TECHNICAL WORKING PAPER SERIES

AN OPTIMIZATION-BASED ECONOMETRIC FRAMEWORK FOR THE EVALUATION OF MONETARY POLICY: EXPANDED VERSION

Julio J. Rotemberg
Michael Woodford

Technical Working Paper 233
http://www.nber.org/papers/T0233

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 1998

We wish to thank Eduardo Loyo and Marc Giannoni for superb research assistance, Ben Bernanke, Gregory Chow, Frank Diebold, Jeff Fuhrer, Bennett McCallum, and Chris Sims for helpful comments, and the NSF for research support through a grant to the National Bureau of Economic Research. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

An Optimization-Based Econometric Framework
for the Evaluation of Monetary Policy: Expanded Version
Julio J. Rotemberg and Michael Woodford
NBER Technical Working Paper No. 233
May 1998
JEL No. E37

ABSTRACT

This paper considers a simple quantitative model of output, interest rate and inflation determination in the United States, and uses it to evaluate alternative rules by which the Fed may set interest rates. The model is derived from optimizing behavior under rational expectations, both on the part of the purchasers of goods and upon that of the sellers. The model matches the estimated responses to a monetary policy shock quite well and, once due account is taken of other disturbances, can account for our data nearly as well as an unrestricted VAR. The monetary policy rule that most reduces inflation variability (and is best on this account) requires very variable interest rates, which in turn is possible only in the case of a high average inflation rate. But even in the case of a constrained-optimal policy, that takes into account some of the costs of average inflation and constrains the variability of interest rates so as to keep average inflation low, inflation would be stabilized considerably more and output stabilized considerably less than under our estimates of current policy. Moreover, this constrained-optimal policy also allows average inflation to be much smaller.

This version contains additional details of our derivations and calculations, including three technical appendices, not included in the version published in *NBER Macroeconomics Annual 1997*.

Julio J. Rotemberg
Morgan Hall
Harvard Business School
Soldiers Field Road
Boston, MA 02163
and NBER
jrotemberg@hbs.edu

Michael Woodford
Department of Economics
Princeton University
Princeton, NJ 08544
and NBER
woodford@princeton.edu

1 Introduction

In this paper, we develop a small, structural econometric model to be used in the quantitative evaluation of proposed rules for monetary policy. The quantitative evaluation of monetary policy rules has gained increased attention in recent years (see, *e.g.*, McCallum, 1988, 1997; Taylor, 1993a; Bryant *et al.*, 1993; Henderson and McKibbin, 1993; Feldstein and Stock, 1994; Leeper and Sims, 1994; Levin, 1996 and Fuhrer 1997a). Our approach differs from these studies in that we derive our econometric specification from an explicit model of intertemporal optimization on the part of both the suppliers and the purchasers of goods and services. [1]

Rigorously grounding our structural relations in optimizing individual behavior has two important advantages. The first is that we are able to respond to the well-known Lucas (1976) critique of econometric policy evaluation. Our analysis of hypothetical policy rules takes full account of the way that an understanding of the change in policy regime ought to affect the decision rules of private agents, and make them different than those that underlie the statistical correlations observed in past data. Much of the recent work cited above does respond to the Lucas critique to at least some extent, by incorporating forward-looking specifications of at least some of the models' structural relations, and assuming rational (or model-consistent) expectations in analyzing alternative policies. But because no attempt is made to derive the complete structural model from internally consistent foundations in terms of individual optimization, doubts must remain as to whether the posited structural relations should genuinely be invariant to changes in the policy regime.

Demanding that one's structural relations be derived from individual optimization also has the advantage that evidence from other sources about the nature of the problems that individuals face can be used to corroborate the quantitative specifications that are used to explain the relations among aggregate time series. Ultimately, this is the only way in which the "observational equivalence" of a multitude of alternative possible structural interpretations of the co-movements of aggregate series can be resolved. We make little attempt here at such validation of our proposed specification. But because our model parameters refer to things like the elasticity of firms' demand curves, or the average length of time that prices remain fixed, that have clear referents apart from the role of these parameters in our structural relations linking interest rates,

[1] Leeper and Sims (1996) represent an ambitious effort of this kind, although, unlike us, they do not derive their price dynamics from producer optimization. Ireland (1996) is another recent study with an aim similar to ours, although our approaches differ considerably in their details.

inflation, and output in the economy as a whole, it becomes possible to consider the reasonableness of our specification on grounds other than simple statistical measures of "goodness of fit".

The second advantage of an optimization-based approach is that the specification of individuals' decision problems that is used to explain the effects of monetary policy can also be used for purposes of welfare analysis. Of course, the analysis of the deadweight losses associated with alternative policies in terms of the individual preferences that account for the predicted responses to a policy change is by now the standard method of the public finance literature. But this method has been little applied to problems of monetary policy, the main exception being analyses of the special issue of the costs of steady inflation (*e.g.*, Lucas, 1993). Analyses of optimal monetary policy – or at least those that are based upon econometric models – consider instead the problem of minimizing one *ad hoc* loss function or another. Here we show, instead, how a utility-based measure of the deadweight loss associated with price-level instability can be derived, and how most of its parameters can be determined from our estimated structural equations. We can use this measure to address questions such as Summers' (1991) suggestion that a positive average rate of inflation is desirable in order to make it possible for nominal interest rates to be lowered as necessary for stabilization purposes. Not only does our econometric model allow us to discuss quantitatively how much variability of nominal interest rates would be necessary for full stabilization, but our welfare measure in principle allows a direct comparison, in comparable units, of the deadweight losses associated with incomplete stabilization on the one hand and higher average inflation on the other.

Our analysis proceeds in five distinct steps. In the first step, we estimate a vector autoregression (VAR) model of the joint process of interest rates, inflation and output. We use this VAR for two purposes. The first is to identify the actual monetary policy rule employed by the Fed. Following Taylor (1993b), we suppose that this rule is a reaction function that sets interest rates as a function of current and past values of output and inflation. The second purpose of this VAR is to estimate the way output, inflation and interest rates respond to a stochastic disturbance to the monetary policy rule. Thus we learn how the economy responds to a monetary shock under the current monetary policy rule, using fairly standard "structural VAR" methodology (*e.g.*, Cochrane, 1995; Sims, Leeper, and Zha, 1996).

In the second step, we postulate a simple theoretical model that can account for the estimated response of output and inflation to monetary policy shocks. In this model, we assume that there are impediments to the free adjustment of prices. In particular, we consider a variant of the Calvo (1983) model in which a firm's opportunity to change its price arrives stochastically and where, if this opportunity does not arise, the firm must keep its price constant. We choose the parameters of this extremely stylized model so that the model's predicted responses to monetary policy shocks match as closely as possible the estimated responses from the VAR.

In the third step, we combine the quantitative specification of the structural model (with the parameters obtained in step two) with the vector autoregression model to identify the shocks to the structural equations. The failure of the VAR to contain any stochastic singularities implies that the model's structural equations have residuals as well, which we can interpret in the context of the model as indicating stochastic variation in preferences and technology. We compute what these disturbances have been over our sample period, and use the VAR representation to determine the stochastic process followed by these real disturbances. One important advantage of our method of analysis is that, once we take into account these constructed disturbances, our model fits the data nearly as well as an unrestricted vector autoregression. This good fit provides an additional rationale for being interested in the models' implications concerning monetary policy.

In the fourth step, we use the quantitative model with parameters estimated in step two and shock processes estimated in step three to simulate the consequences of hypothetical monetary policy rules. Using the estimated historical shock series, we can simulate alternative historical paths for the U.S. economy. In particular, we can compute the realizations of output, inflation and interest rates under counterfactual rules. What is particularly attractive about this exercise is that, under the actual monetary rule estimated in step one, the simulated paths of output, inflation and interest rates are identical to the actual paths. Thus, the simulations with counterfactual rules provide alternative historical paths for the U.S. economy.

In the fifth and final step, we use the parameters estimated in step two to compute the welfare consequences of different monetary rules. Moreover, we derive the rule that would have maximized the utility of our representative households given the shock processes obtained in step three.

2 The Effects of Monetary Policy Shocks Under the Current Policy Regime

In this section, we describe our econometric characterization of the current monetary policy regime, and our estimates of the effects of monetary policy shocks under that regime. By monetary shocks we mean exogenous stochastic shifts in the feedback rule used by the Fed to set the Federal funds rate. Our interest in the effects of such shifts does not derive from a belief that they have played an important role in the generation of fluctuations in either output or inflation in the period with which we are concerned. Rather, we are interested in them because they can be econometrically identified without our having to commit ourselves to detailed assumptions about the true structural relations that determine output and inflation.

The monetary policy shocks and their effects cannot, of course, be identified without at least some weak *a priori* assumptions. In particular, we assume that recent U.S. monetary policy may be described by a feedback rule for the Federal funds rate of the form

$$r_t = r^* + \sum_{k=1}^{n_r} \mu_k [r_{t-k} - r^*] + \sum_{k=0}^{n_\pi} \phi_k [\pi_{t-k} - \pi^*] + \sum_{k=0}^{n_y} \theta_k [y_{t-k} - y^*] + \epsilon_t \tag{2.1}$$

where r_t is the Federal funds rate in period t, π_t is the rate of inflation between periods $t-1$ and t, y_t is the percentage deviation of real GDP from trend, y^* is the long-run average value of y_t, [2] and r^* and π^* are long-run "target" values for the funds rate and the rate of inflation respectively. The ϵ_t's represent exogenous monetary policy shocks which are assumed to be serially uncorrelated. In assuming that monetary policy shocks may be identified with movements in the Federal funds rate that cannot be predicted given the history of the funds rate, or by current and past values of other macro time series such as output and inflation, we follow a large part of the recent "structural VAR" literature on the identification of monetary policy shocks,

[2] One might define the output trend so that $y^* \equiv 0$, but this would not correspond to simple linear detrending of our log GDP series. We feel that our methods are most transparent if we detrend our output series by taking out a log-linear trend that makes the average value of y_t in our sample equal to zero, and then estimate the long-run average value y^*. In this way, our construction of our data series y_t is independent of our estimated VAR. Note also that we write (2.1) as though the central bank is aware of the value of y^* – which, according to the log-linear approximation to our equilibrium conditions that we use below, is independent of monetary policy – and uses this knowledge in setting interest rates. However, we might equally well assume that the central bank implements a rule of the form (2.1), in which y^* is an arbitrary constant chosen by the central bank, that need not equal the long-run average value. Such a rule would be equivalent to a rule of the form (2.1) in which y^* *does* refer to the long-run average value, but in which r^* and π^* have been suitably adjusted. To economize on notation, we let the policy rule be written in this form. For similar reasons, we assume a representation in which r^* and π^* correspond to the long-run average values of the series r_t and π_t. A central bank, of course, might choose "target" values r^* and π^* that are not consistent with the lonr-run average equilibrium real rate of return (that, in our log-linear model, is also independent of monetary policy), and in this case the long-run average values of these series in equilibrium would not correspond to the "target" values. But there would exist an alternative, equivalent representation of the policy rule in which the "target" rates imputed to the central bank coincide with the long-run average values in equilibrium.

beginning with Bernanke and Blinder (1992) and including Cochrane (1995) and Sims, Leeper, and Zha (1996). In our assumption of a feedback rule of the specific form (2.1) as a representation of current policy, we follow the monetary policy "reaction function" literature, especially Taylor (1993b). Taylor (1993b) asserts that at least under the chairmanship of Alan Greenspan (*i.e.*, at least since late 1987), Fed policy has been well described by a rule of this kind. [3]

Identification of the monetary policy shocks $\{\epsilon_t\}$, and estimation of the coefficients in (2.1), requires a further identifying assumption about the correlation between ϵ_t and the period t endogenous variables. Our assumption is that a monetary policy shock at date t has no effect on either output or inflation during period t; the idea is that both pricing and purchasing decisions for period t are made prior to the realization of the shock, *i.e.*, before the period t funds rate is observed. (In the theoretical model proposed in the next section as an interpretation of our VAR results, this assumption is made explicit.) An alternative interpretation of our restriction is that pricing and purchasing decisions for period t are made during period t, but on the basis of incomplete information – in particular, without information about current money-market conditions. The use of such "decision lags" as an identifying assumption is common in the structural VAR literature, beginning with Sims (1986). Under this assumption, equation (2.1) can be estimated using OLS. [4]

Note that this identifying assumption requires that we interpret any correlation between the period t innovations in inflation or output and the period t innovation in the funds rate as due to the way in which the Fed reacts to variations in inflation and output in setting the funds rate. It is therefore necessary that we deny the existence of "information lags" in the monetary policy rule, if we are to avoid imposing any over-identifying restrictions at this stage in our characterization of the data. This is why the ϕ_0 and θ_0 coefficients are allowed to be non-zero in (2.1). We find, in fact, a significant positive coefficient θ_0. This is because there is a significant positive correlation between the funds rate innovation and the contemporaneous output innovation. Since we would expect a negative correlation, if any, if the current output realization had no effect on monetary policy (while a positive correlation is entirely plausible under the assumption that

[3] According to Taylor's much-discussed account of recent policy, $n_r = 0$, $n_\pi = 3$, $n_y = 0$, $\pi^* = 2$ percent per year, $r^* = 4$ percent per year, and the coefficients are given by $\phi_k = 1.5/4$ for $k = 0, 1, 2, 3$, and $\theta_0 = .5$. Taylor presents these values as a rough rule of thumb rather than a precise quantitative specification. It is clear that actual policy involves a greater degree of interest-rate smoothing than the simple "Taylor rule" would predict; hence our allowance for lagged funds rate terms in our generalization (2.1) of Taylor's rule.

[4] Note that the constants r^* and π^* cannot be separately estimated from this equation alone. We are able to estimate them when we estimate our complete VAR model, by assuming that no equation of the three-variable VAR model contains a constant term, if the three state variables are all written in terms of deviations from their "long-run" values.

policy does respond to an output innovation within the quarter), the assumption of a decision lag for the Fed is unattractive, at least as regards the response to output innovations. [5]

We estimate (2.1) as part of a three-variable just-identified VAR model, where the variables included are the funds rate, the inflation rate, and detrended real GDP. [6] We estimate a complete system of this kind so that we obtain not merely an estimated Fed reaction function, but also estimated impulse responses to monetary policy shocks under the policy regime characterized by that reaction function. The three variables included represent a minimal set for our purposes: they are the minimal set needed to allow us to estimate a monetary policy rule of the kind proposed by Taylor (1993b), and they allow us to model the effects of monetary policy on fluctuations in the three variables that central banks are most often supposed to concern themselves with as ultimate goal variables.

The sample period for our estimation of (2.1) runs from 1980;1 through 1995:2. We begin our sample in the first quarter of 1980, because it is widely recognized that a significant change in the U.S. monetary policy regime occurred around that time; thus at least one equation of our model, the monetary policy rule (2.1), cannot be expected to have remained invariant over a longer time period than the one that we use. Many, of course, would doubt that the monetary policy rule has remained unchanged since then. Conventional accounts of the succession of U.S. monetary policy regimes often identify important regime changes in late 1982 (the end of the Fed's experiment with targeting of non-borrowed reserves) and late 1987 (the transition from Volcker to Greenspan) as well. [7] Our choice of the longest among several possible samples is determined by a desire to have long enough time series to allow estimation of an unrestricted VAR model. In fact, most VAR studies of the effects of monetary policy shocks make use of much longer data samples than ours. Our choice of a sample period represents a compromise between these two concerns.

Estimation of a VAR that includes (2.1) as one equation could be carried out in various ways. Perhaps the most obvious would be to estimate a recursive VAR with state vector $[\pi_t, y_t, r_t]'$, in which the "causal" ordering of the variables is the order in which they are listed. An alternative, that we follow here, is to

[5] The assumption made with regard to inflation innovations has little effect upon our results. Table 1 shows that the estimated coefficient ϕ_0 is in any event both small and statistically insignificant.

[6] Our desire to model GDP leads us to use quarterly data. Our $\{r_t\}$ series is the Federal funds rate, annualized and averaged over the quarter. Our $\{\pi_t\}$ series is the quarterly change in the log of the GDP deflator, also annualized. Finally, our $\{y_t\}$ series is the log of real GDP, with a linear time trend removed.

[7] See, e.g., Strongin (1995) and Bernanke and Mihov (1995). Taylor (1993b) proposes his feedback rule only as an account of Fed policy since late 1987.

estimate a recursive VAR with state vector

$$Z_t = [r_t, \pi_{t+1}, y_{t+1}]', \tag{2.2}$$

with the interest rate now first in the "causal" ordering. Specifically, we estimate a system of the form

$$T\bar{Z}_t = A\bar{Z}_{t-1} + \bar{e}_t \ , \tag{2.3}$$

where the vector $\bar{Z}_t$ is the transpose of $[Z_t', Z_{t-1}', Z_{t-2}']$, T is a lower triangular matrix with ones on the diagonal and nonzero off-diagonal elements only in the first three rows, A is a matrix whose first three rows contain estimated coefficients from the VAR. The first three rows of $\bar{e}_t$'s contain the VAR residuals while the other elements are zero. [8]

This notation is unfamiliar in that some data for period $t+1$ are included in the "period t" state vector. The reason is that, according to our assumption about decision lags, these variables that are only observed in period $t+1$ are nonetheless determined on the basis of information that, according to our model, decision-makers have in period t. We prefer this choice of notation because it allows us to describe the information used by decision-makers in terms of the history of the vector $\{Z_t\}$. Specifically, we can refer to the information set consisting of the history $\{Z_{t-j}\}$ for all $j \geq 0$ as the "period t information set". Then the information used in choosing price changes that take effect in period $t+1$ and quantities to be purchased in period $t+1$ may be taken to be exactly the period t information set; similarly, the information used in setting the funds rate in period $t+1$ is the period t information set, except that the Fed's action involves a random disturbance ϵ_{t+1} as well.

The first row of the estimated system (2.3) corresponds to the monetary policy feedback rule (2.1), and the first element of the residual vector e_t represents our identification of the monetary shock ϵ_t. The coefficients of the estimated policy rule are displayed in Table 1. The estimated rule may be described as a generalized "Taylor rule", though the dynamics are more complex than in Taylor's simple specification. One way to measure the overall responsiveness of the funds rate to fluctuations in inflation and output is in terms of "long-run" multipliers, that indicate the eventual increase in the funds rate that would result in the

[8] Three lags in (2.3) suffice to eliminate any significant serial correlation in the residuals, and the estimated coefficients on longer lags are also insignificant. Note that this form differs from the other, more familiar form, in the relative number of lags of the various variables at the point of truncation.

case of permanent changes in the levels of inflation and output. These long-run multipliers are given by [9]

$$r - r^* = 2.13[\pi - \pi^*] + .47y. \tag{2.4}$$

Thus, as in Taylor's rule, an increase in GDP relative to trend raises the funds rate (and our estimated long-run multiplier is essentially the same as that indicated by Taylor), and an increase in inflation relative to its "target" level raises the funds rate by an even greater amount (so that short-term *real* interest rates rise). Our estimate of the sensitivity of the funds rate to inflation fluctuations is even stronger than Taylor's coefficient; this may well be due to our inclusion of the Volcker years in our sample. Differences in our dynamic specification include our finding of significant interest-rate smoothing (the μ_k coefficients are all positive, and sum to .7), our finding that an increase in inflation does not begin to increase the funds rate until the following quarter, and our finding that the short-run multiplier for output is larger than the long-run multiplier (owing to the negative value for θ_2).

The complete estimated system also allows us to compute the response of output, inflation, and interest rates to a monetary policy shock. These impulse responses are plotted in the three panels of Figure 1. In each panel, the central dashed line indicates the point estimate of the impulse response function, while the two outer dot-dash lines indicate a confidence interval for each coefficient (plus and minus two times the standard error), based on analytic derivatives of the responses with respect to the parameters and on the variance-covariance matrix of the parameters.

Responses are plotted for a one-standard-deviation shock, that raises the funds rate unexpectedly by about eight-tenths of a percent. The estimated responses largely agree with conventional wisdom: interest rates are raised only temporarily (according our estimates, only for the first two quarters); output subsequently declines (not noticeably until two quarters later), but eventually returns to normal; and inflation also declines with a lag (with the greatest decline occurring two quarters later). These effects cannot be estimated with much precision (especially the effects on inflation). Nonetheless, they give us an idea of the features that our structural model should possess in order to be consistent with the data. In particular, a monetary tightening should temporarily lower both output and inflation; and these effects should occur only with a lag of a couple of quarters – so that the effects on output and inflation largely occur *after* short-term

[9] Though these cannot be read off from the regression reported in Table 1 alone, we estimate "long-run" values $\pi^* = 3.26\%$ and $r^* = 6.25\%$. We thus estimate a long-run average real funds rate of three percent. These values compare with Taylor's assumption of a two-percent real rate and a two-percent inflation target.

interest rates have returned to their normal level.

These results agree qualitatively with those that emerge from several recent VAR exercises including Christiano, Eichenbaum and Evans (1994) and, more relevantly (because he considers a three-variable VAR similar to ours), Cochrane (1994). In one of the many exercises reported in his paper, Cochrane (1994) estimates a VAR over the period 1959 to 1992 that includes quarterly observations of the federal funds rate as well as of the logarithms of output and of the price level. [10] When he includes a trend and computes the monetary policy shock by supposing that the systematic component of policy lets the federal funds rate at t react to output and the price level at t, he gets very similar impulse responses.

3 A Simple Model of Output and Inflation Determination

In this section we develop a simple equilibrium macroeconomic model, that we propose to use to interpret the fluctuations in the three time series of our VAR. The model is extremely rudimentary; we have reduced it to the essential elements necessary for a general-equilibrium account of the determination of output, inflation and interest rates. The model presented here is intended more as an illustration of the method that we advocate than as a complete model of the US economy. Nonetheless, we believe that it shows that it is possible to account for the main features of these time series in terms of a model derived from optimizing behavior under rational expectations.

The model we use is an extension of the one used in Woodford (1996) to analyze the consequences of interest-rate feedback rules for monetary policy, [11] with additional decision lags (both for price and quantity variables) added in order to allow a better fit with the predictions of the VAR. The optimizing decision-maker in our model is an infinite-lived representative household, which is both a consumer of all the goods produced in the economy and a producer of a single differentiated product. We index each household by i and let i vary continuously from zero to one. The objective of household i, looking forward from date $t = 0$,

[10] We get very similar impulse responses when we consider a longer sample that is similar to his. While the responses of output and inflation are more muted, we obtain qualitatively similar responses if, instead, we start our sample at the beginning of 1982. If the VAR includes only more recent observations, the effects of monetary shocks on output and inflation are much weaker and more poorly determined. This is probably due to the absence of significant monetary disturbances in the more recent period.

[11] Certain aspects of the structure of the model are discussed in more detail there. Note that here, unlike in the previous analysis, we consider only the case of a "Ricardian" fiscal policy, which allows us to avoid discussion of certain equilibrium conditions that play a key role in the earlier paper.

is to maximize

$$E_0 \sum_{t=0}^{\infty} \beta^t [u(C_t^i; \xi_t) - v(y_t^i; \xi_t)], \tag{3.1}$$

where β is a discount factor, y_t^i is the output of the good produced by household i at time t, ξ_t is a vector of random disturbances, and for each value of ξ_t, u is an increasing, concave function, and v is an increasing, convex function. [12] The argument C_t^i represents an index of the household's purchases of all of the continuum of differentiated goods produced in the economy, given by

$$C_t^i = \left[\int_0^1 c_t^i(z)^{\frac{\theta-1}{\theta}} dz \right]^{\frac{\theta}{\theta-1}} \tag{3.2}$$

as in Dixit and Stiglitz (1977), where $c_t^i(z)$ is the quantity purchased of good z, and the constant elasticity of substitution θ is assumed to be greater than one.

For expositional purposes, it is easiest to think of these purchases by the household as purchases of non-durable consumption goods, and our model as one in which all output is non-durable and immediately consumed either by households or by the government. If capital goods are used in production, they are in fixed supply (and non-depreciating), and their allocation across firms (or households) cannot be changed. However, we seek to use the model to explain the dynamics of real GDP, and not just of consumption spending. Thus we actually interpret C_t^i as referring to the household's purchases of investment goods as well as consumption goods, and assume (at the price of an obvious oversimplification) that all such purchases are made to obtain immediate utility, which can be expressed (through a reduced-form or indirect utility function) as an increasing, concave function, of total purchases at a given date. We ignore the effects of investment spending upon the evolution of productive capacity, in the hope that these effects are in any event not too significant at the frequencies with which we are most concerned in evaluating alternative monetary policies. [13] In making this simplifying assumption, of course, we follow a long tradition of macroeconometric modeling based upon more or less elaborate versions of the textbook "IS-LM" model. [14]

Perhaps the most surprising element of our preference specification, given that we are interested in monetary issues, is that we abstract from the liquidity services provided by money. This is no more than a simplification. Our model can be understood as a the limit of a model where real money balances provide

[12] As is discussed further in section 4, we can interpret v as a reduced-form representation of production costs in a model with firms and labor markets. Under this interpretation, v is convex both because there are diminishing returns to labor and because the marginal disutility of labor is increasing in labor.

[13] This will be a subject of quantitative investigation in later work.

[14] Similar simplifications are advocated, for example, in McCallum and Nelson (1997).

utility but where, in the limit, these liquidity services are arbitrarily small. Alternatively, the model can be understood as one where utility is additively separable in real money balances, consumption, and goods supply, as in Woodford (1996). In this case, the model implies an additional first order condition relating real balances to consumption and the interest rate. In the presence of an interest rate rule, which is after all the focus of our analysis, this additional equilibrium condition simply determines the nominal level of money balances. Since this equilibrium condition plays no role in determining inflation, output or interest rates, it can safely be ignored for our purposes.

When allocating a given amount of nominal spending at t, S_t^i, across all the different differentiated goods, the household maximizes (3.2) subject to the constraint that spending on all goods not exceed S_t^i. This leads to the familiar Dixit-Stiglitz demand relations for relative quantities purchased as a function of relative prices. As usual, the total expenditure required to obtain a given quantity of the consumption aggregate (3.2) is given by $S_t^i = P_t C_t^i$, where P_t is the Dixit-Stiglitz price index defined by

$$P_t = \left[\int_0^1 p_t(z)^{1-\theta} dz \right]^{\frac{1}{1-\theta}}. \tag{3.3}$$

The optimal allocation of purchases across the individual goods is then given by

$$c_t^i(z) = C_t^i \left(\frac{p_t(z)}{P_t} \right)^{-\theta} \tag{3.4}$$

for each good z. The intertemporal allocation of spending, in turn, amounts to choosing the path of the aggregate $\{C_t^i\}$ defined in (3.2) to maximize (3.1) subject to an intertemporal budget constraint that be written in terms of aggregate consumption expenditure expenditure each period, given the price P_t of the composite good.

We suppose that there are complete financial markets, [15] and no obstacles to borrowing against future income, so that each household faces a single intertemporal budget constraint. Looking forward from date t, the intertemporal budget constraint of household i is of the form

$$E_t\{\sum_{T=t}^{\infty} \delta_{t,T} S_T^i\} \leq E_t\{\sum_{T=t}^{\infty} \delta_{t,T}[p_t(i) y_t^i - T_t]\} + A_t^i, \tag{3.5}$$

[15] It is important to note that by this we mean simply that a household is able to transfer *purchasing power* freely between dates and states of the world, through the exchange of state-contingent financial claims. We do not suppose that the state-contingent exchanges of real goods and services that occur in our model are contracted at some initial date, as in the Arrow-Debreu model. Nor do we suppose that households are able to arrange more complicated sorts of contracts that would allow them, in effect, to get around the constraints upon price-setting that we assume to exist in the spot markets for goods and services. The state-contingent financial claims that we imagine, however, do include the possibility of insuring oneself against the idiosyncratic income risk that households suffer because they change their prices at different dates.

where A_t^i denotes the nominal value of the household's financial assets at the beginning of period t, T_t denotes its net nominal (lump sum) tax obligation at date t, and $\delta_{t,T}$ is the stochastic discount factor that defines the nominal present value at t of nominal income in any given state at date $T \geq t$. Because of the existence of complete financial markets, these stochastic discount factors are uniquely defined; a financial claim to a random nominal quantity X_T has a nominal value at t of $E_t[\delta_{t,T}X_T]$. In particular, if R_t denotes the (gross) nominal interest rate on a riskless one-period bond purchased in period t, this interest rate must satisfy [16]

$$R_t = E_t[\delta_{t,t+1}]^{-1}. \tag{3.6}$$

We assume that households must choose their index of purchases C_t^i at date $t-2$. [17] We interpret this to mean that actions already taken prior to the realization of date $t-1$ monetary policy make it imperative that certain purchases be made during period t. Certain kinds of interest-sensitive purchases do involve advance commitment of this kind; in particular, many kinds of investment projects take more than one quarter to complete, and abandonment of the project after initiation may be too costly to be contemplated except in the case of quite extreme interim changes in market conditions. [18]

The household's optimal program of purchases then must satisfy

$$E_t\{u_C(C_{t+2}^i; \xi_{t+2})\} = E_t\{\lambda_{t+2}^i P_{t+2}\}, \tag{3.7}$$

at each date, where λ_t^i represents the household's marginal utility of nominal income at date t. [19] The marginal utilities of income at different dates and in different states in turn must satisfy

$$\lambda_t^i \delta_{t,T} = \beta^{T-t}\lambda_T^i \tag{3.8}$$

for any $T \geq t$. Conditions (3.7) and (3.8), and the requirement that (3.5) hold with equality, completely

[16] Other financial claims can similarly be priced, but this is the only one that matters for our purposes, as we assume that the central bank conducts monetary policy by trading in the market for short-term nominal claims of this kind.

[17] In fact, we simply require that C_t^i be determined as of the beginning of period $t-1$, before the monetary policy shock in period $t-1$ is revealed, but at a time at which all period $t-1$ goods transactions have been determined. In terms of our notational convention, this means that C_t^i belongs to the "date $t-2$" information set. Note also that while we require that the index C_t^i be determined at date $t-2$, the purchases of individual goods that are made in order to achieve this are determined only at date $t-1$ – or by the beginning of period t – as these depend upon the period t prices of the individual goods.

[18] Our assumption thus amounts to a limiting case of a "time-to-build" model, in which the bulk of the expenditure connected with a "project" initiated at the beginning of period $t-1$ occurs during period t. Note that we could give an "information lag" interpretation to this restriction upon household purchases: households choose their overall level of purchases in period t with knowledge of period $t-1$ goods market conditions, but before learning about period $t-1$ money market conditions, or about period t conditions in either goods markets or the money markets.

[19] This quantity appears as the Lagrange multiplier associated with constraint (3.5). It is measured in units of period t utility flow per dollar.

determine the household's optimal consumption plan, given its initial wealth, initial predetermined consumption level, and after-tax income expectations.

We furthermore assume that financial markets exist that allow households to insure one another against idiosyncratic income risk (that here results solely from differences in the time at which they change their prices). Assuming that all households have identical initial wealth, they will choose in equilibrium to completely pool their income risk, and we assume an equilibrium of this kind. As a result, in equilibrium the right-hand side of (3.5) has the same value for each household at any date, and households choose identical consumption plans and have identical marginal utilities of income. We can therefore drop the i superscripts in equations (3.7) and (3.8). Note also that it follows from (3.6) and (3.8) that the common marginal utility of income satisfies

$$\lambda_t = \beta E_t[R_t \lambda_{t+1}]. \tag{3.9}$$

In our computation of the equilibrium responses to shocks in subsequent sections, we make use of a log-linear approximation to the equilibrium conditions of our model, expanding in terms of percentage deviations of various stationary state variables from their steady-state values (their constant values in the absence of all stochastic disturbances). In this log-linear approximation, (3.9) becomes

$$\hat{\lambda}_t = E_t[\hat{R}_t - \pi_{t+1} + \hat{\lambda}_{t+1}], \tag{3.10}$$

where $\hat{\lambda}_t$, $\hat{R}_t$, and π_t denote percentage deviations in the stationary variables $\lambda_t P_t$ and R_t respectively, and $\pi_t \equiv \log P_t / P_{t-1}$ is the rate of inflation. (Note that we log-linearize all of our equations around a steady-state equilibrium in which the constant rate of inflation is zero.) We can furthermore solve this forward to obtain

$$\hat{\lambda}_t = \hat{r}_t^l \equiv \sum_{T=t}^{\infty} E_t[\hat{R}_T - \pi_{T+1}] \tag{3.11}$$

where $\hat{r}_t^l$ defines percentage deviations in a long-run real rate of return. (In all of the equilibria that we analyze below, both interest rates and inflation follow stationary ARMA processes, as a result of which the infinite sum in (3.11) converges.)

The corresponding log-linear approximation to (3.7) is given by

$$-\tilde{\sigma} E_t[\hat{C}_{t+2} - \bar{C}_{t+2}] = E_t \hat{r}_{t+2}^l, \tag{3.12}$$

where the elasticity $\tilde{\sigma}$ equals $-u_{CC}C/u_C$ (evaluated at the steady-state level of consumption), $\hat{C}_t$ indicates the percentage deviation of C_t from its steady-state value, and $\bar{C}_t$ is an exogenous disturbance (a certain function of the preference shock ξ_t) indicating the level of consumption required at each point in time to maintain a certain constant marginal utility of consumption. [20] Equation (3.12), together with the stipulation that $\hat{C}_{t+2}$ be determined at date t, indicates how interest-sensitive purchases in period $t+2$ depend upon interest-rate expectations at date t.

Total aggregate demand is assumed to be given by

$$Y_t = C_t + G_t, \tag{3.13}$$

where G_t represents exogenous variation in "autonomous spending". While one natural interpretation of G_t is that it represents exogenous variation in government purchases, there are other possible interpretations as well. For instance, G_t could represent consumption purchases by liquidity-constrained consumers, who spend all their income (as in Campbell and Mankiw, 1989, except that we suppose that the real income of these consumers is an exogenous random process, rather than being a constant fraction of total income). Because of these various possible interpretations, we do not seek to identify G_t with the government purchases variable in the national accounts, any more than we wish to identify C_t with consumer expenditure. Rather, we focus only on the implications of our model for the evolution of Y_t. Note that we assume that G_t is determined by the beginning of period t (*i.e.*, that it belongs to the date $t-1$ information set), for consistency with the assumptions made in our identification of the monetary policy shocks from our VAR in section 2.

Log-linearization of (3.13) yields

$$\hat{Y}_t = s_C\hat{C}_t + \tilde{G}_t,$$

where $\hat{Y}_t$ denotes percentage deviations of Y_t from its steady-state value, $\tilde{G}_t$ denotes the deviation of G_t from its steady-state value expressed as a percentage of the steady state value of Y, and s_C is the steady-state value of the interest-sensitive share C/Y. Substituting this into (3.12) then yields the model's "IS equation"

$$\hat{Y}_t = -\sigma^{-1}E_{t-2}\hat{r}^l_t + \hat{G}_t, \tag{3.14}$$

where $\sigma \equiv \tilde{\sigma}/s_C$, and $\hat{G}_t \equiv \tilde{G}_t + s_C E_{t-2}\bar{C}_t$ is a composite exogenous disturbance. The "aggregate demand block" of our model then consists of the monetary policy rule, the term-structure equation (3.11), and the

[20] Further details of our log-linear approximations are provided in Appendix 3.

IS equation (3.14), [21]

We now turn to our optimizing model of price-setting and aggregate supply. The decision problem of price-setters depends upon the demand that they face for their product. We suppose that autonomous expenditure G_t also represents an aggregate of purchases of individual goods, and we let this aggregate have the same form as (3.2) so that

$$G_t = \left[\int_0^1 g_t(z)^{\frac{\theta-1}{\theta}} dz\right]^{\frac{\theta}{\theta-1}}.$$

Moreover, we suppose that these individual purchases, $g_t(z)$, are chosen to maximize this aggregate for any given level of expenditure. As a result, these purchases are made in the same proportions as those of consumers, described by (3.4). Then the overall demand faced by an individual supplier satisfies

$$y_t^i = Y_t\left(\frac{p_t(i)}{P_t}\right)^{-\theta}. \tag{3.15}$$

As is standard in models of monopolistic competition, we assume that an individual supplier regards itself as unable to affect the evolution of the variables Y_t and P_t, and so chooses its own price taking the evolution of those variables as given.

The source of the real effects of monetary policy in our model is an assumption of decision lags in price-setting. Following Calvo (1983), we assume that prices are changed at exogenous random intervals. Specifically, a fraction $1-\alpha$ of sellers get to choose a new price at the beginning any given period, whereas the others must continue using their old prices. Of those who get to choose a new price, a fraction γ start charging the new price during that period, whereas the remaining fraction $1-\gamma$ must wait until the next period to charge the new price, because (owing to a different organization of the markets for these goods) they must post their prices a quarter in advance. These assumed delays explain why no prices respond in the quarter of the monetary disturbance (as assumed in our identification of the policy shocks), and why the largest response of inflation to a monetary shock takes place only two quarters after the shock. [22]

Let p_t^1 denote the price set by sellers that decide at date $t-1$ upon a new price to take effect at date t, and p_t^2 the price set by sellers that decide at date $t-2$ upon a new price to take effect only two periods later. These prices are chosen to maximize the contributions to expected utility resulting from sales revenues on

[21] Note that this is the same basic structure as is used in the small model of Fuhrer and Moore (1995b). However, because our IS equation is derived from intertemporal optimization, there are some differences; for example, it is the past *expectation* of the *current* long-term real rate that affects current aggregate demand, in our specification, rather than the past value of the long-term real rate itself.

[22] Once again, our decision lags could represent information lags.

the one hand, and the disutility of output supply on the other, at each of the future dates and in each of the future states in which the price commitment still applies. This means that p_t^1 is chosen to maximize

$$\Phi_{t-1}(p) \equiv E_{t-1}\sum_{T=t}^{\infty}(\alpha\beta)^{T-t}\left[\lambda_T(1-\tau)pY_T\left(\frac{p}{P_T}\right)^{-\theta} - v\left(Y_T\left(\frac{p}{P_T}\right)^{-\theta};\xi_T\right)\right] \tag{3.16}$$

over p. Here we have substituted the demand function (3.15) into the household's objective function, written λ_T for the marginal utility (in units of period T utility flow) of additional nominal income during period T, and assumed that revenues each period are taxed at the constant rate τ. [23] The factor α^{T-t} appears as the probability that the price that is charged beginning in period t is still in effect in period $T \geq t$ (which contingency is assumed independent of all aggregate disturbances). Note that our assumption of complete contingent claims markets (including full opportunities for households to insure one another against idiosyncratic risk associated with different timing of their price changes) implies that the marginal utility of income process $\{\lambda_T\}$ is the same for all households, and can be treated as an exogenous stochastic process by an individual household (whose pricing decisions will have only a negligible effect upon aggregate prices, aggregate incomes, and aggregate spending decisions). Similarly, an individual household treats the processes $\{P_T, Y_T\}$ as exogenous in choosing its desired price. The optimizing choice of p_t^1 then must satisfy the first-order condition

$$\Phi'_{t-1}(p_t^1) = 0, \tag{3.17}$$

where the prime denotes the derivative with respect to p in the explicit expression given in (3.16).

As before, we wish to log-linearize this equilibrium condition around a steady state in which $Y_t = \bar{Y}$, $P_t/P_{t-1} = 1$, $p_t^1/P_t = 1$, and $\lambda_t P_t = \bar{\lambda}$ at all times. (The requirement that these constant values satisfy (3.17) when $\xi_t = 0$ at all times determines the steady-state value $\bar{Y}$. [24]) Percentage deviations of $v_y(y_t^j;\xi_t)$ from its steady-state value can be written as $\omega(\hat{y}_t^j - \bar{Y}_t)$, where $\omega \equiv v_{yy}\bar{Y}/v_y$, with partial derivatives evaluated at the steady state, $\hat{y}_t^j$ refers to the percentage deviation of y_t^j from its steady-state value $\bar{Y}$, and $\bar{Y}_t$ is a certain linear function of ξ_t. Using this notation, the log-linear approximation to (3.17) takes the form

$$\begin{aligned} E_{t-1}\sum_{T=t}^{\infty}(\alpha\beta)^{T-t}[(\hat{\lambda}_T \quad &+ \quad \hat{Y}_T - (\theta-1)\hat{p}_{t,T}^1) \\ &-\omega(\hat{Y}_T - \theta\hat{p}_{t,T}^1 - \bar{Y}_T) - (\hat{Y}_T - \theta\hat{p}_{t,T}^1)] = 0, \end{aligned} \tag{3.18}$$

[23] The allowance for non-zero τ is primarily so that we can linearize around a steady state in which the constant level of output is efficient. The convenience of this for our purposes is discussed in Appendix 3.

[24] Under the assumption that $\tau = -(\theta-1)^{-1}$, this requires that $\bar{Y}$ satisfy the equation $u_C(\bar{Y} - \bar{G}; 0) = v_y(\bar{Y}; 0)$, which also defines the efficient level of output.

where in addition $\hat{p}^1_{t,T} \equiv \log(p^1_t/P_T)$. Introducing the notation $\hat{X}_t \equiv \frac{1-\alpha}{\alpha}\log(p^1_t/P_t)$, [25] so that $\hat{p}^1_{t,T} = \frac{\alpha}{1-\alpha}\hat{X}_t - \sum_{s=t+1}^{T}\pi_s$, we can solve (3.18) to obtain

$$\hat{X}_t = \frac{1-\alpha}{\alpha}\frac{1-\alpha\beta}{1+\omega\theta}E_{t-1}\sum_{T=t}^{\infty}(\alpha\beta)^{T-t}[-\hat{\lambda}_T + \omega(\hat{Y}_T - \bar{Y}_T) + (1+\omega\theta)\sum_{s=t+1}^{T}\pi_s] \tag{3.19}$$

as the optimizing choice of the relative price in period t of goods with new prices chosen just the period before.

Note that (3.11) and (3.14) imply that

$$E_{t-2}\hat{\lambda}_t = -\sigma E_{t-2}[\hat{Y}_t - \hat{G}_t]. \tag{3.20}$$

We can use this to eliminate the $E_{t-1}\hat{\lambda}_T$ terms in (3.19), for all $T > t$. Taking the conditional expectation of (3.10) at $t-1$, and again using (3.20), we see that we can also write

$$E_{t-1}\hat{\lambda}_t = \phi_{t-1} - \sigma E_{t-1}[\hat{Y}_t - \hat{G}_t], \tag{3.21}$$

where

$$\begin{aligned}\phi_t &\equiv E_t[\hat{R}_{t+1} - \hat{\pi}_{t+2} - \sigma(\hat{Y}_{t+2} - \hat{G}_{t+2} - \hat{Y}_{t+1} + \hat{G}_{t+1})] \\ &= E_{t-1}\sum_{T=t}^{\infty}(\hat{R}_T - \pi_{T+1}) - E_{t-2}\sum_{T=t}^{\infty}(\hat{R}_T - \pi_{T+1}).\end{aligned} \tag{3.22}$$

Note that the final equality in (3.22) follows from subsitution of (3.14). Then, substituting (3.20) and (3.21) into (3.19), we obtain

$$\hat{X}_t = \frac{1-\alpha}{\alpha}\frac{1-\alpha\beta}{1+\omega\theta}\left[E_{t-1}\sum_{T=t}^{\infty}(\alpha\beta)^{T-t}[(\sigma+\omega)(\hat{Y}_T - \hat{Y}^S_T) + (1+\omega\theta)\sum_{s=t+1}^{T}\pi_s] - \phi_{t-1}\right],$$

$$= \frac{1-\alpha}{\alpha}\frac{1-\alpha\beta}{1+\omega\theta}\left[E_{t-1}\sum_{T=t}^{\infty}(\alpha\beta)^{T-t}[(\sigma+\omega)(\hat{Y}_T - \hat{Y}^S_T) + (1+\omega\theta)\frac{\alpha\beta}{1-\alpha\beta}\pi_{T+1}] - \phi_{t-1}\right], \tag{3.23}$$

where

$$\hat{Y}^S_t \equiv \frac{\omega}{\omega+\sigma}E_{t-1}\bar{Y}_t + \frac{\sigma}{\omega+\sigma}\hat{G}_t$$

is a composite exogenous disturbance. We can think of $\hat{Y}^S_t$ as representing variation in the "natural" or "potential" level of output, since it is expected deviations $\hat{Y} - \hat{Y}^S$, rather than deviations in the level of output relative to trend, that results in a desire by price-setters to increase the *relative* price of their goods,

[25] The factor $(1-\alpha)/\alpha$ turns out to be convenient in giving a simpler form to equations such as (3.25) below.

which in equilibrium requires inflation of the average level of prices. (An equilibrium in which no prices are ever changed is consistent with (3.23) as long as $\hat{Y}_t = \hat{Y}_t^S$ at all times, and interest rates vary so as to ensure that $\phi_t = 0$ at all times. Note that the latter condition ensures that (3.20) and hence (3.14) are also satisfied at all times.)

We turn next to the price-setting decision of sellers that choose a new price p_t^2 at $t-2$ to apply beginning in period t. Because such a price is expected to apply in periods $t+j$ with exactly the same probabilities as for the price p_t^1, the objective of these sellers is simply $E_{t-2}\Phi_{t-1}(p)$, and the first-order condition that determines p_t^2 is given by $E_{t-2}\Phi'_{t-1}(p_t^2) = 0$. Comparison with (3.17) implies that, in our log-linear approximation,

$$\log p_t^2 = E_{t-2} \log p_t^1. \tag{3.24}$$

Finally, our definition of the price index (3.3) implies that this index evolves according to

$$P_t = [\alpha P_{t-1}^{1-\theta} + (1-\alpha)\gamma(p_t^1)^{1-\theta} + (1-\alpha)(1-\gamma)(p_t^2)^{1-\theta}]^{1/1-\theta}.$$

Dividing both sides by P_t, log-linearizing, and substituting (3.24), we obtain

$$\pi_t = \gamma \hat{X}_t + (1-\gamma)[E_{t-2}\hat{X}_t - \frac{1-\alpha}{\alpha}(\pi_t - E_{t-2}\pi_t)].$$

Taking the conditional expectation of both sides at $t-2$, one observes that $E_{t-2}\pi_t = E_{t-2}\hat{X}_t$. Substitution of this then allows the equation to be written in the form

$$\pi_t = \psi \hat{X}_t + (1-\psi)E_{t-2}\hat{X}_t, \tag{3.25}$$

where $\psi \equiv \gamma\alpha/[1-\gamma(1-\alpha)]$. This indicates how aggregate inflation results from the incentives of individual price-setters to choose a higher relative price.

These results may be collected in the form of an implied aggregate supply relation between inflation variation and deviations of output from potential. Equation (3.23) may be expressed in quasi-differenced form as

$$\begin{aligned} \hat{X}_t &= \kappa(\hat{Y}_t - \hat{Y}_t^S) + (1-\alpha)\beta E_{t-1}\pi_{t+1} + \alpha\beta E_{t-1}\hat{X}_{t+1} - \frac{\kappa}{\sigma+\omega}\phi_{t-1} \\ &= \kappa(\hat{Y}_t - \hat{Y}_t^S) + \beta E_{t-1}\hat{X}_{t+1} - \frac{\kappa}{\sigma+\omega}\phi_{t-1}, \end{aligned} \tag{3.26}$$

where $\kappa \equiv (1-\alpha)(1-\alpha\beta)(\omega+\sigma)/\alpha(1+\omega\theta)$, and the second line follows from the fact that (3.25) implies that $E_{t-2}\pi_t = E_{t-2}\hat{X}_t$. Solving this forward, we obtain

$$\hat{X}_t = \kappa E_{t-1}\{\sum_{T=t}^{\infty} \beta^{T-t}(\hat{Y}_T - \hat{Y}_T^S)\} - \frac{\kappa}{\sigma+\omega}\phi_{t-1},$$

where we have used the fact that (3.22) implies that $E_{t-1}\phi_t = 0$. Substitution of this into (3.25) then yields

$$\begin{aligned} \pi_t &= (1-\psi)E_{t-2}\pi_t + \psi\Big[\kappa E_{t-1}\sum_{T=t}^{\infty}\beta^{T-t}(Y_T - Y_T^S) \\ &\quad -\frac{\kappa}{\omega+\sigma}\Big(E_{t-1}\sum_{T=t}^{\infty}(\hat{R}_T - \pi_{T+1}) - E_{t-2}\sum_{T=t}^{\infty}(\hat{R}_T - \pi_{T+1})\Big)\Big]. \end{aligned} \tag{3.27}$$

This is our aggregate supply (AS) equation, relating inflation variation to deviations of output from potential. Because prices are set in advance, expectations of future increases in output relative to Y^S also raise prices. In addition, inflation declines when the long term real interest rate at t is higher than had been expected at $t-1$. The reason for this is that such upwards revisions raise the returns households can expect to earn from their revenues at t. As a result, they are inclined to raise these revenues by cutting their prices. Only surprise variations in the long rate contribute to this term, because only those variations result in changes in the current marginal utility of income that are not reflected in the current level of aggregate consumption demand, and hence in the output gap.

Note that if we take conditional expectations of both (3.25) and (3.26) at date $t-2$, we obtain simply

$$E_{t-2}\pi_t = \kappa E_{t-2}(\hat{Y}_t - \hat{Y}_t^s) + \beta E_{t-2}\pi_{t+1}. \tag{3.28}$$

This is identical, in conditional expectation, to the form of aggregate supply equation obtained from models such as those of Rotemberg (1982) and Calvo (1983), called by Roberts (1995) "the New Keynesian Phillips Curve". In this expectations-augmented Phillips curve, the exogenous disturbance Y_t^s plays the role of a "natural" rate of output for period t. Our aggregate supply relations here have the same implications as regards the relation that must exist between inflation and output fluctuations that can be forecasted sufficiently far in advance; but they allow for a more flexible short-run relationship, due to the existence of additional decision lags beyond those present in the simplest discrete-time version of the Calvo model.

If we similarly take the conditional expectations of (3.11) and (3.14) at date $t-2$, we obtain simply

$$E_{t-2}[\hat{R}_t - \pi_{t+1}] = \sigma E_{t-2}[(\hat{Y}_{t+1} - \hat{Y}_t) - (\hat{G}_{t+1} - \hat{G}_t)]. \tag{3.29}$$

This is identical, in conditional expectation, to the standard consumption Euler equation in a frictionless representative-consumer model. Thus our aggregate demand relations here imply the same kind of relationship as in that simple model between real interest rate and output fluctuations that can be forecasted sufficiently far in advance, but again allow a more flexible short-run relationship. Note that the evolution of the conditional expectations of our state variables, $E_{t-2}\hat{R}_t$, $E_{t-2}\pi_t$, and $E_{t-2}\hat{Y}_t$, is determined entirely by equations (3.28) and (3.29), together with the conditional expectation of (2.1); it is thus the same as in the simpler models used by Rotemberg (1996) and Woodford (1996).

4 Estimation of Model Parameters

We now have a complete model of the determination of output, inflation and interest rates, which consists of equations (2.1), (3.11), (3.14), and (3.27). [26] Apart from the coefficients of the monetary policy rule (2.1), the estimation of which we have discussed in the previous section, and the parameters specifying the stochastic processes for the real disturbances, which we consider in the following section, the model involves six parameters, the structural parameters α, β, γ, σ, θ and ω. Here we consider the estimation of these parameters so as to make the model's predictions regarding the effects of a monetary policy shock fit those estimated by the unrestricted VAR (and shown in Figure 1) as closely as possible.

Before proceeding further, it is perhaps worth commenting briefly on our estimation strategy. [27] We wish to estimate the model parameters so as to match theoretical with measured second moments of our three time series. The second moments of the data can be completely summarized by (i) the variances of the three orthogonal VAR innovations, and (ii) the impulse response functions of the three variables to each of the three orthogonal innovations. [28] Obviously, we can match the observed variances of the three innovations by choosing appropriate variances for the exogenous disturbances in our model. We show in the next section that we can also completely match the estimated impulse responses of all three variables to the

[26] We should perhaps clarify the relation that we assume between the state variables $\hat{R}_t, \hat{Y}_t, \pi_t$ of our log-linear theoretical model and our data. If we define $r_t \equiv \log R_t$ as the instantaneous short-term nominal interest rate, then it is evident that $\hat{R}_t = r_t - \rho$, where $\rho \equiv -\log \beta$ is the steady-state instantaneous real rate of interest, and hence the steady-state nominal rate associated with a zero-inflation steady-state. We assume that our data series π_t measures the theoretical variable π_t and that our data series r_t measures the theoretical variable r_t. Thus our data series r_t, π_t, y_t correspond to the theoretical series $\hat{R}_t, \pi_t, \hat{Y}_t$, up to constants (that result from the fact that the estimated long-run average values correspond to a steady state with inflation of more than three percent per year).

[27] Both the issue of identification, and our approach to estimation, are discussed in further detail in Appendix 1 below.

[28] Note that the possibility of describing the data in this way does not depend upon having a structural interpretation of the orthogonal VAR innovations.

two innovations that are orthogonal to the identified monetary policy shock by appropriately specifying the stochastic processes of the "real" disturbances. This possibility is furthermore completely *independent* of the values we may have assigned to the structural parameters, though it depends upon our use of OLS to estimate the parameters of the monetary policy rule. We do indeed obtain the monetary rule in this way; we use the estimates in the first column of Table 1 with the coefficients on output divided by four so that the estimates are identical to those we would have obtained if we had not annualized our interest rate and inflation series. Given this way of estimating the monetary policy rule, the *only* features of the data that provide any information about the structural parameters are the impulse responses to the monetary policy shock, shown in Figure 1. [29] It follows that it is appropriate to estimate these parameters so as to fit these estimated impulse response functions.

We now discuss the extent to which the structural parameters can be identified from the impulse responses. In fact, they are not all identified. Note first that θ appears only in (3.26), where it matters only for the determination of κ. We may thus propose to estimate κ rather than θ. With this change of variables, neither α nor γ appears except in (3.25), where they matter only through their effect upon the ratio $\psi \equiv \gamma\alpha/(1-\gamma(1-\alpha))$. Thus the identified parameters are at most the five parameters $\beta, \kappa, \sigma, \omega$, and ψ.

In fact, one can show that only four combinations of the structural parameters can be identified from the impulse responses to a monetary policy shock, given a particular feedback rule for the monetary policy. The parameters β, κ and σ are each identified, but only a single function of ω and ψ is, rather than either of these being identified independently. [30] While the impulse responses do not fully identify ψ they *can* be used to put a lower bound on it (given the possible range of variation in ω). Thus our introduction of a subset of producers who must determine price changes two quarters in advance (corresponding to $\gamma < 1$, so that ψ is strictly positive) does allow us to better fit the estimated impulse responses, even if the parameter γ cannot be identified.

The reason that, for a given κ, ψ and ω cannot be identified separately from the impulse response function is the following. The effect of a monetary shock at t on the expected evolution of prices, output and inflation

[29] For clarity of presentation, these are displayed once again using annualized inflation and interest rates even though the equations of the model involve quarterly interest and inflation rates.

[30] The particular combination of parameters that matters for the impulse responses depends on the coefficients of the monetary policy rule. Thus both ω and ψ matter for the model's predictions regarding the consequences of alternative monetary policy rules – they even matter for the model's predictions regarding the effects of *monetary policy shocks* under alternative regimes. However, if we only observe the impulse responses under a *single* policy regime, we are unable to disentangle the two parameters without further information.

starting from their position at $t+1$ is independent of ψ and ω. The reason is that, for $k \geq 1$ the expectation at t of ϕ_{t+k} is zero (so that ω does not affect the dynamic equations for periods beyond $t+1$) while the expectation at t of π_{t+k+1} is equal to the expectation of $E_{t+k+1-2}\pi_{t+k+1}$ (so that ψ does not affect the dynamic equations beyond $t+1$). So, the parameters ψ and ω matter only for determining the response of inflation at $t+1$ (since output at $t+1$ is predetermined). But, the response of just one variable at $t+1$ cannot separately identify two parameters.

The complete set of parameters can thus be determined only if we introduce additional considerations. We propose to "calibrate" two of the unidentified parameters on the basis of independent evidence, allowing us to estimate the remaining parameters from the impulse responses. One parameter we calibrate is α, which determines how frequently on average a producer changes his or her price. Given our assumed constraints upon price-setting our model implies that the mean time that a given price remains in effect is $1/(1-\alpha)$ quarters. We can thus choose a plausible value for this parameter based upon microeconomic evidence regarding the average length of time individual prices remain in effect.

Unfortunately, different studies of individual price adjustment report different estimates for this average length. At one extreme, Cecchetti (1986) reports that the newsstand prices of magazines stay constant for between 1.8 and 14 years. At the other, Dutta, Bergen and Levy (1997) report that the price of orange juice at supermarkets stays constant for between 2 and 10 weeks. The findings of Carlton (1986), Blinder (1994) and Kashyap (1995) fall somewhere in between. Carlton (1986) shows that, depending on the product, prices are constant for between 4 and 13 months. Blinder (1994) reports that his interviewees kept prices constant for an average of 9 months. [31] Finally, Kashyap (1995) shows that products from the L. L. Bean catalogue keep their prices constant for between 11 and 30 months. This diversity of findings leads us to use Blinder's (1994) estimate both because it is relatively conservative and because it covers a broad range of industries. The result is that $1/(1-\alpha)$ equals 3 quarters so that α equals .66.

We can also calibrate ω on the basis of data regarding labor costs. Note that this parameter gives the elasticity of the marginal disutility of producing output with respect to an increase in output. It is thus closely related to the elasticity of marginal cost with respect to output. To relate this parameter to measures

[31] In his Table 4.1, he reports the fraction of his respondents for whom the frequency of price adjustment falls between various threshold levels. We had to make somewhat arbitrary assumptions to convert these frequency categories into mean lengths of constant prices.

of labor costs, we may further specify that output is produced via a production function $Y = f(H)$, where H represents hours worked, and that the representative household has a disutility of working given by $g(H)$, so that $v(Y)$ is equal to $g(f^{-1}(Y))$. The equivalence with our earlier formulation is direct if the household uses its own hours in order to produce, as in the standard "yeoman farmer" model. However, the same equilibrium conditions for output and prices are obtained if we assume that output is actually sold by firms that hire labor in a competitive labor market. [32] In this case, the model also determines the equilibrium wages in each segmented labor market.

This specification implies that $v' = g'/f'$, as a result of which

$$\omega \equiv \frac{v''Y}{v'} = \left(\frac{g''H}{g'} - \frac{f''H}{f'}\right)\frac{f}{f'H}. \tag{4.1}$$

Wage-taking behavior by households implies that the wage in labor market i must satisfy $w^i = g'(H^i)/u'(C)$. Log-linearizing this around the steady-state values of the state variables, and aggregating across labor markets, this yields

$$\frac{g''H}{g'}\frac{f}{f'H} = \epsilon_{wY} - \sigma,$$

where $\frac{g''H}{g'}$ is the inverse of the Frisch elasticity of labor supply with respect to the real wage and ϵ_{wY} is the elasticity of the average real wage with respect to variation in Y, in the case of variations in output that are not associated with shifts in preferences or technology. Various instruments for such changes in economic activity are possible. Christiano, Eichenbaum, and Evans (1996) show, using a structural VAR to identify monetary policy shocks, that an increase in the federal funds rate that leads to a 0.4 percent reduction in output reduces real wages by about 0.1 percent, suggesting an elasticity of 0.25. Rotemberg and Woodford (1996) use a VAR to study the effects of oil price increases, and find that an oil shock that lowers output by about 0.25 percent lowers real wages by 0.1 percent, suggesting an elasticity of 0.4. Given this range of estimates, we set $\epsilon_{wY} = 0.3$. [33]

Finally, assuming an isoelastic (Cobb-Douglas) production function, with an elasticity of output with respect to hours worked given by η, one finds that $-f''f/(f')^2 = (1-\eta)/\eta$. Price-taking behavior on the

[32] It is necessary, under this interpretation, to assume that the firms that sell goods the prices of which are chosen at different dates also hire from distinct labor markets, so that their wages need not move together despite competition in each of the segmented labor markets. This is not inconsistent with the assumption of competition, since there might be several firms at each "location" that share a local labor market, and that all change their prices at the same time.

[33] Note that this elasticity of the real wage with respect to variations in aggregate output agrees with that measured by Solon, Barsky and Parker (1994), who do not instrument for technology or preference shocks.

part of firms implies furthermore that the wage in labor market i must satisfy $w^i = f'(H^i)$. It follows from this that the share of wages in the value of output in that sector should equal η/μ^i, where μ^i is the gross markup of price over marginal cost in sector i (Rotemberg and Woodford, 1997a). If markups are on average modest in size, then the value of η should not be much larger than the average labor share. We accordingly set $\eta = 0.75$, implying that $-f''f/(f')^2 = 0.33$. Substituting this, we find that $\omega = 0.63 - \sigma$. This restriction is imposed in our estimation of the values of these two parameters.

The parameter, β, is "calibrated" as well, not because it is not identifiable from the impulse responses, but because it is identifiable more directly from *first* moments of our data, and thus cannot plausibly be treated as a free parameter when trying to fit the second moments. Our model implies that β^{-1} should equal the gross real rate of return. Since this equals approximately 1.01 on average, we set β equal to 0.99.

We then choose values for our remaining three structural parameters, κ, σ, and ϕ, to ensure that the theoretical impulse responses are as close as possible to the empirical ones. We focus on the responses in the first four quarters after the shock on the ground that we have the most confidence in the estimated impulse responses for the first year, and that it is only for these initial responses that we can reject the hypothesis that monetary policy shocks are irrelevant for our three variables. We thus choose the theoretical parameters that minimize the sum of squared differences between the theoretical and empirical impulse responses of output, inflation and interest rates for quarters 0 through 4 following a monetary policy shock. In this optimization, we give equal weight to the three discrepancies that are depicted in Figure 1.

The parameters that minimize this criterion function are $\kappa = 0.024$, $\sigma = 0.16$ and $\psi = .88$, which in turn imply $\gamma = 0.63$, $\omega = .47$ and $\theta = 7.88$. With the possible exception of ω, these parameter values are all relatively plausible. [34] The estimate of $1/\sigma$ is substantially greater than the intertemporal elasticity of substitution typically found in the literature that analyzes non-durable consumption purchases. However, our estimate of $1/\sigma$ indicates the elasticity of expected *output* growth with respect to the expected real return. Thus, we should expect a lower σ in our model since the purchases of both consumer durables and investment goods are likely to be more interest-sensitive.

We can use (4.1) together with our estimate of ω to obtain an estimate of the Frisch elasticity of labor

[34] The plausibility of these parameters runs counter to the suggestion of Chari, Kehoe and McGrattan (1996) that models of this type are unable to reproduce the empirical persistence of output responses to monetary shocks without implausible parameters.

supply. Given that our assumptions imply that $-f''f/(f')^2$ is equal to .33, the resulting Frisch elasticity is 9.5. This is certainly higher than the estimates obtained from microeconomic studies like Pencavel (1986) and Card (1994). It is important to stress that this high labor supply elasticity is not necessary to match the empirical responses of the three series we have focused on. We could alternatively have imposed a Frisch elasticity of only 1.0, and still have obtained the same theoretical responses for our series, by maintaining the same σ and κ. The parameter ω would then have to equal 1.66, while γ and θ would be 0.71 and 7.61 respectively. Our reason for preferring the higher labor supply elasticity is that it rationalizes the relatively weak observed response of real wages to a monetary disturbance. At the same time, we recognize the need for more research to reconcile the macroeconomic response of real wages with the microeconomic evidence concerning household labor supply.

Finally, the estimated elasticity θ of the demand curve faced by a typical supplier is quite plausible. It implies a degree of market power that results in prices being set at a level 15 percent higher than marginal cost on average. The estimated elasticity of demand is thus neither so low as to imply implausibly large markups, nor so high as to make it implausible for firms to stagger their price adjustments.

In addition to displaying the empirical impulse responses to a monetary policy shock together with their confidence bands, Figure 1 also gives the theoretical responses corresponding to our estimated parameter values. (In each panel, the theoretical response is the solid line, while the estimated response is the middle dashed line.) As this figure indicates, the theoretical responses of output and interest rates match very closely the estimated responses. In particular, it is worth noting that our model accounts for both the magnitude and the degree of persistence of the effects on output of such a shock.

Our ability to match the output response may seem surprising given that the nominal interest rate has essentially returned to its steady state value by the time output finally falls. However, the positive shock to interest rates lowers inflation and, as a result, raises the real interest rate for some time. This makes the output fall consistent with the IS equation (3.14). Because the increase in real interest rates is small relative to the fall in output, the value of σ that rationalizes these relative movements is relatively small.

The response of inflation is matched well for the first few quarters. Subsequently, inflation reverts more quickly to its mean in the theoretical response, whereas the estimated response of inflation is much more protracted, if one can believe the point estimates. The problem is reminiscent of the criticism of the Taylor

(1980) model of overlapping wage contracts by Fuhrer and Moore (1995a) as unable to explain inflation persistence. However, the confidence intervals indicate the response of inflation is very poorly estimated in our sample, so that it is difficult to say that the data reject this aspect of the predictions of our model.

It is also worth remembering that the predictions of our model on this score (as on others) are themselves uncertain. In particular, they depend on the estimated coefficients of the monetary policy rule, which can hardly be estimated with great precision. Our method, which has estimated the monetary policy rule without any reference to the implications of the coefficients of this rule for the nature of the theoretical impulse responses of output and inflation to a monetary shock, makes it particularly unlikely that the theoretical impulse responses will match the unrestricted VAR estimates. A joint estimation strategy (in which the coefficients of the policy rule and the structural parameters are jointly chosen so as to match theoretical with estimated impulse responser) might well improve the fit of the impulse response of inflation.

5 Identification of Shock Processes

In this section, we construct time series for the three stochastic disturbances ϵ_t, $\hat{G}_t$ and $\hat{Y}_t^s$. We further show how the VAR can be used to infer the stochastic process that generates these variables. Finally, we show how to construct the response of our three endogenous variables to the shock processes for any given monetary policy rule. This allows us to construct counterfactual histories that, according to our model, would have taken place if the monetary authority had followed a different rule.

Note first that (2.3) can be premultiplied by T^{-1} to yield

$$\bar{Z}_t = B\bar{Z}_{t-1} + U\bar{e}_t \tag{5.1}$$

where the matrix U consists of zeros except that its upper left 3 by 3 block consists of a lower triangular matrix with ones on the diagonal. Letting i_n denote the vector whose nth element is a one and whose other elements are zero, the historical time series for the monetary policy shock ϵ_t can be derived from the relation

$$\epsilon_t = (i_1)'(\bar{Z}_t - B\bar{Z}_{t-1}). \tag{5.2}$$

We denote by $\tilde{Z}_t$ the vector whose elements are the model's theoretical predictions concerning the elements of $\bar{Z}_t$, the vector of historical time series. (The need for separate notation will become clear when we introduce our simulation method; we will then distinguish between the *historical* law of motion of $\bar{Z}_t$ and the *theoretical* law of motion for $\tilde{Z}_t$.) The structural equations of the previous section, (3.11), (3.14), and (3.27) can be written in terms of the $\tilde{Z}$'s as

$$M'\tilde{Z}_t - N'\sum_{j=1}^{\infty} E_{t-1}\tilde{Z}_{t+j} \quad = \quad \hat{G}_{t+1} \tag{5.3}$$

$$P'E_{t-1}\tilde{Z}_t + R'E_t\tilde{Z}_{t+1} \quad = \quad \hat{Y}^s_{t+1} + \frac{\sigma}{\sigma+\omega}E_t(\hat{G}_{t+2} - \hat{G}_{t+1}) \tag{5.4}$$

Here (5.3) is obtained from (3.14) by substituting (3.11) to eliminate $\hat{r}^l_t$, and recalling that, according to the model, $\hat{Y}_{t+1}$ and $\hat{\pi}_{t+1}$ are in the period t information set. Equation (5.4) is similarly derived from (3.27). Note also that the time subscript is also increased by one in both equations, so that the right-hand sides of both equations represent exogenous shocks in period t.

Assuming that the VAR correctly captures the stochastic process followed by the variables in Z_t, one can thus reconstruct the time series for $\hat{G}_t$ and $\hat{Y}^s_t$ by assuming that agent expectations coincide with the VAR forecasts. [35] For instance, this implies that, under the policy regime that generates the historical data, agents' forecasts $E_{t-1}\tilde{Z}_t$ are equal to $B\bar{Z}_{t-1}$. It then follows that we can reconstruct historical time series for $\hat{G}_t$ and $\hat{Y}^s_t$ using

$$s_t \equiv [\hat{G}_{t+1}\ ,\ \hat{Y}^s_{t+1}]' = C\bar{Z}_{t-1} + D\bar{e}_t, \tag{5.5}$$

$$C = \begin{bmatrix} M' - N'B(I-B)^{-1} \\ P' + R'B - \frac{\sigma}{\sigma+\omega}(N'B - M'(I-B)) \end{bmatrix} B, \qquad D = \begin{bmatrix} M' \\ R'B + \frac{\sigma}{\sigma+\omega}(M'(I-B) + N'B^2(I-B)^{-1}) \end{bmatrix} U.$$

These series are exogenous according to our model so their realizations can be held fixed in simulations of counterfactual history under alternative policy regimes.

Because our model involves forward-looking behavior, such simulations also require that we specify agents' beliefs regarding the stochastic processes generating the shock series. Given our proposed identification of the shock series, there is an obvious model-consistent specification of such beliefs: agents regard the vector of shocks s_t, from some period $t = 1$ onward, as being generated by (5.5), together with the law of motion (5.1) for the stochastic process $\bar{Z}_t$, given a specification of the initial condition $\bar{Z}_0$ and the distribution from which the white-noise innovations $\bar{e}_t$ are drawn each period.

A complete simulation model would therefore consist of a specification of a monetary policy rule of the form (2.1), together with a specified distribution for the monetary policy shocks ϵ_t; the structural equations (5.3) – (5.4); and the law of motion for the real disturbances given by (5.1) and (5.5), together with the

[35] In inferring our shock series from the residuals of our structural equations, we extend the methodology of Parkin (1988).

distribution of the shocks $\bar{e}_t$. In fact, to put the structural equations in a form that can be solved using standard solution algorithms such as that of King and Watson (1995), it is convenient to rewrite (5.3) as

$$M'\tilde{Z}_t = (M' + N')E_{t-1}\tilde{Z}_{t+1} + \hat{G}_{t+1} - E_{t-1}\hat{G}_{t+2}. \tag{5.6}$$

The complete model then consists of equations (2.1), (5.4), and (5.6) to determine the endogenous variables, laws of motion (5.1) and (5.5) for the real disturbances, and specified distributions for the shocks. Such a model would determine the evolution of $\{\tilde{Z}_t, \bar{Z}_t, s_t\}$ given initial conditions $\{\tilde{Z}_0, E_0\tilde{Z}_1, \bar{Z}_0\}$ and the white noise shocks $\{\epsilon_t, \bar{e}_t\}$. The model could be used to simulate counterfactual history if we supply the historical shock series (computed above) for $\{\epsilon_t, \bar{e}_t\}$. In such a simulation, it is natural to specify initial conditions $\tilde{Z}_0 = \bar{Z}_0$ and $E_0\tilde{Z}_1 = B\bar{Z}_0$, where $\bar{Z}_0$ represents the historical data for the period immediately prior to the beginning of the simulation. [36] The simulation model can also be used to generate predictions about second moments of the elements of $\tilde{Z}_t$, in the stationary equilibrium associated with an arbitrary policy regime, by taking expectations over the distribution of possible realizations of the shocks.

This method of simulation would have the property that, if the assumed monetary policy rule is the estimated historical one, and one feeds in the constructed historical shock series, the predicted series $\{\tilde{Z}_t\}$ coincides exactly with the historical data series $\{\bar{Z}_t\}$. Thus this method of identification of the shocks would allow a complete reconstruction of the historical data as the unfolding of a stationary rational expectations equilibrium. The intuition for this result is straightforward. [37] Use of the monetary policy rule implied by the VAR ensures that we can perfectly reconstruct the behavior of interest rates as long as we are also able to match the behavior of output and inflation. Moreover, the proposed method for constructing $\hat{G}_t$ and $\hat{Y}_t^s$ ensures that, when current and past values of output and inflation are equal to their actual values, the model is consistent with the VAR's predictions concerning future movements in these variables. Thus, as long as we start from an initial condition in which the model and the data agree, we are able to rationalize the evolution of all three series.

It is important to stress that this ability to reconstruct the historical series does not imply that our structural model is correct. Indeed, we would be able to reconstruct the historical series of output, inflation

[36] This specification of initial conditions makes sense if we assume that the stationary equilibrium that results in the law of motion indicated by the VAR has been in effect up through period zero, and has been in effect long enough for all of the elements of Z_0 to have been determined under that regime.

[37] See Appendix 2 for a more detailed argument.

and interest rates in the way described even for arbitrary values of the structural parameters. The model does, however, have testable predictions (which is what allows us to identify several of the structural parameters from moments of the data). Specifically, the identified monetary policy shock process ϵ_t should be orthogonal to the real disturbances $\hat{G}_t$ and $\hat{Y}_t^s$ at all leads and lags. This restriction is not imposed in the above method of identification of the historical shock series. In fact, the constructed real and monetary shock series are not orthogonal, unless the matrices in (5.1) and (5.5) happen to imply that the shocks $\bar{e}_{1t}$ have no effect upon s_t. This is not the case for our estimated model, for otherwise the theoretical impulse responses to a monetary shock would perfectly match those estimated using the VAR. If we just wanted to simulate counterfactual histories using the model, we could ignore the correlation between real and monetary disturbances by supposing that $\bar{e}_{1t}$ and ϵ_t are independent random variables, though drawn from identical distributions. But the resulting model of the data generating process is still subject to an internal inconsistency, in that if the model were true, it should *not* be possible to identify the four independent shocks ϵ_t, $\bar{e}_{1t}$, $\bar{e}_{2t}$, and $\bar{e}_{3t}$, from a three-variable VAR of the kind that we use.

This inconsistency can be eliminated by modifying the above procedure for identification of the shocks. Specifically, we assume that the real disturbances s_t are generated by the laws of motion (5.1) and (5.5), but we assume that in these equations $\bar{e}_{1t} = 0$ at all times, whereas $\bar{e}_{2t}$ and $\bar{e}_{3t}$ are again assumed to be white-noise random terms, independent of each other and of the monetary shocks, and drawn each period from distributions identical to those of the corresponding VAR residuals. In other words, we replace the laws of motion (5.1) and (5.5) by

$$Z_t^{\dagger} = BZ_{t-1}^{\dagger} + U^{\dagger} e_t^{\dagger}, \tag{5.7}$$

$$s_t = CZ_{t-1}^{\dagger} + D^{\dagger} e_t^{\dagger}, \tag{5.8}$$

where $e_t^{\dagger} \equiv [\bar{e}_{2t} \ , \ \bar{e}_{3t}]'$, $U^{\dagger}$ is equal to U with the first column deleted, and $D^{\dagger}$ is equal to D with the first column deleted. This alternative stochastic simulation model has only three independent exogenous disturbances each period (two "real" shocks and the monetary policy shock). Under the assumption that such a model is correct, analysis of the VAR according to our method should (at least asymptotically) recover exactly the true stochastic processes generating the various shocks.

Correspondingly, for purposes of counterfactual historical simulations, we do not construct the historical shock series by substituting the historical data $\bar{Z}_{t-1}$ and VAR residuals $\bar{e}_t$ into (5.5). Instead, we use the historical VAR residuals $\bar{e}_{2t}$ and $\bar{e}_{3t}$, to construct the series $e_t^\dagger$, and then simulate (5.7), starting from an initial condition $Z_0^\dagger$ equal to the historical $\bar{Z}_0$, to generate the series $\{Z_t^\dagger\}$. The series $e_t^\dagger$ and $\{Z_t^\dagger\}$ are then substituted into (5.8) to construct the historical sequence of real disturbances. The method thus amounts to using not the residuals of the structural equations (5.3) – (5.4), but rather the component of those residuals that is orthogonal to the identified monetary policy shock and to all its lags. With this modification, the simulated paths using the estimated monetary policy rule no longer exactly equal the actual paths. The extent to which the simulated data track the actual data then becomes a test of the accuracy of our structural model. As we show in the next section, our estimated model (with historical shock series constructed in the way just described) does quite a good job of accounting for the variations in real GDP, inflation, and the Federal funds rate since 1980, despite the specification error that is indicated by its failure to perfectly match the estimated impulse responses in Figure 1.

Another way of assessing the degree of correspondence between our model and the U.S. data is to compare the empirical auto- and cross-correlation functions for our three series with the corresponding predictions of the stochastic simulation model, with the stochastic processes for the shocks specified as above. This comparison is shown in Figure 2, where in each of the nine panels, the solid line indicates the theoretical cross-correlation function, and the dashed line the cross-correlation function implied by the unrestricted VAR characterization of the U.S. data. It is clear that our model accounts for the second moments of the data to essentially the same degree as does the unrestricted VAR. Among other things, our model is able to perfectly reproduce the degree of persistence of inflation despite the criticism of a Taylor-style model of staggered price setting on this score by Fuhrer and Moore (1995a). It is also able to match the negative correlation of output with lagged nominal interest rates that King and Watson (1996) find cannot be explained by an optimizing model with Calvo-style staggered price setting similar to our own. Fuhrer (1997b) also draws attention to this correlation and suggests that a "backwards looking" IS curve is needed to explain it. We suspect that some of the difficulties faced by these previous authors in reconciling models of optimizing consumers and of staggered price setting with these features of the data may relate to the imposition of *a priori* restrictions upon the exogenous shock processes for which there is no theoretical justification.

6 Simulation of Alternative Monetary Policy Rules

In this section we briefly illustrate how the simulation model built up in the previous two sections can be used to predict the consequences of alternative possible monetary policy rules of the form (2.1). We first display in Figure 3 the consequences of following a feedback rule with the coefficients of the estimated historical policy rule. In each panel of this figure, the dashed line represents the actual data for the series in question, the solid line represents the simulation of our model assuming the historical sequence of monetary policy shocks as well as the historical series for the real shocks, and the dash-dot line represents a simulation in which the historical feedback rule for the funds rate is followed, but with the monetary policy shocks ϵ_t set equal to zero.

One observes, first, that the first two plots track one another quite closely in each panel. Thus our model does quite well at accounting for the historical paths of output, inflation, and the funds rate, despite the fact that the theoretical and estimated impulse responses to a monetary shock do not perfectly coincide. The only very noticeable failure of our simulation model is in tracking the level of inflation from 1993 onward. The dash-dot line differs somewhat more from the solid line; this indicates the consequences, according to our model, of the random disturbances to monetary policy. Monetary policy shocks clearly have played some role; in particular, our simulations indicate that unexpectedly tight monetary policy in early 1982 deepened the 1982 recession, and that unexpectedly loose policy stimulated real activity in the period 1992-93.

On the other hand, the simulations imply that relatively little of the variability in output or inflation in this period can be attributed to the monetary policy shocks. Table 2 shows this in a different way by reporting the predicted stationary variances of interest rates, output and inflation under a variety of alternative policy rules. The first two rows of the table give these statistics for two regimes corresponding to the estimated feedback rule with and without the stochastic term. Comparison of the numbers in these two rows shows that, in the simulation of the "historical" policy regime, the monetary policy shocks are responsible, over the long run, for only 5.0% of the variance of deviations of real output from trend, and (perhaps more surprisingly) for only 1.3% of the variance of inflation.

But these results do *not* imply that monetary policy is unimportant. Nor do they necessarily absolve the Fed from any blame for the instability of output or inflation. What they mean is that it is the *systematic* part of recent monetary policy that has been of significance for recent economic performance, not the stochastic variation in Fed policy (which, according to our estimates, has been minimal).

One can gain some understanding of the effect of alternative systematic monetary policies by comparing the predicted consequences of simple feedback rules of the kind discussed by Taylor (1993b),

$$\hat{r}_t = \theta_\pi \hat{\pi}_t + \theta_y \hat{Y}_t,$$

under alternative values for the coefficients θ_π and θ_y. (In this equation, we write $\hat{r}_t \equiv r_t - r^*$, and $\hat{\pi}_t = \pi_t - \pi^*$. [38]) Rows 3-5 of Table 2 report predicted moments of the data for three possible choices of these coefficients. In row 3, we consider a "Taylor rule" with $\theta_\pi = 1.5$, $\theta_y = 0.5$, values which are close to those used by Taylor to characterize current policy (his exact coefficients are in footnote 2). [39] According to our model, adherence to this rule would make a difference since both inflation and interest rates would be significantly more variable.

Even sharper contrasts between policy rules are possible if we vary the coefficients of the "Taylor rule". Row 4 considers a policy in which instead $\theta_\pi = 1$, and $\theta_y = 5$. The increased response to deviations of output from trend is predicted to reduce the variance of output fluctuations to about a tenth of its value under the historical policy regime. This stabilization of output, however, is accompanied by increased volatility of inflation and short-term nominal interest rates. (A counterfactual historical simulation assuming this policy rule is shown in Figure 4.) For purposes of contrast, row 5 of the table considers a "Taylor rule" in which $\theta_\pi = 10$ and $\theta_y = 0$. The increased response to fluctuations in inflation is predicted to reduce the variance of inflation to about an eighth of its value under the historical policy. Inflation stabilization, however, is accompanied by increased volatility of both output and interest rates.

These comparisons show that according to our model, monetary policy matters a great deal for the behavior of both output and inflation since either inflation or output can be stabilized to a much greater extent than it has been historically. This raises the obvious of which policy rule results in more *desirable* patterns of fluctuations. We take this up in the next section.

[38] We suppress the "target" values from our notation for the policy rule, since we are here interested solely in the rule's implications for the kind of fluctuations that occur in response to shocks. One should remember that the policy rule also determines a long-run average rate of inflation π^*, but (subject to a bound discussed in the next section) the choice of π^* is independent of the way that one chooses for economy to respond to shocks. In the simulations reported in the figures, we assume the same values for π^* and r^* as under the estimated historical policy rule.

[39] The rule that we simulate here is not exactly Taylor's, since he assumes that the funds rate responds to the rate of inflation in the current and previous three quarters, while our rule assumes that it responds to the rate of inflation in the current quarter only.

7 The Welfare Loss from Price-Level Instability

We wish to consider the consequences of alternative monetary policy rules for the value achieved in equilibrium by the lifetime utility (3.1) of the representative household. The precise comparison that we propose to make is the following. Associated with any *stationary* rational expectations equilibrium of the kind discussed above (resulting from a time-invariant feedback rule for monetary policy) is an *unconditional* expected value of (3.1), averaging over all the possible histories of shocks that may have occurred prior to date zero. We propose to compare stationary equilibria in terms of the value of this unconditional expected utility. In this way, we take a "long-run" perspective in evaluating alternative policy rules; we do not consider the advantages that a particular rule may have that result from the nature of the particular fortuitous initial conditions that may exist at the time that one contemplates commitment to such a rule.

This objective is easily seen to be equivalent to maximization of the simpler objective function

$$W = E\{u(C_t) - \int_0^1 v(y_t(z))dz\}, \tag{7.1}$$

where E refers to the unconditional expectation. This objective averages the disutility of working across households at a point in time, because, from our "long-run" perspective, any given household is equally likely to be in the situation of any one of the households (who differ, after all, only in the sequence of times at which they have been able to change the prices of the goods that they sell). By including the integral over z in (7.1), we do not need to interpret the expectation operator E as referring to an average over possible histories of opportunities for an individual seller to change its prices, but only an average over possible histories of the aggregate shocks (*i.e.*, the disturbances to preferences and technology).

We furthermore simplify our analysis by taking a second-order Taylor series approximation to our objective (7.1). [40] This has the advantage of allowing us to derive an approximate loss function that can be evaluated using *only* the log-linear approximate characterization of the equilibrium, obtained by solving the equations derived in section 3. Another advantage is that we obtain a loss function that can be written as a weighted sum of contributions from the variances of various endogenous series, which allows direct comparison of our conclusions with the *ad hoc* loss functions typically assumed in the literature.

At the cost of not being able to evaluate the effect of monetary policy on the long run level of output,

[40] See Appendix 3 for details of this calculation.

we suppose that changes in monetary policy are accompanied by changes in the constant income tax rate τ so that this tax is optimal in each case. This ensures that, roughly speaking, the *average* level of output is optimal and independent of monetary policy. [41] Our idea here is to separate the issue of the welfare losses associated with *fluctuations* in output from those due to a sub-optimal *average* level of output, due (for example) to the presence of monopolistic competition or distorting taxes, and to make monetary policy responsible solely for the minimization of the former losses, assuming that other policy instruments will be used to ensure the desired average level of output. This assignment of tasks to policy instruments makes sense if, as a practical matter, the tax code can affect the long run level of output but cannot be adjusted rapidly enough to be used to ensure an optimal response to stochastic disturbances.

In Appendix 3 below, we show that a second-order approximation to W takes the form

$$W = -\frac{1}{2}u_c Y\left[(\sigma+\omega)\text{var}\{E_{t-2}(\hat{Y}_t - \hat{Y}_t^s)\} + (\theta^{-1}+\omega)E\{\text{var}_z\{\log y_t(z)\}\}\right] \tag{7.2}$$

plus terms that are of third or higher order in the amplitude of the shocks, and terms that are unaffected by monetary policy. (Such terms are similarly neglected in the expressions that follow.) Thus our measure of deadweight loss depends upon the variability of aggregate output around the "natural rate", but also upon the dispersion of output levels across producers of different goods. The second term, in turn, depends solely upon the degree of price dispersion, since the demand curve (3.15) implies that

$$\text{var}_z\{\log y_t(z)\} = \theta^2 \text{var}_z\{\log p_t(z)\}. \tag{7.3}$$

Finally, our price-setting equations imply that

$$E\{\text{var}_z\{\log p_t(z)\}\} = \frac{\alpha}{(1-\alpha)^2}\left[\text{var}\{E_{t-2}\pi_t\} + (1+\psi)\text{var}(\pi_t - E_{t-2}\pi_t) + \{E(\pi_t)\}^2\right]. \tag{7.4}$$

Thus the degree of price dispersion that exists on average increases with average inflation, with the variability of the rate of inflation forecasted two quarters in advance, and with the variability of unexpected inflation. Substituting (7.3) and (7.4) into (7.2) one obtains

$$W = -\Omega[L + \pi^{*2}], \tag{7.5}$$

where

$$L = \text{var}\{\pi_t\} + \psi^{-1}\text{var}\{\pi_t - E_{t-2}\pi_t\} + \Lambda\text{var}\{E_{t-2}(\hat{Y}_t - \hat{Y}_t^S)\}, \tag{7.6}$$

[41] Our assumption about the nature of long-run policy is made explicit in Appendix 3.

for certain coefficients $\Omega, \Lambda > 0$. The quantity $L + \pi^{*2}$ represents the measure of deadweight loss due to price-level instability that we shall use to evaluate alternative monetary policies. Here the loss measure L collects the terms that depend solely upon the degree of variability of inflation and the output gap, while π^{*2} is proportional to the deadweight loss due to nonzero inflation, even when it is perfectly steady.

It is worth noting that all three of the terms in (7.6) are directly related, in different ways, to inflation variability. For the analysis of optimal policy below, it is helpful to rewrite L so that it depends only on the stochastic process for the relative price variable $\hat{X}$. We show in the Appendix that the model's structural equations imply that (7.6) may be rewritten in the form

$$L = \text{var}(E_{t-2}\hat{X}_t) + \psi \text{var}[\hat{X}_t - E_{t-2}\hat{X}_t] + \frac{\Lambda}{\kappa^2}\text{var}[E_{t-2}(\hat{X}_t - \beta\hat{X}_{t+1})]. \tag{7.7}$$

This shows that the deadweight losses measured by L are zero if variations in $\hat{X}$ are eliminated (as we show below to be possible in principle). Thus a constant rate of inflation is both necessary and sufficient for achievement of the minimum value of $L = 0$. This means that, even though our proposed welfare criterion (7.2) assigns ultimate importance *only* to the efficiency of the level of real activity in each sector of the economy, it in fact justifies giving complete priority to inflation stabilization as opposed to output stabilization. [42]

We ignore for now the choice of the average inflation rate, and show that there is a policy that ensures that $\hat{X}_t$ is constant so that L is zero. [43] The only effects of a nonzero π^* on our equilibrium are to increase the interest rate in each period by π^* while $\hat{X}$ is increased by a constant as well. Substituting constants for both $\hat{X}_t$ and $\hat{\pi}_t$ into the structural equations (3.11), (3.14), and (3.26), we see that an equilibrium with steady inflation involves

$$\hat{Y}_t = E_{t-2}\hat{Y}_t^s + \hat{G}_t - E_{t-2}\hat{G}_t \tag{7.8}$$

and

$$E_{t-1}\hat{R}_t = \omega[(\hat{G}_t - \hat{Y}_t^s) - E_{t-2}(\hat{G}_t - \hat{Y}_t^s)] + \{\sigma E_{t-1}[(\hat{G}_t - \hat{Y}_t^s) - (\hat{G}_{t+1} - \hat{Y}_{t+1}^s)]\}. \tag{7.9}$$

[42] Note that the complete elimination of deadweight loss, measured by $L + \pi^{*2}$, if possible, would require a constant inflation rate of $\pi^* = 0$, and so a constant price level. Our finding that price stability is optimal in our model is closely related to King and Wolman's (1996) argument that, in their closely related model, price stability leads output to behave as it would if prices were flexible. Note that this conclusion depends upon a number of special features of the model developed here, in particular, upon the assumption that the existence of nominal price rigidities is the only distortion that prevents equilibria from necessarily being optimal.

[43] We do not analyze in this paper the form of the interest rate feedback rule that achieves this stationary equilibrium. The implementation issue is taken up in Rotemberg and Woodford (1997b), in the case of the policy referred to below as the "constrained optimal" policy.

where we have neglected the constants. It is then easily verified that all of the structural equations are satisfied if these two are, and $\hat{X}_t$ and π_t are constant. This establishes the possibility, in principle, of complete inflation stabilization.

Note that (7.9) only determines $E_{t-1}\hat{R}_t$. This is the only restriction upon the path of short-term nominal interest rates implied by price stability. To avoid adding unnecessary noise to interest rates, the central bank should also ensure that the actual value of $\hat{R}_t$ (as opposed to only its expectation at $t-1$) is given by the right-hand side of (7.9). This has the additional advantage that it economizes on the information requirements of the central bank since it makes interest rates at t depend only on the period $t-1$ information set.

Using our estimated processes for the real shocks, we now consider the fluctuations in output and interest rates that would obtain under such a first-best policy. Such an equilibrium would have required output to vary much more than it did under historical policy; Table 2 indicates that this minimum-L policy would have led to a variance of output nearly four times as large as the variance implied by the historical policy. This is mainly due to the highly volatile character of our inferred series for the "supply" disturbances $\hat{Y}_t^s$. According to our model, the reason output movements have been so much smaller under the actual policy is that the actual policy consistently "leans against the wind", so that the interest rate is increased whenever output rises. As has been pointed out by numerous authors (see, *e.g.*, Rotemberg (1983), Ireland (1996), and Aiyagari and Braun (1996)), such countercyclical policy is not appropriate in response to supply shocks of the sort represented by $\hat{Y}_t^s$. With sticky prices an increase in $\hat{Y}_t^s$, which reduces marginal costs, tends to lower prices and thus raise output. However, if interest rates are raised in response to the output increase, prices fall by less than marginal cost so this fall in prices is not sufficient for output to increase by the amount that $\hat{Y}_t^s$ increases. A policy of price stability requires that the monetary authority accommodate the increase in output required by the increase in $\hat{Y}_t^s$. As a result, output becomes more variable if $\hat{Y}_t^s$ is variable.

The path of interest rates that would have achieved complete inflation stabilization involves very large swings in interest rates and is remarkably choppy. In particular, as Table 2 indicates, the variance of the funds rate along this path is 733 (a standard deviation of 27 percentage points), while the variance of the funds rate under the historical policy is only 7.6 (a standard deviation only a tenth that large). One consequence of this is that such a policy is not consistent with a low average interest rate (and inflation rate) unless the nominal interest rate can be negative. Thus, as suggested by Summers (1991), the zero nominal interest-rate floor poses an impediment to stabilization policy with a low average level of inflation. [44] A standard deviation of 27 percentage points for the federal funds rate is clearly impossible unless the mean funds rate, and consequently the average inflation rate π^*, were to be substantially higher than have ever been experienced in the U.S. But that would imply other sorts of welfare losses, both those indicated by π^{*2} in (7.5) that result from the increased dispersion of prices across suppliers, as well as the more conventional "shoe-leather costs" (from which our model abstracts).

Thus our analysis leads to the conclusion that completely *stable* inflation is inconsistent with a low *average* inflation rate. This occurs in our model because we find there to be fairly large fluctuations in $\hat{Y}_t^s$. Insulating prices from the effects of these "supply shocks" requires very large swings in the interest rate if, as seems plausible and as is implied by our parameters, relatively large movements in interest rates are needed to change the prices that firms choose to set. Because of the costs of having to maintain a high average rate of inflation, it is likely to be desirable to accept some degree of inflation variability for the sake of reducing the size of the swings in nominal interest rates required in order to stabilize inflation. It is thus of interest to consider the costs, in terms of a higher value of L, that must be accepted in order to reduce the variability of the Federal funds rate. We thus consider the nature of the equilibrium that achieves the minimum possible value of (7.6) subject to a constraint of the form

$$\mathrm{var}(\hat{R}_t) \leq v_R. \tag{7.10}$$

Figure 5 shows the trade-off between the constraint parameter v_R and the minimum attainable level of the welfare loss L from inflation variability, which, as noted above, is expressed in units of the variance of inflation. This figure, which has the variance of the Federal funds rate on the horizontal axis, shows that the minimum attainable loss L is a convex function of v_R. In particular, the deadweight loss from inflation variation hardly rises as the variance of nominal interest rates is reduced from its optimal value of 733 to a value less than one seventh that size (corresponding to a standard deviation on the order of 10 percentage points). Further reductions in the volatility of interest rates have larger effects on welfare, but even reducing the variance of the funds rate to something near its recent level requires an increase in L that is only a

[44] The discussion of the model in section 3 has ignored this floor because that model abstracts altogether from the fact that money balances are held. If, however, we introduce liquidity services from non-interest-earning money into (3.1), we obtain an additional equilibrium condition representing the demand for money. This equilibrium condition will be inconsistent with an equilibrium nominal interest rate that is negative in any period. Note that the introduction of a demand for money need have no effect other than imposing $R_t \geq 1$ for all t upon the system of equations derived in section 3.

fraction of the deadweight loss associated with current policy, according to our estimates.

The question then becomes which point in Figure 5 is optimal once one recognizes that points with more volatile interest rates require higher average inflation rates. Here we pursue a crude approach to this problem by imposing the constraint that the average federal funds rate must be no smaller than k times the standard deviation of the funds rate, for some $k > 0$. [45] Since the average (or steady-state) funds rate is given by $r^* = \rho + \pi^*$, where ρ is the steady-state real rate of return (determined by the rate of time preference) this constraint can be written in the form

$$\rho + \pi^* \geq k\sigma(\hat{R}). \tag{7.11}$$

This constraint indicates how a higher degree of variability of the funds rate requires a higher target rate of inflation π^*. [46] We can then ask which, among the log-linear approximate equilibria characterized earlier that satisfy (7.11) in addition to the other equilibrium conditions, achieves the highest value of (7.5).

Under our log-linear approximate characterization of equilibrium, the value of L achieved in any stationary equilibrium is independent of the target inflation rate π^* around which inflation fluctuates. (Our approximate equilibrium conditions are derived by log-linearizing around a steady state with zero inflation, but continue to represent a valid approximation as long as π^* is small enough.) Thus we can consider, on the one hand, the lowest value of L consistent with a given value of $\sigma(\hat{R})$, given the structural equations used to derive Figure 5, and on the other hand, the lowest value of π^{*2} consistent with the value of $\sigma(\hat{R})$, given (7.11). The sum of these two terms, expressed as functions of $\sigma(\hat{R})$ is minimized by the unique value of $\sigma(\hat{R})$ for which

$$\{-2\lambda\sigma(\hat{R})\} + \{2k(k\sigma(\hat{R}) - \rho)\} = 0,$$

where the two terms in curly brackets represent, respectively, the derivatives of L and π^{*2} with respect to $\sigma(\hat{R})$, and where $-\lambda$ (with $\lambda > 0$) is the slope of the locus graphed in Figure 5. We assume $\rho = 3$ percent per year, as indicated by the "long-run" values r^* and π^* resulting from historical policy (according to our VAR), and $k = 2.26$, which is the largest value such that the historical equilibrium (according to our VAR) would satisfy (7.11). Then this condition is satisfied at the point in Figure 5 where $\lambda = .22$ and $\text{var}(\hat{R}) = 1.93$,

[45] In the event that the exogenous shocks have bounded supports, this is a *sufficient* condition for non-negativity of the funds rate at all times, and *necessary* within the class of equilibria in which the state variables all are *linear* functions of the shocks. This makes a natural case to consider, given our use of linearization methods here to characterize equilibria. In general, however, the optimal equilibrium subject to the constraint that the funds rate always be non-negative is unlikely to belong to the class of linear solutions.

[46] For example, the degree of funds rate volatility associated with complete inflation stabilization, so that $\sigma(\hat{R}) = 27$ percent, implies that the minimum possible value for π^* would be over 50 percent per year.

which requires a target inflation rate of at least $\pi^* = 0.15$ percent per year. This is a positive rate of inflation, as conjectured by Summers, but a trivially small one. Furthermore, this calculation neglects other costs of inflation, such as the costs of economizing on money balances that are emphasized in much of the literature. Taking account of such costs would only make the optimal average inflation rate even lower, as would the choice of a lower value for k.

Hence, because there exist only small gains in terms of reduction of L from raising the variability of interest rates beyond what is consistent with zero average inflation, the trade-off indicated in Figure 5 is favorable towards keeping inflation low. If there are additional, independent reasons for the Fed to prefer not to have a highly variable funds rate (as discussed, for example, by Goodfriend, 1991), then these would justify choice of a point even further to the left in Figure 5, and hence of an even lower target rate of inflation.

It is of particular interest to compare the constrained-optimal policy that minimizes L while keeping the average inflation rate at 0.15 percent to the actual policy of the Fed. This constrained-optimal policy has the immediate advantages that its average level of inflation as well as its variance of interest rates are considerably smaller than under the historical policy. Moreover, as the figure and Table 2 indicate, the loss from variability (L) under this policy is only about a fourth as large as under the estimated historical policy. The variance of output doubles relative to the historical policy, but according to the model this is desirable as well, since output is kept closer to the "natural rate".

To illustrate how the constrained-optimal policy would differ from actual policy, Figure 6 plots the impulse responses of output, inflation and interest rates to the two real shocks, under the constrained-optimal policy and under our estimate of historical policy. [47] The two shocks are orthogonalized as follows. One period t shock (the one considered in the second column) is defined as the innovation in period $t+1$ autonomous spending – *i.e.*, the part of G_{t+1} that could not be forecasted on the basis of output, inflation and interest rates through period t. The other period t shock (the one considered in the first column) is defined as that component of the innovation in the "natural rate" of output $\hat{Y}_t^s$ that is orthogonal to the innovation in autonomous spending. The specific shock represented in the first column is a "supply shock" that increases equilibrium inflation $\hat{\pi}_{t+1}$ (under the historical policy rule) by one percent more than one

[47] The method that we use to solve for the constrained-optimal policy is discussed further in Rotemberg and Woodford (1997b), where we also consider the nature of an interest-rate feedback rule that implements this policy, and so results in the impulse responses plotted here.

would have expected given the level of autonomous spending; the shock represented in the second column is an "autonomous spending shock" that raises real spending $\hat{Y}_{t+1}$ by one percent.

In each panel of the figure, the dashed line indicates the impulse response to the shock under the historical policy, while the solid line indicates the response that would occur under the constrained-optimal policy. Under historical policy, both types of shocks are inflationary, leading to increases in inflation that persist for many quarters, and so to a large eventual cumulative increase in the price level. Under the constrained-optimal policy, prices also rise slightly on impact. But what is striking about the constrained-optimal policy is that it ultimately leads prices to fall in response to these inflationary shocks. Thus inflationary shocks are accompanied by expected deflation in subsequent quarters. The result is that the constrained-optimal policy not only stabilizes *inflation* to a greater extent than under current policy, it also stabilizes the *price level* to a considerable extent.

Another striking difference between historical and constrained-optimal policy has to do with the effect of "supply shocks" on real activity. Under historical policy, real output is largely insulated from the effects of "supply shocks", which instead result in persistent fluctuations in inflation. Under the constrained-optimal policy, instead, an adverse "supply shock" results in relatively large increase in interest rates and a sharp transitory contraction of real activity. It is for this reason that (as Table 2 indicates) the constrained-optimal policy involves much greater output variability than has occurred under the historical policy.

By contrast, output movements in response to the autonomous spending shock are quite similar under historical and constrained-optimal policy. Historical policy has allowed output to respond to these shocks, but according to our model it is desirable for this to occur. The nominal interest rate rises less under the constrained-optimal policy, and returns more quickly to normal. However, real interest rates are not significantly lower during the transition back to the steady state, because this policy induces deflation as discussed above.

Table 2 also provides a nice contrast between this constrained-optimal policy and the "Taylor rule" whose coefficients are $\theta_\pi = 10$ and $\theta_y = 0$. These two policies induce about the same variance of inflation and output while also having similar losses from variability L. However, the simple "Taylor rule" achieves this by having interest rates react aggressively to inflation and this leads interest rates to be very volatile. Our constrained-optimal rule, by contrast, allows interest rates to be less variable by tailoring the dynamic response to shocks more appropriately.

8 Conclusions

This paper has provided a method for computing optimal monetary policy in the context of an optimizing model that fits the U.S. data nearly as well as an unrestricted vector autoregression. The two basic ingredients of this method are a vector autoregression of the variables of interest and an optimizing model that predicts the evolution of these variables. As long as the model can match the estimated impulse responses of the variables to a monetary shock, the method can be applied easily, because it is straightforward to fit the response of the model to the other shocks. Thus, the method can accommodate much richer vector autoregressions than the one we have considered, as well as more elaborate models. In this paper we have worked with a minimal model, both to show how this method can be applied and to show that even very simple optimizing models can fit the data rather well. Even so, it would be desirable to have a model that deals explicitly with investment and the resulting capital accumulation as well as with labor market variables.

In addition to providing a method of analysis, we have also been able to reach interesting conclusions regarding optimal monetary policy. In particular, the complete stabilization of inflation appears to require fairly large swings in interest rates which, given a zero interest-rate floor, require a high average interest rate and thus a high inflation rate. Thus, the existence of this floor limits the degree to which it is desirable to stabilize inflation. On the other hand, it appears to be possible, at least in principle, both to lower the average inflation rate and to stabilize inflation more than has been done historically in the U.S. While this requires that inflationary shocks still be allowed in increase inflation transitorily, such shocks must be followed by deflation shortly thereafter. The result is that neither surges in autonomous spending nor adverse supply shocks lead to long-run increases in prices.

Our specific conclusions as to the desirable responses of output, inflation, and interest rates to stochastic disturbances may well be sensitive to the particular optimizing model we have considered and, specifically, to the absence of other types of stochastic disturbances, such as time-varying labor-market distortions and changes over time in firms' desired markups of price over marginal cost. These are issues that only further investigation of other, more elaborate optimizing models can settle. Our main hope with this paper is precisely to shift the debate over optimal monetary policy so that it will involve different optimizing models, all of which fit the data reasonably well, instead of involving equations which fit well by construction but which have only a tenuous connection to explicit behavioral hypotheses at the microeconomic level.

9 Appendices

Appendix 1: Identification of Model Parameters

Here we discuss further the extent to which the parameters of our log-linearized model can be identified from the first and second moments of the joint stochastic processes for inflation, detrended output, and the Federal funds rate.

The parameters of our model of the historical policy regime may be usefully broken into three groups. Parameter set Φ_1 consists of the parameters $\{\mu_k, \phi_k, \theta_k\}$, $\sigma_\epsilon^2 \equiv \text{var}(\epsilon_t)$, and

$$c \equiv (1 - \sum_k \mu_k) r^* - \sum_k \phi_k \pi^* - \sum_k \theta_k y^*. \tag{9.1}$$

These are the parameters of the historical monetary policy rule (2.1). [48] Parameter set Φ_2 consists of the "structural" parameters $\alpha, \beta, \gamma, \sigma, \theta, \omega$ and y^*. These parameters specify aspects of the model that are independent both of the monetary policy rule and of the specification of the stochastic processes for the "real" disturbances $\hat{G}_t$ and $\hat{Y}_t^s$. These parameters relate to the preferences of the representative household $(\beta, \sigma, \theta, \omega, y^*)$, and to the lags involved in price setting (α, γ). Finally, parameter set Φ_3 consists of the elements of the matrices $B, U^\dagger, C, D^\dagger$, and the variances $\sigma_2^2 \equiv \text{var}(\bar{e}_{2t})$ and $\sigma_3^2 \equiv \text{var}(\bar{e}_{3t})$. These are the parameters of the laws of motion (5.7) – (5.8) for the exogenous disturbances $\hat{G}_t$ and $\hat{Y}_t^s$. (Recall that the variance-covariance matrix for $e_t^\dagger$ is assumed to be diagonal.)

We seek to identify these parameters, to the extent possible, from the first and second moments of the joint process for the variables r_t, π_t and y_t. Note that these moments may be completely summarized by a specification of (i) the unconditional means (r^*, π^*, y^*), (ii) the matrices T and A specifying the VAR representation (2.3) for the de-meaned series, and (iii) the variances of the three mutually and serially uncorrelated disturbances in the vector $\bar{e}_t$ in (2.3). (We assume in this discussion that the joint process for our three series admits of a finite-order VAR representation, as in (2.3).) These statistics may in turn

[48] Note that we do not list the "target" values (r^*, π^*, y^*) as independent parameters of the policy rule, since different triples that imply the same value of c represent identical policy rules. Because (r^*, π^*, y^*) are not separately identified, we may without loss of generality assume that the values used in representation (2.1) correspond to the actual long-run average values of the three series.

usefully be partitioned into three sets, corresponding to our three parameter sets. Set M_1 consists of the coefficients $\{\mu_k, \phi_k, \theta_k\}$ of the first row of equation (2.3), [49] the variance $\sigma_1^2 \equiv \text{var}(\bar{e}_{1t})$ of the residual in this equation, and the constant c (defined in (9.1), which represents the constant term in this regression when the means are not subtracted from the series. Set M_2 consists of the unconditional means $\rho \equiv r^* - \pi^*$, and y^*, and the predicted impulse responses of the inflation and output series to a unit innovation in $\bar{e}_{1t}$. Note that M_2, together with M_1, allows one to reconstruct the unconditional means of all three series, and to predict the impulse response of the funds rate to an innovation in $\bar{e}_{1t}$ as well. Set M_3 consists of the predicted impulse responses of inflation and output to unit innovations in $\bar{e}_{2t}$ and $\bar{e}_{3t}$, and the variances of these last two innovations. Note that M_3, together with M_1, allows one to predict the impulse response of the funds rate to these two innovations as well. Thus M_1, M_2, and M_3 jointly suffice to describe all first and second moments of the three data series.

The usefulness of these partitions consists in the following. According to our theoretical model, statistics M_1 should depend only upon parameter set Φ_1 – *i.e.*, upon the properties of the monetary policy rule and its stochastic disturbance term. (This is because of the assumed form (2.1) of the monetary policy rule, and because both inflation and output are predetermined variables, and as such unaffected by the current period's monetary policy shock.) Statistics M_2 – and, more generally, the unconditional means of all three series, and the impulse responses of all three variables to a monetary policy shock – should depend only upon parameter sets Φ_1 and Φ_2. (This is because, in our log-linear model, the average values of all variables, and the impulse responses to a monetary policy shock, are independent of one's assumptions about the nature of the other shocks, as long as they are exogenous, mean-zero, and independent of the monetary policy shock.) Finally, the remaining statistics M_3 are determined by the union of parameter sets Φ_1, Φ_2 and Φ_3.

This recursive structure determines our estimation strategy. We first choose the parameters in set Φ_1 so as to match statistics M_1; given these values, we then choose the parameters in set Φ_2 to match statistics M_2 as well as possible; and finally, given values for both Φ_1 and Φ_2, we choose the parameters in set Φ_3 to match statistics M_3. This is probably less efficient than simultaneously choosing all parameters so as to minimize a criterion that involves all of the data moments, but makes it clearer which aspects of the data determine which aspects of our econometric model.

[49] Here we use the notation of (2.1) for this equation, because it is in fact the same equation, except that the constants have been suppressed in (2.3).

The parameters in set Φ_1 are clearly identified, as these correspond directly to the elements of set M_1, as described above; indeed, we have not bothered to introduce distinct notation for coefficients such as μ_k when considered as parameters in Φ_1 as opposed to data statistics in M_1. We turn, then, to the identification of parameter set Φ_2, given Φ_1. These parameters are identified *only* insofar as they are identified in terms of their predictions for the moments M_2; for it is easily seen that there exist sufficient parameters in set Φ_3 to allow moments M_3 to be matched exactly, regardless of what values may have been assigned to the parameters in sets Φ_1 and Φ_2. [50]

The parameters in set Φ_2 may be further partitioned into subsets $\Phi_{2a} = \{\beta, y^*\}$ and $\Phi_{2b} = \{\alpha, \gamma, \sigma, \theta, \omega\}$, corresponding to a partition of M_2 into subsets $M_{2a} = \{\rho, y^*\}$ and M_{2b} consisting of the estimated impulse responses to a monetary policy shock. Once again, our model implies that the elements of M_{2a} should depend only upon the parameters in Φ_{2a} (for (3.8) implies that $\rho = -\log \beta$), and again we follow the recursive estimation strategy of choosing values for the elements of Φ_{2a} so as to match the data moments in M_{2a}, then taking Φ_{2a} as given when choosing values for the elements of Φ_{2b} so as to match the elements of M_{2b}. The elements of Φ_{2a} again directly correspond to the elements of M_{2a}, and so are uniquely identified.

Not all of the elements of Φ_{2b}, however, matter independently for the model's predictions regarding the impulse responses to a monetary policy shock. Note that the complete system of equilibrium conditions, consisting of (2.1), (3.11), (3.14), and (3.27), involves only the five parameters $\beta, \kappa, \omega, \sigma$ and ψ. The parameters α and γ do not matter independently, given the value of $\kappa \equiv (1-\alpha)(1-\alpha\beta)(\sigma+\omega)/\alpha(1+\omega\theta)$, but only through the composite parameter $\psi \equiv (1-\gamma)/\gamma\alpha$. One furthermore observes that the restrictions implied by these equations upon $E_t\hat{R}_{t+j}$, $E_t\pi_{t+j}$, and $E_t\hat{Y}_{t+j}$ for $j \geq 2$ are completely summarized by equations (3.28) and (3.29), which involve only the parameters β, σ, and κ. Thus, among the parameters of Φ_{2b}, only two parameters, σ and κ, have any effect upon the impulse responses for $j \geq 2$. [51]

The impulse responses of output and inflation at $j = 0$ are equal to zero by construction, and this is implied by our theoretical model as well, regardless of parameter values. The model also implies a zero response for output at $j = 1$ (though our estimate of this response if *not* constrained to equal zero, and the

[50] More generally, there exist sufficient parameters in set Φ_3 to allow the impulse responses of all three variables to shocks $\bar{e}_{2t}$ and $\bar{e}_{3t}$ to be matched exactly, given that parameters Φ_1 are chosen to match statistics M_1 exactly. This is why only the impulse responses to the monetary policy shock are presented in the paper, and used as a test of the ability of our model to match the properties of an unrestricted VAR model of the data.

[51] Our ability to fit these impulse responses as well as we do in Figure 1, then, indicates that our basic structure is reasonably consistent with the properties of the data, for we have allowed ourselves only two free parameters with which to adjust the predicted responses for $j \geq 2$.

point estimate is in fact not exactly zero). Thus there is only one additional statistic that can be used to identify any of the remaining parameters in Φ_{2b}, namely, the impulse response of inflation at $j = 1$. That means that we cannot even identify both of the remaining parameters, ω and ψ, that appear independently in the structural equations of the model. As noted in the text, we propose to fix the value of $\omega + \sigma$ on the basis of other information. Thus our strategy is to choose values of σ, κ, and ψ to match the statistics in M_{2b} as closely as possible, given the value of β determined by M_{2a} and the value of $\omega + \sigma$ determined otherwise (from the impulse response of the real wage, in an extended version of the model). Since only σ and κ affect the impulse responses for $j \geq 2$, we choose these two parameters in order to match those impulse responses as well as possible, and then choose ψ to match the impulse response of inflation at $j = 1$. One free parameter turns out to be enough to allow us to match this last statistic perfectly. [52]

The other element of Φ_{2b} that is "calibrated" (*i.e.*, fixed on the basis of information other than the moments of output, inflation, and interest rates) is the parameter α (given which, our estimate of ψ implies an estimate of γ). Unlike the parameter $\omega+\sigma$, which does play an independent role in the structural equations of our model, but is redundant for purposes of matching the impulse responses to a monetary policy shock, the parameter α plays no independent role in the equations that determine output, inflation, and interest rates in our model. Thus we are free to pick it on independent grounds (in this case, microeconomic survey evidence). One might think that, given that α and γ are not independently identified, there is no improvement in the fit of our model from introducing the additional free parameter γ, by assuming that not all suppliers have the same decision lag before their price changes take effect. However, simplifying the model by assuming either that all suppliers must choose their period t prices by the beginning of period t (the $\gamma = 1$ limit) or that all must choose them by the beginning of period $t-1$ (the $\gamma = 0$ limit) reduces our ability to fit the impulse response of inflation at $j = 1$.

Finally, we choose the elements of Φ_3 in a way that allows us to match the statistics in M_3. As noted earlier, we have a sufficient number of free parameters to allow the impulse responses to shocks $\bar{e}_{2t}$ and $\bar{e}_{3t}$ to be matched perfectly, regardless of the values assigned to the parameters in Φ_2. This can be seen as follows. As we show in Appendix 2, it is possible to choose the matrices B, U, C and D in such a way as to make the

[52] This would not necessarily be true, since the model implies the theoretical restriction $0 \leq \psi \leq 1$. But in our case it is possible to find an interior value of ψ that matches the statistic perfectly. Because of this, there is no loss of efficiency in estimating σ and κ on the basis only of the impulse responses for $j \geq 2$.

time series generated by the simulation model coincide exactly with the actual time series, when the VAR residuals $\{\bar{e}_t\}$ are substituted for the disturbances in equations (2.1) and (5.1). In this simulation model, the theoretical impulse responses to shocks $\bar{e}_{2t}$ and $\bar{e}_{3t}$ match perfectly the estimated responses to the VAR innovations $\bar{e}_{2t}$ and $\bar{e}_{3t}$, while the estimated response to VAR innovation $\bar{e}_{1t}$ exactly matches the *sum* of the theoretical responses to unit impulses in $\bar{e}_{1t}$ (a disturbance to equation (5.1) and ϵ_t (a monetary policy shock). The modified simulation model, that we actually use, eliminates the real disturbance $\bar{e}_{1t}$, but lets the effects of $\bar{e}_{2t}$ and $\bar{e}_{3t}$ be the same as in the model that exactly reconstructs the actual time series. In this model, there are exactly three orthogonal underlying disturbances (two real disturbances and the monetary policy disturbance), which can in principle be fully identified given time series observations on the three series, output, inflation, and the interest rate. By construction, the predicted responses to the real shocks $\bar{e}_{2t}$ and $\bar{e}_{3t}$ exactly match those that are implied by the estimated VAR. Because the statistics in M_3 can be perfectly matched, *regardless* of the parameter values that may have been chosen for sets Φ_1 and Φ_2, these additional statistics provide no information that can be used for identification of additional parameters in those earlier sets.

Appendix 2: Specification of Exogenous Disturbance Processes

Here we present further details of the specification of the exogenous disturbance processes $\{\hat{G}_t, \hat{Y}_t^s\}$ in our simulation model. The appendix has two aims. First, we explain why our method for identifying historical time series for the disturbance processes, when we use laws of motion (5.1) and (5.5), allows our simulation model to perfectly reconstruct the historical time series for output, inflation, and interest rates. Second, we give details of the numerical specification of the disturbance processes used in our simulation programs.

We first consider the simulation model that can perfectly reconstruct the historical time series. As explained in the text, the structural equations of this model consist of the estimated monetary policy rule of the form (2.1), structural equations (5.4) and (5.6), and laws of motion (5.1) and (5.5) for the real disturbances. Equations (2.1), (5.4) and (5.6) are a sufficient system of equations to determine the three endogenous variables, and hence the vector of endogenous variables $\tilde{Z}_t$. In the case that these equations have a unique bounded solution for $\tilde{Z}_t$, it can be written in the form

$$\tilde{Z}_t = J\tilde{Z}_{t-1} + K\bar{Z}_{t-1} + LE_{t-1}\tilde{Z}_t + u\epsilon_t + V\bar{e}_t, \tag{9.2}$$

where J, K, L, V are matrices and u is a column vector. We verify numerically, for the parameter values that we have assumed, that the equation system does indeed have a unique stationary solution, using a slightly modified version of the algorithm of King and Watson (1995). We then use the algorithm to compute the coefficients of the solution (9.2).

Now, when the laws of motion of the real disturbances are specified as in equations (5.1) and (5.5), and when the assumed coefficients in the monetary policy rule (2.1) are those implied by the VAR, the solution (9.2) that we obtain has the properties

$$J + K = [I - L]B, \qquad ui_1' + V = U. \tag{9.3}$$

This can be most easily seen as follows. Suppose that we replace equation (5.1) by

$$\tilde{Z}_t = \tilde{B}\bar{Z}_{t-1} + \tilde{U}\bar{e}_t, \tag{9.4}$$

where now the matrices of coefficients $\tilde{B}$ and $\tilde{U}$ need not coincide with the coefficients of the estimated VAR, and are to be solved for so as to satisfy our system of equilibrium conditions. (That is, we now treat both $\tilde{Z}_t$ and $\bar{Z}_t$ as endogenous variables, to be determined by the system of equilibrium conditions.) Instead of stipulating that $\tilde{B} = B$ and $\tilde{U} = U$ (the matrices of VAR coefficients), we substitute the conditions that

$$[I - L]\tilde{B} = J + K, \qquad \tilde{U} = ui_1' + V. \tag{9.5}$$

We then have a system of equations consisting of (2.1), (5.4) – (5.6), and (9.2) – (9.5), to be solved for the matrices of coefficients $J, K, L, V, u, \tilde{B}$, and $\tilde{U}$ that characterize the dynamics of $\tilde{Z}_t$ and $\bar{Z}_t$.

Because of the stipulations (9.5), this last system of equations has the property that if one starts with initial conditions in which $\tilde{Z}_0 = \bar{Z}_0$ and $E_0\tilde{Z}_1 = \tilde{B}\bar{Z}_0$, then *both* $\tilde{Z}_t$ and $\bar{Z}_t$ evolve according to (9.4), so that $\tilde{Z}_t = \bar{Z}_t$ and $E_{t-1}\tilde{Z}_t = E_{t-1}\bar{Z}_t = \tilde{B}\bar{Z}_{t-1}$ for *all* t. (Note that this conclusion is independent of the specification of the matrices C and D in (5.5).) We can then partitition the equilibrium conditions into a first set that determine the common evolution of $\tilde{Z}_t$ and $\bar{Z}_t$ in the case of initial conditions that imply the two series are forever identical, and a second set that determine the separate evolution of $\tilde{Z}_t$ under more general initial conditions. The first set of conditions simply determine the matrices $\tilde{B}$ and $\tilde{U}$, as these are the only coefficients that matter under those special conditions; the second set of conditions then determines the remaining coefficients in J, K, L, V and u. It is furthermore apparent that when we specify the matrices

C and D as given following equation (5.5), at least one solution to the first set of conditions is given by $\tilde{B} = B, \tilde{U} = U$. This is because the matrices C and D were derived exactly so as to satisfy these equilibrium conditions, under the assumption that the joint dynamics of $\tilde{Z}_t$ and $\bar{Z}_t$ (which variables had not at that point been distinguished) were given by the estimated VAR (5.1). Solving the second set of conditions under this assumption, we obtain a complete solution for the dynamics of $\tilde{Z}_t$ and $\bar{Z}_t$, with the property that $\tilde{B} = B, \tilde{U} = U$.

But then it follows that if we stipulate that $\tilde{B} = B, \tilde{U} = U$, *i.e.*, that the dynamics of $\bar{Z}_t$ are given by (5.1), rather than assuming (9.5), the *same* solution continues to work, and it is a solution that satisfies (9.5), or equivalently, (9.3). Thus our original system of equations possesses a solution that satisfies (9.3), as asserted. We furthermore verify numerically that, at least in the case of our assumed parameter values, this is the unique stationary solution. Finally, given a solution of this form, if we simulate the model starting from initial conditions in which the vector $\bar{Z}_0$ takes its historical value, and in which $\tilde{Z}_0 = \bar{Z}_0$ and $E_0\tilde{Z}_1 = B\bar{Z}_0$, and we use shock series $\bar{e}_t$ given by the VAR residuals, and set $\epsilon = \bar{e}_{1t}$ each period, we obtain a solution in which $\tilde{Z}_t = \bar{Z}_t$ at all dates, and both series evolve according to (5.1). Given that we start with the historical initial condition $\bar{Z}_0$ and use the historical residuals for the $\bar{e}_t$ series, simulation of (5.1) simply gives us back the historical data for the predicted path of both $\bar{Z}_t$ and $\tilde{Z}_t$.

The model that we actually use for our stochastic simulations, of course, is not this one, but the one that replaces the laws of motion (5.1) and (5.5) by (5.7) and (5.8), as explained in the the text. The resulting simulation model differs from the one described above only in the suppression of the $\bar{e}_{1t}$ shock. This change has no effect upon the predicted impulse responses to the $\bar{e}_{2t}$ and $\bar{e}_{3t}$ shocks (which are in any event orthogonal to the $\bar{e}_{1t}$ shock), nor upon the validity of our identification of those shocks with the second and third orthogonalized VAR residuals (assuming the data to be generated by the theoretical model). Hence this method of construction of the laws of motion for the real disturbances guarantees that the predicted impulse responses to shocks $\bar{e}_{2t}$ and $\bar{e}_{3t}$ (the two shocks that are orthogonal to the monetary policy shock, in our structural VAR) match identically our estimated impulse responses. Hence the ability to match impulse responses other than the responses to the monetary policy shock provides no test of the validity of our model structure.

The resulting laws of motion for the exogenous disturbances imply a significant degree of variability in these terms. For example, the standard deviations of $\hat{G}_t$ and $\hat{Y}_t^s$ are 29.5 and 13.7 respectively. (These values are considerably larger, for example, than the estimated standard deviation of $\hat{Y}_t$, which is only 2.1.) These two disturbances are slightly negatively correlated with one another: corr($\hat{G}_t$, $\hat{Y}_t^s$) = -.03. They are both somewhat positively correlated with output fluctuations: corr($\hat{G}_t$, $\hat{Y}_t$) = .35, corr($\hat{Y}_t^s$, $\hat{Y}_t^s$) = .15.

The high volatility of the $\hat{G}_t$ process does not mean that innovations in the $\hat{G}_t$ process are large; rather, the fluctuations exhibit a high degree of persistence (the coefficient of serial correlation is .92). The volatility of the $\hat{Y}_t^s$ process, instead, is largely due to a highly volatile transitory component. The coefficient of serial correlation of this process is only .15. Even more strikingly, the standard deviation of $E_{t-2}\hat{Y}_t^s$ is only 4.3; more than 90% of the variance of $\hat{Y}_t^s$ is due to the component that is not forecastable a quarter in advance. The less volatile, forecastable component of variations in "potential" output is much more strongly correlated with deviations of output from trend: corr($E_{t-2}\hat{Y}_t^s$, $\hat{Y}_t$) = .59. The fact that, according to our accounting for recent U.S. aggregate fluctuations, deviations of output from trend are associated to a large extent with fluctuations of the forecastable component of "potential" output is probably the reason why our welfare calculations indicate that policy rules that would do better at stabilizing output would not represent welfare improvements.

Appendix 3: Derivation of the Utility-Based Loss Function

Here we present further details of the derivation of equations (7.2), (7.5), (7.6), and (7.7), which describe our utility-based loss function $L + \pi^{*2}$. We begin with the derivation of (7.2) as a second-order Taylor series approximation to (7.1). Note that our objective function is of the form $W \equiv E[w_t]$, where w_t is the average utility flow (integrating over the continuum of households) each period. This utility flow may be written as a function solely of the pattern of real activity $\{y_t(z)\}$ within a period, and the exogenous shocks:

$$w_t = u(Y_t - G_t; \xi_t) - \int_0^1 v(y_t(z); \xi_t) dz. \qquad (9.6)$$

We begin by considering a Taylor series expansion for each of the two terms in this expression, expanding around the levels of output $y_t(z) = \bar{Y}$ for each z, and the values $G_t = \bar{G}, \xi_t = 0$ for the exogenous shocks. Here $\bar{Y}$ represents the level of output in an *optimal* steady state; it represents the constant equilibrium level of output in an equilibrium with no variation in the values of G_t and ξ_t around their steady-state values, a constant price level, and a tax rate $\tau = \tau^* \equiv -(\theta - 1)^{-1}$ that perfectly offsets the distortion resulting from firms' monopoly power. (As we shall see, our loss function takes an especially simple form in this case, and we wish to direct attention to the terms in it that survive even under these ideal circumstances. We leave for further work the analysis of how the welfare effects of monetary policy change when one considers possible interactions between monetary policy and distortions other than the one resulting from sluggish nominal price adjustment.) The steady-state value $\bar{G}$ is chosen to equal $E[G_t]$, and the shocks ξ_t are normalized so that $E[\xi_t] = 0$; thus the steady-state values of the exogenous variables equal their unconditional means.

A second-order Taylor series expansion for the first term on the right-hand side of (9.6) is given by

$$\begin{aligned} u &= u(\bar{C};0) + u_C \cdot (C_t - \bar{C}) + u_\xi \cdot \xi_t \\ &\quad + \frac{1}{2}u_{CC}(C_t - \bar{C})^2 + u_{C\xi} \cdot (C_t - \bar{C})\xi_t + \frac{1}{2}u_{\xi\xi}\xi_t^2 + \mathcal{O}(\|\xi\|^3) && (9.7) \\ &= u(\bar{C};0) + u_C\bar{Y} \cdot (\hat{Y}_t + \frac{1}{2}\hat{Y}_t^2 - \tilde{G}_t) + u_\xi \cdot \xi_t \\ &\quad + \frac{1}{2}u_{CC}\bar{Y}^2 \cdot (\hat{Y}_t - \tilde{G}_t)^2 + u_{C\xi}\bar{Y} \cdot (\hat{Y}_t - \tilde{G}_t)\xi_t + \frac{1}{2}u_{\xi\xi}\xi_t^2 + \mathcal{O}(\|\xi\|^3) && (9.8) \\ &= u_C\bar{Y} \cdot \hat{Y}_t + \frac{1}{2}[u_C\bar{Y} + u_{CC}\bar{Y}^2] \cdot \hat{Y}^2 \\ &\quad - u_{CC}\hat{Y}^2 \cdot [\tilde{G}_t + s_C\tilde{C}_t]\hat{Y}_t + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3) && (9.9) \\ &= u_C\bar{Y} \cdot \hat{Y}_t + \frac{1}{2}[u_C\bar{Y} + u_{CC}\bar{Y}^2] \cdot \hat{Y}^2 \\ &\quad - u_{CC}\hat{Y}^2 \cdot \hat{G}_t\hat{Y}_t + \text{t.i.p.} + \text{unf.} + \mathcal{O}(\|\xi\|^3). && (9.10) \end{aligned}$$

In (9.7), we simply expand in terms of the index of aggregate consumption C_t, where $\bar{C} \equiv \bar{Y} - \bar{G}$, and each of the partial derivatives is evaluated at the steady-state values $(\bar{C};0)$. Here the term $\mathcal{O}(\|\xi\|^3)$ indicates that we neglect terms that are of third or higher order in the deviations of the various variables from their steady-state values. In the case of a monetary policy rule that implies $\pi^* = 0$ and a tax rate $\tau = \tau^*$, the variables will deviate from these values in an equilibrium only because of fluctuations in the shocks G_t and ξ_t around their steady-state values. In this case, the omitted terms are all of third or higher order in the size of the exogenous shocks (and we use $\|\xi\|$ to indicate the a measure of the size of these shocks, where the size of fluctuations in G_t is intended to be included). More generally, the omitted terms also include terms that are of third or higher order in deviations of π^* from the value zero and of τ from the value τ^*; but we shall (for now) retain terms that are of first or second order in perturbations of those assumptions about long-run aspects of policy. In (9.8), we rewrite the expressions in terms of $\hat{Y}_t \equiv \log(Y_t/\bar{Y})$ and $\tilde{G}_t \equiv (G_t - \bar{G})/\bar{Y}$, using the Taylor expansion

$$Y_t = \bar{Y} \cdot [1 + \hat{Y}_t + \frac{1}{2}\hat{Y}_t^2] + \mathcal{O}(\|\xi\|^3).$$

In (9.9), we suppress the terms that are independent of policy (because they involve only constants and exogenous disturbances), denoted "t.i.p." as in the text, and make use of the definition $u_{C\xi} \cdot \xi_t = -u_{CC}\bar{C}\tilde{C}_t$ to obtain a scalar representation of the disturbance to the marginal utility of consumption. Finally, in (9.10), we recall the notation

$$\hat{G}_t \equiv \tilde{G}_t + s_C E_{t-2}\bar{C}_t = \tilde{G}_t + s_C\bar{C}_t + \text{unf.},$$

where "unf." stands for an unforecastable term (i.e., a term x_t with the property that $E_{t-2}x_t = 0$). Unforecastable terms may be neglected because we are ultimately interested only in the unconditional expectation of each of the terms in (9.10).

Similarly, a second-order Taylor series expansion of household z's disutility of working is given by

$$\begin{aligned} v &= v(\bar{Y};0) + v_y \cdot (y_t(z) - \bar{Y}) + +v_\xi \cdot \xi_t \\ &\quad + \frac{1}{2}v_{yy} \cdot (y_t(z) - \bar{Y})^2 + v_{y\xi} \cdot (y_t(z) - \bar{Y})\xi_t + \frac{1}{2}v_{\xi\xi}\xi_t^2 + \mathcal{O}(\|\xi\|^3) \\ &= v_y\bar{Y} \cdot \hat{y}_t(z) + \frac{1}{2}[v_y\bar{Y} + v_{yy}\bar{Y}^2] \cdot \hat{y}_t(z)^2 - v_{yy}\bar{Y}^2 \cdot \hat{y}_t(z)\bar{Y}_t + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3), \end{aligned} \tag{9.11}$$

where now $\hat{y}_t(z) \equiv \log(y_t(z)/\bar{Y})$, and $\bar{Y}_t$, defined by the relation $v_{y\xi}\xi_t = -v_{yy}\bar{Y}\bar{Y}_t$, provides a scalar measure of disturbances to the marginal disutility of supply. Integrating (9.11) over z we obtain

$$\begin{aligned} \int_0^1 v(y_t(z);\xi_t)dz &= v_y\bar{Y} \cdot E_z\hat{y}_t(z) + \frac{1}{2}[v_y\bar{Y} + v_{yy}\bar{Y}^2] \cdot [E_z\hat{y}_t(z)]^2 + \frac{1}{2}[v_y\bar{Y} + v_{yy}\bar{Y}^2] \cdot var_z\hat{y}_t(z) \\ &\quad - v_{yy}\bar{Y}^2 \cdot E_z\hat{y}_t(z)\bar{Y}_t + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3). \end{aligned} \tag{9.12}$$

Next we wish to express the terms in (9.12) involving the population average $E_z\hat{y}_t(z)$ in terms of the Dixit-Stiglitz output aggregate $\hat{Y}_t$ instead. To do so, we first compute a Taylor series expansion for the right-hand side of the aggregator equation (3.2), obtaining

$$\hat{Y}_t = E_z\hat{y}_t(z) + \frac{1}{2}\Big(\frac{\theta-1}{\theta}\Big)\text{var}_z\hat{y}_t(z) + \mathcal{O}(\|\xi\|^3).$$

Solving this equation for $E_z\hat{y}_t(z)$ and substituting into (9.12) yields

$$\begin{aligned} \int_0^1 v(y_t(z);\xi_t)dz &= v_y\bar{Y} \cdot \hat{Y}_t + \frac{1}{2}[v_y\bar{Y} + v_{yy}\bar{Y}^2] \cdot \hat{Y}_t^2 + \frac{1}{2}[\theta^{-1}v_y\bar{Y} + v_{yy}\bar{Y}^2] \cdot var_z\hat{y}_t(z) \\ &\quad - v_{yy}\bar{Y}^2 \cdot \hat{Y}_t\bar{Y}_t + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3). \end{aligned} \tag{9.13}$$

Substituting (9.10) and (9.13) into (9.6), we obtain

$$\begin{aligned} w_t &= u_C\bar{Y}\cdot[\hat{Y}_t + \frac{1}{2}(1-\sigma)\hat{Y}_t^2 + \sigma\hat{G}_t\hat{Y}_t] \\ &\quad -v_y\bar{Y}\cdot[\hat{Y}_t + \frac{1}{2}(1+\omega)\hat{Y}_t^2 + \frac{1}{2}(\theta^{-1}+\omega)\text{var}_z\hat{y}_t(z)] + \text{t.i.p.} + \text{unf.} + \mathcal{O}(\|\xi\|^3) \\ &= -\frac{1}{2}u_C\hat{Y}\cdot[(\sigma+\omega)\hat{Y}_t^2 - 2(\sigma+\omega)\hat{Y}_t^S\hat{Y}_t + (\theta^{-1}+\omega)\text{var}_z\hat{y}_t(z)] \\ &\quad +\text{t.i.p.} + \text{unf.} + \mathcal{O}(\|\xi\|^3). \end{aligned} \qquad (9.14)$$

Note that in deriving (9.14) from the line above we have (at last) used the assumption that $\bar{Y}$ is the efficient level of output, so that $u_C = v_y$, and the definition

$$\hat{Y}_t^S \equiv (\sigma+\omega)^{-1}[\sigma\hat{G}_t + \omega E_{t-1}\bar{Y}_t] = (\sigma+\omega)^{-1}[\sigma\hat{G}_t + \omega\bar{Y}_t] + \text{unf.}$$

Then taking the unconditional expectation of (9.14), we obtain

$$\begin{aligned} W &= -\frac{1}{2}u_C\hat{Y}\cdot[(\sigma+\omega)\text{var}\{\hat{Y}_t - \hat{Y}_t^S\} + (\sigma+\omega)[E\{\hat{Y}_t\}]^2 + (\theta^{-1}+\omega)E\{\text{var}_z\hat{y}_t(z)\}] \\ &\quad +\text{t.i.p.} + \mathcal{O}(\|\xi\|^3). \end{aligned} \qquad (9.15)$$

As promised, we have obtained a welfare measure that allows us to compute all second-order or lower terms in W using only a first-order (log-linear) approximation to the equilibrium solution for the pattern of activity $\{y_t(z)\}$, since no terms of order $\mathcal{O}(\|\xi\|^2)$ in the solution for $y_t(z)$ have any effect upon terms of order lower than $\mathcal{O}(\|\xi\|^3)$ in (9.15). [53] If we furthermore assume that the tax rate τ (or some other aspect of "long-run" policy) is adjusted so as to guarantee that $E\{\hat{Y}_t\} = 0$ (i.e., $\log Y_t$ equals $\log\bar{Y}$ on average) *regardless* of the monetary policy rule, then the $[E\{\hat{Y}_t\}]^2$ term in (9.15) can also be suppressed, as this term is also independent of the monetary policy rule. Our decision to assume this results from a belief, as discussed in the text, that monetary policy is not an appropriate instrument with which to seek to affect the long-run average level of economic activity, given the existence of other instruments with which policymakers may more directly seek to offset the distortion resulting from suppliers' market power. Finally, noting that $\hat{Y}_t$ equals $E_{t-2}\hat{Y}_t$ plus a forecast error term that is both unforecastable and independent of monetary policy,

[53] This result depends upon our having linearized around the efficient $\bar{Y}$, since otherwise our expression for W would contain a term that is linear in $E\{\hat{Y}_t\}$. However, even without this choice, we could have obtained the same result by assuming that tax policy adjusts in response to any change in the monetary policy rule in order to preserve a particular value for $E\{\hat{Y}_t\}$, which quantity then becomes one of the terms independent of the monetary policy rule. In fact, we make such an assumption in the work reported here in any event.

one can show that

$$\mathrm{var}\{\hat{Y}_t - \hat{Y}_t^S\} = \mathrm{var}\{E_{t-2}(\hat{Y}_t - \hat{Y}_t^S)\} + \text{t.i.p.}$$

Substitution of this into (9.15), along with the stipulation that $E\{\hat{Y}_t\} = 0$ regardless of monetary policy, then yields (7.2).

Some might prefer instead an analysis that would assume a tax rate τ that remained invariant under alternative monetary policy rules. In this case, we would not be able to drop the $[E\{\hat{Y}_t\}]^2$ term in (9.15). However, our model implies that

$$E\{\hat{Y}_t\} = \frac{1-\beta}{\kappa} E\{\hat{X}_t\} + \mathcal{O}(\|\xi\|^2) = \frac{1-\beta}{\kappa} E\{\pi_t\} + \mathcal{O}(\|\xi\|^2),$$

where the first equality follows from taking the unconditional expectation of all terms in (3.26), and the second from taking the unconditional expectation of all terms in (3.25). (Note that in the log-linear approximations to the model equations reported in the paper, we routinely suppress terms of order $\mathcal{O}(\|\xi\|^2)$.) Thus the only difference in the alternative case would be the presence of an additional negative term in π^{*2} in (7.5). This would have no effect upon the definition of the loss from incomplete stabilization L in (7.6), but it would mean that in (7.5) we would have $L + \mu\pi^{*2}$ instead of $L + \pi^{*2}$, for a certain $\mu > 1$, as our overall deadweight loss measure. This would imply that the optimal point on the $L - \pi^*$ frontier (discussed and graphed in Rotemberg and Woodford, 1997) would be slightly different from the one that we assume here, involving a slightly smaller, though still slightly positive, value of π^*. Such a change makes no qualitative difference, however, for our conclusions regarding the nature of optimal policy.

Next we turn to the derivation of (7.5) from (7.2). As noted in the text, we need to show that the dispersion of levels of production across differentiated goods is a function of the degree of variability of the aggregate price level. We begin by noting that output dispersion follows from price dispersion, since (3.15) implies (7.3), and hence that

$$E\mathrm{var}_z\{\log y_t(z)\} = \theta^2 E\mathrm{var}_z\{\log p_t(z)\}. \tag{9.16}$$

To relate the cross-sectional variance of prices to the variability over time of the price index P_t, we recall that in any period t, a fraction α of suppliers charge the same price as at $t-1$ (and the distribution of their prices is the same as the distribution of period $t-1$ prices); a fraction $(1-\alpha)\gamma$ charge a common new price p_t^1 chosen at $t-1$, and a fraction $(1-\alpha)(1-\gamma)$ charge a common new price p_t^2 chosen at $t-2$. Then, introducing the notation, $\bar{p}_t \equiv E_z \log p_t(z)$, we obtain

$$\begin{aligned}
\text{var}_z\{\log p_t(z)\} &= \text{var}_z\{\log p_t(z) - \bar{p}_{t-1}\} = E_z\{[\log p_t(z) - \bar{p}_{t-1}]^2\} + (\Delta\bar{p}_t)^2 \\
&= \alpha E_z\{[\log p_{t-1}(z) - \bar{p}_{t-1}]^2\} + (1-\alpha)\gamma[\log p_t^1 - \bar{p}_{t-1}]^2 \\
&\quad +(1-\alpha)(1-\gamma)[\log p_t^2 - \bar{p}_{t-1}]^2 + (\Delta\bar{p}_t)^2 \\
&= \alpha \text{var}_z\{\log p_{t-1}(z)\} + (1-\alpha)\gamma[\log p_t^1 - \bar{p}_{t-1}]^2 \\
&\quad +(1-\alpha)(1-\gamma)[\log p_t^2 - \bar{p}_{t-1}]^2 + (\Delta\bar{p}_t)^2. \qquad (9.17)
\end{aligned}$$

Taking the unconditional expectation of both sides of (9.17) then yields

$$\begin{aligned}
E[\text{var}_z\{\log p_t(z)\}] &= \gamma E[(\log p_t^1 - \bar{p}_{t-1})^2] + (1-\gamma)E[(\log p_t^2 - \bar{p}_{t-1})^2] \\
&\quad +(1-\alpha)^{-1}E[(\Delta\bar{p}_t)^2]. \qquad (9.18)
\end{aligned}$$

Similar reasoning as is used in deriving (9.17) also yields

$$\begin{aligned}
\bar{p}_t - \bar{p}_{t-1} &= E_z\{\log p_t(z) - \bar{p}_{t-1}\} \\
&= \alpha E_z\{\log p_{t-1}(z) - \bar{p}_{t-1}\} + (1-\alpha)\gamma[\log p_t^1 - \bar{p}_{t-1}] \\
&\quad +(1-\alpha)(1-\gamma)[\log p_t^2 - \bar{p}_{t-1}] \\
&= (1-\alpha)\gamma[\log p_t^1 - \bar{p}_{t-1}] + (1-\alpha)(1-\gamma)[\log p_t^2 - \bar{p}_{t-1}]. \qquad (9.19)
\end{aligned}$$

Taking the expectation of (9.19) conditional upon date $t-2$ information, one obtains

$$E_{t-2}(\bar{p}_t - \bar{p}_{t-1}) = (1-\alpha)[\log p_t^2 - \bar{p}_{t-1}] + \mathcal{O}(\|\xi\|^2), \qquad (9.20)$$

using the fact that $\bar{p}_t^2 = E_{t-2}\bar{p}_t^1 + \mathcal{O}(\|\xi\|^2)$ and that all date $t-1$ prices are known at $t-2$. This combined with (9.19) implies that

$$(\bar{p}_t - \bar{p}_{t-1}) - (1-\gamma)E_{t-2}(\bar{p}_t - \bar{p}_{t-1}) = (1-\alpha)\gamma[\log p_t^1 - \bar{p}_{t-1}] + \mathcal{O}(\|\xi\|^2). \qquad (9.21)$$

Furthermore, given that we are expanding around a steady state with zero inflation, the right hand sides of both (9.20) and (9.21) consist solely of terms of order $\mathcal{O}(\|\xi\|)$. Thus by squaring (9.20) and taking the conditional expectation, we obtain

$$\begin{aligned}
E[(\log p_t^2 - \bar{p}_{t-1})^2] &= (1-\alpha)^{-2}E[(E_{t-2}\Delta\bar{p}_t)^2] + \mathcal{O}(\|\xi\|^3) \\
&= (1-\alpha)^{-2}\text{var}\{E_{t-2}\Delta\bar{p}_t\} + (1-\alpha)^{-2}[E\Delta\bar{p}_t]^2 + \mathcal{O}(\|\xi\|^3).
\end{aligned}$$

A similar expression for $E[(\log p_t^1 - \bar{p}_{t-1})^2]$ is implied by (9.21). Substituting these expressions into (9.18) then yields

$$E[\text{var}_z\{\log p_t(z)\}] = \frac{\alpha}{(1-\alpha)^2}\{\text{var}\{E_{t-2}\Delta\bar{p}_t\} + [E\Delta\bar{p}_t]^2\} + \frac{1-\gamma(1-\alpha)}{\gamma(1-\alpha)^2}\text{var}\{\Delta\bar{p}_t - E_{t-2}\Delta\bar{p}_t\} + \mathcal{O}(\|\xi\|^3). \quad (9.22)$$

Finally, the definition of the price index (3.3) implies that

$$\bar{p}_t = \log P_t + \mathcal{O}(\|\xi\|^2).$$

Making this substitution in (9.22), we obtain (7.4), neglecting a remainder of order $\mathcal{O}(\|\xi\|^3)$. Substitution of (7.4) and (9.16) into (7.2), and using the fact that $E\pi_t = \pi^* + \mathcal{O}(\|\xi\|^2)$(as a consequence of our definition of π^*), then yields the desired expression for W. This expression can be written in the form (7.5), where L is defined by (7.6), and

$$\Omega \equiv -\frac{1}{2}u_C\bar{Y}\theta(1+\theta\omega)\frac{\alpha}{(1-\alpha)^2}, \qquad \Lambda \equiv \frac{(1-\alpha)\kappa}{(1-\alpha\beta)\theta}.$$

Finally, we can rewrite L so that it depends only on the stochastic process for the relative price variable $\hat{X}$. To do this, note first that (3.26) implies that

$$E_{t-2}(\hat{Y}_t - \hat{Y}_t^S) = (1/\kappa)E_{t-2}(\hat{X}_t - \beta\hat{X}_{t+1}). \quad (9.23)$$

At the same time, (3.25) implies that

$$\pi_t - E_{t-2}\pi_t = \psi(\hat{X}_t - E_{t-2}\hat{X}_t), \quad (9.24)$$

so that

$$\text{var}(\pi_t) = \text{var}(E_{t-2}\hat{X}_t) + \psi^2\text{var}(\hat{X}_t - E_{t-2}\hat{X}_t). \quad (9.25)$$

Substituting the expressions in (9.23), (9.24) and (9.25) into the term in square brackets in (7.6), we obtain (7.7).

10 References

Aiyagari, S. R. and R. A. Braun. (1996). Some Explorations into Optimal Cyclical Monetary Policy. Minneapolis: Federal Reserve Bank of Minneapolis Working Paper no. 565.

Bernanke, B. S. and A. S. Blinder. (1992). The Federal Funds Rate and the Channels of Monetary Policy *American Economic Review*, 82(4):901-921.

Bernanke B. S. and I. Mihov. (1995). Measuring Monetary Policy, Cambridge, MA: NBER Working Paper 5145.

Bernanke, B. S. and M. Woodford. (1997) Inflation Forecasts and Monetary Policy *Journal of Money, Credit and Banking*, forthcoming.

Blinder, A. S. (1994) On Sticky prices: Academic Theories Meet the Real World. In N. Gregory Mankiw ed. *Monetary Policy*, Chicago: University of Chicago Press.

Bryant, R. C., P. Hooper and C. L. Mann. (1993). *Evaluating Policy Regimes*, Washington: Brookings Institution.

Calvo, G. (1983). Staggered Prices in a Utility-Maximizing Framework *Journal of Monetary Economics*, 12(3):383-98.

Campbell, J. Y. and N. G. Mankiw. (1989). Consumption, Income and Interest Rates: Reinterpreting the Time-Series Evidence *NBER Macroeconomics Annual*, 185-216.

Card, D. (1994). Intertemporal Labor Supply: An Assessment. In C. Sims ed. *Advances in Econometrics*, New York: Cambridge University Press.

Carlton, D. W. (1986). The Rigidity of Prices *American Economic Review* 76(4):637-658.

Cecchetti, S. G. (1986). The Frequency of Price Adjustment: A Study of the Newsstand Prices of Magazines *Journal of Econometrics*, 31(3):255-274.

Chari, V.V., P. J. Kehoe and E. R. McGrattan. (1996). Sticky Price Models of the Business cycle: Can the Contract Multiplier Solve the Persistence Problem, Cambridge MA: NBER Working Paper 5809.

Christiano, L., M. Eichenbaum and C. Evans. (1994). Identification and the Effects of a Monetary Policy Shock. In M. Blejer et al., *Financial Factors in Economic Stabilization and Growth*, Cambridge: Cambridge University Press, 1994.

—, — and — . (1996). The Effects of Monetary Policy Shocks: Evidence from the Flow of Funds *Review of Economics and Statistics* 78(1):16-34.

Cochrane, J. (1994). Shocks *Carnegie-Rochester Conference Series on Public Policy*, 41:295-364.

Dixit, A. and J. Stiglitz (1977). Monopolistic Competition and Optimum Product Diversity *American Economic Review*, 67(3):297-308.

Dutta, S., M. Bergen and D. Levy. (1997). The Flexibility of Retail Prices and the Rigidity of Wholesale Prices in Response to Exogenous Cost Changes: Evidence from Store-Level Data *Quarterly Journal of Economics*, forthcoming.

Feldstein, M. and J. H. Stock. (1994). The Use of a Monetary Aggregate to Target Nominal GDP. In N. G. Mankiw ed. *Monetary Policy*, Chicago; University of Chicago Press.

Fuhrer J. C. (1997a) Inflation/Output variance Trade-Offs and Optimal Monetary Policy *Journal of Money, Credit and Banking* 29(2):214-234.

— (1997b). Towards a Compact, Empirically-Verified Rational Expectations Model for Monetary Policy Analysis *Carnegie-Rochester Conference Series on Public Policy*, forthcoming.

Fuhrer, J. C. and G. Moore. (1995a) Inflation Persistence *Quarterly Journal of Economics*, 110(1):127-160.

— and — . (1995b). Monetary Policy Trade-offs and the Correlation between Nominal Interest Rates and Real Output *American Economic Review*, 85(1):219-239.

Goodfriend, M. (1991). Interest Rates and the Conduct of Monetary Policy *Carnegie-Rochester Conference Series on Public Policy*, 34:7-30.

Henderson, D. W. and W. J. McKibbin. (1993). A Comparison of Some Basic Monetary Policy Regimes for Open Economies *Carnegie-Rochester Conference Series on Public Policy*, 39:221-317.

Ireland, P. (1997). A Small, Structural, Quarterly Model for Monetary Policy Evaluation *Carnegie-Rochester Conference Series on Public Policy*, forthcoming.

Kashyap, A. K. (1995). Sticky Prices: New Evidence from Retail Catalogues *Quarterly Journal of Economics*, 110(1):245-74.

King, R. G. and M. W. Watson. (1995). System Reduction and Solution Algorithms for Singular Linear Difference Systems under Rational Expectations, unpublished, University of Virginia, May 1995.

— and —. (1996). Money, Prices, Interest Rates and the Business Cycle *Review of Economics and Statistics* 78(1):35-53.

King, R. G. and A. Wolman. (1996). Inflation Targeting in a St. Louis Model of the Twenty-First Century, Cambridge, MA: NBER Working Paper 5507.

Leeper, E. M. and C. A. Sims, (1994). Towards a Modern Macroeconomic Model Usable for Policy Analysis *NBER Macroeconomics Annual*, 81-117.

Leeper, E. M., C. A. Sims, and T. Zha. (1996) What Does Monetary Policy Do? *Brookings Papers on*

Economic Activity, 1:1-63.

Levin, A. (1996) A Comparison of Alternative Monetary Policy Rules in the FRB Multi-Country Model. Mimeo. Washington: Federal Reserve Board of Governors.

Lucas R. E. Jr. (1976). Econometric Policy Evaluation: A critique *Carnegie-Rochester Conference Series on Public Policy*, 1:19-46.

—. (1993). On the Welfare Cost of Inflation, Mimeo, University of Chicago.

McCallum, B. T. (1988). Robustness Properties of a Rule for Monetary Policy *Carnegie-Rochester Series on Public Policy* 29:173-204.

——. (1997) Issues in the Design of Monetary Policy Rules, Mimeo, Carnegie-Mellon University.

—— and E. Nelson. (1997). An Optimizing IS-LM Specification for Monetary Policy and Business Cycle Analysis, Cambridge, MA: NBER Working Paper 5875.

Parkin, M. (1988). A Method for Determining whether Parameters in Aggregative Models are Structural *Carnegie-Rochester Conference Series on Public Policy* 29:215-252.

Pencavel, J. H. (1986). Labor Supply of Men: A Survey. In O. Ashenfelter and R. Layard eds. *Handbook of Labor Economics*, Amsterdam: North Holland.

Roberts, J. M. (1995). New Keynesian Economics and the Phillips Curve *Journal of Money, Credit and Banking*, 27(4):975-984.

Rotemberg, J. J. (1982). Sticky Prices in the United States *Journal of Political Economy*, 90(6):1187-1211.

—. (1983). Supply Shocks, Sticky Prices and Monetary Policy *Journal of Money, Credit and Banking*, 15(4):433-435.

—. (1996) Prices, Output and Hours: An Empirical Analysis Based on a Sticky Price Model *Journal of Monetary Economics* 37:505-533.

Rotemberg, J. J. and M. Woodford. (1996). Imperfect Competition and the Effects of Energy Price Increases on Economic Activity *Journal of Money, Credit and Banking*, 28(4):549-577.

—— and ——. (1997a) The Cyclical Behavior of Prices and Costs, unpublished, Princeton University, February.

—— and ——. (1997b) Interest-Rate Rules in an Estimated Sticky-Price Model, unpublished, Princeton

University, December.

Sims, C. A. (1986). Are Forecasting Models Usable for Policy Analysis? *Federal Reserve Bank of Minneapolis Quarterly Review* 2-16.

Solon, G., R. Barsky and J. A. Parker. (1994). Measuring the Cyclicality of Real Wages: How Important is Composition Bias? *Quarterly Journal of Economics*, 109(1):1-26.

Strongin, S. (1995). The Identification of Monetary Policy Disturbances: Explaining the Liquidity Puzzle *Journal of Monetary Economics* 35(3):463-498.

Summers, L. (1991). How Should Long Term Monetary Policy Be Determined? *Journal of Money, Credit and Banking*, 23(3):625-631.

Taylor, J. B. (1980). Aggregate Dynamics and Staggered Contracts *Journal of Political Economy* 88(1):1-23.

—. (1993a). *Macroeconomic Policy in a World Economy*, New York: W.W. Norton.

—. (1993b). Discretion Versus Policy Rules in Practice *Carnegie-Rochester Conference Series on Public Policy* 39:195-214.

Woodford, M. (1996). Control of the Public Debt: A Requirement for Price Stability? Cambridge, MA: NBER Working Paper 5684.

Yun, T. (1996). Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles *Journal of Monetary Economics*, 37(2):345-370.

Table 1
The Vector Autoregression
Sample: 1980:1-1995:2

Independent var:	R_t	π_{t+1}	Y_{t+1}
π_{t+1}			0.111 0.109
R_t		-0.085 0.133	-0.030 0.104
R_{t-1}	0.507 0.129	-0.078 0.141	-0.316 0.110
R_{t-2}	0.023 0.145	0.194 0.139	0.382 0.110
R_{t-3}	0.158 0.109	-0.096 0.106	-0.086 0.083
π_t	-0.062 0.130	0.540 0.125	-0.099 0.113
π_{t-1}	0.427 0.152	0.077 0.156	0.174 0.121
π_{t-2}	0.300 0.157	0.379 0.156	-0.135 0.128
Y_t	0.624 0.175	0.260 0.187	1.335 0.148
Y_{t-1}	-0.059 0.306	-0.297 0.293	-0.163 0.230
Y_{t-2}	-0.418 0.189	0.143 0.189	-0.258 0.148
R^2	0.947	0.834	0.933
D.W.	2.07	2.13	1.83

Standard Errors below estimates

Table 2
Variances of Output, Inflation and
Interest Rates
Under Different Monetary Rules

	Var(R)	Var(Y)	Var(π)	Var(π-Eπ)	Var{E($Y - Y^s$)}	Loss from Variability (L)
Historical Policy with shocks	7.64	4.79	2.28	.66	12.14	3.43
Historical Policy without shocks	6.73	4.55	2.25	.65	11.89	3.39
$\theta_\pi = 1.5$ $\theta_y = .5$	17.14	3.87	7.34	.81	13.86	8.72
$\theta_\pi = 1$ $\theta_y = 5$	22.95	.51	6.45	.91	17.84	8.10
$\theta_\pi = 10$ $\theta_y = 0$	30.11	12.61	.30	.25	4.58	.74
Minimum L Policy	732.9	18.77	0	0	0	0
Constrained Optimal Policy	1.93	11.30	.39	.20	7.57	.93

Note: For these computations, the interest rate and inflation are measured in annualized percentage points while output is measured in percentage deviations from trend.

Figure 1
Estimated and Theoretical Responses to a
Monetary Policy Shock

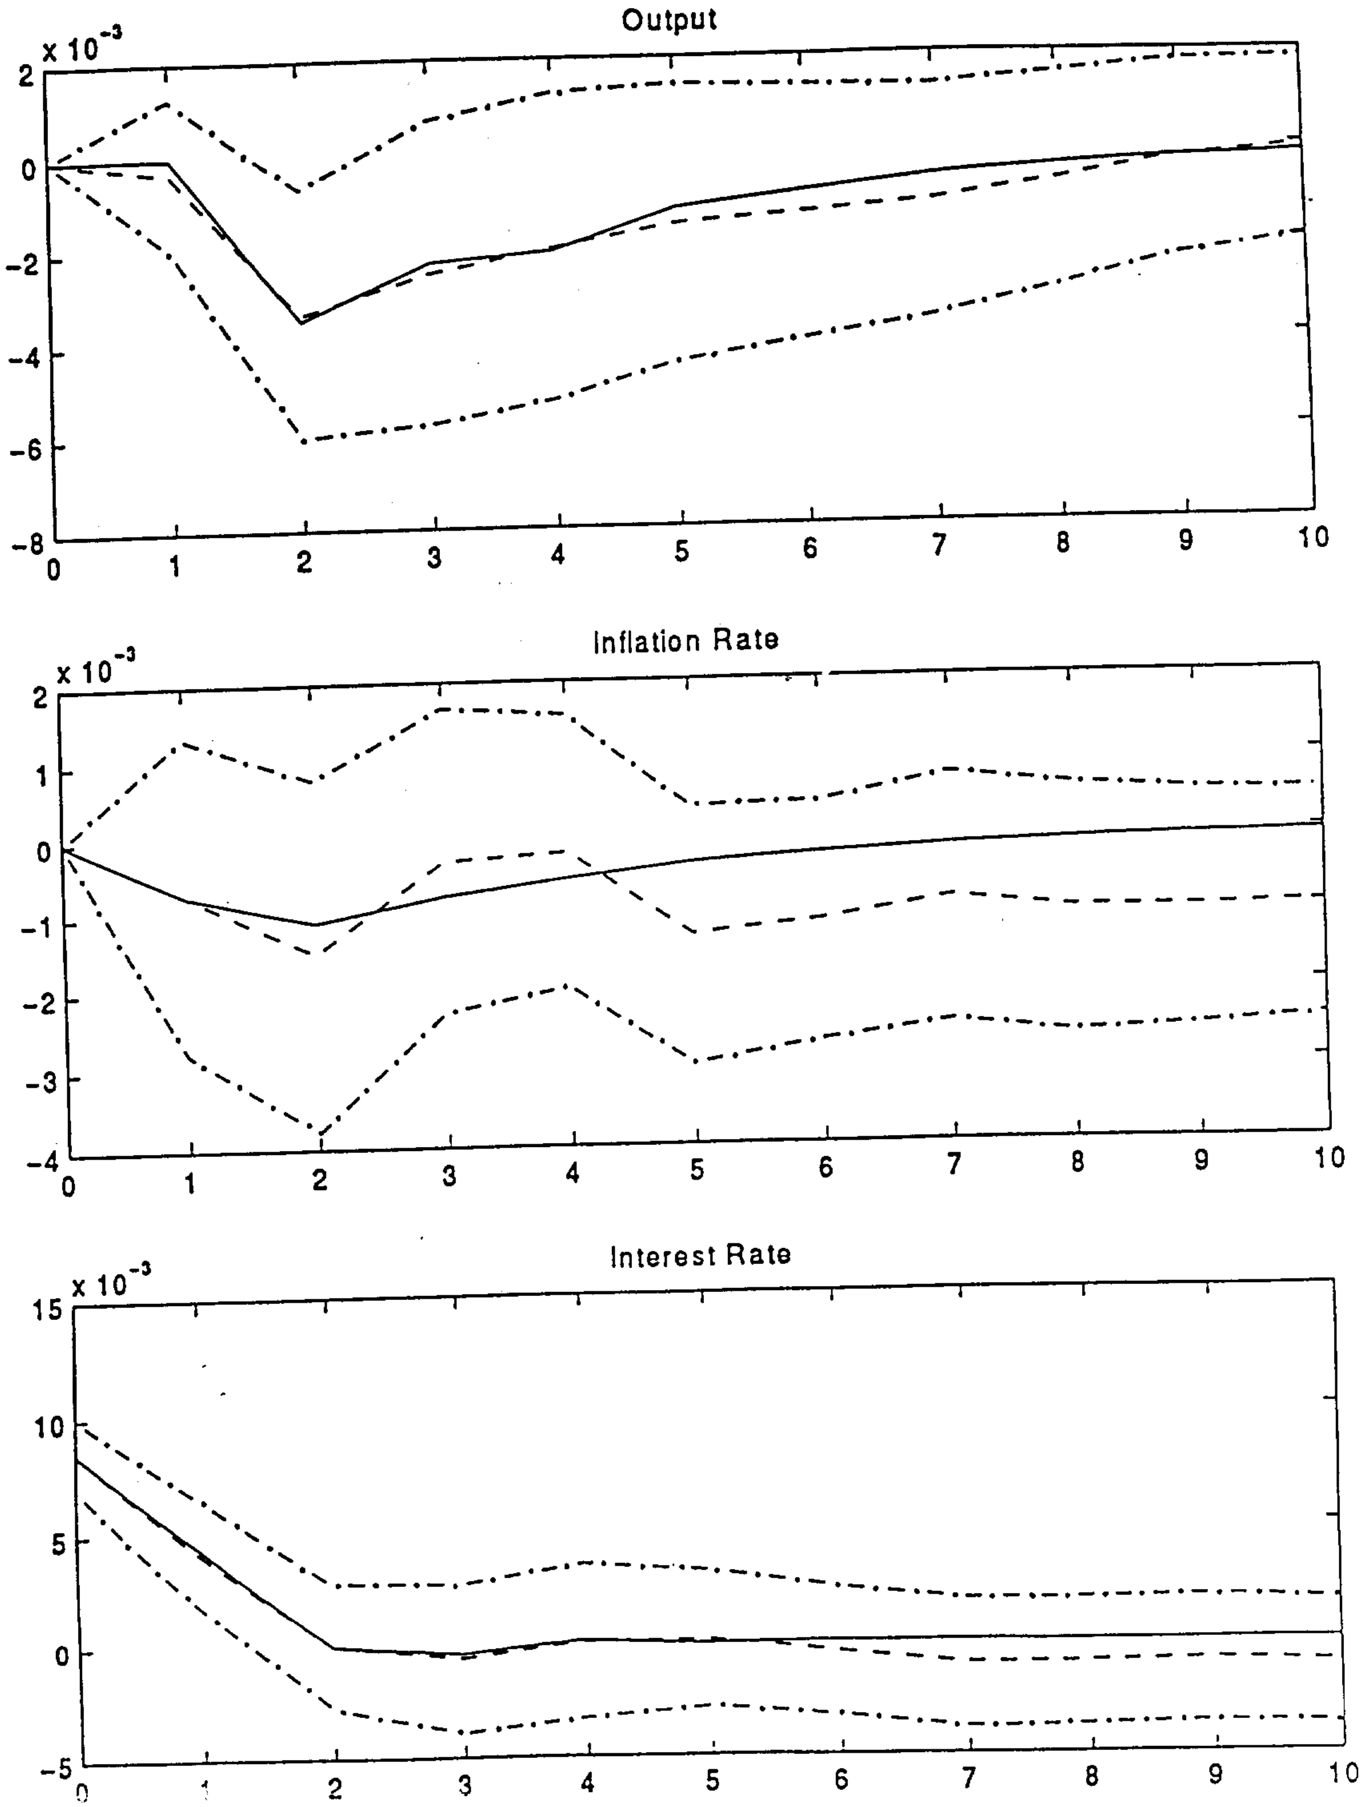

Figure 2
Empirical and Theoretical Auto and Cross-Correlations

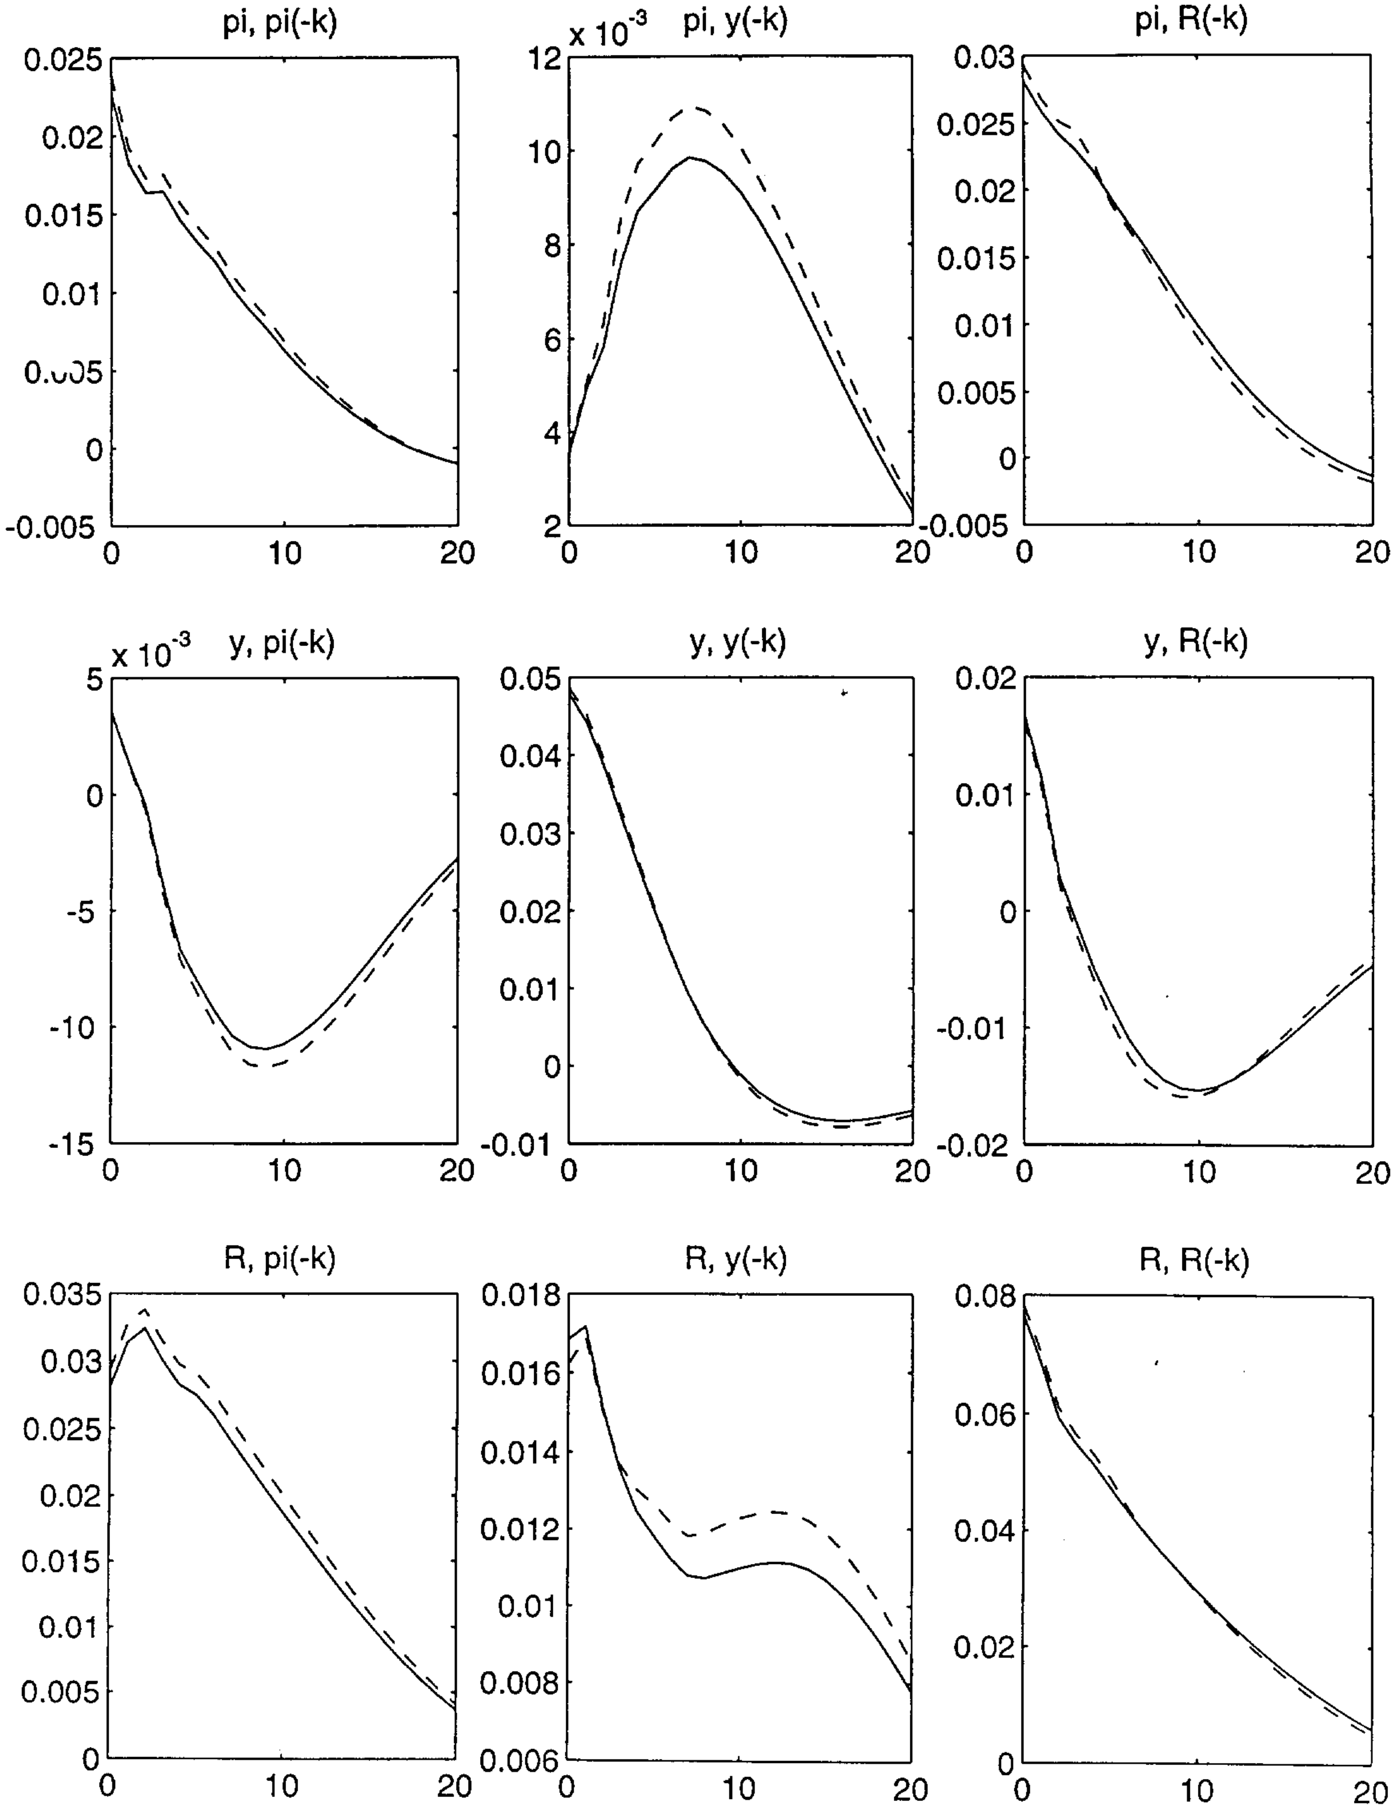

Figure 3
Actual and Simulated Paths

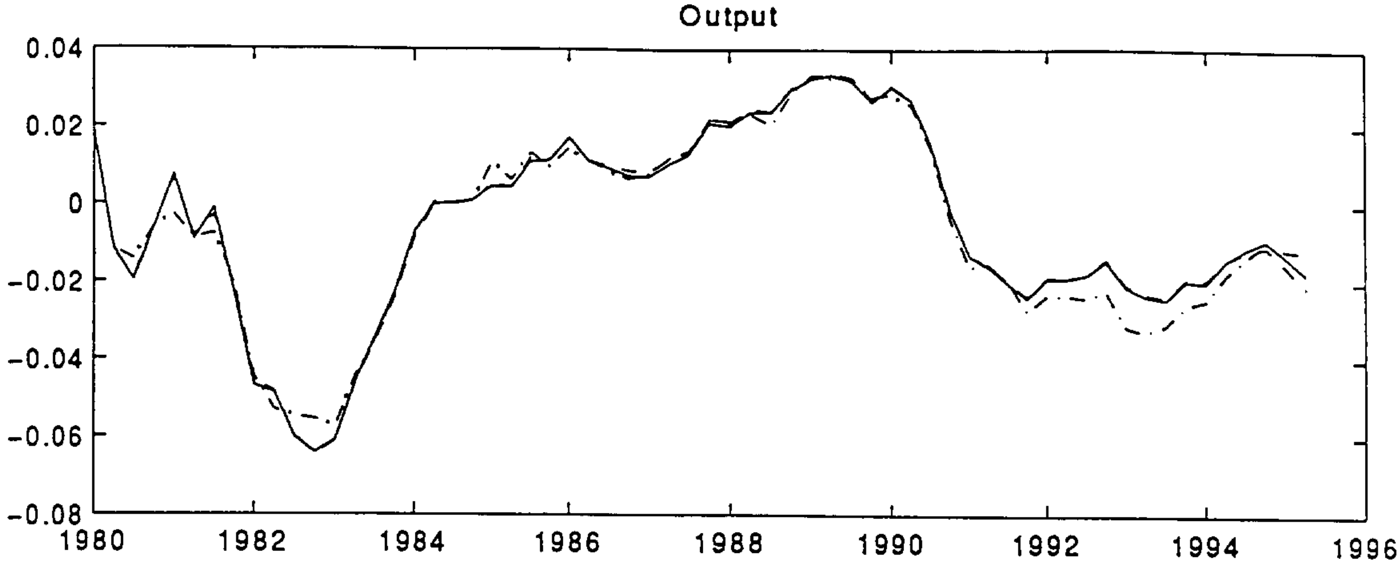

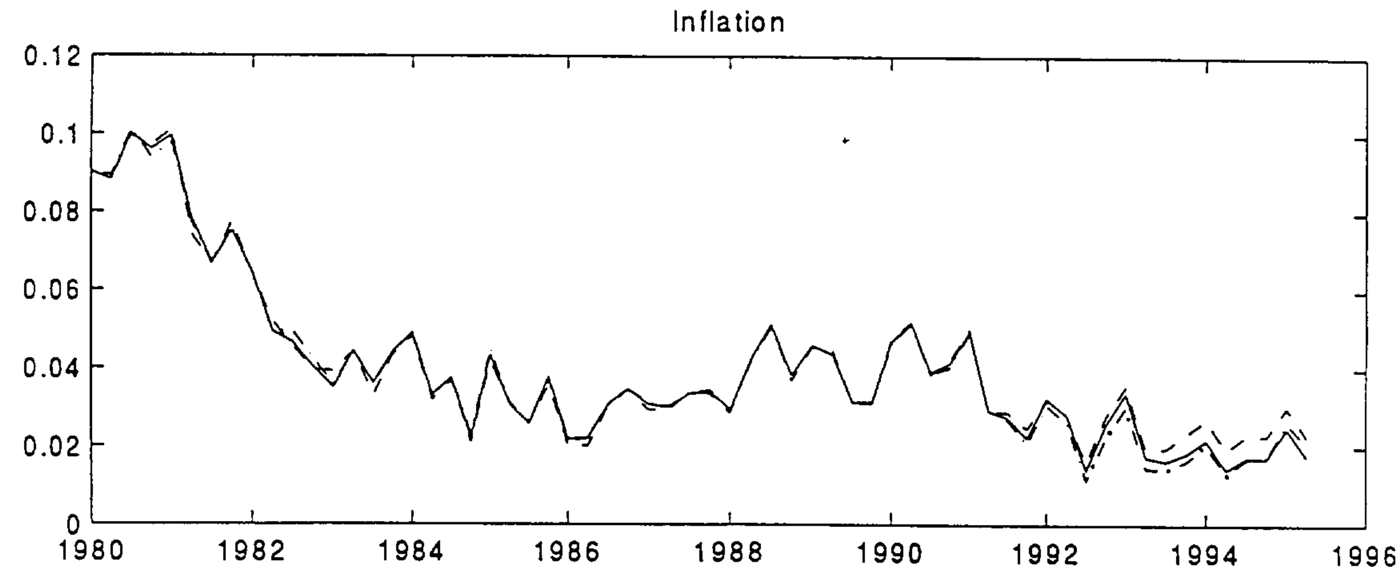

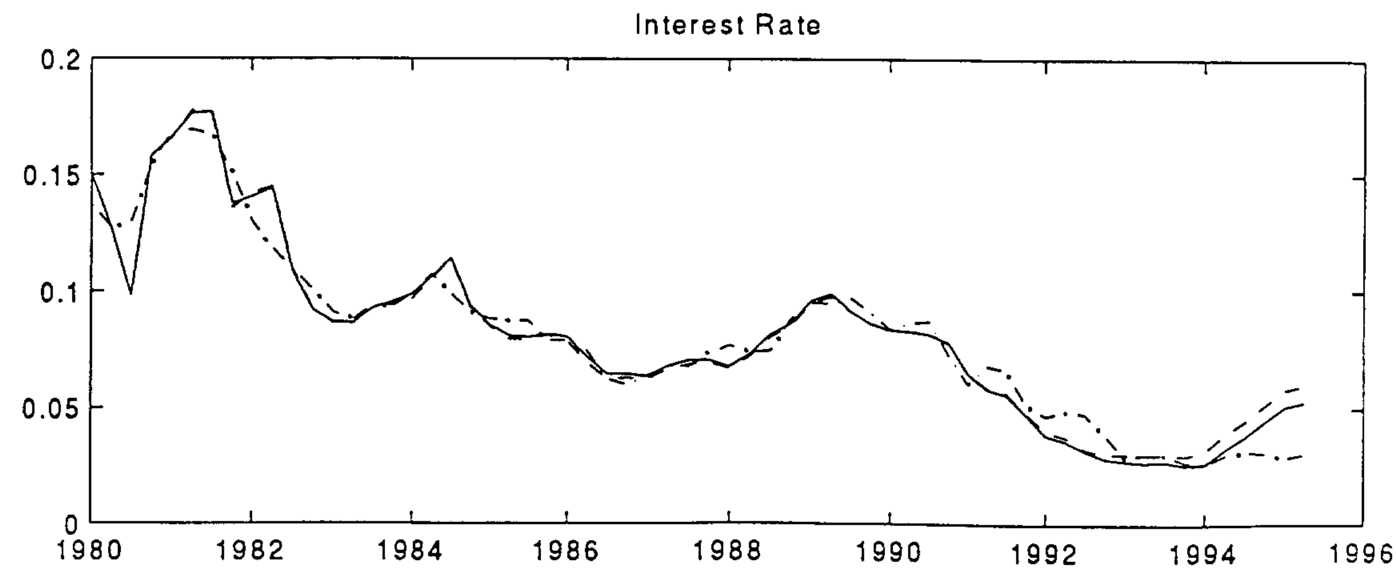

Figure 4

Simulated "Taylor Rule" ($\theta_\pi = 1$, $\theta_y = 5$)

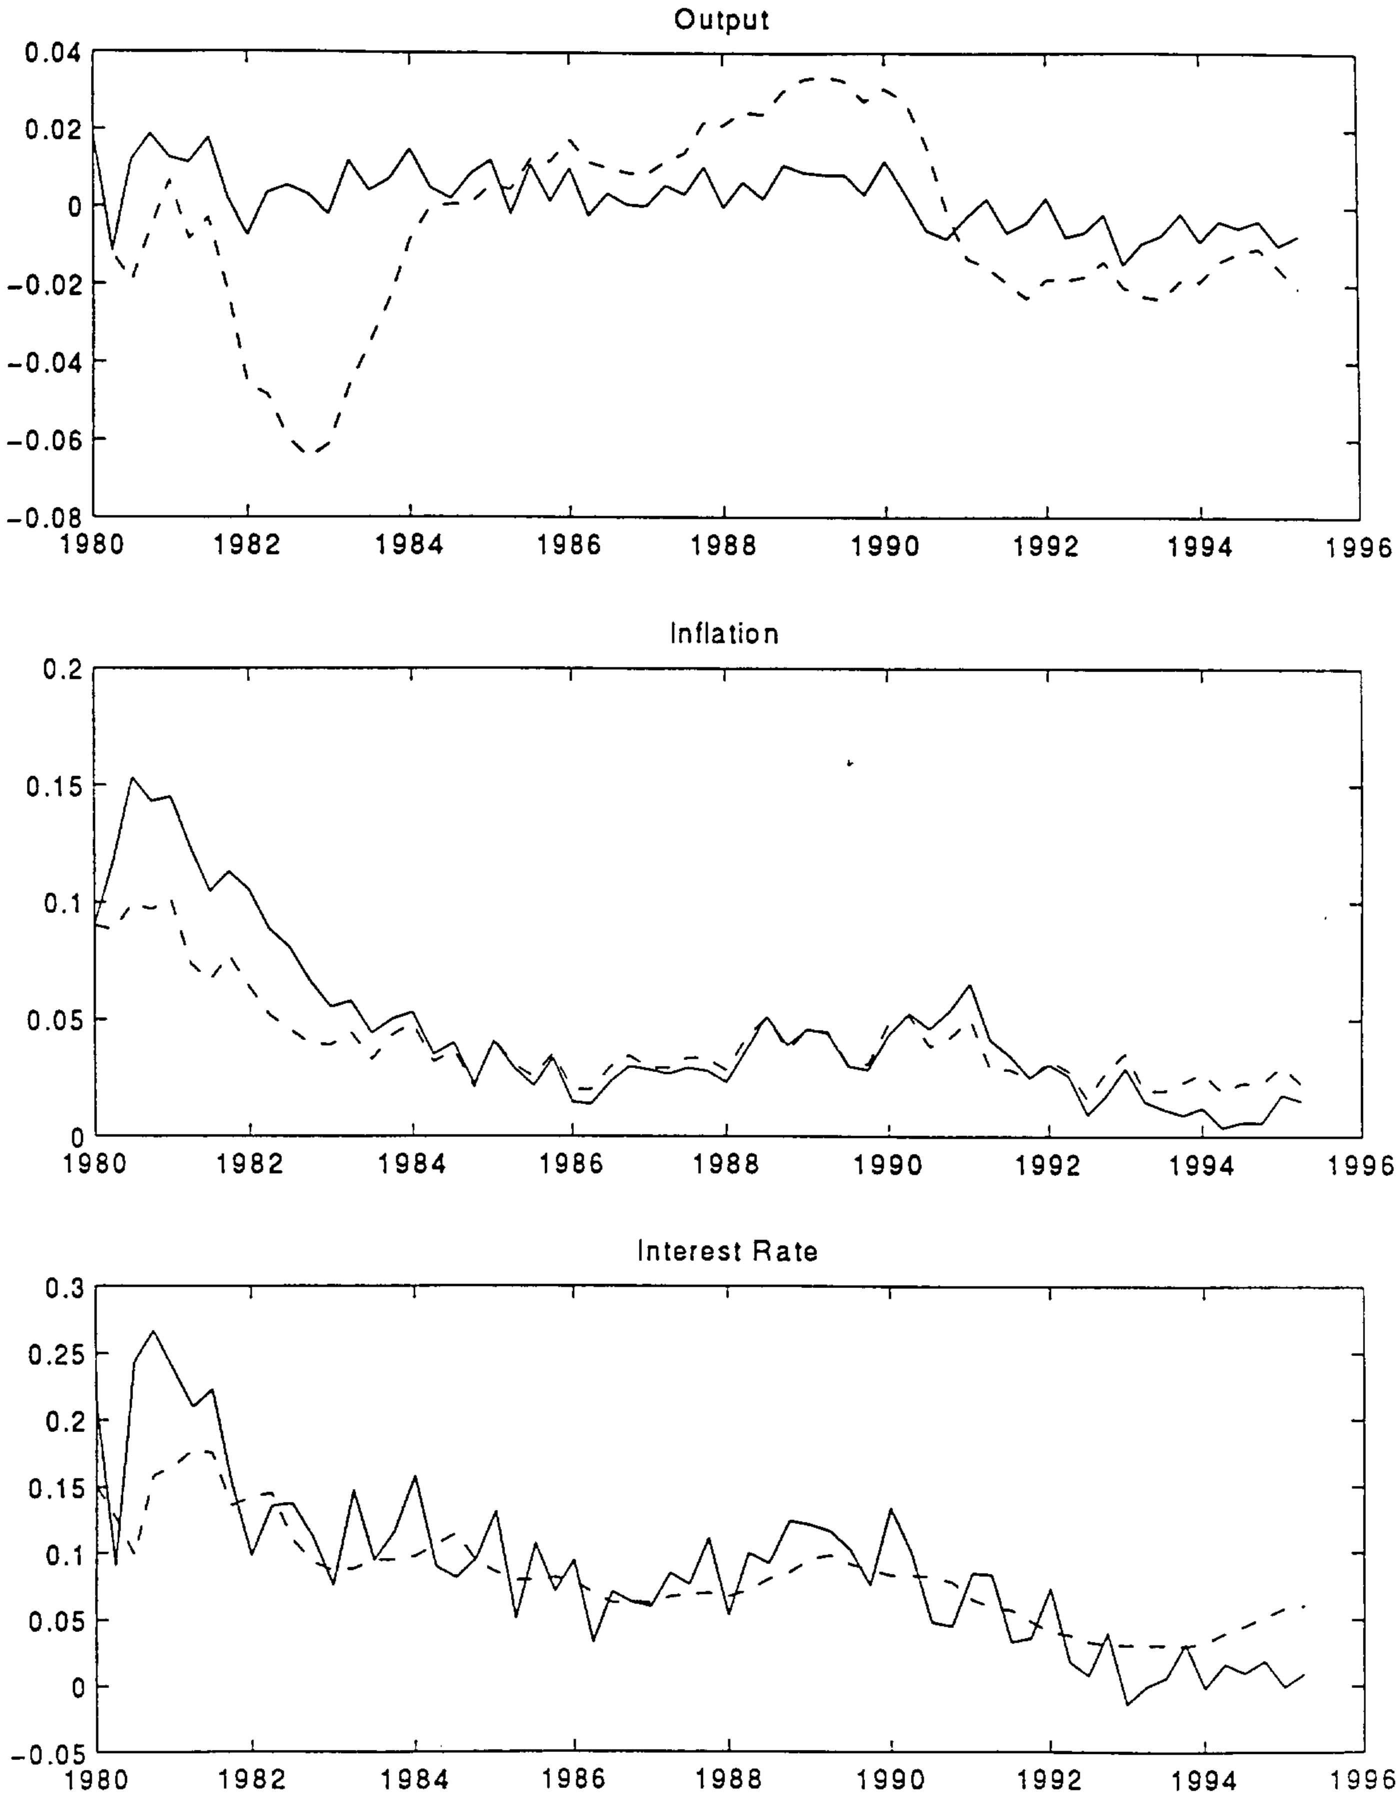

Figure 5
The Tradeoff Between Interest-Rate Volatility and Welfare Loss L

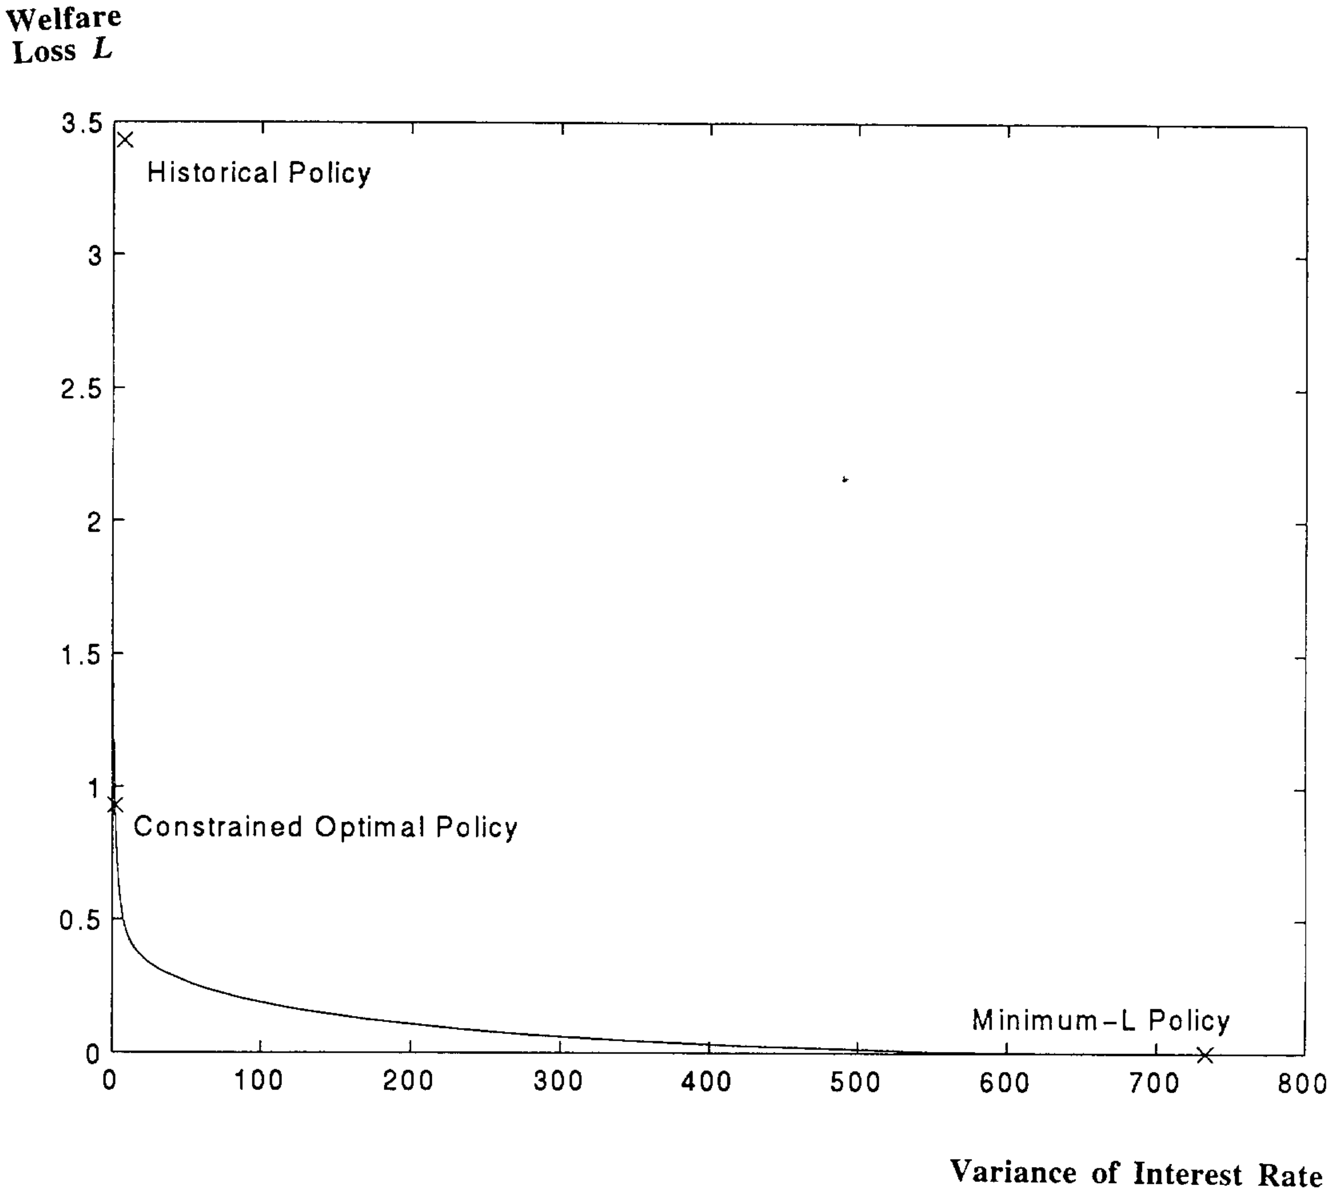

Figure 6
Impulse Responses Under Actual and Constrained-Optimal Policy