AI, Expectations, and Optimal Interest Rates

Alessandro T. Villa[1]

October 2026

Abstract

An AI productivity boom can warrant lower nominal rates even when most gains are anticipated. I derive analytical representations of optimal rates under commitment and discretion in an economy with sticky prices and wages and hand-to-mouth households. Holding the realized productivity path fixed, I vary expectations. Anticipated gains raise the natural rate, but below-target inflation supports real wages and keeps nominal rates lower. Surprises leave the natural rate unchanged, calling for cuts. After surprises, expected redistribution toward hand-to-mouth households calls for deeper rate cuts to sustain savers' spending. Raising rates requires more anticipation when wages are stickier relative to prices.

JEL codes: E52, E31, D84, E25, O33.

Keywords: Optimal monetary policy; productivity expectations; commitment and discretion; sticky prices and wages; hand-to-mouth households.

Will artificial intelligence call for higher or lower interest rates? Building on Goolsbee's Hoover Institution remarks (Goolsbee 2026), Barlevy, Fisher, Pennington, and Villa (2026) show that under a Taylor rule tracking the natural rate, the answer depends on what agents expect.[2] If households see faster productivity growth coming, they spend against future income, the natural rate rises, and the rule calls for higher rates. If each gain comes as a surprise, the natural rate does not move, inflation falls, and the rule calls for lower rates. In this paper, I show that optimal monetary policy, with or without commitment, can call for lower nominal rates even when most productivity gains are anticipated. With hand-to-mouth households, surprises can warrant even deeper rate cuts than in the representative-household economy, despite nearly unchanged optimal inflation and output. I derive analytical Taylor-type representations of the optimal nominal interest rate in an economy with sticky prices and wages, savers, and hand-to-mouth households. I then compare an announced boom with a sequence of surprises along the same productivity path in an illustrative calibrated economy. Three results emerge.

First, in the calibrated economy, whether optimal policy calls for higher or lower nominal rates depends on how much of the boom is anticipated, both with and without commitment. In the representative-household economy, the optimal real rate closely tracks the natural rate: it rises under full anticipation and barely moves when gains arrive as surprises. Because nominal wages adjust

[1]Federal Reserve Bank of Chicago, 230 South LaSalle Street, Chicago, IL 60604-1413, USA. Email: alessandro.villa@chi.frb.org. I thank Gadi Barlevy for helpful comments. The views expressed are those of the author and do not necessarily reflect those of the Federal Reserve Bank of Chicago or the Federal Reserve System.

[2]Workplace evidence shows that generative AI can raise productivity (Brynjolfsson, Li, and Raymond 2025), while its aggregate effects depend on the scope for gains across tasks (Acemoglu 2025).

slowly, optimal policy allows below-target inflation to help real wages catch up with productivity. Under full anticipation, below-target inflation offsets much of the increase in the natural rate, yielding smaller nominal-rate increases than under the benchmark Taylor rule. After surprises, below-target inflation instead calls for nominal rate cuts. The same qualitative optimal nominal-rate pattern holds with hand-to-mouth households.

Second, hand-to-mouth households affect the size of these rate responses primarily through expected redistribution between wages and profits, especially when productivity gains arrive as surprises. Unexpected productivity gains temporarily raise firm profits relative to labor income. Expecting no further productivity gains, savers anticipate this advantage fading as wages catch up and, at unchanged rates, would want to save more. Lower real rates support their current spending, sustaining nearly the same aggregate allocation. This mechanism operates under both commitment and discretion. At the midpoint of the surprise-driven boom, a 30 percent hand-to-mouth share lowers the optimal nominal rate by about 34 basis points relative to the representative-household economy, while inflation and the output gap remain nearly unchanged. When gains are anticipated, redistribution is limited, and household heterogeneity has little effect on optimal policy.

Third, I examine how much of the productivity boom must be anticipated for optimal monetary policy to call for higher nominal interest rates. Holding the realized productivity path fixed, I vary the share of future productivity gains that agents foresee, from a fully announced boom to a sequence of entirely unanticipated gains. At the midpoint of the boom, nominal rates rise only when the anticipated share exceeds 74 percent under commitment and 76 percent under discretion in the representative-household economy. Thus, optimal monetary policy may call for lower nominal interest rates even when most productivity gains are anticipated. With hand-to-mouth households, these thresholds increase to 82 and 83 percent, respectively. Through the mechanism described in the second result, expected redistribution lowers optimal rates when at least some gains arrive as surprises, while barely changing the response under full anticipation. A larger anticipated share than in the representative-household economy is therefore needed for the rise in the natural rate to offset this additional downward pressure on rates. The threshold rises as wage rigidity increases relative to price rigidity, because optimal policy then allows inflation to fall further below target. When wages are about as flexible as prices, rates rise once a much smaller share of the boom is anticipated, but hand-to-mouth households raise the threshold in every calibration I consider. I establish this result analytically in the limit of a sustained boom under commitment: a larger hand-to-mouth share lowers nominal rates whenever some gains are unanticipated, without changing inflation.

The analytical rules also help explain how real wages adjust across policy regimes. Under discretion, the central bank takes future policymakers' responses to inherited real-wage gaps as given. When real wages are expected to keep lagging productivity, firms anticipate low future marginal costs and set lower prices today. In these experiments, after the initial adjustment, discretion relies more on price declines and less on nominal wage growth than commitment.

Related literature. Gilchrist and Leahy (2002) study productivity news in a financial-accelerator model with sticky prices but flexible wages. Christiano, Ilut, Motto, and Rostagno (2010) show how

productivity news interacts with sticky wages and inflation-focused monetary policy to generate inefficient booms. Offick and Wohltmann (2016) examine how anticipation changes stabilization losses from cost-push shocks under optimal policy. I study how expectations shape the optimal interest-rate response to a sustained productivity boom with hand-to-mouth households and sticky prices and wages. Holding the realized productivity path fixed, I isolate the roles of the natural rate, below-target inflation, and expected redistribution, and show that nominal rates can fall even when most gains are anticipated.

Limited asset-market participation and redistribution shape monetary transmission (Bilbiie 2008; Auclert 2019). Debortoli and Galí (2025) establish how expected changes in consumption heterogeneity enter aggregate demand. McKay and Wolf (2026) distinguish optimal aggregate stabilization from the interest-rate path needed to implement it. I show how this distinction depends on productivity expectations: surprise gains temporarily shift income toward profits, and the expected reversal calls for deeper rate cuts with hand-to-mouth households, even when optimal inflation and output barely change relative to the representative-household economy.

Analytical interest-rate rules with sticky prices and wages build on Giannoni and Woodford (2003), while Dávila and Schaab (2023) study commitment and discretion with heterogeneous households. Acharya, Challe, and Dogra (2023) show how consumption inequality enters the optimal stabilization trade-off; Hansen, Lin, and Mano (2023) study optimal policy and labor-share targeting after technology shocks in a two-agent model with sticky prices and wages. Miyazaki (2026a,b) studies profit redistribution and inherited wage states in two-agent models with type-specific wages. I derive explicit Taylor-type representations of optimal nominal interest rates under commitment and discretion with common sticky wages and hand-to-mouth households. These representations identify how the natural rate, nominal rigidities, inherited states, and expected redistribution jointly shape the response to anticipated and unanticipated productivity gains.

I. The Economy

Time is discrete. A fraction $\eta \in [0, 1)$ of households is hand-to-mouth, denoted H; the remaining households are savers, denoted S. Both types have discount factor β and period utility $u(C) - v(N)$, with $u'(C) = C^{-\sigma}$ and $v'(N) = \chi N^{\varphi}$, where $\sigma, \varphi, \chi > 0$. Savers own firms and trade nominal risk-free bonds in zero net supply. There is no capital. Hand-to-mouth households consume labor income each period.

Differentiated firms produce with labor and common productivity. Goods and labor varieties have substitution elasticities $\epsilon, \epsilon_w > 1$. Prices and wages remain unchanged with Calvo probabilities θ_p, θ_w. Each labor union represents the same population mix of households and assigns equal hours to the two types. Labor income is pooled across occupations within each type. Production and wage subsidies eliminate steady-state markups. Fiscal incidence is chosen so that each type has the same consumption in the efficient zero-inflation steady state; savers receive residual profits and finance the production subsidy. Mathematical details are provided in Appendix A.

Lowercase quantities are log deviations from this steady state. Let a_t be productivity, $g_t = a_t - a_{t-1}$ its growth, y_t output, and w_t the real wage. Real marginal cost is $m_t = w_t - a_t$. To first order, output equals consumption and satisfies $y_t = a_t + n_t$. The household budgets imply

$$c_t^H = y_t + m_t, \qquad c_t^S = y_t - \nu m_t, \qquad \nu \equiv \frac{\eta}{1-\eta}. \tag{1}$$

Thus m_t also measures the deviation of labor income from its steady-state share. If real wages trail productivity, $m_t < 0$: hand-to-mouth consumption is below aggregate consumption and savers receive the offsetting profits.

Productivity is the only aggregate exogenous shock. Households, firms, unions, and the central bank share the same information: they observe current productivity and any news about future growth before making current decisions.

Write $\psi = \sigma + \varphi$, and let $x_t = y_t - y_t^n$ be output relative to the efficient flexible-price, flexible-wage allocation. That allocation and its natural real rate are

$$y_t^n = \frac{1+\varphi}{\psi} a_t, \qquad r_t^n \equiv r_t^* = \frac{\sigma(1+\varphi)}{\psi} \mathrm{E}_t g_{t+1}.$$

Here and below, interest rates are deviations from their steady-state values. Under flexible prices and wages, both consumption types equal y_t^n. Consequently r_t^n does not depend on η under this fiscal and labor-market structure.

Price inflation π_t and nominal wage inflation π_t^w satisfy

$$\pi_t = \beta \mathrm{E}_t \pi_{t+1} + \kappa_p m_t, \tag{2}$$

$$\pi_t^w = \beta \mathrm{E}_t \pi_{t+1}^w - \kappa_w (m_t - \psi x_t), \tag{3}$$

$$m_t - m_{t-1} = \pi_t^w - \pi_t - g_t. \tag{4}$$

The slopes are $\kappa_p = (1-\theta_p)(1-\beta\theta_p)/\theta_p$ and $\kappa_w = (1-\theta_w)(1-\beta\theta_w)/[\theta_w(1+\epsilon_w\varphi)]$. The log deviation of the common union's population-weighted marginal utility is $-\sigma y_t$. Heterogeneity therefore leaves these slopes unchanged.

Let R_t be the nominal rate and $r_t = R_t - \mathrm{E}_t \pi_{t+1}$ the ex ante real rate. The saver Euler equation gives

$$R_t = r_t^n + \mathrm{E}_t \pi_{t+1} + \sigma \mathrm{E}_t \Delta x_{t+1} - \mathcal{H}_t, \qquad \mathcal{H}_t \equiv \sigma \nu \mathrm{E}_t \Delta m_{t+1}. \tag{5}$$

Throughout, $\Delta z_t = z_t - z_{t-1}$. Defining the consumption difference $d_t = c_t^H - c_t^S = m_t/(1-\eta)$ gives $\mathcal{H}_t = \sigma\eta \mathrm{E}_t \Delta d_{t+1}$. Expected recovery of hand-to-mouth consumption relative to saver consumption lowers the rate that supports a given output-gap path.

The planner gives equal welfare weight to each household. A second-order approximation around the efficient symmetric steady state yields the loss

$$\frac{1}{2}\mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t [\lambda_x x_t^2 + \lambda_p \pi_t^2 + \lambda_w (\pi_t^w)^2 + \lambda_m m_t^2], \tag{6}$$

where

$$\lambda_x = \psi, \quad \lambda_p = \frac{\epsilon}{\kappa_p}, \quad \lambda_w = \frac{\epsilon_w}{\kappa_w}, \quad \lambda_m = \sigma\nu. \tag{7}$$

The last term equals $\sigma\eta(1-\eta)d_t^2$: it is the cost of unequal consumption, as derived from household utility.[3] Appendix A derives all four weights.

II. Optimal Policy

The planner chooses an allocation satisfying (2)–(4) and uses (5) to implement it. In what follows, define the real-rate gap $s_t = r_t - r_t^n$ and the expected output-gap change $u_t = \sigma \mathrm{E}_t \Delta x_{t+1} = s_t + \mathcal{H}_t$.

A. *Commitment*

Under commitment, the central bank chooses a state-contingent plan at date zero and honors it after future shocks. There are no inherited promises initially.

Proposition 1 (Optimal policy with commitment). *Starting from the symmetric steady state with no inherited promises, the Ramsey allocation satisfies at $t = 0$*

$$R_0 = r_0^n + \mathrm{E}_0\pi_1 + \bar{a}_x x_0 + a_\pi \pi_0 + a_w \pi_0^w + a_g g_0 - \mathcal{H}_0.$$

For $t > 0$, the interest-rate representation is

$$\begin{aligned} R_t = r_t^n + \mathrm{E}_t\pi_{t+1} + s_{t-1} + \bar{a}_x \Delta x_t - \frac{\sigma}{\beta}\Delta x_{t-1} \\ + a_\pi \pi_t + a_w \pi_t^w + a_g g_t + \mathcal{H}_{t-1} - \mathcal{H}_t. \end{aligned} \tag{8}$$

The coefficients are given in (B1). *Initial values are $m_{-1} = 0$, $x_{-2} = x_{-1} = 0$, and $s_{-1} = \mathcal{H}_{-1} = 0$. For $t \geq 1$, $\mathcal{H}_{t-1}$ is the forecast formed at $t-1$.*

Appendix B proves the proposition by eliminating the Ramsey multipliers. The rule depends on changes in the output gap and inherited policy history. Promises made earlier affect current choices because they previously influenced wage and price setting. Heterogeneity changes the optimal allocation through λ_m and its implementing rate through $\mathcal{H}_t$.

[3] Equal steady-state consumption eliminates a linear redistribution term and an output–inequality cross term.

B. Discretion

Without commitment, the central bank reoptimizes each period, taking future policy functions as given. It inherits m_{t-1}, but no promises. For an announced productivity path, future policy also depends on the known continuation of that path; for stochastic growth, it depends on the exogenous information state.

Proposition 2 (Optimal policy without commitment). *Fix an admissible linear Markov equilibrium, with $M = \partial m_t / \partial m_{t-1}$, in which the exogenous productivity process is independent of the inherited endogenous state. Discretionary policy has the representation*

$$\begin{aligned} R_t = r_t^n + \mathrm{E}_t \pi_{t+1} + b_x(M) x_t + b_\pi(M) \pi_t + b_w(M) \pi_t^w \\ + b_g(g_t - m_{t-1}) - \mathcal{H}_t. \end{aligned} \tag{9}$$

The coefficients are given in (B12)*, with b_g independent of M; M solves* (B9)*. The rule applies from $t = 0$ with $m_{-1} = 0$.*

Appendix B provides the proof and admissibility conditions. The lagged real-wage gap matters because today's policymaker cannot undo inherited nominal wages. Unlike commitment, discretion does not attach value to previously promised responses. Both policies nevertheless contain the same expected-redistribution correction $-\mathcal{H}_t$.

C. Wage Rigidity and Redistribution

When productivity grows, real wages can catch up through nominal wage growth, price deflation, or both, as (4) shows. With staggered adjustment, wage changes create dispersion across workers and price changes create dispersion across firms. Optimal policy balances these costs with output-gap stabilization. Allowing some deflation reduces the adjustment required of nominal wages and can lower the overall welfare cost.

Hand-to-mouth households add consumption inequality to this trade-off. When real wages lag behind productivity, income shifts from wages toward profits. Lower prices support the purchasing power of wages, but the effect on household consumption also depends on employment. Heterogeneity therefore need not imply more deflation. It also affects the rate needed to implement the chosen allocation: an expected recovery of labor income relative to profits implies $\mathcal{H}_t > 0$, lowering that rate through (5). This correction changes sign if labor income is expected to fall further behind.

Two limiting cases clarify these mechanisms. Without hand-to-mouth households, both optimal policy rules reduce to those of the representative-household model ($\eta = 0$, so $\lambda_m = \mathcal{H}_t = 0$). With fully flexible wages, wage setting gives $m_t = \psi x_t$. Both policy regimes attain $x_t = m_t = \pi_t = 0$, with nominal wage growth $\pi_t^w = g_t$ absorbing productivity growth. Since wage adjustment then entails no dispersion cost, this allocation achieves zero loss. Equation (5) gives $R_t = r_t^n = r_t^*$ for every $\eta < 1$.

III. An Illustrative Productivity Boom

A quarter is one period. I use conventional preference and markup parameters: $\beta = 0.99$, logarithmic consumption utility, and a unit Frisch elasticity ($\sigma = \varphi = 1$). The substitution elasticities $\epsilon = \epsilon_w = 6$ imply price and wage markups of 20 percent before subsidies. Firms reoptimize prices roughly every five quarters on average ($\theta_p \simeq 0.8035$), yielding a price Phillips-curve slope of $\kappa_p = (1-\theta_p)(1-\beta\theta_p)/\theta_p = 0.05$, the value used by Barlevy, Fisher, Pennington, and Villa (2026).[4] Wages are reoptimized every four quarters on average ($\theta_w = 0.75$), yielding a wage Phillips-curve slope of $\kappa_w = (1-\theta_w)(1-\beta\theta_w)/[\theta_w(1+\epsilon_w\varphi)] \simeq 0.01226$.[5] The wage Phillips curve is flatter despite more frequent wage resets: the factor $1+\epsilon_w\varphi = 7$ reflects the response of hours to relative wages (Erceg, Henderson, and Levin 2000). I compare economies with $\eta = 0$ and $\eta = 0.30$, the latter broadly consistent with U.S. evidence on hand-to-mouth households (Kaplan, Violante, and Weidner 2014). The comparison policy is the Taylor rule $R_t = r_t^n + 1.5\pi_t + 0.5x_t$.

Following Barlevy, Fisher, Pennington, and Villa (2026), productivity growth rises by 0.0025 per quarter for 40 quarters and then decays at rate 0.70. This is an annualized increase of one percentage point for ten years, followed by a gradual return to baseline growth. Productivity remains permanently higher.

A. *An Announced Path*

In Figure 1, the entire productivity path unexpectedly becomes common knowledge at $t = 0$ (quarter 1); quarter 0 is the pre-shock observation. This MIT shock is followed by perfect foresight and no further innovations, so $\mathrm{E}_t g_{t+j} = g_{t+j}$ for $j \geq 1$. Expected growth raises the natural rate by one percentage point in both economies while next-quarter annualized growth remains one percentage point above baseline.

Optimal policy allows more deflation than the Taylor benchmark and keeps output close to its efficient level. At quarter 20, the representative-household Ramsey policy has an annualized price inflation of about −0.80 percentage points and an output gap of 0.003 percent. The Taylor rule delivers −0.54 percentage points and a positive gap of 0.14 percent. Allowing the larger price decline supports real-wage adjustment without the same expansion above efficient output.

At quarter 20, the Ramsey nominal rate rises by about 20 basis points, compared with 46 under the Taylor rule and 19 under discretion. Real rates remain within one basis point of the 100-basis-point natural-rate increase across policies and household shares. Nominal-rate differences therefore mainly reflect expected deflation.

Adding hand-to-mouth households changes these paths little. At quarter 20, real wages have largely caught up with productivity, and the labor share is expected to change little. The expected-

[4] Micro evidence points to more frequent regular price changes (Nakamura and Steinsson 2008). More flexible prices raise the anticipation thresholds in Section III.C (Appendix C), so the benchmark is conservative in this respect.

[5] The implied quarterly reset probability of 25 percent is consistent with the 21.1–26.6 percent wage-change frequency in Barattieri, Basu, and Gottschalk (2014). The annual reset probability, $1-\theta_w^4 \simeq 0.68$, is close to the 65 percent of workers with a year-over-year base-wage change in Grigsby, Hurst, and Yildirmaz (2021).

redistribution correction to the interest rate is therefore small. Discretion produces similar inflation and rates, with a somewhat larger output gap.

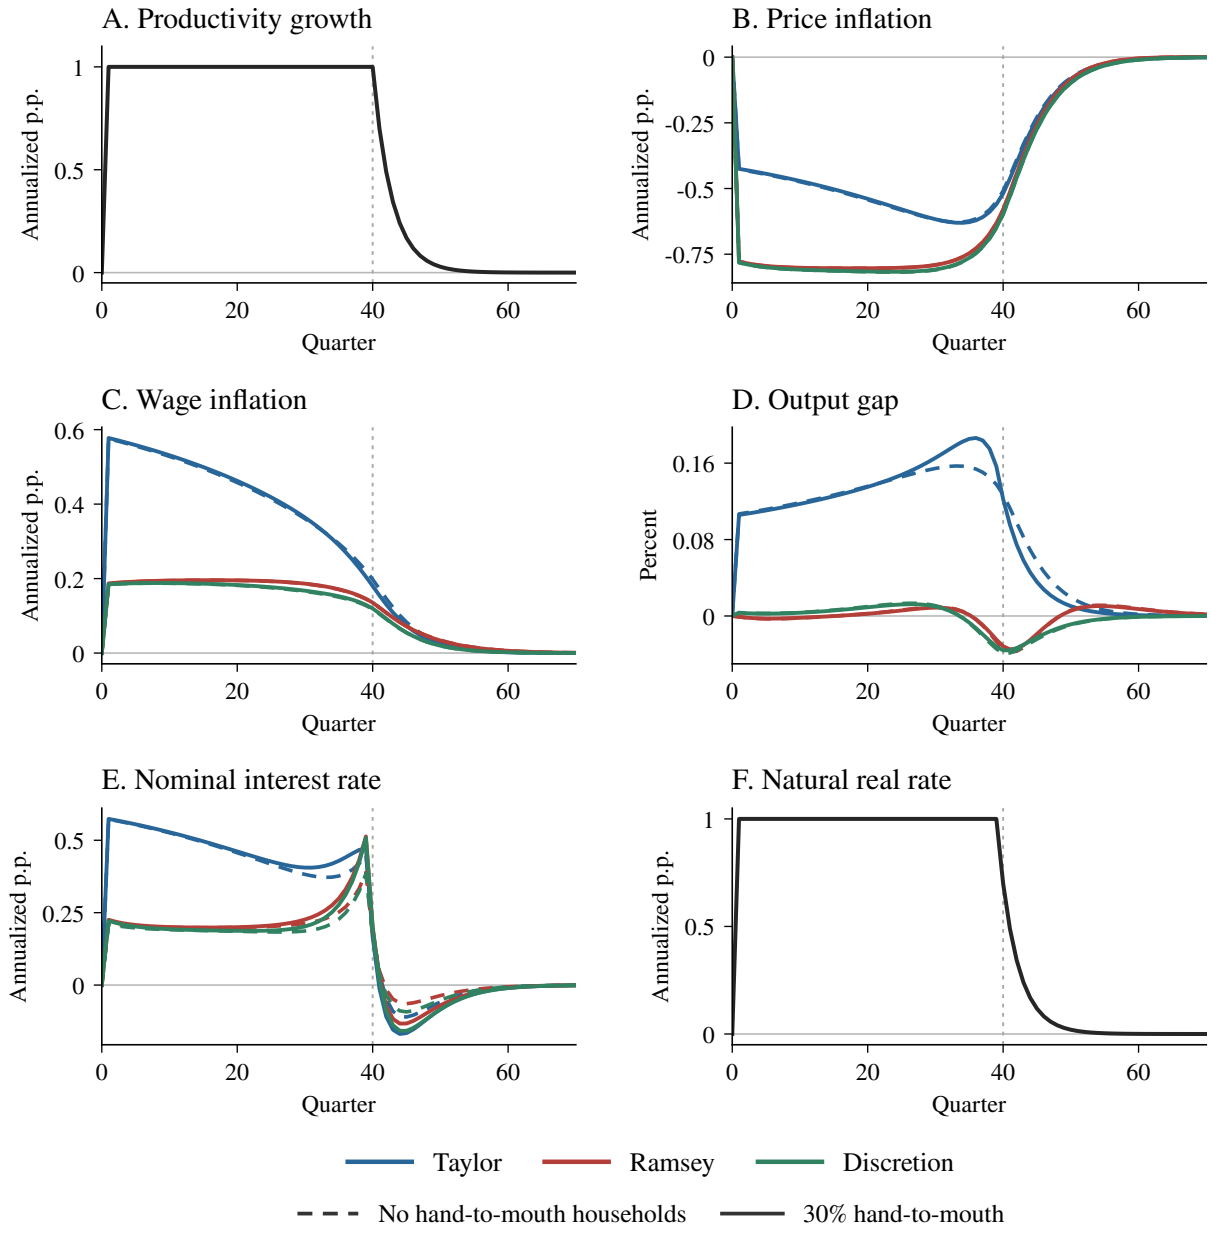

FIGURE 1. AN ANNOUNCED PRODUCTIVITY BOOM

Notes: The full productivity path becomes known in quarter 1. Growth is one annualized percentage point above baseline for 40 quarters and then decays at rate 0.70. Dashed lines have no hand-to-mouth households; solid lines have a 30 percent share. Ramsey denotes commitment. Black lines in panels A and F are common to all policies and both shares. All panels show deviations from steady state, except the output gap, which is relative to efficient output. Rates and inflation are annualized percentage points. Quarter 0 is the pre-shock observation. The vertical line marks quarter 40. Source: Author's model calculations.

B. Repeated Surprises

Figure 2 uses the same realized productivity path, evaluated under $g_t = \xi_t$, where ξ_t is i.i.d. with mean zero. All agents know this process and observe each innovation when it arrives, so $\mathrm{E}_t g_{t+j} = 0$ for every $j \geq 1$. The plotted sequence selects positive innovations that reproduce the announced path, but future realizations are never revealed in advance. Conditional forecasts therefore differ from subsequent realizations along this path. The natural-rate deviation is zero throughout. Commitment remains a contingent plan with inherited promises.

Both optimal policies again tolerate deflation. The presence of hand-to-mouth households barely changes its size: at quarter 20, Ramsey inflation is about −0.80 annualized percentage points in both economies. The output gap moves from −0.109 percent without hand-to-mouth households to −0.101 percent with them. Yet the nominal-rate cut increases from 56 to 90 basis points. Under discretion, it increases from 58 to 92 basis points. Without hand-to-mouth households, the ex ante real rate at that date is 7.01 basis points above the unchanged natural rate under Ramsey and 7.25 under discretion: near tracking does not imply exact invariance.

The Euler equation explains the large rate difference. In the heterogeneous Ramsey economy, marginal cost at quarter 20 is about −0.92 percent. Hand-to-mouth consumption is 1.31 percent below saver consumption. Since agents expect no further productivity growth, they expect nominal wages and labor income to catch up. The expected consumption difference next quarter is only −1.03 percent. Expressed in annualized basis points, (5) gives

$$\underbrace{r_t - r_t^n}_{-27.07} = \underbrace{\sigma \mathrm{E}_t \Delta x_{t+1}}_{6.50} \underbrace{-\sigma\eta \mathrm{E}_t \Delta d_{t+1}}_{-33.58}.$$

Expected redistribution accounts for nearly all of the 34-basis-point nominal-rate difference across household economies. Savers expect to lose some of their temporarily high profit income as labor income recovers. The lower rate supports their spending while aggregate output follows almost the same optimal path. Discretion has the same mechanism and a similar quantitative decomposition.

The Taylor rate also falls more with hand-to-mouth households, by about 27 basis points at quarter 20, but its prescribed response is different. Roughly 21 basis points come from its output-gap term and 5 from its inflation term. Under optimal policy, the large additional rate cut mainly implements a similar allocation in an economy with different saver consumption growth.

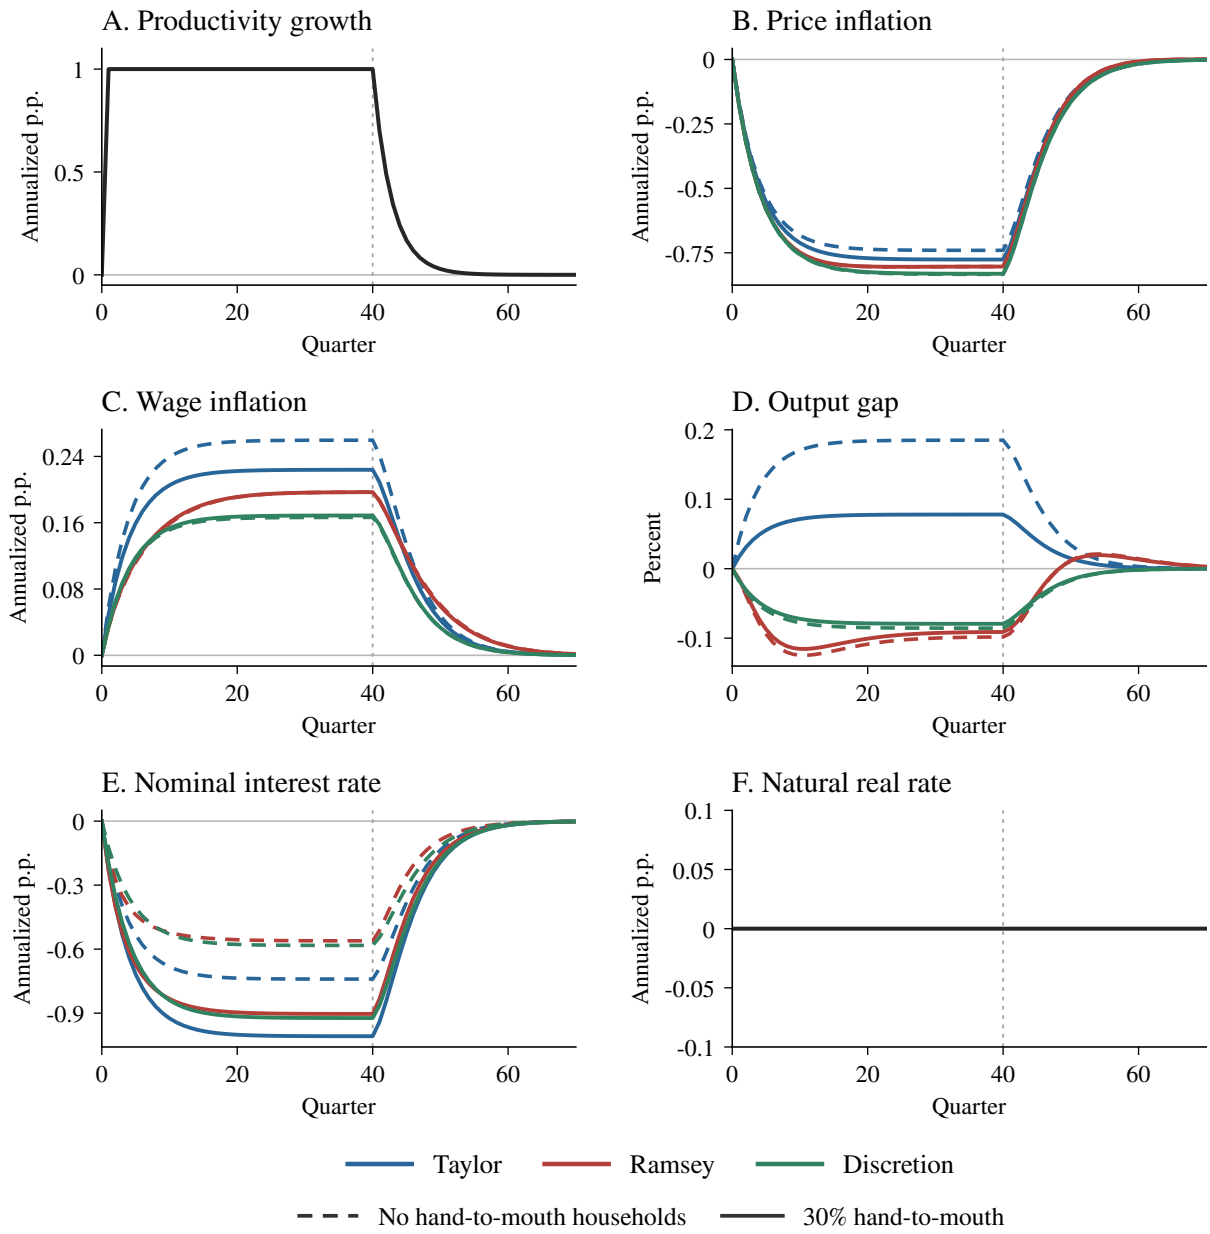

Figure 2. The Same Boom as Repeated Productivity Surprises

Notes: Realized productivity is identical to Figure 1. Each increase is an unanticipated realization under an i.i.d. growth process: $E_t g_{t+1} = 0$. The natural-rate deviation is therefore zero. Ramsey retains the promises made under its contingent plan. Colors, household shares, units, and the quarter-40 marker are as in Figure 1. Dashed and solid optimal-policy paths largely overlap in panels B and C. Source: Author's model calculations.

C. Partial Anticipation

Let $\bar{g}_t$ be the same realized growth path and $\lambda \in [0, 1]$ its anticipated share. In quarter 1, agents learn the deterministic component $\lambda \bar{g}_t$. Independent innovations subsequently realize $(1 - \lambda)\bar{g}_t$ each

quarter, with future innovations expected to be zero. Thus $E_t g_{t+j} = \lambda \bar{g}_{t+j}$ for $j \geq 1$, and the natural rate is λ times its announced-path value. The endpoints reproduce the two preceding experiments.

At first order, the nominal rate is a weighted average of its responses under full anticipation and repeated surprises: $R_t(\lambda) = \lambda R_t(1) + (1 - \lambda) R_t(0)$.

Figure 3 plots the quarter-20 nominal rate for the finite boom. With $\eta = 0$, the anticipated share must exceed 0.74 under commitment and 0.76 under discretion for the rate to rise. With $\eta = 0.30$, these thresholds rise to 0.82 and 0.83. Expected redistribution lowers the optimal rate in the surprise component and contributes little in the announced component. A larger anticipated share is therefore needed to offset the deeper surprise-induced cut. The Taylor thresholds are lower, 0.62 and 0.68: the rule resists deflation and raises rates for a smaller anticipated share.[6]

Across quarters 1–40, the Ramsey thresholds range from 0.43 to 0.77 without hand-to-mouth households and from 0.51 to 0.82 with them. Quarter 20 is close to the middle-of-boom optimal thresholds, but overstates those at the onset; Appendix C reports all ranges.

The threshold depends mainly on wage rigidity relative to price rigidity, as measured by the slopes of the two Phillips curves. The flatter the wage Phillips curve is relative to the price Phillips curve, the more optimal policy relies on falling prices to raise real wages. A larger deflation then offsets more of the higher natural rate, and a larger anticipated share is needed for the nominal rate to rise. Appendix C varies the wage reset probability and the slope of the price Phillips curve one at a time. The quarter-20 Ramsey threshold rises to 0.90 with $\theta_w = 0.85$ and falls to 0.52 with $\theta_w = 0.65$. With $\theta_w = 0.50$, when wages are reset about every two quarters and are about as flexible as prices, it falls to 0.17. Across these cases, the threshold moves with the optimal deflation under full anticipation, which ranges from 0.41 to 0.92 annualized percentage points. Hand-to-mouth households raise the threshold in every case. In the benchmark calibration, optimal rates can fall even when most gains are anticipated.

[6]At quarter 20, optimal nominal rates remain negative for every anticipated share with $\sigma = 0.5$; with $\sigma = 2$, all thresholds fall (Appendix C).

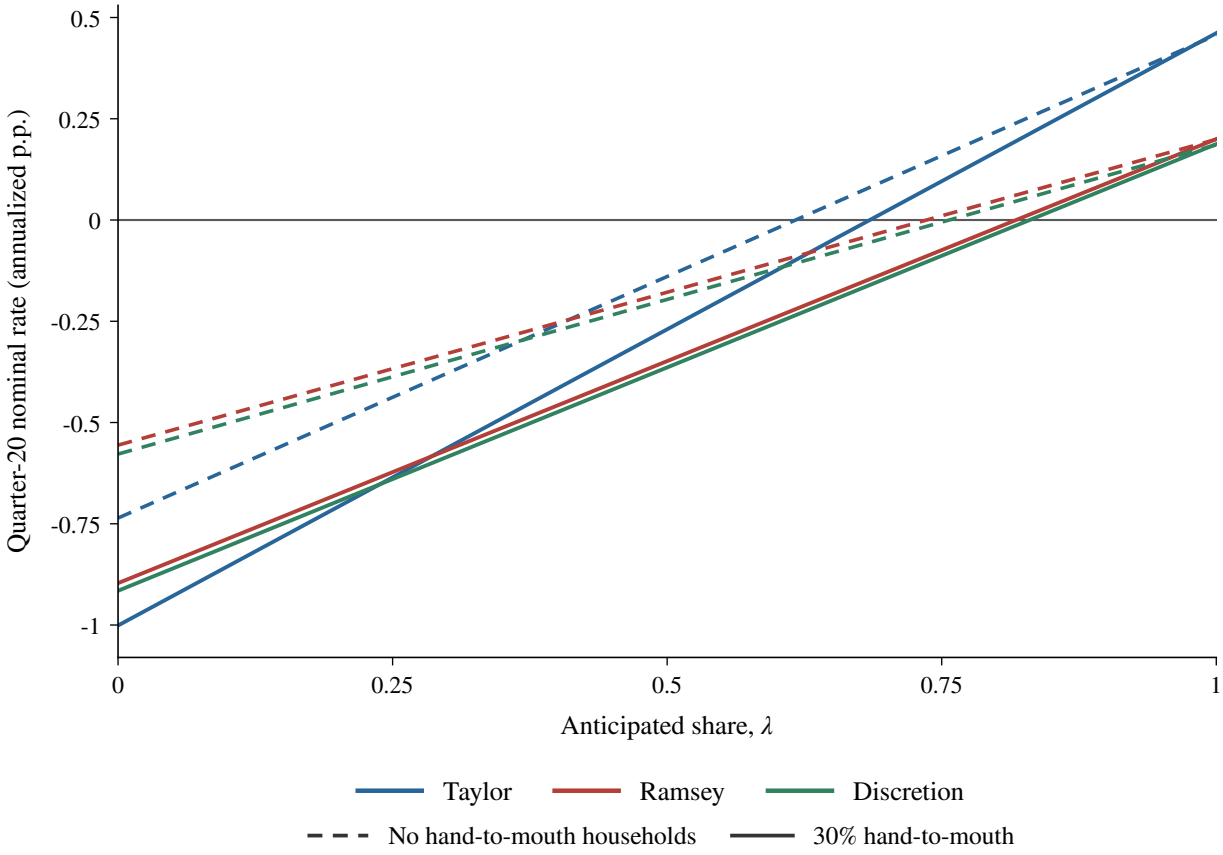

FIGURE 3. PARTIAL ANTICIPATION AND THE NOMINAL RATE

Notes: The nominal-rate deviation in quarter 20 against the share of productivity growth anticipated from quarter 1. Realized growth is the same at every share. At $\lambda = 0$ all increases are surprises; at $\lambda = 1$ the full path is announced. The calibration has $\sigma = 1$. Rates are annualized percentage points relative to steady state; the horizontal line marks zero. Colors and household line styles match Figures 1 and 2. Each curve uses directly solved conditional equilibria, with inherited multipliers under Ramsey. Source: Author's model calculations.

The following proposition gives a more general analytical counterpart to these numerical findings under commitment, using the limit of a sustained boom. It links the inflation adjustment to relative nominal rigidities and the anticipation threshold to household heterogeneity.

Proposition 3 (Anticipation in a sustained boom). *Under commitment, consider the realized path $g_t = g > 0$, with $\mathrm{E}_t g_{t+j} = \lambda g$ for $j \geq 1$. Assume gaps and Ramsey multipliers converge and innovation responses are absolutely summable. Let $\omega = \lambda_w/(\lambda_p + \lambda_w)$. Limiting price inflation is independent of λ and η:*

$$\pi^{\infty}(\lambda, \eta) = -\omega g.$$

Under full anticipation ($\lambda = 1$), the limiting nominal rate is also independent of η:

$$R^{\infty}(1, \eta) = \left[\frac{\sigma(1+\varphi)}{\psi} - \omega\right] g.$$

Holding preferences and substitution elasticities fixed, a higher κ_p/κ_w lowers limiting inflation and the fully anticipated nominal rate. When some gains arrive as surprises ($\lambda < 1$), increasing η strictly lowers $R^{\infty}(\lambda, \eta)$ without changing limiting inflation. If $R^{\infty}(0, \eta) < 0 < R^{\infty}(1, \eta)$, the anticipated share at which the rate crosses zero strictly increases with η.

Appendix B provides the proof. Under discretion, the fully anticipated rate also varies with η; the same threshold ranking holds in the numerical experiments above.

D. Commitment and Discretion

Figures 1 and 2 reveal a systematic difference in how real wages adjust. Discretion generates lower wage inflation and greater price deflation throughout the announced experiment and after the initial adjustment in the surprise experiment. Through (4), real-wage adjustment therefore relies more on falling prices and less on nominal wage growth.

Under commitment, promises about future policy influence current price and wage setting and subsequently constrain policy, as reflected in (8). The temporary positive output gap during the recovery in both figures also illustrates this history dependence.

Under discretion, current policy instead determines the real-wage gap inherited by future policymakers. Their anticipated responses affect future marginal costs and hence current price setting. When a wage shortfall is expected to persist, firms anticipate low marginal costs beyond the current period, reinforcing downward pressure on prices today. This feedback enters through the M-dependent coefficients in (9). Greater persistence M strengthens this feedback, raising the effective marginal-cost coefficient in the price Phillips curve to $A(M) = \kappa_p/(1 - \beta M) > \kappa_p$ for $0 < M < 1$ (Appendix B). This feedback helps explain the greater reliance on deflation in these experiments when the policymaker anticipates future policy responses but cannot choose them independently of its successors.

Finally, in this illustration, Ramsey policy reduces discounted quadratic loss relative to the Taylor rule by about one-third along the announced path. With repeated surprises and hand-to-mouth

households, the reduction is only 0.06 percent. Commitment lowers loss relative to discretion by less than 0.2 percent in all four cases (Appendix C).

IV. Conclusion

I show that an AI productivity boom can warrant lower nominal rates even when most gains are anticipated. Anticipated gains raise the natural rate, but below-target inflation supports real-wage adjustment and limits nominal-rate increases. Surprises leave the natural rate unchanged and call for cuts.

With hand-to-mouth households, surprises temporarily shift income toward profits. The expected reversal requires deeper cuts to sustain savers' spending despite similar aggregate allocations. In the benchmark calibration, nominal rates rise at the boom's midpoint once the anticipated share exceeds roughly three-quarters without hand-to-mouth households and just over four-fifths with them. The threshold rises as wages become stickier relative to prices. Beyond the initial adjustment, discretion relies more on falling prices and less on nominal wage growth than commitment.

A. Model Details and Welfare

The final consumption good is a CES aggregate of differentiated goods with elasticity ϵ. Firm j produces $Y_t(j) = \exp(a_t)N_t(j)$, where $N_t(j)$ is a CES aggregate of labor varieties with elasticity ϵ_w. Household utility is $u(C_t) - \int_0^1 v(N_t(h))\,\mathrm{d}h$. Each occupation h contains the same fraction η of hand-to-mouth households. Hours $N_t(h)$ are identical across types within an occupation. Within-type pooling insures occupational earnings risk but does not insure households against differences between labor income and profits.

Let L_t be the aggregate real labor bill. A revenue subsidy $\tau_p = \epsilon/(\epsilon - 1) - 1$ removes the goods markup. A wage subsidy $\tau_w = \epsilon_w/(\epsilon_w - 1) - 1$ removes the labor markup. Each household receives its pooled wage subsidy and pays a lump-sum levy $\tau_w L_t$. An atomistic wage setter takes aggregate taxes as given. Savers receive firm profits $[(1 + \tau_p)Y_t - L_t]/(1 - \eta)$ and pay the production-subsidy bill $\tau_p Y_t/(1 - \eta)$. Thus equilibrium budgets, after imposing zero aggregate bond holdings, are

$$C_t^H = L_t, \qquad C_t^S = L_t + \frac{Y_t - L_t}{1 - \eta}, \qquad \eta C_t^H + (1 - \eta)C_t^S = Y_t. \tag{A1}$$

The efficient steady state has $L = Y = C^H = C^S$ and real marginal cost one. To first order, the labor share has log deviation m_t; Calvo dispersion is second order. Log-linearizing (A1) gives (1). Substituting saver consumption into $c_t^S = \mathrm{E}_t c_{t+1}^S - (R_t - \mathrm{E}_t \pi_{t+1})/\sigma$ gives (5) after subtracting natural consumption growth.

A union chooses its nominal wage to maximize the population-weighted utility of members, subject to labor demand and the probability that its wage remains fixed. Its marginal utility of income is $\eta u'(C_t^H) + (1 - \eta)u'(C_t^S)$. The log deviation is $-\sigma y_t$. Combining this with labor disutility gives the desired real-wage deviation $\sigma y_t + \varphi n_t$. Relative to the actual real wage, the wage-setting gap is therefore $\psi x_t - m_t$. Linearizing the Calvo reset-wage condition gives (3); the factor $1 + \epsilon_w \varphi$ accounts for hours responding to an occupation's relative wage. The analogous firm reset-price condition gives (2). The identity $w_t = a_t + m_t$ gives (4).

For welfare, let $C_t = \eta C_t^H + (1 - \eta)C_t^S$ in levels. A second-order expansion around common steady-state consumption gives

$$\frac{\eta u(C_t^H) + (1 - \eta)u(C_t^S) - u(C_t)}{u'(\bar{C})\bar{C}} = -\frac{\sigma}{2}\eta(1 - \eta)d_t^2 + O(3) = -\frac{\sigma\nu}{2}m_t^2 + O(3). \tag{A2}$$

Type-specific consumption deviations from their level-weighted mean sum to zero. This eliminates both the linear distributional term and its cross term with aggregate consumption. Labor disutility is common across types. Efficient surplus has normalized curvature $-\psi$ with respect to log output, giving the loss $\psi x_t^2/2$.

Let $\mathcal{V}_t^p$ and $\mathcal{V}_t^w$ denote cross-sectional variances of log prices and wages. The CES aggregators imply dispersion losses $\epsilon\mathcal{V}_t^p/2$ and $\epsilon_w(1+\epsilon_w\varphi)\mathcal{V}_t^w/2$. For $j \in \{p, w\}$, with $\pi_t^p = \pi_t$, the Calvo recursion is

$$\mathcal{V}_t^j = \theta_j \mathcal{V}_{t-1}^j + \frac{\theta_j}{1-\theta_j}(\pi_t^j)^2 + O(3).$$

Discounted summation yields coefficients ϵ/κ_p and ϵ_w/κ_w on squared inflation. Combining these losses with (A2) gives (6)–(7), up to policy-independent and third-order terms. Inherited dispersion contributes only a predetermined term. Equal steady-state consumption, common labor assignments, and the stated fiscal incidence are essential for this particular welfare expression.

B. Proofs

Proof of Proposition 1. Define $K = \lambda_x/(\kappa_w\psi)$ and $F = \sigma/(\beta K)$. The commitment coefficients are

$$\begin{aligned} a_\pi &= F(\kappa_p\lambda_p + \kappa_w\lambda_w - \lambda_m), \\ a_w &= F[\lambda_w(\kappa_p - \kappa_w) + \lambda_m], \\ a_g &= F(\kappa_w\lambda_w - \lambda_m), \\ \bar{a}_x &= F[\kappa_w\psi\lambda_w + K(\kappa_p + \kappa_w)] + \frac{\sigma}{\beta}. \end{aligned} \tag{B1}$$

Attach multipliers $\alpha_t, \gamma_t, \delta_t$ to the price Phillips curve, wage Phillips curve, and real-wage law, writing each constraint with its current endogenous variable on the left. The first-order conditions are

$$\lambda_x x_t - \kappa_w\psi\gamma_t = 0, \tag{B2}$$

$$\lambda_p\pi_t + \alpha_t - \alpha_{t-1} + \delta_t = 0, \tag{B3}$$

$$\lambda_w\pi_t^w + \gamma_t - \gamma_{t-1} - \delta_t = 0, \tag{B4}$$

$$\lambda_m m_t - \kappa_p\alpha_t + \kappa_w\gamma_t + \delta_t - \beta \mathrm{E}_t\delta_{t+1} = 0. \tag{B5}$$

Initial multipliers satisfy $\alpha_{-1} = \gamma_{-1} = 0$. The positive quadratic objective, affine constraints, and transversality conditions make these conditions sufficient for the LQ optimum.

Equations (B2)–(B4) imply $\gamma_t = Kx_t$, $\delta_t = \lambda_w\pi_t^w + K\Delta x_t$, and $\Delta\alpha_t = -\lambda_p\pi_t - \delta_t$. For $t \geq 1$, subtract the date-$t-1$ version of (B5) from its date-t version. Keeping each conditional expectation at its original date gives

$$\begin{aligned} 0 = \lambda_m\Delta m_t + \kappa_p\lambda_p\pi_t + (\kappa_p + \kappa_w)K\Delta x_t + \kappa_p\lambda_w\pi_t^w \\ + \delta_t - \delta_{t-1} - \beta\mathrm{E}_t\delta_{t+1} + \beta\mathrm{E}_{t-1}\delta_t. \end{aligned} \tag{B6}$$

Use the wage Phillips curves at t and $t-1$, and $u_t = \sigma \mathrm{E}_t \Delta x_{t+1}$. Equation (B6) becomes

$$\begin{aligned}\frac{\beta K}{\sigma} u_t = \frac{\beta K}{\sigma} u_{t-1} &+ [K(1+\kappa_p+\kappa_w) + \kappa_w \psi \lambda_w]\Delta x_t - K\Delta x_{t-1} \\ &+ \kappa_p \lambda_p \pi_t + \kappa_p \lambda_w \pi_t^w + (\lambda_m - \kappa_w \lambda_w)\Delta m_t.\end{aligned}$$

Substitute $\Delta m_t = \pi_t^w - \pi_t - g_t$ and use (B1) to obtain

$$u_t = u_{t-1} + \bar{a}_x \Delta x_t - \frac{\sigma}{\beta}\Delta x_{t-1} + a_\pi \pi_t + a_w \pi_t^w + a_g g_t. \tag{B7}$$

At $t = 0$, the undifferenced first-order conditions give $u_0 = \bar{a}_x x_0 + a_\pi \pi_0 + a_w \pi_0^w + a_g g_0$. Thus (B7) holds from impact under the stated initial conventions. Substituting $u_t = s_t + \mathcal{H}_t$ proves (8). No realized outcome has replaced an earlier forecast, so the argument applies to both announced paths and repeated surprises. Together with the initial and terminal conditions, the elimination can be reversed to recover the multipliers. □

Proof of Proposition 2. Let $\mathcal{G}_t$ denote the exogenous information state, including the known future growth path when appropriate. The discretionary value function is

$$V_t(m_{t-1}; \mathcal{G}_t) = \min_{x_t, \pi_t, \pi_t^w, m_t} \left\{ \frac{\lambda_x x_t^2 + \lambda_p \pi_t^2 + \lambda_w (\pi_t^w)^2 + \lambda_m m_t^2}{2} + \beta \mathrm{E}_t V_{t+1}(m_t; \mathcal{G}_{t+1}) \right\},$$

subject to (2)–(4), with future policies taken as functions of m_t and the exogenous state. Define their slopes with respect to the inherited real-wage gap as $p = \partial \pi_t / \partial m_{t-1}$, $w = \partial \pi_t^w / \partial m_{t-1}$, and $h = \partial x_t / \partial m_{t-1}$. Here w without a time subscript denotes a policy slope, not the real wage. Differentiating the three constraints gives

$$\begin{aligned}p(M) &= \frac{\kappa_p M}{1-\beta M}, \qquad w(M) = M - 1 + p(M), \\ h(M) &= \frac{-1 + (1+\beta+\kappa_p+\kappa_w)M - \beta M^2}{\kappa_w \psi}, \\ A(M) &= \kappa_p + \beta p(M) = \frac{\kappa_p}{1-\beta M}, \qquad B(M) = \kappa_w - \beta w(M).\end{aligned}$$

The current first-order conditions give $\gamma_t = K x_t$, $\delta_t = \lambda_w \pi_t^w + K x_t$, and $\alpha_t = -\lambda_p \pi_t - \delta_t$. The envelope condition is $V_{t+1,m} = -\delta_{t+1}$. The remaining first-order condition therefore becomes

$$\begin{aligned}0 = A\lambda_p \pi_t &+ \lambda_w [(1+A)\pi_t^w - \beta \mathrm{E}_t \pi_{t+1}^w] \\ &+ K[(1+A+B)x_t - \beta \mathrm{E}_t x_{t+1}] + \lambda_m m_t.\end{aligned} \tag{B8}$$

Differentiating this condition with respect to m_{t-1} gives the scalar equation

$$0 = \mathfrak{F}(M) \equiv A\lambda_p p + \lambda_w(1 + A - \beta M)w + K(1 + A + B - \beta M)h + \lambda_m M. \tag{B9}$$

All functions on the right are evaluated at M. An admissible root has $|M| < 1$ and continuation curvature

$$P = \frac{\lambda_x h^2 + \lambda_p p^2 + \lambda_w w^2 + \lambda_m M^2}{1 - \beta M^2} = -Kh - \lambda_w w \geq 0. \tag{B10}$$

Set $k = \kappa_w \psi$ and $C_M = 1 + A + B$. Substitution into the one-period problem gives curvature

$$Q = \lambda_x (C_M/k)^2 + \lambda_p A^2 + \lambda_w (1 + A)^2 + \lambda_m + \beta P > 0. \tag{B11}$$

The homogeneous best response is $M = [\lambda_x C_M/k^2 + \lambda_w(1 + A)]/Q$. Equations (B9)–(B11) impose both optimization and consistency of the continuation value. The calculations follow the admissible branch continuous from $\eta = 0$; the proposition does not require global uniqueness of Markov equilibria.

Using the wage Phillips curve to eliminate expected wage inflation in (B8), then substituting (4), gives $u_t = b_x x_t + b_\pi \pi_t + b_w \pi_t^w + b_g(g_t - m_{t-1})$, with

$$\begin{aligned} b_\pi &= F(A\lambda_p + \kappa_w \lambda_w - \lambda_m), \\ b_w &= F[\lambda_w(A - \kappa_w) + \lambda_m], \qquad b_g = F(\kappa_w \lambda_w - \lambda_m), \\ b_x &= \frac{\sigma}{\beta}\left[1 + \kappa_p + \kappa_w - \beta M + \frac{\kappa_w^2 \psi^2 \lambda_w}{\lambda_x}\right]. \end{aligned} \tag{B12}$$

Finally, (5) converts u_t into the nominal rate in (9). □

For completeness, affine terms close the solution for any announced growth path. Write $m_t = M m_{t-1} + \mu_t$, $\pi_t = p m_{t-1} + \bar{\pi}_t$, $\pi_t^w = w m_{t-1} + \bar{\pi}_t^w$, $x_t = h m_{t-1} + \bar{x}_t$, and $V_t(s) = Ps^2/2 + J_t s + \text{constant}$. Let $f_t = \beta \bar{\pi}_{t+1}$, $q_t = g_t + f_t$, and $e_t = g_t + \beta(\bar{\pi}_{t+1} - \bar{\pi}_{t+1}^w)$. Backward recursion gives

$$\begin{aligned} \mu_t &= -\frac{\lambda_x C_M e_t/k^2 + \lambda_p A f_t + \lambda_w (1 + A) q_t + \beta J_{t+1}}{Q}, \\ \bar{\pi}_t &= A\mu_t + f_t, \qquad \bar{\pi}_t^w = (1 + A)\mu_t + q_t, \\ \bar{x}_t &= (C_M \mu_t + e_t)/k, \qquad J_t = -\lambda_x \bar{x}_t/k - \lambda_w \bar{\pi}_t^w. \end{aligned}$$

For a stochastic exogenous process, replace future affine terms by their conditional expectations. The homogeneous slopes are unchanged.

Proof of Proposition 3. Summing (B3) and (B4) in the stationary limit gives $\lambda_p \pi^\infty + \lambda_w (\pi^w)^\infty = 0$. Equation (4) gives $(\pi^w)^\infty - \pi^\infty = g$, establishing the inflation formula for every λ and η. Under full

anticipation, expected changes in m_t and x_t vanish, so (5) gives the stated nominal rate. The rigidity comparison follows from

$$\omega = \left[1 + \frac{\epsilon}{\epsilon_w}\frac{\kappa_w}{\kappa_p}\right]^{-1}.$$

To establish the comparative static in η, set $q = \lambda_m = \sigma\eta/(1-\eta)$ and let $\pi_0(q), x_0(q), m_0(q)$ denote the impact responses to $g_0 = 1$, with zero inherited states and $g_t = 0$ thereafter. By linearity, these are also the incremental responses to each later innovation, without resetting inherited promises. Let $R^S = R^\infty(0,\eta)$ and $R^A = R^\infty(1,\eta)$. At the repeated-surprise limit, next period's realization differs from its conditional forecast by the impact response times g. Thus (5) implies

$$\frac{R^S}{g} = -\omega - \pi_0 - \sigma x_0 + q m_0.$$

It therefore suffices to show

$$\pi_0'(q) > 0, \qquad x_0'(q) > 0, \qquad [q m_0(q)]' < 0,$$

where primes denote derivatives with respect to q. These are responses to a single innovation; limiting inflation along the repeated-surprise path remains $-\omega g$.

To establish the signs, first solve over dates $0, \ldots, L$, setting $\pi_{L+1} = \pi^w_{L+1} = 0$, and then let L increase. Write $\widetilde{z}_t = \beta^{t/2} z_t$ and collect each variable's discounted values in a column vector. Let D_L be the lower bidiagonal matrix with ones on its diagonal and $-\sqrt{\beta}$ immediately below it, and let $e = (1, 0, \ldots, 0)^\top$. Define $T_L = D_L D_L^\top$, $B_L = I + T_L/\kappa_p$, $k = \kappa_w \psi$, and $c = \kappa_w/\kappa_p$. The private equilibrium conditions become

$$\widetilde{m} = \frac{D_L^\top \widetilde{\pi}}{\kappa_p}, \qquad \widetilde{\pi}^w = B_L \widetilde{\pi} + e, \qquad \widetilde{x} = \frac{D_L^\top[(B_L + cI)\widetilde{\pi} + e]}{k}.$$

The discount factors are now absorbed into the vectors, so (6) is an ordinary sum of squared entries.

The symmetric matrix T_L has orthonormal eigenvectors v_j and eigenvalues s_j satisfying

$$0 < \underline{s} \equiv (1-\sqrt{\beta})^2 \le s_j \le (1+\sqrt{\beta})^2 \equiv \overline{s}.$$

These bounds follow from $(1-\sqrt{\beta})\|v\| \le \|D_L^\top v\| \le (1+\sqrt{\beta})\|v\|$. Set $\rho_j = (v_j^\top e)^2$, so $\rho_j \ge 0$ and $\sum_j \rho_j = 1$, and define the positive coefficients

$$\begin{aligned} b_j &= 1 + s_j/\kappa_p, \\ n_j &= \lambda_w b_j + \frac{\lambda_x}{k^2} s_j (b_j + c), \\ h_j(q) &= \lambda_p + \lambda_w b_j^2 + \frac{\lambda_x}{k^2} s_j (b_j + c)^2 + \frac{q s_j}{\kappa_p^2}. \end{aligned}$$

Writing $u_j = v_j^\top \widetilde{\pi}$, substitution into the loss gives, up to a constant,

$$\frac{1}{2}\sum_j \left[h_j(q)u_j^2 + 2n_j(v_j^\top e)u_j\right].$$

Hence $u_j = -n_j(v_j^\top e)/h_j(q)$. Since $D_L^\top e = e$, the finite-horizon impact responses are

$$\begin{aligned}
\pi_{0,L} &= -\sum_j \rho_j \frac{n_j}{h_j(q)}, \\
x_{0,L} &= \frac{1}{k}\left[1 - \sum_j \rho_j \frac{s_j(b_j + c)n_j}{h_j(q)}\right], \\
qm_{0,L} &= -\frac{1}{\kappa_p}\sum_j \rho_j s_j n_j \frac{q}{h_j(q)}.
\end{aligned}$$

Every coefficient other than $h_j(q)$ is independent of q. Since

$$h_j'(q) = \frac{s_j}{\kappa_p^2} > 0, \qquad \frac{d}{dq}\frac{q}{h_j(q)} = \frac{h_j(0)}{h_j(q)^2} > 0,$$

the three required signs hold at every finite horizon.

The signs persist as $L \to \infty$. All eigenvalues lie in the same compact interval $[\underline{s}, \overline{s}]$, and $h_j(q) \geq \lambda_p > 0$. The scalar functions in these sums and their q-derivatives admit uniform polynomial approximations on that interval, locally uniformly in q. For any fixed polynomial, each fixed entry of that polynomial in T_L stops changing once L is sufficiently large, because T_L is tridiagonal. Thus the solution and its q-derivative converge at each fixed date. For each fixed q, the derivative terms after removing ρ_j are bounded strictly away from zero with the stated signs, so the inequalities remain strict in the limit. The limiting allocation satisfies the infinite-horizon first-order conditions and has finite discounted loss: the feasible choice $\pi_t = 0$ bounds the finite-horizon optimum by $(\lambda_w + \lambda_x/k^2)/2$. Strict convexity therefore identifies the limit as the Ramsey allocation.

Consequently,

$$\frac{1}{g}\frac{\partial R^S}{\partial q} = -\pi_0'(q) - \sigma x_0'(q) + [qm_0(q)]' < 0.$$

Since $dq/d\eta = \sigma/(1-\eta)^2 > 0$, R^S strictly decreases with η.

Linearity gives $R^\infty(\lambda, \eta) = \lambda R^A + (1-\lambda)R^S$. Because R^A is independent of η, the rate strictly decreases with η whenever $\lambda < 1$. If $R^S < 0 < R^A$, the unique crossing share is

$$\lambda^* = \frac{-R^S}{R^A - R^S}, \qquad \frac{d\lambda^*}{d\eta} = -\frac{R^A}{(R^A - R^S)^2}\frac{dR^S}{d\eta} > 0.$$

□

C. Numerical Solution and Robustness

Announced MIT shock. At date zero, agents learn the full productivity path. I solve the model under perfect foresight over 240 quarters, starting from steady state and preserving the permanent productivity gain in the terminal conditions. The benchmark uses the specified Taylor rule, which is determinate in every reported calibration. Under commitment, I impose the Ramsey first-order conditions with zero initial multipliers and recover the nominal rate from (5). Under discretion, I first solve (B9) numerically for M, selecting the admissible solution described in Appendix B. I then solve the model with (9) as the policy rule. In every reported calibration, the resulting system in gaps satisfies the Blanchard–Kahn conditions: its unique stable solution reproduces the selected Markov equilibrium of Appendix B, with stable root M.

Repeated surprises. Repeated surprises form a sequence of MIT shocks. After each innovation, the conditional perfect-foresight path starts from the inherited state and assumes zero future productivity growth. Under commitment, inherited Ramsey multipliers preserve past promises. The next surprise updates the state and the conditional path. With partial anticipation, each continuation instead includes announced future growth $\lambda\bar{g}$ and assumes no further innovations. The same M applies across information experiments within each calibration.

Welfare. Discounted losses evaluate (6) along realized paths over 240 quarters from the first shock, comparing policies at fixed household shares and information. Listing Taylor, Ramsey, and discretion, announced losses are $(0.01442812, 0.00966240, 0.00967910)$ at $\eta = 0$ and $(0.01452288, 0.00967799, 0.00969237)$ at $\eta = 0.30$. Repeated-surprise losses are $(0.00937686, 0.00903969, 0.00904965)$ and $(0.00958265, 0.00957734, 0.00958200)$, respectively.

Thresholds. Across quarters 1–40, the anticipated shares at which nominal-rate deviations cross zero vary as follows. Listing Taylor, Ramsey, and discretion, at $\sigma = 1$ the ranges are (0.225–0.795, 0.433–0.765, 0.372–0.775) for $\eta = 0$ and (0.279–0.858, 0.507–0.818, 0.470–0.831) for $\eta = 0.30$. Across quarters 10–30, the optimal thresholds stay within 2.3 percentage points of their quarter-20 values.

Preferences. At $\sigma = 0.5$, both optimal nominal rates remain negative at quarter 20 even under full anticipation. Taylor thresholds are 0.903 for $\eta = 0$ and 0.918 for $\eta = 0.30$. At $\sigma = 2$, quarter-20 thresholds are $(0.482, 0.487, 0.505)$ for $\eta = 0$ and $(0.585, 0.691, 0.698)$ for $\eta = 0.30$, listing Taylor, Ramsey, and discretion. These exercises hold other structural parameters fixed and recompute welfare weights and policy solutions; the Taylor benchmark remains determinate.

Rigidity. Rigidity exercises vary one parameter at a time with $\sigma = 1$, recomputing weights and policies. Quarter-20 thresholds list Taylor, Ramsey, and discretion for $\eta = 0$ and 0.30. With $\theta_w = 0.50$ ($\kappa_w \simeq 0.0721$), they are (0.29, 0.17, 0.30) and (0.41, 0.44, 0.50). With $\theta_w = 0.65$

($\kappa_w \simeq 0.0274$), (0.46, 0.52, 0.57) and (0.55, 0.67, 0.69). With $\theta_w = 0.85$ ($\kappa_w \simeq 0.0040$), (0.81, 0.90, 0.91) and (0.84, 0.93, 0.93). With $\kappa_p = 0.02$, (0.50, 0.53, 0.57) and (0.61, 0.67, 0.70). With $\kappa_p = 0.10$, (0.66, 0.84, 0.85) and (0.71, 0.89, 0.90). The optimal thresholds exceed the Taylor threshold in every case except $\theta_w = 0.50$ with $\eta = 0$. Quarter-20 Ramsey price inflation under full anticipation is −0.41, −0.65, −0.92, −0.62, and −0.89 annualized percentage points in the same order, against −0.80 in the benchmark.

References

Acemoglu, Daron. 2025. "The Simple Macroeconomics of AI." *Economic Policy* 40 (121): 13–58. https://doi.org/10.1093/epolic/eiae042.

Acharya, Sushant, Edouard Challe, and Keshav Dogra. 2023. "Optimal Monetary Policy According to HANK." *American Economic Review* 113 (7): 1741–82. https://doi.org/10.1257/aer.20200239.

Auclert, Adrien. 2019. "Monetary Policy and the Redistribution Channel." *American Economic Review* 109 (6): 2333–67. https://doi.org/10.1257/aer.20160137.

Barattieri, Alessandro, Susanto Basu, and Peter Gottschalk. 2014. "Some Evidence on the Importance of Sticky Wages." *American Economic Journal: Macroeconomics* 6 (1): 70–101. https://doi.org/10.1257/mac.6.1.70.

Barlevy, Gadi, Jonas D. M. Fisher, Will Pennington, and Alessandro Villa. 2026. "The AI Revolution and Monetary Policy." *Chicago Fed Letter* 528. https://doi.org/10.21033/cfl-2026-528.

Bilbiie, Florin O. 2008. "Limited Asset Markets Participation, Monetary Policy and (Inverted) Aggregate Demand Logic." *Journal of Economic Theory* 140 (1): 162–96. https://doi.org/10.1016/j.jet.2007.07.008.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2025. "Generative AI at Work." *The Quarterly Journal of Economics* 140 (2): 889–942. https://doi.org/10.1093/qje/qjae044.

Christiano, Lawrence J., Cosmin L. Ilut, Roberto Motto, and Massimo Rostagno. 2010. "Monetary Policy and Stock Market Booms." NBER Working Paper 16402. https://doi.org/10.3386/w16402.

Dávila, Eduardo, and Andreas Schaab. 2023. "Optimal Monetary Policy with Heterogeneous Agents: Discretion, Commitment, and Timeless Policy." NBER Working Paper 30961. https://doi.org/10.3386/w30961.

Debortoli, Davide, and Jordi Galí. 2025. "Heterogeneity and Aggregate Fluctuations: Insights from TANK Models." *NBER Macroeconomics Annual* 39: 307–57. https://doi.org/10.1086/735272.

Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin. 2000. "Optimal Monetary Policy with Staggered Wage and Price Contracts." *Journal of Monetary Economics* 46 (2): 281–313. https://doi.org/10.1016/S0304-3932(00)00028-3.

Giannoni, Marc P., and Michael Woodford. 2003. "How Forward-Looking Is Optimal Monetary Policy?" *Journal of Money, Credit and Banking* 35 (6, Part 2): 1425–69. https://www.jstor.org/stable/3649892.

Gilchrist, Simon, and John V. Leahy. 2002. "Monetary Policy and Asset Prices." *Journal of Monetary Economics* 49 (1): 75–97. https://doi.org/10.1016/S0304-3932(01)00093-9.

Goolsbee, Austan. 2026. "Hoover Institution Monetary Policy Conference 2026." Remarks at the Hoover Institution, Stanford, CA, May 8. Federal Reserve Bank of Chicago.

Grigsby, John, Erik Hurst, and Ahu Yildirmaz. 2021. "Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data." *American Economic Review* 111 (2): 428–71. https://doi.org/10.1257/aer.20190318.

Hansen, Niels-Jakob H., Alessandro Lin, and Rui C. Mano. 2023. "Should Inequality Factor into Central Banks' Decisions?" Bank of Italy Temi di Discussione 1410. https://doi.org/10.32057/0.TD.2023.1410.

Kaplan, Greg, Giovanni L. Violante, and Justin Weidner. 2014. "The Wealthy Hand-to-Mouth." *Brookings Papers on Economic Activity* Spring: 77–138. https://doi.org/10.1353/eca.2014.0002.

McKay, Alisdair, and Christian K. Wolf. 2026. "Optimal Policy Rules in HANK." *Review of Economic Studies*, forthcoming.

Miyazaki, Kenji. 2026a. "Inherited Wage Dispersion and Optimal Discretion in a Dual-Rigidity TANK Model." arXiv:2609.39070v1, September 30.

Miyazaki, Kenji. 2026b. "Profit Redistribution in a TANK Model with Price and Wage Rigidities: An Analytical Benchmark and an Own-Lag Comparison." arXiv:2508.12073v4, August 25. https://doi.org/10.48550/arXiv.2508.12073.

Nakamura, Emi, and Jón Steinsson. 2008. "Five Facts about Prices: A Reevaluation of Menu Cost Models." *The Quarterly Journal of Economics* 123 (4): 1415–64. https://doi.org/10.1162/qjec.2008.123.4.1415.

Offick, Sven, and Hans-Werner Wohltmann. 2016. "Volatility Effects of News Shocks in New Keynesian Models with Optimal Monetary Policy." *Economics Letters* 147: 78–82. https://doi.org/10.1016/j.econlet.2016.08.018.