ELSEVIER Journal of Monetary Economics 41 (1998) 277–300

JOURNAL OF Monetary ECONOMICS

What do the VARs mean? Measuring the output effects of monetary policy

John H. Cochrane[a,b,c,*]

[a] Graduate School of Business, University of Chicago, Chicago, IL 60637, USA
[b] Federal Reserve Bank of Chicago, Chicago, IL 60604, USA
[c] National Bureau of Economic Research, Cambridge, MA 02138, USA

Received 15 April 1996; received in revised form 2 June 1997; accepted 2 July 1997

Abstract

VARs describe the history of output and other variables following monetary shocks. To measure the *effects* of monetary shocks, one must add economic identifying assumptions. I specify the relative effects of anticipated and unanticipated money, and I calculate how VAR-based measures of the effect of money on output change as one varies this assumption. The anticipated/unanticipated assumption influences measured output effects as much or more than the variable selection and shock orthogonalization assumptions on which the VAR literature focuses. Assuming that anticipated monetary policy can have some effect on output results in much shorter, smaller, and perhaps more believable estimates of the output response to monetary shocks. © 1998 Elsevier Science B.V. All rights reserved.

JEL classification: E52; C32; C52

Keywords: Monetary policy; Vector autoregression; Identification

1. Introduction

What are the real effects of monetary policy? A huge recent literature has returned to this classic question, and has measured the history of real and nominal variables following monetary shocks.[1] By thinking carefully about how

* Correspondence address: Graduate School of Business, University of Chicago, 1101 E. 58th Street, Chicago, IL 60637. Tel.: 773-702-3059; fax: 773-702-0458; e-mail: john.cochrane@gsb.uchicago.edu.

[1] See among many others Bernanke and Blinder (1988), Christiano and Eichenbaum (1995), Christiano, Eichenbaum and Evans (1996), Gordon and Leeper (1993), Sims (1992) and Strongin (1992) in the VAR tradition, and Romer and Romer (1994) for an estimate based on reading the FOMC minutes to identify money supply shocks.

PII S0304-3932(97)00075-5

to identify monetary policy shocks (or, a critic might charge, by extensive fishing), this literature has at last produced impulse-response functions that capture common views about monetary policy.

Fig. 1 presents two examples, taken from Cochrane (1994a) which surveys this literature. (The specification and estimation of these VARs is described in the Appendix.) The top row presents responses to a one standard deviation M2 shock. An unexpected rise in M2 is followed by a protracted rise in M2. The federal funds rate responds with an initial decline and then a protracted rise. This response can be interpreted as a 'liquidity effect' followed by an 'inflation effect'. Consumption and output rise temporarily, and prices rise. The long-run consumption and output effects are much less than one standard error from zero. In the bottom row, federal funds rate innovations are used to identify money supply shocks. In response to a surprise increase in the federal funds rate, consumption and output decline temporarily, prices decline, and M1 declines. The federal funds rate also stays high for quite some time after the initial shock.

The general patterns in Fig. 1 are just what Friedman (1968) taught us to expect. However, the magnitudes are a bit surprising. The output responses are protracted, hump-shaped and large. Output peaks two *years* after the M2 shock, and takes five years to die out. Output rises by about 0.5% following a one-standard deviation shock, which is also about 0.5%; since M2 is much smaller than output, and the base is again much smaller than M2, this response implies a large amplification in dollar terms. And I selected these specifications in part to minimize the size and length of output responses. Simpler specifications (M2, output and price level for example) produce even larger and more protracted responses.

These impulse-response functions capture *history*, the average values of money, output, interest rates and other variables following a money supply shock. What does this *history* tell us about the *effects* of monetary policy? What does it tell us, for example, about the course of events we should expect if there is a monetary shock *not* followed by the customary further expansion of money? Put another way, how important is the proviso, widely neglected in the VAR literature, that the VAR measures the effects of policy shocks, *only* if they are followed by the customary further policy actions? What does this history tell us about monetary *theory*, about which kinds of monetary models one needs to construct in order to understand and guide monetary policy? In particular, do the responses mean that we must construct monetary theories with important delay and propagation mechanisms?

It is tempting to think that, once monetary policy shocks are correctly identified, the impulse-response functions provide direct answers to these questions, direct measures of the effects of monetary policy. The extensive economic interpretations of impulse-response functions in the VAR literature can easily leave readers with this impression. But Sargent's (1976) 'Observational equivalence' warns us that many different theories of the *effects* of monetary policy are

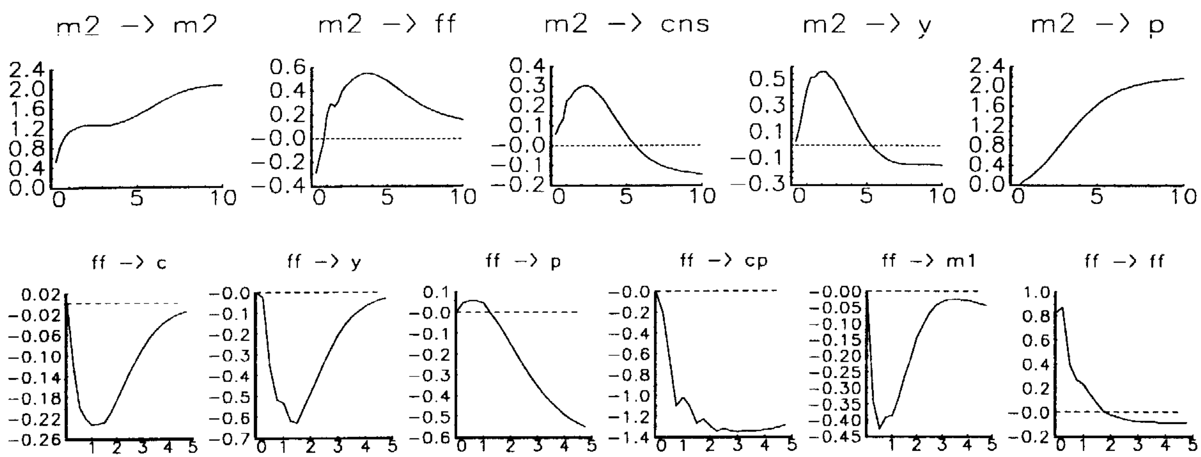

Fig. 1. Responses to monetary shocks in two VARs. Horizontal axis in years, variables are 100 * logs except ff in percent. Top panel: response to M2 shocks. Bottom panel: response to ff shocks. See the Appendix for details on the VAR specification and estimation.

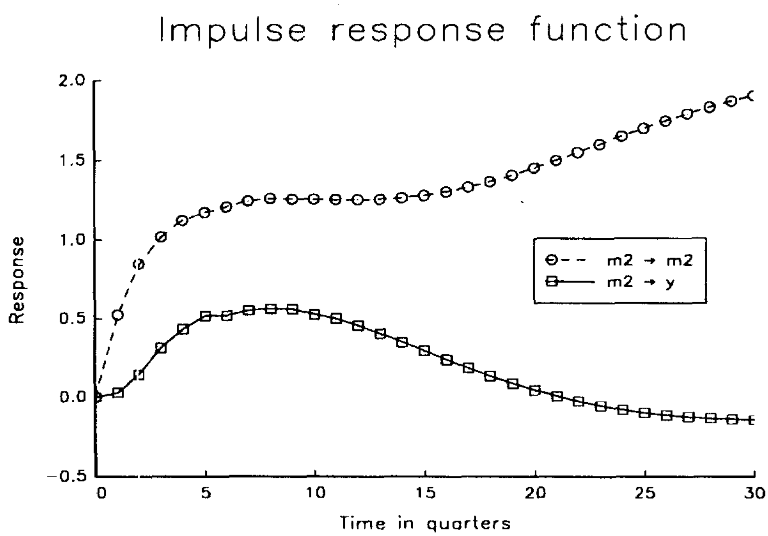

Fig. 2. M2 and output responses from M2 VAR of Fig. 1, plotted on same scale.

in fact consistent with the *history* of monetary policy as captured by a VAR, even if the shocks are correctly identified. A further, theoretical, identification is necessary to measure the effects of monetary policy.

To be specific, focus on the money → money and money → output responses plotted together in Fig. 2. If only unanticipated monetary shocks can affect output, the output response does measure the dynamic effect of the initial monetary shock. Precisely, in this case the output response is 'policy invariant': A money shock will have this effect on output in any regime, or no matter what the path of money following the shock.

But suppose instead that *anticipated* monetary shocks can also affect output. Now, the continued expansion of money that is expected to follow a shock can, *when it happens*, cause the prolonged expansion of output following the shock. The *effects* of a monetary policy shock on output, i.e. the policy-invariant quantity, or the path of output following a shock that is not followed by the habitual further expansion of money, may be in fact small, short and decaying (rather than hump-shaped). If a model produces anticipated money effects, it may not require amplification and delay mechanisms to be consistent with the data.

Of course, one wants numbers or graphs; estimates not stories. Exactly what measure of the effects of monetary policy does one get from assuming that anticipated money can matter? How short and small are the true response functions in this case? In the body of this paper, I adopt alternative explicit identifying assumptions in which anticipated money can have some effect on output. From the estimated impulse-response functions, I calculate estimates of the effect of money on output under the alternative identifying assumptions. I find that the estimated output responses vary a lot as one changes identifying assumption. The anticipated-unanticipated identifying assumption is as important to the output effect measures as the variable selection and orthogonalization assumptions on which the VAR literature focuses. I also find that the anticipated money models do result in output effect estimates that are much shorter and smaller than the impulse-response function.

The point and contribution of this paper is what to do with impulse response functions after they have been estimated. Therefore, I use a fairly standard specification and do not focus on the variable selection, specification, and orthogonalization issues on which the VAR literature rightly focuses. The central general pattern of 'correct' signs, persistent policy, and persistent responses to shocks is not at all sensitive to the details of specification, and is common to almost all the monetary VARs in the literature. Therefore, the qualitative results apply to most of that literature.

1.1. Expected money?

Should an empiricist even consider the possibility that anticipated monetary can have real effects? Lucas (1972, 1973, and in a forceful summary, 1996) says no. In Lucas' model, only *unexpected* monetary policy shocks can have real effects. Lucas' view is attractive, since it can explain the fact that monetary policy sometimes seems to have large output effects but at other times seems to have small or no effects. The experience of 1980–1982 is often seen as a classic example of a large effect. The debate over the great depression seems now to be over *how* monetary policy caused the depression rather than *whether* it did so. But in the ends of hyperinflations (Sargent, 1986) and currency revaluations, money growth or its stock can change by factors of thousands, literally overnight, with no real effect at all.

Lucas' reconciliation of these different regimes has a dramatic implication: *systematic* monetary policies – policies taken predictably by central banks to offset recessions – have no real effects, a point made forcefully by Sargent and Wallace (1975). This implication may help to explain the difficulties that central banks have faced in fine tuning output or exchange rates by trying to systematically offset real shocks.

However, Lucas' view has not been universally accepted. Many authors state or reveal views that anticipated money has real effects or (equivalently) that systematic policy matters. For example, Romer and Romer (1994) claim that systematic monetary policy ended postwar recessions. The argument against currency unions and fixed exchange rates is that they prevent national governments from pursuing systematic monetary policy to offset country-specific real shocks. The literature that evaluates nominal GNP targeting (Feldstein and

Stock, 1994; Hall and Mankiw, 1994, are recent examples) is predicated on the idea that better systematic policies can reduce the variance of output.

Monetary theorists have also constructed models in which anticipated monetary shocks can have real effects. Overlapping contract models (Taylor, 1979), sticky price models (e.g., Rotemberg 1982, 1994; see also Blanchard, 1990, for a review with references), limited participation models (Grossman and Weiss, 1983; Rotemberg, 1984; more recently Alvarez and Atkeson, 1996) are prime examples. Even Lucas' (1972) model can generate effects of anticipated money if money is not injected by proportional transfer. In Lucas and Stokey's (1983) cash in advance model only anticipated monetary policy has real, inflation-tax, effects (though admittedly of the wrong sign). Cash-in-advance models with adjustment costs (Fuerst, 1992; Christiano and Eichenbaum, 1992, 1995) produce more traditional real effects of anticipated and unanticipated money.

For these reasons, it seems that an empiricist should at least consider the possibility that anticipated monetary changes can have real effects, rather than ruling out such a view a priori.

2. Models of the anticipated/unanticipated split

2.1. Overview

I specify two simple models for the relation between money and output that capture anticipated money effects. The first model allows expected and unexpected money to affect output. It is

$$y_t = a^*(L)[\lambda m_t + (1 - \lambda)(m_t - E_{t-1}m_t)] + b^*(L)\delta_t.$$

Asterisks on $a^*(L)$ and $b^*(L)$ denote structural lag polynomials. λ is a prespecified parameter that varies between 0 and 1. $b^*(L)\delta_t$ captures non-monetary output disturbances. As $\lambda \to 0$ this model specifies that only unanticipated money matters. As $\lambda \to 1$ there is no difference between anticipated and unanticipated money.

The second model is a slight modification of a standard sticky-price model (Rotemberg, 1982, 1994),

$$y_t = a^*(L)\left[m_t - \frac{1-\alpha}{1-\alpha L}E_{t-1}\left(\frac{1-\alpha\beta}{1-\alpha\beta L^{-1}}m_t\right)\right] + b^*(L)\delta_t.$$

β is a discount factor, slightly less than 1, and α between 0 and 1 measures the costs of price adjustment. The sticky price-model captures the idea that prices and hence output respond to expected future money. The distributed *lead* of future money following the expectation captures this effect, absent from the first model. As the price-stickiness parameter $\alpha \to 0$, this model also reduces to the unexpected money model ($\lambda = 0$), and as $\alpha \to 1$ it reduces to the mechanistic model ($\lambda = 1$). It gives a model in between these two extremes that is more complicated but somewhat more grounded in economic theory than the first model.

I assume that these relations between money and output are invariant to the policy regime; once we have estimated the parameters $a^*(L)$, we can calculate the response of output to arbitrary monetary experiments. Lucas (1976) argued that a good source – but not the only source – of policy-invariant relations are relations carefully grounded in economic theory. Both of the above models have some grounding in economic theory, but the dynamics $a^*(L)$ are ad-hoc. The models are intended as a flexible yet precise way of capturing easily interpretable anticipated–unanticipated distinctions, as found in many more formal models and discussions of monetary policy.

In the remainder of this section, I develop and motivate these models, and I show how to identify the parameters $a^*(L)$ from the impulse response function.

2.2. Anticipated–unanticipated model

2.2.1. Model

First, suppose that only unanticipated money affects output. The standard model (as, e.g., in Sargent's 1987 textbook) is a variant of Lucas (1973),

$$y_t = \theta(m_t - E_{t-1}m_t). \tag{1}$$

Of course, even this representation is not structural or policy invariant. (That is the point of Lucas' 1973 paper.) The parameter θ depends on the relative variance of aggregate and idiosyncratic price shocks. However, regime changes that do not alter this ratio will leave θ unchanged, so we can at least evaluate a limited set of regime changes.

This simple model does not allow for serially correlated output. Therefore, empirical specifications allow lagged effects of monetary shocks and serially correlated non-monetary output disturbances. Thus, I start with the model

$$y_t = a^*(L)[m_t - E_{t-1}m_t] + b^*(L)\delta_t. \tag{2}$$

I assume that the non-monetary shocks δ_t are orthogonal to monetary shocks $[m_t - E_{t-1}m_t]$.

Second, suppose there is no distinction between anticipated and unanticipated money, so our structural view is

$$y_t = a^*(L)m_t + b^*(L)\delta_t.$$

Typically one hopes or imposes that $a^*(1) = 0$, so that the level of money has no long-run effect on output.

One may again complain about micro-foundations. However, this is a model with a great historical tradition in empirical work, from the St. Lois Fed regressions (Anderson and Jordan, 1968) and its many antecedents to the dynamic multipliers calculated by Romer and Romer (1994). Most importantly, this model is implicit in any discussion that does not explicitly distinguish effects of anticipated vs. unanticipated monetary policy. Since almost no policy discussions make this distinction, even among academics, it seems worth interpreting the data with this view.

Finally, we want a model that assumes that anticipated money can have some effect, though unanticipated money might have stronger effects. As a way to capture this idea,[2] I assume values for λ in

$$y_t = a^*(L)[\lambda m_t + (1-\lambda)(m_t - E_{t-1}m_t)] + b^*(L)\delta_t. \tag{3}$$

2.2.2. *Identification*

Denote the joint moving average representation of output and money (or federal funds rate), as one might recover from a VAR *after* all the orthogonalization assumptions have been imposed, as follows.

$$\begin{bmatrix} m_t \\ y_t \end{bmatrix} = \begin{bmatrix} c_{mm}(L) & c_{my}(L) \\ c_{ym}(L) & c_{yy}(L) \end{bmatrix} \begin{bmatrix} \varepsilon_{mt} \\ \varepsilon_{yt} \end{bmatrix}; \qquad E\left(\begin{bmatrix} \varepsilon_{mt} \\ \varepsilon_{yt} \end{bmatrix} [\varepsilon_{mt} \quad \varepsilon_{yt}] \right) = I. \tag{4}$$

Every identification formula below goes through unchanged in larger VARs. I drop constants and possible time trends, so that y and m typically represent the detrended log of output and a monetary aggregate. I focus on the output

---

[2] Many other schemes are possible. One might try to allow separate dynamic effects of anticipated and unanticipated money,

$$y_t = a_a^*(L)m_t - a_u^*(L)(m_t - E_{t-1}m_t) + b^*(L)\delta_t.$$

Alas, a_a^* and a_u^* are not separately identified. (See Cochrane, 1994b, for a proof in this specific case; this is the point of Sargent, 1976, more generally.) The model in the text adds to this specification an ad-hoc assumption that shape of a_a^* and a_u^* are the same, so that so that $a_a^*(L) = \lambda/(1-\lambda)a_u^*(L)$. One parameter choice λ identifies the two polynomials.

Alternatively, motivated by information lags, we might join anticipated and unanticipated models by allowing only k period expectations to affect output,

$$y_t = a^*(L)[m_t - E_{t-k}m_t] + b^*(L)\delta_t.$$

After plugging in the moving average representation for money, we obtain

$$a^*(L) = [1 + c_{mm,1}L + \cdots + c_{mm,k-1}L^{k-1}]^{-1}c_{ym}(L).$$

As $k \to \infty$, we recover the anticipated money case, as $k = 1$, we recover the unanticipated money case. This model led to k period oscillations in $a^*(L)$, so I do not present the results. Models along this line must allow different agents to have different information lags, in order to avoid k-period dynamics.

response to monetary policy, so I do not question variable selection, specification, shock variable choice, orthogonalization, and sampling error questions that (rightly) pervade the VAR literature.

In order to identify $a^*(L)$, substitute the moving average representations for y_t and m_t from Eq. (4) into Eq. (3). Equating coefficients on ε_{mt}, we obtain

$$c_{ym}(L) = a^*(L)[\lambda c_{mm}(L) + (1-\lambda)c_{mm}(0)]. \tag{5}$$

The $b^*(L)$ are irrelevant; $a^*(L)$ is identified from the output and money responses $c_{ym}(L)$, $c_{mm}(L)$ alone. I match powers of L in Eq. (5) to find the $\{a_j^*\}$ from $\{c_{ym,j}\}$ and $\{c_{mm,j}\}$. Expanding and matching powers of L, we obtain

$$a_0^* = \frac{c_{ym,0}}{c_{mm,0}}; \qquad a_j^* = \frac{c_{ym,j} - \lambda\sum_{k=0}^{j-1} a_k^* c_{mm,j-k}}{c_{mm,0}} \quad j > 0.$$

In the special case that only unanticipated money matters, $\lambda = 0$, we do in fact recover the structural coefficients $a^*(L)$ from the impulse-response function. (One can also see this fact from Eq. (2) directly.) Eq. (5) simplifies to

$$a^*(L) = c_{ym}(L)/c_{mm}(0). \tag{6}$$

If there is no distinction between anticipated and unanticipated money, $\lambda = 1$, we use the dynamic response of output to a unit impulse to *money*, m_t rather than to the money *innovation*, ε_{mt}. Eq. (5) simplifies to

$$a^*(L) = c_{ym}(L)/c_{mm}(L). \tag{7}$$

In this case, the regression coefficients of y on m are invariant to policy, not the impulse-response function. For example, Romer and Romer (1994) analyze systematic policy with a view that there is no expected–unexpected distinction. Correctly, given this view, they present a dynamic simulation of the regression of y on lagged y and lagged m to measure the effects of monetary policy rather than an impulse-response function.

The Lucas model also makes predictions about prices. In its textbook form, it is derived from a relation between output and price surprises

$$y_t = \frac{\theta}{1-\theta}[p_t - E_{t-1}p_t]$$

together with

$$m_t = y_t + p_t.$$

I do not use measured price responses in the calculations that follow. My focus is on the effect of the expected–unexpected identifying assumption on the money to output relation. For this focus, I do not want to impose model price restrictions on the VARs, or further generalize the models so there are no price restrictions. Also, results that include prices are sensitive to the choice of index. (For example, see the difference between the commodity price and GDP deflator responses in Fig. 1.) I do not want price measurement issues to interfere with the interpretation of the money–output relation. Finally, there are different models that give the same quantity predictions but differ on price predictions. (For example, the Lucas and Woodford, 1994, model in which firms set their supply curve in advance also generates output effects based only on unanticipated money shocks, but with different price predictions than the Lucas, 1972, model.) However, for the purpose of evaluating a *given* model and its descriptions of alternative policies, one would want to look at all important variables including prices, and impose and test any restrictions.

2.3. Sticky price model

2.3.1. Model

A mechanistic relation between money and output, $y_t = a(L)m_t$ is not considered a serious anticipated money model by theorists, despite its continued popularity in policy discussions (the absence of any distinction between expected and unexpected policy actions) and empirical work. Instead, models in which prices are fixed in advance by nominal contracts or other frictions are a common and tractable framework that can generate anticipated money effects.

One must of course specify what prices are set *to* in advance. I use a slight modification of an explicit sticky price model due to Rotemberg, 1982, 1994. Its specification and central predictions are typical of many sticky price and nominal contract models. Rotemberg posits a detailed and rigorous microeconomic structure. However, his log-linearized first order conditions are the same as those of a representative price-setter who maximizes an objective function with a price adjustment cost term,

$$\max_{\{p_t\}} -\frac{1}{2}E\sum_t \beta^t[(1-\alpha)(1-\alpha\beta)(p_t - m_t)^2 + \alpha(p_t - p_{t-1})^2]. \tag{8}$$

α is a parameter between 0 and 1 that measures the costs of price adjustment. Scaling the first term by $1-\alpha\beta$ has no effect on the problem, but simplifies the following expressions. The ratio of level costs to price change costs is all that matters. That cost ratio, $\alpha/[(1-\alpha)(1-\alpha\beta)]$, goes monotonically from 0 to ∞ as α goes from 0 to 1. I modify Rotemberg's model slightly by requiring p_t to be set as a function of time $t-1$ information. In this way, the limit $\alpha \to 0$ leads to the unexpected money model rather than just $p_t = m_t$ and $y_t =$ constant.

The first order condition is

$$(1+\alpha^2\beta)p_t - \alpha p_{t-1} - \alpha\beta E_{t-1}p_{t+1} = (1-\alpha)(1-\alpha\beta)E_{t-1}m_t.$$

Factoring the lag polynomial and solving for prices, we obtain

$$p_t = \alpha p_{t-1} + (1-\alpha)(1-\alpha\beta)E_{t-1}\sum_{j=0}^{\infty}(\alpha\beta)^j m_{t+j}. \tag{9}$$

To model the relation between output and money, I add money demand $m_t = p_t + y_t$, leading to

$$y_t = m_t - p_t = \left[m_t - \frac{1-\alpha}{1-\alpha L}E_{t-1}\frac{1-\alpha\beta}{1-\alpha\beta L^{-1}}m_t\right].$$

This simple structure captures effects common to most sticky price and overlapping contract models. Prices can rise and hence output can decline in response to *expected future* monetary increases. Prices adapt slowly to monetary shocks. The model imposes long-run neutrality: $y = 0$ for any constant level of m.

Above, we added ad-hoc dynamics and real shocks to the Lucas model so that it could be used to interpret the impulse-response function from a VAR. The natural analogy is to add ad-hoc dynamics to the sticky price model, giving[3]

$$y_t = a^*(L)\left[m_t - \frac{1-\alpha}{1-\alpha L}E_{t-1}\left(\frac{1-\alpha\beta}{1-\alpha\beta L^{-1}}m_t\right)\right] + b^*(L)\delta_t. \tag{10}$$

We can think of this model as allowing velocity to respond to monetary shocks, i.e. $y_t = a^*(L)(m_t - p_t)$.

2.3.2. Identification

I prespecify α and β. β is a discount factor, which I specify as $\beta = 0.98$. The results are insensitive to β so long as it is nearer to 1 than to 0. I vary the

[3] I investigated an alternative way to enrich the dynamics, by adding higher derivatives to the objective function,

$$-\frac{1}{2}E\sum_t \beta^t\left[(p_t - m_t)^2 + \sum_{k=1}^{r}\eta_k(p_t - p_{t-k})^2\right].$$

The flexibility in $\eta(L)$ allows the model to produce more complex money and output impulse-response functions. However, all impulse-response functions are not possible, so this model still imposes restrictions on $c_{ym}(L)$ and $c_{mm}(L)$ that are not satisfied by the point estimates of the VARs investigated below. One could estimate this model by maximum likelihood, imposing its restrictions on the data. However, the costs of higher derivatives of price change do not seem much more compelling than the $a^*(L)$ polynomial as a policy-invariant description of technology.

It is tempting to generalize the model Eq. (10) by simply allowing general lagged effects of a distributed lead of future money,

$$a^*(L)y_t = E_t d^*(L^{-1})m_t.$$

This form captures the majority of sticky price or nominal contract models. However, one cannot separately identify a^* and b^*. Furthermore, the point of the sticky price model is that the forward and backward lag polynomials are linked by common 'technological' parameters.

parameter α from 0 to 1 and track the results, as we examined the sensitivity to the parameter λ in the anticipated–unanticipated model Eq. (3). We can then identify $a^*(L)$ from the impulse-response function.

Plugging the VAR responses $c_{mm}(L)$, $c_{ym}(L)$ to money shocks into Eq. (10), we obtain[4]

$$\begin{aligned}(1-\alpha L)(1-\alpha\beta L^{-1})c_{ym}(L) \\ = a^*(L)[\alpha(1-L)(1-\beta L^{-1})c_{mm}(L) + (1-\alpha)(1-\alpha\beta)c_{mm}(\alpha\beta)]\end{aligned} \tag{11}$$

As $\alpha \to 0$, Eq. (11) recovers the impulse-response function as in the $\lambda = 0$ case, Eq. (6). As $\alpha \to 1$, we recover the same dynamic multiplier of the $\lambda = 1$ case, Eq. (7). In the latter case, the firm sets price p to a constant. For $\alpha \in (0,1)$, I solve Eq. (11) recursively for the $\{a_j^*\}$ from values for $\{c_{mm,j}, c_{ym,j}\}$ by matching the coefficients on each power of L. To calculate the output response to monetary experiments, I first find the price path. I solve Eq. (9) for prices given a path of money[5] and I then find output from

$$y_t = a^*(L)(m_t - p_t).$$

This model also makes predictions for prices. As above, I focus on the relationship between money and output, ignoring extra restrictions that come from price data.

---

[4] We obtain directly

$$c_{ym}(L)\varepsilon_{mt} = a^*(L)\left[1 - \frac{1-\alpha}{1-\alpha L}E_{t-1}\frac{1-\alpha\beta}{1-\alpha\beta L^{-1}}\right]c_{mm}(L)\varepsilon_{mt}.$$

Now,

$$E_{t-1}\frac{1}{1-\alpha\beta L^{-1}}c_{mm}(L)\varepsilon_{mt} = \frac{c_{mm}(L) - c_{mm}(\alpha\beta)}{1-\alpha\beta L^{-1}}\varepsilon_{mt}.$$

(You can check directly that the $c_{mm}(\alpha\beta)$ terms removes the current and future ε_{mt}. See Sargent, 1987, p. 304.) Hence,

$$c_{ym}(L) = a^*(L)\left\{c_{mm}(L) - \left(\frac{1-\alpha}{1-\alpha L}\right)\left(\frac{1-\alpha\beta}{1-\alpha\beta L^{-1}}\right)[c_{mm}(L) - c_{mm}(\alpha\beta)]\right\}.$$

Eq. (11) follows by multiplying both sides by $(1-\alpha L)(1-\alpha\beta L^{-1})$ and simplifying the coefficient on $c_{mm}(L)$.

[5] I find p_1 from the requirement that prices not explode as $t \to -\infty$, and then find other prices recursively. The solutions are

1. Anticipated step, $t \leq 1{:}\ p_t = \frac{1-\alpha}{1-\alpha^2\beta}(\alpha\beta)^{1-t}$; $t \geq 1{:}\ p_t = 1 - \left(\frac{1-\alpha\beta}{1-\alpha^2\beta}\right)\alpha^t$.
2. Unanticipated step, $t \leq 0{:}\ p_t = 0$; $t \geq 1{:}\ p_t = 1 - \alpha^{t-1}$.
3. Unanticipated blip, $p_t = 0$.
4. Anticipated blip, $t \leq 1{:}\ p_t = \frac{(1-\alpha)(1-\alpha\beta)}{1-\alpha^2\beta}(\alpha\beta)^{1-t}$; $t > 1{:}\ p_t = \frac{(1-\alpha)(1-\alpha\beta)}{1-\alpha^2\beta}\alpha^t$

3. VAR results

3.1. M2 VAR, anticipated–unanticipated model

Fig. 1 plots the responses to the money shock for the M2 VAR, Fig. 2 collects the money and output responses to a money shock. The appendix describes the M2 VAR in detail.

The top panel of Fig. 3 presents the lag polynomials $a^*(L)$ that we estimate assuming that various fractions of anticipated and unanticipated money affect output. This is also the effect on output of a particular monetary intervention, a unit innovation in money that lasts one period, or an 'unanticipated blip'.

The effects of this experiment are the same across identifying assumptions for the first quarter. But then the output effect calculated with the unanticipated-money assumption increases dramatically, and decays back to zero very slowly. This effect is, of course, exactly the money to output impulse-response function plotted in Fig. 1 and Fig. 2.

The output effect calculated with no anticipated/unanticipated distinction is very small, and reverts to zero after three quarters. If this were the true model, the impulse-response function would give a *very* misleading guide to the effects of this monetary intervention! Furthermore, one only needs to assume that anticipated money matters slightly in order to identify short-lived effects of this monetary intervention. The $\lambda = 0.2$ assumption produces a response pattern much closer to the completely anticipated ($\lambda = 1$) assumption than to the completely unanticipated ($\lambda = 0$) assumption.

Of course, all the models (values of λ) give identical output responses to a money innovation that is followed by further increases in money prescribed by the money $\rightarrow$ money response function. They differ on the question, 'What if the Fed shocks money, and then follows a path of future money *different* from the historical pattern?'

An *announced* blip in money is another interesting monetary experiment, since any policy can be decomposed into a sum of expected and unexpected blips. The bottom panel of Fig. 3 presents calculations of the response to an announced blip. Now, the anticipated money model response is the largest (least small), and the same as in the top panel. As unanticipated money matters more, the response gets smaller. If only unanticipated money matters, there is of course no response to this anticipated blip.

In every case, the output responses one sees are much *shorter* and *smaller* than the impulse-response function if one allows anticipated money to even have some affect output. The theoretical identification is quantitatively and economically important; this is not a minor issue of forgotten 1970s methodological debates.

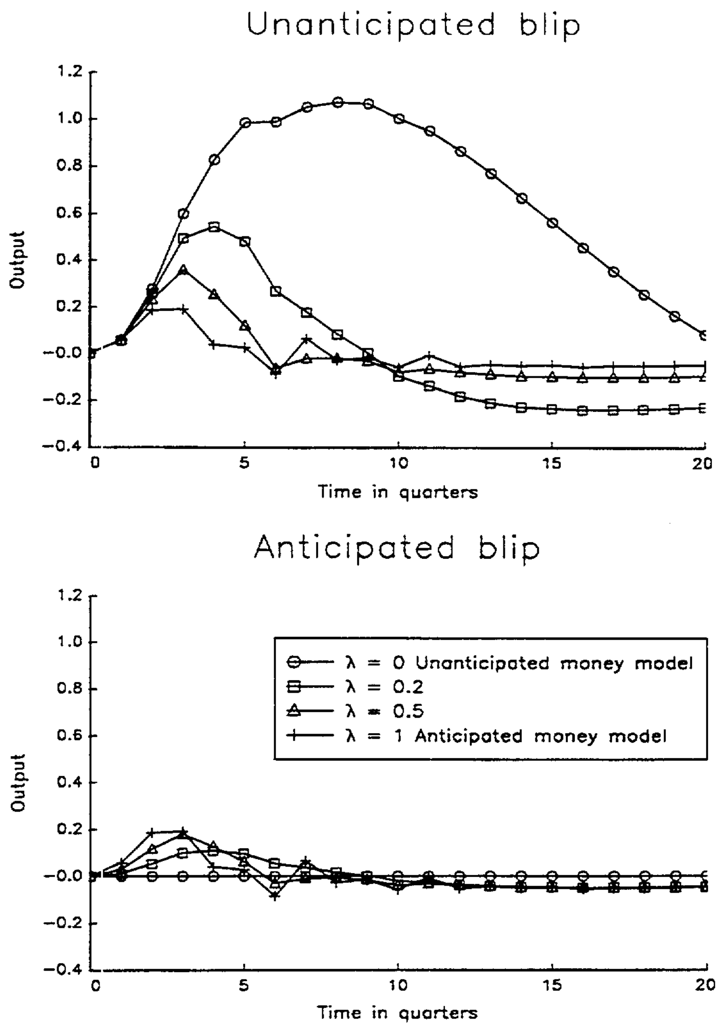

Fig. 3. Output effects of two monetary experiments, under various assumptions about the effects of anticipated vs. unanticipated money. Calculated from M2 VAR.

3.2. M2 VAR, sticky price model

Fig. 4 presents the lag polynomial $a^*(L)$ inferred from the impulse response function with the sticky price model, via Eq. (11), for several values of the price stickiness parameter, α. Again, $\alpha = 0$ recovers the impulse-response function, and $\alpha = 1$ recovers the same response as the mechanistic model with $\lambda = 1$. Intermediate values of α give intermediate results.

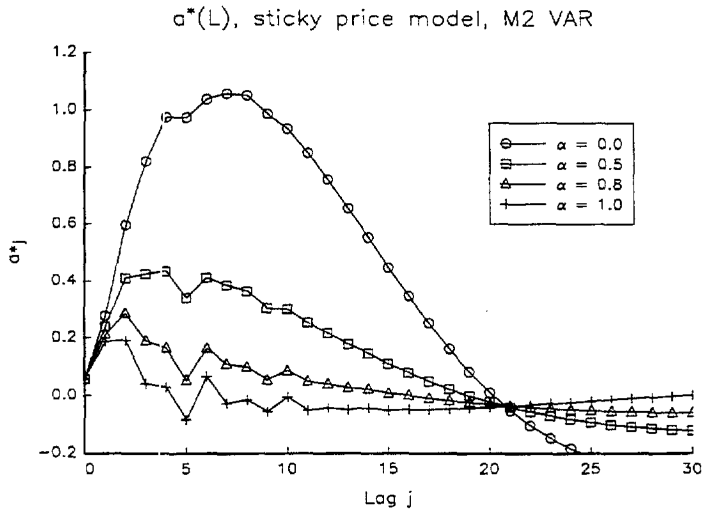

Fig. 4. Lag polynomial $a^*(L)$ inferred from M2 impulse-response function using sticky price model. $\alpha =$ price stickiness parameter.

For a low value of price stickiness α, we need a drawn-out lag polynomial $a^*(L)$ in order to match the VAR impulse-response function. For higher values of α the lag polynomial $a^*(L)$ becomes much smaller and shorter. More of the dynamic relation between money and output can be accounted for by the stickiness in the model, requiring less work from the ad-hoc lags $a^*(L)$. As Rotemberg (1994) argues, this model provides a reasonable account of the data even with no ad-hoc lags, $a^*(L) = 1$, so long as one imposes a high value of price stickiness, above 0.8 or 0.9.

One must question whether such a high value of the price stickiness parameter is reasonable. $\alpha = 0.8$ implies that the relative costs of price change and price level deviations is $\alpha/[(1-\alpha)(1-\alpha\beta)] = 17$ (see the firm's problem, Eq. (8)). A 1% quarterly price increase has the same costs as prices 17% away from their proper level. Rotemberg, 1982, 1994 argues that these relative costs are not so unreasonable, in that the costs of price-level deviations may be low. He posits a model of disaggregated monopolistically competitive firms, in which the costs of having a price level different from the industry average are high, but the industry average can deviate from m_t by a great deal with little cost.

Fig. 5 presents the calculated responses to anticipated and unanticipated steps and blips in money for price stickiness $\alpha = 0.8$. The $\alpha = 0$ and $\alpha = 1$ limits are the same as the anticipated $\lambda = 0$ and unanticipated $\lambda = 1$ models investigated above. $\alpha = 0.8$ makes an interesting intermediate case. I include the price paths, since they make clear how the output responses are generated.

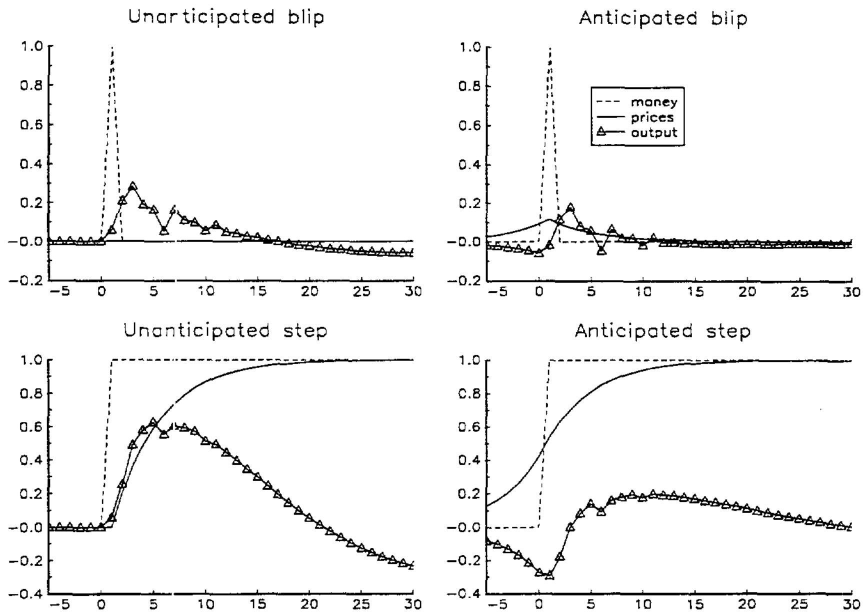

Fig. 5. Output and price responses to monetary experiments in sticky price model, using m2 VAR impulse-response function. Stickiness parameter $\alpha = 0.8$.

Start with the top left corner, an unanticipated blip in money. Prices do not move at all in this case. They cannot move until one period after the blip is known to have happened. But by this time, money has returned to its original level and is expected to remain there. Since $y_t = a^*(L)(m_t - p_t)$, output follows this one-period impulse $m_t - p_t$ with the lag pattern $a^*(L)$. Since the lags $a^*(L)$ are small and short, the model predicts a small and short output response.

Notice the counterintuitive result: Assuming *stickier* prices implies *shorter and smaller* effects of these monetary interventions than is given by the impulse-response functions. The assumption of stickier prices forces us to estimate a much smaller and shorter dynamic relation $a^*(L)$ to remain consistent to the VAR representation of the data.

Next look at the effects of an *an*ticipated money blip in the top right corner. Now prices rise somewhat in advance of the blip, but not much since the blip is expected to last only one period and it is costly to raise prices. Therefore, the output response is similar to, but slightly smaller than the pattern found for the unanticipated blip.

The bottom row, left, presents an unanticipated step in money. After the money innovation, prices slowly increase. The output response looks almost like the impulse-response function. It should, since the unanticipated step is similar to the actual policy regime. The mechanism by which this model reproduces a drawn-out response is quite different from the unanticipated money model however. The difference between m and p is now strung out over many quarters. $y_t = a^*(L)(m_t - p_t)$ thus gives a drawn out effect of this many-period impulse through the relatively short structural lag $a^*(L)$.

When the money step is anticipated (bottom row, right), prices rise in anticipation of the monetary expansion. The price rise is slow and smooth, since the price stickiness parameter is large. The fact that β is near one makes the price path nearly symmetrical. Output *declines* in advance of the money step, and recovers when the increase in money actually occurs. The major effect of an anticipated step in money is a *depressing* effect on output!

Comparing across columns, it is clear that models with price stickiness cannot be used to justify mechanical relations between output and money that do not distinguish anticipated and unanticipated effects. While a mechanical relation is true for *infinite* stickiness ($\alpha = 1$ and constant prices) even with stickiness $\alpha = 0.8$ or price change costs 17 times the cost of price level errors, one obtains dramatically different predictions for an anticipated vs. unanticipated money step. The reason is simple and general: intertemporally optimizing agents faced with some friction will attempt to adjust *before* an expected change in policy. In fact, the larger the frictions one presents to economic agents, the farther ahead they will start to adjust to the expected change in policy.

3.3. Federal funds VAR

The monetary VAR literature has recently tried on variables other than monetary aggregates as indicators of a change in monetary policy. Bernanke and Blinder (1988) first used the federal funds rate. Their idea is that the Fed controls the funds rate on a day-to-day basis, accommodating shifts in money demand. A shift in money supply is seen when the Fed changes the federal funds rate.

The bottom row of Fig. 1 plots responses to federal funds shocks. Following Christiano, Eichenbaum and Evans (1996) I orthogonalize federal funds last and include a commodity price index in the VAR. Federal funds VARs without this feature produce a 'price puzzle' – increases in the federal funds rate produce *increases* in prices. Christiano, Eichenbaum and Evans explain that the Fed contracts on news of future inflation, inducing a spurious correlation between contractionary shocks and subsequent price rises. By including a commodity price index and orthogonalizing federal funds last, they control for the Fed's information. More pragmatically, these choices produce better-looking pictures. The appendix details the specification and estimation of the VAR.

Fig. 6 presents output responses to federal funds shocks, again imposing that various fractions of anticipated and unanticipated shocks can affect output. We see the usual pattern that responses calculated assuming that anticipated money

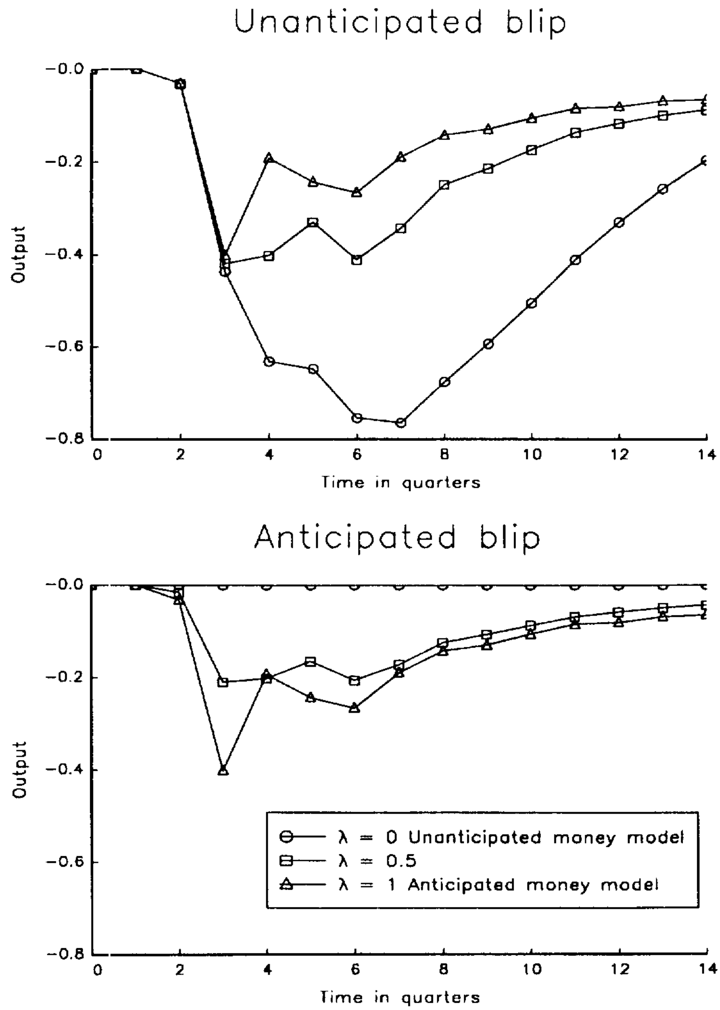

Fig. 6. Output effects of two monetary experiments, under various assumptions about the effects of anticipated vs. unanticipated money. Calculated from federal funds VAR.

matters or that the shock is anticipated are much smaller and shorter than the impulse-response function suggests. The anticipated-money effects of the shock peak in three quarters, while the unanticipated-money effects peak in 7 quarters and then die off slowly.

I do not present calculations for the sticky price model because it is unclear how to apply that model to something other than an actual money aggregate. (Even in the unanticipated model, it is unclear what a federal funds shock means.

If only unanticipated monetary policy has non-neutral effects – the identifying assumption behind impulse-response functions – what does the Fed reaction function equation of the VAR mean?)

4. Interpretation

4.1. Which view is right?

As we have seen, the anticipated/unanticipated identifying assumption makes a large difference to the results. Which identifying assumption is correct? Again, the data (from one regime) cannot tell us. However, we can evaluate the plausibility of the results. First, in the unanticipated money view, the striking similarity of the output and money responses is a pure and unlikely coincidence, since the output response would be the same for any money response. In the anticipated money view, the output response has its large, protracted and hump shape precisely *because* the money response has the same shape. Less persistent monetary policies would give rise to less persistent output responses. One may dislike theories that require such coincidences in the data. Second, the unanticipated money view requires that one construct a theory of the long delayed and hump-shaped response, or swallow these dynamics as an ad-hoc patch. The anticipated money views result in estimates of a short, almost contemporaneous effect, of the sort generated by most current monetary theories. Both of these considerations suggest (but of course cannot prove) that there is at least some role for anticipated, systematic monetary policy.

4.2. Implications for monetary theory

Eventually, we want formal theoretical models, specified at the level of policy-invariant tastes, technology, monetary and real frictions, rather than anticipated/unanticipated distinctions. The models should explain the VAR evidence, and they should be consistent with at least the broad brush of experiences with different regimes.

The 'theories' I imposed on the data are a long way from this ideal. The basic unanticipated money model $y_t = a^*(L)(m_t - E_{t-1}m_t)$ does best on consistency across regimes. But the heart of the model for explaining postwar US time series is the long, large, and hump-shaped distributed lag $a^*(L)$. That lag is questionably policy-invariant; distrust of this lag polynomial is exactly my motivation for looking at anticipated money models. The mechanistic model $y_t = a^*(L)m_t$ has of course been subject to a generation of derision. The sticky price model I examine, like most anticipated-money models, is also questionably policy-invariant. Imagine using the Rotemberg (1982, 1994) sticky price model, Taylor's (1979) overlapping contract model or a cash in advance model with frictions such as Christiano and Eichenbaum (1992) to understand a hyperinflation. In that regime, contract lengths are optimally chosen to be much shorter, one would doubtless find almost no price stickiness, and the timing conventions and portfolio frictions of a CIA model would surely disappear. Obviously, the contract length, price stickiness, timing conventions and portfolio frictions are not really fixed, policy invariant features of technology; one might also expect them to change in response to less drastic changes in regime. Hence the 'deep' parameters the model needs to specify, and the poor empiricist has to try to estimate, relate to the technology of contracting, or the costs of adopting different timing conventions or portfolio habits. Such models are not yet written, let alone ready to estimate.

Thus, existing explicit monetary theories, specified at the level of tastes, technology and frictions, are also a long way from the ideal as well. The results of this paper are, I hope, a useful data summary for the construction of better explicit models. The impulse responses naturally suggest that one add amplification and persistence mechanisms to an unanticipated money model. For example, Christiano and Eichenbaum (1992) add portfolio adjustment costs to a liquidity effect cash in advance model, in order to produce an extended response function. The results suggest an alternative path: if one constructs models with anticipated money effects, they will not require such extensive amplification and persistence mechanisms.

4.3. Policy advice

Empirical work is used for policy advice as well as in guiding the theoretical development of monetary economics. The Fed (and its critics) basically want to know from academics, 'what will happen if we raise the funds rate today?' This is regarded as an essentially empirical question, and one to which economists should be able to give some answer beyond 'wait until we finish constructing explicit dynamic general equilibrium monetary models with frictions'. However, this is not a purely empirical question, and the theoretical identification one brings to the issue, even at the simplistic level of distinguishing anticipated from unanticipated effects that I pursue here, has a large impact on the answer one gives.

The unanticipated money assumption results in one set of answers: If the funds rate increase is anticipated – if this is the normal time or set of events for a funds rate increase – there will be no output effect. If the funds rate increase is not anticipated, there will be a long, large, and hump-shaped effect on output as given by the impulse response function. What subsequent actions the Fed takes are irrelevant to the answer.

The anticipated money assumption results in quite different answers: Policy 'actions' – no longer only surprises – result in short and small output responses. A long-lasting and large output rise can only occur if the Fed continues expansionary policy actions in the future, as it typically has in the past.

Two opinions run through much debate over the Fed's interest rate policy: (1) A federal funds rate rise now causes output declines that peak in one and a half to two years; (2) It does not matter that most federal funds rate changes are part of a slow and predictable tightening or loosening, i.e. federal funds rate changes are largely anticipated. At a minimum, the results in Fig. 6 shows that these two views are inconsistent, in that opposite theoretical assumptions must be made for the two conclusions. If the anticipated gradual tightenings have *any* effect, they have the short and small effects graphed for $\lambda = 1$. If one believes in a large and 1.5–2 year delayed effect, then only surprise movements in federal funds will have any effect at all.

5. Summary and conclusion

This paper elaborates two stories for the large and drawn out output responses we see in VARs. Either (1) the large and drawn out output responses really do represent large and drawn out effects of the initial shock, effects that would be present no matter what the path of monetary policy after the shock, or (2) the large and drawn out output responses occur *because* monetary policy also has a large and drawn out response to the initial shock. If monetary policy were less persistent, the output response functions would be less persistent as well. However, this view requires that anticipated monetary policy can affect output.

The sensitivity of the output effect calculations to the anticipated/unanticipated identifying assumption is a strong, quantitative reminder of the methodological warning Sargent (1976) gave more than 20 years ago. Estimates of the effects of monetary policy, and calculations such as 'how would output have behaved if the Fed followed a different rule?' such as nominal GNP targeting, are almost entirely driven by the *theoretical* assumptions one makes. The data alone do not answer these questions.

The results one obtain by imposing different views of the effects of anticipated money are subtle. Sticky price models do not justify a view that one can ignore the distinction between anticipated and unanticipated components of monetary policy. Furthermore, the estimated size and persistence of output responses may be smaller when one assumes stickier prices, because one estimates a shorter dynamic relation when anticipated money effects are allowed.

Finally, the output responses calculated assuming that anticipated money can affect output are much smaller and shorter than those recovered from the impulse-response function. If your priors are in favor of small short effects of monetary policy, or models with less reliance on ad-hoc dynamics, this fact may suggest to you that anticipated money and systematic monetary policy can in fact affect output.

Acknowledgements

I thank Larry Christiano, Tim Cogley, Martin Eichenbaum, Lars Hansen, Anil Kashyap, Charles Plosser, Tom Sargent, Ken West and especially Mark Watson, Mike Woodford and an anonymous referee for helpful comments, and I thank Alexander Reyfman for research assistance. This research was supported by a grant from the NSF administered by the NBER, and by the University of Chicago Graduate School of Business.

Appendix A.

This appendix summarizes the specification of the M2 and Federal Funds VAR presented in Fig. 1 and the basis of all the other calculations. For further details, including the long specification search that led to these 'best case' VARs, standard errors, and so forth, see Cochrane (1994a). The data and all programs used in this paper are available via anonymous ftp from ftp-finance-gsb.uchicago.edu, in /pub/cochrane/vars.

Both VARs use quarterly data, 59:1-92:4. The data used in the M2 VAR are, with citibase codes in parentheses, m2 = log M2 (FM2); ff=federal funds rate (FYFF), quarterly average; cns=log of nondurables+services consumption (GCNQ + GCSQ) + 0.65 times government purchases (GGEQ); y = log GDP (GDPQ); p=log GDP deflator (GDPD). I run the M2 VAR with two cointegrating vectors, m2-y-p and y-c. I run the VAR in error correction form with four lags. For example the first line is

$$\Delta m2_{t+1} = a_0 + a_1(m2_t - y_t - p_t) + a_2(y_t - cns_t) + \sum_{j=1}^{4} b_j \Delta m2_{t-j} + \sum_{j=1}^{4} c_j ff_{t-j} + \sum_{j=1}^{4} d_j \Delta cns_{t-j} + \sum_{j=1}^{4} e_j \Delta y_{t-j} + \sum_{j=1}^{4} e_j \Delta p_{t-j} + \varepsilon_{t-1}.$$

The errors are Choleski orthogonalized in the given order. Equivalently, no current variables occur in the $\Delta m2$ equation, current $\Delta m2$ appears in the ff equation and so forth. The impulse response function graphs the impled response of the log *levels* of all variables except ff.

The variables used in the ff VAR are c = log consumption of nondurables + services (GCNQ + GCSQ), y = log private GDP, or GDP less government purchases (GDPQ − GGEQ), p = log GDP deflator (GDPD), cp = log sensitive commodities price index used in Christiano et al. (1995), provided by Charles Evans, m1 = log of M1 (FM1) and ff=quarterly average of federal funds rate (FYFF). I run the VAR in log levels with four lags. The shocks are orthogonalized in the given order.

Including government purchases as above helps to define a stationary consumption/output ratio and hence improves the long-run properties of the VARs. Government purchases have increased as a share of GDP over the postwar period. Assigning 0.65 of government purchases to consumption, or taking the ratio of consumption to private GDP results in a more stationary consumption/output ratio. This is a refinement, the results are not drastically changed if one omits the government purchases.

References

Alvarez, F., Atkeson, A., 1996. Liquidity, inflation, interest rates and exchange rates. Manuscript, University of Pennsylvania.

Anderson, L.C., Jordan, J.L., 1968. Monetary and fiscal actions: a test of their relative importance in economic stabilization. Federal Bank of St. Louis Review 50, 11–23.

Bernanke, B.S., Blinder, A.S., 1988. The federal funds rate and the channels of monetary transmission. American Economic Review 82, 901–921.

Blanchard, O.J., 1990. Why does money affect output: a survey. In: Friedman, B.M., Hahn, F.H., (Eds.), Handbook of Monetary Economics. vol. II, Elsevier, Amsterdam, pp. 779–835.

Christiano, L.J., Eichenbaum, M., 1992. Liquidity effects and the monetary transmission mechanism. American Economic Review 82, 346–353.

Christiano, L.J., Eichenbaum, M., 1995. Liquidity effects, monetary policy and the business cycle. Journal of Money, Credit and Banking 27, Part I, 1113–1136.

Christiano, L.J., Eichenbaum, M., Evans, C., 1996. The effects of monetary policy shocks: some evidence from the flow of funds. Review of Economics and Statistics 78, 16–34.

Cochrane, J.H., 1994a. Shocks. Carnegie–Rochester Conference Series on Public Policy 41, 295–364.

Cochrane, J.H., 1994b. Comment on Romer and Romer. In: Fischer, S., Rotemberg, J.J., (Eds.), 1994 NBER Macroeconomics Annual. MIT Press, Cambridge, pp. 58–74.

Eichenbaum, M., 1992. Comments on 'Interpreting the macroeconomic time series facts: the effects of monetary policy' by Christopher Sims. European Economic Review 36, 1001–1011.

Feldstein, M., Stock, J.H., 1994. The use of a monetary aggregate to target nominal GDP. In: Gregory Mankiw, N., (Ed.), Monetary Policy. University of Chicago Press, Chicago.

Friedman, M., 1968. The role of monetary policy. American Economic Review 58, 1–17.

Fuerst, T., 1992. Liquidity, loanable funds and real activity. Journal of Monetary Economics 29, 3–24.

Gordon, D.B., Leeper, E.M., 1993. The dynamic impacts of monetary policy: an exercise in tentative identification. Journal of Political Economy 102, 1228–1247.

Grossman, S., Weiss, L., 1983. A transactions-based model of the monetary transmission mechanism. American Economic Review 73, 871–880.

Hall, R.E., Gregory Mankiw, N., 1994. Nominal Income Targeting. In: Gregory Mankiw, N., (Ed.), Monetary Policy. University of Chicago Press, Chicago.

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

Lucas, R.E. Jr., 1973. Some international evidence on output–inflation trade-offs. American Economic Review 63, 326–334.

Lucas, R.E. Jr., 1976. Econometric policy evaluation: a critique. Carnegie–Rochester Conference Series on Public Policy 1, 19–46.

Lucas, R.E. Jr., 1996. Monetary neutrality. Nobel prize lecture manuscript, University of Chicago.
Lucas, R.E. Jr., Stokey, N.L., 1983. Money and interest in a cash-in-advance economy. Econometrica 55, 491–513.
Lucas, R.E. Jr., Woodford, M., 1994. Real effects of monetary shocks in an economy with sequential purchases. Manuscript, University of Chicago.
Romer, D.H., Romer, C.D., 1994. What ends recessions? In: Fischer, S., Rotemberg, J.J., (Eds.), 1994 NBER Macroeconomics Annual. MIT Press, Cambridge, pp. 13–57.
Rotemberg, J.J., 1982. Sticky prices in the United States. Journal of Political Economy 90, 1187–1211.
Rotemberg, J.J., 1984. A monetary equilibrium with transactions costs. Journal of Political Economy 92, 40–58.
Rotemberg, J.J., 1994. Prices and output: an empirical analysis based on a sticky price model. Working paper, Massachusetts Institute of Technology.
Sargent, T.J., 1976. Observational equivalence of natural and unnatural rate theories of macroeconomics. Journal of Political Economy 84.
Sargent, T.J., 1986. Rational Expectations and Inflation. Harper and Row, New York.
Sargent, T.J., 1987. Macroeconomic Theory. Academic Press, Orlando.
Sargent, T.J., Wallace, N., 1975. Rational expectations, the optimal monetary instrument and the optimal money supply rule. Journal of Political Economy 83, 241–254.
Sims, C.A., 1992. Interpreting the macroeconomic time series facts: the effects of monetary policy. European Economic Review 36, 975–1000.
Strongin, S., 1992. The identification of monetary policy disturbances: explaining the liquidity puzzle. Federal Reserve Bank of Chicago Working Paper WP-92-27.
Taylor, J.B., 1979. Estimation and control of a macroeconomic model with rational expectations. Econometrica 47, 1267–1286.