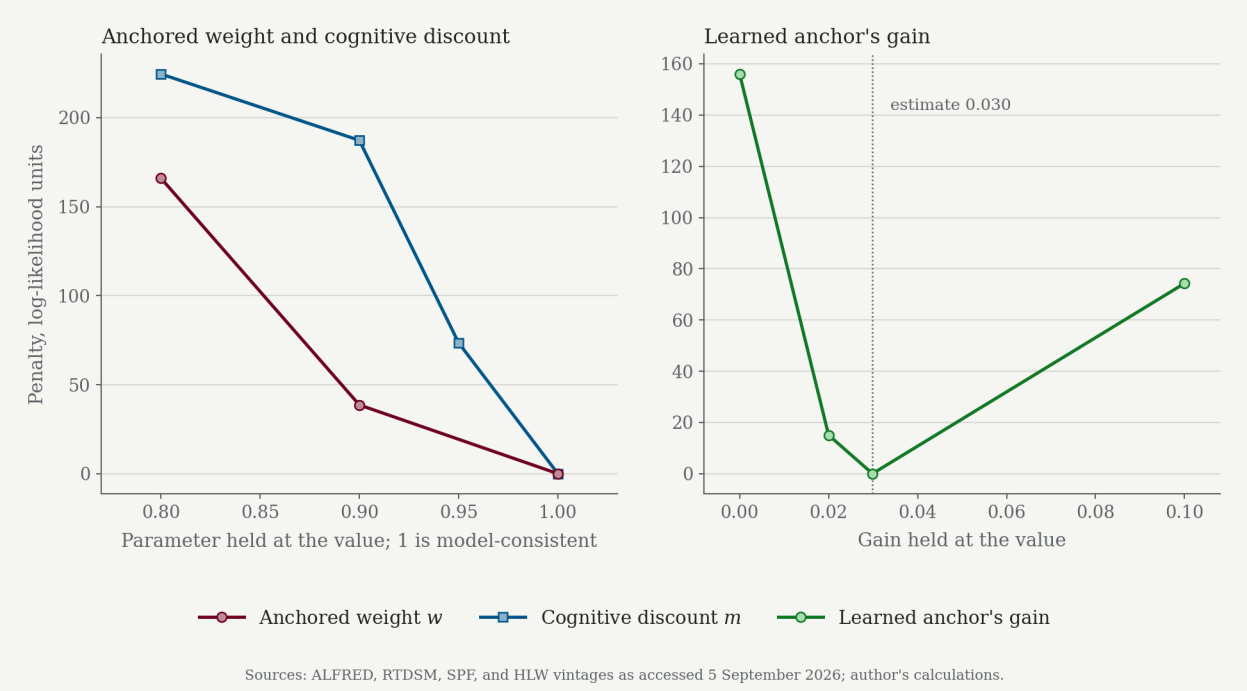


Figure 4. The likelihood profiles of the three expectations families on 1996:Q1–2026:Q2. Left panel: the fall in the log-likelihood over the 836 observation rows (vertical axis) when the anchored family's weight w with the level as measured (maroon circles) or the cognitive-discounting family's discount m (deep blue squares) is held at the value on the horizontal axis, below one, with every other free parameter re-optimized; both are zero at one, where each family is model-consistent, and rise to 166.1 and 224.6 at 0.8. Right panel: the same fall, over the observation rows with the ten-year survey added, when the learned anchor's gain (green circles, horizontal axis) is held at zero, 0.02, and 0.10, against its estimate of 0.030 (dotted line): 156 at zero, 15 at 0.02, and 74 at 0.10. Each point is the deeper of the optima reached from the baseline and the alternative estimates, with c_π at its normalized value. Sources: as Table 1; author's calculations.

Figure 4 plots the three profiles. Figure 5 plots the survey against the learned level and realized inflation. The survey lies below the learned level in the early years of the sample and, from 2023:Q1 on, is nearer the target than the learned level is in every quarter. Over the sample the survey varies less than the learned level, with standard deviations of 0.21 and 0.32 pp, but from 2020:Q3 through 2022:Q4 it rises by 0.92 pp, against 0.68 pp for the learned level. The survey lies more than its standard deviation over the sample, 0.21 pp, below the learned level in 43 quarters, all of them before 2010:Q1 or from 2023:Q1 on, and within that distance of it in each of the 52 quarters from 2010:Q1 through 2022:Q4.

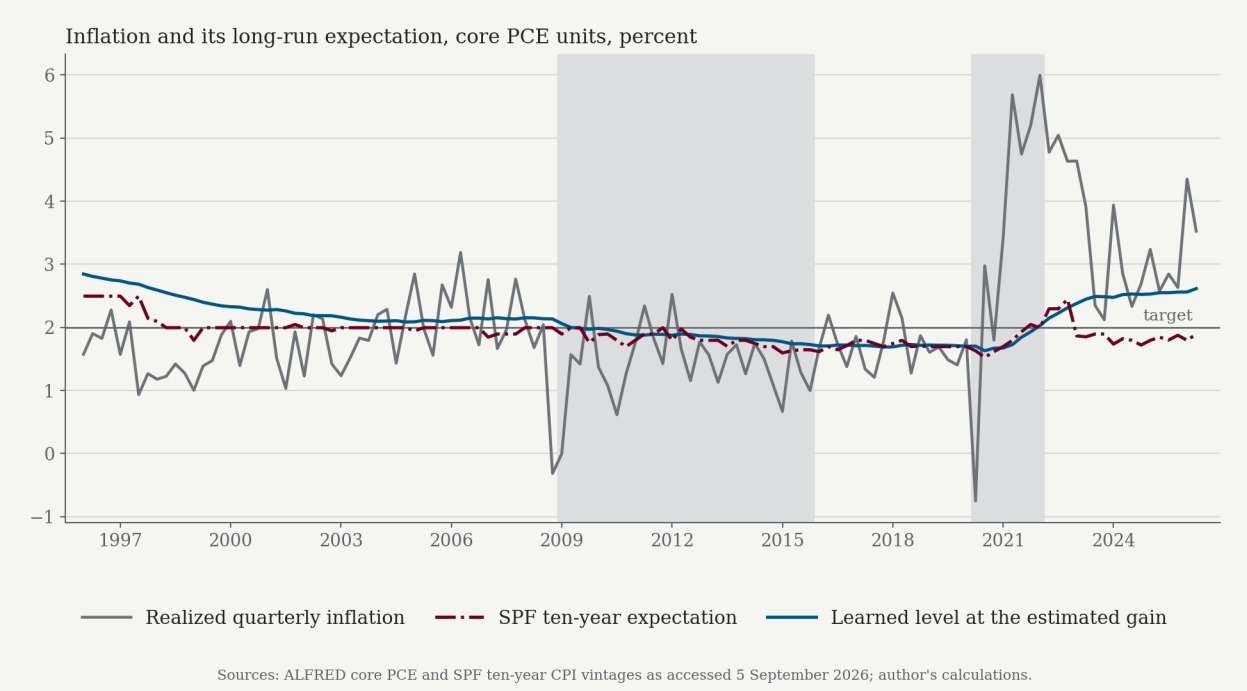


Figure 5. The long-run survey expectation and the learned level, 1996:Q1–2026:Q2, in percent. The maroon dash-dot line is the SPF ten-year expectation of CPI inflation in core PCE units, the survey less the sample's mean wedge between headline CPI and core PCE inflation; the deep blue line is the learned level, a constant-gain mean of realized quarterly core PCE inflation at the estimated gain of 0.030, starting at 2.85 percent and ending at 2.61 percent; the gray line is realized quarterly core PCE inflation at an annual rate; the thin gray line is the inflation target; shaded areas mark the quarters at the ELB. Sources: ALFRED core PCE and SPF ten-year CPI vintages as accessed 5 September 2026; author's calculations.

6 Residual Components and a Null Distribution with an ELB Episode

The IS and Phillips equations leave residuals, and a nonparametric correction can be fitted to them. For each of the two filtered shocks, the IS shock and the Phillips shock, a Gaussian process (GP) component is fitted on four features of the filtered state: the lagged output gap, mean inflation over the past four quarters, the lagged policy rate, and the neutral rate. Each component has a squared-exponential kernel with one length scale for the four features, log-normal priors on its hyperparameters, maximum a posteriori (MAP) estimates of them, and a signal standard deviation capped at a detection threshold. The threshold is calibrated on the simulated economies with an ELB episode of the next paragraph: it is the standard deviation of the smallest planted component whose test statistic exceeds, on at least four of five economies, the largest statistic among the economies with no planted component, and it caps each component's signal standard deviation at 0.33 times the standard deviation of its filtered shock.

The statistic that tests a component for signal is its leave-one-out (LOO) log-score excess: the mean over the quarters of the GP's leave-one-out log predictive density of the filtered shock, less that of a zero-mean Gaussian with the shock's sample variance, in natural-log units per quarter. It is computed on a fit of the component without the cap on its signal standard deviation, for the

sample and for every simulated economy alike, the cap being set from the simulated economies' statistics. Its distribution under the null of no signal is obtained from 25 economies simulated at known parameters with no residual signal, each with an ELB episode and the sample's pattern of missing observation rows (condition O16 of Appendix A). The simulated economies differ from the sample in three respects: each has 96 quarters where the sample has 122; their shocks are filtered at the parameters that generated them, where the sample's are filtered at the estimate; and their parameters are the volume's calibration with an anchoring weight of 0.80, not the baseline estimate. The paper measures a further difference, the length of the ELB episode, for five of them, the samples of condition O16 in the Monte Carlo study: these have 7 to 17 quarters at the ELB, also the range of conditions O15 and O16 together in Appendix A, against the 28 quarters of the sample's first episode, which falls in the sample's first 96 quarters. The paper does not measure the length for the other 20. For the Phillips component, the null's median is 0.217 and its 95th percentile 0.385; on 25 simulated economies of 52 quarters, the length of the sample before its first ELB episode, without an ELB episode, the 95th percentile is 0.079. For the IS component, the 25 economies with an ELB episode give a 95th percentile of 0.145. An ELB episode makes the filtered shocks depend on the state even at the true parameters: the policy-rate row is censored at the ELB, the lagged rate in the state marks those quarters, and a GP on the state's features fits that dependence. A null simulated without an ELB episode would therefore count the dependence as signal.

On the baseline estimate, the Phillips component's excess, 0.215, lies at the 48th percentile of its null, and the Phillips component shows no evidence of signal beyond what an ELB episode induces. The IS component's excess, 0.239, exceeds all 25 simulated statistics of its null, the largest of which is 0.155; the excess is evidence of signal in the IS residual (Section 9 states it as a limitation of the IS equation). The capped components, which are the ones that enter a simulation, have a lower excess: 0.092 for the IS component and 0.092 for the Phillips component. Against the same statistic on the 25 simulated economies, each component there capped as the sample's is, the capped IS component's excess again exceeds every simulated statistic, the largest of which is 0.064 and the 95th percentile 0.047, and the capped Phillips component's excess lies at the 48th percentile of its null, whose 95th percentile is 0.205.

The baseline agent's training environment does not include the components; Section 7.2 trains an agent with them and compares the two. Where the components enter a simulation, the solution at each quarter adds each component's prediction to that quarter's shock and to no later quarter's. The closed-loop simulation of 24 quarters from the 2014:Q4 state of Section 3.1, under the volume's IFB rule in the model at the baseline estimate, stays inside the 99th-percentile Mahalanobis region of the fitted states, the states whose four features lie at a Mahalanobis distance from the fitted states below the 99th percentile of a chi-square distribution with four degrees of freedom: the state of every one of the 24 quarters lies inside it, so the components are not extrapolated on that path. Of the fitted states, 36 are at the ELB. On that path, with the components, the equilibrium exists at every quarter under the volume's IFB rule and under each of nine other rules on a grid of inflation and gap responses.

7 The Policy Agent

The policy agent is the rule (4) with the shock u_t set to zero, its four coefficients trained by gradient descent in the estimated model as Section 3 describes. This section states the loss and the training

environment, compares the baseline agent with an agent trained with the residual components, and selects the baseline agent's coefficients from eight training runs.

7.1 Loss and training environment

The agent minimizes the loss (1) at the weights of Adrian, Laxton, and Obstfeld's volume: α of one on squared deviations of four-quarter inflation from the target, β of one on the squared output gap, and γ of one half on squared quarterly changes in the policy rate. Its coefficients (ρ, c_π, c_y, c_0) are trained on episodes. An episode is a simulation of the model from the real-time state at one decision point over 24 quarters with the structural shocks at zero, the equilibrium at each quarter solved exactly over the model's forward horizon of 48 quarters with the ELB as a complementarity condition; its loss is (1) summed over those 24 quarters. The episodes start from the real-time states of the sample, 118 decision points from 1997:Q1 through 2026:Q2, 36 of them at the ELB. The training environment is the IS and Phillips equations of the baseline estimate, with four of their coefficients drawn for each episode around the estimate (Section 7.3), and without the residual components of Section 6.

7.2 Training with and without the residual components

A policy optimized against a model can gain from the model's errors as well as from its structure. This subsection trains an agent with the residual components of Section 6, the GP corrections to the IS and Phillips residuals, in its environment and evaluates it in the estimated model without them. Against the null with an ELB episode, the IS component's excess of 0.239 exceeds all 25 simulated statistics of its null, and the Phillips component's excess lies at the 48th percentile of its null. The evidence of signal holds for the IS component and not for the Phillips component.

A second agent is trained on the IS and Phillips equations of the baseline estimate with the components in the environment, each episode drawing the components from their posterior, and is selected among eight training runs as the baseline agent is (Section 7.3). It has ρ of 0.767, c_π of 1.918, c_y of 3.900, and c_0 of -2.176, against the baseline agent's 0.736, 3.353, 3.159, and -0.580.

At the 118 decision points, its prescribed rate has a mean difference from the realized federal funds rate of +0.637 pp, with a standard error of 0.218 pp, and a root-mean-square (RMS) difference of 0.987 pp (Section 1 defines the difference and Section 2 the realized rate), against the baseline agent's mean difference of +0.723 pp and RMS difference of 1.035 pp. Before the first ELB episode, its RMS difference is 0.76 pp against the baseline agent's 0.81 pp, and it prescribes the ELB on 30 of the 36 quarters at the ELB against the baseline agent's 29.

Its loss over 24 quarters from the 2014:Q4 state, in the model at the baseline estimate without the components, is 12.853, against 9.944 for the volume's IFB rule and 6.106 for the baseline agent. With the components' posterior means in the model, from the same state, its loss is 13.296 and the baseline agent's 17.755, so that each of the two rules has the lower loss in the environment in which it was trained; the agent trained with the components still has a lower loss without them than with them, and the volume's IFB rule has a loss of 159.1 with them. With the gap at zero, no policy shock, and the rate at the neutral rate plus inflation, the rule (4) is at rest at inflation of $\pi^* - c_0/c_\pi$, which is 2.17 percent for the baseline agent and 3.13 percent for the agent trained with the components.

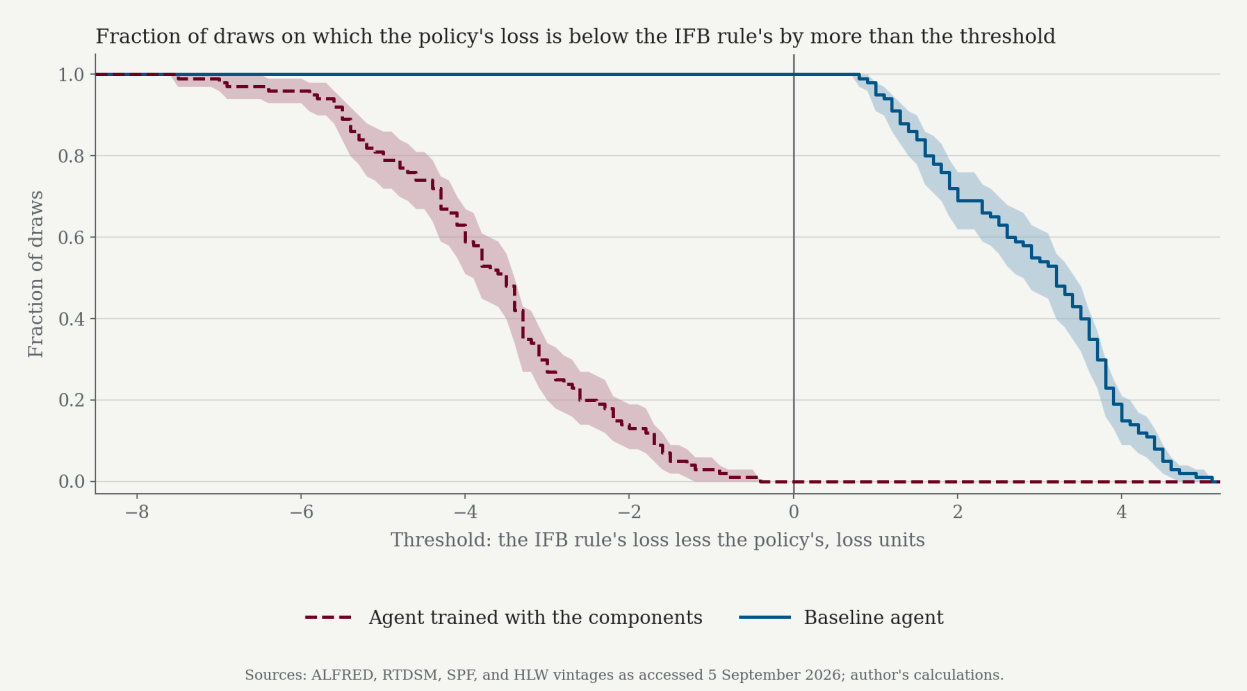


Figure 6. Performance of the trained rules across draws around the baseline estimate. For each of 100 draws of the four estimated coefficients of the IS and Phillips equations (a_{lag}, a_{lead}, λ_1, and κ) around the baseline estimate, the loss (1) is summed over 24 quarters from the 2014:Q4 state, with no shocks, in the model at the draw, for the baseline agent (deep blue, solid) and the agent trained with the residual components (maroon, dashed), and for the volume's IFB rule, the reference, which is not drawn. Each curve is the fraction of draws (vertical axis) on which the policy's loss lies below the IFB rule's by more than the threshold on the horizontal axis, in the loss's units; the shaded band is the pointwise 90 percent interval from 1,000 bootstrap resamples of the draws; the vertical line marks zero. Sources: as Table 1; author's calculations.

Figure 6 reports performance across draws around the estimate rather than at the point estimate. Over 100 draws of the four coefficients around the baseline estimate, the baseline agent's loss from the 2014:Q4 state is below the volume's IFB rule's on 100 draws, by 2.89 on average.

On the same draws, the loss of the agent trained with the components is below the IFB rule's on none: it is higher on every draw, by between 0.45 and 7.58 and by 3.72 on average.

In the estimated model without the components, an agent trained with them thus has a higher loss than the volume's IFB rule at the baseline estimate and on every draw around it, and the agent trained without them a lower loss than the IFB rule on every draw. The evaluation is conditional on the estimated equations: it does not measure the loss in a model that contains what the IS component captures. With the components in the model, at the one state and estimate evaluated, the agent trained with them has the lower loss of the two agents.

7.3 Training runs and the selected checkpoint

The coefficients at which gradient descent stops depend on the number of its steps, and those of a single training run depend on its random draws: another run, with its own random-number stream for the draws and the starting states, ends at other coefficients. A step is one gradient update of

the four coefficients on a set of episodes, a training run is a sequence of steps under one random-number stream, and a checkpoint is the rule's coefficients at one step of one run. Each training episode draws the four coefficients of the IS and Phillips equations (a_{lag}, a_{lead}, λ_1, and κ) around the estimate, each from an independent normal distribution centered on the estimate's value with its standard error and clipped at the wider bounds of Table A.2, so the policy minimizes the expected loss over those draws rather than the loss at the point estimate. The draws are not draws from the sampling distribution of the estimator: the four coefficients are drawn independently of one another; the two IS weights lie on the cap on their sum, where their standard errors are computed as if the weights were free (Section 4.3); and every other parameter, the real-rate coefficient, the policy equation, and the shocks' standard deviations among them, is held at the estimate. A draw whose two IS weights sum above 0.97, the cap the estimator places on their sum (Section 4.3), is projected onto the cap by the shortest move within those bounds. The draws of the validation, test, and fresh sets below, those of Figure 6, those of the band of the agent's rate in Section 8.1, and those of the rolling windows' validation, test, and fresh sets in Section 8.3 are made the same way.

The baseline agent is selected from eight training runs of 2,000 steps each. The validation set is 32 draws of the four coefficients and eight starting states; its draws enter no training episode, and its starting states are decision points from which training episodes also start. The validation loss is the mean of the episode loss over its 32 times eight episodes; it is computed every 10 steps on every run. The selected checkpoint is the one with the lowest validation loss over the runs and their steps, and where that checkpoint falls in the final 200 of the 2,000 steps, every run is extended once, to 4,000 steps, and the checkpoint is selected over the extended runs. The validation starts are drawn uniformly from the decision points, for the baseline agent 2000:Q1, 2010:Q4, 2011:Q2, 2013:Q3, 2016:Q1, 2016:Q4, 2020:Q4, and 2022:Q2, seven of them from 2008:Q4 on. The test set, eight further starts from the remaining decision points and 32 further draws around the estimate, enters no selection of the reported checkpoint.

On the baseline estimate, the lowest validation loss over the first 2,000 steps is that of run 5 at step 1,820, within the final 200 steps, so the runs are extended to 4,000 steps. The validation loss falls from 44.11 at the initial coefficients to 30.74 at the selected checkpoint, run 4 at step 3,410 (Figure 7). Run 4's validation losses at the steps on either side of the selected checkpoint, 3,400 and 3,420, are 30.80 and 30.76, and its validation losses after the first tenth of its steps run from 30.74 to 31.32. On the test set the selected run ranks eighth of eight. The runs' ranking on the validation set is nearly the reverse of their ranking on the test set. Over the baseline agent's eight runs, the Spearman rank correlation of the checkpoints' validation and test losses is -0.98. Besides the rolling windows of Section 8.3, the paper has eleven searches, each a selection among eight training runs: the baseline agent's, the four under other random-number streams of Table 11, the four under other loss weights of Table 10, that of an agent trained on the alternative estimate, and that of the frozen agent (Section 8.2). The correlation is positive in none of them; it runs from -0.95 to -0.86 across the four searches under other streams and from -0.98 to -0.33 across the four loss-weight variants, and it is -0.62 for the agent trained on the alternative estimate and -0.62 for the frozen agent. These correlations are not independent: nine of the eleven searches, the baseline agent's, those under other streams, and the loss-weight variants', share one estimate, one validation set, and one test set, and the correlation for the agent trained on the alternative estimate rests on test losses within 0.6 percent of one another. Three further searches are reported with their comparisons and are not among the eleven: that of the agent trained with the components (Section 7.2), whose selected run ranks fifth of eight on its test set, with a correlation of +0.36, and those of the two rules that exchange the Phillips slope between the estimates (Section 8.2), whose selected runs rank sixth and first, with correlations of -0.12 and -0.26.

The eight runs' checkpoints have validation losses from 30.74 to 30.81 and test losses from 14.83 to 14.92. The two levels differ with the two sets' starting states: the validation set's lie farther from the target, its largest output gap in absolute value being −6.28 percent of potential output, at 2010:Q4, against −4.56 percent in the test set, and the episodes that start far from the target contribute most to a mean loss. Section 8.3 measures, on the 82 searches of the rolling agent, how a ranking of eight runs on a set of this size carries over to a fresh set of starting states and draws, defined below: the median Spearman correlation with the fresh set's ranking is +0.20 for the test-set ranking and −0.24 for the validation-set ranking. The rules at the eight checkpoints differ by more than their losses: across the baseline agent's runs, c_π runs from 3.215 to 3.698, c_y from 3.140 to 3.223, and ρ from 0.717 to 0.736, while their test losses differ by 0.7 percent of the lowest. The loss varies little across a set of rules, the selection picks one of them, and the baseline coefficients are read as that selection, not as the only rule the procedure supports.

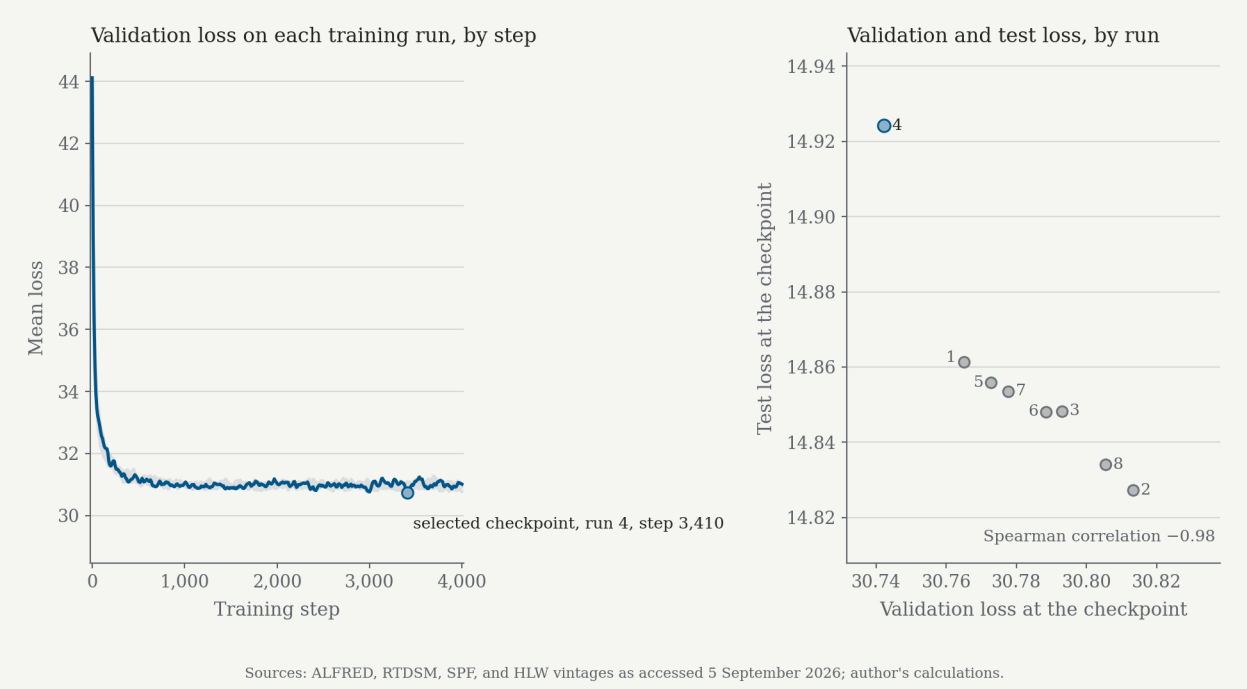


Figure 7. The search over training runs on the baseline estimate. Left panel: the validation loss (vertical axis), computed every 10 steps (horizontal axis) as the mean loss over the fixed validation set of eight real-time starts times 32 draws of the four coefficients of the IS and Phillips equations, the draws entering no training episode, on each of the eight runs over 4,000 steps; the selected run, 4, is deep blue and the others light gray, and the marker is the selected checkpoint at step 3,410. Right panel: each run's checkpoint, labeled by its run, its validation loss (horizontal axis) against its loss on the test set (vertical axis), the selected checkpoint in deep blue; the Spearman rank correlation of the two is −0.98. Losses are in the units of the loss (1). Sources: as Figure 3; author's calculations.

A second criterion, the mean of a checkpoint's validation and test losses, selects run 1 for the baseline agent, where the validation set selects run 4. It selects a different run from the validation set's in ten of the eleven searches, all but that of Repetition 1 of Table 11. Each search is also evaluated on a fresh set, eight starting states drawn from the decision points outside its validation and test sets and 32 draws around its estimate that enter no training episode; the fresh set enters no selection. In each of the ten searches where the two criteria differ, the run the second criterion

selects ranks higher on the fresh set than the run selected on the validation set, and the run selected on the validation set ranks eighth of eight on the fresh set in eight of the eleven. The paper reports the checkpoint selected on the validation set throughout. The search on the alternative estimate's equations runs to 2,000 steps; its selected checkpoint is at step 1,240 of run 3, and that agent's mean difference from the realized federal funds rate over the 118 decision points is +0.744 pp.

The baseline agent has ρ of 0.736, c_π of 3.353, c_y of 3.159, and c_0 of -0.580. Its loss over 24 quarters from the 2014:Q4 state, with no shocks, is 6.106 in the model at the baseline estimate, against the volume's IFB rule's 9.944, and 10.07 in the volume's calibration, against the IFB rule's 14.38. An agent trained in a single run on the same equations at the point estimate, without draws around the estimate, has a loss of 11.01 in the volume's calibration. The experiment of Section 3.1 sums the loss over sixteen quarters and adds the loss of the terminal plan, so its losses differ from these.

8 The Learned Rule and the FOMC

This section compares the rate the baseline agent prescribes at each decision point with the realized federal funds rate, with a band from draws around the baseline estimate and beside three published rules. It then repeats the comparison for an agent estimated and trained on data through the end of the first ELB episode and for an agent re-estimated and retrained each quarter, and it simulates the outcomes each rule would have produced under the structural shocks recovered from the realized data.

8.1 The baseline agent and the realized federal funds rate

At each of the 118 decision points from 1997:Q1 through 2026:Q2, the rate the baseline agent prescribes for the decision quarter is computed in the model at the baseline estimate. It is compared with the realized federal funds rate at that decision point (Section 2) and with the rate of the volume's IFB rule. The band around the agent's rate is the 10th–90th percentile band across 20 draws around the baseline estimate; with 20 draws, each percentile rests on few observations. The solver finds an equilibrium within its convergence tolerance at every one of the 2,360 pairs of a decision point and a draw. Figure 8 plots the three rates and the band, and Table 4 reports the comparison by period.

Over all 118 quarters the mean difference (the agent's prescribed rate minus the realized federal funds rate, Section 2) is +0.723 pp, with a standard error of 6.8 bp under independence. The quarterly differences are serially correlated, with a first autocorrelation of 0.89: the Newey–West standard error with 4 lags is 0.141 pp, and a moving-block bootstrap of 2,000 resamples in blocks of 8 quarters gives a 95 percent interval of +0.372 to +1.038 pp. A bandwidth of 4 lags understates the standard error at that persistence. The equal-weighted cosine estimator of the long-run variance, with the 9 terms that the rule of Lazarus, Lewis, Stock, and Watson gives for 118 quarters, puts the standard error at 0.213 pp; the ratio of the mean difference to it, 3.39, exceeds 2.26, the critical value of a two-sided test at the 5 percent level under Student's t distribution with 9 degrees of freedom. Over 1997:Q1 through 2008:Q4 the standard error is 0.173 pp and the ratio 3.81 against a critical value of 2.57; over 2016:Q1 through 2020:Q1, where the mean difference is +0.659 pp, they are 0.151 pp and 4.35 against 4.30, and over 2022:Q2 through 2026:Q2, where it is +2.041 pp, 0.211 pp and 9.68 against 4.30, the quarters from 2016:Q1 outside the ELB being two separate spans that

are read apart, since the estimator is defined for consecutive quarters; and over 2021:Q1 through 2023:Q4, with two terms, 0.544 pp and 3.49 against 4.30.[12] The rule is derived for series no more persistent than a first-order autoregression with a coefficient of 0.7, which the differences over all quarters exceed. Every other standard error of a mean difference or of a mean loss differential in the text and the tables is by the equal-weighted cosine estimator, with the number of terms the rule gives for the period's length, and is read against the critical value of Student's t distribution with that many degrees of freedom. A ratio is stated in absolute value, the test being two-sided. No standard error is reported for a period in which most quarters are at the ELB, where the difference cannot be negative. Some of the other periods over which a standard error of a mean difference is reported include quarters at the ELB beside quarters outside it. All 118 quarters include both episodes, and the 42 quarters from 2016:Q1 include the eight quarters of the second; 2016:Q1 through 2020:Q4 includes three of them, and 2021:Q1 through 2023:Q4 and 2021:Q1 through 2026:Q2 include five each.

Table 4. The baseline agent and the realized federal funds rate, 1997:Q1–2026:Q2. Mean is the mean difference between the rate the policy prescribes for the decision quarter and the realized federal funds rate (Section 2), in percentage points; RMS is the root mean square of that difference, in percentage points; the row 2009:Q1–2015:Q4 is the first ELB episode and the row 2020:Q2–2022:Q1 the second; the row from 2016:Q1 with the ELB excluded pools 2016:Q1–2020:Q1 and 2022:Q2–2026:Q2, which the text reads apart; the final row, 2021:Q1–2023:Q4, overlaps the two rows above it; coverage is the share of quarters at which the realized rate lies within the agent's 10th–90th percentile band across 20 draws around the baseline estimate; below is the number of quarters at which the realized rate lies below that band. The IFB rule is the volume's, at the volume's coefficients. Table C.1 reports the ratio of the baseline agent's mean difference to its standard error by period. Sources: as Table 1; author's calculations.

Period	Quarters	Agent		IFB rule		Coverage	Below
		Mean	RMS	Mean	RMS		
All quarters	118	+0.723	1.035	+0.405	0.602	0.25	83
1997:Q1–2008:Q4	48	+0.659	0.813	+0.450	0.665	0.02	42
2009:Q1–2015:Q4	28	+0.013	0.053	+0.227	0.343	0.93	2
2020:Q2–2022:Q1	8	+0.924	1.243	+1.086	1.169	0.38	5
From 2016:Q1, ELB excluded	34	+1.350	1.555	+0.328	0.461	0.00	34
2021:Q1–2023:Q4	12	+1.898	2.075	+0.524	0.813	0.08	11

The band covers the realized rate in a share of 0.25 of the quarters, and the realized rate lies below the band in 83 of the 118. Over the first ELB episode, 2009:Q1 through 2015:Q4, the lower edge of the agent's band is at the ELB in 26 of 28 quarters, the band covers the realized rate in a share of 0.93 of them, and the agent's RMS difference is 0.053 pp against the volume's IFB rule's 0.343 pp; its mean difference of +1.3 bp is the smallest in absolute value of any period in Table 4. Over the second ELB episode, 2020:Q2 through 2022:Q1, the agent's mean difference is +0.924 pp, the realized rate being at the ELB, and over 2021:Q1 through 2023:Q4 it is +1.898 pp, with the realized rate below the band in 11 of its 12 quarters. These are properties of the learned rule under the loss (1) in the estimated model, conditional on the loss weights and on the model.

The state of Section 2 sets the decision quarter's IS and Phillips shocks to zero, and its filter reads the decision quarter's inflation. Three other treatments of the state follow; a reading is the ratio of

[12] Eben Lazarus et al., "HAR Inference: Recommendations for Practice," *Journal of Business & Economic Statistics* 36, no. 4 (2018): 541–559.

a mean difference to its standard error set against its critical value. With the decision quarter's IS and Phillips shocks at the filter's estimates in place of zero (Section 2), the mean difference over all 118 quarters is +0.649 pp, and the prescribed rate is at the ELB at three further decision points, 2008:Q4, 2015:Q3, and 2020:Q2; over 2016:Q1 through 2026:Q2 the baseline agent's mean difference is then +1.230 pp against +1.269 pp, and that of the frozen agent of Section 8.2 +0.284 pp against +0.212 pp (Table C.1 collects the frozen agent's means under the four treatments). Over 2021:Q1 through 2023:Q4 the baseline agent's is then +2.010 pp against +1.898 pp, and the frozen agent's +0.640 pp against +0.278 pp.

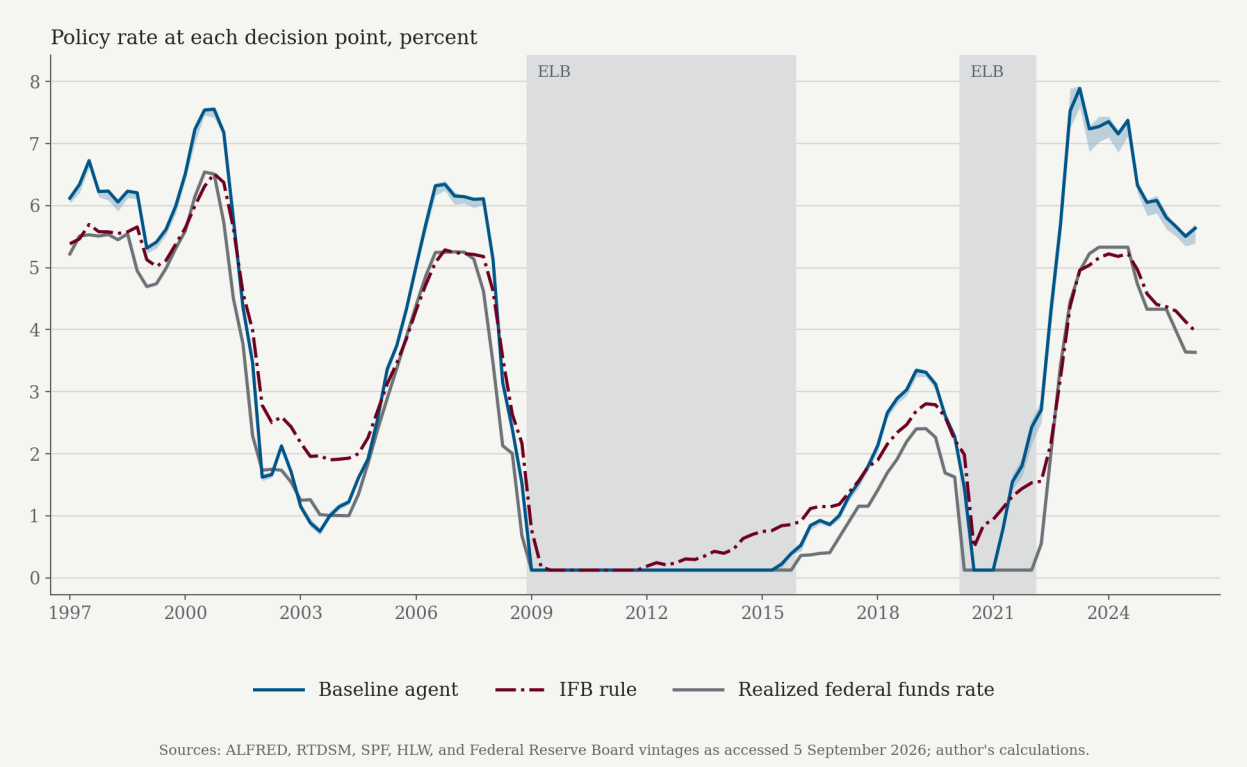


Figure 8. The baseline agent and the realized federal funds rate at the 118 decision points from 1997:Q1 through 2026:Q2. Vertical axis: policy rate, in percent; horizontal axis: decision quarter. The deep blue line is the rate the baseline agent prescribes for the decision quarter in the model at the baseline estimate, and the blue band its 10th–90th percentile band across 20 draws around the estimate; the maroon dash-dot line is the rate of the volume's IFB rule at the volume's coefficients; the gray line is the realized federal funds rate (Section 2), set to the model's lower bound of 0.125 percent in the two ELB episodes (shaded). Sources: ALFRED, RTDSM, SPF, HLW, and Federal Reserve Board vintages as accessed 5 September 2026; author's calculations.

With the decision quarter's inflation withheld from the filter at each decision point, and every other observation of that quarter and of the quarters before it retained, the lagged gap of the state changes by 0.12 percentage points of potential output root mean square and the baseline agent's prescribed rate by 0.09 pp; the prescribed rate changes most at 2024:Q1, by 0.33 pp. The mean difference over all 118 quarters is then +0.720 pp against +0.723 pp, a change of −0.3 bp whose standard error, computed on the changes decision point by decision point, is 1.4 bp, a ratio of 0.24 that does not exceed 2.26; from 2021:Q1 the change is −6.5 bp, with a standard error of 1.8 bp and a ratio of 3.59, which exceeds 3.18. The frozen agent's mean difference over 2016:Q1 through

2026:Q2 is then +0.200 pp against +0.212 pp. The model is not re-estimated without the decision quarters' inflation.

The remaining treatment adds the vintage of the price index. At each decision point whose vintage covers the four preceding quarters, 103 of the 118 (all but 1997:Q1 through 2000:Q2 and 2020:Q1), the inflation of the ten or eleven quarters that the vintage holds is read from that vintage, in the filter's observations and in the state's four lags, and the decision quarter's inflation is withheld; the inflation of the quarters before those is read from the later vintage of Section 2, and the other decision points carry the state with the decision quarter's inflation withheld. Over the 118 decision points the baseline agent's prescribed rate then differs from its rate at the state of Section 2 by 0.16 pp root mean square, and most at 2004:Q1, by 0.48 pp. Its mean difference over all quarters is +0.666 pp (standard error 0.206 pp; ratio 3.24 against 2.26). Against the mean difference of +0.723 pp at the state of Section 2, that is a change of −5.6 bp. The change has a standard error of 3.4 bp and a ratio of 1.66, which does not exceed 2.26. From 2021:Q1 the change is −16.3 bp, with a standard error of 3.8 bp and a ratio of 4.33, which exceeds 3.18. With the vintages' inflation in the four lags alone and the filtered gap unchanged, the mean difference over all quarters is +0.717 pp. With the vintages' inflation in the filter and in the lags, the frozen agent's mean difference over 2016:Q1 through 2026:Q2 is +0.206 pp.

Table C.1 in Appendix C sets the readings at the state of Section 2 beside those under the three other treatments. The reading over all quarters and those over 2022:Q2 through 2026:Q2 and over 2021:Q1 through 2026:Q2 exceed their critical values under all four, and the reading over 2021:Q1 through 2023:Q4 under none. The reading over the 42 quarters from 2016:Q1 exceeds its critical value under all four as well; the smallest of the four ratios, with the decision quarter's shocks at the filter's estimates, is 2.79 against 2.78. The reading over 1997:Q1 through 2008:Q4 does not exceed its critical value with the decision quarter's shocks at the filter's estimates (2.20 against 2.57), and the reading over 2016:Q1 through 2020:Q1 does not with those shocks (2.78 against 4.30) or with the decision quarter's inflation withheld (4.19). Over 2016:Q1 through 2020:Q4, the earlier period of the division of Section 8.2, which adds the first three quarters of the second ELB episode to 2016:Q1 through 2020:Q1, the reading with the shocks at zero does not exceed its critical value (4.20 against 4.30), and of the four treatments only the vintages' inflation gives a reading that exceeds it (4.53).

Table 5 places the agent beside three of the rules presented in the Board of Governors' *Monetary Policy Report* of February 2019, which writes them with the unemployment gap, each evaluated at the states at which the agent's rate is computed, in output-gap form on the state's four-quarter core PCE inflation, filtered lagged gap, and lagged rate, and bounded below by the model's lower bound: the Taylor (1993) rule, the balanced-approach rule with a unit response to the gap, and the first-difference rule.[13] The two level rules use a neutral real rate of 2 percent, while the model's neutral rate is the HLW path of Section 2, which runs from 0.57 to 3.90 percent over the sample. Over all quarters the mean difference is +1.483 pp for the Taylor (1993) rule and +1.603 pp for the balanced-approach rule, against +0.723 pp for the agent. Over 2021:Q1 through 2023:Q4 the two level rules' mean differences are +6.117 pp and +7.244 pp, each rule prescribing a rate above the realized rate in all 12 quarters, where the agent's mean difference is +1.898 pp. The agent's mean difference is smaller than both level rules' over all quarters, over each ELB episode, and over the quarters from 2016:Q1 outside the ELB, but not before the first ELB episode: over 1997:Q1 through 2008:Q4 the mean difference is +0.130 pp for the Taylor (1993) rule and +0.285 pp for

[13] Board of Governors of the Federal Reserve System, *Monetary Policy Report* (Washington, DC: Board of Governors of the Federal Reserve System, February 2019), Part 2, box "Monetary Policy Rules and Systematic Monetary Policy."

the balanced-approach rule, against +0.659 pp for the agent. The level of the agent's rate has a correlation of 0.98 with the first-difference rule's and 0.97 with the volume's IFB rule's, against 0.72 with the Taylor (1993) rule's; these are correlations of levels between rules that, the Taylor (1993) rule excepted, share the lagged rate, and they do not establish the rule's form. The first-difference rule's mean difference is below the agent's in every period of Table 5 except the two ELB episodes.

Table 5. The baseline agent, the volume's IFB rule, and three published rules against the realized federal funds rate, 1997:Q1–2026:Q2. Mean and RMS are the mean and root-mean-square difference between the rule's prescribed rate and the realized federal funds rate (Section 2), in percentage points, for every rule, over the periods and quarters of Table 4. Taylor (1993) is $i = 2 + \pi + 0.5(\pi - 2) + 0.5y$ on four-quarter core PCE inflation π and the previous quarter's filtered gap y; the balanced-approach rule has a unit gap response; the first-difference rule adds $0.5(\pi - 2) + 0.5(y - y_{-4})$ to the lagged rate; each is bounded below by the model's lower bound of 0.125 percent. The IFB rule is the volume's, at the volume's coefficients. Sources: as Table 1; author's calculations.

Period	Agent		IFB rule		Taylor (1993)		Balanced		First difference	
	Mean	RMS	Mean	RMS	Mean	RMS	Mean	RMS	Mean	RMS
All quarters	+0.723	1.035	+0.405	0.602	+1.483	2.615	+1.603	2.967	+0.223	0.750
1997:Q1–2008:Q4	+0.659	0.813	+0.450	0.665	+0.130	1.404	+0.285	1.096	−0.037	0.511
2009:Q1–2015:Q4	+0.013	0.053	+0.227	0.343	+1.042	1.322	+0.284	0.633	+0.192	0.284
2020:Q2–2022:Q1	+0.924	1.243	+1.086	1.169	+4.070	4.863	+3.805	4.923	+1.000	1.570
From 2016:Q1, ELB excluded	+1.350	1.555	+0.328	0.461	+3.149	3.735	+4.031	4.778	+0.432	0.968
2021:Q1–2023:Q4	+1.898	2.075	+0.524	0.813	+6.117	6.510	+7.244	7.761	+1.752	2.014

8.2 The agent frozen at the end of the first ELB episode

The baseline agent is estimated and trained on the same sample over which its prescriptions are compared with the realized federal funds rate. The frozen agent is the comparison without that use of later data: it is estimated and trained on data through 2015:Q4, the final quarter of the first ELB episode, and then held fixed. The sample is truncated to its first 80 quarters (28 of them at the ELB). The wedge between headline and core inflation, through which the ten-year survey is expressed in core PCE units, is re-measured on that window, at a mean of 0.5400 pp against the full window's 0.5052 pp. The model is estimated with model-consistent expectations (the anchoring weight at one) from ten random starting values within the wider parameter bounds of Table A.2, the residual components are re-fitted, and the agent is trained and selected by the procedure of Section 7.3 on the frozen estimate, with episodes starting at decision points through that quarter, the cutoff, and with a validation set of its own. The moments of the measurement errors are those known at the cutoff (Section 2). Apart from the three inputs that are not real-time, the ex post HLW path, the later vintage of the price index, and the later release of the term premium (Section 2), no observation after 2015:Q4 enters the frozen agent's estimation or training. Its specification is the one settled on the full sample: the normalization of c_π, the cap on the IS weights, the parameter bounds, and the two assumed standard deviations of Section 4.2. The frozen agent then prescribes a rate at each of the 42 decision points from 2016:Q1 through 2026:Q2, at the real-time state filtered under the frozen estimate.

Table 6 places the frozen estimate beside the full-sample estimate. Of the ten starting values, seven reach optima with λ_1 from 0.687 to 0.693 and a Phillips slope from 0.0164 to 0.0174. Their

objectives run from -0.2110 to -0.2085 on the frozen window's 596 observation rows, 1.5 log-likelihood units apart, and the deepest is the frozen estimate, with λ_1 at 0.693 and a Phillips slope of 0.0172 (standard error 0.0037). The other three reach optima with a Phillips slope of zero, at objectives from -0.1511 to -0.1361. A search on the frozen window from the baseline and the alternative estimates as starting values reaches no deeper optimum: the better of its two optima lies 64.6 log-likelihood units below the frozen estimate. The full-sample slope, 0.0596, is 3.5 times the frozen estimate's and lies 11.5 of the frozen estimate's standard errors from it. The frozen estimate's policy equation holds c_π at the normalization's value, as every estimate in the paper does. The frozen agent's selected checkpoint stands at step 1,760 of its training, at ρ of 0.639, c_π of 1.924, c_y of 3.084, and c_0 of -0.270: its inflation and gap responses are 0.57 and 0.98 times the baseline agent's. Its training run ranks eighth of eight on the frozen training's test set (Section 7.3).

Table 6. The estimate frozen at 2015:Q4 beside the baseline estimate. The frozen estimate is the optimum with the lowest objective among ten random starting values of the estimator, on the first 80 quarters of the sample (the row Quarters), within the wider parameter bounds of Table A.2, with model-consistent expectations (the anchoring weight at one) and c_π held at the normalization's value of 3.0, which carries no standard error; the full-sample estimate is the baseline estimate of Table 1. Standard errors are computed from the finite-difference Jacobian of the standardized innovations, as in Table 1. In both estimates the two IS weights lie on the cap of 0.97 on their sum, and their standard errors are computed as if the weights were free (Section 4.3). SD is a standard deviation; the intercept c_0 and σ_η are in percentage points. Sources: as Table 1; author's calculations.

Block parameter	Frozen at 2015:Q4		Full sample	
	Estimate	SE	Estimate	SE
IS: lagged gap, a_{lag}	0.883	0.074	0.771	0.030
IS: expected gap, a_{lead}	0.087	0.083	0.199	0.034
Phillips: expectation weight, λ_1	0.693	0.024	0.946	0.041
Phillips: slope, κ	0.0172	0.0037	0.0596	0.0166
Rule: smoothing, ρ	0.774	0.031	0.656	0.027
Rule: inflation response, c_π	3.00	—	3.00	—
Rule: gap response, c_y	1.31	0.27	0.84	0.28
Rule: intercept (pp), c_0	0.58	0.21	3.00	0.80
Rule shock: persistence, ϕ_u	0.572	0.067	0.959	0.016
Rule shock: SD (pp), σ_η	0.324	0.047	0.947	0.103
Quarters	80		122	

Table 7 gives the result, with the rolling agent of Section 8.3 beside the two agents. Over the 42 quarters from 2016:Q1 through 2026:Q2, the frozen agent's mean difference from the realized federal funds rate is +0.212 pp (standard error 0.199 pp; the ratio of the two, 1.07, is below the critical value of 2.78), with an RMS difference of 0.547 pp, against a mean difference of +1.269 pp for the baseline agent (standard error 0.348 pp; ratio 3.64, which exceeds the same critical value, as the ratio does under the three other treatments; Table C.1). The frozen agent's ratio is below its critical value in each of the six periods of Table 7 other than the second ELB episode, at 0.05 to 1.47. The paired difference, the frozen agent's prescription less the baseline agent's at the same decision point, from which the realized rate drops out, has a mean of -1.057 pp and a standard error of 0.449 pp, a ratio of 2.35 against a critical value of 2.78, the value for the four terms the estimator's rule gives for 42 quarters. Table C.2 in Appendix C reports the paired differences; the absolute value

of a paired difference's mean is the distance between the two agents' mean differences. The 42 quarters divide at 2020:Q4, the final quarter for which the staff's estimates of the gap are among the observations (Section 2). Through it the paired difference has a mean of −0.195 pp and a standard error of 0.152 pp, a ratio of 1.28 against 4.30; from 2021:Q1 through 2026:Q2 it has a mean of −1.840 pp and a standard error of 0.262 pp, a ratio of 7.03 against 3.18. The ratio of the paired difference's mean to its standard error thus does not exceed its critical value over the 42 quarters or through 2020:Q4, and it exceeds it from 2021:Q1. Of the sum of the paired differences over the 42 quarters, 91 percent falls in the 22 quarters from 2021:Q1. Between the two periods the baseline agent's mean difference rises from +0.627 to +1.853 pp, and the frozen agent's falls from +0.432 to +0.013 pp (Table 7); of the change in the paired difference's mean, the baseline agent's rise is the larger part.

The division at 2020:Q4 is one of 35 that leave four quarters or more on each side, four quarters being the fewest for which the estimator's rule gives a term; the earlier period ends at a quarter from 2016:Q4 through 2025:Q2. The ratios that follow are at the state of Section 2 until a treatment is named. For each of the 24 divisions from 2019:Q3 through 2025:Q2 the ratio does not exceed its critical value over the earlier period and exceeds it over the later. For the ten divisions from 2017:Q1 through 2019:Q2 neither period's ratio exceeds its critical value, and at 2016:Q4, where the earlier period holds four quarters and the estimator one term, the earlier period's ratio exceeds its critical value and the later period's does not. From the division at 2023:Q3 the later period holds fewer than 12 quarters and the estimator one term, as the earlier period does through 2018:Q3. The earlier period's ratio does not exceed its critical value at any division after 2020:Q4, where that period includes from one to 18 of the 22 quarters from 2021:Q1; at 2023:Q3 it is 1.56 against 3.18. The first of the 24 divisions and the division before it carry ratios near their critical value: the later period's is 3.30 against 3.18 at 2019:Q3 and 3.13 against the same value at 2019:Q2. A change therefore cannot be dated at 2020:Q4 from these ratios, and the paper reports no test of the change between the two periods' means. With the decision quarter's inflation withheld from the states (Section 8.1), the ratios over the 42 quarters, through 2020:Q4, and from 2021:Q1 are, in that order, 2.42, 1.61, and 6.55, and at the states that carry the vintages' inflation 2.43, 1.52, and 6.21, each on the same side of its critical value. Under each of the two treatments the earlier period's ratio does not exceed its critical value and the later period's exceeds it at the same 24 divisions. Before them the later period's ratio at 2018:Q2 stands near its critical value of 2.78: it is 2.75 at the state of Section 2, 2.81 with the inflation withheld, which exceeds that value, and 2.78 with the vintages' inflation, equal to the critical value at two decimals, so that the paper does not place it on either side; the earlier period's ratio there does not exceed its critical value at any of the three, so that with the inflation withheld the same pair of results holds at 2018:Q2 as well. With the inflation withheld the division at which the earlier period's ratio alone exceeds its critical value is 2017:Q1 and not 2016:Q4. With the decision quarter's shocks at the filter's estimates the frozen and the baseline agent's mean differences over the 42 quarters are +0.284 and +1.230 pp (Section 8.1), and the paper does not compute the paired difference's standard error under that treatment. At the state of Section 2, the paired difference over 2021:Q1 through 2023:Q4 alone, with two terms, has a mean of −1.620 pp and a standard error of 0.380 pp, a ratio of 4.26 against 4.30.

The two agents' prescribed rates differ by 1.41 pp root mean square and have a correlation of 0.95. The frozen agent prescribes the ELB from 2020:Q3 through 2021:Q4, as the realized rate stands at it, and prescribes a rate above it from 2022:Q1, by 1.8 bp in that quarter; with the decision quarter's inflation withheld from the states, its prescribed rate in that quarter is at the ELB. The realized rate, which is measured before each quarter's final meeting (Section 2), stands above the ELB from 2022:Q2, one quarter later. The frozen agent's prescribed rate then exceeds the realized

rate in every quarter through 2023:Q2, by at most 1.02 pp; over 2021:Q1 through 2023:Q4 it lies above the realized rate in 7 of 12 quarters, with a mean difference of +0.278 pp against the baseline agent's +1.898 pp. From 2023:Q3 it lies below the realized rate in 10 of 12 quarters, with a mean difference of −0.272 pp, and over 2022:Q2 through 2026:Q2 its mean difference is +0.015 pp.

Table 7. The frozen agent, the rolling agent, the baseline agent, and the volume's IFB rule against the realized federal funds rate after 2015:Q4. Mean and RMS are the mean and root-mean-square difference between the rate the rule prescribes for the decision quarter and the realized federal funds rate (Section 2), in percentage points; the frozen agent's rate is computed in the model at the frozen estimate, the rolling agent's in the model re-estimated through each decision point with its rule retrained there, and the other two in the model at the baseline estimate, each at the real-time state of the quarter. The IFB rule is the volume's. Low and High are the lowest and highest of the eight run numbers' mean differences, each run number's checkpoints taken across the windows (the runs of every window being numbered one through eight); the rolling agent's mean is that of the checkpoint selected in each window, and since the selected run changes from window to window, that mean can lie outside the range of the eight run numbers' mean differences, as it does over 2022:Q2–2026:Q2. The period 2020:Q2–2022:Q1 is the second ELB episode, and the final two rows divide the quarters at 2020:Q4, the final quarter for which the staff's estimates of the gap are among the observations. Table C.2 reports the paired differences between the agents with their standard errors. Sources: as Table 1; author's calculations.

Period	Quarters	Frozen agent		Rolling agent				Baseline agent		IFB rule	
		Mean	RMS	Mean	RMS	Low	High	Mean	RMS	Mean	RMS
2016:Q1–2026:Q2	42	+0.212	0.547	+1.366	1.730	+1.357	+1.369	+1.269	1.501	+0.472	0.657
2016:Q1–2020:Q1	17	+0.447	0.650	+0.569	0.779	+0.567	+0.581	+0.659	0.705	+0.541	0.566
2020:Q2–2022:Q1	8	+0.132	0.369	+0.807	1.160	+0.802	+0.817	+0.924	1.243	+1.086	1.169
2021:Q1–2023:Q4	12	+0.278	0.493	+2.105	2.394	+2.084	+2.105	+1.898	2.075	+0.524	0.813
2022:Q2–2026:Q2	17	+0.015	0.504	+2.425	2.480	+2.404	+2.421	+2.041	2.083	+0.114	0.323
2016:Q1–2020:Q4	20	+0.432	0.643	+0.560	0.795	+0.558	+0.570	+0.627	0.715	+0.608	0.692
2021:Q1–2026:Q2	22	+0.013	0.443	+2.098	2.266	+2.084	+2.098	+1.853	1.958	+0.349	0.624

Figure 9 draws the paths. At 2020:Q2, the first quarter of the second ELB episode, the frozen agent and the baseline agent prescribe rates above the ELB, at which the realized rate stands, and both prescribe the ELB the quarter after. The frozen agent then stays at the ELB with the realized rate, leaves it a quarter earlier, and from 2023:Q3 stays near the realized rate, where the baseline agent's mean difference over 2022:Q2 through 2026:Q2 is +2.041 pp.

The frozen estimate's Phillips slope is 0.0172 and the full-sample estimate's 0.0596. The two estimates also differ in the Phillips equation's weight on expected inflation, λ_1 (0.693 against 0.946), in the IS weights a_{lag} and a_{lead} (0.883 and 0.087 against 0.771 and 0.199), in the estimated policy equation's gap response c_y (1.31 against 0.84), smoothing ρ, intercept c_0, and shock persistence ϕ_u and standard deviation σ_η (Table 6), and in the residual components and the measurement-error moments with which they are estimated.

The distance between the two agents' mean differences has three sources, which can be exchanged one at a time: the trained rule; the IS and Phillips parameters with which the model is solved forward from a state; and the lagged output gap of that state, which each estimate's filter infers. The four lags of inflation, the lagged rate, and the neutral-rate path of a state are the same for the two agents. Table 8 reports the mean difference from the realized federal funds rate for the eight combinations. With the rule and the parameters held at either agent's, replacing the baseline estimate's filtered gap by the frozen estimate's lowers the mean difference over the 42 quarters by 1.226 to 1.642 pp, which is more than the 1.057 pp between the two agents. Replacing the baseline

agent's rule by the frozen agent's raises it by 0.024 to 0.375 pp, and replacing the baseline estimate's parameters by the frozen estimate's raises it by 0.145 to 0.249 pp. Averaged over the six orders in which the three can be exchanged, the filtered gap lowers the mean difference by 1.428 pp, and the rule and the parameters raise it by 0.187 pp and 0.184 pp. A combination that pairs one estimate's filtered gap with the other estimate's parameters is an accounting of the distance and not a model that either estimate describes.

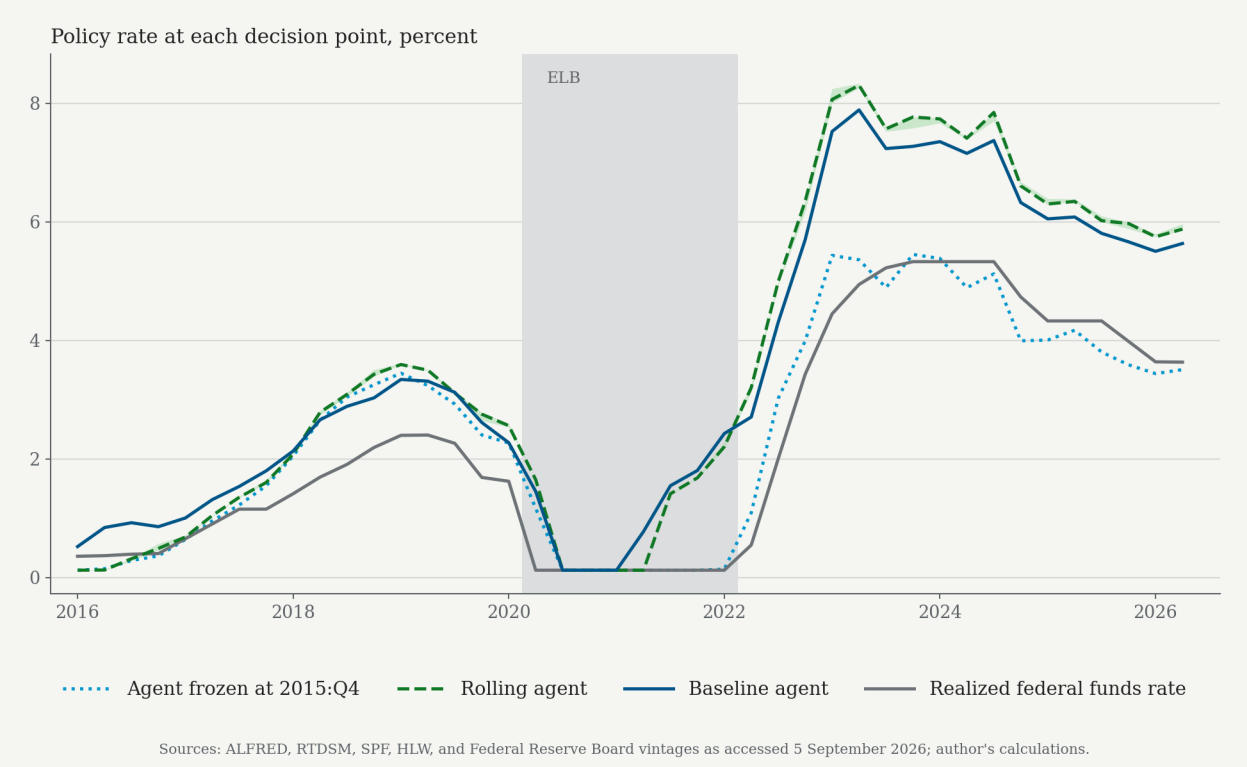


Figure 9. The agent frozen at 2015:Q4, the rolling agent, and the baseline agent against the realized federal funds rate at the 42 decision points from 2016:Q1 through 2026:Q2. Vertical axis: policy rate, in percent; horizontal axis: decision quarter. The bright blue dotted line is the rate the frozen agent prescribes for the decision quarter in the model at the frozen estimate; the deep blue solid line is the baseline agent's rate in the model at the baseline estimate, as in Figure 8; the green dashed line is the rolling agent's rate, re-estimated and retrained at each decision point on the data through that point, and the green band the range of its eight checkpoints' rates at that point; the gray line is the realized federal funds rate (Section 2), set to the model's lower bound of 0.125 percent in the second ELB episode (shaded). The frozen agent's mean difference from the realized rate is +0.212 pp (standard error 0.199 pp) against the baseline agent's +1.269 pp on the same quarters; the text reads the distance between the two by period. Sources: ALFRED, RTDSM, SPF, HLW, and Federal Reserve Board vintages as accessed 5 September 2026; author's calculations.

The two filters agree on the lagged gap to within 0.62 percentage points of potential output at every decision point from 2016:Q1 through 2020:Q4, the final quarter for which the staff's estimates of the gap are among the observations (Section 2). At the decision points of 2021:Q1 and 2021:Q2 the baseline estimate's filtered gap exceeds the frozen estimate's by 0.39 and 1.46 percentage points of potential output; from 2021:Q3 the excess is 2.76 to 4.60 percentage points of potential output. From 2022:Q3 through 2026:Q2 the baseline estimate's filtered gap lies between

2.43 and 5.42 percent of potential output and the frozen estimate's between −0.52 and 1.04 percent. No observation of the gap enters either filter in those quarters: the gap is inferred through each estimate's equations from inflation, the federal funds rate, the two-year rate, and the survey. In Table 8 the distance between the two agents' mean differences from 2021:Q1 is accounted for by the two estimates' inferences of the output gap. Averaged over the six orders of exchange, the filtered gap lowers the mean difference by 2.572 pp over 2021:Q1 through 2026:Q2, against a distance of 1.840 pp between the two agents, and by 2.583 pp over 2021:Q1 through 2023:Q4, against 1.620 pp; the rule raises it by 0.206 and 0.279 pp, and the parameters raise it by 0.526 and 0.684 pp. The distance is not a difference between what the two rules prescribe at one state: at the baseline estimate's states the frozen agent's rule prescribes the higher rate on average. Over 2016:Q1 through 2020:Q4, when the two agents' mean differences stand 0.195 pp apart, the filtered gap and the parameters lower the mean difference by 0.169 and 0.192 pp, and the rule raises it by 0.166 pp.

Table 8. The exchange of the trained rule, the model's parameters, and the filtered output gap between the baseline agent and the frozen agent. Each row is the mean difference between the prescribed rate and the realized federal funds rate (Section 2), in percentage points, over the decision points of the period, with the rule, the IS and Phillips parameters with which the model is solved, and the lagged output gap of each state taken from the agent named; the columns are the 42 quarters from 2016:Q1, the two periods of their division at 2020:Q4, and 2021:Q1 through 2023:Q4. The first row is the baseline agent and the final row the frozen agent. Sources: as Table 1; author's calculations.

Rule	Parameters	Filtered gap	2016:Q1–2026:Q2	2016:Q1–2020:Q4	2021:Q1–2026:Q2	2021:Q1–2023:Q4
Baseline	Baseline	Baseline	+1.269	+0.627	+1.853	+1.898
Baseline	Baseline	Frozen	+0.043	+0.474	−0.349	−0.278
Baseline	Frozen	Baseline	+1.480	+0.411	+2.452	+2.706
Baseline	Frozen	Frozen	+0.188	+0.271	+0.114	+0.313
Frozen	Baseline	Baseline	+1.606	+0.795	+2.343	+2.466
Frozen	Baseline	Frozen	+0.067	+0.603	−0.419	−0.259
Frozen	Frozen	Baseline	+1.855	+0.621	+2.977	+3.292
Frozen	Frozen	Frozen	+0.212	+0.432	+0.013	+0.278

Exchanged between the two estimates, both in the rule's training and in the model with which its rate is computed, the Phillips slope does not account for the distance between the two agents' mean differences. Two further rules exchange the slope between the estimates, each trained as its agent is and evaluated in its agent's estimate with the slope so replaced. The first is trained on the baseline estimate with the Phillips slope replaced by the frozen estimate's; its mean difference over the 42 quarters is +1.421 pp (standard error 0.385 pp), 0.152 pp above the baseline agent's, a paired difference with a standard error of 0.041 pp and a ratio of 3.71 against 2.78. The second is trained on the frozen estimate with the slope replaced by the baseline estimate's; its mean difference is +0.117 pp (standard error 0.119 pp), 0.096 pp below the frozen agent's, a paired difference with a standard error of 0.115 pp and a ratio of 0.83 against the same critical value. The two agents' mean differences stand 1.057 pp apart; the exchanges move the mean difference by 14 and 9 percent of that distance, each away from the other agent's. Each rule is evaluated at the states that its own agent's estimate filters, so the exchange of the slope leaves each estimate's filtered gap of Table 8 unchanged. The paper does not separate which of the two estimates' differences, in the parameters (the slope among them) and in the error moments and the survey's wedge that their filters read,

produce the divergence of the two filtered gaps. The first rule's response to the inflation gap is 0.980, against the baseline agent's 3.353, and the second's is 2.113, against the frozen agent's 1.924. The first rule's selected checkpoint is at step 3,990 of 4,000, within the final 200 steps of its search. The procedure of Section 7.3 extends the runs once, and the first rule's runs stand at that extension; whether a longer search would select a rule with another response to the inflation gap, and with it another mean difference, is not established.

The full-sample comparison states what the loss (1) and the estimated model would have prescribed with the data of the whole sample; the frozen comparison states what the same procedure would have prescribed with the data available at 2015:Q4. From 2021:Q1 each comparison also carries its own estimate's inference of an output gap that no observation enters.

8.3 The agent re-estimated and retrained each quarter

The frozen agent conditions every prescription on data through one date, and the baseline agent conditions every prescription on the whole sample. The rolling agent conditions each prescription on the data available at its decision point: at each of the 82 quarterly decision points from 2006:Q1 through 2026:Q2, the model is re-estimated on the data through that point by the baseline estimation procedure, from 10 random starting values within the wider bounds of Table A.2 and from the previous decision point's optimum; the rule is retrained on that estimate by the procedure of Section 7.3; and the rate it prescribes for the decision quarter is computed at that point's real-time state, filtered under that estimate. The neutral-rate path is real-time from 2016:Q1: each quarter's path is the HLW vintage available at its decision point, the final vintage before the suspension of the HLW estimates' publication during the pandemic being held through the suspension. Before 2016:Q1 the path is the ex post one, since no real-time vintage exists. The rolling agent is therefore real-time in its estimate, its state, and its filter throughout and in the neutral rate from 2016:Q1, apart from core PCE inflation and the term premium, which it reads as every agent does (Section 2). Its specification is that of every agent in the paper; what differs is its information set, the data available at each decision point with the exceptions just stated. The first decision point, 2006:Q1, falls forty quarters into the sample, and the span includes the financial crisis and the first ELB episode.

Table 7 and Figure 9 report the rolling agent beside the other two after the cutoff. Over the 42 quarters from 2016:Q1 through 2026:Q2, its mean difference from the realized federal funds rate is +1.366 pp (standard error 0.474 pp), against the frozen agent's +0.212 pp and the baseline agent's +1.269 pp. The ratio of the rolling agent's mean difference to its standard error, 2.88, exceeds the critical value of 2.78. Its mean difference lies 0.097 pp above the baseline agent's, a paired difference with a standard error of 0.144 pp and a ratio of 0.67, which does not exceed that critical value (Table C.2 reports it by period, and the ratio exceeds its critical value in none of the periods the table gives for this pair). The rolling agent's mean difference is +0.569 pp over 2016:Q1 through 2020:Q1, against +0.567 to +0.581 pp for the eight run numbers taken across the windows (below), and +2.105 pp over 2021:Q1 through 2023:Q4, whose windows' samples include 2021, against +2.084 to +2.105 pp; the standard errors of the two means are 0.360 and 0.780 pp, and neither ratio, 1.58 and 2.70, exceeds the critical value of 4.30. The paper reports no test of the change between those two means. Over 2022:Q2 through 2026:Q2 and over 2021:Q1 through 2026:Q2 the rolling agent's ratios are 8.12 against 4.30 and 4.52 against 3.18, and each exceeds its critical value. The frozen agent's mean differences over 2016:Q1 through 2020:Q1 and over 2021:Q1 through 2023:Q4 are +0.447 pp and +0.278 pp (Table 7). Over the 40 quarters from 2006:Q1 through 2015:Q4, which include the financial crisis and the first ELB episode, the rolling agent's mean difference is +1.7 bp (RMS difference 0.19 pp); 28 of these quarters are at the ELB,

and no standard error is reported for the span. Across the windows the estimate's λ_1 rises from 0.69 at 2021:Q2 to 0.99 at 2021:Q4 as the data of 2021 enter the sample, and it stays below its structural limit of one in every window, as the baseline estimate's does; at the final window the optimum stands at 0.1335 per observation row, with λ_1 at 0.93 and κ at 0.0498, under the real-time neutral-rate path. The final window uses the same data as the baseline estimate except the neutral-rate path, which is the real-time one, and it fits better than the baseline estimate by 9.15 log-likelihood units, with ϕ_u at its bound; the baseline estimate keeps the ex post path (Section 2). Appendix B takes up two questions on the rolling agent and one on the rolling estimate's λ_1.

Each rolling window's agent is selected as in Section 7.3, from eight training runs by a validation set, so the 82 windows show how the selection performs across many estimates. The selection is read in three ways. The first is whether a window's validation set ranks its eight checkpoints as its test set does. At the median across the windows, it does not: the median Spearman correlation of the checkpoints' validation and test losses is -0.31, the correlation is positive in 30 of the 82 windows, and the selected run ranks in the better half on the test set in 32. The second is whether a larger selection set, the validation and test sets together, ranks the checkpoints better. It is answered against a fresh set of the same size as the validation set, drawn for every window, whose draws enter no training episode and which enters no selection. The larger set ranks the checkpoints better than the validation set and less well than the test set: across the windows, the median Spearman correlation of a checkpoint's loss on the fresh set is 0.00 with the mean of its validation and test losses, against +0.20 with its test loss alone and -0.24 with its validation loss alone. A selection on the mean of the validation and test losses ranks in the better half on the fresh set in 39 windows, and in the 69 windows where it differs from the validation selection it has the lower loss on the fresh set in 40; an even chance reaches or exceeds these counts with probabilities of 0.71 and 0.11. The validation selection ranks in the better half on the fresh set in 32 of the 82 windows, fewer than the 41 a random choice among the eight would give on average; an even chance falls to or below that count with a probability of 0.03. On the fresh set, at the median across the windows, the highest of the eight checkpoints' losses lies within 0.19 percent of the lowest, and the selected checkpoint's loss exceeds the lowest by 0.11 percent. The third is whether the rolling agent's prescriptions depend on the selection. At each window each of the eight checkpoints prescribes a rate at the window's state; run k of every window, for k from one through eight, is trained on the k-th of the window's eight random-number streams, which differ from window to window. The mean differences are computed for each run number across the windows. For all eight run numbers, the mean difference over 2006:Q1 through 2015:Q4 lies between +1.6 and +1.7 bp; the mean difference over 2016:Q1 through 2020:Q1 (+0.567 to +0.581 pp) is smaller than that over 2021:Q1 through 2023:Q4 (+2.084 to +2.105 pp); and the mean difference from 2016:Q1 (+1.357 to +1.369 pp) lies above the baseline agent's +1.269 pp. The lowest and highest of the eight prescribed rates are a median of 0 bp apart over the windows through 2015:Q4 and 4 bp over 2016:Q1 through 2020:Q1. Over 2021:Q1 through 2023:Q4 the median is 8 bp, and of all the windows they lie farthest apart at that of 2023:Q1, by 0.24 pp. Figure 9 shades this range from 2016:Q1, and Table 7 gives the lowest and highest of the eight run numbers' mean differences. The rolling agent's prescriptions thus depend on the selection by at most 0.24 pp at any window, and the mean differences carry the same signs and order for every run number.

The rolling agent's prescriptions stay near the realized federal funds rate through 2015:Q4 and lie above it on average from 2016:Q1, by +0.569 pp over 2016:Q1 through 2020:Q1 and by +2.105 pp over 2021:Q1 through 2023:Q4, where the baseline agent's mean differences are +0.659 and +1.898 pp. Neither of the rolling agent's means over those two periods has a ratio that exceeds its critical value, and no test of the change between them is reported. Its ratios from 2016:Q1,

over 2022:Q2 through 2026:Q2, and over 2021:Q1 through 2026:Q2 exceed their critical values, and its mean difference from 2016:Q1 is not distinguished from the baseline agent's by the paired difference. The realized rate after 2015 lies closer, in mean and in root-mean-square difference, to the prescriptions of the rule estimated and trained on data through 2015:Q4 (mean difference +0.212 pp) than to those of the rule re-estimated and retrained each quarter. The rolling agent's prescription less the frozen agent's has a mean of +1.153 pp and a standard error of 0.545 pp over the 42 quarters, a ratio of 2.12 against 2.78, which it does not exceed. Neither the frozen agent's comparison nor the rolling agent's measures what the FOMC learned; both compare its decisions with rules that condition on stated information sets.

8.4 Counterfactual outcomes

Table 9 reports the outcomes each rule would have produced, under a procedure whose limits are stated first. The structural residuals of the IS and Phillips equations are recovered from the realized data at the estimate: at each quarter, the filter's residual states are the innovations the model needs to reproduce the observed inflation and federal funds rate given the filtered gap. Fed back one quarter at a time under the estimated policy equation with its persistent disturbance, from each quarter's filtered lagged gap and its observed lagged inflation and rate, the residuals reproduce the data over the 122 quarters from 1996:Q1: the observed inflation to a root-mean-square error of 0.9 bp (largest error 4.6 bp), and the rate to within 1.03 bp on the 86 quarters at which it is observed above the ELB. They do not reproduce the data on a path whose state advances on the model's own outputs, because the filter revises the lagged gap every quarter and the revisions are not among the residuals: simulated in that way under the estimated policy equation with its disturbance, the baseline estimate's path departs from the observed quarterly inflation by 0.41 pp root mean square and from the filtered gap by 1.58 percentage points of potential output. That path differs from a rule's simulated path in the policy equation alone, the estimated equation with its recovered disturbance standing where the rule stands, and Table 9 reports it for each estimate beside the realized path. The model, without the residual components as in the comparison of Section 8.1, is then simulated from 1997:Q1 under the agent's rule, with those residuals as each quarter's shocks, expectations model-consistent under the rule, future shocks at zero, and the state advancing on the model's own outputs. The persistent disturbance of the estimated policy equation, which is the FOMC's deviation from its estimated reaction function, is not fed back; over 2021:Q1 through 2023:Q4 it averages −5.14 pp, the realized rate standing below the estimated policy equation's prescription. Three conditions bound the comparison. The private sector's expectations are model-consistent under the simulated rule throughout. The residuals are treated as invariant to the rule, the condition the Lucas critique questions.[14] And the gap on the realized path is the filter's fitted gap, because the true gap is latent, so models estimated on different samples measure the realized gap differently. The simulated paths are outcomes in the estimated model, not outcomes the FOMC could have chosen. The rates on a simulated path also differ from the prescriptions of Sections 8.1 through 8.3, which are computed at the realized states: over 2021:Q1 through 2023:Q4 the baseline agent's rate averages 2.78 percent on its simulated path against the realized 2.21 percent, while its prescriptions at the realized states over the same quarters have a mean difference of +1.898 pp.

Under the baseline agent in the model at the baseline estimate, four-quarter core PCE inflation over 2021:Q1 through 2023:Q4 averages 2.33 percent against the realized 4.29 percent, the gap 0.09 percent of potential output against 2.72 percent, and the federal funds rate 2.78 percent

[14] Robert E. Lucas Jr., "Econometric Policy Evaluation: A Critique," *Carnegie-Rochester Conference Series on Public Policy* 1 (1976): 19–46.

against 2.21 percent. In the model, the Phillips equation places a weight of 0.946 on expected inflation, which is model-consistent under the simulated rule; the paper does not decompose the reduction in inflation by channel. The loss (1) summed over 1997:Q1 through 2026:Q2 is 237.2 under the baseline agent, against 518.9 on the path simulated under the estimated policy equation and 989.3 on the realized path. Against the simulated path, the baseline agent's loss is the lower from 2016:Q1 and the higher before: 38.8 against 332.7 over 2016:Q1 through 2026:Q2, and 198.4 against 186.2 over 1997:Q1 through 2015:Q4. Over 2021:Q1 through 2023:Q4 the path simulated under the estimated policy equation has inflation of 4.28 percent and a gap of 2.73 percent of potential output, so the lower inflation on the baseline agent's path comes with the change of rule. Under the frozen agent in the model at the frozen estimate, whose filter puts the realized 2021:Q1 through 2023:Q4 gap at -0.87 percent of potential output rather than 2.72 percent, inflation over the same quarters averages 4.46 percent against the realized 4.29 percent, with a rate of 2.63 percent on the simulated path. The frozen agent's rule thus leaves 2021:Q1 through 2023:Q4 inflation 0.17 pp above the realized level, and 0.25 pp above that of the path simulated under the frozen estimate's policy equation, where the baseline agent's rule lowers it by 1.95 pp against the path simulated under the baseline estimate's policy equation. Over 2016:Q1 through 2026:Q2 the frozen agent's loss is 158.4, against 158.4 on the path simulated under the frozen estimate's policy equation, the two alike at the table's precision, and the realized path's 164.6 under the frozen model's gap; over 1997:Q1 through 2015:Q4, the quarters on which the frozen rule was trained, it is 172.3 against 222.6 and the realized 595.5. The frozen estimate's residuals are recovered by its own filter, which reads the measurement-error moments known at 2015:Q4 (Section 2). The baseline agent's 38.8 is computed with the baseline model's residuals and gap, and the two agents' losses are not ranked across the two models.

The dependence of the reduction in inflation on the expectations in the Phillips equation is measured by re-estimating the model with λ_1 held at the volume's 0.70 and every other free parameter re-estimated, from the baseline estimate and from the alternative estimate as starting values, and running the baseline rule, unchanged, in that model with the residuals recovered at that estimate. The sensitivity estimate is the deeper of the two optima, reached from the alternative estimate; the other lies 1.0 log-likelihood unit below it, at an objective of 0.1820. The sensitivity estimate's objective is 0.1807 against the baseline estimate's 0.1445, a log-likelihood lower by about 30, so the likelihood is higher at the baseline estimate's λ_1; the Phillips slope is 0.0297 (standard error 0.0092) against the baseline estimate's 0.0596, and the policy smoothing 0.714. In the model with $\lambda_1 = 0.70$, 2021:Q1 through 2023:Q4 inflation averages 2.40 percent under the baseline rule against the realized 4.29 percent, above the 2.33 percent of the baseline model and 1.88 pp below the realized level; against the path simulated under that model's estimated policy equation the reduction is 1.93 pp, beside 1.95 pp in the baseline model, with the loss over the sample at 192.3 against 403.6 on the path simulated under that model's estimated policy equation and the realized path's 905.5 under that model's gap; in that model the baseline rule's loss is the lower of the two in both periods, 153.1 against 156.0 over 1997:Q1 through 2015:Q4 and 39.2 against 247.6 from 2016:Q1. At the other optimum, 2021:Q1 through 2023:Q4 inflation averages 2.29 percent under the baseline rule.

A further measure of the gap comes from outside the model: potential output as the Congressional Budget Office (CBO) projected it in its vintage of 27 February 2026, against real gross domestic product (GDP) at its vintage of 15 August 2026, both from the vintage data of Section 2. The realized paths of Table 9 differ only in the gap, and the table's final row takes the gap from the CBO measure. Under the CBO measure the 2021:Q1 through 2023:Q4 gap averages 0.93 percent of potential output, between the frozen model's -0.87 percent and the baseline model's 2.72 percent,

and its mean over 1997:Q1 through 2015:Q4 is −0.76 percent; the realized path's loss over 1997:Q1 through 2026:Q2 under the CBO gap is 594.4. A counterfactual path's gap is read under the CBO measure as the CBO gap of the quarter plus the path's deviation from the realized path under its own model's filter, a reading that assumes the model's counterfactual deviation of the gap carries over to the CBO measure unchanged; the inflation and rate terms of the loss are the path's own. Table C.3 in Appendix C reports the losses under the CBO measure, the paths simulated under the two estimated policy equations read in the same way. Under the CBO measure, the baseline agent's loss lies below the realized path's and below that of the path simulated under the baseline estimate's policy equation before 2016, from 2016:Q1, and over the sample, and the frozen agent's loss lies above that of the path simulated under the frozen estimate's policy equation in each of the three. The CBO measure is a further measure of the gap and not a benchmark for the two filtered gaps.

Table 9. Counterfactual outcomes under the baseline agent and under the frozen agent, each in the model at its own estimate with the residuals recovered at that estimate, and under the baseline agent in the model re-estimated with λ_1 held at 0.70, beside the realized path as each model measures it. Inflation is four-quarter core PCE inflation, in percent; the gap is in percent of potential output; the rate is the federal funds rate, in percent; each is averaged over 2021:Q1–2023:Q4. The loss is the quarterly loss (1) summed over 1997:Q1–2015:Q4, 2016:Q1–2026:Q2, and 1997:Q1–2026:Q2 (total), in the units of equation (1). The realized rows share the observed inflation and rate and differ in the gap, which is each model's filtered gap. A row labeled estimated policy equation is the path simulated by the procedure of a rule's path, with that estimate's own policy equation and its disturbance in the rule's place. In the final row the gap is measured as the log ratio of real GDP to CBO potential output, in percent. Sources: as Table 1; CBO potential output (ALFRED, series GDPPOT, vintage of 27 February 2026) and real GDP (RTDSM, ROUTPUT, vintage of 15 August 2026); author's calculations.

Path	2021:Q1–2023:Q4, means			Loss		
	Inflation	Gap	Rate	1997–2015	2016–2026	Total
Realized, gap of the baseline model	4.29	+2.72	2.21	645.8	343.5	989.3
Estimated policy equation, baseline model	4.28	+2.73	2.29	186.2	332.7	518.9
Baseline agent, baseline model	2.33	+0.09	2.78	198.4	38.8	237.2
Realized, gap of the frozen model	4.29	−0.87	2.21	595.5	164.6	760.1
Estimated policy equation, frozen model	4.21	−1.70	2.32	222.6	158.4	381.1
Frozen agent, frozen model	4.46	−1.33	2.63	172.3	158.4	330.7
Realized, gap of the $\lambda_1 = 0.70$ model	4.29	+1.88	2.21	665.4	240.2	905.5
Estimated policy equation, $\lambda_1 = 0.70$ model	4.34	+2.33	2.55	156.0	247.6	403.6
Baseline agent, $\lambda_1 = 0.70$ model	2.40	−0.26	2.13	153.1	39.2	192.3
Realized, gap from CBO potential output	4.29	+0.93	2.21	370.9	223.5	594.4

9 Sensitivity and Limitations

The loss weights are the volume's and are not estimated; Table 10 reports how the trained rule and its mean difference from the realized federal funds rate change with them. The baseline agent's training and selection were repeated under four variants of the weights: β halved and raised by half, and γ halved and raised by half. The first of these is the volume's own robustness check, DM2, the dual-mandate loss with the output-gap weight halved, a halving that the volume reports "makes hardly any difference to the outcome" of its policy simulations.[15] In every variant α is held

[15] Adrian, Laxton, and Obstfeld, *Advancing the Frontiers*, 223.

at 1.0 as the normalization of the loss, and the draws around the estimate, the episodes' starting states, the validation and test sets, and the random-number streams of the baseline training are unchanged, so that a variant differs from the baseline agent in the weights alone; each variant's training selects its own checkpoint under its own weights, and each is evaluated as the baseline agent is.

Table 10. The baseline agent (row Baseline) and four agents trained under other loss weights, everything else unchanged. The weights are (α, β, γ) with $\alpha = 1.0$; the baseline weights are $\beta = 1.0$ and $\gamma = 0.50$, and each variant changes the one it names. Run and step are the selected training run among the row's eight and the step of its checkpoint, every row's runs being extended from 2,000 to 4,000 steps; ρ, c_π, and c_y are the smoothing, inflation, and gap responses; the mean difference is the agent's prescribed rate minus the realized federal funds rate over the 118 decision points, in percentage points, with its standard error by the equal-weighted cosine estimator of Section 8.1 (SE) and the RMS difference, in percentage points. Loss, own is the loss (1) on the simulated path from the volume's 2014:Q4 state in the model at the baseline estimate, under the row's own weights; IFB, own is the volume's IFB rule's loss on its own path from the same state under the same weights; loss, baseline is the agent's loss under the baseline weights, where the IFB rule's is 9.944. Sources: as Table 1; author's calculations.

Weights	Run	Step	ρ	c_π	c_y	Mean diff.	SE	RMS	Loss, own	IFB, own	Loss, baseline
Baseline	4	3,410	0.736	3.353	3.159	+0.723	0.213	1.035	6.106	9.944	6.106
$\beta = 0.5$	1	3,860	0.764	2.278	2.576	+0.580	0.152	0.794	3.592	5.424	6.180
$\beta = 1.5$	4	3,410	0.716	4.574	3.602	+0.839	0.266	1.244	8.484	14.464	6.058
$\gamma = 0.25$	4	3,410	0.716	4.434	4.189	+0.906	0.303	1.386	5.710	9.806	6.013
$\gamma = 0.75$	4	3,410	0.746	2.670	2.714	+0.638	0.174	0.883	6.463	10.082	6.215

Raising the weight on the gap, or lowering the weight on the rate change, lowers the smoothing and raises both responses, and in each of the three coefficients the baseline lies between the two variants of each weight. As β rises from 0.5 to 1.5, ρ falls from 0.764 to 0.716, and c_y rises from 2.576 to 3.602; as γ falls from 0.75 to 0.25, ρ falls from 0.746 to 0.716, and c_y rises from 2.714 to 4.189. Every move of ρ, c_π, or c_y from the baseline to one of the four variants exceeds the coefficient's spread across the five trainings at the baseline weights (Table 11) but one, that of ρ from $\gamma = 0.50$ to $\gamma = 0.75$, which lies within it. Across the grid c_y runs from 2.576 to 4.189, and across the five trainings from 3.114 to 3.201: the grid's range in the gap response is 19 times the trainings', so at this grid's spacing the weights' effect on that coefficient is resolved, that is, distinguishable from the dispersion across repetitions of the training. The prescriptions move less than the coefficients. The mean difference from the realized federal funds rate over the 118 decision points (Section 2) runs from +0.580 pp at $\beta = 0.5$ to +0.906 pp at $\gamma = 0.25$; every variant lies within 0.19 pp of the baseline +0.723 pp, and every variant's ratio of the mean difference to its standard error exceeds its critical value, the smallest being 2.98 against a critical value of 2.26. Every variant has a lower loss than the volume's IFB rule from the volume's 2014:Q4 state in the model at the baseline estimate under its own weights. Under the baseline weights, two variants have a lower loss there than the baseline agent's 6.106, the variant with the gap weight raised by half at 6.058 and the variant with the rate-change weight halved at 6.013, each within the 0.095 by which the five trainings' losses spread at the baseline weights: the baseline agent minimizes its own objective to within the dispersion of the training procedure, and not more precisely. Closeness to the realized federal funds rate is not a criterion for the weights, the objective being the loss as stated, and on that measure the variants differ: the variant nearest the realized rate is the one with the gap weight halved, with an RMS difference of 0.794 pp against the baseline 1.035 pp, and the farthest the one with the rate-change weight halved, at 1.386 pp. The variant with $\beta = 0.5$ stands at a checkpoint within the final 200

steps of its search, step 3,860 of 4,000, as the first rule of the slope exchange does (Section 8.2); whether a longer search would select another rule under those weights is not established. No other row of Table 10 or of Table 11 stands at such a checkpoint.

Training is a further source of variation, and Table 11 reports it. The baseline agent's training and selection were repeated under four further sets of random-number streams, everything else unchanged, each over eight training runs of 2,000 steps under the same extension rule, which extends the runs to 4,000 steps in Repetitions 1 through 3 and not in Repetition 4, and each evaluated as the baseline agent is. The coefficients vary: c_π spans 3.006 to 3.614 and c_y 3.114 to 3.201 across the five, with ρ between 0.725 and 0.736. The prescriptions vary less: the mean difference from the realized federal funds rate over the 118 decision points runs from +0.723 pp to +0.755 pp, a range of 3.2 bp, with standard errors of 0.213 to 0.225 pp, and every training's loss from the volume's 2014:Q4 state in the model at the baseline estimate lies between 6.011 and 6.106, against the volume's IFB rule's 9.944. The mean difference from the realized federal funds rate is a property of the estimated model and the loss, not of one training.

Table 11. The baseline agent (row Baseline agent) and four agents trained under other random-number streams (rows Repetition 1 through Repetition 4), everything else unchanged. Run and step are the selected training run among each repetition's eight and the step of its checkpoint, the runs being extended from 2,000 to 4,000 steps for the baseline agent and Repetitions 1 through 3 and not for Repetition 4; test rank is the selected run's rank of eight on the shared test set; ρ, c_π, and c_y are the smoothing, inflation, and gap responses; the mean difference is the agent's prescribed rate minus the realized federal funds rate over the 118 decision points, in percentage points, with its standard error by the equal-weighted cosine estimator of Section 8.1 (SE); the loss is the loss (1) on the simulated path from the volume's 2014:Q4 state in the model at the baseline estimate, where the volume's IFB rule's is 9.944. Sources: as Table 1; author's calculations.

Agent	Run	Step	Test rank	ρ	c_π	c_y	Mean diff.	SE	Loss
Baseline agent	4	3,410	8	0.736	3.353	3.159	+0.723	0.213	6.106
Repetition 1	5	3,790	6	0.725	3.584	3.176	+0.755	0.224	6.011
Repetition 2	4	3,420	7	0.726	3.614	3.201	+0.754	0.225	6.056
Repetition 3	5	3,790	7	0.735	3.498	3.178	+0.729	0.216	6.064
Repetition 4	4	1,550	8	0.728	3.006	3.114	+0.731	0.214	6.035

The following limitations apply. First, under the baseline estimate the IS residual carries structure beyond the null distribution that includes an ELB episode (Section 6), so the IS equation is incomplete in a way the residual components measure and the baseline agent's training environment does not include. Second, the baseline estimate's Phillips equation is nearly forward-looking, with λ_1 at 0.946, so that inflation in the estimated model depends mainly on model-consistent expectations of it. The counterfactual reduction in inflation of Section 8.4 is nearly the same with λ_1 held at the volume's 0.70, where the baseline rule lowers 2021:Q1 through 2023:Q4 inflation by 1.93 pp against the path simulated under that model's estimated policy equation, beside 1.95 pp at the baseline estimate; the paper measures the reduction at these two values of λ_1 alone, and the counterfactual does not feed back the persistent policy disturbance. Third, the loss has the target-deviation form of the volume's chapter 12. The framework the FOMC adopted in 2020 added history dependence, an aim of inflation that averages 2 percent over time, and the updated Statement on Longer-Run Goals and Monetary Policy Strategy that the FOMC approved on 22 August 2025, inside the sample, omits that aim.[16] The agent's prescriptions over 2021:Q1 through

[16] Board of Governors of the Federal Reserve System, "Federal Open Market Committee Announces Approval of Updates to Its Statement on Longer-Run Goals and Monetary Policy Strategy," press release (Washington, DC: Board of Governors of the Federal Reserve System, August 22, 2025).

2023:Q4 are therefore a statement about the loss (1) as much as about the agent. Fourth, the ELB episodes of the Monte Carlo study last 7 to 17 quarters under the anchored specification, 5 to 12 under cognitive discounting, and 14 to 43 under the learned anchor, against the sample's first episode of 28 quarters. The null distribution of Section 6 is simulated under condition O16, and the five of its economies whose episodes the paper measures are the study's samples of that condition. The study's estimates of the anchor's weight, the cognitive discount, and the learned anchor's gain are for episodes of those lengths, four of the five estimates of the cognitive discount lie above its true value, and the study does not show the recovery of the sum of the two IS weights or of the policy equation's inflation and gap responses (Sections 4.2 and 5 and Appendix A). Fifth, the estimated policy equation is the FOMC's reaction function as estimated on the data, and the agent's rate at each state is computed under expectations consistent with the agent's rule, so the comparison holds the private sector's way of forming expectations, model-consistent under each policy, fixed across the two policies; this is the condition of Section 8.4. Sixth, the scale of the estimated policy equation, which the likelihood determines only weakly, is fixed by normalizing its inflation response to $c_\pi = 3.0$, not estimated. Along the profile of Section 4.3 the objective falls to 0.1349 at $c_\pi = 8$, a log-likelihood higher by about 8 than at the normalization. Every agent is trained on an estimate at the normalization, and the paper does not report how the agents' prescriptions change with the normalized value; the normalization is a choice the paper states and does not estimate. Seventh, the forward solution at a decision point reads the neutral-rate path of Section 2 from that quarter on, so the neutral rate on which expectations are formed in the states of Section 8 is not real-time, except for the rolling agent of Section 8.3, whose neutral rate is real-time from 2016:Q1 and ex post before. The path runs from 0.57 to 3.90 percent over the sample, and the paper does not construct a real-time path for the baseline or the frozen estimate. The final rolling window, on the same data as the baseline estimate except the real-time neutral-rate path, fits better than the baseline estimate by 9.15 log-likelihood units (Section 8.3); the baseline estimate keeps the ex post path. Eighth, the selected checkpoint is the best of eight training runs on a validation set of 32 draws and eight starting states, and on each search's test set the eight runs' ranking does not follow the validation ranking in any of the eleven searches of Section 7.3, the Spearman correlation of the two being positive in none of them; across the rolling windows of Section 8.3, the validation selection ranks in the better half on a fresh set in 32 of the 82. The prescriptions the paper reports are those of the selected checkpoint, and Table 11 reports how the prescriptions vary with the random-number streams of the training. Ninth, no observation of the output gap enters the model after 2020:Q4, and from 2021:Q1 the gap is inferred through the model's equations. Over 2021:Q1 through 2023:Q4 the baseline estimate's filtered gap averages 2.72 percent of potential output, the frozen estimate's -0.87 percent, and the CBO measure 0.93 percent (Section 8.4); from 2022:Q3 the baseline estimate's filtered gap lies between 2.43 and 5.42 percent and the frozen estimate's between -0.52 and 1.04 percent (Section 8.2). The prescriptions of every agent after 2020, and the difference between the baseline and the frozen agent's prescriptions there, depend on that inference, and the paper does not estimate the model with an observation of the gap in those quarters. The published rules of Section 8.1 read the same filtered gap, the baseline agent's band varies four coefficients and holds the filtered state, and the rolling agent's larger mean difference after 2020, for which no test of the change is reported, falls in the same quarters; none of the three is separated from that inference. Tenth, the standard deviations of the IS and Phillips shocks are assumed (Section 4.2): the estimates, the likelihood comparisons of Sections 4.3, 5, 8.2, 8.3, and 8.4, and the standard errors are conditional on the two values, which the paper does not vary on the sample (they are free only in condition O12 of the Monte Carlo study). The moments of the staff's errors that the filter reads are also inputs fixed before the estimation; at

the baseline estimate, in the rolling windows from 2021:Q1 (from 2022:Q2 for the inflation projection's errors), and in the simulated economies of Sections 4.2 and 6, they include quarters whose errors are computed from the staff's paths of 2020:Q4 at later decision points (Section 2), and the paper does not estimate the model, or simulate those economies, with the moments of the other quarters alone. Eleventh, core PCE inflation is read from one vintage of the price index, the filter reads the decision quarter's own inflation, which no meeting of the quarter had in full, and the term premium is read from one release (Section 2), so the states, the estimates, and the frozen and rolling agents carry later information in those inputs as well as in the neutral-rate path (the rolling agent's only before 2016:Q1). Withholding the decision quarter's inflation from the states changes the baseline agent's mean difference by -0.3 bp over all quarters and by -6.5 bp from 2021:Q1, and reading the preceding quarters' inflation from the vintage available at each decision point as well changes it by -5.6 bp over all quarters, to $+0.666$ pp, and by -16.3 bp from 2021:Q1; from 2021:Q1 the two changes' ratios to their standard errors exceed their critical values, and over all quarters they do not (Section 8.1). Beyond what Sections 8.1 and 8.2 and Table C.1 report under the two treatments, nothing is computed under either: the estimates, the band of Figure 8, the published rules, the exchange of Table 8, the counterfactual paths, and the rolling agent's states stand at the later vintage and with the decision quarter's inflation. Twelfth, the standard errors of the mean differences come from one estimator of the long-run variance, whose rule for the number of terms is derived for series less persistent than the differences over all quarters (Section 8.1). A reading over a period of 12 to 22 quarters rests on two or three terms, and one over fewer than 12 quarters on one. At the state of Section 2, the following ratios stand near their critical values: 4.35 against 4.30 for the baseline agent over 2016:Q1 through 2020:Q1; 4.26 against 4.30 for the paired difference between the frozen and the baseline agent over 2021:Q1 through 2023:Q4; and, on four terms over 42 quarters, 2.88 against 2.78 for the rolling agent from 2016:Q1. The first of these does not exceed its critical value with the decision quarter's shocks at the filter's estimates or with the decision quarter's inflation withheld, and over 2016:Q1 through 2020:Q4 the baseline agent's ratio, 4.20 against 4.30 at the state of Section 2, does not exceed its critical value under three of the four treatments (Table C.1). With the decision quarter's shocks at the filter's estimates the baseline agent's ratio over the 42 quarters from 2016:Q1 is 2.79 against 2.78. The paired difference between the frozen and the baseline agent has three further ratios near their critical values, each the later period's at a division of Section 8.2: 3.30 at 2019:Q3, the first of the 24 divisions at which the earlier period's ratio does not exceed its critical value and the later period's does, and at the division before it, 3.13 at 2019:Q2, both against 3.18; and 2.75 against 2.78 at 2018:Q2, where with the inflation withheld the ratio is 2.81 and exceeds that value. A change cannot be dated at 2020:Q4 from the ratios at those 24 divisions, and no test of the change is reported. The divisions through 2018:Q3 leave the earlier period one term and those from 2023:Q3 the later. Thirteenth, the agent is trained inside a model whose forecasts have a higher root-mean-square error than the Federal Reserve Board staff's on every line of Table 2: for four-quarter inflation four quarters ahead the model's RMSE over the line's 83 quarters is 1.32 times the staff's. The ratios of the loss differentials of the four one-quarter lines to their standard errors do not exceed their critical values; that of the four-quarter line does over its 75 consecutive quarters through 2018:Q3 and does not over the eight quarters after its missing quarter, on one term (Section 4.4). The comparison does not separate a difference in information from a difference in model, and the paper does not measure how the agent's prescriptions depend on the model's forecast accuracy.

10 Conclusion

The volume edited by Adrian, Laxton, and Obstfeld sets, in its chapter on the United States, a policy that minimizes a loss against a calibrated model beside a policy that follows a simple rule. This paper asks what that policy is when the model must be estimated from data, and it measures the four learners the question involves. On the full sample, each of three alternatives to model-consistent expectations in the IS and Phillips equations, a fixed anchor, cognitive discounting of forward expectations, and an anchor learned from experience, is estimated at the model-consistent value, the upper bound of its parameter. In the Monte Carlo study, on simulated economies with an ELB episode and at one true value away from model-consistent expectations for each of the three, the estimates of the fixed anchor's weight and of the learned anchor's gain lie on both sides of their true values, and four of the five estimates of the cognitive discount lie above its true value. The learned level, a constant-gain mean of realized inflation at a gain of 0.030 estimated with the long-run survey expectation taken as an observation of it, varies more over the sample than the survey does, and the survey's measurement error about it has an estimated standard deviation of 0.30 pp, larger than the survey's own; the survey lies below the learned level early in the sample, within 0.21 pp of it from 2010:Q1 through 2022:Q4, and, from 2023:Q1 on, nearer the target than the learned level is. The likelihood determines the scale of the estimated policy equation only weakly, and the paper fixes it by a normalization. Under the normalization, the optimum with the highest likelihood is reached from two of the four non-random starting values and from none of the 40 random starting values, and the estimated model differs between the baseline estimate and the alternative estimate, the optimum reached from the data-driven starting value: the alternative estimate has a Phillips equation with a weight of 0.662 on expected inflation and a slope of 0.0314, and the baseline estimate a nearly forward-looking Phillips equation, with λ_1 at 0.946, a slope of 0.0596 (standard error 0.0166), and a persistent policy shock. The baseline estimate's forecasts have a higher root-mean-square error than the Federal Reserve Board staff's Tealbook projections on every line scored; the ratio of the mean loss differential to its standard error exceeds its critical value on one line alone, that of four-quarter inflation four quarters ahead over its 75 consecutive quarters through 2018:Q3 (Section 4.4). An ELB episode induces structure in the filtered shocks at the true parameters; against a null distribution that includes an ELB episode, the Phillips component shows no evidence of signal beyond what the episode induces, while the IS component's excess, 0.239, exceeds the null's 95th percentile of 0.145. In the paper's reproduction of the volume's model, the learned rule lifts off from the 2014:Q4 state one quarter after the DM1 path computed there, whose rate in 2016:Q1 stands 0.5 bp above the ELB. It has a lower loss than the volume's IFB rule in the estimated model on 100 of 100 draws around the estimate and in the volume's calibration, and in the model at the baseline estimate it prescribes rates above the realized federal funds rate on average in every period of Table 4, over the first ELB episode (2009:Q1 through 2015:Q4) by a mean difference of +1.3 bp; in the two ELB episodes the difference cannot be negative. The ratio of the mean difference to its standard error exceeds its critical value over all quarters and does not over 2021:Q1 through 2023:Q4, a period of 12 quarters; over 2021:Q1 through 2026:Q2 it does. With the decision quarter's shocks at the filter's estimates, with the decision quarter's inflation withheld, and with the preceding quarters' inflation read from the vintage available at each decision point as well, the ratios over all quarters, over 2022:Q2 through 2026:Q2, and over 2021:Q1 through 2026:Q2 exceed their critical values as they do at the state of Section 2; those over 1997:Q1 through 2008:Q4 and over 2016:Q1 through 2020:Q1, which exceed them at that state, do so under two and under one of those three (Section 8.1). Estimated on data through 2015:Q4 and trained by the procedure of Section 7.3, the frozen agent prescribes rates with

a mean difference of +0.212 pp over the 42 quarters that follow, against +1.269 pp for the baseline agent; the ratio of the frozen agent's mean difference to its standard error is below its critical value, and the baseline agent's exceeds it under each of the four treatments. The frozen estimate's Phillips slope, 0.0172, is smaller than the full-sample estimate's, 0.0596, and the two estimates also differ in other parameters; exchanged between the two estimates, both in the rule's training and in the model with which its rate is computed, the slope does not account for the difference between the two agents' prescriptions. The rule trained on the baseline estimate with the frozen estimate's slope stands at a checkpoint within the final 200 steps of its search, and whether a longer search would change that rule's mean difference is not established (Section 8.2). The exchange of the slope leaves each estimate's filtered gap unchanged, and the paper does not separate which of the estimates' differences produce the two gaps' divergence. The rolling agent, re-estimated and retrained each quarter, has a mean difference of +1.366 pp from 2016:Q1; its ratio to its standard error exceeds its critical value, and the paired difference does not distinguish it from the baseline agent's (Section 8.3).

Training the policy on draws around the estimate optimizes the expected loss across those draws rather than the loss at the point estimate. On those draws, an agent trained with the residual components has a higher loss than the volume's IFB rule in the estimated model without the components, and the baseline agent a lower one; with the components in the model, at the one state and estimate evaluated, the agent trained with them has the lower loss of the two (Section 7.2). Optimal monetary policy when the economy must be learned is, in this paper, a rule whose optimality is conditional on the four learners' estimates: the expectations specification, the estimated model, the residual components, and the selected checkpoint. The baseline agent's higher rates are a property of that conditional rule and not an evaluation of monetary policy as the FOMC conducted it. Within the grid of Section 9, the loss weights move the mean difference over all quarters by at most 0.19 pp and do not remove it. The ratio of the mean of the frozen agent's prescription less the baseline agent's to its standard error does not exceed its critical value over 2016:Q1 through 2026:Q2 or over 2016:Q1 through 2020:Q4 and exceeds it from 2021:Q1; at the state of Section 2 the same holds for the earlier and the later period at each division of those quarters from 2019:Q3 through 2025:Q2 and at none before, the later period resting on one term of the estimator from 2023:Q3, so that a change cannot be dated at 2020:Q4 from these ratios, and the paper reports no test of the change between the two periods' means. From 2021:Q1 the difference between the two agents' prescriptions is accounted for, in the exchange of Section 8.2, by the output gaps the two estimates infer in quarters in which no observation of the gap enters either filter, and not by the frozen agent's rule, which at the baseline estimate's states prescribes the higher rate on average, or by the parameters with which the model is solved, whose exchange raises the mean difference on average. The counterfactual outcomes depend on the estimate as well: against the path simulated under each estimate's own policy equation, with the gap that estimate's filter measures, the baseline rule lowers the loss from 2016:Q1 and not before, and the frozen rule lowers it over the quarters on which it was trained, after which the two losses are alike at the precision of Table 9. In the model with λ_1 held at the volume's value the baseline rule has the lower loss in both periods; under the CBO measure of the gap, read at the two estimates, it has the lower loss in both and the frozen rule the higher (Section 8.4). The higher rates rest on the expectations specification at which the likelihood is highest (Section 5) and on the limitations of Section 9, among them a loss without the history dependence of the framework the FOMC adopted in 2020 and the normalized scale of the estimated policy equation.

A Conditions of the Monte Carlo Study

Table A.1. The conditions of the Monte Carlo study. Quarters is the length of each simulated sample and runs the number of estimation runs, one simulated sample each; condition O2 reads the samples of O3, and O5 its samples and its estimates. The true values of the IS, Phillips, and rule coefficients are those of the final column of Table 1, with the volume's IFB rule; the policy shock's persistence and standard deviation match the one-quarter unpredictability of the effective federal funds rate. Source: author's calculations.

Condition	Design	Parameter	True value	Quarters	Runs
O1	Gap observed without error; the rule known	a_r	0.10	96	5
O2	O3's samples; the IS equation by single-equation least squares	a_r	0.10	96	—
O3	Gap observed with the measured errors; the rule estimated; the two-year yield observed	a_r	0.10	96	5
O4	O3 without the two-year yield	SE of a_r	—	96	5
O5	O3's samples and estimates	κ	0.10	96	—
O6	O3 on shorter samples	a_r	0.10	79	5
O7	O3 with a scripted ELB episode, under which the simulated economy has no equilibrium; not estimated	a_r	0.10	96	—
O8	O3 with the three errors of the staff's gap correlated as measured	a_r	0.10	96	5
O9	O3 with the neutral rate observed with the measured real-time error	a_r	0.10	96	5
O10	O3 on the length of the sample before its first ELB episode	a_r	0.10	52	5
O11	O10 with the staff's inflation projection absent where the sample lacks it	a_r	0.10	52	5
O12	O11 with the shock variances of the IS and Phillips equations estimated	σ_y (and σ_π, 0.25)	0.40	52	5
O13	O11 with anchored expectations and the survey row	w	0.60	52	5
O14	O13 with a_r held at its calibrated value	w	0.60	52	5
O15	O14 on full-length samples with an ELB episode	w	0.80	96	5
O16	O15 with the sample's pattern of missing rows	w	0.80	96	5
O17	O16 with cognitive discounting in place of the anchor	m	0.85	96	5
O18	O16 with the anchor learned from experience and the ten-year survey row	gain	0.05	96	5

The Monte Carlo study of Section 4.2 comprises 18 conditions, O1 through O18, and 75 estimation runs on economies simulated at known parameters. Each condition simulates the economy and applies the estimator to the simulated data, with three exceptions: O2 applies single-equation least squares to O3's samples, O5 reads O3's samples and estimates, and O7 is not estimated, its simulated economy having no equilibrium under the scripted ELB episode. Table A.1 lists the conditions, each with what it changes, the parameter it studies at its true value, the length of the simulated samples, and the number of estimation runs; Table A.2 lists the parameter bounds of the estimator on the real-time data. Conditions O1 through O14, O7 excepted, concern the estimator on simulated trajectories without an ELB episode, its treatment of measurement error, data revision, and missing observations, and the anchored expectations block. Conditions O15 through O18 add an ELB episode to the simulated trajectories, from O16 on together with the pattern of missing

observations of the real-time data, and concern the anchored, cognitive-discounting, and learned-anchor blocks in turn. With an ELB episode and the staff's gap rows present in every quarter (O15), the anchoring weight is estimated at 0.769 to 0.809 on five simulated samples against a true value of 0.80, with standard errors of 0.037 to 0.046. With the staff observations also absent in the final 22 quarters (O16), it is estimated at 0.737 to 0.805, with standard errors of 0.039 to 0.048. The cognitive discount (O17) is estimated at 0.828 to 0.931 against a true value of 0.85, with standard errors of 0.022 to 0.033, and the gain of the learned anchor (O18) at 0.0483 to 0.0534 against a true value of 0.05, with standard errors below 0.0028. The ELB episodes of the simulated samples last 7 to 17 quarters in O15 and O16, 5 to 12 in O17, and 14 to 43 in O18.

Table A.2. The parameter bounds of the estimator on the real-time data. Search gives the bounds of the multi-start search of Section 4.3; wider gives the bounds under which the baseline estimate is re-optimized, those of the frozen estimate and of the rolling windows, and those at which the draws of every agent's training are clipped. The weights a_{lag} and a_{lead} also satisfy $a_{\text{lag}} + a_{\text{lead}} \leq 0.97$ under both, and c_π is held at its normalized value of 3.0. Source: author's calculations.

Parameter	Symbol	Search		Wider	
		Lower	Upper	Lower	Upper
IS: lagged gap	a_{lag}	0.00	1.00	0.00	1.00
IS: expected gap	a_{lead}	0.00	0.40	0.00	0.40
Phillips: expectation weight	λ_1	0.05	0.95	0.05	1.00
Phillips: slope	κ	0.00	0.50	0.00	0.50
Rule: smoothing	ρ	0.05	0.98	0.05	0.98
Rule: gap response	c_y	−0.50	2.00	−0.50	2.00
Rule: intercept (pp)	c_0	−3	3	−8	8
Rule shock: persistence	ϕ_u	0.00	0.97	0.00	0.97
Rule shock: SD (pp)	σ_η	0.02	2.00	0.02	2.00
Survey noise: SD (pp)	σ_{spf}	0.02	2.00	0.02	2.00

B Five Questions on the Estimates and the Agents

Five questions bear on the estimates and the agents: the first on the full-sample estimate's Phillips curve, the second on the frozen agent's responses, the third and fourth on the rolling agent, and the fifth on the rolling estimate's λ_1. Each is stated with its evidence and its conclusion.

B1. Is the Phillips curve of the full-sample estimate purely forward-looking? The curve is purely forward-looking at the structural limit $\lambda_1 = 1$. The evidence: λ_1 stands at 0.946, with a standard error of 0.041 computed near the bound of the parameter (Section 4.3); the limit lies within two of its standard errors of the estimate and not within one. The two IS weights lie on the cap $a_{\text{lag}} + a_{\text{lead}} \leq 0.97$ that the estimator imposes for a stationary solution, so that their moves are equal and opposite. The conclusion: The estimate places nearly all of the weight on expected inflation and lies below the limit.

B2. Are the frozen agent's responses to inflation and to the gap smaller than the baseline agent's? The evidence: The frozen agent's inflation and gap responses stand at 0.57 and 0.98 times the baseline agent's (Section 8.2). The inflation response of each of the eight training runs of the frozen agent's search lies below that of each of the eight training runs of the baseline agent's search, the

largest of the first at 0.60 times the smallest of the second; the gap responses of the two searches' runs overlap, those of the frozen agent's runs lying from 2.961 to 3.149 and those of the baseline agent's runs from 3.140 to 3.223. The conclusion: The frozen agent's inflation response is the smaller for every pair of training runs, one of each search; its gap response is the smaller at the selected checkpoints, and the two searches' runs overlap in it.

B3. How large is the rolling agent's mean difference over 2006:Q1 through 2015:Q4? The evidence: Over the 40 decision points of that span, of which 28 are at the ELB, the mean difference is +1.7 bp, and for each of the eight run numbers taken across the windows it lies between +1.6 and +1.7 bp. The conclusion: The mean difference is +1.7 bp on the selected checkpoints and between +1.6 and +1.7 bp for the run numbers; no standard error is reported for a span in which most quarters are at the ELB.

B4. Is the rolling agent's mean difference positive from 2016:Q1, and larger over 2021:Q1 through 2023:Q4 than over 2016:Q1 through 2020:Q1? The evidence: The mean differences are +0.569 pp over 2016:Q1 through 2020:Q1 and +2.105 pp over 2021:Q1 through 2023:Q4, two of the periods of Table 7; for each run number taken across the windows, they lie between +0.567 pp and +0.581 pp over the first period and between +2.084 pp and +2.105 pp over the second. The conclusion: The mean difference is positive over both periods and larger over the second, for every run number. Neither period's ratio of the selected agent's mean difference to its standard error exceeds its critical value, and no test of the change between the two periods' means is reported; the selected agent's ratios from 2016:Q1, over 2022:Q2 through 2026:Q2, and over 2021:Q1 through 2026:Q2 exceed their critical values (Section 8.3).

B5. Does the rolling estimate's λ_1 reach its structural limit of one? The evidence: λ_1 rises from 0.69 at 2021:Q2 to 0.99 at 2021:Q4, as the inflation of 2021 enters the data, and stands at 0.93 in the final window. The conclusion: The estimate reaches the limit in no window, and neither does the full-sample estimate.

C Mean Differences by Treatment, Paired Differences, and Losses under the CBO Measure

This appendix holds three tables by period. Table C.1 reports the baseline agent's mean difference from the realized federal funds rate, and its ratio to its standard error, under four treatments of the state, with the frozen agent's mean difference from 2016:Q1 under the same four. Table C.2 reports the paired differences, the differences between two agents' prescriptions at the same decision points, with their standard errors (Sections 8.2 and 8.3), and Table C.3 the losses of Section 8.4 under the CBO measure of the gap. The four treatments are these. Under the first, the treatment of Section 2, the IS and Phillips shocks of the decision quarter are zero. Under the second those two shocks are at the filter's estimates. Under the third the decision quarter's inflation is withheld from the filter, and under the fourth the inflation of the preceding quarters is, in addition, read from the vintage of the price index available at the decision point (Section 8.1). Under the third and the fourth the decision quarter's shocks are at zero, as under the first.

Table C.1. The baseline agent's mean difference from the realized federal funds rate under four treatments of the state, by period, and the frozen agent's over 2016:Q1–2026:Q2. Upper panel: the mean of the baseline agent's prescribed rate minus the realized federal funds rate (Section 2) over the decision points of the period, in percentage points. Middle panel: the ratio of that mean to its standard error under the equal-weighted cosine estimator (Section 8.1), and the critical value of a two-sided test at the 5 percent level under Student's t distribution with the number of terms the estimator's rule gives for the period's length. Lower panel: the mean of the frozen agent's prescribed rate minus the realized federal funds rate over the decision points from 2016:Q1 through 2026:Q2, under the same four treatments, in percentage points. Quarters is the number of decision points in the period. Shocks at zero: the decision quarter's IS and Phillips shocks at zero; shocks at estimates: those shocks at the filter's estimates; inflation withheld: the decision quarter's inflation withheld from the filter; vintage inflation: the inflation of the ten or eleven preceding quarters read, in addition, from the vintage of the price index available at the decision point; at 15 of the 118 decision points, 14 of them among the 48 of the row 1997:Q1–2008:Q4 and the other, 2020:Q1, in the rows 2016:Q1–2020:Q1, 2016:Q1–2020:Q4, and 2016:Q1–2026:Q2, the vintage does not cover the four preceding quarters and the column carries the state with the decision quarter's inflation withheld. With the inflation withheld and with the vintages' inflation, the decision quarter's shocks are at zero. Sources: as Table 1; author's calculations.

Period	Quarters	Shocks at zero	Shocks at estimates	Inflation withheld	Vintage inflation	Critical value
Mean difference, pp						
All quarters	118	+0.723	+0.649	+0.720	+0.666	—
1997:Q1–2008:Q4	48	+0.659	+0.518	+0.664	+0.581	—
2016:Q1–2020:Q1	17	+0.659	+0.462	+0.694	+0.685	—
2016:Q1–2020:Q4	20	+0.627	+0.393	+0.665	+0.656	—
2021:Q1–2023:Q4	12	+1.898	+2.010	+1.855	+1.718	—
2022:Q2–2026:Q2	17	+2.041	+2.232	+1.985	+1.876	—
2021:Q1–2026:Q2	22	+1.853	+1.992	+1.788	+1.690	—
2016:Q1–2026:Q2	42	+1.269	+1.230	+1.253	+1.198	—
Ratio of the mean difference to its standard error						
All quarters	118	3.39	2.67	3.45	3.24	2.26
1997:Q1–2008:Q4	48	3.81	2.20	3.62	2.71	2.57
2016:Q1–2020:Q1	17	4.35	2.78	4.19	4.45	4.30
2016:Q1–2020:Q4	20	4.20	2.07	4.12	4.53	4.30
2021:Q1–2023:Q4	12	3.49	3.02	3.33	3.34	4.30
2022:Q2–2026:Q2	17	9.68	8.65	8.89	12.09	4.30
2021:Q1–2026:Q2	22	5.72	5.11	5.39	5.50	3.18
2016:Q1–2026:Q2	42	3.64	2.79	3.77	3.91	2.78
Frozen agent, mean difference, pp						
2016:Q1–2026:Q2	42	+0.212	+0.284	+0.200	+0.206	—

Table C.2. Paired differences: the differences between two agents' prescribed rates at the same decision points. Mean is the mean of the first agent's prescribed rate minus the second's over the decision points of the period, in percentage points, the realized federal funds rate dropping out of the difference; SE is its standard error under the equal-weighted cosine estimator (Section 8.1), in percentage points; Ratio is the absolute value of the mean divided by its standard error; the critical value is that of a two-sided test at the 5 percent level under Student's t distribution with the number of terms the estimator's rule gives for the period's length. Quarters is the number of decision points in the period; 2020:Q4 is the final quarter for which the staff's estimates of the gap are among the observations. Sources: as Table 1; author's calculations.

Period	Quarters	Mean	SE	Ratio	Critical value
Frozen agent less baseline agent					
2016:Q1–2026:Q2	42	−1.057	0.449	2.35	2.78
2016:Q1–2020:Q4	20	−0.195	0.152	1.28	4.30
2021:Q1–2026:Q2	22	−1.840	0.262	7.03	3.18
2021:Q1–2023:Q4	12	−1.620	0.380	4.26	4.30
Rolling agent less baseline agent					
2016:Q1–2026:Q2	42	+0.097	0.144	0.67	2.78
2016:Q1–2020:Q4	20	−0.067	0.197	0.34	4.30
2021:Q1–2026:Q2	22	+0.246	0.146	1.69	3.18
2021:Q1–2023:Q4	12	+0.207	0.240	0.86	4.30
Rolling agent less frozen agent					
2016:Q1–2026:Q2	42	+1.153	0.545	2.12	2.78

Table C.3. Losses under the CBO measure of the output gap. Each entry is the sum over the quarters of the period of the loss (1) on the path named, with the path's gap read under the CBO measure as Section 8.4 states: the CBO gap of the quarter plus the path's deviation from the realized path under its own model's filter. The agent's path and the path under the estimated policy equation are simulated in the model at the estimate named, with the residuals recovered at that estimate; the realized path's loss is the same under both estimates. Sources: as Table 1; Congressional Budget Office; author's calculations.

Path	1997:Q1–2015:Q4	2016:Q1–2026:Q2	1997:Q1–2026:Q2
Baseline estimate			
Baseline agent	115.9	168.2	284.0
Estimated policy equation	159.6	202.5	362.1
Frozen estimate			
Frozen agent	129.1	215.0	344.1
Estimated policy equation	128.2	185.0	313.2
Realized path	370.9	223.5	594.4

Bibliography

Adrian, Tobias, Richard K. Crump, and Emanuel Moench. "Pricing the Term Structure with Linear Regressions." *Journal of Financial Economics* 110, no. 1 (2013): 110–138.

Adrian, Tobias, Douglas Laxton, and Maurice Obstfeld, eds. *Advancing the Frontiers of Monetary Policy*. Washington, DC: International Monetary Fund, 2018.

Al-Mashat, Rania, Kevin Clinton, Douglas Laxton, and Hou Wang. "United States: Federal Reserve's Dual Mandate." In *Advancing the Frontiers of Monetary Policy*, edited by Tobias Adrian, Douglas Laxton, and Maurice Obstfeld, 209–237. Washington, DC: International Monetary Fund, 2018.

Alichi, Ali, Kevin Clinton, Charles Freedman, Ondra Kamenik, Michel Juillard, Douglas Laxton, Jarkko Turunen, and Hou Wang. "Avoiding Dark Corners: A Robust Monetary Policy Framework for the United States." IMF Working Paper 15/134. Washington, DC: International Monetary Fund, 2015.

Blanchard, Olivier. "Where Danger Lurks." *Finance & Development* 51, no. 3 (2014): 28–31.

Board of Governors of the Federal Reserve System. "Federal Open Market Committee Announces Approval of Updates to Its Statement on Longer-Run Goals and Monetary Policy Strategy." Press release. Washington, DC: Board of Governors of the Federal Reserve System, August 22, 2025.

Board of Governors of the Federal Reserve System. *Monetary Policy Report*. Washington, DC: Board of Governors of the Federal Reserve System, February 2019.

Brainard, William C. "Uncertainty and the Effectiveness of Policy." *American Economic Review* 57, no. 2 (1967): 411–425.

Brayton, Flint, Thomas Laubach, and David Reifschneider. "Optimal-Control Monetary Policy in the FRB/US Model." FEDS Notes. Washington, DC: Board of Governors of the Federal Reserve System, November 21, 2014.

Chen, Mingli, Rama Cont, Andreas Joseph, Michael Kumhof, Xinlei Pan, Wei Xiong, and Xuan Zhou. "Deep Reinforcement Learning in a Monetary Model." Staff Working Paper 1,142. London: Bank of England, 2025.

Edge, Rochelle M., Thomas Laubach, and John C. Williams. "Welfare-Maximizing Monetary Policy under Parameter Uncertainty." *Journal of Applied Econometrics* 25, no. 1 (2010): 129–143.

Evans, George W., and Seppo Honkapohja. *Learning and Expectations in Macroeconomics*. Princeton, NJ: Princeton University Press, 2001.

Gabaix, Xavier. "A Behavioral New Keynesian Model." *American Economic Review* 110, no. 8 (2020): 2271–2327.

Ha, David, and Jürgen Schmidhuber. "Recurrent World Models Facilitate Policy Evolution." In *Advances in Neural Information Processing Systems* 31 (2018): 2450–2462.

Hansen, Lars Peter, and Thomas J. Sargent. *Robustness*. Princeton, NJ: Princeton University Press, 2008.

Hinterlang, Natascha, and Alina Tänzer. "Optimal Monetary Policy Using Reinforcement Learning." Discussion Paper 51/2021. Frankfurt am Main: Deutsche Bundesbank, 2021.

Holston, Kathryn, Thomas Laubach, and John C. Williams. "Measuring the Natural Rate of Interest: International Trends and Determinants." *Journal of International Economics* 108, supplement 1 (2017): S59–S75.

Holston, Kathryn, Thomas Laubach, and John C. Williams. "Measuring the Natural Rate of Interest after COVID-19." Federal Reserve Bank of New York Staff Reports, no. 1063, June 2023.

Krishnamurthy, Arvind, and Miguel Chumbo. "Disappearing Dollar Convenience?" NBER Working Paper 35742. Cambridge, MA: National Bureau of Economic Research, 2026.

Lazarus, Eben, Daniel J. Lewis, James H. Stock, and Mark W. Watson. "HAR Inference: Recommendations for Practice." *Journal of Business & Economic Statistics* 36, no. 4 (2018): 541–559.

Lucas, Robert E., Jr. "Econometric Policy Evaluation: A Critique." *Carnegie-Rochester Conference Series on Public Policy* 1 (1976): 19–46.

Orphanides, Athanasios. "Monetary Policy Rules Based on Real-Time Data." *American Economic Review* 91, no. 4 (2001): 964–985.

Primiceri, Giorgio E. "Why Inflation Rose and Fell: Policy-Makers' Beliefs and U.S. Postwar Stabilization Policy." *Quarterly Journal of Economics* 121, no. 3 (2006): 867–901.

Sargent, Thomas J. *The Conquest of American Inflation*. Princeton, NJ: Princeton University Press, 1999.

Sutton, Richard S. "Dyna, an Integrated Architecture for Learning, Planning, and Reacting." *ACM SIGART Bulletin* 2, no. 4 (1991): 160–163.

Taylor, John B. "Discretion versus Policy Rules in Practice." *Carnegie-Rochester Conference Series on Public Policy* 39 (1993): 195–214.

Taylor, John B., and John C. Williams. "Simple and Robust Rules for Monetary Policy." In *Handbook of Monetary Economics*, vol. 3B, edited by Benjamin M. Friedman and Michael Woodford, 829–859. Amsterdam: Elsevier, 2010.

Wieland, Volker. "Monetary Policy, Parameter Uncertainty and Optimal Learning." *Journal of Monetary Economics* 46, no. 1 (2000): 199–228.