STABILITY AND INSTABILITY IN US MONETARY POLICY BEHAVIOR

CHRISTOPHER A. SIMS

ABSTRACT. A monetary policy reaction function is estimated, allowing for several possible patterns of time variation in both its coefficients and its disturbance variances. A clear improvement in fit over a fixed-parameter linear model is found. The strongest effect on likelihood is from time variation in variances, but there are also improvements in fit from allowing coefficient variation. The variation is estimated as evolving in a stochastic, repeating pattern, not as evolution from one style of policy at mid-century and a new style in the 90's. The "regime shifts" that are estimated to occur do not last very long, and appear to reflect temporary shifts in the level of policy activism, not systematic improvement.

I. INTRODUCTION

One explanation of the rise in inflation in the US during 1950-1980 and of its subsequent return to low levels is that monetary policy was different, and worse, before 1980 than it has been since. The inflation was a result of policy errors, in this view, and the improved policy has both eliminated the inflation of the 70's and set up defenses against any repeat of the inflationary episode. This view has some statistical support (Clarida, Galí, and Gertler, 2000; Cogley and Sargent, 2001) and fits well with many macroeconomists' understanding of the history. But the evidence for substantial change in policy is weak, and seems to weaken further the more carefully the statistical analysis is carried out. Orphanides (2001) shows that when real-time data[1] are used to estimate a policy reaction function, the evidence for a change in policy after 1980 is weak. Hanson (2001) has examined VAR's for stability, finding that to the extent there is evidence of changed structure, it is stronger for the non-policy equations in the system than for the policy equations.

This paper looks for change in the monetary policy reaction function while taking account of time-varying variances in its disturbances. A constant-parameter reaction function, with a short interest rate on the left-hand side and estimated with monthly data US data,

Date: October 26, 2001.

Key words and phrases. time-varying parameters, hidden markov chain, regime switch, monetary policy.

Giorgio Primiceri provided excellent research assistance on this project. Discussions with Tao Zha have helped shape my thinking on this topic.

[1]That is, data that were available to policy makers at the time they were making their policy decisions.

is likely to show some estimated disturbances that are 6 or even 10 standard deviations away from zero. The disturbances are therefore so non-normal that standard statistical methods that ignore the heteroscedasticity may be very misleading. This point is especially important when the possibility of changes in coefficients is the center of interest. A specification that ignores heteroskedasticity or other sources of non-normality, while allowing for changes in regression coefficients, is very likely spuriously to find significant changes in coefficients, even if the only actual changes are in equation residual variances.

The paper finds very strong evidence of time-varying disturbance variances in the reaction function, and some evidence of time variation in the regression coefficients. However, the changes in coefficients in the model with the best fit are not at all like the "regime switches" of conventional thinking about policy. Instead of "post-1980" and "pre-1980" regimes, the data favor "activist" and "smoothing" regimes. The activist regimes occur rarely and last for only a few months, but arise both before and after 1980, and are associated with times of vigorous monetary policy action, most commonly when policy is countering the early stages of a recession. The rest of the time policy has much-dampened reactions to inflation and output growth.

II. The model

Most of the models that have been estimated for this paper are special cases of the following one.

$$r_t = c(S_t) + \sum_{i=1}^{k} \alpha_i(S_t) r_{t-i} + \sum_{j=1}^{m} b_{yj}(S_t) y_{t-j} + \sum_{\ell=1}^{m} b_{p\ell}(S_t) p_{t-\ell} + \varepsilon_t \tag{1}$$

$$\varepsilon(t) \,|\, \{S_v, y_v, p_v, v < t\} \sim N\big(0, \sigma^2(S_t)\big) \tag{2}$$

$$P\left[S_t = s_t \,\middle|\, \big\{S_v, y_{v-1}, p_{v-1}, v \leq t\big\}\right] = \pi(s_t \,|\, s_{t-1}), \tag{3}$$

where r is an interest rate, y is industrial production, p is a price index, and the time unit is monthly. In words, we estimate a regression equation with time-varying coefficients and time-varying disturbance variances, with the time variation governed by an exogenously evolving unobserved discrete state.

Though there are papers in the literature that estimate monetary policy reaction functions as single regression equations, there is good reason to think that this sort of specification has its limits. In any period in which policy authorities make reserves or any other monetary aggregate a short-run target, interest rates will respond quickly to disturbances elsewhere in the economy. As a result, regression equations with interest rates on the left will reflect a mixture of monetary policy behavior and private sector behavior. While the Federal Reserve has for most of our sample period arguably concentrated on a short run interest rate target, accommodating short run fluctuations in reserve demand, during part of 1979-92 it stated publicly that it was targeting reserves. While a change in our regression equation specific to that period can be interpreted as reflecting a change in policy, the actual coefficients estimated in that period may be unreliable as a description of policy behavior. Even outside that period, it would be better to have a multiple-equation specification that allowed for simultaneity and that included a monetary aggregate in the policy behavior equation. In another paper, joint with Tao Zha, we are working with such a model. However, because of its greater inherent complexity, it does not allow as flexible an exploration of forms of heteroskedasticity and time variation.

Other researchers (Clarida, Galí, and Gertler, 2000, e.g.) have in some instances estimated monetary reaction functions with explicitly forward-looking components, with current policy actions depending on expectations of future inflation and output growth. Such specifications, if estimated by single-equation methods, imply an equivalent backward-looking specification, obtained by substituting for expected future values a function of current information. They therefore do not conflict with the specification in this paper. They do imply the possibility of a stable behavioral rule in terms of expected future inflation and output growth, that appears unstable in a backward-looking specification. Expectational formulas could change, while reactions to expectations themselves remain stable. However, the papers in the literature that take this approach do not find stability of the reaction function. Furthermore, the type of instability found in this paper is not plausibly accounted for by simple shifts in expectations functions. And finally, the type of instability found in this paper implies that expectation-formation functions might be highly nonlinear, a possibility not allowed for in the literature.

The exogeneity of the state-evolution process is unappealing intuitively. Stories about the evolution of US monetary policy (in Cogley and Sargent (2001), e.g.) emphasize the possibility of changes in policy behavior responding to the history of the economy. The reserve-targeting period of the early 80's seems likely to have been a response to accelerating inflation, not a randomly timed shift. The data analysis for this paper has included examination of some models that allow for endogeneity of the change in state, and these will be described briefly below. However, the fit of these attempts has not been competitive with that of the exogenous-state specifications to which we devote most attention below. This accords with this paper's conclusion that there is little evidence of the kind of regime shift, from weakly counter-inflationary to strongly counter-inflationary, that dominates informal discussion in the literature. Such regime shifts would be expected to be endogenous. The changes in behavior actually found in this paper, between brief activist and more persistent smoothing regimes, may indeed be close to randomly timed.

Despite its limitations, this paper's specification seems worth exploring, so long as we bear in mind that its results represent a starting point, rather than the last word in this area.

III. RESULTS FROM THE BEST FIT

The best fit obtained to this point is with the following restricted form of the model:

$$\begin{aligned} r(t) = c(S_{1t}) + \sum_{j=1}^{6} \alpha_j r(t-j) \\ + b_y(S_{1t})\left(y_{t-1} - y_{t-4}\right) + b_p(S_{1t})\left(p_{t-1} - p_{t-4}\right) + \sigma(S_{2t})\varepsilon(t) \end{aligned} \tag{4}$$

$$\begin{gathered} S_1, S_2 \text{ evolve independently.} \quad S_{1t} \in \{1,2\}, S_{2t} \in \{1,2,3\} \\ P[S_{2t} = 3 \mid S_{2,t-1} = 1] = 0,\, P(S_{2t} = 1 \mid S_{2,t-1} = 3] = 0, \\ P[S_{2t} = 1 \mid S_{2,t-1} = 2] = P[S_{2t} = 3 \mid S_{2,t-1} = 2], \\ \sigma(j) \text{ normalized to increase in } j. \end{gathered} \tag{5}$$

This specification has one state, S_1, controlling regression coefficients, and another, S_2, evolving independently to control disturbance variances. It is a special case of our general model with one state variable, because the two states jointly determine 6 states, and the evolution of this one state with six values then has transition probabilities restricted by the requirement that S_1 and S_2 be independent and satisfy the restrictions (5).

The data are the natural logs of the 3 month Treasury Bill rate, industrial production, and either the CPI or the NAPM commodity price index, all over the period January 1948 to April 2001. We used the log of the interest rate because doing so consistently improved likelihood (of course accounting for the Jacobian term in comparing likelihoods). We used the Treasury Bill Rate to allow us to use data predating the active Federal Funds market, though we did also experiment with using the Federal Funds rate, with no important differences in results. We used simple differences of log prices and log industrial production, because our initial experiments with more flexible lag structures suggested there was little to be lost from a specification in differences.

Model estimates are reported in Table 1. The estimates use CPI as the price variable.[2] The first two entries in the π vector are the probabilities of staying in the two mean states. So it is apparent that the second mean state is modeled as having no chance of persisting more than one period. Both the other states have similar persistence parameters, tending to repeat about 96% of the time, and thus to last about two years. The first mean state gives modest sizes to b_y and b_p, and indeed in neither state is the b_p coefficient strongly significantly different from zero.[3]

---

[2]The model was estimated also with commodity prices as the price variable. Results were qualitatively similar.

[3]However, it should be noted that all the standard errors shown here are taken from the numerical approximation to the Hessian of the log likelihood built up by the optimization routine. These approximations are not guaranteed to be accurate.

parameter	estimate	standard error
a(s)	0.0133	0.004
	−0.0311	0.0143
b_y	0.2329	0.0887
	7.3729	0.642
b_p	0.3939	0.2632
	−3.158	1.6597
	1.2745	0.0422
	−0.3127	0.066
α	0.0106	0.0654
	0.0763	0.0677
	−0.0957	0.0639
	0.0355	0.0358
$\sum(\alpha)$	.9885	
	0.016940467	0.002598844
σ	0.040896941	0.002764809
	0.094737061	0.006964339
	0.961313133	0.016244964
	0.000000000	0.000231803
π	0.946601725	0.040633185
	0.95806119	0.019050042
	0.961164097	0.0211569
#params	20	
log LH	976.2188	

TABLE 1. Model 1: 3 variance states, 2 mean states, CPI

In the persistent mean state 1, the impact response of the *level* of the interest rate to a one percentage point change in the inflation rate at an annual rate, when the interest rate itself starts at 5%, is $.3939 \times 5 \times .0025$ or 8 basis points, a very small response. However, since the sum of coefficients on lagged r is near 1, the implied long run response is large, about 1.75 percentage points. The response to output growth in this state is smaller, though more sharply estimated. In the second mean state, the response to output growth is 30 times bigger on impact (which is all that matters here, because the state lasts only one period), so that there is a 1 percentage point interest rate response to each 1% change in output growth. The response to prices is also large in this state, but negative. It is only marginally significant statistically.

So we have here state 1 that implies slow, smooth response of r to right-hand side variables, with the responses in the expected direction and strong enough to imply uniqueness of the price level under conventional assumptions on fiscal policy. Another state 2 involves very fast, strong reactions to output growth, but is estimated to last only 1 month when it occurs. Figure 1 shows the estimated probabilities of the second coefficient state, together with a plot of the history of the interest rate.[4] It is apparent that state 2 tends to occur at times of rapid adjustment in interest rates, most often reductions in reaction to recessions, but also occasionally rapid increases. There are two long periods of sustained "state 1" behavior: the 70's and the 90's. The 90's do have some periods with probabilities of state 2 above 20%, whereas the 70's have no such indications of ambiguity.

While this model arguably offers the best fit of those explored in this project, others with somewhat different implications also fit fairly well. The leading contenders for good fit are listed in Table 2. Comparing the first and third rows of the table, we see that allowing more states for the mean, while reducing states for the variance, reduces likelihood while increasing parameter count. A model that is more similar to that in an earlier related paper (Sims, 1999) increases parameter count even more sharply. This model includes more lags of P and y, allows for 3 states, and has coefficients and variances moving jointly across states. It does produce a higher log likelihood, but only by 6, while adding 10 free parameters. This is not enough to allow even a classical 95% asymptotic χ^2 test to reject the simpler 3m,2v model in favor of the bigger model.[5]

[4]What is shown is the conditional probability of each state at each date, given the full data set and ignoring uncertainty in the maximum likelihood parameter estimates. The graph shows separately (in its color version) probabilities for the three variance states times the probability of the second mean state, but since times when more than one variance state has high probability are rare, the graph is nearly the same as one that simply plots the probability of the second mean state.

[5]The classical test statistic is justified only when the models are nested, which is not true here. The Bayesian reasoning underlying the Schwarz criterion does not require that the models be nested, however, so the failure to reject can be interpreted as an indication that even with priors rather strongly favoring the less restricted model, the posterior would be likely to favor the simpler one.

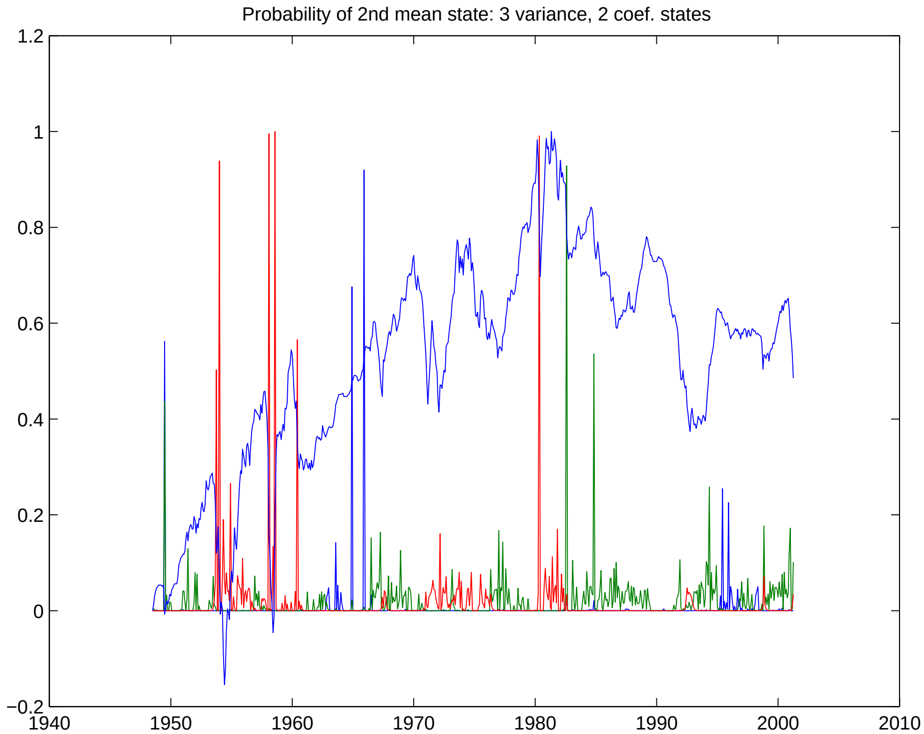

FIGURE 1

TABLE 2. Comparing model fits

	log LH	# parameters	SC	Akaike
3m, 2v	977.8	20	913.3	967.8
3m, 3v (Pcom)	983.1	25	905.7	970.1
2m, 3v	964.4	22	893.4	952.4
joint, 6 lag	983.8	30	887.0	953.8

The close competitor in this list is the second row, which simply adds an additional mean state. This increases the parameter count by 5 and increases likelihood by 5.3. This is not enough to allow rejection of the simpler model (which in this case is nested) by the classical asymptotic χ^2, and the Schwarz criterion (column SC) firmly favors the smaller model, but the Akaike criterion favors the larger model (as it always does when log likelihood increases by more than parameter count).[6]

Estimates for the 3-mean, 3-variance state model are reported in Table III. Here there is one state, state 2, that involves strong reaction to both output and price, in the expected

[6]Results are shown here for commodity prices as the price variable. In all the other models, CPI produced slightly higher likelihood, but here commodity prices did.

direction and with a well-determined coefficient. The much smaller coefficient on inflation than output growth in part reflects the much greater volatility of commodity prices. State 1 has coefficients of the expected sign, but they are small and ill-determined. State 3 has a well-determined coefficient of perverse sign on output and a small, ill-determined coefficient on inflation. The most salient feature of state 3 is its large positive constant term.

Figures 2-4 show the probabilities of the three mean states for this model, starting with the second, well-behaved one. From Figure 2 we see that this strongly reacting behavior is not very persistent, tending to show up at times of rapid changes in the interest rate, particularly around recessions. It occurs at intervals throughout the sample period, and a period of this type of period is estimated to have begun in early 2001. The first policy state, which might be characterized as “inactive”, is more persistent, and also occurs throughout the sample period. The third, which might be characterized as “purely contractionary”, It also occurs throughout the period, though there is a relatively long stretch in the 70’s where it did not occur.

My own conclusion is that this three by three modeling framework is not an improvement over the three by two. It creates non-persistent states with ill-determined coefficients, and “explains” quite a bit of behavior as exogenous shifts in a target interest rate. So it neither adds to fit nor suggests useful interpretations.

[There are further aspects of the fit that deserve some discussion, but time constraints limit their treatment here. The pure constant parameter model produces residuals with extreme non-normality. Both the 3×3 and the 3×2 models, but especially the latter, succeed in “standardizing” the disturbances. That is, they make the distribution of prediction errors divided by expected standard deviation look well-behaved. I should also catalog here the variant models that were tried but are not presented in detail because of poor fit.]

REFERENCES

CLARIDA, R., J. GALÍ, AND M. GERTLER (2000): “Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory,” *Quarterly Journal of Economics*, CXV, 147–180.

COGLEY, T., AND T. J. SARGENT (2001): “Evolving US Post-World War II Inflation Dynamics,” *NBER Macroeconomics Annual*.

HANSON, M. (2001): “Varying Monetary Policy Regimes: A Vector Autoregressive Investigation,” Discussion paper, Wesleyan University.

ORPHANIDES, A. (2001): “Monetary Policy Rules, Macroeconomic Stability, and Inflation: A View from the Trenches,” Discussion paper, Board of Governors of the Federal Reserve System.

parameter	estimate	standard error
	0.0168	0.0042
a	-0.0002	0.0067
	0.0674	0.007
	0.1319	0.107
b_y	2.5795	0.4363
	-1.1826	0.2385
	0.0157	0.0091
b_p	0.1034	0.0317
	0.0338	0.0436
	1.1344	0.046
	-0.2165	0.0659
α	0.0701	0.0642
	0.0097	0.0632
	-0.0333	0.0545
	0.0234	0.0318
$\Sigma(\alpha)$	0.9878	
	0.010931574	0.001306428
σ	0.035764482	0.002315015
	0.101860607	0.007570382
	0.952565091	0.019293033
	0.700294367	0.073906049
π	0.806558299	0.059912261
	0.867024231	0.044263643
	0.920209727	0.022809501
	0.92477035	0.024993024
#params	24	
log LH	983.138	

TABLE 3. Model 2: 3 variance states, 3 mean states, Pcom

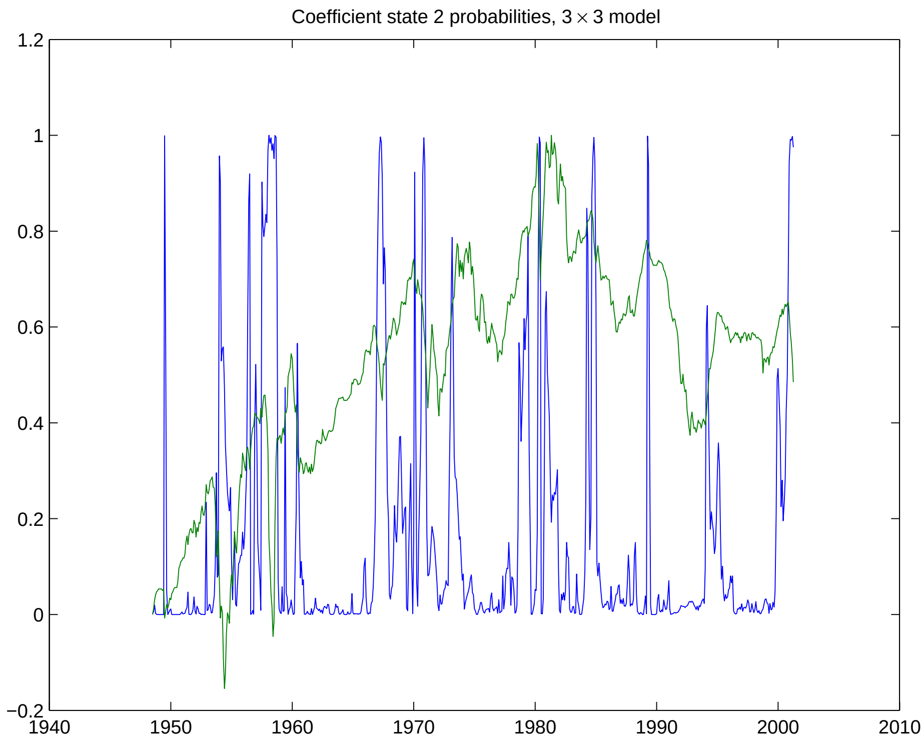

FIGURE 2

SIMS, C. A. (1999): "Drift and Breaks in Monetary Policy," Discussion paper, Princeton University, `http://www.princeton.edu/~sims/`, Presented at a plenary session of the July, 1999 meetings of the Econometric Society, Australasian region.

DEPARTMENT OF ECONOMICS, PRINCETON UNIVERSITY
E-mail address: `sims@princeton.edu`

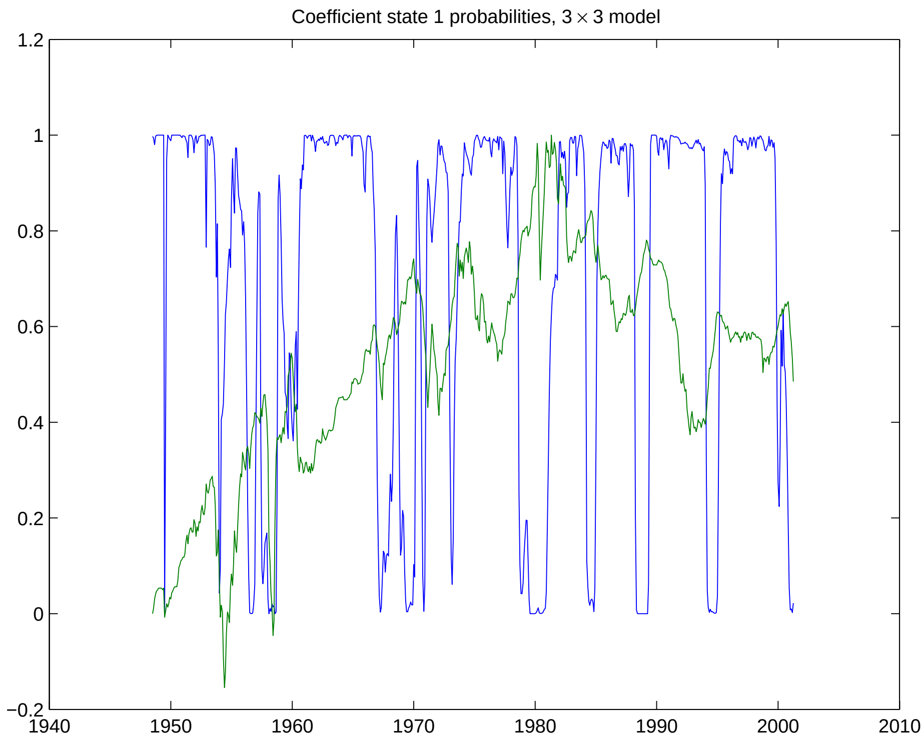
Coefficient state 1 probabilities, 3 × 3 model
1.2
1
0.8
0.6
0.4
0.2
0
−0.2
1940
1950
1960
1970
1980
1990
2000
2010

FIGURE 3

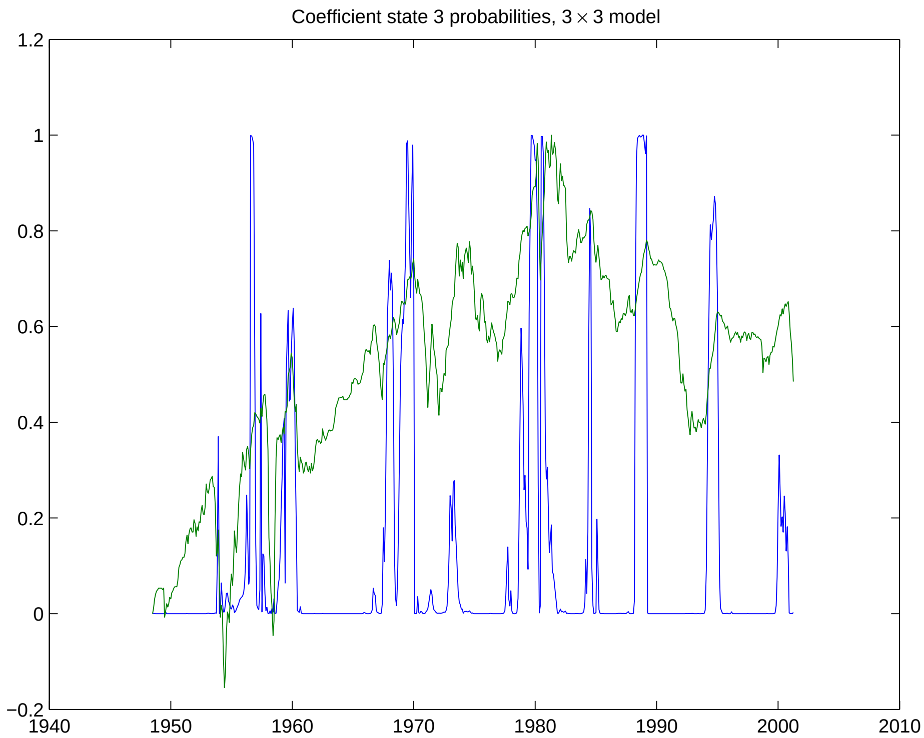
Coefficient state 3 probabilities, 3 × 3 model
1.2
1
0.8
0.6
0.4
0.2
0
−0.2
1940
1950
1960
1970
1980
1990
2000
2010

FIGURE 4