Available online at www.sciencedirect.com
SCIENCE DIRECT®

Review of Economic Dynamics 8 (2005) 528–563

Review of Economic Dynamics

www.elsevier.com/locate/red

The conquest of US inflation: Learning and robustness to model uncertainty

Timothy Cogley [a], Thomas J. Sargent [b,c,*]

[a] *University of California, Davis, CA 95616, USA*
[b] *New York University, New York, NY 10003, USA*
[c] *Hoover Institution, Stanford, CA 94305, USA*

Received 1 January 2005

Available online 26 February 2005

Abstract

Previous studies have interpreted the rise and fall of US inflation after World War II in terms of the Fed's changing views about the natural rate hypothesis but have left an important question unanswered. Why was the Fed so slow to implement the low-inflation policy recommended by a natural rate model even after economists had developed statistical evidence strongly in its favor? Our answer features model uncertainty. Each period a central bank sets the systematic part of the inflation rate in light of updated probabilities that it assigns to three competing models of the Phillips curve. Cautious behavior induced by model uncertainty can explain why the central bank presided over the inflation of the 1970s even after the data had convinced it to place much the highest probability on the natural rate model.

Keywords: Natural unemployement rate; Phillips curve; Bayes' law; Anticipated utility; Robustness

* Correspondence to: Thomas J. Sargent, Department of Economics, New York University, 269 Mercer Street, 7th Floor, New York, NY 10003, USA.

E-mail address: Thomas.Sargent@nyu.edu (T.J. Sargent).

1094-2025/$ – see front matter
doi:10.1016/j.red.2005.02.001

1. Introduction

This paper uses a model of an adaptive monetary authority to interpret the rise and fall of US inflation in the 1960s, 1970s, and 1980s.[1] Following DeLong (1997), Taylor (1997), and Sargent (1999), we explore whether the rise and fall of US inflation can be attributed to policy makers' changing beliefs about the natural rate hypothesis. One story that emphasizes changing beliefs goes as follows. Samuelson and Solow (1960) taught policy makers that there was an exploitable long-run tradeoff between inflation and unemployment, and inflation rose as the authorities tried to exploit the tradeoff. But that misguided policy experiment generated observations that taught the authorities the natural rate hypothesis, which eventually convinced them to reduce inflation.

The adverb "eventually" significantly qualifies this story because the data indicate that the authorities should have learned the natural rate model long before they acted on it. Sims (2001) and Cogley and Sargent (2001) demonstrate that the natural rate hypothesis should have been learned by the early 1970s, yet inflation remained high until the early 1980s.[2] If the rise and fall of US inflation reflected only changing beliefs about the natural rate hypothesis, and not, say, altered purposes or decision making arrangements, then it is puzzling that inflation remained high for a decade after substantial statistical evidence favoring the natural rate hypothesis had accumulated. By the early 1970s, average inflation was on the rise, yet average unemployment had not fallen, contrary to the Samuelson–Solow model. The events of the early 1970s turned the economics profession away from the Samuelson–Solow model. Why did policy makers wait a decade to act on this lesson?

A number of alternative explanations for the high US inflation of the 1970s refrain from assigning an important role to policy makers' changing beliefs about the natural rate hypothesis. For example, DeLong (1997) questions the motives of Arthur Burns. Parkin (1993) and Ireland (1999) say that discretionary policy making combined with a higher natural rate of unemployment resulted in a higher inflationary bias. Chari et al. (1998) and Albanesi et al. (2003) ascribe the high inflation to an expectations trap. Orphanides (2003), Lansing (1999), and Romer and Romer (2002) emphasize that the Federal Reserve was slow to detect the productivity slowdown and the rise in the natural rate of unemployment.[3] Primiceri (2003) emphasizes evolution in the Fed's estimates not only of the natural rate but also of the slope of the short-term tradeoff between inflation and unemployment.[4] We believe that there is some truth in all of these ideas but explore a different explanation. We show that concerns about the robustness of a proposed inflation-stabilization policy across

[1] The model is what David Kreps (1998) called an anticipated utility model. In an anticipated utility model, a decision maker recurrently maximizes the expected utility of a stream of future outcomes with respect to a model that is recurrently re-estimated. Although an anticipated utility agent optimizes and learns, he does not purposefully experiment. But as data accrue, he adapts and possibly even respecifies the model that he uses to evaluate expected utility.

[2] Sims credits Albert Ando for first making this point.

[3] See Christiano and Fitzgerald (2003) and Velde (2004) for critical surveys of these and other theories of the Great Inflation. See John Taylor (2002) for an account that emphasizes the Fed's learning about theories of the natural unemployment rate.

[4] Erceg and Levin (2003) investigate how the public's learning about the Fed's motives can explain the recession that accompanied the reduction of inflation that the Fed engineered under Paul Volcker.

alternative models would have induced policy makers to choose high inflation even though the data favored a model that recommended low inflation.[5]

Our calculations confirm that by the mid 1970s, zero inflation would have been optimal according to the model that was most probable among the ones that we consider. But because the optimal policy takes model uncertainty into account via Bayesian model averaging, that was not enough to convince our Bayesian policy maker to abstain from trying to exploit the Phillips curve. In two competing models that had smaller but still non-zero probability weights, a policy of quickly reducing inflation would have been calamitous. If very bad results are associated with a particular policy according to *any* of the models that retain a positive but small posterior probability, our Bayesian policy maker refrains from that policy. When outcomes from following a recommendation to stabilize inflation become worse under some particular worst-case model, the ultimate decisions tend to track the recommendations of that worst-case model. In this way, our statistical model rationalizes the idea that the high inflation of the 1970s reflected the Fed's desire to guard against the bad outcomes that would have come from a low-inflation policy under models that by the 1970s should have had low posterior probabilities.

2. Learning and policy making with multiple approximating models

2.1. Preliminary adaptive setup with a single approximating model

We model the US monetary authority as an adaptive decision maker like ones described by Kreps (1998) and Sargent (1999) in the context of models of the Phillips curve like ones in Sargent (1999), augmented with features that introduce a concern for robustness to model uncertainty. The models in Sargent's monograph work as follows. In truth, but unbeknownst to the central bank, a natural rate version of a Phillips curve relates unemployment to surprise inflation,

$$u_t - u_t^* = -\theta(y_t - x_t) + \rho(L)u_{t-1} + \eta_t, \tag{1}$$

where u_t is unemployment, u_t^* is the natural rate, y_t is actual inflation, x_t is both the systematic part of inflation and the rate of inflation expected by the public, and η_t is an i.i.d. normal shock with mean zero and variance σ_η^2. The central bank sets x_t with a decision rule that is described below; x_t is related to actual inflation according to

$$y_t = x_t + \xi_t, \tag{2}$$

where ξ_t is an i.i.d. normal shock with mean zero and variance σ_ξ^2.

The central bank does not know Eq. (1), and instead bases its decisions on a statistical approximating model,

$$Y_t = X_t'\theta + \nu_t, \tag{3}$$

[5] Blinder (1998, pp. 12–13) recounts how he used multiple models to evaluate alternative policies when he was Vice Chairman of the Federal Reserve Board.

where X_t and Y_t represent generic right- and left-hand variables in a regression, and θ is a vector of regression parameters.

The central bank starts date t with an estimated approximating model carried over from date $t-1$. It forms a decision rule for x_t by minimizing a discounted quadratic loss function,[6]

$$\mathcal{L} = E_t \sum_{j=0}^{\infty} \beta^j \left(u_{t+j}^2 + \lambda y_{t+j}^2\right), \tag{4}$$

subject to a constraint implied by its time $t-1$ estimates of model (3). This induces a best-response mapping from the loss function and the time $t-1$ estimates of the parameters of its approximating model to those of a time t policy rule, the first period outcome of which is a time t policy action x_t. To emphasize its dependence on predetermined state variables and parameter estimates, we denote it as $x_{t\,|\,t-1}$. The policy action $x_t = x_{t\,|\,t-1}$ influences outcomes through Eq. (1). After observing outcomes Y_t, X_t at t, the central bank re-estimates the parameters of its policy model (Eq. (3)), preparing to repeat the same decision process in the next period.

As the central bank's beliefs (i.e., the parameters of its approximating model) evolve with the accumulation of data, so too does its policy rule.

2.2. Our three-model model

In contrast to Sargent (1999), we take no stand on the true data generating process. We drop Sargent's (1999) assumption that the central bank has a unique approximating model and instead assume that the central bank acknowledges multiple models.[7] Model uncertainty affects the bank's deliberations and induces a Bayesian form of robustness of decision rules at least across the domain of alternative models that are on the table. Our central bank adopts a 'blue-collar' version of the approach to policy making presented in Brock et al. (2003).

We choose three approximating models to represent Phillips curve specifications that have been influential since the mid to late 1960s. The first model allows a permanently exploitable Phillips curve, the second a temporarily exploitable one, and the third a statistical Phillips curve that is not exploitable even temporarily. The first statistical model is inspired by Samuelson and Solow (1960),

$$y_t = \gamma_0 + \gamma_1(L)y_{t-1} + \gamma_2(L)u_t + \eta_{1t}, \tag{5}$$

where η_{1t} is i.i.d. $N(0, \sigma_1^2)$. Its key feature is that it permits a long-run tradeoff unless the parameter values take special configurations. We assume that the central bank considers an unrestricted parameter configuration in order to allow the possibility of a long-run tradeoff.

[6] Our linear-quadratic framework abstracts from concerns about the zero bound on nominal interest rates. That could make preferences asymmetric with respect to inflation and deflation and cause policy makers to set an inflation target above zero.

[7] The sense in which the central bank has 'a model' is as a weighted average across several models.

The bank also makes the identifying assumption that η_{1t} is orthogonal to the right side variables, including current unemployment.[8] This makes Eq. (5) a regression.

Model 2 is a restricted form of model 1, inspired by Solow's (1968) and Tobin's (1968) suggestion about how to represent the natural rate hypothesis:

$$\Delta y_t = \delta_1(L)\Delta y_{t-1} + \delta_2(L)\big(u_t - u_t^*\big) + \eta_{2t}. \tag{6}$$

Like the Samuelson–Solow model, this features an exploitable short-run tradeoff between inflation and unemployment, but it recognizes a distinction between actual unemployment and the natural rate, and it imposes Solow and Tobin's version of long-run neutrality. The restriction that the intercept is zero and that the sum of the lag weights on y_t equals one enforces that the long-run Phillips curve is vertical, located at u_t^* if the roots of $\delta_2(L)$ are outside the unit circle. We assume that the authorities make the identifying assumption that η_{2t} is uncorrelated with current unemployment, so that Eq. (6) is also a regression.

Our third model is inspired by Lucas (1972) and Sargent (1973), and it enforces both long-run neutrality and a policy ineffectiveness proposition,

$$u_t - u_t^* = \phi_1(y_t - x_{t\,|\,t-1}) + \phi_2(L)\big(u_{t-1} - u_{t-1}^*\big) + \eta_{3t}. \tag{7}$$

This model says that only unexpected inflation matters for unemployment. When updating estimates of this model, the latent variable $x_{t\,|\,t-1}$ is measured recursively as the optimal policy of a Bayesian linear regulator who solves a discounted quadratic control problem, taking all three approximating models into account. We describe how the bank chooses $x_{t\,|\,t-1}$ in detail below.

This model also reverses the direction of fit, putting unemployment on the left side and inflation on the right. When entertaining this representation, the central bank makes the identifying assumption that η_{3t} is orthogonal to current inflation. This identifying assumption differs from those of the other two models, but what matters for statistical updating is the internal logic of each. Thus, we can proceed with regression updates for this model as well.

As we demonstrate later, the assumption that the Keynesian models are estimated with a Keynesian direction of fit is important for our story because it influences estimates of sacrifice ratios. Models estimated in the Keynesian direction imply a high sacrifice ratio in the 1970s, while those estimated in a classical direction predict a low sacrifice ratio. We could incorporate uncertainty about the direction of fit by expanding our three-model model to include versions of the Samuelson–Solow and Solow–Tobin models estimated with a classical direction of fit. But this seems revisionist, and in any case does not alter the results. Policy choices for the five-model model are essentially the same as those for the three-model model.[9] Accordingly, we focus on the more parsimonious three-model representation.

[8] Sargent (1999) had more success tracking actual inflation with this 'Keynesian' direction of fit. In contrast, a 'classical' direction of fit would put unemployment on the left-hand side, as in Eq. (1). See King and Watson (1994) and Sargent (1999) for background on how the direction of fit matters for the properties of Phillips curves. By adopting the model of government learning of Sargent and Williams (2003), Sargent et al. (2004) obtain much better econometric explanations of the US inflation path than were obtained by Sargent (1999).

[9] This happens because the additional models do not alter the identity of the 'worst-case' model. The significance of this fact will become clear later on.

Associated with each model and date is a posterior probability, α_{it}, $i = 1, 2, 3$. We assume the authorities entertain no other possibilities, so $\alpha_{1t} + \alpha_{2t} + \alpha_{3t} = 1$ for all t. One obvious shortcoming of this approach is the assumption of an exhaustive list of explicitly specified possible models. This has led Hansen and Sargent (2001, 2002) and others to specify a single approximating model, to surround it with an uncountable cloud of alternative models each of whose entropy relative to the approximating model is bounded, and to use a minimax decision rule because the decision maker declines to put a unique prior over that huge set of models. Relative to Hansen and Sargent's robust decision maker and his continuum of vaguely specified models, our three-model decision maker knows much more. He can construct a posterior over his three models, then make decisions by an appropriate form of model averaging. Nevertheless, a minimax flavor emerges from our Bayesian calculations with three models.

The central bank's decision process is similar to the one in Sargent's (1999) model that we described above, but with an extra step. At each date t, the central bank implements the action $x_{t\,|\,t-1}$ recommended by last period's policy rule. Next, it gets new data y_t, u_t and updates the estimates of each of its three models according to Bayes's theorem. It also recalculates the model weights α_{it} by evaluating the marginal likelihood associated with each model. Then it re-optimizes in light of its revised view of the world, preparing a contingency plan $x_{t+1\,|\,t}$ for the next period.[10]

Notice that with this timing protocol, the Lucas–Sargent model always recommends $x_{t\,|\,t-1} = 0$. There is no inflation bias because the central bank moves first, setting policy for date t based on information available at $t - 1$, and there is no reason to vary $x_{t\,|\,t-1}$ to stabilize u_t because systematic policy is neutral. Because variation in y_t is costly and there are no offsetting benefits in terms of reduced variation in u_t, the optimal policy within this model is $x^{LS}_{t\,|\,t-1} = 0$.[11]

[10] The observational equivalence of natural and unnatural rate models pointed out by Sargent (1976) pertains to our calculations in an important way. Using the Wold decomposition theorem, Sargent demonstrated that under a time-invariant decision rule for the systematic part of inflation, models that make unemployment depend on a distributed lag of inflation are observationally equivalent with ones that make unemployment depend on a distributed lag of *surprises* in inflation. That result seems to suggest that our three models should be difficult to distinguish, but in fact there are two sources of information that inform the posterior probabilities α_{it}. One involves a distinction between variation within and across monetary regimes. Although our models are observationally equivalent within a monetary regime, they are not equivalent across regimes, because the Lucas–Sargent model says that the equilibrium law of motion for unemployment is invariant across policy rules while the Keynesian models say it is not. In the calculations reported below, $x_{t\,|\,t-1}$ is formed not from a time-invariant decision rule but from one that changes from period to period. The weight on the Lucas–Sargent model rises when new data support the invariance property, and it falls when they do not. The second force for change in α_{it} is parsimony. The Solow–Tobin model is nested within the Samuelson–Solow model but has fewer parameters, and our calculations credit that parsimony. If new data suggest the Solow–Tobin model fits about as well, then its posterior weight will rise relative to the Samuelson–Solow model. During an age before the advent of non-parametric (i.e., infinite-parametric) models, Arthur Goldberger is reputed often to have warned against proliferating free parameters. Robert Lucas (1981, p. 188) said that "... it is useful, in a general way, to be hostile toward theorists bearing free parameters." Our Bayesian calculations are hostile toward additional free parameters.

[11] See Stokey (1989) for a discussion of time inconsistency problems in terms of alternative timing protocols. See Sargent (1999, Chapter 3) for an application of Stokey's analysis to the Phillips curve.

2.3. The central bank's decision making process

Now we turn to the details of the central bank's decision process. The first of the bank's three tasks is to update parameter estimates for each of its approximating models. Within each model, the central bank's identifying assumptions make Eqs. (5)–(7) regressions. In each case, we assume the bank adopts a normal-inverse gamma prior,

$$p(\theta, \sigma^2) = p(\theta \mid \sigma^2) p(\sigma^2), \tag{8}$$

where σ^2 is the variance of Phillips curve residuals. The marginal prior $p(\sigma^2)$ makes the error variance an inverse gamma variate, and the conditional prior $p(\theta \mid \sigma^2)$ makes the regression parameters a normal random vector. Along with this, we assume that the Phillips curve residuals η_{it} are identically and independently distributed and conditionally normal given the right-hand variables. These assumptions make the conditional likelihood function normal. With a normal-inverse gamma prior and a Gaussian conditional likelihood, the posterior also belongs to the normal-inverse gamma family, and its parameters can be updated recursively.

In particular, let Z^t summarize the joint history of (X_t, Y_t) up to date t. Before seeing data at t, the central bank's prior is

$$\begin{aligned} p(\theta \mid \sigma^2, Z^{t-1}) &= N(\theta_{t-1}, \sigma^2 P_{t-1}^{-1}), \\ p(\sigma^2 \mid Z^{t-1}) &= IG(s_{t-1}, v_{t-1}), \end{aligned} \tag{9}$$

where θ_{t-1}, P_{t-1}, s_{t-1}, and v_{t-1} represent estimates based on data through period $t-1$. The variable P_{t-1} is a precision matrix, s_{t-1} is a scale parameter for the inverse-gamma density, and v_{t-1} counts degrees of freedom. The estimate of σ^2 is just s_{t-1}/v_{t-1}. After seeing outcomes at t, the central bank's updated beliefs are

$$\begin{aligned} p(\theta \mid \sigma^2, Z^{t}) &= N(\theta_{t}, \sigma^2 P_{t}^{-1}), \\ p(\sigma^2 \mid Z^{t}) &= IG(s_{t}, v_{t}), \end{aligned} \tag{10}$$

where

$$\begin{aligned} P_t &= P_{t-1} + X_t X_t', \\ \theta_t &= P_t^{-1}(P_{t-1}\theta_{t-1} + X_t Y_t), \\ s_t &= s_{t-1} + Y_t' Y_t + \theta_{t-1}' P_{t-1} \theta_{t-1} - \theta_t' P_t \theta_t, \\ v_t &= v_{t-1} + 1. \end{aligned} \tag{11}$$

The posterior for date t becomes the prior for date $t+1$.

The second task the central bank performs each period is to revise the model probability weights. For the normal-inverse gamma family, this can also be done recursively. Let $\alpha_{i0} = p(M_i)$ represent the prior probability on model i. The posterior weight on model i is defined as

$$\alpha_{it} = \frac{w_{it}}{w_{1t} + w_{2t} + w_{3t}}, \tag{12}$$

where w_{it} is an unnormalized model weight. In an unpublished appendix available on the RED web site, we show that Bayes's theorem implies the following recursion for w_{it},

$$\log w_{it+1} = \log w_{it} + \log p\big(Y_{it+1} \mid X_{it+1}, \theta_i, \sigma_i^2\big) - \log \frac{p(\theta_i, \sigma_i^2 \mid Z_i^{t+1})}{p(\theta_i, \sigma_i^2 \mid Z_i^t)}. \tag{13}$$

The term $\log p(Y_{it+1} \mid X_{it+1}, \theta_i, \sigma_i^2)$ is the conditional log-likelihood for observation $t+1$, and $\log p(\theta_i, \sigma_i^2 \mid Z_i^{t+1}) - \log p(\theta_i, \sigma_i^2 \mid Z_i^t)$ is the change in the log posterior that results from a new observation. These are easy to calculate for the normal-inverse gamma family.[12] The central bank uses the probability weights α_{it} in its policy deliberations.

From conditions (11) and (9), one can see that the central bank's approximating models are built for tractability. So long as the models remain within the normal-inverse gamma family, updates of parameters and model weights are quite simple. Stepping outside this family would significantly complicate the calculations because then at every date we would have to integrate numerically a high-dimensional posterior density for each model, and that would be very costly to compute. For tractability, we want to remain within a conjugate family, and the normal-inverse gamma family is a natural choice for this problem.

The central bank's third task each period is to solve an optimal control problem that takes the form of a discounted stochastic linear quadratic dynamic programming problem, i.e., a so-called optimal linear regulator problem. In forming this problem, we ascribe the type of adaptive behavior contained in models of Kreps (1998) and Sargent (1999). In particular, when reformulating its policy rule each period, we assume the central bank treats the estimated parameters of its approximating models as if they were constants rather than random variables that come from a sequential estimation process. This behavioral assumption has two consequences. First, it means that decision rules depend only on point estimates and not posterior variances or higher moments, so it deactivates a concern for parameter uncertainty within each approximating model, an element of the problem emphasized by Brainard (1967). Second, it also ignores the connection between today's policy and tomorrow's information flow, a link that provides a motive for experimentation in the models of Wieland (2000a, 2000b) and Beck and Wieland (2002). Thus, our central bank adheres to the prescriptions of Blinder (1998) and Lucas (1981) that the central bank should resist the temptation to run experiments that will help it learn about the structure of the economy.[13] Nevertheless, experimentation that emerges from benevolent motives but mistaken beliefs *is* a decisive feature of our story.

To cast the central bank's problem as an optimal linear regulator, we first write the approximating models as

$$S_{it+j} = A_i(t-1)S_{it+j-1} + B_i(t-1)x_{t+j\,|\,t-1} + C_i(t-1)\eta_{it+j} \tag{14}$$

[12] Analytical expressions for these terms can be found in Appendix A.

[13] Blinder (1998, p. 11) states "while there are some fairly sophisticated techniques for dealing with parameter uncertainty in optimal control models with learning, those methods have not attracted the attention of either macroeconomists or policymakers. There is a good reason for this inattention, I think: You don't conduct policy experiments on a real economy solely to sharpen your econometric estimates." Lucas (1981, p. 288) remarks: "Social experiments on the grand scale may be instructive and admirable, but they are best admired at a distance. The idea, if the marginal social product of economics is positive, must be to gain some confidence that the component parts of the program are in some sense reliable prior to running it at the expense of our neighbors."

where $(S_{it}, A_i(t-1), B_i(t-1), C_i(t-1))$ are the state vector and system arrays for model i at time t. The system arrays are evaluated at the point estimates that emerge from Eq. (11). The details for each model are spelled out in Appendix B. One point that bears emphasis is that arriving at this specification involves inverting the Keynesian Phillips curves to express them in terms of the classical direction of fit. In other words, after estimating (5) and (6), we rearrange to put unemployment on the left-hand side and current inflation on the right. This puts the Keynesian models in a form in which it is sensible to imagine controlling u_t and y_t via settings for $x_{t\,|\,t-1}$. We emphasize that although the regressions that produce the Samuelson–Solow and the Solow–Tobin models both have y_t on the left side, the government always regards inflation as under its control via Eq. (2).[14] Thus, $x_{t\,|\,t-1}$ is the common time t instrument of the monetary authority in all three of the models. The mechanics are spelled out in Appendix B.

The loss function for model i can be expressed as

$$\mathcal{L}(\mathcal{M}_i) = E_t \sum_{j=0}^{\infty} \beta^j \left(S'_{it+j} M'_{s_i} Q M_{s_i} S_{it+j} + x'_{t+j\,|\,t-1} R x_{t+j\,|\,t-1} \right), \tag{15}$$

where $\mathcal{M}_i$ denotes model i, M_{s_i} is a selection matrix that picks out the targets (y_t, u_t) from the state vector S_{it}, Q and R are positive semidefinite weighting matrices that penalize deviations from the target and variation in the instrument, and $x_{t+j\,|\,t-1}$ is the expected value of x_{t+j} conditioned on the model estimated at time $t-1$. We set

$$Q = \begin{pmatrix} 1 & 0 \\ 0 & \lambda \end{pmatrix}, \tag{16}$$

to reflect the relative weights on unemployment and inflation, respectively. We also set R equal to 0.001, mostly for computational reasons.[15]

To put the problem in the standard form for an optimal linear regulator, we stack the model-specific constraints as

$$S_{Et+j} = A_E(t-1) S_{Et+j-1} + B_E(t-1) x_{t+j\,|\,t-1} + C_E(t-1)\eta_t, \tag{17}$$

where $S_{Et} = [S'_{1t}, S'_{2t}, S'_{3t}]'$, $\eta_t = [\eta_{1t}, \eta_{2t}, \eta_{3t}]'$, and

$$A_E(t-1) = \begin{bmatrix} A_1(t-1) & 0 & 0 \\ 0 & A_2(t-1) & 0 \\ 0 & 0 & A_3(t-1) \end{bmatrix}, \tag{18}$$

$$B_E(t-1) = \begin{bmatrix} B_1(t-1) \\ B_2(t-1) \\ B_3(t-1) \end{bmatrix},$$

[14] Again, see King and Watson (1994) for a discussion of the consequences and interpretations of alternative directions of fit in the empirical Phillips curve literature.

[15] This allows us to use a doubling algorithm to solve the optimal linear regulator problem. The doubling algorithm requires a positive definite R, and we would have to use slower algorithms if R were 0. See Anderson et al. (1995) for a discussion of the doubling algorithm and its properties *vis a vis* alternative algorithms.

$$C_E(t-1) = \begin{bmatrix} C_1(t-1) & 0 & 0 \\ 0 & C_2(t-1) & 0 \\ 0 & 0 & C_3(t-1) \end{bmatrix}.$$

The composite transition equation (17) encompasses the three submodels.

The central bank's loss function can also be written in this notation. After averaging across models, the expected loss is

$$\begin{aligned} \mathcal{L}_E &= \alpha_{1t}\mathcal{L}(\mathcal{M}_1) + \alpha_{2t}\mathcal{L}(\mathcal{M}_2) + \alpha_{3t}\mathcal{L}(\mathcal{M}_3), \\ &= E_t \sum_{j=0}^{\infty} \beta^j \left(S'_{E,t+j} Q_{Et} S_{E,t+j} + x'_{t+j\,|\,t-1} R x_{t+j\,|\,t-1}\right), \end{aligned} \tag{19}$$

where

$$Q_{Et} = \begin{bmatrix} \alpha_{1t} M'_{s_1} Q M_{s_1} & 0 & 0 \\ 0 & \alpha_{2t} M'_{s_2} Q M_{s_2} & 0 \\ 0 & 0 & \alpha_{3t} M'_{s_3} Q M_{s_3} \end{bmatrix}. \tag{20}$$

Notice how Q_{Et}, the weighting matrix in the complete model, counts each model in proportion to its posterior probability. The weights vary from period to period as the models become more or less plausible in light of new data.

The central bank chooses a decision rule for $x_{t\,|\,t-1}$ by minimizing its expected loss (Eq. (19)) subject to the constraint implied by the composite transition equation (Eq. (17)). The central bank's Bellman equation at time t is

$$v_t(S_E) = \max_x \left\{-S'_E Q_{Et} S_E - x' R x + \beta E_{t-1} v_t\left(S_E^*\right)\right\} \tag{21}$$

where the maximization is subject to (17) and the expectation E_{t-1} is with respect to the distribution of η_t in (17). This problem takes the form of an optimal linear regulator. The optimal decision rule takes the form

$$\begin{aligned} x_{t\,|\,t-1} &= -f_E(t-1) \cdot S_{Et-1} \\ &= -f_E(t-1)^1 S_{1t-1} - f_E^2(t-1) S_{2t-1} - f_E^3(t-1) S_{3t-1}. \end{aligned} \tag{22}$$

If the composite model is 'detectable' and 'stabilizable,' then the policy rule f_E can be computed using standard algorithms (e.g., see Sargent, 1980 and Anderson et al., 1995).[16]

Detectability and stabilizability also guarantee that the closed-loop matrix $(A_E - B_E f_E)$ has eigenvalues less than $\beta^{-1/2}$ in magnitude, thus ensuring that $\mathcal{L}_E$ is finite. Assuming that the posterior probability weights are all strictly positive, this means that each submodel also has finite expected loss under the optimal rule. In other words, an optimal policy simultaneously stabilizes all the submodels.

This is not necessarily the case under an arbitrary policy rule. If a policy fails to stabilize one of the submodels, so that the closed-loop matrix associated with that model has an eigenvalue greater than $\beta^{-1/2}$ in absolute value, then the expected loss is infinite, both for that submodel and for the complete model. A Bayesian linear regulator avoids such rules,

[16] Our computer programs always check these conditions.

even if the unstable submodel has a low probability weight, because a small probability weight cannot counterbalance an infinite expected loss.

From this observation there emerges a connection with the minimax approach. An unstabilized submodel is a bad outcome against which a Bayesian linear regulator wants to guard. That submodel exerts an influence on the choice of policy that is disproportionate to its probability weight. As in the minimax approach, preventing disasters, even those expected to occur with low probability, is the first priority for policy.

A Bayesian linear regulator who is permitted endlessly to proliferate submodels could become paralyzed because the encompassing model may become unstabilizable if too many exotic components are added to the mix. That is, if the composite model contains several unusual elements, there may exist no f that simultaneously stabilizes them all. One solution, proposed by Madigan and Raftery (1994), uses 'Occam's window' to exclude submodels with posterior probability below some threshold. This threshold is the analog to the multiplier parameter θ that restrains the activity of the 'evil agent' in the robust control approach described by Hansen and Sargent (2001).

The problem of paralysis does not arise in our application because we consider three conventional Phillips curve specifications. For the calculations reported below, the composite model is always detectable and stabilizable. But it could be relevant in other applications.

3. A statistical history of thought and policy

The free parameters in this model are the initial probability weights, α_{i0}, the central bank's initial priors on parameters of each approximating model, and the parameters β and λ that govern its loss function. Everything else is updated recursively, via Eqs. (11) and (13). In principle, these parameters could be estimated by GMM or MLE, but our empirical exploration consists of a calibration. We set plausible values for the free parameters, turn on the recursions with actual data for inflation and unemployment, and see what we get.

The data are standard. Inflation is measured by the log difference in the chain-weighted GDP deflator, and unemployment is the civilian unemployment rate.[17] Both series are quarterly and seasonally adjusted, and they span the period 1948.Q1 through 2002.Q4.

To set initial priors for the central bank's approximating models, we use estimates taken from the first 12 years of data, 1948–1959, with an allowance for lags at the beginning of the sample.[18] The lag order for each model is summarized in Table 1. The current value plus 4 lags of inflation enter the two Keynesian models along with the current value

[17] We also tried CPI inflation and the unemployment rate for white males aged 20 years or more, but they resulted in undetectable, unstabilizable systems, which undermines theorems guaranteeing convergence of Riccati equation iterations.

[18] The Lucas–Sargent model requires that we make an assumption about the evolution of x_t in the training sample. Within the training sample, we generate x_t from $x_t = x_{t-1} + 0.075(y_t - x_{t-1})$ with initial condition x_0 set to the initial rate of inflation y_0 in the sample. Outside the training sample, we use our model-generated $x_{t\,|\,t-1}$ as x_t in the Lucas–Sargent Phillips curve.

Table 1
Lag order in central bank approximating models

	Inflation	Unemployment
Samuelson–Solow	$\gamma_1 : 4$	$\gamma_2 : 2$
Solow–Tobin	$\delta_1 : 3$	$\delta_2 : 2$
Lucas–Sargent	$\phi_1 : 0$	$\phi_2 : 2$

plus two lags of unemployment. In the Lucas–Sargent model, unemployment is assumed to be $AR(2)$, perturbed by the current value of unexpected inflation, $y_t - x_{t\,|\,t-1}$. These choices of lag length emerged after some experimentation. They compromise parsimony and fit.[19]

The Solow–Tobin and Lucas–Sargent specifications involve $u_t - u_t^*$, the gap between actual unemployment and the natural rate. It would be best to treat u_t^* as unobservable, adding an optimal filtering step to the bank's deliberations, but doing so would substantially complicate our updating procedure. Instead, we prefer to construct an observable proxy. Thus, we measure the natural rate of unemployment by exponentially smoothing actual unemployment,

$$u_t^* = u_{t-1}^* + g\big(u_t - u_{t-1}^*\big), \tag{23}$$

with the gain parameter $g = 0.075$.[20] The natural rate series was initialized by setting $u_t^* = u_t$ in 1948.Q1. This is a rough and ready way to track movements in the natural rate.[21] In its defense, one important feature is that it is a one-sided low-pass filter that preserves the integrity of the information flow to the central bank. In contrast, a two-sided filter would allow the authorities to peek at the future of unemployment.

The parameters $\theta_{i0},\ i = 1, 2, 3,$ were set equal to the point estimates from the initial regressions, the precision matrix P_{i0} equal to the appropriate $X'X$ matrix; s_{i0} is the residual sum of squares and ν_{i0} is the degrees of freedom. Since the first 12 years of the sample are used to set priors, 1959.Q4 becomes date zero, and the recursions begin in 1960.Q1.

In addition to setting the parameters of the central bank's prior, we must also initialize the model weights, α_{i0}.[22] Because the Solow–Tobin and Lucas–Sargent models were yet to be invented, we put most of the initial weight on the Samuelson–Solow model. Thus, we set $\alpha_{10} = 0.98$ and $\alpha_{20} = \alpha_{30} = 0.01$. The results are insensitive to this choice, primarily because the data quickly come to dominate posterior probabilities.

[19] In principle, one could account for uncertainty about lag lengths by proliferating submodels, taking one for each possible lag specification, but that would result in many submodels. It would probably also multiply the number of worst-case models at any time, for when one lag specification is ill-behaved other nearby specifications are also likely to be ill-behaved. Occam's window is likely to be helpful in such cases. How best to implement that is an open question.

[20] Note that this is simply a measurement equation that we use. It is not part of the model used by the central bank. In particular, the central bank does not believe that it can manipulate u_t^*.

[21] Hall (1999) recommends a sample average as a robust estimator of the natural unemployment rate.

[22] The initial regression output cannot be used to set α_{i0} because it represents a posterior derived from a flat prior. Relative model weights are indeterminate in this case, an instance of Lindley's paradox. Thus, α_{i0} must be set a priori.

Finally, we adopt standard values for the parameters of the central bank's loss function. The discount factor, β, is set at $1.04^{-1/4}$, reflecting an annual discount rate of 4 percent. The weight on inflation, λ, is set equal to 16, reflecting an equal weight with unemployment when inflation is measured at an annual rate. The results are not sensitive to plausible changes in λ or β.[23]

These parameters initialize the recursions. On the basis of information available through 1959.Q4, our hypothetical central bank prepares a contingency plan for 1960.Q1 and sets x_1 accordingly. Then we give it data on actual inflation and unemployment for 1960.Q1. The central bank re-estimates its models in light of the new information, re-evaluates their probability weights, and revises its policy rule for $x_{t|t-1}$. The bank continues in this fashion, updating one quarter at a time, through 2002.Q4. The results of their calculations are summarized below.

Figure 1 illustrates the puzzle. The top panel shows the history of inflation, and the bottom portrays the evolution of model weights. The weight on the Lucas–Sargent model is depicted by a solid line, the weight on the Samuelson–Solow by a dashed line, and that on the Solow–Tobin model by a dashed and dotted line.

The main features of the top panel are familiar. Inflation rose gradually during the 1960s, was high and volatile in the 1970s, and fell sharply in the early 1980s. The run-up in inflation occurred at a time when the Samuelson–Solow model was dominant. Indeed, its probability weight is visually indistinguishable from 1 between 1960 and 1970. But in the early 1970s, evidence began to pile up against the model, and within 5 years its probability weight had fallen almost to zero. The data from those years were persuasive because they contradicted a key prediction of the Samuelson–Solow model, viz. that lower average unemployment could be attained at the cost of higher average inflation. Inflation was trending higher in the early 1970s, but so was unemployment. The policy experiment being run in those years was very informative about the Samuelson–Solow model.

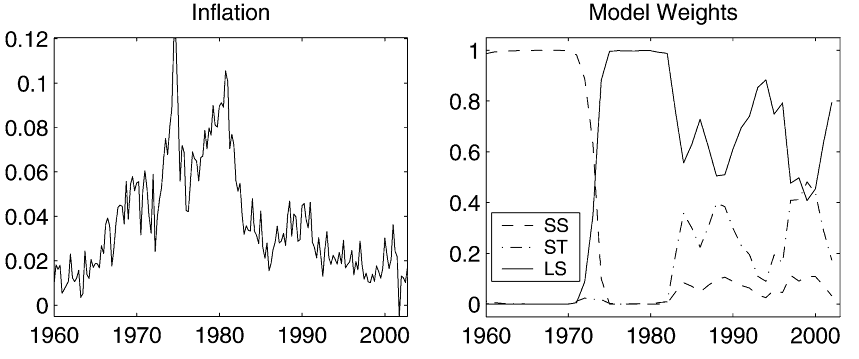

Fig. 1. Inflation and posterior model probabilities.

[23] The discount factor β matters more if one is willing to entertain nonstandard values, for reasons explained below.

The Lucas–Sargent model emerged to take its place. Its model weight rose from almost zero in the late 1960s, to approximately 0.5 by 1973, and then to almost 1 by 1975. It remained dominant until the early 1980s, after which it shared top billing with the Solow–Tobin model. Yet the period of its dominance was also the time when inflation was highest. As explained above, our version of the Lucas–Sargent model always recommends zero inflation. How is it that the Federal Reserve, which presumably weighs policy models according to their empirical plausibility, chose to follow a high inflation policy, at a time when the most plausible model recommended low inflation?

Although zero inflation is always the recommendation of the Lucas–Sargent model, it is not the recommendation of the system as a whole. The policy actions recommended by our Bayesian linear regulator are shown as dashed lines in Fig. 2, along with the history of actual inflation, which is portrayed by a solid line. For the most part, the path for $x_{t|t-1}$ tracks the general profile of y_t, but in the critical, middle period it is often higher and more volatile than actual inflation. Indeed, for the 1970s, the recommendations of the Bayesian linear regulator are far from those of the Lucas–Sargent model, even though its weight was close to 1. If anything, this figure deepens the puzzle.

The next two figures provide clues about why the Bayesian linear regulator behaves this way. Figure 3 shows the expected loss associated with a zero inflation policy (i.e., $f = 0$) in each of the submodels and for the system as a whole. Solid lines record the expected loss when it is finite, and open spaces represent an infinite expected loss. The open spaces are key features.

A zero inflation policy may have been optimal for the Lucas–Sargent model, but it would have been dreadful if implemented in the Keynesian models. Not until 1985 was the expected loss finite under zero inflation in both the Samuelson–Solow and Solow–Tobin models. The expected loss for the system as a whole is a probability weighted average

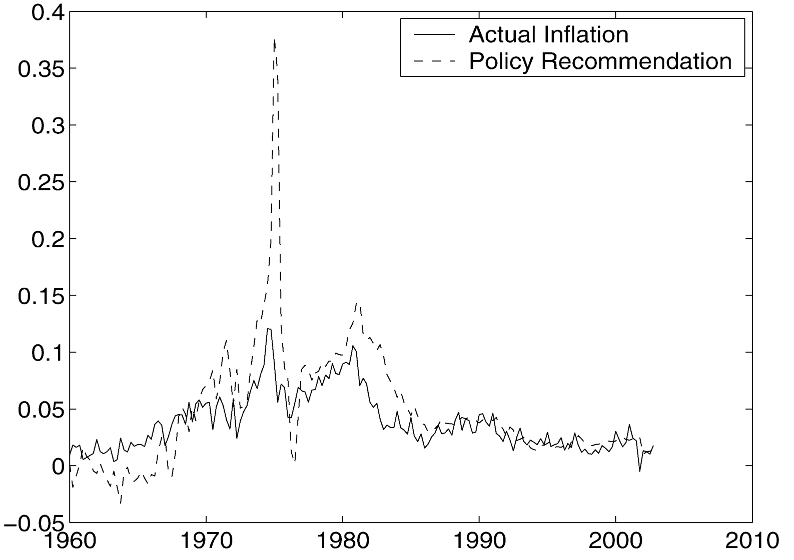

Fig. 2. Inflation and optimal policy.

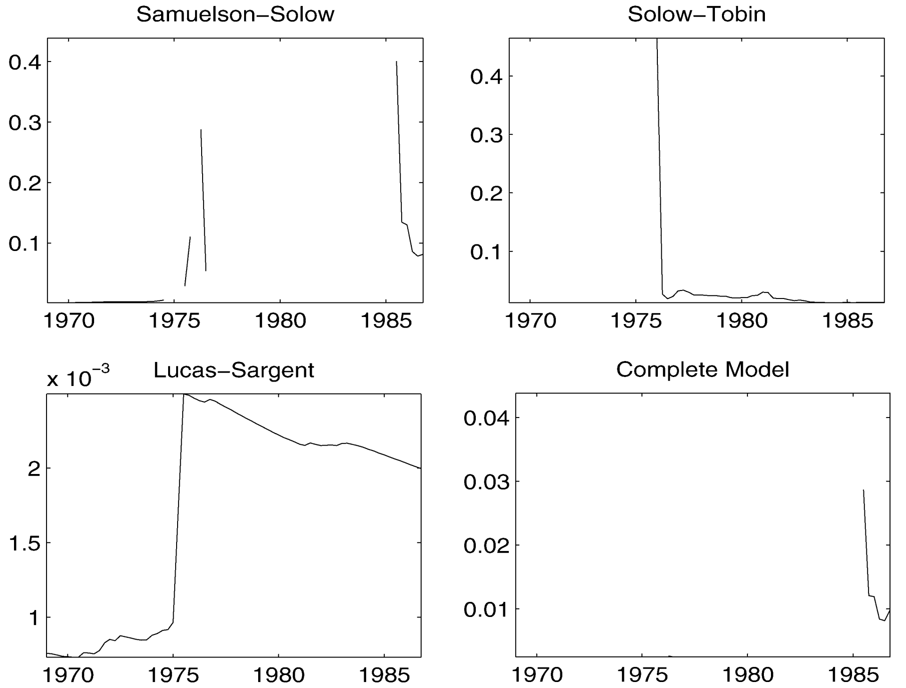

Fig. 3. Expected loss of a zero inflation policy.

of the losses in the submodels, and it is infinite whenever one of the submodel losses is. Therefore, the expected loss under the complete model was also infinite for most of this period, despite the Lucas–Sargent model's high probability weight.

With $f = 0$, an infinite expected loss occurs when an eigenvalue of A_i exceeds $\beta^{-1/2}$ in magnitude. Figure 4 shows recursive estimates of the dominant eigenvalue for each submodel, along with the stability boundary $\beta^{-1/2}$. The dominant eigenvalue of A_E is the upper envelope of the values for the submodels. The Solow–Tobin model was unstable under the zero inflation $f = 0$ policy prior to 1976, and the Samuelson–Solow model was unstable for most of the period from 1975 until 1985. Hence, for most of the period prior to 1985, one or both of the Keynesian submodels would have been unstable under a zero inflation policy, and so the complete model also would have been unstable.[24]

According to the Keynesian submodels of that vintage, attaining desirable unemployment outcomes required that policy feedback in the proper way on the state, and that failing to run policy in an appropriate activist way could destabilize unemployment. DeLong (1997) describes this view as a legacy of the Great Depression. This view emerges from

[24] Figure 4 reveals that during the 1970s there is one date at which the dominant eigenvalues of the Samuelson–Solow and the Solow–Tobin models both exceed $\beta^{-1/2}$. At that date, $x_{t\,|\,t-1}$ spikes upward in Fig. 2. A few quarters later, there is a single date at which the eigenvalues of both of those models become less than $\beta^{-1/2}$. At that lonely date, $x_{t\,|\,t-1}$ closely approximates the zero-inflation recommendation of the Lucas model.

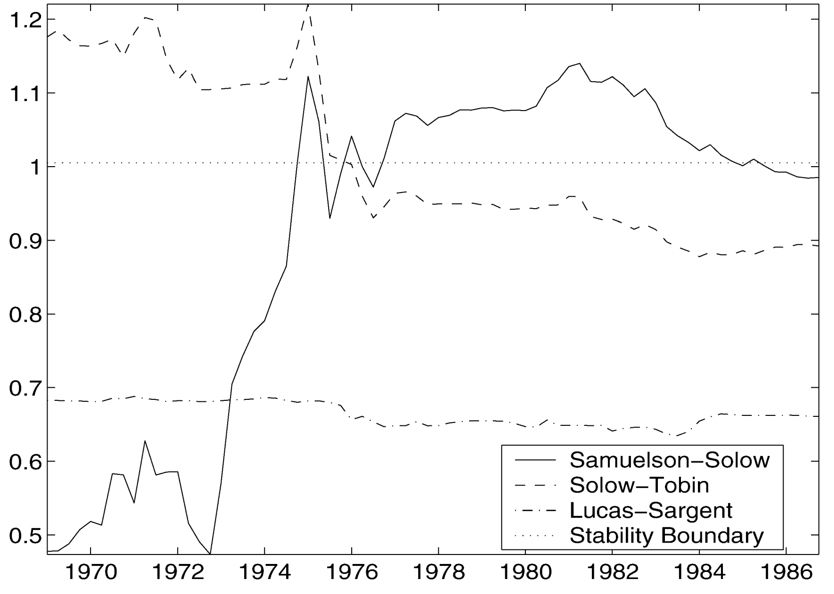

Fig. 4. Dominant eigenvalue under zero inflation.

the following pair of figures that compare expected outcomes under the optimal model-weighted policy[25] with those under zero inflation at two points in the 1970s. In Fig. 5, the forecast date is 1975.Q4, and in Fig. 6 it is 1979.Q4. Solid lines portray forecasts of unemployment under zero inflation, dashed lines depict forecasts of unemployment under the optimal policy, and the dashed and dotted lines show expected inflation under the optimal model-weighted policy.[26]

In the Lucas–Sargent model, expected unemployment is the same under the two rules, reflecting the irrelevance of the systematic part of policy for real variables. In the Keynesian submodels, however, projections for unemployment differ dramatically across rules. In the Samuelson–Solow and Solow–Tobin models, a zero-inflation policy triggers Depression levels of unemployment, an outcome that is very costly indeed. In contrast, Keynesian projections of unemployment are more ordinary under the optimal policy.

The projection for optimal inflation[27] calls for a gradual disinflation. For example, starting in 1975.Q4 inflation is expected to decline from 7.75 to 6.25 percent over a period of 2 years. Slightly more disinflation is expected in 1979.Q4, with inflation forecasts falling from 9.75 to 6.25 percent. But in both cases, inflation remains well above the Lucas–Sargent optimal value even after 2 years. The central bank prefers gradualism because it dramatically improves unemployment outcomes in the Keynesian scenarios albeit at the cost of worsening the inflation outlook in the Lucas–Sargent model. The Bayesian linear

[25] I.e., the policy produced by our Bayesian linear regulator.

[26] Expected inflation is of course zero under a zero inflation policy.

[27] This is the same in all submodels because there is a single policy rule.

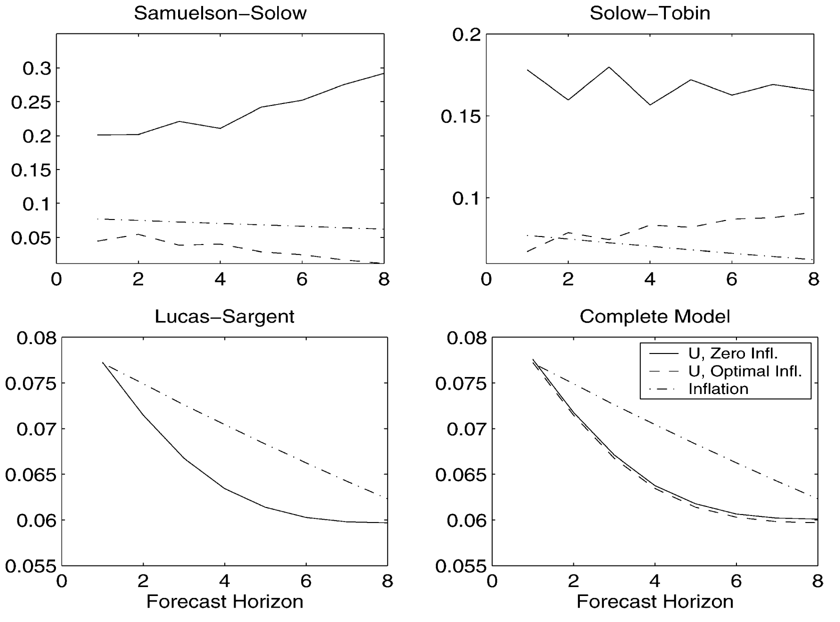

Fig. 5. Optimal policy v. zero inflation, 1975.Q4.

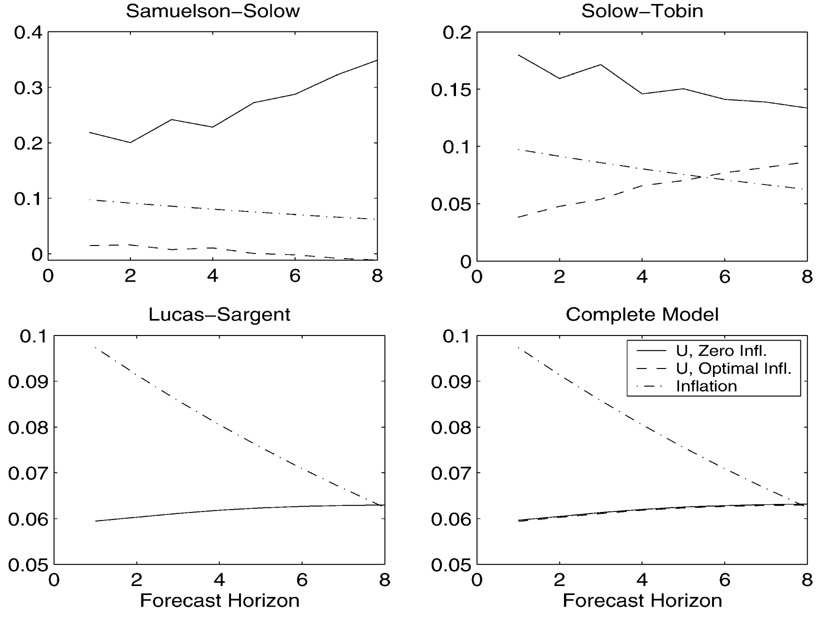

Fig. 6. Optimal policy v. zero inflation, 1979.Q4.

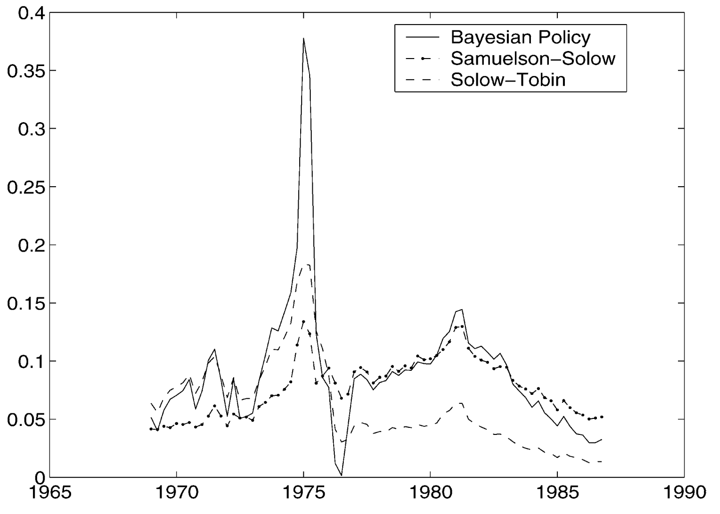

Fig. 7. Optimal policy and policy for worst-case scenarios.

regulator accepts this tradeoff. From his perspective, a worsening of the inflation outlook is a price worth paying to prevent a recurrence of the Depression.[28]

Furthermore, notice how closely the composite forecast tracks the Lucas–Sargent model projection. This reflects the high probability weight on the Lucas–Sargent model at that time. The Bayesian linear regulator accepts higher inflation to prevent a Depression, even though he thinks it will occur with low probability, because a Depression is a truly dreadful outcome. One may object that this concern was misplaced because it was so unlikely, and judging by the posterior model weights it was very unlikely. Yet our Bayesian linear regulator was unwilling to risk the possibility in exchange for the comparatively modest benefits of low inflation. The Bayesian linear regulator disregards the recommendation of the Lucas–Sargent model because there is a positive probability that it would have resulted in unbounded loss. The Lucas–Sargent optimal policy was not robust to the possibility that the data were generated by a Keynesian model.

In this instance, the Bayesian linear regulator behaves like a minimax controller, putting more weight on the recommendations of worst-case models than on that of the most likely model. Figure 7 illustrates this by comparing the optimal Bayesian policy with that of each submodel. For a discounted quadratic loss function such as ours, the worst-case scenario under a given policy rule is an unstabilized submodel. Before 1975, the Solow–Tobin model was the worst case, as it was the only submodel unstable under zero inflation. After 1977, the Samuelson–Solow model became the worst case, because then it was the

[28] This preference for gradualism helps resolve a puzzle about the correlation between the mean and persistence of inflation. The escape route models of Sargent (1999) and Cho et al. (2002) predict an inverse correlation, but Cogley and Sargent (2001, 2005) estimate a positive correlation. The model sketched here is broadly consistent with that estimate. The Bayesian linear regulator wants to reduce inflation in the 1970s, but he moves very slowly. Thus, when inflation was highest, the optimal policy called for a very gradual adjustment toward the target, making deviations from the target persistent.

unique submodel unstable under zero inflation. Notice how the optimal Bayesian policy first tracks the recommendations of the Solow–Tobin model when it is the worst case, and then switches to that of the Samuelson–Solow model when it becomes unstable. The most likely model has little influence compared with the worst-case models.[29]

In the model, it is literally an explosive root that causes the Bayesian linear regulator to shy away from a zero-inflation policy. The policy rule must cancel explosive roots to stabilize a system, so the Bayesian linear regulator pays special attention to the blocks of A_E in Eq. (17) in which explosive roots reside. All other considerations are secondary in comparison with losses arising from an unstabilized submodel. But this should be interpreted metaphorically. We suspect that conditional forecasts like those in Figs. 5 and 6 would be enough to deter real-world policy makers. Expectations that unemployment would exceed 15 percent for several years would probably be sufficient to dissuade the authorities from pursuing a cold-turkey disinflation.

3.1. Direction of fit

The direction-of-fit issue discussed by King and Watson (1994) is important in understanding how some of Primiceri's (2003) results relate to ours.[30] Like us, Primiceri emphasizes how monetary policy changed as the authorities updated their estimates, and he also attributes the inflation of the 1970s to the high perceived sacrifice ratio that Keynesian Phillips curve models presented to policy makers. But Primiceri assumes that the Fed relied exclusively on a version of the Solow–Tobin model and does not address why they disregarded the recommendations of the Lucas–Sargent model.

The central element of his story—the high perceived cost of disinflation—is not a robust prediction across models. On the contrary, it depends critically on the direction of fit. For example, Table 2 reports sacrifice ratios estimated from our three models, along with two other specifications that invert the Samuelson–Solow and Solow–Tobin models to put unemployment on the left-hand side and current inflation on the right. Here the sacrifice ratio refers to the cumulative output loss associated with reducing inflation by one percentage point relative to the inherited value and holding it there for 8 quarters. We project the consequences for unemployment in each model and then approximate the output loss using Okun's law that each percentage point of extra unemployment corresponds to 2.5 percent of foregone GDP.

Perceptions about the cost of disinflation vary a lot across models. The two Keynesian specifications, labeled *SS-K* and *ST-K*, have huge sacrifice ratios, ranging from around one quarter of a year's GDP to almost three fifths. On the other hand, sacrifice ratios are close to zero in Phillips curve models estimated with the classical direction of fit. Notice that this is true not only of the Lucas–Sargent model but also of the inverted Samuelson–Solow and

[29] The spike in x_t in 1975 represents the central bank's best response to a special challenge. The spike occurs during a brief window when both Keynesian submodels were unstable under zero inflation, forcing the central bank to address two 'worst-case' models at the same time. An extreme setting for x_t was needed simultaneously to stabilize both. The spike might seem implausible until one remembers that inflation actually did reach levels like this in other developed economies. For example, inflation in the UK peaked at 32 percent around that time.

[30] Again, see King and Watson (1994) for a critical discussion of Keynesian and Classical directions of fit.

Table 2
Sacrifice ratios and the direction of fit

	SS-K	*ST-K*	*SS-C*	*ST-C*	LS
1970.Q4	0.295	0.607	0.036	0.014	0
1975.Q4	0.578	0.249	−0.009	0.006	0
1979.Q4	0.636	0.227	−0.008	0.007	0

Note: Percent output loss associated with a 1 percentage point reduction in inflation sustained for 8 quarters.

Solow–Tobin representations, which are labeled *SS-C* and *ST-C*, respectively. Indeed, the direction of fit matters more than the qualitative nature of the tradeoff.

The reason that the sacrifice ratios differ so much has to do with how the models interpret a near-zero contemporaneous covariance between inflation and unemployment. In the early 1970s, this covariance moved downward toward zero, altering perceptions about the cost of disinflation. In a Keynesian Phillips curve, this diminished covariance flattens the short-term tradeoff, making the authorities believe that a long spell of high unemployment would be needed to bring down inflation, prompting Keynesian modelers to be less inclined to disinflate. But for a classical Phillips curve, the shift toward a zero covariance steepens the short-term tradeoff, making the authorities believe that inflation could be reduced at less cost in terms of higher unemployment. Thus, a classically-oriented policy maker would be more inclined to disinflate.

Classical versions of the Phillips curve were competitive in terms of fit, so it appears puzzling that so much weight would be put on the lessons of one model and so little on the others. Our story explains this by emphasizing a concern for robustness. Models predicting a high sacrifice ratio are weighed heavily for policy not because they dominate in terms of fit (they do not), but because they represent worst-case scenarios against which a prudent central bank wants to guard.

3.2. Sensitivity to discounting

When calculating expected loss at date t, the Bayesian linear regulator imagines that he will follow today's rule forever, despite knowing that the rule will be revised tomorrow in light of new information. In other words, he pretends a permanent commitment to the rule, when he knows it is only a temporary attachment. A short-term perspective might better reflect the mutability of policy rules. One way to shorten the planner's perspective, while remaining within the framework of an optimal linear regulator, is to reduce the discount factor β. This shifts weight away from losses expected to occur far into the future, when today's policy rule is less likely to be in effect, toward outcomes expected in the near future, when today's rule is likely to be more relevant.

In principal, this can matter for robustness because it expands the region within which a submodel has finite expected loss, thus enlarging the role of model probability weights. Figure 8 shows, however, that the basic picture is insensitive to changes in β. The figure portrays the Bayesian choice of $x_{t\,|\,t-1}$ for discount rates of 4, 8, 12, 16, and 20 percent per annum, respectively.

The general contour of optimal policy is altered only slightly by an increase in the discount rate. One difference is that a brief window opens in 1975–1976 during which

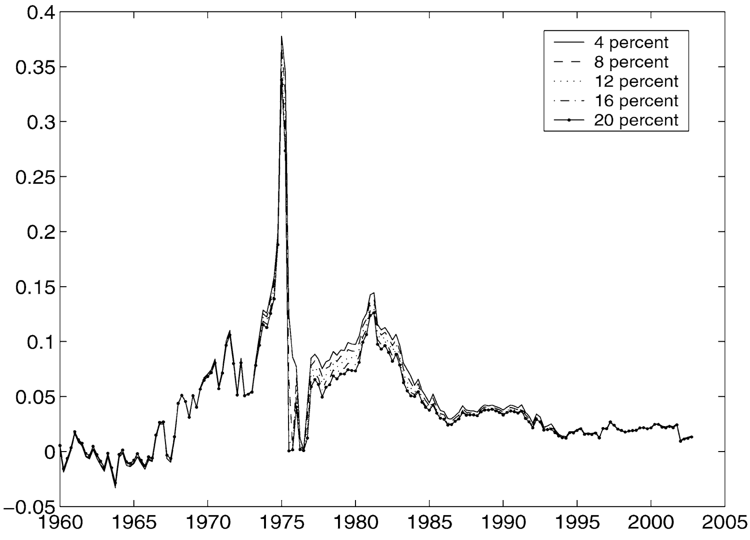

Fig. 8. Sensitivity of Bayesian policy to the discount factor.

zero inflation becomes more attractive. This occurs because both Keynesian submodels cross into the stable region for a few quarters, making the Bayesian linear regulator less concerned about achieving a bounded loss and more concerned about performance. As a consequence, he can pursue the Lucas–Sargent policy with less concern for its down-side risk. For lower values of β, $x_{t\,|\,t-1}$ is also slightly lower in the late 1970s and slightly higher after 1985, but none of these differences is especially important relative to the big picture.

One can of course eliminate a concern for robustness by discounting at a sufficiently high rate. This would require setting the discount factor so that $\beta^{1/2}$ times the largest eigenvalue of A_E is always less than 1. In the mid-1970s, the largest eigenvalue of A_E was approximately 1.2, so β would have to be around $1.2^{-2} = 0.69$ to accomplish this. But this corresponds to a 45 percent *quarterly* discount rate, which strains the interpretation that policy is set by intertemporal optimization. If we take the recursive model estimates as given, it seems difficult to dismiss a concern about unbounded losses in this way.

3.3. *Sensitivity to prior model weights*

The model's policy recommendations are also insensitive to the prior model weights. The baseline results are predicated on the assumption that $\alpha_{SS}(0) = 0.98$ and $\alpha_{ST}(0) = \alpha_{LS}(0) = 0.01$, reflecting that the Solow–Tobin and Lucas–Sargent models were developed later. Here we examine how the results change when the Solow–Tobin model is given greater initial weight. We hold $\alpha_{LS}(0)$ constant at 0.01 and increase $\alpha_{ST}(0)$ by borrowing from $\alpha_{SS}(0)$. Figures 9 and 10 illustrate the results.

The prior is informative, so changes in $\alpha_{ST}(0)$ alter views about which model is most likely at a given time. This influence is depicted in Fig. 9, which portrays posterior model weights for four scenarios. The dashed, dashed-dotted, and solid lines represent, respectively, the posterior weights on the Samuelson–Solow, Solow–Tobin, and Lucas–Sargent

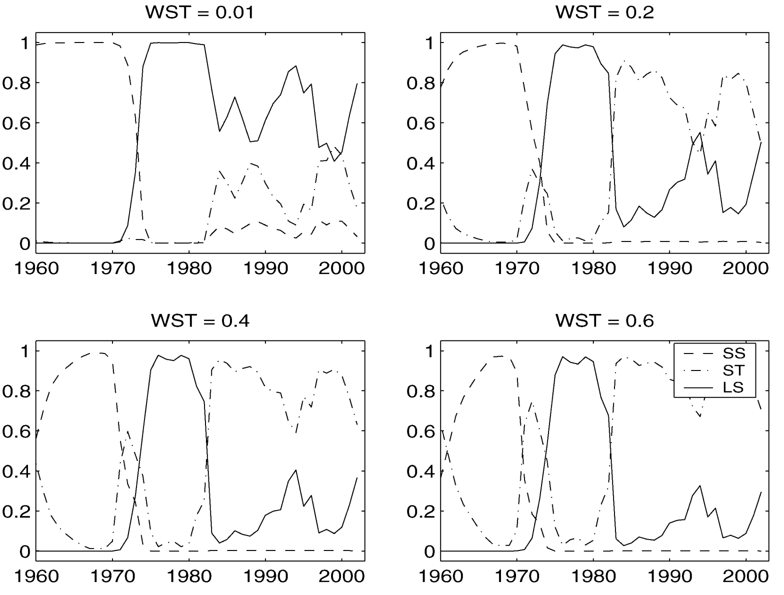

Fig. 9. Sensitivity of posterior model weights to prior model weights.

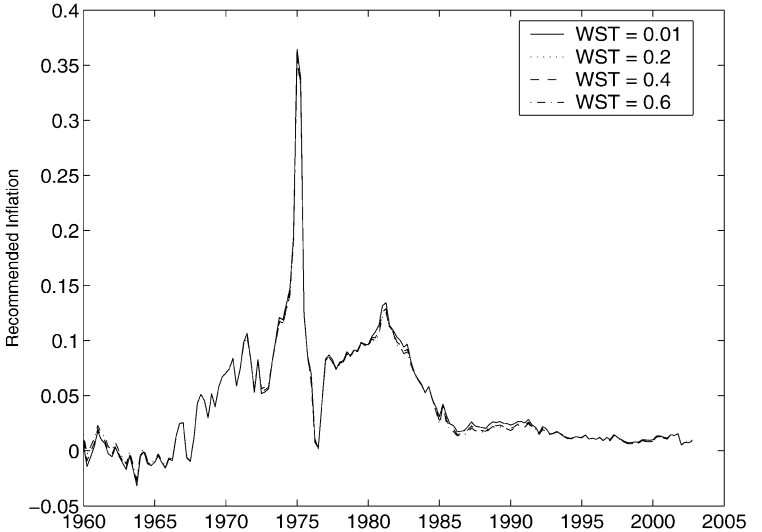

Fig. 10. Sensitivity of Bayesian policy to prior model weights.

models. The top-left panel reproduces the baseline case from Fig. 1, and the other three panels increase $\alpha_{ST}(0)$ to 0.2, 0.4, and 0.6, respectively. Thus, the baseline calibration assumes even prior odds on the Solow–Tobin and Lucas–Sargent models. In the other panels, the prior odds ratio increases to 20, 40, and 60, respectively, in favor of the Solow–Tobin model.

Two features are robust to this change, but one is not. In all four cases, the Samuelson–Solow model remains dominant in the 1960s, and the Lucas–Sargent model still has the highest probability weight for most of the 1970s and early 1980s. What changes is the ordering after the Volcker disinflation. As the prior odds on the Solow–Tobin model increase, its posterior weight catches up with and surpasses that on the Lucas–Sargent model. In this respect, the results are sensitive to initial values.

Figure 10 shows how this alters the Bayesian policy recommendation. The solid line reproduces the baseline calculation, and the other lines show how $x_{t\,|\,t-1}$ varies as $\alpha_{ST}(0)$ increases. There is virtually no difference, especially during the Great Inflation.[31] At that time, the Bayesian regulator was heavily influenced by concerns about worst-case scenarios, and since the identity of worst-case models does not depend on probability weights, it follows that changes in model weights have little influence on policy recommendations. In this respect, the results are insensitive to assumptions about the central bank's prior.

3.4. Anatomy of stabilization

Why, then, did inflation finally fall? In part, the answer is that recursive estimates of the Keynesian submodels eventually crossed into a region in which zero inflation was safe, so

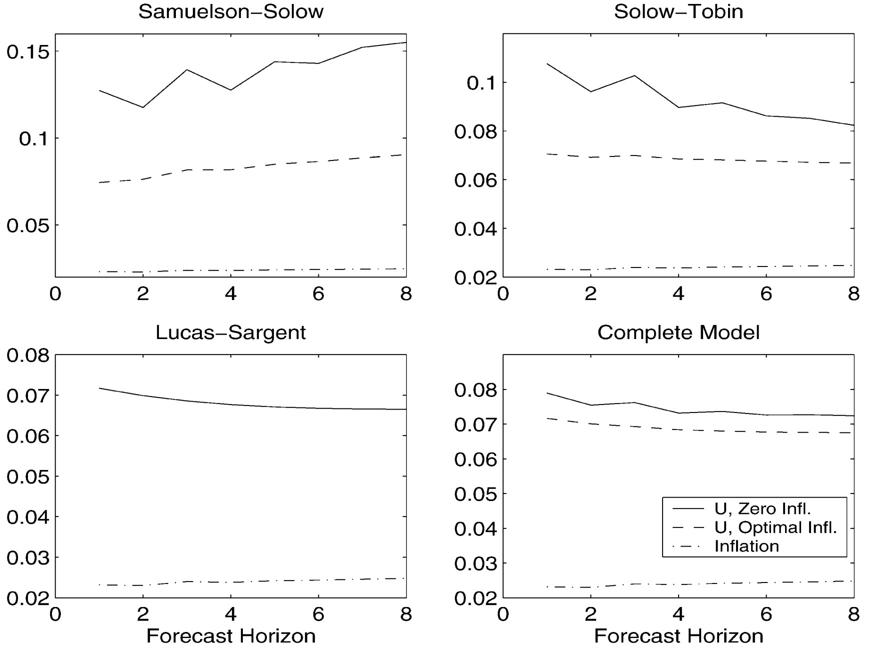

Fig. 11. Optimal policy v. zero inflation, 1992.Q4.

31 Inflation is slightly lower after the Volcker disinflation because the increase in $\alpha_{ST}(0)$ comes at the expense of a decrease in $\alpha_{SS}(0)$. This drives the posterior probability on the Samuelson–Solow model close to zero in the 1980s and 1990s, and that pushes down $x_{t\,|\,t-1}$.

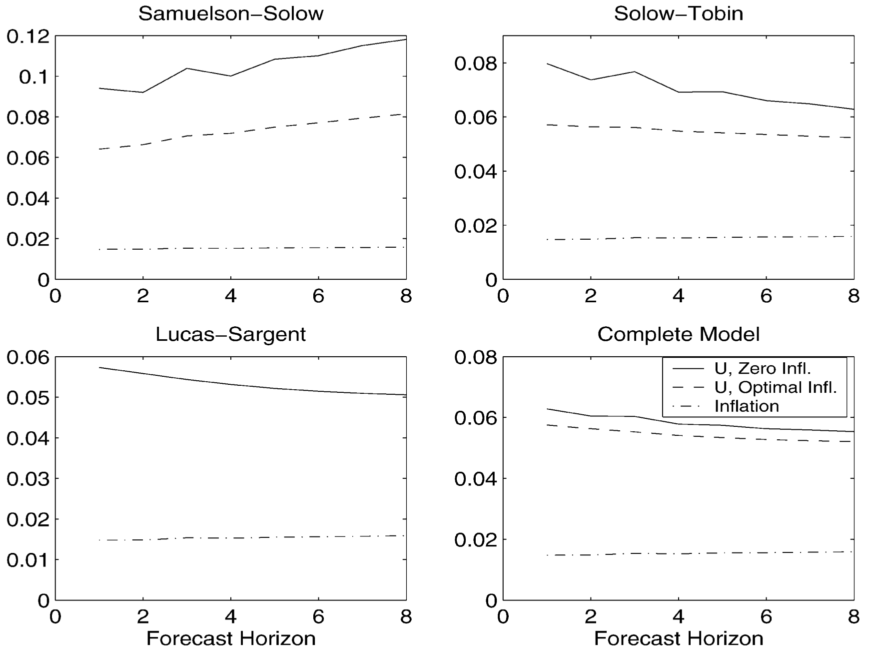

Fig. 12. Optimal policy v. zero inflation, 2002.Q4.

that concerns about robustness no longer dominated the choice of policy rule. By the 1990s, estimates of the Samuelson–Solow and Solow–Tobin models no longer predicted that low inflation would result in Depression-level unemployment. This is shown in Figs. 11 and 12, which provide two examples of forecasts from the Greenspan era. They have the same format as Figs. 5 and 6, but they advance the forecast dates to 1992.Q4 and 2002.Q4, respectively. By then, circumstances had become quite different from those in the 1970s.

Unemployment outcomes still differ across policy rules in the Keynesian submodels, but the difference is much less dramatic. Unemployment is higher under zero inflation, but now only by a few percentage points, and the forecasts do not approach the levels that earlier vintages of the same models predicted. Because the Bayesian linear regulator is less concerned about the risk of Depression, he is no longer willing to tolerate high and persistent inflation. Instead, he moves closer to zero inflation. The Lucas–Sargent model no longer has a probability close to one, however, so the optimal policy represents a compromise between zero inflation and Keynesian choices.

According to our calculations, the key difference between the 1970s and 1990s relates to concerns about a recurrence of the Depression. By the 1990s, this concern was substantially alleviated, allowing the Bayesian linear regulator to focus more on inflation.

4. 'Triumph' versus 'vindication'

Our findings blend aspects of two competing stories that Sargent (1999) told about Volcker's conquest of US inflation. The 'triumph of the natural rate' story has the monetary authority commit itself to keeping inflation low once it has accepted the rational expectations version of the natural rate theory. The 'vindication of econometric policy evaluation' story has the monetary authority remain unaware of the rational expectations version of the natural rate theory but nevertheless be induced to fight inflation after data render revised estimates of a Samuelson–Solow Phillips curve consistent with an imperfect version of the natural rate hypothesis. The imperfect version does not nest the rational expectations version, but it is good enough to cause an optimal control problem called the Phelps problem[32] to recommend low inflation.[33]

The triumph story has the central banker being persuaded by theoretical arguments to embrace a dogmatic prior favoring the rational expectations version of the natural rate hypothesis. The vindication story is about how the US inflation experience of the 1960s and 1970s induces an erroneously specified econometric model coupled with a control problem that violates the Lucas critique to give approximately correct advice.

Our empirical results weave together aspects of both stories because even after the evidence in favor of the Lucas version of the natural rate hypothesis becomes very strong, our Bayesian linear regulator does not become dogmatic—he still attaches positive, albeit small, probabilities to mostly discredited hypotheses. As we have stressed, those discredited hypotheses continue to exert a powerful influence on policy so long as they predict that very adverse outcomes would follow from adopting the policy recommendation that would flow from attaching probability one to the Lucas model. Our story requires that the Samuelson–Solow and the Solow–Tobin specifications both have to indicate less than disastrous outcomes before the Bayesian linear regulator can embrace recommendations that flow from the Lucas theory.[34]

5. Conclusion

One popular interpretation of the rise and fall of US inflation during the 1960s, 1970s, and 1980s emphasizes the central bank's changing beliefs about the natural rate hypothesis: the central bank conquered inflation because data generated by its own earlier misguided attempts to exploit the Phillips curve convinced it to accept the natural rate theory. But

[32] When the probability assigned to the Samuelson–Solow model is one, our Bayesian linear regulator problem becomes identical with the Phelps problem.

[33] Sargent's 'vindication' story took for granted that the Samuelson–Solow Phillips curve is misspecified, that the Phelps problem is subject to the Lucas critique, and that the data truly are generated by the rational expectations version of the natural rate hypothesis.

[34] The present paper is silent about self-confirming equilibria and escape routes, important ingredients of the analyses in Sargent (1999), Cho et al. (2002), Sargent and Williams (2003), and Sargent et al. (2004). Those papers endow the government with a single model, the Samuelson–Solow model. That model is incorrect out of equilibrium but correct in a self-confirming equilibrium. To conquer inflation, it is necessary to escape from the self-confirming equilibrium. Such escapes recur, punctuated by returns to the self-confirming equilibrium. Sargent (1999) attempted to interpret the conquest of US inflation as an escape episode that is bound to be temporary unless the government adopts a better specification of the Phillips curve. It would be an interesting and non-trivial exercise to extend that work to a setting in which the government is endowed with the mixture of models possessed by our Bayesian linear regulator. In a self-confirming equilibrium, the Lucas and Samuelson–Solow models are observationally equivalence. That might make it possible that the Samuelson–Solow model would not be discarded in the long run.

if the evolution of the central bank's beliefs were all that mattered, what postponed the conquest of inflation until the 1980s? The data had revealed the natural rate property by the early 1970s.

Our paper assembles evidence that confirms this timing puzzle. Using recursive Bayesian techniques, we find that posterior probabilities strongly favored a version of the Lucas–Sargent model as early as 1975. Nevertheless, the central bank did not implement that model's recommendations for inflation. This paper shows that a concern for robustness across a variety of models can explain why. According to our calculations, despite its high probability weight, the Lucas–Sargent model had little influence on policy because its recommendations were not robust across some other recently popular models. The central bank could agree with Lucas and Sargent about the workings of the economy, yet it could also refrain from adopting their policy recommendations because of its fears about the downside risk. In light of what we know now, say as represented by smoothed estimates of the approximating models, those fears may seem silly, but they would not have seemed so silly to decision makers armed with vintage-1970s estimates of Keynesian approximating models. On the contrary, those models warned that high and rising unemployment would accompany low inflation.

In this connection, it is useful to read again the analyses of stagflation that leading policy economists presented in the late 1970s. For example, the contributors to Okun and Perry's (1978) edited collection of essays all assign high probability to what we would categorize as either the Samuelson–Solow or the Solow–Tobin specification, and all of them take for granted that using monetary policy to reduce inflation would entail very large costs in terms of unemployment. Therefore they advocated alternative policy interventions (e.g., so called tax-based incomes policies). Okun and Perry (1978) summarize things as follows:

> "Thus, the mainline model and its empirical findings reaffirm that there is a slow-growth, high unemployment cure for inflation, but that it is an extremely expensive one. ... Using one of Perry's successful equations as an example, an extra percentage point of unemployment would lower the inflation rate by only about 0.3 percentage point after one year and by 0.7 percentage point if maintained for three years. That extra point of unemployment would cost over a million jobs and some $60 billion of real production each year" (p. 5).

Okun and Perry also summarize Perry's reasons for rejecting Fellner's suggestion that much lower costs in terms of unemployment could be attained through a credible disinflationary policy:

> Perry "believes that much of the [inflation] inertia is backward-looking rather than forward looking, and so is not susceptible to even convincing demonstrations that demand will be restrained in the future. [Perry's] own empirical evidence shows that wage developments are better explained in terms of the recent past history of wages and prices than on any assumption that people are predicting the future course of wages and prices in a way that differs from the past" (p. 6).

Perry (1978, pp. 50–51) forcefully elaborates on his argument against an expectational interpretation of Phillips curve dynamics. Okun (1978a, p. 284) says that "recession will slow inflation, but only at the absurd cost in production of roughly $200 billion per point."[35] At that time, $200 billion amounted to roughly 10 percent of GDP. Inflation averaged 7.4 percent from 1974 to 1979, and extrapolating to zero inflation implies a total cost of almost three quarters of a year's GDP.

Our Bayesian linear regulator attaches much lower probability to such outcomes than Okun and Perry did. But even so, a prudent central bank would also have been concerned about these outcomes and would have designed a policy that put a bound on its losses. Our calculations suggest that the high inflation of the 1970s was part of such a policy, given the models of the Phillips curve that research in the 1960s had presented to the central bank.

Our calculations also point to a connection between robust control theory and Bayesian model averaging. Historically, robust control theory was motivated by concerns about the stability of a system under a given decision rule and a perturbation to an approximating model. Practical decision makers who used ordinary control theory had found that controls that should have been optimal under the approximating model actually destabilized a system, resulting in bad payoffs. Because instability has catastrophic implications for a typical intertemporal objective function with little or no discounting, control theorists sought decision rules that would assure stability under a largest possible set of perturbations to an approximating model. This quest led directly to H_∞ control theory. Concerns about system stability are also foremost on the mind of our Bayesian linear regulator.

Acknowledgments

For comments and suggestions, we thank Willem Buiter, Bennett McCallum, Ulf Söderström, Harald Uhlig, an anonymous referee, and seminar participants at Georgetown, Harvard, Bocconi, the European Central Bank, the London Business School, and the University of Pennsylvania.

Appendix A. A recursive formula for the posterior model weights

Each period the central bank must update the model probability weights. For the normal-inverse gamma family, this can be done recursively. Let $\alpha_{i0} = p(M_i)$ represent the prior probability on model i. According to Bayes's theorem, the posterior probability is

$$p\big(M_i \mid Y^t, X^t\big) \propto m_{it} \cdot p(M_i) \equiv w_{it}, \tag{1}$$

where m_{it} is the marginalized likelihood function for model i at date t. The conditional likelihood for model i through date t is defined via a prediction error decomposition as

$$l\big(Y^t, X^t, \theta, \sigma^2\big) = \prod_{s=1}^{t} p\big(Y_s \mid X_s, \theta, \sigma^2\big). \tag{2}$$

[35] Okun (1978b) lists authors of the models on which this estimate is based. Many distinguished economists are included.

Then the marginalized likelihood is

$$m_{it} = \iint l\big(Y_i^t, X_i^t, \theta_i, \sigma_i^2\big) p\big(\theta_i, \sigma_i^2\big)\, \mathrm{d}\theta_i\, \mathrm{d}\sigma_i^2. \tag{3}$$

Notice that m_{it} is the normalizing constant in Bayes's theorem; i.e., the posterior density is

$$p\big(\theta_i, \sigma_i^2 \mid Z_i^t\big) = \frac{l(Y_i^t, X_i^t, \theta_i, \sigma_i^2) p(\theta_i, \sigma_i^2)}{\iint l(Y_i^t, X_i^t, \theta_i, \sigma_i^2) p(\theta_i, \sigma_i^2)\, \mathrm{d}\theta_i\, \mathrm{d}\sigma_i^2}. \tag{4}$$

By inverting this expression, we can express m_{it} as

$$m_{it} = \frac{l(Y_i^t, X_i^t, \theta_i, \sigma_i^2) p(\theta_i, \sigma_i^2)}{p(\theta_i, \sigma_i^2 \mid Z_i^t)}. \tag{5}$$

For the models described above, analytical expressions are available for all the functions on the right-hand side, so the marginal likelihood can be calculated simply by evaluating (5) at any point. We use the posterior mean of θ_i, σ_i^2, but any point would give the same answer. This is helpful because it side-steps the integration in (3).

To develop a recursion for the unnormalized model weights w_{it}, take the ratio of w_{it+1} to w_{it},

$$\begin{aligned}\frac{w_{it+1}}{w_{it}} &= \frac{m_{it+1} \cdot p(M_i)}{m_{it} \cdot p(M_i)} \\ &= \frac{l(Y_i^{t+1}, X_i^{t+1}, \theta_i, \sigma_i^2) p(\theta_i, \sigma_i^2)}{p(\theta_i, \sigma_i^2 \mid Z_i^{t+1})} \frac{p(\theta_i, \sigma_i^2 \mid Z_i^t)}{l(Y_i^t, X_i^t, \theta_i, \sigma_i^2) p(\theta_i, \sigma_i^2)}.\end{aligned} \tag{6}$$

Next, factor the conditional likelihood as

$$l\big(Y^{t+1}, X^{t+1}, \theta, \sigma^2\big) = p\big(Y_{t+1} \mid X_{t+1}, \theta, \sigma^2\big) l\big(Y^t, X^t, \theta, \sigma^2\big). \tag{7}$$

After substituting this expression into the previous equation and simplifying, one finds

$$\frac{w_{it+1}}{w_{it}} = p\big(Y_{it+1} \mid X_{it+1}, \theta_i, \sigma_i^2\big) \frac{p(\theta_i, \sigma_i^2 \mid Z_i^t)}{p(\theta_i, \sigma_i^2 \mid Z_i^{t+1})}, \tag{8}$$

or

$$\log w_{it+1} = \log w_{it} + \log p\big(Y_{it+1} \mid X_{it+1}, \theta_i, \sigma_i^2\big) - \log \frac{p(\theta_i, \sigma_i^2 \mid Z_i^{t+1})}{p(\theta_i, \sigma_i^2 \mid Z_i^t)}. \tag{9}$$

The term $\log p(Y_{it+1} \mid X_{it+1}, \theta_i, \sigma_i^2)$ is the conditional log-likelihood for observation $t+1$, and $\log p(\theta_i, \sigma_i^2 \mid Z_i^{t+1}) - \log p(\theta_i, \sigma_i^2 \mid Z_i^t)$ is the change in the log posterior that results from a new observation. These terms are easy to evaluate. The conditional likelihood is just

$$p\big(Y_{it} \mid X_{it}, \theta_i, \sigma_i^2\big) = \frac{1}{\sqrt{2\pi}\sigma_i} \exp\left[-\frac{1}{2} \frac{(Y_{it} - X_{it}'\theta_i)^2}{\sigma_i^2}\right], \tag{10}$$

and the posterior density is

$$p\big(\theta_i, \sigma_i^2 \mid Z_i^t\big) = C_{Ng}^{-1} \sigma_i^{-(\nu_{it}+k_i+2)} \exp\left[-\frac{s_{it} + (\theta_i - \hat{\theta}_{it})' X_{it}' X_{it} (\theta_i - \hat{\theta}_{it})}{2\sigma_i^2}\right], \tag{11}$$

where

$$C_{Ng} = \Gamma\left(\frac{\nu_{it}}{2}\right)\left(\frac{2}{s_{it}}\right)^{\nu_{it}/2}(2\pi)^{k_i/2}\left|X_i^{t\prime}X_i^t\right|^{-1/2}. \tag{12}$$

$\Gamma(\cdot)$ is the gamma function, k_i is the dimension of θ_i, ν_{it} is the posterior degrees of freedom for estimating σ_i^2, s_{it} is the sum-of-squared errors in the regression, $\hat{\theta}_{it}$ is the least-squares point estimate, and X_i^t is the matrix of regressors through date t.

After updating w_{it}, a renormalization enforces that the model weights sum to 1,

$$\alpha_{it} = \frac{w_{it}}{w_{1t} + w_{2t} + w_{3t}}. \tag{13}$$

An equivalent expression for α_{it} relates it to the sum of log posterior odds ratios,

$$\alpha_{it} = \left[\exp R_{1i}(t) + \exp R_{2i}(t) + \exp R_{3i}(t)\right]^{-1}, \tag{14}$$

where $R_{ji}(t) = (\log w_{jt} - \log w_{it})$ summarizes the weight of the evidence favoring model j relative to model i. Equation (14) says that α_{it} is inversely proportional to the weight of the evidence that can be advanced against model i.

Appendix B. State-space arrays for the policy models

This section provides details on the specification of the transition equations for each submodel. Along the way, we also clarify how the Keynesian approximating models are transformed into a ‘classical’ form, for the purpose of choosing a decision rule for $x_{t\,|\,t-1}$ in order to control y_t and u_t.

B.1. Samuelson–Solow model

We estimate a Samuelson–Solow Phillips curve of the form

$$\begin{aligned} y_{t+1} = {} & \gamma_0 + \gamma_1 u_{t+1} + \gamma_2 u_t + \gamma_3 u_{t-1} + \gamma_4 y_t + \gamma_5 y_{t-1} + \gamma_6 y_{t-2} \\ & + \gamma_7 y_{t-3} + \eta_{t+1}^{ss}. \end{aligned} \tag{15}$$

This can be expressed more compactly as

$$A_{K0}S_{t+1} = A_{K1}S_t + A_{K2}x_{t+1\,|\,t} + \eta_{K,t+1}, \tag{16}$$

where the state vector is defined as

$$S_t = [u_t, u_{t-1}, y_t, y_{t-1}, y_{t-2}, y_{t-3}, 1]'. \tag{17}$$

The system matrices A_{K0}, A_{K1}, and A_{K2} are

$$A_0^K = \begin{bmatrix} -\gamma_1 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \tag{18}$$

$$A_{K1} = \begin{bmatrix} \gamma_2 & \gamma_3 & \gamma_4 & \gamma_5 & \gamma_6 & \gamma_7 & \gamma_0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \tag{19}$$

and

$$A_{K2} = [\,0 \quad 0 \quad 1 \quad 0 \quad 0 \quad 0 \quad 0\,]', \tag{20}$$

respectively. The matrices A_{K0} and A_{K1} are updated recursively using the latest point estimates of (15). The residual vector is

$$\eta_{K,t+1} = \begin{bmatrix} \eta_{t+1}^{SS} & 0 & \xi_{t+1} & 0 & 0 & 0 & 0 \end{bmatrix}'. \tag{21}$$

Before solving the control problem, the model is transformed from a Keynesian to a classical direction of fit. This is accomplished by dividing the first row of the system by $-\gamma_1$, which delivers

$$A_{C0} = \begin{bmatrix} 1 & 0 & -1/\gamma_1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \tag{22}$$

$$A_{C1} = \begin{bmatrix} -\gamma_2/\gamma_1 & -\gamma_3/\gamma_1 & -\gamma_4/\gamma_1 & -\gamma_5/\gamma_1 & -\gamma_6/\gamma_1 & -\gamma_7/\gamma_1 & -\gamma_0/\gamma_1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \tag{23}$$

$$A_{C2} = A_{K2} = \begin{bmatrix} 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}',$$

and

$$\eta_{C,t+1} = \begin{bmatrix} -\eta_{t+1}^{SS}/\gamma_1 & 0 & \xi_{t+1} & 0 & 0 & 0 & 0 \end{bmatrix}'. \tag{24}$$

In the text, the reduced form transition equation is expressed as

$$S_{t+1} = AS_t + Bx_{t+1|t} + C\eta_{K,t+1}, \tag{25}$$

where $A = A_{C0}^{-1}A_{C1} = A_{K0}^{-1}A_{K1}$, $B = A_{C0}^{-1}A_{C2} = A_{K0}^{-1}A_{K2}$, and $C = A_{K0}^{-1}$.[36] The matrix A_{K0} is invertible if current unemployment has a non-zero coefficient in the Keynesian

[36] The extra step involving inversion to a classical form is not necessary for arriving at (25), but it clarifies the sense in which x_t is chosen to control y_t and u_t.

direction of fit; i.e., if $\gamma_1 \neq 0$. Finally, the targets are related to the state vector according to $(u_t, y_t)' = M_s S_t$, where the matrix M_s is defined as

$$M_s = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}. \tag{26}$$

B.2. Solow–Tobin model

The Solow–Tobin Phillips curve is a restricted version of the Samuelson–Solow model, specified so that the lag weights on inflation sum to one. The model also introduces a distinction between actual unemployment and the natural rate. We estimate the following version,

$$\Delta y_{t+1} = \delta_1 g_{t+1} + \delta_2 g_t + \delta_3 g_{t-1} + \delta_4 \Delta y_t + \delta_5 \Delta y_{t-1} + \delta_6 \Delta y_{t-2} + \eta^{ST}_{t+1}, \tag{27}$$

where $g_t = u_t - u^*_t$ is the unemployment gap. The natural rate u^*_t is measured by exponentially smoothed unemployment, as described in the text. With some recycling of notation, the model can be represented as

$$A_{K0} S_{t+1} = A_{K1} S_t + A_{K2} x_{t+1|t} + \eta_{t+1}, \tag{28}$$

where

$$S_t = [g_t, g_{t-1}, y_t, y_{t-1}, y_{t-2}, y_{t-3}, 1]', \tag{29}$$

$$A_{K0} = \begin{bmatrix} -\delta_1 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \tag{30}$$

$$A_{K1} = \begin{bmatrix} \delta_2 & \delta_3 & 1+\delta_4 & \delta_5-\delta_4 & \delta_6-\delta_5 & -\delta_6 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \tag{31}$$

and

$$A_{K2} = [\,0 \quad 0 \quad 1 \quad 0 \quad 0 \quad 0 \quad 0\,]'. \tag{32}$$

The matrices A_{K0} and A_{K1} are updated recursively using estimates of (27). The residual vector is

$$\eta_{K,t+1} = \begin{bmatrix} \eta^{ST}_{t+1} & 0 & \xi_{t+1} & 0 & 0 & 0 & 0 \end{bmatrix}'. \tag{33}$$

Once again, we transform to a classical representation before solving the control problem. The system matrices for the inverted, classical form are

$$A_{C0} = \begin{bmatrix} 1 & 0 & -1/\delta_1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \tag{34}$$

$$A_{C1} = \begin{bmatrix} -\delta_2/\delta_1 & -\delta_3 & -(1+\delta_4)/\delta_1 & -(\delta_5-\delta_4)/\delta_1 & -(\delta_6-\delta_5)//\delta_1 & \delta_6/\delta_1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \tag{35}$$

and

$$A_{C2} = A_{K2}. \tag{36}$$

The transformed residual vector becomes

$$\eta_{C,t+1} = \begin{bmatrix} -\eta^{ST}_{t+1}/\delta_1 & 0 & \xi_{t+1} & 0 & 0 & 0 & 0 \end{bmatrix}'. \tag{37}$$

The reduced form transition equation is

$$S_{t+1} = AS_t + Bx_{t+1|t} + C\eta_{K,t+1}, \tag{38}$$

where $A = A^{-1}_{C0}A_{C1} = A^{-1}_{K0}A_{K1}$, $B = A^{-1}_{C0}A_{C2} = A^{-1}_{K0}A_{K2}$, and $C = A^{-1}_{K0}$. The matrix A_{K0} is invertible if δ_1 is non-zero. The relation between targets and instruments differs slightly from that in the previous model, reflecting the distinction between actual unemployment and the natural rate. Now we have $(u_t, y_t)' = M_s S_t$, where

$$M_s = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & u^*_t \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}. \tag{39}$$

In this model, M_s is also updated each period to reflect the latest estimate of the natural rate.

B.3. Lucas–Sargent model

The Lucas–Sargent model is estimated from a classical direction of fit,

$$g_{t+1} = \phi_1(y_{t+1} - x_{t+1|t}) + \phi_2 g_t + \phi_3 g_{t-1} + \eta^{LS}_{t+1}, \tag{40}$$

where $g_t = u_t - u_t^*$ again measures the difference between actual unemployment and the exponentially smoothed proxy for the natural rate. The variable $x_{t\,|\,t-1}$ is generated recursively from the solution of the Bayesian linear regulator problem. The transition equation for this model can be written as

$$A_{C0} S_{t+1} = A_{C1} S_t + A_{C2} x_{t+1\,|\,t} + \eta_{C,t+1}, \tag{41}$$

where

$$S_t = [g_t, g_{t-1}, y_t, 1]', \tag{42}$$

and

$$\eta_{C,t+1} = \begin{bmatrix} \eta_{t+1}^{LS} & 0 & \xi_{t+1} & 0 \end{bmatrix}'. \tag{43}$$

The system arrays are

$$A_{C0} = \begin{bmatrix} 1 & 0 & -\phi_1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \tag{44}$$

$$A_{C1} = \begin{bmatrix} \phi_2 & \phi_3 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \tag{45}$$

and

$$A_{C2} = [\, -\phi_1 \quad 0 \quad 1 \quad 0\,]', \tag{46}$$

and they are updated using recursive estimates of (40). The residual vector is

$$\eta_{C,t+1} = \begin{bmatrix} \eta_{t+1}^{LS} & 0 & \xi_{t+1} & 0 \end{bmatrix}'. \tag{47}$$

The reduced form transition equation is

$$S_{t+1} = A S_t + B x_{t+1\,|\,t} + C \eta_{C,t+1}, \tag{48}$$

where $A = A_{C0}^{-1} A_{C1}$, $B = A_{C0}^{-1} A_{C2}$, and $C = I$. The relation between targets and the state is $(u_t, y_t)' = M_s S_t$, where

$$M_s = \begin{bmatrix} 1 & 0 & 0 & u_t^* \\ 0 & 0 & 1 & 0 \end{bmatrix}. \tag{49}$$

B.4. The composite model

The composite model collects the arrays (A, B, C, M_s) from each submodel and stacks them as follows:

$$\begin{bmatrix} S_{1t+j} \\ S_{2t+j} \\ S_{3t+j} \end{bmatrix} = \begin{bmatrix} A_1(t-1) & 0 & 0 \\ 0 & A_2(t-1) & 0 \\ 0 & 0 & A_3(t-1) \end{bmatrix} \begin{bmatrix} S_{1t+j-1} \\ S_{2t+j-1} \\ S_{3t+j-1} \end{bmatrix}$$

$$+\begin{bmatrix} B_1(t-1) \\ B_2(t-1) \\ B_3(t-1) \end{bmatrix} x_{t+j\,|\,t-1}$$
$$+\begin{bmatrix} C_1(t-1) & 0 & 0 \\ 0 & C_2(t-1) & 0 \\ 0 & 0 & C_3(t-1) \end{bmatrix} \begin{bmatrix} \eta_{1t} \\ \eta_{2t} \\ \eta_{3t} \end{bmatrix} \tag{50}$$

or in a more compact notation

$$S_{Et+j} = A_E(t-1)S_{Et+j-1} + B_E(t-1)x_{t+j\,|\,t-1} + C_E(t-1)\eta_t, \tag{51}$$

where $S_{Et} = [S'_{1t}, S'_{2t}, S'_{3t}]'$ and so on. The central bank's loss function can also be written in this notation. After averaging across models, the expected loss is

$$\begin{aligned} \mathcal{L}_E &= \alpha_{1t}\mathcal{L}(M_1) + \alpha_{2t}\mathcal{L}(M_2) + \alpha_{3t}\mathcal{L}(M_3), \\ &= E_t \sum_{i=1}^{3} \alpha_{it} \sum_{j=0}^{\infty} \beta^j \left(S'_{it+j} M'_{s_i} Q M_{s_i} S_{it+j} + x'_{t+j\,|\,t-1} R x_{t+j\,|\,t-1} \right), \\ &= E_t \sum_{j=0}^{\infty} \beta^j \left(S'_{E,t+j} Q'_{Et} S_{E,t+j} + x'_{t+j\,|\,t-1} R x_{t+j\,|\,t-1} \right), \end{aligned} \tag{52}$$

where

$$Q_{Et} = \begin{bmatrix} \alpha_{1t} M'_{s_1} Q M_{s_1} & 0 & 0 \\ 0 & \alpha_{2t} M'_{s_2} Q M_{s_2} & 0 \\ 0 & 0 & \alpha_{3t} M'_{s_3} Q M_{s_3} \end{bmatrix}. \tag{53}$$

This puts the model in the form of an optimal linear regulator problem, so that the optimal policy can be computed using standard algorithms.

References

Albanesi, S., Chari, V.V., Christiano, L.J., 2003. Expectations traps and monetary policy. Review of Economic Studies 70, 715–741.

Anderson, E.W., Hansen, L.P., McGratten, E.R., Sargent, T.J., 1995. Mechanics of forming and estimating dynamic linear economies. In: Rust, J. (Ed.), Handbook of Computational Economics. North-Holland.

Beck, G., Wieland, V., 2002. Learning and control in a changing environment. Journal of Economic Dynamics and Control 26, 1359–1377.

Blinder, A.S., 1998. Central Banking in Theory and Practice. MIT Press, Cambridge, MA.

Brainard, W.C., 1967. Uncertainty and the effectiveness of monetary policy. American Economic Review, Papers and Proceedings 57, 411–425.

Brock, W.A., Durlauf, S.N., West, K.D., 2003. Policy evaluation in uncertain economic environments. Brookings Papers on Economic Activity 1, 235–301.

Chari, V.V., Christiano, L.J., Eichenbaum, M., 1998. Expectations, traps and discretion. Journal of Economic Theory 81, 462–492.

Cho, I.K., Williams, N., Sargent, T.J., 2002. Escaping Nash inflation. Review of Economic Studies 69, 1–40.

Christiano, L.J., Fitzgerald, T., 2003. Inflation and monetary policy in the 20th century. Federal Reserve Bank of Chicago Economic Perspectives 27, 22–45.

Cogley, T., Sargent, T.J., 2001. Evolving post world war II US inflation dynamics. In: NBER Macroeconomics Annual, vol. 16, pp. 331–373.

Cogley, T., Sargent, T.J., 2005. Drifts and volatilities: Monetary policies and outcomes in the post-war US. Review of Economic Dynamics 8.

DeLong, J.B., 1997. America's only peacetime inflation: the 1970s. In: Romer, C.D., Romer, D.H. (Eds.), Reducing Inflation NBER Studies in Business Cycles, vol. 30. Univ. of Chicago Press, Chicago.

Erceg, C.J., Levin, A., 2003. Imperfect credibility and inflation persistence. Journal of Monetary Economics 50, 915–944.

Hall, R.E., 1999. Comment on rethinking the role of the NAIRU in monetary policy: implications of model formulation and uncertainty. In: Taylor, J.B. (Ed.), Monetary Policy Rules. NBER conference report. Univ. Chicago Press, Chicago.

Hansen, L.P., Sargent, T.J., 2001. Acknowledging misspecification in macroeconomic theory. Review of Economic Dynamics 4, 519–535.

Hansen, L.P., Sargent, T.J., 2002. Elements of robust control and filtering for macroeconomics. Unpublished manuscript. University of Chicago and New York University.

Ireland, P., 1999. Does the time-consistency problem explain the behavior of inflation in the United States? Journal of Monetary Economics 44, 279–292.

King, R.G., Watson, M.W., 1994. The post-war US Phillips curve: A revisionist econometric history. Carnegie-Rochester Conference Series on Public Policy 41, 157–219.

Kreps, D., 1998. Anticipated utility and dynamic choice, a 1997 Schwartz lecture. In: Jacobs, D.P., Kalai, E., Kamien, M. (Eds.), Frontiers of Research in Economic Theory. Cambridge Univ. Press, Cambridge, UK.

Lansing, K.J., 1999. Learning about a shift in trend output: implications for monetary policy. Unpublished manuscript. Federal Reserve Bank of San Francisco.

Lucas Jr., R.E., 1972. Expectations and the neutrality of money. Journal of Economic Theory 4, 103–124.

Lucas Jr., R.E., 1981. Methods and problems in business cycle theory. In: Lucas Jr., R.E. (Ed.), Studies in Business-Cycle Theory. MIT Press, Cambridge, MA.

Madigan, D., Raftery, A.E., 1994. Model selection and accounting for model uncertainty in graphical models using Occam's window. Journal of the American Statistical Association 89, 1535–1546.

Okun, A.M., 1978a. Discussion. In: Okun, A.M., Perry, G. (Eds.), Curing Chronic Inflation. The Brookings Institution, Washington, DC.

Okun, A.M., 1978b. Efficient disinflationary policies. American Economic Review, Papers and Proceedings 68, 348–352.

Okun, A.M., Perry, G., 1978. Editor's summary. In: Okun, A.M., Perry, G. (Eds.), Curing Chronic Inflation. The Brookings Institution, Washington, DC.

Orphanides, A., 2003. The quest for prosperity without inflation. Journal of Monetary Economics 50, 633–663.

Parkin, M., 1993. Inflation in North America. In: Shigehara, K. (Ed.), Price Stabilization in the 1990s. McMillan, London.

Perry, G., 1978. Slowing the wage–price spiral: the macroeconomic view. In: Okun, A.M., Perry, G. (Eds.), Curing Chronic Inflation. The Brookings Institution, Washington, DC.

Primiceri, G.E., 2003. Why inflation rose and fell: Policymakers' beliefs and US postwar stabilization policy. Unpublished manuscript. Princeton University.

Romer, C.D., Romer, D.H., 2002. The evolution of economic understanding and postwar stabilization policy. In: Rethinking Stabilization Policy. Federal Reserve Bank of Kansas City.

Samuelson, P.A., Solow, R.M., 1960. Analytical aspects of anti-inflation policy. American Economic Review, Papers and Proceedings 50, 177–184.

Sargent, T.J., 1973. Rational expectations, the real rate of interest, and the natural rate of unemployment. Brookings Papers on Economic Activity, 429–472.

Sargent, T.J., 1976. The observational equivalence of natural and unnatural rate theories of macroeconomics. Journal of Political Economy 84, 631–640.

Sargent, T.J., 1980. Lecture notes on filtering, control, and rational expectations. Unpublished manuscript. University of Minnesota, Minneapolis.

Sargent, T.J., 1999. The Conquest of American Inflation. Princeton Univ. Press, Princeton, NJ.

Sargent, T.J., Williams, N., 2003. Impacts of priors on convergence and escapes from Nash inflation. Review of Economic Dynamics. In press.

Sargent, T.J., Williams, N., Zha, T., 2004. Shocks and government beliefs: The rise and fall of American inflation. Working paper 10764. NBER.

Sims, C.A., 2001. Comment on Sargent and Cogley's 'Evolving post World War II US inflation dynamics'. In: NBER Macroeconomics Annual, vol. 16, pp. 373–379.
Solow, R.M., 1968. Recent controversy on the theory of inflation: an eclectic view. In: Rousseaus, S. (Ed.), Proceedings of a Symposium on Inflation: Its Causes, Consequences, and Control. New York University, New York.
Stokey, N.L., 1989. Reputation and time consistency. American Economic Review, Papers and Proceedings 79, 134–139.
Taylor, J.B., 1997. Comment on America's only peacetime inflation: the 1970s. In: Romer, C.D., Romer, D.H. (Eds.), Reducing Inflation, NBER Studies in Business Cycles, vol. 30. Univ. Chicago Press, Chicago.
Taylor, J.B., 2002. A half-century of changes in monetary policy. Written remarks delivered at the conference in honor of Milton Friedman, November 8.
Tobin, J., 1968. Discussion. In: Rousseaus, S. (Ed.), Proceedings of a Symposium on Inflation: Its Causes, Consequences, and Control. New York University, New York.
Velde, F., 2004. Poor hand or poor play? Alternative views of high inflation in the US. Federal Reserve Bank of Chicago Economic Perspectives 28, 34–51.
Wieland, V., 2000a. Monetary policy, parameter uncertainty, and optimal learning. Journal of Monetary Economics 46, 199–228.
Wieland, V., 2000b. Learning by doing and the value of optimal experimentation. Journal of Economic Dynamics and Control 24, 501–534.