

POLICY GAMES

1. THE SAN JOSE MODEL

The policy authority believes

$$u_t = \theta_0 - \theta_1 \pi_t + \varepsilon_t. \quad (1)$$

In fact, though,

$$u_t = \bar{u} - \alpha \cdot (\pi_t - E_{t-1} \pi_t) + \xi_t. \quad (2)$$

2. POLICYMAKERS' BEHAVIOR

They minimize

$$\frac{1}{2} E \left[\sum_{t=0}^{\infty} \beta^t (u_t^2 + \omega \pi_t^2) \right]. \quad (3)$$

At each t , they estimate θ_0 and θ_1 — by a method that may allow for variation over time in these parameters (the Kalman Filter). They do not control π precisely, but instead control g_t , with

$$\pi_t = g_{t-1} + v_t. \quad (4)$$

3. EQUILIBRIUM

They do not take account of their own learning pattern, but instead just optimize at each t as if their current estimates were true values that would remain constant forever. With λ the LM on the false Phillips curve and μ the Lagrange multiplier on the equation connecting g to π , the FOC's for this problem are

$$\begin{aligned} \partial u : & & u_t &= \lambda_t \\ \partial \pi : & & \omega \pi_t &= \theta_1 \lambda_t + \mu_t \\ \partial g : & & E_t \mu_{t+1} &= 0. \end{aligned}$$

4. EQUILIBRIUM, CONT.

These, together with the Phillips curve, yield policy as

$$g_t = \frac{\theta_1 \theta_0}{\omega + \theta_1^2}. \quad (5)$$

Substituting this expression into the true Phillips curve (2) and matching coefficients tells us that OLS applied to data from this situation and to the false model (1) would deliver

$$\begin{aligned}\theta_1 &= \alpha \\ \theta_0 &= \bar{u} + \alpha \frac{\theta_1 \theta_0}{\omega + \theta_1^2}.\end{aligned}$$

5. DYNAMICS

Solving these equations for θ_0 and θ_1 tells us the equilibrium position of the false Phillips Curve. It implies that in steady state $g_t \equiv \alpha \bar{u} / \omega$. As we would expect, equilibrium inflation is higher the greater the natural rate, the greater the apparent effect of inflation on unemployment in the Phillips Curve, and the smaller the weight on inflation in the objective function.

But how do we get there? The theory so far only considers what estimation will deliver if OLS is used and g is held constant. But to progress from low inflation to the equilibrium, the policy authority will have to change g . When it does so, it will generate data in which π is changing without producing any effect on unemployment. This will make the α appear smaller, and reduce the apparent gains to inflation. So progress to the Kydland-Prescott equilibrium is slow.

6. FIGURES

The figures that follow are from (Sims, 1988). They are not those that appeared in the original article, but replacements that appear in the web version. The models and discussion of how the charts were generated are described in the web version of the paper. All the figures show simulated time series from economies in which there is a natural rate Phillips curve and the Kydland-Prescott equilibrium level of inflation is 6%. Figures 1 and 2 illustrate the fact that such simulations can produce very different results depending on the first few observations. Figure 3 shows a typical simulation with time variation modeled by the policy authorities as equal on constant and slope, starting from low inflation. Over the 1000 year span of the graph (assuming annual data is used in the regression updates), inflation stays permanently low, never moving toward the Kydland-Prescott equilibrium. Figure 4, in contrast shows the economy near the Kydland-Prescott equilibrium most of the time, with only “brief” (100 year or so) deviations from it. This is the typical outcome when the policy authority attributes most time variation to shifts in the constant term — i.e. the “natural rate”. Figure 5 shows what happens when the same beliefs on the part of the policy authority as in Figure 3 prevail, but the economy starts near the K-P equilibrium. This figure may be misleading, in that it shows a break away from the KP equilibrium after a few hundred years, while the model was actually run for 2000 years before the start of the chart that

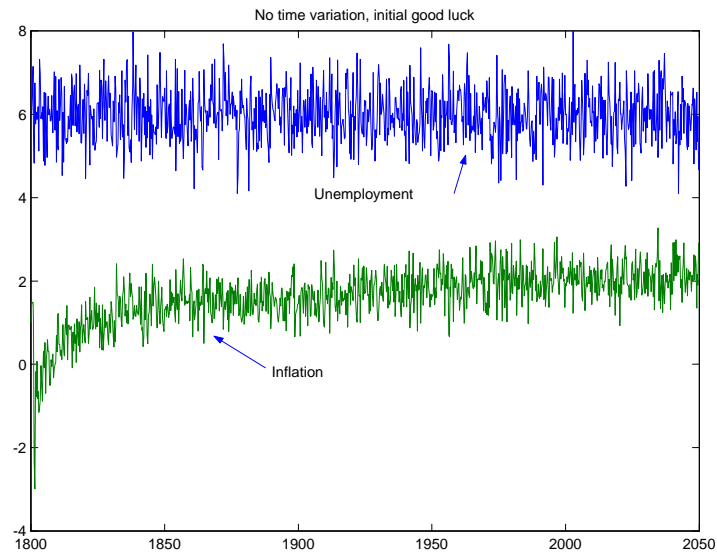


FIGURE 1




FIGURE 2

is displayed. The simulation was also continued for over 10,000 years after the end of the period displayed, and never returned to the neighborhood of KP equilibrium in that span.

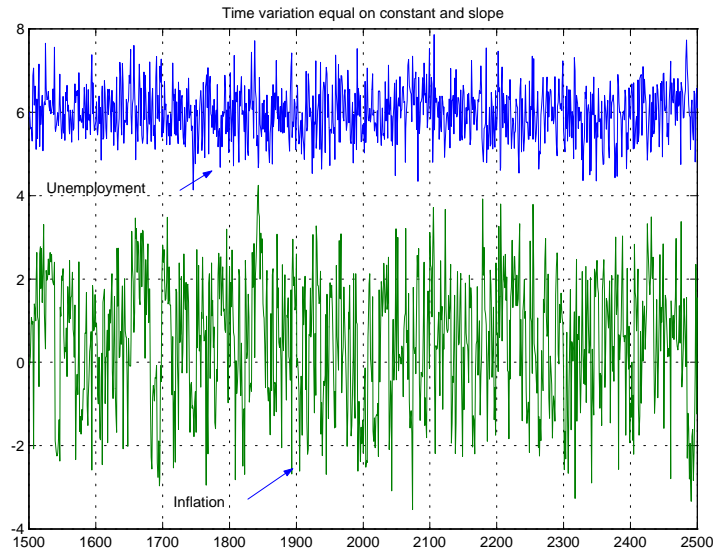


FIGURE 3




FIGURE 4

7. "MEAN" AND "ESCAPE" DYNAMICS, AND THEIR INTERPRETATION

- Sargent describes the dynamics as a "mean dynamics" drawing the economy toward Kydland-Prescott equilibrium, occasionally "punctuated" by episodes of "escape dynamics".
- He does not use the Kalman filter, and our charts show that this affects his conclusions. The point that "escape dynamics" prevail over brief periods in

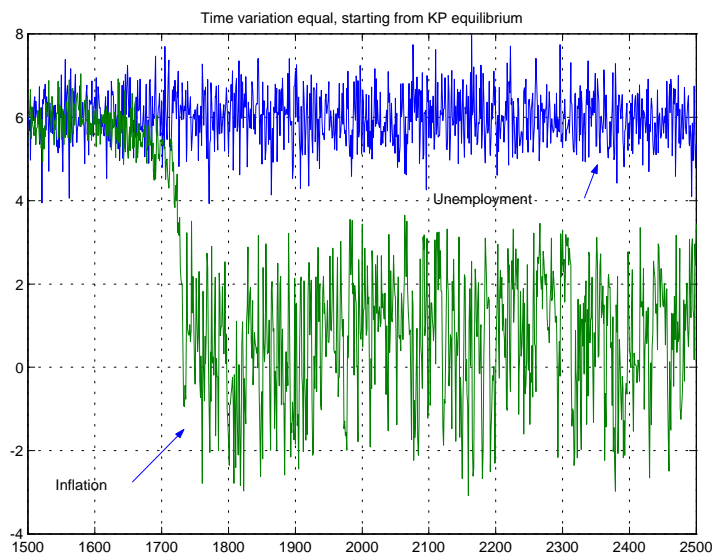


FIGURE 5

which the nature of the process changes radically is correct, but there is no necessary tendency to drift toward K-P equilibrium.

- Sargent's is one of several competing stories that explain why reliance on empirical models that do not embody the received wisdom of natural rate theory could lead to a *temporary* episode of good policy that is constantly in danger of being undermined by new, but spurious empirical results.
- An alternate view: The natural rate theory, like any other simple orthodoxy, is at best partially correct and at worst can end up an albatross weighing down any attempt to arrive at understanding of new policy challenges. (How much of Japan's problem is an effect of natural rate thinking?) Good empirical models can lead to good policy *even if they do not exactly embody the truth*.

8. FULL COMMITMENT, TIME CONSISTENT AND GAME-THEORETIC (BARRO-GORDON) APPROACHES

Full Commitment: We set the problem up as an optimization problem of the usual form, with the private sector's behavioral equations, which generally include expectational equations (Euler equations or, in our current case, the true Phillips curve (2)) among the constraints. If we maintain the assumption that the policy authority must choose g in advance, so it has no **information advantage** over the private sector, we will get the uninteresting and obvious conclusion that the optimal policy is $g_t \equiv 0$. [Why is it obvious?]. So we will examine the case where the policy authority can pick π_t directly at time

t . This implies that the policy authority can surprise the private sector, or equivalently that it has an information advantage.

No-Commitment: • The Full Commitment solution generally implies that actions taken at time 0 are different from those taken at later dates in otherwise similar conditions. This occurs because what the authority *promises* at time 0 to do at time $s > 0$ affects welfare, and because the full commitment solution requires that promises be fulfilled and believed. But if the optimization problem can be “restarted” at time $s > 0$, promises made earlier will be broken. It will be tempting to do this, particularly if the “policy authority” is actually a sequence of different office-holders. This is the “time-inconsistency of optimal plans” pointed out by Kydland and Prescott.

- A **time consistent** policy is one in which policy depends on the state of the economy at time t only, not on the date. There is no unique way to define the “state” however. What is usually called the time-consistent solution is one in which the state does not include past policy behavior, but only “exogenous states” — variables that are uninfluenced by policy behavior. This implies that nothing the policy-maker does at t can influence the behavior of policy-makers at future dates. Calling this the “no commitment” solution is better terminology.

Barro and Gordon (1983): In one section of their paper they point out that it could be that the public’s expectations of future behavior are affected by current policy behavior even if the policy authority is not believed when it makes announcements. If this were true, it would make no sense for the policy authority to ignore it. And in this case even a policy authority that cannot make commitments may find it optimal to act in the same way as an authority that can make commitments.

9. DETAILED DISCUSSION OF OUR EXAMPLE MODEL

Full Commitment: Same objective function (3), with the constraint the true Phillips curve (2). To get the problem into standard form (so expectation operators apply only to entire constraints, not individual variables), we need to define $w_t = E_t \pi_{t+1}$ and add this definitional equation to the list of constraints. With multipliers λ, μ on the Phillips curve constraint and the definitional constraint, the Euler equations are then

$$\partial \pi: \quad 2\omega \pi_t = \alpha \lambda_t + \beta^{-1} \mu_{t-1} \quad (6)$$

$$\partial w: \quad \beta \alpha E_t \lambda_{t+1} + \mu_t = 0 \quad (7)$$

$$\partial u: \quad 2u_t = \lambda_t \quad (8)$$

for $t > 0$. For $t = 0$, the μ_{t-1} term in (6) does not appear, because at the initial date the policy authority is not constrained to act in accordance with expectations as of time $t = -1$.

Then from these equations we can conclude

$$2\omega\pi_t = 2\alpha u_t - 2\alpha E_{t-1}u_t \quad (9)$$

$$\therefore \omega\pi_t = \alpha \cdot (-\alpha \cdot (\pi_t - E_{t-1}\pi_t) + \xi_t) \quad (10)$$

$$\therefore \pi_t = \frac{\alpha}{\omega + \alpha^2} \bar{\xi}_t, \quad (11)$$

where in deriving (11) we have used the Phillips curve (2) and the assumption $E_t \bar{\xi}_{t+1} = 0$.

At time $t = 0$, we instead arrive at

$$\pi_0 = \frac{\alpha}{\omega + \alpha^2} \bar{u} + \frac{\alpha^2}{\omega + \alpha^2} E_{-1}\pi_0 + \frac{\alpha \bar{\xi}_0}{\omega + \alpha^2}. \quad (12)$$

This formula implies that even when $\bar{\xi}_0 = E_{-1}\pi_0 = 0$, optimal π_0 is positive.

No commitment: If the private sector assumes that the policy authority will always act as if it is solving the full commitment problem afresh, then it will expect (12) to prevail at every date. In that case, if expectations are rational there is only one possible value for $E_{t-1}\pi_t$, which we can find by taking E_{t-1} of (12) and solving to get $E_{t-1}\pi_t = (\alpha/\omega)\bar{u}$, and this leads to

$$\pi_t = \frac{\alpha}{\omega} \bar{u} + \frac{\alpha \bar{\xi}_t}{\omega + \alpha^2}. \quad (13)$$

This is the Kydland-Prescott, no-commitment, time-consistent equilibrium policy. Though it is described here as arising from a policy authority that at each date solves the full-commitment problem *de novo*, this is true here only because the policy-maker's choices at t do not in fact influence the state at $t + 1$ under our assumptions. More generally, the no-commitment solution is one in which in which the policy authority chooses its action optimally as a function of the state, recognizing that future policy authorities' behavior will be functions of future states. Of course in equilibrium, authorities at all dates choose the same function of the state, even though these choices are made date by date, with each date's authority assuming that its own choice of policy rule has no effect on choices at other dates.

Barro-Gordon: It is plausible that the public does not perfectly understand what the policy authority is doing (even that the policy authority does not perfectly understand what it is doing itself). The public therefore might "model" policy behavior, projecting future policy actions on the basis of observed history, ignoring the plans and announcements of the policy authority. In that case, it might be that $\hat{E}_t \pi_{t+1} = f(\pi_{t-s}, u_{t-s}, s \geq 0)$. If so, we can

substitute f for the $\hat{E}_{t-1}\pi_t$ in (2), and the policy authority's problem becomes a standard dynamic optimization. The result is what is known as a **self-confirming equilibrium** if the expectation function f turns out, when the policy authority optimizes, to deliver accurate forecasts. There are in general many such equilibria. Barro and Gordon pointed to one that can produce zero inflation:

$$f(\pi_t) = \begin{cases} 0 & \pi_t = 0 \\ \frac{\alpha}{\omega}\bar{u} & \pi_t \neq 0. \end{cases} \quad (14)$$

If lagged inflation was non-zero, expectations will be consistent with the no-commitment equilibrium. In this case, if the policy authority in fact chooses nonzero inflation again, it will choose π according to the myopic solution (13) because the policy authority can do no better than this KP solution this period, and all non-zero values of π have the same implications for future expected inflation. In this case the current period losses will be

$$L_{11} = \left(\bar{u} + \frac{\omega\xi_t}{\omega + \alpha^2} \right)^2 + \omega \left(\frac{\alpha\bar{u}}{\omega} + \frac{\alpha\xi_t}{\omega + \alpha^2} \right)^2.$$

However, the policy maker might be tempted to lower inflation to zero, to get the benefits of low future inflation. Since this would be a surprise deflation, it would raise unemployment in the current period, creating total current losses of

$$L_{10} = \left(\bar{u} + \alpha \left(\frac{\alpha\bar{u}}{\omega} \right) + \xi_t \right)^2.$$

If lagged inflation was zero, expectations will be that zero inflation will persist. If the policy authority does in fact persist with zero inflation, current period losses will be

$$L_{00} = (\bar{u} + \xi_t)^2.$$

If instead the authority unexpectedly creates non-zero inflation, it will set inflation at the level implied by (12), the time-zero policy under commitment. It will do so because, as in the commitment solution at time zero, it sees no connection between its choice of π and future expectations of π . It will therefore generate current period losses of

$$L_{01} = \frac{\omega(\bar{u} + \xi_t)^2}{\omega + \alpha^2}.$$

We would like now to verify that there can be an equilibrium in which private agents have these beliefs and policy makers stay forever with either the no-commitment solution or the $\pi_t \equiv 0$ solution. Clearly if policy-makers behave this way, the private sector's forecasting rule is accurate: inflation

does in fact stay at zero if it was zero in the past and does in fact have an expected value of $\bar{u}\alpha/\omega$ if it was non-zero in the past. What remains to be checked is that a policy maker who understands the private forecasting rule and believes that future policy makers (or future incarnations of himself) will stick with one of these two policies, himself has no incentive to make an unexpected change in policy.

The conditions that guarantee this result are

$$L_{00} + \frac{\beta}{1-\beta}E[L_{00}] < L_{01} + \frac{\beta}{1-\beta}E[L_{11}] \quad (15)$$

$$L_{11} + \frac{\beta}{1-\beta}E[L_{11}] < L_{10} + \frac{\beta}{1-\beta}E[L_{00}]. \quad (16)$$

These conditions will not be met automatically. Whether they are depends on parameter values, and since they have to be met for every possible value of the shock ξ_t , we have to have bounds available on ξ_t also in order to make the equilibrium viable. For example, with $\bar{u} = 5$, $\sigma_\xi^2 = 1$, $\alpha = 1 = \omega$ and $\beta = .9$, the equilibrium is *not* viable, because it will appear optimal for the policy maker to revert permanently to $\pi_t = 0$ when lagged inflation is non-zero. Note that this does not imply that an f which predicts $\pi_t = 0$ in every state produces an equilibrium, because with these beliefs it would be optimal for policy makers to create surprise inflation.

An equilibrium is possible, indeed likely with plausible bounds on ξ_t , if ω is as small as 0.1, with the other parameters set as above. Note the somewhat paradoxical result here: for the public to be convinced that a deviation from $\pi = 0$ will persist, it must believe that the policy authority puts such heavy weight on unemployment that it would not be willing to pay the price in unemployment to get back to zero inflation once it deviated. This gives the policy authority enough of a credibility problem to make the consequences of deviation severe enough to sustain equilibrium.

10. DISCUSSION

“weak” results?: There are so many self-confirming equilibria, and they can be so different, that this fact is sometimes taken as suggesting that the theory of reputation-based equilibria is “weak”, or “uninformative. But a better interpretation is that the widely studied no-commitment equilibria are on very shaky grounds, theoretically.

Adaptive expectations equilibria: Suppose the public uses a rule like

$$E_{t-1}\pi_t = \hat{\pi}_t = \rho\hat{\pi}_{t-1} + (1-\rho)\pi_{t-1}.$$

These expectations are not “rational”, but they do guarantee that in a steady state with constant inflation $\hat{\pi}_t \equiv \bar{\pi}$. If government follows a policy that converges to a steady state inflation rate, these irrational expectations will not be provably incorrect based on regression estimates. This setup results in relatively easy to solve dynamic models. With $\rho = 0$, for example, $\pi_t \rightarrow (1 - \beta)\omega\bar{u}/\alpha$, which for $\beta \cong 1$ is much smaller than the no-commitment value of $\omega\bar{u}/\alpha$.

(Arifovic and Sargent, 2003): An experimental study, suggesting that it is likely that the public adapts its forecasts of π to historical data rather quickly, and that policy authorities might recognize this.

Sargent-Williams-Zha, Primiceri: Sargent, Williams, and Zha (2006), Primiceri (2006)

REFERENCES

- ARIFOVIC, J., AND T. J. SARGENT (2003): “Laboratory Experiments with an Expectational Phillips Curve,” in *Evolution and Procedures in Central Banking*, ed. by D. E. Altig, and B. D. Smith, pp. 23–55. Cambridge University Press.
- BARRO, R. J., AND D. B. GORDON (1983): “A Positive Theory of Monetary Policy in a Natural Rate Model,” *Journal of Political Economy*, 91(4), 589–610.
- PRIMICERI, G. (2006): “Why Inflation Rose and Fell: Policymakers’ Beliefs and US Postwar Stabilization Policy,” *Quarterly Journal of Economics*, 121, 867–901.
- SARGENT, T. J., N. WILLIAMS, AND T. ZHA (2006): “Shocks and Government Beliefs: The Rise and Fall of American Inflation,” *American Economic Review*, 96(4), 1193–1224, NBER Working Paper w10764.
- SIMS, C. A. (1988): “Projecting Policy Effects with Statistical Models,” *Revista de Analisis Economico*, 3(2), 3–20, www.princeton.edu/~sims.