

What is loose commitment?

Nicolas Franz-Pattillo
2024-05-01

Table of contents

The economy	2
Households	2
Firms and the non-linear Phillips curve	2
The private-sector equilibrium	3
The benchmark: flexible prices and an efficient subsidy	4
The efficient allocation	4
Flexible prices with market power	4
The subsidy that undoes the markup	5
What this means for monetary policy	5
Why optimal policy is time-inconsistent	6
Three policy problems	6
Discretion	7
Commitment	7
Loose commitment	9
Measuring welfare	10
Solving it without linearizing	10
Where this leaves us	11

Textbook monetary policy comes in two flavors. Either the central bank follows a rule, such as a Taylor rule, or it chooses policy optimally, in which case one has to say whether it can **commit** to its promises or acts under **discretion** (Kydland and Prescott (1977)). Outside of the textbook, we may find something in between.

This note explains the machinery behind a model of that middle ground, **loose commitment**, as used in my paper *Optimal Monetary Policy under Loose Commitment and the Zero Lower Bound*. Here the goal is to show how the math works: the model, the three policy problems, and why solving them requires a global (non-linear) method.

The economy

The model is a small New Keynesian economy with three ingredients: sticky prices Rotemberg (1982), monopolistic competition and a lower bound on the nominal interest rate. It is deliberately small; the point is to isolate how commitment interacts with those three frictions, not to fit the data.

Households

A representative household chooses consumption c_t , hours n_t and nominal bonds to maximize

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \Delta_t \left[\frac{c_t^{1-\sigma} - 1}{1-\sigma} - \chi \frac{n_t^{1+\eta}}{1+\eta} \right], \quad \Delta_{t+1} = \delta_t \Delta_t,$$

where δ_t is a preference shock: when δ_t is high the household becomes more patient and wants to save. It follows $\log \delta_t = \rho_\delta \log \delta_{t-1} + \sigma_\delta \epsilon_t$. This is the shock that pushes the economy to the zero lower bound, because a sudden desire to save lowers the interest rate the economy needs.

The household's choices give two familiar conditions, the Euler equation and the labor supply condition:

$$\frac{c_t^{-\sigma}}{R_t} = \delta_t \beta \mathbb{E}_t \left[\frac{c_{t+1}^{-\sigma}}{\pi_{t+1}} \right], \quad \chi n_t^\eta = c_t^{-\sigma} w_t (1 + \tau),$$

where R_t is the gross nominal interest rate, $\pi_t = p_t/p_{t-1}$ gross inflation, w_t the real wage and τ a labor subsidy.

Firms and the non-linear Phillips curve

A competitive final-goods firm combines a continuum of intermediate goods with a constant elasticity of substitution $1/\nu$. Intermediate firm j therefore faces the demand $y_t(j) = [p_t(j)/p_t]^{-1/\nu} y_t$, produces with $y_t(j) = z_t n_t(j)$ and, following Rotemberg, pays a quadratic cost to change its price.¹

$$AC_t(j) = \frac{\phi}{2} \left[\frac{p_t(j)}{p_{t-1}(j)} - \pi_f \right]^2 y_t.$$

¹I use Rotemberg pricing because it gives a smooth, fully nonlinear price-adjustment problem that is convenient for global solution and for studying an occasionally binding lower bound (Rotemberg (1982)). It also avoids introducing an additional endogenous distribution of firms' prices as a state variable in this formulation. The familiar Calvo model instead gives each firm a fixed probability of being able to reoptimize its price in each period (Calvo (1983)), which generally requires tracking price dispersion. The two formulations generate the same New Keynesian Phillips curve only locally, after a first-order linearization around a steady state and with an appropriate mapping between parameters; they are not the same model away from that approximation. In particular, their nonlinear welfare costs, price dispersion, and responses near the ZLB can differ. A third possibility is state-dependent pricing, in which firms adjust when their individual state makes it worthwhile rather than according to a fixed probability (Dotsey et al. (1999)), adding a distribution of pricing states to the aggregate state.

Two parameters matter here. ϕ governs how sticky prices are. π_f is the inflation rate that is free of cost; the literature usually sets $\pi_f = 1$ (zero net inflation), and so does the benchmark calibration.

Each firm chooses its price path to maximize the discounted value of profits, using the household's stochastic discount factor (Yun (2005)). In a symmetric equilibrium this yields the **non-linear Phillips curve**:

$$\begin{aligned} 1 - v - \frac{w_t}{z_t} + \phi v \pi_t (\pi_t - \pi_f) \\ = \phi v \delta_t \beta \mathbb{E}_t \left[\left(\frac{c_{t+1}}{c_t} \right)^{-\sigma} \frac{y_{t+1}}{y_t} (\pi_{t+1} - \pi_f) \pi_{t+1} \right]. \end{aligned}$$

i Derivation: from the firm's problem to the Phillips curve

Firm j solves

$$\max_{\{p_{t+s}(j)\}} \mathbb{E}_t \sum_{s \geq 0} Q_{t+s|t} \left[p_{t+s}(j) y_{t+s}(j) - p_{t+s} \frac{w_{t+s}}{z_{t+s}} y_{t+s}(j) - p_{t+s} AC_{t+s}(j) \right],$$

with $y_t(j) = [p_t(j)/p_t]^{-1/v} y_t$ and the nominal discount factor $Q_{t+1|t} = \delta_t \beta (c_{t+1}/c_t)^{-\sigma} / \pi_{t+1}$.

The price $p_t(j)$ appears in period t (revenue, labor cost and today's adjustment cost) and in period $t + 1$ (tomorrow's adjustment cost, through $p_{t+1}(j)/p_t(j)$). Differentiating and imposing symmetry, $p_t(j) = p_t$:

- revenue: $\left(1 - \frac{1}{v}\right) y_t$;
- labor cost: $+\frac{1}{v} \frac{w_t}{z_t} y_t$;
- today's adjustment cost: $-\phi(\pi_t - \pi_f) \pi_t y_t$;
- tomorrow's adjustment cost: $+\mathbb{E}_t \left[Q_{t+1|t} \frac{p_{t+1}}{p_t} \phi(\pi_{t+1} - \pi_f) \pi_{t+1} y_{t+1} \right]$.

Setting the sum to zero, using $Q_{t+1|t} p_{t+1}/p_t = \delta_t \beta (c_{t+1}/c_t)^{-\sigma}$ and multiplying by $-v/y_t$ gives the equation above. With flexible prices ($\phi = 0$) it reduces to $w_t/z_t = 1 - v$: firms charge a constant markup $1/(1 - v)$ over marginal cost.

The private-sector equilibrium

Using market clearing and the resource constraint, which subtracts the resources used up in changing prices, the economy is summarized by

$$\begin{aligned}
\text{Euler:} \quad & \frac{c_t^{-\sigma}}{R_t} = \delta_t \beta \mathbb{E}_t \left[\frac{c_{t+1}^{-\sigma}}{\pi_{t+1}} \right] \\
\text{Phillips:} \quad & c_t^{-\sigma} y_t \left[1 - \nu - \frac{w_t}{Z_t} + \phi \nu \pi_t (\pi_t - \pi_f) \right] \\
& = \phi \nu \delta_t \beta \mathbb{E}_t \left[c_{t+1}^{-\sigma} y_{t+1} (\pi_{t+1} - \pi_f) \pi_{t+1} \right] \\
\text{Labor:} \quad & \chi n_t^\eta = c_t^{-\sigma} w_t (1 + \tau) \\
\text{Resources:} \quad & c_t = y_t \left[1 - \frac{\phi}{2} (\pi_t - \pi_f)^2 \right], \quad y_t = Z_t n_t \\
\text{ZLB:} \quad & R_t \geq 1
\end{aligned}$$

Only the Euler equation and the Phillips curve involve expectations. That will turn out to be the whole story: the only reason policy is hard is that today's outcomes depend on what people expect policy to do tomorrow.

The benchmark: flexible prices and an efficient subsidy

Before asking what a central bank can achieve, it helps to know what the best possible outcome is, and which frictions stand in the way.

The efficient allocation

A planner who could simply assign consumption and hours, with no prices involved, would maximize $u(c_t, n_t)$ subject to $c_t = z_t n_t$. At the optimum, the marginal disutility of an hour of work equals the consumption it produces:

$$\chi n_t^\eta = c_t^{-\sigma} z_t.$$

Combined with $c_t = z_t n_t$, this gives the efficient allocation:

$$n_t^* = \left[\frac{z_t^{1-\sigma}}{\chi} \right]^{\frac{1}{\eta+\sigma}}, \quad c_t^* = \left[\frac{z_t^{1+\eta}}{\chi} \right]^{\frac{1}{\eta+\sigma}}.$$

The market economy reaches this point only if prices are flexible ($\phi = 0$) and markets are competitive ($\nu = 0$). Monetary policy is then irrelevant.

Flexible prices with market power

Now keep prices flexible but let firms have market power ($\nu > 0$). With $\phi = 0$ the Phillips curve says that firms set a constant markup, so the real wage is below the marginal product of labor:

$$\frac{w_t}{z_t} = 1 - v.$$

Substituting into the household's labor supply condition, $\chi n_t^\eta = c_t^{-\sigma} w_t (1 + \tau)$, gives

$$\chi n_t^\eta = (1 - v)(1 + \tau) c_t^{-\sigma} z_t,$$

which differs from the efficient condition by the factor $(1 - v)(1 + \tau)$. The flexible-price allocation is therefore

$$n_t^{\text{flex}} = \left[\frac{(1 - v)(1 + \tau) z_t^{1-\sigma}}{\chi} \right]^{\frac{1}{\eta+\sigma}}, \quad c_t^{\text{flex}} = \left[\frac{(1 - v)(1 + \tau) z_t^{1+\eta}}{\chi} \right]^{\frac{1}{\eta+\sigma}}.$$

Without a subsidy ($\tau = 0$), the markup acts like a tax on work: people work too little and the economy produces too little. With the benchmark calibration ($v \approx 0.087$, $\sigma = 1$, $\eta = 0.5$), output is about 6% below its efficient level even with perfectly flexible prices.

The subsidy that undoes the markup

A labor subsidy can close this gap exactly. Setting

$$\tau^* = \frac{v}{1 - v} \implies (1 - v)(1 + \tau^*) = 1$$

makes the flexible-price allocation efficient. The government pays the subsidy on wages and finances it with a lump-sum tax of the same size, so its budget balances and the tax itself distorts nothing. This is the scheme that much of the optimal-policy literature assumes, for example Siu (2004) and Yun (2005).

What this means for monetary policy

With the subsidy in place and sticky prices ($\phi > 0$), the central bank can reproduce the efficient allocation by keeping inflation at its free-of-cost level every period, $\pi_t = \pi_f$. Adjustment costs are then zero, both sides of the Phillips curve collapse to $w_t/z_t = 1 - v$, and the economy behaves as if prices were flexible. The interest rate that delivers this follows from the Euler equation:

$$R_t = \frac{\pi_f c_t^{-\sigma}}{\delta_t \beta \mathbb{E}_t[c_{t+1}^{-\sigma}]}.$$

Two things can break this. First, a large enough preference shock makes the required R_t fall below one, which the lower bound rules out. Second, without the subsidy, keeping $\pi_t = \pi_f$ delivers the

flexible-price allocation, not the efficient one, and the central bank is tempted to push output up with inflation.

The paper keeps $\tau = 0$ on purpose. The subsidy would remove the second problem, and with it the inflation bias. Leaving the markup in place is what lets the model study how the two biases, inflationary in normal times and deflationary at the lower bound, trade off as commitment changes.

Why optimal policy is time-inconsistent

The central bank sets R_t to maximize the household's welfare, taking this equilibrium as a constraint. It would like to close three gaps, or **wedges**, but it cannot close all of them at once:

1. **Price adjustment.** Any inflation different from π_f wastes resources. This wedge says: keep $\pi_t = \pi_f$.
2. **Monopolistic competition.** Markups keep output below its efficient level. In the benchmark there is no subsidy ($\tau = 0$; see the previous section), so this wedge is open. Surprise inflation raises output, so this wedge says: inflate.
3. **The zero lower bound.** When the economy needs a negative nominal rate, the bound keeps real rates too high. This wedge says: promise higher inflation later, so that real rates fall today.

The second and third wedges create opposite temptations, and both are temptations to break a promise.

- **Inflation bias.** Suppose the central bank promises $\pi_{t+1} = \pi_f$ and firms believe it. Once expectations are set, a bit of surprise inflation raises output at no cost in credibility for today. Firms anticipate this, so in equilibrium inflation ends up above π_f with no gain in output.
- **Deflation bias.** At the lower bound, the central bank would like to promise higher future inflation to lower real rates now. But when the future arrives, the crisis is over and it would rather deliver π_f . If people anticipate that, the promise is worthless, and the economy spends longer at the bound with lower inflation.

This is Kydland and Prescott's time inconsistency in its cleanest form. What the central bank can do therefore depends on what it can promise, which is what the three policy problems below differ in.

Three policy problems

For all three, write $u(x_t) = \frac{c_t^{1-\sigma}-1}{1-\sigma} - \chi \frac{n_t^{1+\eta}}{1+\eta}$ and split each forward-looking equation into the part that depends on today's variables and the part that depends on tomorrow's:

$$\begin{aligned}
A(x_t) &= \frac{C_t^{-\sigma}}{R_t}, \\
B(x_{t+1}) &= \frac{C_{t+1}^{-\sigma}}{\pi_{t+1}}, \\
C(x_t) &= c_t^{-\sigma} y_t \left[1 - v - \frac{w_t}{z_t} + \phi v \pi_t (\pi_t - \pi_f) \right], \\
D(x_{t+1}) &= \phi v c_{t+1}^{-\sigma} y_{t+1} (\pi_{t+1} - \pi_f) \pi_{t+1}.
\end{aligned}$$

The Euler equation is then $A(x_t) = \delta_t \beta \mathbb{E}_t B(x_{t+1})$ and the Phillips curve is $C(x_t) = \delta_t \beta \mathbb{E}_t D(x_{t+1})$.

Discretion

Under discretion no promise binds: each period the central bank re-optimizes, taking how it will behave in the future as given. Because it has no incentive to preserve a promise when the next period arrives, promises have no independent force; it is as if they did not exist. The state therefore carries no burden from previous promises: policy depends only on the current exogenous shock, not on an inherited policy plan. The right-hand sides $\mathbb{E}_t B$ and $\mathbb{E}_t D$ are then outside its control, and the problem is static:

$$\max_{x_t} u(x_t)$$

subject to the equilibrium conditions, with $\mathbb{E}_t B$ and $\mathbb{E}_t D$ taken as given.

Discretionary models have more than one equilibrium. I focus on the Markov-perfect equilibrium in which policy depends only on the current shock and the lower bound binds occasionally; there is also a “deflation trap” equilibrium in which low expectations keep the economy at the bound, which I rule out.

The first-order condition for R_t gives the most useful insight in the whole problem. With multipliers λ_t^E on the Euler equation and $\lambda_t^R \geq 0$ on the lower bound,

$$\lambda_t^E \frac{C_t^{-\sigma}}{R_t^2} = \lambda_t^R.$$

Away from the bound, $\lambda_t^R = 0$, so $\lambda_t^E = 0$: **the Euler equation is not a constraint in normal times**. The central bank can always set the rate the Euler equation requires. Aggregate demand only becomes a binding problem at the lower bound. Keep this in mind for the next section.

Commitment

Under full commitment the central bank chooses, once and for all, a plan for every possible history of shocks. Attach multipliers λ_t^E and λ_t^P to the Euler equation and the Phillips curve, and λ_t^R to the

lower bound. The Lagrangian is

$$\begin{aligned} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \Delta_t \{ & u(x_t) + \lambda_t^E [A(x_t) - \delta_t \beta \mathbb{E}_t B(x_{t+1})] \\ & + \lambda_t^P [C(x_t) - \delta_t \beta \mathbb{E}_t D(x_{t+1})] \\ & + \lambda_t^R (R_t - 1) \}. \end{aligned}$$

The difficulty is that x_{t+1} appears in the constraints of period t . Marcet and Marimon's trick is to move those terms into period $t + 1$, where they belong. Because the discount weight of period $t + 1$ is $\beta^{t+1} \Delta_{t+1} = \beta^t \Delta_t \cdot \delta_t \beta$, the term $-\beta^t \Delta_t \lambda_t^E \delta_t \beta B(x_{t+1})$ is exactly $-\beta^{t+1} \Delta_{t+1} \lambda_t^E B(x_{t+1})$. So in period $t + 1$ the problem contains yesterday's multiplier. Define

$$\mu_{t+1} \equiv \lambda_t, \quad \mu_0 = 0.$$

The problem becomes recursive, a **saddle-point functional equation** with the lagged multipliers as a new state:

$$\begin{aligned} W(\delta, \mu) = \min_{\lambda} \max_x \{ & u(x) + \lambda^E A(x) - \mu^E B(x) \\ & + \lambda^P C(x) - \mu^P D(x) + \lambda^R (R - 1) \} \\ & + \delta \beta \mathbb{E} W(\delta', \lambda). \end{aligned}$$

This is where the word *promise* becomes literal. μ^E and μ^P are yesterday's promises about today; they enter today's problem as costs of deviating from what was announced. Today's multipliers become tomorrow's promises. The state now has two endogenous dimensions on top of the shock.

Two consequences follow from the first-order conditions.

No inflation bias in the long run. In a steady state without the bound, promises are constant, so $\mu = \lambda$. The terms $\lambda^P C(x) - \mu^P D(x)$ then pull in opposite directions with equal weight, and the first-order condition for inflation reduces to

$$\lambda^{\text{res}} n \phi(\pi - \pi_f) = 0,$$

where $\lambda^{\text{res}} > 0$ is the multiplier on the resource constraint. Hence $\pi = \pi_f$. Under discretion there is no μ^P term to offset λ^P , and inflation settles above π_f : the inflation bias.

Forward guidance lives in μ^E . As under discretion, $\lambda_t^E = 0$ whenever the bound does not bind. So the Euler promise μ^E is positive only after the economy has been at the bound. When it is, the first-order condition for inflation acquires an extra term,

$$\dots = \mu_t^E c_t^{-\sigma} \pi_t^{-2} + \dots,$$

which raises inflation after the episode ends. This is the credible promise of higher future inflation that makes commitment so effective at the bound.

Loose commitment

Loose commitment, from Debortoli and Nunes (2010), sits between the two. Each period an i.i.d. coin is tossed:

- with probability ξ the plan continues (*commit*);
- with probability $1 - \xi$ the central bank reneges and a new plan is drawn up from scratch (*renege*).

A promise made today therefore survives until tomorrow with probability ξ . When the central bank reneges, it starts again with no inherited promises, so the state jumps back to $\mu = 0$. The parameter ξ moves the economy continuously from discretion ($\xi = 0$) to commitment ($\xi \rightarrow 1$). Promises last $1/(1 - \xi)$ periods on average.

Private agents know this, so their expectations average over the two branches. The Euler equation of the committed plan becomes

$$A(x_t^c) = \delta_t \beta \mathbb{E}_t \left[\xi B(x_{t+1}^c) + (1 - \xi) B(x_{t+1}^r) \right],$$

and similarly for the Phillips curve, where x^r is what a re-optimizing central bank will do. The same regrouping trick as before moves the ξ part to the next period of the committed branch. The $(1 - \xi)$ part stays where it is, because the central bank cannot promise anything about a plan that will have been replaced. The result is the loose-commitment saddle-point problem:

$$\begin{aligned} V^c(\delta, \mu) = \min_{\lambda} \max_x \{ & u(x) + \lambda^R(R - 1) \\ & + \lambda^E [A(x) - (1 - \xi) \delta \beta \mathbb{E} B(x^r)] - \mu^E B(x) \\ & + \lambda^P [C(x) - (1 - \xi) \delta \beta \mathbb{E} D(x^r)] - \mu^P D(x) \} \\ & + \delta \beta \mathbb{E} [\xi V^c(\delta', \lambda) + (1 - \xi) V^r(\delta')]. \end{aligned}$$

The reneging branch is not free to choose. Whoever re-optimizes tomorrow faces exactly the same problem with a clean slate, so equilibrium requires a fixed point:

$$V^r(\delta) = V^c(\delta, 0), \quad x^r(\delta) = x^c(\delta, 0).$$

A loose-commitment equilibrium is a set of functions V^c , V^r , x^c and x^r that solve the problem above and satisfy this fixed point. Setting $\xi = 1$ recovers full commitment; setting $\xi = 0$ removes every promise and recovers the discretionary problem.

i Why there is no steady state

Under full commitment the promise state starts at zero and drifts to its long-run level. Under loose commitment it is reset to zero every time the coin comes up *renege*, so the economy keeps falling back to the origin and climbing away from it. Even without shocks it never settles. Its long-run distribution is a mixture of commitment spells of different lengths, with spell length geometric in ξ , and it changes with ξ itself.

Measuring welfare

Welfare is reported as a consumption equivalent: the percentage Λ of consumption that would make the household indifferent between the economy with the lower bound and the efficient allocation (c^*, n^*) derived above,

$$u(c_{zlb}(1 + \Lambda), n_{zlb}) = u(c^*, n^*),$$

and its unconditional version Λ^u , which takes expectations over the model's long-run distribution on both sides.

Solving it without linearizing

The usual shortcut, a linear-quadratic approximation around a steady state, fails here for three reasons:

- **The lower bound is a kink.** Policy functions bend sharply where the bound starts to bind.
- **There is no single point to approximate around.** Average inflation moves from π_f under commitment to π_f plus the inflation bias under discretion, so any approximation point is wrong for most values of ξ . And under loose commitment there is no steady state at all.
- **The results are about the shape of things.** The main result is that the probability of reaching the bound is hump-shaped in ξ . That depends on how the mean *and* the volatility of inflation shift along the commitment line, which is exactly what a linearization around one point throws away.

So the model is solved globally, with **time iteration** (Coleman (1991)): guess policy functions on a grid of states, solve the first-order conditions at each grid point given the guess for tomorrow, update, and repeat until nothing changes.

1. At each grid point, compute the expectations in the Euler equation and the Phillips curve. Tomorrow's promise state is λ with probability ξ and 0 with probability $1 - \xi$; the shock is integrated with Gauss–Hermite quadrature.
2. Solve the first-order conditions assuming the bound does not bind.
3. If the implied R is below one, re-solve with $R = 1$.
4. Update the policy guess with some dampening, and stop when successive guesses agree to

within a tolerance.

Discretion is the easy case: the only state is the exogenous shock, whose distribution is known, so the grid is simple. The only refinement is to find the point where the bound starts to bind by bisection and place extra nodes around it.

Loose commitment is the hard case. The state has three dimensions (the shock and two promises), and the relevant part of the promise space is unknown before solving the model, and different for every ξ . A dense rectangular grid is both slow and unstable, because it includes promise levels the economy never reaches. The solution follows Maliar and Maliar (2015): let the model tell you where it lives.

1. Solve on a coarse grid.
2. Use that solution to approximate the long-run distribution of the promise state. Every spell starts at $\mu = 0$; mass is pushed forward along the policy functions, and spells survive each period with probability ξ .
3. Keep only the grid nodes with non-negligible probability, which gives an irregular grid concentrated where the economy actually spends its time.
4. Re-solve on the new grid, interpolating between the scattered nodes with a Delaunay triangulation, and repeat until the grid stops changing.

Where this leaves us

The machinery comes down to one idea: **a promise is a state variable**. Under discretion there are no promises, and the central bank suffers from an inflation bias in normal times and a deflation bias at the lower bound. Under commitment, promises are permanent and both biases disappear. Under loose commitment, promises decay at rate $1 - \xi$, and with them the value of making them.

Because the model is solved globally, one can move ξ smoothly from zero to one, holding everything else fixed, and examine how inflation, welfare and the lower bound change along the way.

The recursive formulation of the policy problem follows the loose-commitment and recursive-contracts literature (Debortoli and Nunes (2010); Marcet and Marimon (2019)), while the treatment of the lower bound and the welfare comparison is related to Siu (2004).

This is my personal website. The views and opinions expressed here are my own and do not necessarily represent those of the Central Bank of Chile or its Board. Any errors, unfinished thoughts, questionable code, and overly complicated solutions found here are, of course, my own.

CALVO, G. A. (1983): "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, 12, 383–98.

COLEMAN, W. J. (1991): "Equilibrium in a production economy with an income tax," *Econometrica*, 59, 1091–1104.

DEBORTOLI, D., AND R. NUNES. (2010): "Fiscal policy under loose commitment," *Journal of Economic*

Theory, 145, 1005–32.

DOTSEY, M., R. G. KING, AND A. L. WOLMAN. (1999): “State-dependent pricing and the general equilibrium dynamics of money and output,” *The Quarterly Journal of Economics*, 114, 655–90.

KYDLAND, F. E., AND E. C. PRESCOTT. (1977): “Rules rather than discretion: The inconsistency of optimal plans,” *Journal of Political Economy*, 85, 473–91.

MALIAR, L., AND S. MALIAR. (2015): “Merging simulation and projection approaches to solve high-dimensional problems with an application to a new keynesian model,” *Quantitative Economics*, 6, 1–47.

MARCET, A., AND R. MARIMON. (2019): “Recursive contracts,” *Econometrica*, 87, 1589–1631.

ROTEMBERG, J. J. (1982): “Sticky prices in the united states,” *Journal of Political Economy*, 90, 1187–1211.

SIU, H. E. (2004): “Optimal fiscal and monetary policy with sticky prices,” *Journal of Monetary Economics*, 51, 575–607.

YUN, T. (2005): “Optimal monetary policy with relative price distortions,” *American Economic Review*, 95, 89–109.