This is to certify that the
dissertation entitled

Three Papers in Bayesian Empirical Macroeconomics

presented by

Paul Richard Corrigan

has been accepted towards fulfillment
of the requirements for the

Doctoral degree in Economics

Major Professor's Signature

1 July 2004

Date

THREE PAPERS IN BAYESIAN EMPIRICAL MACROECONOMICS

By

Paul Richard Corrigan

A DISSERTATION

Submitted to
Michigan State University
In partial fulfillment of the requirements
For the degree of

DOCTOR OF PHILOSOPHY

Department of Economics

2004

ABSTRACT

THREE PAPERS IN BAYESIAN EMPIRICAL MACROECONOMICS

By

Paul Richard Corrigan

These three related papers use a variety of Bayesian methods to examine the specification of "New Keynesian" dynamic stochastic general equilibrium (DSGE) models and their uses in forecasting and business cycle analysis.

In the first paper, "Forecasting output and inflation with a Bayesian VAR using a New Keynesian prior," I evaluate the forecasting performance of a simple New Keynesian DSGE model. I do this by using the model to calculate a prior mean and covariance matrix for the coefficients of a VAR forecasting model for output, inflation and interest rates, using Ingram and Whiteman's (1994) technique for using DSGE models to calculate priors for Bayesian vector autoregression (BVAR) forecasting models. The resulting BVAR generates forecasts of inflation competitive with those from a BVAR with a atheoretical prior. However, the New Keynesian BVAR results in very poor output forecasts, particularly in the short run, suggesting some source of misspecification in the New Keynesian DSGE model.

In the second paper, "Loss-based evaluation of a New Keynesian DSGE model," I evaluate and compare the in-sample specification of a simple New Keynesian DSGE model to a cash-credit model with flexible prices and to an identified Bayesian VAR, in a manner similar to Schorfheide's (2000) study of the portfolio adjustment cost model. I calculate Bayes factors for each model so as to calculate posterior probabilities for all three models, and construct a benchmark distribution for correlations and impulse

response functions to which I can compare the correlations and impulse responses of the New Keynesian model, according to a variety of loss functions. I find that with a realistic monetary rule, both the flexible and New Keynesian DSGEs are competitive on a posterior odds basis with each other, as well as with the identified VAR. Very different levels of price stickiness, with very different policy implications, are compatible with the same data, making use of outside prior information crucial in assessing the role of the Phillips curve and of "supply-side" and "demand-side" shocks in business cycles.

This point is further underlined in the third paper, "Technology shocks versus monetary shocks: Identifying sources of business cycle fluctuations with a New Keynesian DSGE Model." Following recent papers such as that of Smets and Wouters (2002), I estimate by Bayesian methods similar to that in the previous paper a three-variable New Keynesian DSGE model of output, inflation and interest rates, including habit formation in preferences to improve the model's dynamics as well as stochastic price rigidity. I calculate Bayes factors for the model along with several Bayesian VAR models, and find it competitive with a four-lag BVAR. I use the DSGE model to estimate series for monetary shocks, technology/supply shocks, and autonomous demand shocks for the United States since 1965. I find that monetary policy shocks have probably been of only secondary importance, as compared to that of supply shocks, in movements in US inflation, as well as output, business cycles since 1965. My results conflict with those of Smets and Wouters for Europe, who found a larger role for monetary shocks; I argue that our differing prior assumptions regarding price stickiness are more likely to account for our different results than institutional differences between the US and Europe.

To the memory of Kathleen Flood
1930-2002
for whom I will always be thankful
and look forward to one day seeing again
RIP

ACKNOWLEDGEMENTS

Perhaps you have seen a certain children's book called *I Am a Pencil,* which is designed to teach readers about the marvels of the free market, describing the many people, tasks, and materials that go into making a lead pencil, the wonder of it all being the fact that the pencils get made but, says the pencil, "Nobody knows how to make me." I doubt anyone knows how to make an economist either; economists are more complex than pencils. Certainly I don't know how to make one, least of all single-handedly, and I couldn't have done this alone. However, an exhaustive list of all the people who helped me finish this dissertation and all their contributions would be as long as the dissertation. So I'll save that for my memoirs and just mention a few of the more important ones.

On top of the list surely must come my committee chair, Dr. Rowena Pecchenino. Many acknowledgements praise committee chairs first; after all, it is the committee chair that one has to talk into signing the dissertation. However, I truly am convinced I'd never have completed this dissertation without Dr. Pecchenino; she not only forced me to think hard about macroeconomics and my research, but kept faith in me even when I had begun to believe that I would never finish and would die penniless and alone working in an East Lansing burger joint (I got into these states of mind frequently in graduate school—if you don't, or didn't, you're made of pretty stern stuff). I am also grateful to the rest of my committee, Dr. Ana Maria Herrera, Dr. Jeffrey Wooldridge and Dr. Christine Amsler, for further indulgences, as a student (in the case of Drs. Herrera and Wooldridge) and as a teaching assistant (in the case of Dr. Amsler), as well as an advisee.

For input and comments on various portions of this dissertation in its draft stages as well as other sources of encouragement and support, I am grateful to my fellow students in the Department of Economics at MSU, including (among others) Scott Adams, Linda Bailey, Vinit Jagdish, Mark Jahn, Jongbyung Jun, Alina Luca, Iva Petrova, Artem Prokhorov, Monika Tothova, and David Wetzell. Outside the department, Dr. Gerhard Glomm of Indiana University encouraged and stimulated my interest in macroeconomic research, as well as helping Dr. Pecchenino help me keep the faith. I am grateful to Dr. Glomm and to Dr. Eric Leeper of Indiana University for comments on my research, as well as to seminar participants at MSU, at the Bank of Canada (especially Eva Ortega and Greg Tkacz), and at the 2003 and 2004 meetings of the Midwest Economic Association. One non-economist academic I'm also grateful to for comments (and attending my defense as Dean's Representative) is Dr. Dennis Gilliland of the MSU Department of Statistics.

Going back a bit farther, I am also grateful to Drs. Michael Dowd and David C. Black of the University of Toledo who first encouraged my study of economics, and to Dr. James Lesage of UT, my first teacher in (especially Bayesian) econometrics. Without them I'd have never started graduate school, never mind finished; my research presented here is in a real sense a continued exploration of the topics I became interested in as an undergraduate at UT.

Food (ideally sushi; Taco Bell is a highly imperfect substitute), drink (Beaner's coffee, not beer) and shelter (in Spartan Village) are also vital inputs for an economist. In a free-market monetary economy of the sort I study on most days, this requires cash. Hence, I will remain forever grateful to MSU's University Distinguished Fellowship program for financial support during most of my graduate school career.

The non-academic types who contributed physical, emotional and spiritual support are numerous enough that they'll have to be content with slightly briefer mentions. They fall for the most part into three categories:

My family: For everything. Especially: my mother; who had faith I could do it; my father, who forced me to prove it; my older brother Robert, who gave me something to aspire to (including but not limited to a real job); and my younger brother and sister, Oliver and Amelia, who let me return the favor. (Amelia was eight when I started college; she turned eighteen in May 2004. So this dissertation has been more than half her life in the making, and in all that time I still haven't managed to explain to her what I do all day. Now you know, Amelia. You still sure you want to go to grad. school?)

Animation fans (don't laugh): For giving me something to do that wasn't research and hence keeping me from going crazy. Most of these are members of the MSU animation club that I have had the privilege of advising for most of my graduate school career (on my own time, naturally); these include (among others) Mike and Stephanie Brandl, Jae-Hwang and Michelle Park, Jacinta Raphael, Al Roderick, Lisa and Melanie Schoen, Jen Stepanski, Denise Wyatt, and Douglass Weeks. Some non-members to whom I owe as great a debt include Jonathan Clements and Randall and Rachel Larson.

William G. Wilson and his friends: I would elaborate, but this acknowledgments section is already too long, and in any case friends of Bill prefer to remain anonymous.

Thank you and God bless you all.

Paul Corrigan

East Lansing, Michigan

Canada Day, July 1, 2004

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

INTRODUCTION

In an introduction to a book, especially one distilled from a dissertation for academic readers, one generally gives a roadmap to a book and explains what each chapter is about, when in fact that's what you would think an abstract and a table of contents listing the chapter titles were for. At any rate, the papers have each their own introductions explaining what they're actually about, so I suggest you read those instead. I must say I'm flattered you've managed to even read this far, mind you, so as a reward, and/or on the off chance I win the Nobel Prize and any historians of thought happen to go back and read this (vain hope on either count) I'll give you a sneak preview of my forthcoming memoirs and give a bit of background as to how the papers that make up this dissertation came to be written.

My research program thus far has mostly been devoted to satisfying myself that there is a marginal value to economic theory in prediction of macroeconomic activity and that I hadn't been wasting my time taking bunches of courses in economics when a simple autoregression could tell me all I needed to know about what the economy was going to do next week. Obviously mathematical macroeconomic models are grossly oversimplified pictures of macroeconomic reality, but surely not all were not so oversimplified that they could not have more than pedagogical value. When as a child (at the University of Toledo) I learned at the knee of James Lesage the rudiments of economic forecasting using Bayesian vector autoregressions, I quickly grew dissatisfied with the atheoretical priors that were so often used in these models; surely a simple RBC model, even if it didn't fit the data, could be used to add information to a VAR that could help in forecasting economic fluctuations.

I could swear that it was Gerhard Glomm who referred me to the Ingram and Whiteman (1994) paper, describing a technique for using an RBC model to derive a prior for a BVAR, though Dr. Glomm himself doesn't remember doing so. At any rate, I ran with it, and my third-year paper at MSU, which eventually became the first chapter of this dissertation, was an attempt to apply Ingram and Whiteman's method, to see if New Keynesian models could help in forecasting output and inflation.

It was an early experiment, and (as will become clear from reading it) not totally successful. Eventually a more successful exercise, and a somewhat more elegant one, came to my attention in the form of del Negro and Schorfheide (2003), just as I was finishing the first paper. However, it was thinking about how and why the forecasting experiments with the New Keynesian prior failed (was it a prior information problem? A specification problem?) led me to try to take the New Keynesian model more seriously and see what would be required to get the New Keynesian model to fit the data properly, rather than just write it off as useful only as a guide for a less theoretical model.

What resulted was the second chapter of this dissertation. It was started before the third, obviously, but finished not until a little bit after. In part this was because I didn't fully trust what my results were telling me, so, being the perfectionist I am, I wound up tweaking and analyzing the sensitivity of my results until the very last moment (much to the chagrin of Rowena Pecchenino, who'd surely have preferred I just put it to bed). It answered the question of how the model in the first paper had failed—poor dynamics that were easily corrected with addition of another friction—and the result was a DSGE model that actually fit the data better and gave more precise results than a vector autoregression regarding the effects of monetary shocks on output (namely, quite small overall). What I did find puzzling was that the output and inflation data itself didn't seem to have much information at all regarding the amount of price stickiness in the economy—which was rather worrying, given how many studies had essentially mined just the output and inflation data for evidence on price stickiness in the US economy. More information was necessary, which meant a bigger model. It also meant disentangling inflation from other demand shocks and comparing their effects and importance on the business cycle.

The second chapter was basically an application of Schorfheide's (2000) method of assessing the goodness of fit of DSGE models and evaluating the deficiencies thereof as compared to other less theoretical models like VARs (though ironically now it was the VARs that were deficient, in part because they weren't parsimonious or precise enough). The third chapter was more analogous to Dejong et al. (2000b), which assessed the role of various types of non-monetary shocks in economic fluctuations, in investment as well as output. The third chapter is a similar exercise, as far as that goes, but I was uncomfortable with not checking the goodness of fit of my DSGE model before trying to construct shocks with my model (if the model doesn't fit the data, it's not clear how meaningful the "shocks" it spits out really are). So I was sure to do that before constructing series of shocks, and finding that the experiments of the second and third chapters of my dissertation were consistent in this: that monetary shocks were relatively unimportant in output fluctuations, and that technology shocks explain business cycles best for reasonable levels of price stickiness. A DSGE model with New Keynesian features (surprisingly for many people, including myself) is not necessarily incompatible with a supply-shock driven, essentially "old-style RBC" explanation of business cycles.

What is also perhaps surprising, but more comforting, is that these DSGE models are well-specified enough that they can be taken seriously when they try to tell us all this. Praise be. My economics courses were worthwhile after all.

With Bayesian DSGE models proven to be competitive in forecasting and goodness of fit with VAR macroeconomic models, and offer greater precision in analyzing the nature of business cycles than the VARs can, there is much interest in them currently (in 2004) at such research institutions as the Federal Reserve Bank of Atlanta and the Bank of Canada. I probably shouldn't speculate on whether such models are the wave of the future in empirical macro—forecasting of any sort is a formidable task. Research in this area does look promising enough, though, that I do plan to make more contributions in this area that (please God) will be of use to macroeconomic researchers, perhaps (vain hope) even policymakers. For now, though, this opus will have to do.

On a related note (thanks to Dennis Gilliland for reminding me): Bayesian DSGE research is a young enough literature (in 2004) that routines for estimating such models aren't available "canned" in statistical packages. Therefore being partially out of my mind I decided to cook entirely from scratch my Mathematica 4.1 code for estimating Bayesian VARs with New Keynesian priors, as well as Bayesian DSGE models, instead of stealing Ingram and Whiteman's or Schorfheide's like Ana María Herrera wanted me to. So if you have Mathematica 4.1 or later and want to amuse yourself on a rainy day by trying to replicate my results, or just want a sample routine for estimating Bayesian DSGE models that you can adapt to suit your own research needs, please contact me care of the MSU Department of Economics and I'll be happy to send you copies of the code.

Eh bien, continuons.

CHAPTER 1

FORECASTING OUTPUT AND INFLATION WITH A BAYESIAN VAR USING A NEW KEYNESIAN PRIOR

Introduction

So-called New Keynesian pricing models have become the principal models used by economists to model inflation and its effects on real variables such as output and employment. New Keynesian pricing models incorporate "old Keynesian" features into modern general equilibrium macroeconomic models by incorporating price setting by imperfectly competitive firms that do not have perfect flexibility in their price-setting, either because they face menu costs of changing their price in any given period (Rotemberg 1982, 1987), simply do not have a chance to adjust their price every period (Calvo 1983), or both. Both infrequent opportunities for price setting, pace Calvo, or menu costs, pace Rotemberg, lead to processes for inflation of the structure $\pi_t = \beta E_t \pi_{t+1} + \lambda mc_t$, where π_t is inflation and mc_t is real marginal cost. Inflation in Calvo-Rotemberg type models is then the present value of real marginal cost, assumed, in models with imperfect competition, to be high when output is high.

Part of the popularity of New Keynesian pricing models is that such models can be used to rationalize expectations-augmented Phillips curves, which are widely depended on as useful models of the output-inflation process in applied macroeconomic work (e.g. Stock and Watson 1999). However, attempts to estimate the implied New Keynesian Phillips curve have had at best mixed results, in part because of the lack of a good measure of deviations of marginal cost from trend. Early studies like that of Roberts (1995) used a measure of the output gap (deviation of real GDP from a calculated trend) as a proxy for marginal cost. Roberts found average price fixity for the US economy of about five quarters. Fuhrer and Moore (1995) criticized Roberts' study in part because it attempted to estimate the New Keynesian Phillips' curve in isolation. Fuhrer and Moore's own findings, using full-information techniques including the Phillips curve as pat of a system also including an equation describing behavior of the output gap, found that a standard New Keynesian Phillips curve did not fit the output-inflation data very well; they suggested that the forward-looking New Keynesian Phillips curve be augmented with elements of backward-looking price setting.

A more sophisticated limited-information study by Gali and Gertler (1999) used labor income share, rather than deviation of output from trend, as their measure of marginal cost. They, like Roberts, found average price fixity in the rage of four to six quarters. However, like Roberts, Gali and Gertler estimated the New Keynesian Philips curve in isolation. Attempts to estimate a more complete model, examining in particular the implications for output dynamics of an inflation process that fits Calvo-Rotemberg specifications, are not nearly as favorable to the model. The findings of Kurmann (2001) are typical. Kurmann, using full-information econometric techniques, found that the rational expectations restrictions implied by a Calvo-Rotemberg model that generated reasonable inflation dynamics implied an effect of inflation on future real marginal cost (measured by labor income share) that was very much at odds with that implied by an unrestricted estimation of the forecasting process for real marginal cost would have implied. Specifically, the unrestricted model of the process for marginal cost (and by implication output and employment) implied a small and insignificant effect of inflation. With the rational expectations restrictions imposed, on the other hand, the inflation effect on marginal cost became implausibly large and the inflation terms summed to near zero, implying that changes in inflation might be better predictors of marginal cost/output, as Fuhrer and Moore (1995) argue.

Kurmann cites two main suspects that might be to blame for the failure of the Calvo-Rotemberg model under full-information. Either the Calvo-Rotemberg process for inflation is misspecified (because, in particular, it ignores the possibility of backward-looking inflation setting pace Fuhrer and Moore), or the labor income share, like the output gap, is a poor proxy for real marginal cost, or both. Almost certainly the answer is "both." If we assume some imperfect competition in the economy, average labor productivity will not equal the marginal product of labor, and real marginal cost will not be proportional to labor income share. Part of the problem, though, might come from taking the Calvo-Rotemberg model of inflation too literally, taking it to be a "true" model rather than as a first approximation to a more complex, unknown inflation process, and imposing it exactly on the data.

It might be helpful to impose the restrictions on output and inflation dynamics more loosely on the data, to see if the Calvo-Rotemberg model can give us any guidance on their probable true form, even if it does not fit exactly. Also, a professional forecaster less interested in determining the "true" model of inflation, which is unknown (and likely to remain so) than in making accurate forecasts might be interested in determining the marginal value of New Keynesian theories of the output-inflation relationship in making predictions about the future behavior of the real and nominal sides of the economy.

In this paper, I attempt to evaluate the forecasting performance of a simple computable competitive general equilibrium model (DSGE or "RBC" model) that allows for inflation to affect output by including price stickiness of the Calvo-Rotemberg type. I do this by using such a model, with structural parameters corresponding to the values in U.S. data, to calculate a prior mean and covariance matrix for the coefficients of a VAR forecasting model for output, inflation and interest rates, using a technique described by Ingram and Whiteman (1994) for using DSGE models to calculate priors for Bayesian VAR forecasting models. The Calvo-Rotemberg inflation process, in particular, is calibrated to the specifications implied by Gali and Gertler (1999). As no good proxies for real marginal cost exist, I do not attempt to specify a forecasting process for marginal cost, but rather spin from the New Keynesian model prior information regarding the dynamics of the output gap and the nominal interest rate, as well as output.

Allowing for a sufficiently wide range for the structural parameters of the model permits me to add this prior information to the data without imposing exact restrictions. In particular, the Calvo-Rotemberg inflation process itself is only regarded as a first approximation to a more general inflation process, which could include longer lags of output than just the first (as in the Calvo-Rotemberg model), allowing for backward looking inflation setting of the Fuhrer-Moore type. I do find that that prior information does allow me to generate forecasts of inflation that are at least competitive, and often significantly better, than those from a leading, atheoretical VAR forecasting model. However, I also find that even the loose imposition of the output dynamics from the New Keynesian model results in very poor output forecasts, particularly in the short run, suggesting that as a model of the output-inflation relationship such models are not especially promising even as a first approximation.

The structure of the rest of the paper is as follows. Section 2 describes the DSGE model with New Keynesian features, incorporating Calvo-Rotemberg type price stickiness to allow inflation to affect output, that I use to generate a New Keynesian prior for forecasting. Section 3 describes the technique for generating the New Keynesian prior and imposing it on a Bayesian VAR (BVAR) forecasting model, a technique closely following Ingram and Whiteman (1994). Section 4 will describe the performance of the New Keynesian BVAR relative to competing forecasting models, in particular to a BVAR using a modified version of the popular "Minnesota" prior developed by Litterman (1984). Section 5 concludes with a discussion of the implications of the results for theory and for practical forecasting applications.

Model

The economy I use to derive the prior is similar in structure to those used in King and Wolman (1996), Yun (1996), and others in the New Keynesian literature. The economy is inhabited by a large number of consumers that supply labor and capital to monopolistically competitive firms, using the proceeds from wages and rental income to buy consumer goods and acquire real money balances. I assume consumers demand real money balances to save them time in making transactions and increase their leisure time. The typical consumer i has preferences on his own consumption c_{it}^{d} of a composite of goods supplied by the monopolistically competitive firms, labor h_{it} and holdings of real money balances M_{it}^{d} / P_t where M_{it}^{d} is his nominal money holdings and P_t is the price level, a composite of the prices of all goods supplied by the firms. (The means by which goods and prices are transformed into the composite consumer good and price level will be defined more clearly below.) Those preferences are

$$E_0 \sum_{t=0}^{\infty} \beta^t \left(\ln c_{it}^d + a \ln \left(1 - h_{it}^s - \frac{\chi}{1-\chi} \varepsilon_t \left(\frac{P_t c_{it}^d}{M_{it}^d} \right)^{\frac{1-\chi}{\chi}} \right) \right)$$

where β is a discount factor assumed to be less than one. The taste shock ε_t follows an autoregressive process such that $\ln \varepsilon_t = (1-\zeta) \ln \varepsilon_0 + \zeta \ln \varepsilon_{t-1} + \upsilon_{\varepsilon t}$, where ζ is less than or equal to one. Allowing for variations in ε_t is a crude way of modeling changes in financial technology (a fall in ε_t reduces the amount of time one must devote to shopping, given a certain ratio of money balances to required transactions, allowing one to reduce one's money balances).

In each period the consumer starts with $M_{it-1}^d / P_t = (M_{it-1}^d / P_{t-1}) \times (P_{t-1} / P_t)$ worth of real money balances, k_{it-1}^s units of capital and $B_{it-1}^d / P_t = (B_{it-1}^d / P_{t-1}) \times (P_{t-1} / P_t)$ worth of bonds, carried over from the previous period. The consumer rents out his k_{it-1}^s units of capital and h_{it}^s units of labor to firms in competitive labor and capital markets, getting in return $w_t h_{it}^s$ in wages and a gross return on his capital of $(1-\delta + r_t) k_{it-1}^s$. He also receives a "helicopter drop" of fresh money balances from the monetary authority of value x_t / P_t, along with revenue from maturation of bonds of B_{it-1}^d / P_t. The consumer can now use his income and his holdings of real money balances to purchase c_{it}^d units of consumer goods, M_{it}^d / P_t worth of real money balances and k_{it}^s units of capital, along with B_{it}^d / P_t worth of bonds at

$1/(1+R_t)$ dollars per dollar worth of bonds. The consumer's budget constraint is therefore

$$c_{it}^d + k_{it}^s - (1-\delta)k_{it-1}^s + \left(\frac{M_{it}^d}{P_t}\right) - \frac{P_{t-1}}{P_t}\left(\frac{M_{it-1}^d}{P_{t-1}}\right) - \frac{x_t}{P_t} + \frac{1}{1+R_t}\left(\frac{B_{it}^d}{P_t}\right) - \frac{P_{t-1}}{P_t}\left(\frac{B_{it-1}^d}{P_{t-1}}\right) - r_t k_{it-1}^s - w_t h_{it}^s \le 0,$$

allowing for free disposal.

Define λ_{it} as the marginal utility of consumption. The first order conditions of the consumer's problem are

(C1: cons. dmd.)

$$\beta^t\left(\frac{1}{c_{it}^d} - a\varepsilon_t c_{it}^{d\,\frac{1-2\chi}{\chi}}\left(\frac{M_{it}^d}{P_t}\right)^{-\frac{1-\chi}{\chi}}\left(1-h_{it}^s - \frac{\chi}{1-\chi}\varepsilon_t\left(\frac{P_t c_{it}^d}{M_{it}^d}\right)^{\frac{1-\chi}{\chi}}\right)^{-1}\right) = \lambda_{it}$$

(C2: labor supply) $\quad a\beta^t\left(1-h_{it}^s - \frac{\sigma}{1-\sigma}\varepsilon_t\left(\frac{P_t c_{it}^d}{M_{it}^d}\right)^{\frac{1-\chi}{\chi}}\right)^{-1} = \lambda_{it} w_t$

(C3: money demand)

$$a\beta^t\varepsilon_t c_{it}^{d\,\frac{1-\chi}{\chi}}\left(\frac{M_{it}^d}{P_t}\right)^{-\frac{1}{\chi}}\left(1-h_{it}^s - \frac{\sigma}{1-\sigma}\varepsilon_t\left(\frac{P_t c_{it}^d}{M_{it}^d}\right)^{\frac{1-\chi}{\chi}}\right)^{-1} = \lambda_{it} - E_t\left(\frac{P_t}{P_{t+1}}\lambda_{it+1}\right)$$

(C4: capital supply) $\quad \lambda_{it} = E_t((1-\delta+r_{t+1})\lambda_{it+1})$

(C5: dmd. for bonds) $\quad \frac{1}{1+R_t}\lambda_{it} = E_t\left(\frac{P_t}{P_{t+1}}\lambda_{it+1}\right)$

Some rearrangement and combining (C1) with (C2), and (C3) with (C2) and (C5), yields the more manageable consumption and money demand functions

$$\text{(C1')} \qquad \beta^t (1/c_{it}^d) = \left(1 + \varepsilon_t c_{it}^{d\,\frac{1-2\chi}{\chi}} (M_{it}^d / P_t)^{-\frac{1-\chi}{\chi}} w_t \right) \lambda_{it}$$

$$\text{(C3')} \qquad \frac{M_{it}^d}{P_t} = \varepsilon_t^\chi c_{it}^{d^{1-\chi}} w_t^\chi \left(\frac{1+R_t}{R_t} \right)^\chi$$

The composite consumer good demanded by individuals is a composite, as noted above, of the individual goods produced by a large number of monopolistically competitive firms, which I shall index $i \in [0,1]$. The outputs Y_{it} of each firm i are aggregated into a single output Y_t in the fashion $Y_t = \left(\int_0^1 Y_{it}{}^{(\omega-1)/\omega} di \right)^{\omega/(\omega-1)}$, where ω>1 is the elasticity of substitution between the two goods. Minimizing the cost of $P_t Y_t \equiv \int_0^1 P_{it} Y_{it} di$, given Y_t, results in a demand function for firm i of $Y_{it} = (P_{it} / P_t)^{-\omega} Y_t$. Substituting into the cost function yields the composite price index $P_t = \left(\int_0^1 P_{it}{}^{1-\omega} di \right)^{1/(1-\omega)}$.

The typical firm's production function implies increasing returns to scale; there is a fixed cost to production of h_0 units of labor (which can be thought of as the entrepreneur's labor input), but the function, I'll assume, is otherwise Cobb-Douglas in capital input k_{it}^d and "net" labor input $h_{it}^d - h_0$. The firm purchases capital at r_t a unit and labor (including the entrepreneur running the firm) at w_t. The firm's cost minimization problem, given output of Y_{it}, is $tc_t(i) = \min_{h_{it}^d, k_{it}^d} r_t k_{it}^d + w_t h_{it}^d$ such

that $Y_{it} \geq k_{it}^{d\,1-\theta}(z_t(h_{it}^d - h_0))^\theta$, where z_t is a measure of technology following the autoregressive process $\ln z_t = \gamma_0 + (1-\rho)\gamma t + \rho \ln z_{t-1} + \upsilon_{zt}$ for $\rho < 1$ and $\ln z_t = \gamma_0 + \ln z_{t-1} + \upsilon_{zt}$ for $\rho = 1$. Solving the problem results in a cost function of $tc_{it} = w_t h_0 + mc_t Y_{it}$, where marginal cost $mc_t = z_t^{-\theta} r_t^{1-\theta} w_t^\theta \theta^{\theta-1}(1-\theta)^{-\theta}$. The resulting factor demands, along with the production function, are

(F1: labor demand) $$h_{it}^d = h_0 + Y_{it} z_t^{-\theta} r_t^{1-\theta} w_t^{\theta-1} \theta^\theta (1-\theta)^{-\theta} = h_0 + mc_t Y_{it} / w_t$$

(F2: capital demand) $$k_{it}^d = Y_{it} z_t^{-\theta} r_t^{-\theta} w_t^\theta \theta^{\theta-1}(1-\theta)^{1-\theta} = mc_t Y_{it} / r_t$$

(F3: production function) $$Y_{it} = k_{it}^{d\,1-\theta}(z_t(h_{it}^d - h_0))^\theta$$

The markup at time t, μ_t, is $1/mc_t$. Given that $r_t k_{it}^d = (1-\theta) mc_t Y_{it}$, it follows that, given free entry, profits are zero in the long run, the labor share of output is $\phi \equiv wh/Y = (1-(1-\theta)/\mu)$, and so $\theta = 1 - \mu(1-\phi)$, where μ is the long run level of the markup (to be derived below) and h is the average amount of labor effort. As in the long run $mcY = Y/\mu = w(h/h_0) + rk$, $wh_0 = (1-1/\mu)Y$, and so the ratio of h_0 to h is given by $h_0/h = (\mu-1)/\mu\phi$.

The firm, with its output in hand, can now sell it for the amount $(P_{it}/P_t)Y_{it} = (P_{it}/P_t)^{1-\omega} Y_t$; profits in the current period are hence given by

$$((P_{it}/P_t) - mc_t)Y_{it} - w_t h_0 = ((P_{it}/P_t)^{1-\omega} - mc_t(P_{it}/P_t)^{-\omega})Y_t - w_t h_0 .$$

If prices were perfectly flexible, profit maximization would be as simple for the firm as maximizing this function over (P_{it} / P_t) and getting $(P_{it} / P_t) = \omega /(\omega - 1) mc_t$; furthermore, as all firms face the same problem, it would be the case that $P_{it} = P_t$ and so $mc_t = (\omega - 1)/\omega$ for all t. (This gives us the long-run value of the markup, $\mu = 1/mc = \omega/(\omega - 1)$.) However, the whole point of the model is that prices are not perfectly flexible. Specifically, following Calvo (1983), in each period firms face a probability $\alpha<1$ that they will not be allowed to adjust their price freely in that period, but will instead be forced to raise their price by a factor of π, the long run inflation rate. When a firm does get a chance to adjust its price, of which there is a probability $1-\alpha$ (so that on average a firm gets to adjust its price every $1/(1-\alpha)$ periods, it must face the probability α^j of that the price it sets will remain fixed (except for the rigid adjustments by a factor of π) for j periods, and adjust the price accordingly so as to maximize the discounted value of lifetime profits, given P_{it}. The firm solves the expected profit maximization problem

$$\max_{p_{it}} E_t \sum_0^\infty (\alpha\beta)^j (\lambda_{t+j} / \lambda_t)(((\pi^j P_{it} / P_{t+j})^{1-\omega} - mc_t (\pi^j P_{it} / P_{t+j})^{-\omega}) Y_{t+j} - w_{t+j} h_0)$$

.

If we define P_t^* as the price chosen by the proportion $1-\alpha$ of firms allowed to adjust their prices freely in each period, we get

$$\text{(F4: Price-setting equation)}\ P_t^* = \frac{\omega}{\omega - 1} \frac{E_t \sum_0^\infty (\alpha\beta)^j (\lambda_{t+j} / \lambda_t)(Y_{t+j} / Y_t) mc_{t+j} P_{t+j}^{\varpi} \pi^{-j\omega}}{E_t \sum_0^\infty (\alpha\beta)^j (\lambda_{t+j} / \lambda_t)(Y_{t+j} / Y_t) P_{t+j}^{\varpi - 1} \pi^{-j(\omega-1)}}$$

Meanwhile, the form for the price index suggests the function for the evolution of the price $P_t = \left(\int_0^1 P_{it}{}^{1-\omega} di \right)^{1/(1-\omega)}$ is

(M1: Price index process) $$P_t^{1-\omega} = (1-\alpha)P_t^{*1-\omega} + \alpha(\pi P_{t-1})^{1-\omega}.$$

Finally, defining q_t as the economy-wide aggregate for allocation q, the market clearing conditions for this economy is

(M2: Goods market-clearing condition) $$c_t^d + k_t^s - (1-\delta)k_{t-1}^s = Y_t.$$

(M3: Factor market clearing conditions) $$h_t^d = h_t^s, k_t^d = k_{t-1}^s$$

(M4: Money market-clearing condition) $$M_t^d = M_t^s = M_t = M_{t-1} + x_t.$$

To close the system completely, I need to define the rule by which the monetary authority selects the cash payments x_t / P_t; the rule, however, is stated in terms of a log-linearized form of the model I have just described, and so I will put off describing the rule until later. An equilibrium for this economy is a collection of allocations $\{Y_t, C_t, K_t, h_t, B_t, (M/P)_t\}$ and prices $\{w_t, r_t, P_t, P_t^*, R_t\}$ such that the consumers' and firms' problems are solved and all markets clear.

To get the prior for the Bayesian VAR, Ingram and Whiteman (1994) suggest log-linearizing the first-order conditions for the consumers' and firms' problems and the market-clearing conditions and using those to solve for the approximate decision rules and state-variable transition matrix. Then, the VAR representation of the variables of interest can be derived and a prior mean and covariance matrix calculated.

Log-linearizing the consumer's first-order conditions is straightforward, as is the goods market clearing condition. Letting a circumflex indicate deviation from trend,

(C1: cons. dmd.)

$$-\hat{c}_t = \frac{R_0(M/PY)}{(C/Y)(1+R_0)+(M/PY)}\left(\hat{\varepsilon}_t + \frac{1-2\chi}{\chi}\hat{c}_t - \frac{1-\chi}{\chi}(M/P)_t + \hat{w}_t\right) + \hat{\lambda}_t$$

(C2: labor supply) $\quad \hat{\lambda}_t + \hat{w}_t = -\frac{h}{l}h_t - \frac{1-h-l}{l}\left(\hat{\varepsilon}_t + \frac{1-\chi}{\chi}(c_t - (M/P)_t)\right)$

(C3: money demand) $(M\,\hat{/}\,P)_t = \chi\hat{\varepsilon}_t + (1-\chi)\hat{c}_t + \chi\hat{w}_t + \frac{\chi}{R_0}(1\,\hat{+}\,R_t)$

(C4: capital supply) $\quad \hat{\lambda}_t = E_t\hat{\lambda}_{t+1} + \frac{r}{1-\delta+r}E_t\hat{r}_{t+1}$

(C5: dmd. for bonds) $\quad \hat{\lambda}_{it} - (1\,\hat{+}\,R_t) = E_t\hat{\lambda}_{t+1} - E_t\hat{\pi}_{t+1}$

Above R_0 is the long-run average nominal interest rate, $r = (\gamma/\beta)-1+\delta$ is the long-run average gross rate of return on capital, (M/PY) is the inverse of velocity, (C/Y) is the average ratio of consumption to output, and $\pi_t \equiv P_t / P_{t-1}$ is the inflation rate. l, the average portion of each day not devoted to market activities (including shopping), and h is the average portion of each day spent working; hence $1-h-l$ will be the average amount of time spent shopping, viz. $\frac{\chi}{1-\chi}\varepsilon\left(\frac{PC}{M}\right)^{\frac{1-\chi}{\chi}} = \frac{\chi}{1-\chi}\varepsilon\left(\frac{C/Y}{M/PY}\right)^{\frac{1-\chi}{\chi}}$. The aggregate money demand function implies $\varepsilon = \left(\frac{M/PY}{C/Y}\right)^{1/\chi}\left(\frac{hC}{\phi Y}\right)\frac{R}{1+R}$, and so

$l = 1 - h\left(1 + \frac{\chi}{1-\chi}\frac{1}{\phi}\frac{R_0}{1+R_0}\frac{M}{PY}\right)$, or equivalently, $h = \frac{1-l}{\left(1+\frac{\chi}{1-\chi}\frac{1}{\phi}\frac{R_0}{1+R_0}\frac{M}{PY}\right)}$.

Meanwhile, the log-linearized market-clearing condition for the goods market is

(M2: Goods market-clearing condition) $\quad \frac{C}{Y}\hat{c}_t + \frac{K}{Y}\hat{k}_t - \frac{1-\delta}{\gamma}\frac{K}{Y}\hat{k}_{t-1} = \hat{Y}_t$.

As $r=(\gamma/\beta)-1+\delta$, K/Y will equal $(1-\phi)/r=\beta(1-\phi)/(\gamma-\beta(1-\delta))$. It follows that $C/Y=1-I/Y=1-(\gamma-1+\delta)K/Y$.

The log-linearized factor demands are

(F1: labor demand) $$\hat{w}_t = \hat{mc}_t + \hat{Y}_t - \frac{h}{h-h_0}\hat{h}_t = \hat{mc}_t + \hat{Y}_t - \frac{\mu\phi}{1-\mu(1-\phi)}\hat{h}_t$$

(F2: capital demand) $$\hat{r}_t = \hat{mc}_t + \hat{Y}_t - \hat{k}_{t-1}$$

(F3: production function) $$\hat{Y}_t = (1-\mu(1-\phi))\hat{z}_t + \mu(1-\phi)\hat{k}_{t-1} + \mu\phi\hat{h}_t$$

Log-linearizing (F4) gives us

$$\hat{P}_t^* = (1-\alpha\beta)\left(\sum_0^\infty(\alpha\beta)^j\left(E_t\hat{mc}_{t+j} + E_t\hat{\lambda}_{t+j} + E_t\hat{Y}_{t+j} + \omega E_t\hat{P}_{t+j}\right)\right)$$
$$-(1-\alpha\beta)\left(\sum_0^\infty(\alpha\beta)^j\left(E_t\hat{mc}_{t+j} + E_t\hat{\lambda}_{t+j} + E_t\hat{Y}_{t+j} + (\omega-1)E_t\hat{P}_{t+j}\right)\right)$$
$$= (1-\alpha\beta)\left(\sum_0^\infty(\alpha\beta)^j\left(E_t\hat{mc}_{t+j} + E_t\hat{P}_{t+j}\right)\right),$$

implying $\hat{P}_t^* - \alpha\beta E_t P_{t+1}^* = (1-\alpha\beta)(\hat{mc}_t + \hat{P}_t)$.

Log-linearizing the price index process (M1) gives us $\hat{P}_t = (1-\alpha)P_t^* + \alpha\hat{P}_{t-1}$, given that in the long run $P_t^* = P_t$. We can use the fact that $\hat{\pi}_t = \hat{P}_t - \hat{P}_{t-1}$ to get $\hat{\pi}_t = (1-\alpha)(\hat{P}_t^* - \hat{P}_t)$. Substituting into (F4) (and a little rearrangement) finally gets us our Calvo-Rotemberg type Phillips curve:

(F4': Phillips curve) $$\hat{\pi}_t = \beta E_t\hat{\pi}_{t+1} + \frac{(1-\alpha)(1-\alpha\beta)}{\alpha}\hat{mc}_t.$$

It remains to define the monetary policy rule for this economy. Economic agents assume (correctly) that the central bank sets the nominal interest rate according to the rule:

(P1: Policy rule)

$$E_t(1\hat{+}R)_{t+1} = \xi(1\hat{+}R)_t + (1-\xi)(R_y E_t(\hat{Y}_{t+1} - \hat{Y}^f_{t+1}) + R_\pi(E_t\hat{\pi}_{t+1} + E_t\hat{Y}_{t+1} - E_t\hat{Y}_t))$$

This is a combination of a Taylor rule, relating the nominal interest rate target to deviation of output from "potential" (here assumed to be $\hat{Y}^f_t$, the level of output that would prevail in the absence of price-stickiness) and inflation, and of a nominal income target rule. A rule resembling this one was estimated by McCallum and Nelson (1999) in the course of estimating a New Keynesian structural macro model for use in policy analysis. The ξ parameter is a measure of the degree to which the monetary authority smoothes interest rates, adjusting the nominal interest rate only gradually towards the target $R_y(\hat{Y}_t - \hat{Y}^f_t) + R_\pi(\hat{\pi}_t + \hat{Y}_t - \hat{Y}_{t-1})$. To be sure, as McCallum and Nelson freely admit, actual Fed behavior is best described as discretionary rather than as following any necessarily simple rule like this one. The "rule" used here is probably best thought of as a crude attempt to summarize a Fed reaction function of a form unknown even to the policymakers themselves.

Implementation

Given values for the structural parameters $\{\alpha, \beta, \chi, \delta, \phi, \gamma, \mu, \rho, \xi, R_y, R_\pi, (M/PY), R_0\}$, I can derive an approximate solution for the decision rules for variables of interest, as well as the processes for the state variables, capital k_t, technology z_t, the taste shock for money demand ε_t, and the nominal interest rate $(1+R_t)$. In the interest of making the model parsimonious, and because I am not interested in estimating money demand separately from output, I set ζ=0 and ε_t equal to a constant for all t. That leaves three state variables, capital, technology, and interest rates, allowing separate estimation of three endogenous variables. I estimate a three-variable VAR in output Y_t, inflation π_t, and the nominal interest rate $(1+R_t)$, which will serve as my measure of monetary policy. Detrended capital, technology and the interest rate are expected to follow the process

$$\text{(1)} \qquad \begin{bmatrix} \hat{k}_t \\ E_t\hat{z}_{t+1} \\ E_t(1\hat{+}R)_{t+1} \end{bmatrix} = \mathbf{S} \begin{bmatrix} \hat{k}_{t-1} \\ \hat{z}_t \\ (1\hat{+}R)_t \end{bmatrix} \text{where } \mathbf{S} = \begin{bmatrix} \varphi_{kk} & \varphi_{kz} & \varphi_{kR} \\ 0 & \rho & 0 \\ \varphi_{Rk} & \varphi_{Rz} & \varphi_{RR} \end{bmatrix}.$$

Meanwhile, the mapping from the detrended state variables to the detrended endogenous variables is

$$\text{(2)} \qquad \begin{bmatrix} \hat{Y}_t \\ \hat{\pi}_t \\ (1\hat{+}R)_t \end{bmatrix} = ? \begin{bmatrix} \hat{k}_t \\ \hat{z}_t \\ (1\hat{+}R)_t \end{bmatrix} \text{where } ? = \begin{bmatrix} \varphi_{yk} & \varphi_{yz} & \varphi_{yR} \\ \varphi_{\pi k} & \varphi_{\pi z} & \varphi_{\pi R} \\ 0 & 0 & 1 \end{bmatrix}$$

Let $\hat{\mathbf{y}}_t = [Y_t \quad \hat{\pi}_t \quad (1\hat{+}R)_t]'$ and define $E_t\hat{\mathbf{y}}_{t+1} = \mathbf{F}\hat{\mathbf{y}}_t$. Using (1) and (2), it is straightforward to show that **F**=**ΠSΠ**$^{-1}$. The entries in **F** are completely determined by the structural parameters of the underlying log-linearized model.

A sixth-order vector autoregression of $\hat{y}_{it}$, the ith component of $\hat{\mathbf{y}}_t$, can be written as $\hat{y}_{it} = \sum_0^6 \mathbf{F}_{ij}\hat{\mathbf{y}}_{t-j}$, $t = 1,\ldots,T$, $i = 1,2,3$, where T is the number of observations. My prior mean for the elements $\mathbf{F}_{i1}$, the vector corresponding to the first lag coefficients of equation i, is given by the mean of the associated elements in the ith row of **F**; my prior mean for the elements of $\mathbf{F}_{ij}$ for j>1 are all zero. The prior covariance matrix for the first lag coefficients is given by the covariance matrix for the corresponding row of **F**; the covariance matrix for jth lag coefficients, where j>1, are equal to that for the first lag coefficients divided by a factor of j^2. I combine the prior with the data to obtain a posterior estimate for each equation by simple Theil-Goldberger mixed estimation (Theil and Goldberger 1961; Theil 1971) Let **X**, the data set used in estimation, be the $T\times 18$ matrix consisting of observations on $\hat{\mathbf{y}}_{t-1},\ldots\hat{\mathbf{y}}_{t-6}$, and let $\hat{y}_i$ be the $T\times 1$ vector of observation on $\hat{y}_{it}$. Then we can write the VAR(6) model as $\hat{y}_i = \mathbf{X}\mathbf{B}_i + e_i$. Letting $\overline{\mathbf{B}}_i$ be the prior mean for $\overline{\mathbf{B}}_i$ and Ω_i be its prior covariance matrix, the mixed estimate of $\mathbf{B}_i$ is

$$\tilde{\mathbf{B}}_i = \left[\mathbf{X}'\mathbf{X}/\sigma_i^2 + \Omega_i^{-1}\right]^{-1}\left(\mathbf{X}'\hat{y}_i/\sigma_i^2 + \Omega_i^{-1}\overline{\mathbf{B}}_i\right)^{-1}$$

where σ_i^2 is the estimated variance of the residuals from the corresponding equation estimated in unrestricted VAR.

It remains to specify a prior mean and covariance matrix for **F**. I do that by specifying a prior mean value for the vector of structural parameters $\mathbf{m} = \{\alpha,\beta,\chi,\delta,\phi,\gamma,\mu,\rho,\xi,R_y,R_\pi,(M/PY),R_0\}$; my prior mean for **F** will be the value implied by the prior means of the structural parameters. To calculate the prior covariance, I assume **m** has a diagonal covariance matrix with non-zero entries $\{\sigma_\alpha^2,\sigma_\beta^2,\sigma_\chi^2,\sigma_\delta^2,\sigma_\phi^2,\sigma_\gamma^2,\sigma_\mu^2,\sigma_\rho^2,\sigma_\xi^2,\sigma_{Ry}^2,\sigma_{R\pi}^2,\sigma_{(M/PY)}^2,\sigma_{R0}^2\}$. I then can take a first-order Taylor expansion of **F** about the prior mean. **F** is then approximately normally distributed with mean $\mathbf{F}(\overline{\mathbf{m}})$ and variance $\nabla\mathbf{F}(\overline{\mathbf{m}}) * \text{var}(\overline{\mathbf{m}}) * \nabla\mathbf{F}(\overline{\mathbf{m}})^T$.

The baseline values for the means and standard deviations of the New Keynesian model's structural parameters are given in Table 1.1. The mean values for the discount factor β, the depreciation rate δ, the labor share ϕ, and the persistence parameter for the technology shock ρ, as well as their standard deviations, are taken from a previous study by Dejong et al. (2000a), in which a Bayesian technique was used to calibrate a somewhat different DSGE model. These are well within the range for these variables used in the RBC literature, and the standard deviations allow for ranges for the variables the include most of the values used in that literature. The mean value for the gross rate of per capita economic growth, γ, is also standard; its standard error was calculated as that of the trend rate of growth of real GDP estimated over the period 1959:I to 1983:IV, from an OLS regression (corrected for serial correlation) of logged real GDP on a constant and a time trend. For the markup μ, estimates found in the literature vary widely, from 1.2 to around 2.0; see Rotemberg and Woodford (1995) for a survey. For the prior mean, I picked 1.4, near the average for the values found in the literature (and the value used by Rotemberg and Woodford in their discussion of imperfect competition in DSGE models), along with a standard error allowing for a fairly wide range (from 1.2 to 1.6) for this parameter.

Moving on to the structural parameters for the "demand side" of the economy, the means and standard deviations for the mean nominal interest rate R_0 and the inverse of mean velocity, $(M / PY)_0$, were taken to be the historical means and standard errors over the 1959:I to 1983:IV period for the (quarterly) Federal Funds rate and the inverse of the (quarterly) velocity of the Board of Governors adjusted monetary base. The mean value for the interest elasticity of money demand, χ, is that found for the US economy by

Goldfeld (1973); the standard error allows for the wide range of values for χ found in the literature. (See Goldfeld and Sichel (1990) for a survey.) I should add that for the three variable VAR discussed here, omitting money, the results simply are not sensitive to differences in these parameters.

The parameters of the policy rule, ξ, R_y and R_π, are within the range estimated for the US economy for the post-1979 period by McCallum and Nelson (1999); however, the standard errors are much larger than theirs, to reflect the greater uncertainty I have regarding how well this (or any) assumed policy rule accurately describes the actual policy behavior of the Federal Reserve.

It remains to set the inflation persistence parameter α. Any given value for α suggests that opportunities for price adjustment are presented to the typical firm every $\eta = 1/(1-\alpha)$ quarters. Estimates for the length of time the typical firm takes to adjust their prices varies widely in the literature. On the industry level, most overall estimates are from two to six quarters, but estimates vary widely among industries (see Taylor 1999 for a survey). For the economy as a whole, Roberts (1995) derived an estimate for $(1-\alpha)(1-\alpha\beta)/\alpha$ of about 0.08, corresponding to a value of α of about 0.75. Gali and Gertler's (1999) later study derives a number of estimates of α (their θ) that average about 0.8; as this is of the same order of magnitude as Roberts' estimate, this is the estimate I picked for the prior mean, implying average price fixity of about 5 quarters. However, Gali and Gertler's results for various subsamples suggest a higher α (around 0.85) for the peiod since 1980 than for before 1980 (for which their estimates of α are nearer Roberts' estimate of 0.75, derived with a slightly earlier data set). The prior standard deviation for α I picked permits a range for the parameter from 0.75 to 0.85.

The baseline values for the structural parameters imply a prior mean for **F** of

$$\begin{bmatrix} 1.028 & -2.582 & 0.874 \\ 0.00904 & 0.524 & 0.117 \\ 0.00757 & -0.403 & 1.058 \end{bmatrix}$$

The second row, corresponding to the inflation equation, suggests that expected inflation is as a first approximation best modeled as a function of the output gap and current inflation, *pace* Stock and Watson (1999), and that inflation itself should be modeled as stationary. The third row suggests that the real interest rate $(1+R_t)-\pi_t$ should be modeled as highly persistent but stationary, while the nominal interest rate $(1+R_t)$ is closely tied to the inflation rate (as it should be in any reasonable model).

Note that the prior implies that the behavior of output, inflation and the interest rate is well described by a VAR(1) process, with the values of longer lags assumed a priori to be zero. For inflation in particular, the Calvo-Rotemberg Phillips curve (F4') implies expected inflation at time t+1 to be a function only of inflation and the output gap/marginal cost at time t. However, with the prior standard deviations for the structural parameters specified here, the imposition of zero values for longer lags is loose enough to permit a role for longer lags of inflation to enter the process, allowing for a backward-looking element to that process.

I wish to compare the performance of the VARs using the prior derived from the New Keynesian DSGE model (hereafter the "New Keynesian BVAR") to a Bayesian VAR using a standard atheoretical prior. One popular such prior is the so-called "Minnesota" prior described by Doan, Litterman and Sims (1984). This time, let **X** be the $T\times 19$ matrix consisting of observations on a $T\times 1$ vector of ones and vectors

$\mathbf{y}_{t-1},\ldots \mathbf{y}_{t-6}$, where $\mathbf{y}_t$ is the vector of observations on the levels (not the deviations from trend) of the three variables, and let y_i be the $T\times 1$ vector of observation on the level of variable *i*. For the *i*th equation in the VAR, the Minnesota prior imposes a prior mean of one on the first autoregressive lag, and a prior mean of zero on all further autoregressive and all non-autoregressive lags; the implication is that each variable in the VAR is best described as a first approximation as a random walk. The prior covariance matrix is a diagonal matrix, with the prior variance for the *k*th coefficient of variable *j* in equation *i* being $\sigma_{ijk}^2 = t_1^2 / k^2$ for *i*=*j* and $\sigma_{ijk}^2 = (t_1^2 t_2^2 / k^2)\times(\sigma_j^2 / \sigma_i^2)$ for *i*≠*j*, where t_1 is an overall tightness parameter, $t_2 < 1$ is a parameter determining how strictly imposed is the prior restriction of zero values for non-autoregressive terms, and s_i^2 is the estimated variance of the residuals from a sixth-order AR model of variable *i*, s_j^2 / s_i^2 serving as a scaling factor accounting for the differences in magnitude between variables *i* and *j*. I used values of 0.2 for both t_1 and t_2.

Doan, Litterman and Sims (1984) suggest forecasts can be further improved by including prior information on the sums of coefficients. Such information is usually included by adding dummy observations to the data set, before mixed estimation. For example, in a three variable VAR with four lags of output, employment and interest rates, I add to **X** the dummy observations

$$t_3\begin{bmatrix} 0 & y_0 & 0 & 0 & y_0 & 0 & 0 & y_0 & 0 & 0 & y_0 & 0 & 0 \\ 0 & 0 & \pi_0 & 0 & 0 & \pi_0 & 0 & 0 & \pi_0 & 0 & 0 & \pi_0 & 0 \\ 0 & 0 & 0 & R_0 & 0 & 0 & R_0 & 0 & 0 & R_0 & 0 & 0 & R_0 \end{bmatrix}.$$

When estimating the output equation, the corresponding dummy observations on current output are $t_3[y_0 \quad 0 \quad 0]'$; similarly, the corresponding dummy observations on inflation for the inflation equation are $t_3[0 \quad \pi_0 \quad 0]'$, and those for the interest rate equation are $t_3[0 \quad 0 \quad R_0]'$. These dummy observations imply that the autoregressive terms in each equation sum to one (i.e. that each has a unit root), and that the terms for lags of all other variables sum to zero (i.e. the changes in each variable, assumed to be permanent, have no permanent effect on any other others).

Even with the sums of coefficients restrictions imposed, the disadvantage of the Minnesota prior is that it does not take into account cointegrating relationships between variables, as in this case one would expect between (at least) inflation and the nominal interest rate. To compensate for this somewhat, Sims (1992) suggests adding to **X** the dummy observation (again assuming three variables and four lags)

$$t_4[1 \quad y_0 \quad \pi_0 \quad R_0 \quad y_0 \quad \pi_0 \quad R_0 \quad y_0 \quad \pi_0 \quad R_0 \quad y_0 \quad \pi_0 \quad R_0],$$

with corresponding dummy observation on output $t_4[y_0]'$, on inflation $t_4[\pi_0]'$, and on the interest rate $t_4[R_0]'$. This dummy observation implies a cointegrating relationship between the output gap, the inflation rate and the nominal interest rate. t_3 and t_4 are parameters setting the tightness by which the prior information given by the dummy observations is imposed. y_0, π_0, and R_0, the implied trend values of output, inflation and interest rates, are measured by the means of the *k* presample observations, where *k* is the number of lags.

Robertson and Tallman (1999), in a recent paper on forecasting with VARs, argue that a Bayesian VAR estimated imposing the Minnesota/Litterman prior with mixed estimation, and including the Doan-Litterman-Sims (1984) and Sims (1992) sums-of coefficients restrictions, do about as well in forecasting than Bayesian VARs estimated using more complex methods of imposing prior information (e.g. Sims and Zha 1998). Hence, it is this modified Litterman BVAR (as I shall call it from now on) that I used as the benchmark for the New Keynesian BVAR. For the modified Litterman BVAR's parameter values, I used the values used by Robertson and Tallman in their experiments, i.e. $t_1 = 0.2, t_2 = 0.2, t_3 = 5.0, t_4 = 5.0$.

Results

My data set consisted of quarterly US data on output (seasonally adjusted real GDP), inflation (quarterly rate of change in GDP chain-type price deflator), and a nominal interest rate, the Federal Funds rate, for the period 1959:I to 2002:II. The Federal Funds rate was used for the nominal interest rate, as it is this interest rate that the Federal Reserve uses in practice as its policy instrument; hence the interest rate does double duty as a measure of current monetary policy. (I did experiment with other short-term interest rates, such as the 3-month Treasury Bill rate, but it made little difference; this is not surprising, as "market" short-term interest rates tend to track the "administered" federal funds rate closely.)

The Bayesian VARs using the modified Litterman prior contained six lags of the log of each variable and a constant. The VARs using the New Keynesian prior derived from the New Keynesian DSGE model (hereafter the "New Keynesian BVAR") were estimated without a constant, using six lags of detrended output, inflation and interest rates. Detrended output was measured as the deviation of log output from a linear time trend estimated by OLS over the period of the data set used to estimate the VAR, while detrended inflation and interest rates were measured as the deviation of inflation and interest rates from their sample mean over the period of the data set. Such a VAR can be used to generate out-of-sample forecasts of the deviations of output, inflation and interest rates from trend. To get forecasts of the level inflation and interest rates, I simply added the sample mean to the forecast detrended value of each variable; to get forecasts of level output, I added the forecast deviation of output from trend to the forecast level of trend output implied by the estimated trend.

Tables 1.2a through 1.2c presents performance statistics, the Theil-U, for the two competing models (modified Litterman and New Keynesian). The statistics are calculated as follows. Each model is estimated over the first 100 periods (1959:I to 1983:IV), and forecasts are calculated for one quarter, two quarters, four quarters (one year), eight quarters (two years) and twelve quarters (three years) ahead of 1983:IV. The actual observation for 1984:I is then added to the data set, forecasts are calculated for one, two, four, eight and twelve quarters ahead of 1984:I, and so forth; in all 63 sets of one, two, four, eight and twelve step ahead forecasts were generated (the last VAR being estimated over the 1959:I to 1999:II). The Theil-U is given by the ratio of the forecast mean squared error from the model in question to that from a naïve forecasting model (random walk with drift); a value above one indicates that the model forecasts are worse than those from a naïve random-walk model. Lower Theil-U's indicate better forecasting performance.

Along with the Theil-U's, I also calculated the test statistics for equality of loss-differentials for each model, to both the naïve model and each other (corrected for serial correlation in the loss differentials, i.e. differences in squared forecast errors between each model), as suggested by Diebold and Mariano (1995) as a measure of improvement or worsening of forecast performance. Under the null that forecast performance is equal, this statistic should follow a *t*-distribution; hence large Diebold-Mariano (DM) statistics suggest rejection of the hypothesis that forecast performance is equal.

Table 1.2a gives the forecast performance statistics for the entire period (1984:I-1986:IV through 1999:III-2002:II). For the nominal side of the economy, the New Keynesian prior does surprisingly well. For inflation and interest rates, the New Keynesian BVAR easily beats the random walk model, and for inflation, the New Keynesian BVAR consistently beats the modified Litterman BVAR, with the largest DM statistics being those for one year to two years ahead. For the nominal interest rate, the New Keynesian prior is also competitive with the modified Litterman BVAR, beating it at horizons of a year or more. While none of the DM statistics are significant, this suggests that the New Keynesian BVAR is at least competitive with the modified Litterman BVAR for forecasting nominal variables.

There, however, is where the good news ends. The forecasts generated by the New Keynesian BVAR for output are consistently worse than those from a naïve model, let alone those from the modified Litterman BVAR. (The difference is statistically significant at near horizons, and substantial at all.) Some comparison of the Litterman and New Keynesian prior is in order here. The Litterman prior constrains the inflation and interest terms in the output equation to zero fairly tightly. The sums of coefficients dummy observations, moreover, constrain the long-run effect of level changes in inflation to zero; this can be interpreted as implying that if anything, it is changes in inflation, *pace* Fuhrer and Moore, that affect output, not necessarily the level of inflation as such. The

New Keynesian prior, on the other hand, allows the inflation and interest terms to have coefficients that have values that are individually, and sum to, far from zero.

The prior mean values for the first lags (and, hence, the details of how the New Keynesian pricing model predicts inflation will affect output) are probably less important here for the results than the prior covariance matrix from the New Keynesian BVAR, which puts much looser restrictions on the values of the inflation and interest terms. In effect, the New Keynesian prior assumes much less information on just what those values should be. As a result, these terms are much more determined by the data, when estimating the BVAR, than would be the case if a Litterman prior were used. Given that the Litterman prior (which holds that "(level) inflation doesn't matter" for output, and restricts its effect to zero) does better than the New Keynesian prior (which holds that "inflation matters", and allows its estimated effect to be large) in predicting output, a reasonable conclusion to draw is that "(level) inflation doesn't matter," or at any rate its effect is exaggerated by the Calvo-Rotemberg model.

To check the stability of the results, I also calculated Theil U's and DM statistics for two sub-periods, 1984:I-1986:IV through 1991:IV-1994:III (32 sets of forecasts) and 1992:I-1994:IV through 1999:III-2002:II (31 sets of forecasts). For the earlier period, the advantage of the New Keynesian BVAR over the modified Litterman BVAR for inflation is even greater; the Theil-U's for the New Keynesian BVAR are significantly below one at any reasonable confidence level, and forecasts of inflation a year or more ahead are significantly better than those from the modified Litterman BVAR. However, the New Keynesian BVAR is still significantly worse at forecasting inflation than a naïve model at shorter horizons (and certainly than the Litterman model).

For the later period, the picture is somewhat different. The Internet boom and inflation rates far below (postwar) historical levels would be bound to hamper forecasting using models estimated by necessity with data from earlier periods. For inflation, the 1992:I-1994:IV through 1999:III-2002:II period is not a banner period for either model, with Theil-U's well above one for periods of a year or more ahead. However, the New Keynesian BVAR fares relatively worse, actually being beaten by the Litterman model for forecasts of inflation a year or more ahead. This is not surprising, as the New Keynesian BVAR (as well as the detrending method used to estimate it) assume inflation is stationary, while the Litterman prior assumes a unit root in inflation. Given that, the Litterman prior can be expected to fare better in the face of an apparent trend shift in inflation. For output, again, neither model fares well, with Theil U's well above one; however, the New Keynesian model's relative disadvantage remains (though the DM statistics no longer indicate significance).

A few general conclusions can be drawn. The good news (for those favoring New Keynesian models of inflation) is that the New Keynesian BVAR does surprisingly well at forecasting inflation. Generally speaking, it can be trusted to be at least competitive with a leading atheoretical forecasting model (barring trend shifts in inflation). The bad news is that the forecasts for output generated by the New Keynesian BVAR cannot even beat a random walk model, let alone more sophisticated atheoretical models. This result is robust to all the specification changes I have tried, including, but not limited to, large changes in the prior means for the "supply side" structural parameters (β, δ, ϕ, μ, and ρ) and increases in their prior standard deviations (particularly for ρ), to loosen the prior for the output equation. Starkly put, a New Keynesian prior using reasonable values for the prior means of the structural parameters and standard deviations, and generating reasonable forecasts for nominal variables, cannot generate adequate forecasts for output. The implication is that a New Keynesian prior specified so as to generate good inflation forecasts implies a relationship between output and inflation that is inconsistent with the data.

To get some measure of the "overall fit" of the New Keynesian BVAR, I followed Ingram and Whiteman (themselves following Doan, Litterman and Sims (1984)) by calculating the log determinants of the forecast covariance matrices for the naïve model, modified Litterman BVAR and "baseline" New Keynesian BVAR at each horizon. These are reported in Table 1.3, along with the log determinants of the forecast covariance matrices for New Keynesian BVARs with lesser (α=0.75) and greater (α=0.85) price stickiness.

In general, the New Keynesian BVARs, while certainly beating naïve models, are inferior to the modified Litterman BVARs in overall forecasting performance, except at very long forecast horizons (over which the advantage of the Litterman over the New Keynesian BVAR, or any model, tends to vanish). Among the New Keynesian BVARs, those with higher α's tend to do better at shorter horizons (over which prices are more likely to be sticky), while those with lower α's do better at farther horizons (where price stickiness is likely to be less important).

More generally, the New Keynesian BVAR tends to catch up with the modified Litterman BVAR as the forecast horizon increases, a tendency robust to sample period. (Looking back at the DM statistics confirms that for each variable the advantage of the New Keynesian BVAR over the modified Litterman BVAR increases with forecast horizon, for inflation and interest rates, and its relative disadvantage decreases.) My own conjecture is that the New Keynesian BVAR does a better job of capturing long-run tendencies in the data (say, between inflation and interest rates) than does a Litterman BVAR, even supplemented with Sims-type dummy observations to allow for cointegration. However, this is clearly insufficient to allow the New Keynesian BVAR's overall forecasting performance to be up to the standard of the competition.

Conclusions

The results of the forecasting experiments here suggest, at least at first sight, that a "reasonable" New Keynesian model for predicting inflation is incompatible with a reasonable model for predicting output. This finding parallels that of recent research on New Keynesian models of the output-inflation relationship such Kurmann (2001), who found that a New Keynesian model of the Calvo-Rotemberg type that generated reasonable inflation dynamics was incompatible with reasonable dynamics for real marginal cost, and, by extension, output and employment.

The Bayesian (or quasi-Bayesian) techniques used here to impose loosely the restrictions on output and inflation dynamics implied by the model allow much more room for misspecification than did Kurmann's (classical) full-information techniques. The Ingram-Whiteman technique imposes only inexactly the restriction on output dynamics implied by a DSGE model with Calvo-Rotemberg type New Keynesian features. However, even these loose restrictions clearly degrade forecasts of output, to the point where even a naïve random walk does better; this result is robust over a variety of forecast periods. More generally, any theoretical model that suggests a large amount of information content in inflation (and nominal variables in general) for forecasting output seems to do poorly compared to models that do not (like the modified Litterman prior, which assumes a priori no such role). Given that, the relevance of Calvo-Rotemberg type New Keynesian models for modeling the output-inflation relationship is put into question by these results.

How disastrous one finds all this, of course, will depend on one's loss function. For those more interested in accurate forecasts than proving one theoretical model or other to be "true" in some sense, the New Keynesian BVAR does generate forecasts competitive, and, provided inflation can be reasonably assumed to be stationary, superior to those from competing atheoretical models. (Altering the BVAR to allow for trend shifts in inflation, as apparently occurred in the 1990's, should be straightforward; alternatively, one could simply correct for secular bias in the forecasts generated by the model in formulating actual forecasts.) If one is interested in forecasting output as well, a BVAR for forecasting could very easily be estimated using a New Keynesian prior for the inflation equation, and an atheoretical prior along the lines of the modified Litterman BVAR for the output equation. An obvious objection on the grounds of logical consistency could be defended along these lines. The New Keynesian prior, in a VAR model with more than one lag, does allow for a non-zero term for long inflation lags, permitting a backward-looking element to the inflation process along the lines of Fuhrer and Moore (1995). The sums-of-coefficients restrictions on the inflation terms in the Doan-Litterman-Sims dummy observations, as well, do allow for changes in inflation, though not levels, to affect output. The hybrid model could then be defended as an attempt to convey the idea that what effects inflation has on output occur along the lines of a Fuhrer-Moore sticky-inflation model.

In a sense, what this amounts to is accepting the New Keynesian model as the best model of inflation economics has to offer, and accepting that one should try to forecast outpout without pretending to have too much theory. What is admittedly not at all obvious from this exercise, however, is that the failure of the New Keynesian DSGE model necessarily has to do with the New Keynesian elements, as opposed to other elements of the DSGE model that make for rather poor dynamics and could easily be corrected. For instance, adjustment costs in capital or investment could improve the dynamics of the model. It is also possible that my prior information is bad; Gali and Gertler's estimates of price stickiness are on the high end of those estimated in the literature, and possibly the lower estimates suggested by survey data (e.g. Blinder 1994) would be appropriate. Examining which modifications of the model would get the New Keynesian DSGE prior to improve performance on the real side (to see, as it were, whether it's a prior problem or a specification problem) would be an interesting exercise, but if one is going to go to that trouble, one might as well take the plunge and try to estimate a Bayesian New Keynesian DSGE model, and see what chances it has of being as true a model as the BVAR. It is that approach, checking what modifications the DSGE model needs to "fit the data," that I will be taking in the next chapter.

Table 1.1: Prior means and standard deviations for structural parameters of New Keynesian DSGE model

Parameter	Mean	S.D.
β	0.988	0.001
δ	0.016	0.004
ϕ	0.72	0.025
ρ	0.97	0.015
γ	1.004	0.0005
μ	1.4	0.1
R_0	0.016	0.004
$(M/PY)_0$	0.40	0.032
χ	0.3	0.1
ξ	0.8	0.2
R_y	0.2	0.2
R_π	1.5	0.5
α	0.8	0.025

Table 1.2a

Theil U and Diebold Mariano-statistics for forecasting models, 1984:I-1986:IV through 1999:III-2002:II

		Modified Litterman		New Keynesian		
Variable	Forecast horizon	Theil U	DM stat., H_0: MSE(Litt.)= MSE(random walk)	Theil U	DM stat., H_0: MSE(NK)= MSE(Random walk)	DM stat., H0: MSE(Litt.)=MSE(NK)
Output	1 qtr.	0.83	1.32	1.24	-1.87	-2.77
	2 qtrs.	0.8	0.93	1.31	-1.35	-2.18
	4 qtrs.	0.8	0.59	1.32	-0.9	-1.63
	8 qtrs.	0.8	0.42	1.16	-0.34	-0.77
	12 qtrs.	0.83	0.33	1.13	-0.37	-0.59
GDP deflator	1 qtr.	0.86	3.32	0.82	2.28	0.97
	2 qtrs.	0.81	3.56	0.73	2.81	1.47
	4 qtrs.	0.94	0.66	0.83	1.35	1.73
	8 qtrs.	1.01	-0.05	0.82	0.72	1.51
	12 qtrs.	1.02	-0.07	0.95	0.17	0.43
Federal funds	1 qtr.	0.79	2.81	0.86	0.91	-0.75
	2 qtrs.	0.77	1.95	0.81	0.85	-0.32
	4 qtrs.	0.65	2.31	0.61	1.79	0.55
	8 qtrs.	0.56	2.02	0.49	1.75	0.97
	12 qtrs.	0.52	2.38	0.44	2.08	1.07

Theil U statistics for each model are calculated as the ratio of the forecast mean squared error to the mean squared error of a random-walk with drift forecast. DM statistics are the t-statistics generated by an asymptotic test of a null hypothesis of zero loss-differential consistent in the presence of (k-1)th order serial correlation in the loss differentials; for details, see Diebold and Mariano (1995).

Table 1.2b
Theil U and Diebold Mariano-statistics for forecasting models, 1984:I-1986:IV through 1991:IV-1994:III

		Modified Litterman		New Keynesian		
Variable	Forecast horizon	Theil U	DM stat., H_0: MSE(Litt.)= MSE(random walk)	Theil U	DM stat., H_0: MSE(NK)= MSE(Random walk)	DM stat., H0: MSE(Litt.)=MSE(NK)
Output	1 qtr.	0.68	1.87	1.14	-1.23	-2.87
	2 qtrs.	0.66	1.28	1.15	-0.69	-2.08
	4 qtrs.	0.66	0.89	1.06	-0.21	-1.54
	8 qtrs.	0.63	0.73	0.85	0.47	-0.69
	12 qtrs.	0.63	0.67	0.84	1.23	-0.53
GDP deflator	1 qtr.	0.87	2.56	0.82	1.73	0.86
	2 qtrs.	0.8	4.41	0.68	3.32	1.89
	4 qtrs.	0.86	2	0.67	6.93	4.18
	8 qtrs.	0.84	Complex	0.57	9.8	2.18
	12 qtrs.	0.7	10.17	0.45	10.31	3.71
Federal funds	1 qtr.	0.8	2.25	0.86	0.76	-0.47
	2 qtrs.	0.77	1.66	0.8	0.7	-0.21
	4 qtrs.	0.63	2.24	0.58	1.66	0.57
	8 qtrs.	0.57	1.98	0.48	1.66	0.91
	12 qtrs.	0.51	3.11	0.42	2.45	1.18

Table 1.2c

Theil U and Diebold Mariano-statistics for forecasting models, 1992:I-1994:IV through 1999:III-2002:II

		Modified Litterman		New Keynesian		
Variable	Forecast horizon	Theil U	DM stat., H_0: MSE(Litt.)= MSE(random walk)	Theil U	DM stat., H_0: MSE(NK)= MSE(Random walk)	DM stat., H0: MSE(Litt.)=MSE(NK)
Output	1 qtr.	1.17	-1.85	1.43	-1.41	-0.94
	2 qtrs.	1.24	-2.36	1.81	-1.24	-0.94
	4 qtrs.	1.39	-2.6	2.36	-1.12	-0.83
	8 qtrs.	1.71	-2.5	2.75	-0.98	-0.52
	12 qtrs.	1.6	-2.76	2.19	-0.99	-0.45
GDP deflator	1 qtr.	0.84	2.14	0.82	1.57	0.45
	2 qtrs.	0.84	1.32	0.84	0.88	0.07
	4 qtrs.	1.19	-1.15	1.3	-1.41	-1.1
	8 qtrs.	1.47	-1.08	1.52	-1.14	-1.23
	12 qtrs.	1.94	-6.88	2.43	-16.32	-2.56
Federal funds	1 qtr.	0.74	2.42	0.88	0.9	-1.72
	2 qtrs.	0.75	1.52	0.82	1.05	-0.92
	4 qtrs.	0.7	1.22	0.7	1.19	-0.06
	8 qtrs.	0.51	1.33	0.46	1.33	1.63
	12 qtrs.	0.54	1.51	0.55	1.26	-0.13

Table 1.3: Multivariate forecasting performance statistics

Forecast period	Model	One qtr.		Two qtrs.		Four qtrs.		Eight qtrs.		Twelve qtrs.	
1984:I-1986:IV thru 1999:III-2002:II		ABS	REL	ABS	REL	ABS	REL	ABS	REL	ABS	REL
	Naïve	36.11		33.99		32.28		30.36		29.31	
	Mod. Litt.	36.62	8.51	34.56	9.54	32.93	10.71	30.96	9.98	29.95	10.65
	NK, α=0.8	36.28	2.86	34.28	4.83	32.57	4.83	30.74	6.36	29.84	8.93
	α=0.75	36.25	2.3	34.25	4.26	32.58	5.05	30.8	7.48	29.97	10.98
	α=0.85	36.3	3.18	34.3	5.06	32.55	4.4	30.7	5.43	29.79	8.01
1984:I-1986:IV thru 1991:IV-1994:III		ABS	REL	ABS	REL	ABS	REL	ABS	REL	ABS	REL
	Naïve	35.26		33.03		31.24		29.51		28.6	
	Mod. Litt.	35.88	10.5	33.74	11.76	32.17	15.49	30.44	15.4	30	23.24
	NK, α=0.8	35.5	4.06	33.54	8.35	32.08	14	30.53	17.07	30.21	26.83
	α=0.75	35.45	3.17	33.48	7.42	32.1	14.45	30.66	19.07	30.45	30.87
	α=0.85	35.53	4.45	33.55	8.55	32	12.82	30.42	15.14	30.06	24.27
1992:I-1994:IV thru 1999:III-2002:II		ABS	REL	ABS	REL	ABS	REL	ABS	REL	ABS	REL
	Naïve	37.74		35.85		34.34		32.7		31.07	
	Mod. Litt.	38.09	5.81	36.12	4.66	34.26	-1.27	32.68	-0.47	30.83	-4.12
	NK, α=0.8	37.69	-0.83	35.65	-3.25	33.6	-12.26	32.07	-10.57	30.38	-11.53
	α=0.75	37.68	-1.03	35.64	-3.38	33.58	-12.6	32.1	-10.13	30.37	-11.59
	α=0.85	37.68	-1.01	35.65	-3.23	33.62	-11.88	32.05	-10.87	30.34	-12.12

ABS denotes the absolute value of the log determinant of the forecast error covariance matrix. REL denotes the average percentage improvement in forecasting accuracy over the random walk with drift model, i.e. the difference in log determinants divided by 3 (the number of variables), divided by 2 (to convert variances to standard errors), multiplied by 100 (to convert to percent).

CHAPTER 2

LOSS-BASED EVALUATION OF A NEW KEYNESIAN DSGE MODEL

Introduction

Part of the reason that I only loosely imposed the restrictions on the data implied by my DSGE model, rather than just estimating a fully-specified New Keynesian DSGE model, was an a priori belief that "DSGE models (unaided) just don't fit the data." Certainly the restrictions on economic data imposed DSGE models, especially the older, simpler ones, are frequently rejected by classical hypothesis tests. If the null is rejected, however, a simple hypothesis test does not provide a superior alternative model to which the DSGE model can be compared, so it cannot tell us where to go next or what to correct.

Many recent studies (like Watson (1993) and Diebold *et al.* (1998)) of DSGE models assume the DSGE is certainly misspecified (i.e. the probability that the DSGE model is the "true" model is assumed to be zero), and in essence propose measures of just how far off the mark the DSGE models are from an atheoretical reference model (like an identified VAR). Watson measures misspecification by how much uncertainty must be introduced to the model to make it match sample moments (such as correlations and impulse responses); Diebold et al. measure misspecification using loss functions penalizing deviations of sample from model moments. However, such an approach assumes, in effect, that the model can itself teach us nothing that the sample can't about how (say) the impulse response of output and inflation to various shocks should look. Given the improvements in the dynamics of DSGE models over the earlier, simpler ones that failed classical hypothesis tests, it is no longer possible to make that assumption with confidence, unless one really believes that a given DSGE model cannot possibly be the "true" model of the economy, or at any rate a good description of a subset of time series.

A study by Schorfheide (2000) of a cash-in-advance model with partial adjustment costs uses just such a method. Schorfheide takes a middle-of-the-road approach, allowing for a non-zero probability that the DSGE model fits the data better than a given competitor, a probability that can be updated after looking at the data. Specifically, he compares a standard cash-in-advance model and a cash-in-advance model with portfolio adjustment costs, both including (permanent) technology shocks and (temporary) monetary shocks, to an identified VAR with disturbances disaggregated into permanent and transitory shocks *pace* Blanchard and Quah (1989). The specifications of the cash-in-advance model and the model with portfolio adjustment costs provide likelihood functions for the data which can be evaluated using a Kalman filter algorithm; combining the likelihood function with prior distributions for the structural parameters gives a posterior distribution for the parameters. Schorfheide calculates posterior probabilities that each model describes the data and overall posterior estimates of the population characteristics, using so-called Bayes factors (e.g. Kass and Raftery 1995). Unlike classical model selection criteria, which are generally functions of the maximum likelihood of a given model, the Bayes factor is a measure of the *average* likelihood of the model, given the prior density on the model parameters.

Given posterior probabilities, Schorfheide can construct a probabilistic representation of the data, a weighted mixture of the DSGE models and the identified VAR, as a benchmark to which he can compare the DSGE models. He can then calculate the posterior distributions of various moments of interest, such as correlations and impulse response functions. He goes on to assess the ability of the DSGE models to track US output and inflation time series, generate plausible correlation patterns between output and inflation, and describe the response of output growth to a monetary expansion; this ability is measured by the values of loss functions penalizing the deviations of the DSGE model's moments from those of the benchmark.

The use of a mixture of the DSGE models and the identified VAR in constructing the benchmark for comparison allows us in effect to learn from the models regarding how the data should look. In Schorfheide's actual application, however, the posterior probabilities that the cash-in-advance model and the model with portfolio adjustment costs were true turned out to be vanishingly small. This indicated that we could learn very little from these models about inflation and output dynamics in the actual US economy. Given that the New Keynesian model is compatible with a negative correlation of output growth and inflation of a magnitude more in line with the data, as King and Watson (1996) argue, and with a potentially much larger positive response of output to a monetary shock, a natural question to ask is whether adding New Keynesian features would not correct the discrepancy between model and posterior.

My purpose in this paper, then, is to evaluate and compare a simple New Keynesian monetary DSGE model, assessing the discrepancy between its predictions and a posterior distribution of population characteristics such as cross-correlations and impulse-response functions, similar to that done by Schorfheide (2000) for the portfolio adjustment cost model. A necessary first step in this process is to estimate a monetary DSGE model with a New Keynesian Phillips curve using US data for output and inflation, using Bayesian methods to impose prior information regarding the values of the structural parameters. In particular, I use a prior assumption for the measure of price stickiness in the economy consistent with that found in survey data such as Blinder (1994). In so doing I obtain an estimate of the degree of price-stickiness in the US economy, conditional on the model's being correctly specified. This is an interesting exercise in itself, as far as it goes, corresponding to the exercise conducted by Dejong et al. (2000a, 2000b) for the Greenwood et al. (1988) variable capital utilization model.

However, given that there is a non-zero chance that the New Keynesian DSGE model is misspecified, it is only a first step towards a full evaluation of the model; my real purpose is to evaluate just how well specified such a model, estimated with a prior using conventional values for the parameters, really is (and so how "plausible" such values are) and whether it can succeed where the portfolio adjustment cost model failed in teaching us something about the output-inflation relationship. In particular, I want to see if assuming prices to be somewhat sticky (and as a first approximation as sticky as limited information studies suggest) actually correct some of the problems that cash-in-advance models with flexible prices have in matching inflation dynamics and the effect of monetary shocks on output. To make my results comparable to Schorfheide's, I compare the New Keynesian DSGE model to a cash-in-advance model without price stickiness and to an identified Bayesian VAR. I estimate all three models to allow me to calculate posterior probabilities for each model, allowing me to construct a benchmark distribution for correlations and impulse response functions to which I can compare the correlations and impulse responses of the New Keynesian model, according to a variety of loss functions.

Among the things I find in what follows are that a New Keynesian DSGE model with sufficiently rich dynamics, contrary to Schorfheide's results, is competitive and in fact superior in fit to a Bayesian VAR in output and inflation, with the DSGE being the model with a much greater posterior probability than the VAR. I find, however, that the goodness of fit derives from modifications of the basic model—a more realistic monetary policy in terms of inflation rather than money growth as Schorfheide used, and well as costs of adjusting investment—that have little to do with the New Keynesian features of the model as such. The output and inflation data for the US since the mid 1960's, while certainly compatible with an amount of price stickiness found in survey data, turn out to not be especially informative themselves regarding the amount of price stickiness. This is at least partly because inflation shocks can account for very little of US output fluctuations, and so there is not likely to be much information about the impact of an inflation or money shock on output in the data. What the DSGE model does allow us to do is to add a good deal of precision regarding out estimation of the output-inflation relationship and the impact of a monetary/inflation shock on output, given our best prior information on price stickiness.

Loss-Function Evaluation of DSGE Models

Before I begin, I should describe the Schorfheide method of loss-function evaluation of DSGE models in more detail. (My discussion will be somewhat informal; for the finer details, see Schorfheide (2000).)

Before I do so, some shorthand is in order. Let M_1 denote model 1 (say, the flexible-price cash-credit model) and M_2 denote model 2 (say, the New Keynesian model). Let $\theta_{(i)}$ be the parameter vector for model M_2. Also, let y_t be the $n \times 1$ vector of the values of the n observables to be used in estimation at time t, and let $Y_T = [y_1,\ldots,y_T]'$ be the full $n \times$ T data set. The specification of each model M_i provides a likelihood function for Y_T, denoted hereafter as $p(Y_T \mid \theta_{(i)}, M_i)$; the product of this likelihood and the prior distribution for $\theta_{(i)}$, $p(\theta_{(i)} \mid M_i)$, is proportional to the posterior density for $\theta_{(i)}$, $p(\theta_{(i)} \mid Y_T, M_i) \propto p(Y_T \mid \theta_{(i)}, M_i)p(\theta_{(i)} \mid M_i)$. As there is a good chance that both DSGE models are misspecified, I also consider an identified VAR as a reference model, which I shall call model 0 or M_0; hence the parameter vector for the identified VAR will be labeled $\theta_{(0)}$, its likelihood function $p(Y_T \mid \theta_{(0)}, M_0)$, and its prior density $p(\theta_{(0)} \mid M_0)$.

There are three steps to the model evaluation:

Step 1

With the posterior densities $p(\theta_{(i)} \mid Y_T, M_i)$ in hand for each model, we can compute posterior distributions for each model's parameters $\theta_{(i)}$. For model 0, the identified VAR, this can be done analytically; for the DSGE models, however, the posterior densities will be analytically intractable, and so the posterior distributions for $\theta_{(1)}$ and $\theta_{(2)}$ will have to be derived numerically. The means by which the DSGE model's parameter distributions are computed will be discussed in Section 4.

Let $\pi_{0,0}$, $\pi_{1,0}$, and $\pi_{2,0}$ be the prior probabilities we place on models 0, 1, and 2 being the "correct" model, before estimating the models. By Bayes' rule, the posterior model probabilities $\pi_{0,T}$, $\pi_{1,T}$, and $\pi_{2,T}$ are

$$\pi_{i,T} \equiv p(M_i \mid Y_T) = \frac{\pi_{i,0} p(Y_T \mid M_i)}{\sum_{i=0}^{2} \pi_{i,0} p(Y_T \mid M_i)},$$

where $p(Y_T \mid M_i) = \int p(Y_T \mid \theta_{(i)}, M_i) p(\theta_{(i)} \mid M_i) d\theta_i$ is the marginal data density, or Bayes factor, given model M_i. Intuitively $p(Y_T \mid M_i)$, which is the average likelihood of model M_i given the prior distribution $p(\theta_{(i)} \mid M_i)$, can be thought of as the probability that an economy that was well described by model M_i could have produced time series Y_T; a relatively large/small $p(Y_T \mid M_i)$ will result in a relatively large/small posterior probability $\pi_{i,T}$ is the correct model.

Step 2

The population characteristics (correlations, impulse response functions and so forth) that we obtain from each model will, of course, be functions of their parameters $\theta_{(i)}$. Let ω be an $m \times 1$ vector of population characteristics of interest. With the posterior distributions of $\theta_{(i)}$ in hand, we can obtain posterior distributions of ω conditional on model M_i, for short $p(\omega \mid Y_T)$, for each model. The overall posterior distribution of φ, given the data Y_T, is a weighted average of the densities $p(\varphi \mid Y_T, M_i)$, namely $p(\varphi \mid Y_T) = \sum_{i=0}^{2} \pi_{i,T} p(\varphi \mid Y_T, M_i)$.

Step 3:

With the posterior distribution of ω in hand, we can now calculate the minimum expected loss $L_x(\omega, \hat{\omega})$ associated with each model in terms of the deviation of the "best estimate" of value of ω (which I will label $\hat{\omega}$) given by each model from the "actual" value ω.

(a) For each model M_i, find the value $\hat{\omega}_{x,i}$ for the population characteristic vector φ that minimizes the expected value of loss function $L_x(\omega, \hat{\omega})$, given that M_i is the true model. That is,

$$\omega_{x,i} = \arg\min_{\omega=\Omega} \int L_x(\omega, \hat{\omega}) p(\omega \mid Y_T, M_i) d\omega$$

(b) Then, the expected loss or risk $R_x(\hat{\omega}_{x,i} \mid Y_T)$ associated with the "optimal" estimate $\hat{\omega}_{x,i}$ for model M_i is the expected value of $L_x(\omega, \hat{\omega}_{x,i})$, given the overall posterior distribution $p(\varphi \mid Y_T)$ derived in step two. In other words,

$$R_x(\hat{\omega}_{x,i} \mid Y_T) = \int L_x(\omega, \omega_{x,i}) p(\omega \mid Y_T) d\omega$$

Following Schorfheide (2000), I use three loss functions. Of course, other loss functions are possible. The virtue shared of the three used here is that the $\hat{\omega}_x$'s they imply are reasonably easy to compute. (See Judge et al. (1985, chap. 4), for further discussion of loss functions commonly used in Bayesian applications.) The first, L_p (for "*L*-*p*-value"), is

$$L_p(\omega, \hat{\omega}) = I\{p(\hat{\omega} \mid Y_T) > p(\hat{\omega} \mid Y_T)\}$$

where $I\{x\}$ is an indicator function that equals one when condition x is true and zero otherwise. The expected L_p-loss provides a measure of how far $\hat{\omega}$ lies in the tail of the posterior distribution for ω.

The second loss function, L_{χ^2} (for "L-chi-square") is defined as follows. Let V_ω be the posterior covariance matrix of ω given density $p(\omega \mid Y_T)$. Define the function $C_{\chi^2}(\varphi \mid Y_T)$ as a weighted deviation of ω from its posterior mean $E(\omega \mid Y_T)$, namely

$$C_{\chi^2}(\omega \mid Y_T) = \left(\omega - E(\omega \mid Y_T)\right)' V_\omega^{-1}\left(\omega - E(\omega \mid Y_T)\right)$$

The loss function L_{χ^2} itself is

$$L_{\chi^2}(\omega, \hat{\omega}) = I\left\{C_{\chi^2}(\omega \mid Y_T) > C_{\chi^2}(\hat{\omega} \mid Y_T)\right\}$$

Expected L_{χ^2} -loss amounts to a measure of the weighted discrepancy of $\hat{\varphi}$ from the posterior mean $E(\omega \mid Y_T)$. The third loss function L_q ("*L*-quadratic"), is

$$L_q(\varphi, \hat{\varphi}) = (\varphi - \hat{\varphi})'W(\varphi - \hat{\varphi})$$

where W is a subjective weighting matrix. One virtue of L_q is that it has a straightforward solution for the expected risk, viz.

$R_q(\hat{\omega}_{q,i} \mid Y_T) = (\hat{\omega}_{q,i} - E(\omega \mid Y_T))'W(\hat{\omega}_{q,i} - E(\omega \mid Y_T))$.

It can be shown that the estimate $\hat{\omega}_{p,i}$ that minimizes expected L_p -loss given model M_i is the posterior mode of $p(\omega \mid Y_T, M_i)$, while expected L_{χ^2} - and L_q -loss are minimized by the posterior mean estimate $\hat{\varphi}_{q,i} = E(\varphi \mid Y_T, M_i)$.

Flexible Prices Versus Sticky Prices: Two Monetary DSGE Models

To estimate and compare the monetary DSGE models, it is necessary to describe them in more detail.

Once upon a time there were two economies. Both were inhabited by a large number of individuals (normalized to one) who supplied labor and rented out capital (in perfectly competitive resource markets) to a large number (also normalized to one) of imperfectly competitive firms. Each firm used labor and capital to produce a particular differentiated good. All of these goods were then aggregated to produce a single good which individuals could use either as a consumer good or as a capital good to be added to their existing stocks of capital. The amount of the composite good y that could be produced, given amounts of the individual goods y_i produced by each firm i, was

$$y = \left(\int_0^1 y_i^{(\varepsilon-1)/\varepsilon} di \right)^{\varepsilon/(\varepsilon-1)},$$

where $\varepsilon > 1$. It is a standard result (Dixit and Stiglitz 1977) that minimizing the cost of producing y units of the composite good results in a quantity demanded y_i of the good of firm i of y_i

$$y_i = (P_i / P)^{-\varepsilon} y,$$

where P is the aggregate price level, given by

$$P = \left(\int_0^1 P_i^{1-\varepsilon} di \right)^{1/(1-\varepsilon)}.$$

Consumer's problem

In both economies, the preferences of each individual in terms of consumption and labor could be described in terms of the lifetime utility function

$$\sum_{t=0}^{\infty} \beta^t (\chi \ln c_{1t} + (1-\chi) \ln c_{2t} - \xi h_t),$$

where h is labor supply and c is consumption (in terms of the composite good). As in all good model economies, each consumer attempted to maximize the expected value of their lifetime utility function in each period. To ensure that there was a demand for money (following Cooley and Hansen 1989), the custom prevailed in both countries that a subset of the consumer goods (labeled c_1) could only be bought with cash, while the others (c_2) could be bought either with cash or on credit. The weight of cash goods in the preference function is given by χ.

At the beginning of each period t each individual carried forward bonds with a nominal value of b_{t-1} dollars, money balances with a nominal value of m_{t-1} dollars, and k_{t-1} units of the (composite) capital good from the previous period $t-1$. He (for so I shall call the individual for short) would sell h_t units of labor in the competitive labor market at the real wage rate w_t per hour, and would rent out his capital stock to firms in a competitive capital market, getting back $r_t k_{t-1}$ units, r_t being the market rental rate. Finally, the individual received a helicopter drop of x_t dollars on top of the receipts from maturing bonds of b_{t-1} dollars, leaving him with total cash balances of $b_{t-1}+m_{t-1}+x_t$ dollars. He could then purchase consumer goods (either cash or credit), investment goods i_t, bonds and money balances subject to the budget constraint

$$c_{1t}+c_{2t}+i_t+\Phi\left(\frac{i_t}{i_{t-1}}\right)+\frac{1}{1+R_t}\frac{b_t}{P_t}-\frac{b_{t-1}}{P_t}+\frac{m_t}{P_t}-\frac{m_{t-1}}{P_t}-\frac{x_t}{P_t}-w_t h_t-r_t k_{t-1}\leq 0,$$

where R_t is the prevailing nominal interest rate. Following Smets and Wouters (2002), I add a cost of adjustment of investment, with the cost function Φ equaling zero in the steady state and having a constant second derivative, so that $\Phi(\gamma)=0, \dfrac{d^2\Phi(x)}{(dx)^2}=\varphi>0$.

DSGE models of all kinds are notorious for their poor dynamics; adding these adjustment costs improves the dynamics considerably in the DSGE models used here, and hence, as will become evident below, their fit to the data (and posterior probabilities).

We have related that cash goods c_1 could only be bought with real money balances. The custom also prevailed bonds could only be bought with cash; so, individuals were also subject to the cash-in-advance constraint

$$c_{1t}+\frac{1}{1+R_t}\frac{b_t}{P_t}-\frac{b_{t-1}}{P_t}-\frac{m_{t-1}}{P_t}-\frac{x_t}{P_t}\leq 0$$

in both economies, ensuring a demand for money.

Firm's problem

Each firm i purchased labor and capital in perfectly competitive labor markets to produces good y_i. The amount of good y_i that firm i could produce in period t with labor input h_{it} and capital input k_{it} was given by the production function

$$y_{it}=k_{it}^{1-\varphi}\left(z_t(h_{it}-h_0)\right)^{\varphi},$$

h_0 being a fixed cost of production in terms of labor (which serves as a barrier to entry) and z_t was the level of technology. z_t evolved according to the random walk process

$$\ln z_t-\ln z_{t-1}=\ln\gamma+\varepsilon_{zt},$$

where the technology shock ε_{zt} is white noise, so that technology grew on average at a rate of ? per period. The cost minimization problem of firm i was

$$\min tc_{it}=w_t h_{it}+r_t k_{it} \text{ such that } y_{it}=k_{it}^{1-\varphi}\left(z_t(h_{it}-h_0)\right)^{\varphi}.$$

The solution of that problem can be expressed (Yun 1996) in the form

$tc_{it} = w_t h_0 + mc_t y_{it}$, where mc is the real marginal cost of producing an additional unit of y_i. The instantaneous real profit of firm i at time t was then

$$\begin{aligned}\Pi_{it} &= (P_{it}/P_t)y_{it} - tc_{it} = (P_{it}/P_t)y_{it} - w_t h_0 - mc_t y_{it} \\ &= (P_{it}/P_t)^{1-\varepsilon} y_t - w_t h_0 - mc_t (P_{it}/P_t)^{-\varepsilon} y_t\end{aligned}.$$

The difference between the two economies lay in the flexibility with which firms could set their prices. In one economy, the "flexible-price" economy, firms could adjust their prices as often as they pleased, and did so as to maximize profit. As all firms were identical, it always turned out that $P_{it} = P_t$ and $mc = (\varepsilon - 1)/\varepsilon$ in each period. Hence, in the flexible-price economy, real marginal cost was always a constant.

However, in the other economy, the "New Keynesian" economy, prices were not perfectly flexible, and the real marginal cost of output could vary over the business cycle. In the "New Keynesian" economy, following a queer custom first described in Calvo (1983), a fraction $\alpha > 0$ of firms, chosen by lot, were not allowed to freely adjust their price, but only to raise it by the average rate of inflation π_0. This is a standard assumption in the New Keynesian literature, meant to convey the idea that firms might be locked into long-lasting price contracts. Its chief drawback is not allowing the degree of price flexibility to change over time, particularly in response to long-lasting changes in the average rate of inflation (ruled out by assumption, as the long-run rate of inflation is assumed given). The remaining fraction $1-\alpha$ of firms were allowed to adjust their prices, and did so so as to maximize the expected present value of their profits, allowing for the chance that the chance was α^k that the firm might not be able to change that price for k periods in the future, no matter what the price level was doing.

Now as the lucky $1-\alpha$ firms in the group that could adjust their prices were identical to each other, as were the *a* unlucky firms in the group that could not, the price charged by each firm in each group was the same, so that the price level at time *t* was given by

$$P_t^{1-\varepsilon} = (1-\alpha)\left(P_t^*\right)^{1-\varepsilon} + \alpha P_{t-1}^{1-\varepsilon},$$

where P_t^* was the price charged by the firms allowed to adjust their prices in period *t*.

Monetary policy

We do not know exactly what was the method by which the helicopters determined x_t, but several methods come to mind. The choice of methods calls for wisdom. Schorfheide (2000) assumed a rule of money growth without feedback, so that the rate of money growth followed an AR(1) process such that

$$\ln \mu_t = (1-\rho_m)\ln \mu_0 + \rho_m \ln \mu_{t-1} + \varepsilon_{mt}, \varepsilon_{mt} \sim N(0,\sigma_m^2), 0 \le \rho_m < 1.$$

While such a policy is very simple, a money growth model without feedback does not at all resemble US monetary policy during the post-1959 period (with the possible exception of the "monetarist" experiment from 1979 to 1982). Among the side effects of assuming such a policy is that, with no feedback to account for supply shocks (in either direction), is an effect on the price level from a permanent supply shock that is far too large, which Schorfheide (2000) observed in his experiments with monetary DSGE models, as well as a spike in the price level from a monetary shock. These are generally not observed in VARs, however monetary and supply shocks are measured. Also, in practice, monetary policy accommodates supply shocks to at least some degree.

In many papers in the New Keynesian literature, monetary policy would be specified in terms of a Taylor rule, with the nominal interest rate being written in terms of the output gap and inflation. However, given that I will be estimating my models with only output and inflation data, I cannot expect the data to be all that informative regarding the Taylor rule. For that reason, and the fact that all the real effects of money in New Keynesian DSGE models come from its effects on inflation, I will specify monetary policy in terms of the resulting behavior of inflation. Inspection of a VAR process for output and inflation suggests that logged inflation, ln π_t, is well described by an AR(1) process, so that

$$\ln \pi_t = (1-\rho)\pi_0 + \rho \pi_{t-1} + \varepsilon_{\pi t}$$

where $\ln \pi_0 = \ln \mu_0 - \ln \gamma$ and ε_π is an inflation term with variance σ_π^2 that is correlated with the technology shock ε_z. Specifically, $\varepsilon_\pi = \pi_z \varepsilon_z + \varepsilon_m$, where $\pi_z < 0$ is the contemporaneous response of inflation to a technology shock and ε_m is the portion of inflation shocks uncorrelated with technology, which I assume to be the monetary shock. With the correlation of technology and inflation shocks $corr(z,\pi)$ in hand it is straightforward to calculate $\pi_z = corr(\varepsilon_z, \varepsilon_\pi)\dfrac{\sigma_\pi}{\sigma_z}$.

Empirical Analysis

Given that there are only two shocks in my DSGE models (technology/"aggregate supply" and money/inflation/"aggregate demand"), to ensure non-degenerate distributions for the data I can only estimate my models for two time series at a time. As my interest is primarily in the relationship between output and inflation, I estimated each of the competing models with data on output (real GDP deflated by population) and the price level (chain-type GDP price deflator), logged and differenced to obtain series on (per capita) economic growth and inflation. The data is quarterly data for the United States from 1964:I to 2003:I.

As a reference model, I will be using a VAR in output growth and inflation:

$$\mathbf{B}(L)y_t = \mathbf{u}_t = \mathbf{A}\mathbf{e}_t$$

where $y_t = [d\ln GDP_t, \pi_t]'$, $\mathbf{B}(L)$ is a lag polynomial and $\mathbf{e}_t = [\varepsilon_{zt}, \varepsilon_{mt}]'$ is a 2x1 vector of a technology shock and a monetary shock, related to the prediction errors $\mathbf{u}_t$ for output and inflation by the 2x2 matrix **A**. To identify the four elements of A requires one restriction above and beyond those given by the correlations of the prediction errors from the VAR. I do so here by imposing the restriction that the long-run effect on output of the temporary monetary shocks is zero, as implied by both the cash-in-advance and New Keynesian models. So, following Blanchard and Quah (1989), I disaggregate the prediction error into a permanent element, with a long-run effect on output, and a temporary element, which does not, and assume that the permanent shock measures technology and the temporary shock money. (I do this a bit reluctantly, as the temporary shock could easily be measuring non-monetary types of demand shocks, such as government spending shocks.

Key to my analysis is calculating posterior model probabilities, which requires that I estimate the VAR with a proper Bayesian prior. There are several ways of doing this; here, I supplemented my VAR models of output and inflation with a version of the Normal-Wishart variant on the Minnesota prior suggested by Kadiyala and Karllson (1997). Kadiyala and Karllson (1997) assumed a priori that the prior covariance matrix of the disturbances from the VAR equations followed an inverse Wishart distribution with low degrees of freedom, while the coefficients of the VAR conditional on the covariance matrix of the disturbances followed a multivariate normal distribution centered on the random-walk process assumed by Litterman (i.e. first own lag coefficient equals one, all other lag coefficients equal zero). They set the parameters of their prior so that prior variances of the lag coefficients in the BVAR satisfied:

π_1 / k for the *k*th own lag;

$\pi_1 s_i^2 / k s_j^2$ for the *k*th lag of variable *j*; and

$\pi_3 s_i^2$ for the constant term in equation *i*.

Following Kadiyala and Karllson π_1 was set to 0.012, suggesting a range of the own lags from 0.8 to 1.2. For π_3, a value of 0.3 was used. The inverse Wishart distribution used four degrees of freedom. The prior means of the BVAR coefficients differed slightly from those of Kadiyala and Karllson's original Kadiyala-Karllson prior. The prior mean of the constant in the output equation was set at 0.005 (not zero). The prior for the first lagged term in the BVAR output equation was set to zero, not one, to imply a random walk in output (not output growth). The last modification was to discard, when drawing from the BVAR, any draw that implied non-stationarity in the system.

The prior distributions for the DSGE model parameters were chosen to fit the domain of the structural parameters. As calculated posterior probabilities can be somewhat sensitive to prior distributions, I imposed highly informative priors only on parameters for which I was fairly sure there was little if any information in the data; if I judged the data were highly informative, I used as loose a prior as I could while keeping the prior proper (which is necessary for calculating posterior probabilities). A couple of parameters were fixed a priori. When the weight of cash goods χ in the preference function was allowed to float between zero and one, the data strongly favored values near zero; I set it a priori to 0, thus shutting down inflation tax effects which the data suggest are not very important (Very similar results would have obtained for small but positive values for χ.) The data also strongly favor very low values for the markup μ, tending towards one (perfect competition); as there is much evidence in recent years that markups are fairly low (cite someone here), but I need to keep the markup above one to ensure monopolistic competition, I set μ to 1.01.

The prior distributions for the rest of the DSGE model parameters are given in table 2.1a. The mean and standard deviation for the normal priors of the growth rate ln ? and inflation rate ln p center the distributions near the sample means for the period, while allowing a wide range for the parameters, so as not to make these priors too informative. With a view to a minimally informative prior on the inflation process, I also use a flat (uniform) prior on the autoregressive term for the inflation process ρ (between 0 and 1), the correlation between technology shock and inflation disturbance $corr(\varepsilon_z, \varepsilon_\pi)$, (between –1 and 0, given that this correlation is almost certainly negative). The variance terms for technology σ_z^2 and inflation σ_π^2 are given inverse gamma prior distributions, with 2 degrees of freedom (resulting in infinite prior variances for these terms). For the investment adjustment cost parameter φ, I wanted a reasonably diffuse but still proper prior. To ensure that, instead of estimating φ itself, I estimated $\varphi/(1-\varphi)$, which is constrained to fall between zero and one, assuming a prior distribution for $\varphi/(1-\varphi)$ uniform between zero and one.

The data have very little information for the depreciation rate d and the labor share of output ϕ; I chose priors that allow most values for these parameters used in the RBC literature to fall comfortably within two standard deviations of the mean. The discount rate β is also not identified. The best analogy to the nominal interest rate in my model is a short-term interest rate such as federal funds, the real mean value over most of the sample period would suggest a $\beta = \dfrac{\gamma\pi}{1+R}$ of just about one; I set the prior mean in the end at 0.999. This results in a capital-output ratio that is probably too high; however, the results presented here are not sensitive to reasonable prior distributions for δ, ϕ or β.

As for price stickiness, I have stated the prior in terms of the average period of price stickiness in quarters, $1/(1-\alpha)$. Most micro surveys of firms' price behavior suggest a range for this parameter of two to four quarters; an estimate of 2.5 to 3.5 is obtained by a macro study by Sbordone (2002) using a measure of marginal cost instead of output, so I picked a prior constraining $1/(1-\alpha)$. to be above one and assuming the most probable values (within two standard errors) to be two to four. Gali and Gerlter (1999) found rather higher estimates in the 5 to 6 range, but they conjecture that these are probably biased upward, and are sensitive to sample period; at any rate, as labor contracts and the "menus" of various firms are usually revised at most annually, there is no reason to expect $1/(1-\alpha)$ to be much above four.

Step 1

As the posterior distributions of the Bayesian VARs are Normal-Wishart like the priors, it is straightforward to take draws from them, and calculating moments of interest (as well as posterior probabilities). However, the posterior distribution of the DSGE models, being the product of the likelihood (evaluated using a Kalman filter algorithm) and prior distributions of various shapes, clearly isn't analytically tractable, so I cannot generate draws from them directly. So, I used a Metropolis-Hastings algorithm to generate parameter draws from the distribution. The results reported here are based on 80,000 draws from that algorithm. The posterior means and standard errors for the parameter draws from the Metropolis algorithm, for both the flexible-price and New Keynesian models, are given in Table 2.1b. Studying these gives some preliminary clues regarding the performance (good or bad) of the DSGE models.

Starting on a positive note, for both models the posterior distributions for the parameters suggest not a great deal of misspecification. For the structural parameters, neither model updates the prior distributions very much, suggesting that for both models the data is quite compatible with reasonable values of δ, ϕ or β. For both models the measure of investment adjustment costs $\varphi/(1-\varphi)$ is substantial, though the estimated range is quite wide. The inflation process in both models is fairly well behaved as well; the persistence term ρ in both models is around 0.9, with a fairly small error. $corr(z,\pi)$ is constrained to be negative; this constraint seems to be binding, but not unduly so, for the flexible price model (zero falling within two standard errors of the mean), while for the New Keynesian model it seems not to be too binding at all.

In the New Keynesian model, the posterior mean average price fixity $1/(1-\alpha)$ is almost spot on at three quarters. Hence the data on output and inflation and their behavior are compatible with survey data on the stickiness of prices. However, given no obvious signs of misspecification in the flexible-price model, it is fair to ask whether the data really suggest an average price fixity of three quarters, or whether the data are simply uninformative about average price fixity and so are compatible with a range of priors.

This is confirmed by the calculated posterior odds for the various models, given in Table 2.2. I assigned prior probabilities of one-third to both DSGE models (flexible price versus sticky price/New Keynesian), leaving a probability of one-third for the reference model (the VAR). Choice of lag length is always a matter of uncertainty with VAR models, so I split the one-third probability four ways, with one-twelfth probability each going to VARs with one, two, three and four lags.

The posterior model probabilities are proportional to the Bayes factors, for which, again, the formula is $p(Y_T \mid M_i) = \int p(Y_T \mid \theta_{(i)}, M_i) p(\theta_{(i)} \mid M_i) d\theta_i$. For each model I calculated a modified harmonic mean estimator (Geweke 1999) of the Bayes factor of the form

$$p_{HM}(Y_T \mid M_i) = \left[\frac{1}{n_{sim}} \sum_{s=1}^{n_{sim}} \frac{f(\theta_{(i)}^{(s)})}{p(Y_T \mid \theta_{(i)}^{(s)}, M_i) p(\theta_{(i)}^{(s)} \mid M_i)} \right]^{-1}$$

where $f(\theta) = (2\pi)^{-d_i/2} \left| V_{\theta,i} \right|^{-1/2} \exp\left[-0.5 \left(\theta - \bar{\theta}_i \right)' V_{\theta,i}^{-1} \left(\theta - \bar{\theta}_i \right) \right]$, $\bar{\theta}_{(i)}$ being the calculated posterior mean of $\theta_{(i)}$ and $V_{\theta,i}$ being the posterior covariance matrix of $\theta_{(i)}$ generated by the Metropolis algorithm. The $\left| V_{\theta,i} \right|^{1/2}$ term can be interpreted as penalizing high dimensionality, as do more common model selection criteria as the Schwarz criterion (which can be interpreted as a rough approximation to the Bayes factor; see the next chapter below and Kass and Raftery 1995 for more details).

One clear result is that the DSGE models dominate the Bayesian VARs; the combined probabilities of the two DSGE models compared to the VARs suggest a

DSGE/BVAR posterior odds of 104:1. Changes in the BVAR prior distribution might alter those odds somewhat, but what is clear is that the DSGE models are not only competitive in fit but also possibly superior. Among the advantages of the DSGE models are greater parsimony (11 free parameters for the New Keynesian model versus 21 for the best of the BVARs, the BVAR(4)) allowing for greater precision in estimation of moments.

Here, however, the good news for New Keynesians ends. As expected, the posterior probability is greater for the New Keynesian than the flexible-price model. However, the posterior odds of the sticky-price model are only about three to one; as it were, if we allow price flexibility and price stickiness to be equally likely a priori (as my prior probabilities imply), we wind up not being able to "reject the null" of perfect price flexibility at any reasonable level (above 90%, that is, a nine to one ratio). To be sure, to put a one-half prior probability on price flexibility implies a prior distribution assigning a probability of one-half on $1/(1-\alpha)$ equaling one, when my actually prior beliefs are the gamma distribution on $1/(1-\alpha)$ centered at three and giving a prior probability of $1/(1-\alpha)$ equaling one of zero. More consistent with such a prior belief would have been to set the prior probability of the flexible price model equal to zero. Relaxing that assumption, however, brings into sharper relief just how little information the inflation and output data really possess; this calls into question the usefulness of such macro studies as those of Gali and Gertler (1999) or Sbordone (2002), who essentially mine the output and inflation data in search of the most plausible value of $1/(1-\alpha)$.

To underline the point that the evidence from the output and inflation data alone that there is substantial price stickiness and/or New Keynesian elements in the business cycle, I re-estimated the New Keynesian model, this time with a flat prior on α (i.e. a uniform prior between 0 and 1). The posterior mode of α was about 0.77, suggesting average price fixity of slightly over a year, and consistent with Gali and Gertler's results using a classical generalized method of moments estimator. However, the posterior distribution was very flat, with values for α from as low as 0.2 to as high as 0.85 not possible to rule out. The modified harmonic mean estimator of the Bayes factor for the "flat prior" New Keynesian model was about 1188.22, below that for the flexible price model that restricted α to zero; Occam's razor would then suggest taking the flexible-price model as a benchmark, and assuming no price-stickiness at all. It is clear that prior information on the amount price-stickiness has to be imposed from without for us to be able to learn much about inflation's impact on output, given that the data are compatible with such a wide range of degree of price stickiness.

The difference between the results here and Schorfheide's (2000) results that found monetary DSGE models come far less from the imposition of New Keynesian effects as with other changes to the model. By stating monetary policy in terms of inflation instead of monetary growth, the "spike" in inflation generated by a persistent change in monetary growth, which is inconsistent with the data, is removed here. The addition of investment adjustment costs also improves output dynamics to the degree that the overall fit of the model is much improved. If φ is restricted to zero, the Bayes factor of the flexible price model falls to 1180.74 and that of the New Keynesian model to 1182.93, leaving posterior odds of about 7:1 in favor of the BVAR(4).

Step 2

The good in-sample fit of (correctly-specified) DSGE models is encouraging, but probably of more interest is the behavior these models imply for the output-inflation relationship. It would also be helpful to know along what lines the "bad" models (the VARs) fail. To do this, I calculated loss functions based on predicted population characteristics. Specifically, following Schorheide, I considered: (a) correlations between output growth and inflation $corr(\Delta \ln GDP_t, \Delta \ln P_{t+h}), h = -2,\ldots,2$, and (b) the responses of output growth $d\Delta \ln GDP / d\varepsilon_z$ and inflation $d\Delta \ln P / d\varepsilon_z$ to a permanent technology shock that increases technology by one standard deviation, and the responses of output growth $d\Delta \ln GDP / d\varepsilon_m$ and $d\Delta \ln P / d\varepsilon_m$ to a one standard deviation monetary shock.

To calculate the posterior distribution of the correlations and the IRFs, I obtain the means and covariance matrices of the vectors of correlations and IRFs from each model. Then I assume the posterior density of the correlations and IRFs to be approximately normal with mean

$$E[\omega \mid Y_T] = \sum_{i=0}^{2} \pi_{i,T} E[\omega \mid Y_T, M_i]$$

and covariance matrix

$$V_{\omega|Y_T} = \left(\sum_{i=0}^{2} \pi_{i,T} \left(V_{\omega|Y_T,M_i} + E[\omega'\omega \mid Y_T, M_i] \right) \right) - E[\omega'\omega \mid Y_T].$$

I have argued above why my prior beliefs suggest I should set the prior (and so posterior) probability of the flexible price model equal to zero, and so ignore it in calculating posterior distributions of moments. Since the posterior probabilities of the Bayesian VAR models are also very low, I can also ignore those, and take the posterior distributions from the New Keynesian model as my posterior distribution. In what follows, then, the loss functions calculations taking the distributions from the New Keynesian model to be the "true" one can be regarded as the "true" values of the loss functions. However, I also calculate for most moments the loss functions assuming the BVAR(4) to be the "true" model, for two reasons. First, it is of some interest to examine just how the DSGE and BVAR models differ. It will turn out that, in general, the distribution of the moments is tighter under the DSGE models than under the BVAR, resulting the DSGE moments fitting nicely within the BVAR moment distributions, but the BVAR moments having minimal probability under the DSGE models. Also, in most of the literature classical and Bayesian VARs are the models actually used to study the effects of monetary shocks, and in Schorfheide (2000) the VAR was taken as the "true" or benchmark model, so that the loss functions for the DSGEs, calculated taking the BVAR(4) to be the "true" or reference model, are most closely analogous to those for the DSGE models examined in his paper.

Step 3

Correlations

The posterior distributions for the cross correlations, and the L_p -losses for the correlations implied by the DSGE models taken individually, are given in table 2.3a. At each lead and lag the posterior (that is, the New Keynesian model) suggests a correlation between output growth and inflation between zero and -0.2 at all leads and lags. By contrast, the flexible price model given a range between zero and –0.1 for correlation of current output with leads of inflation, but a tight constraint around zero for correlations of current output with lagged inflation. The lower absolute values for the flexible price model are in part due to the lower estimated correlation between technology and inflation disturbances; the zero correlations with lagged inflation are due to the effect of a monetary/inflation shock being constrained to zero.

In contrast to the narrower range for correlations of the flexible price model (a special case of the New Keynesian model), the BVAR(4) suggests a much wider range of the correlations, between zero and –0.4 for all leads and lags. The distribution of the correlations from the BVAR(4) nest those for the New Keynesian and the flexible-price models comfortably. The flexible-price correlations are on the high end of the range suggested by the BVAR(4), but at no lead or lag are the L_p-risks higher than 0.95; in the case of the New Keynesian model, all are well below 0.9. However, the mean of the correlations for the BVAR (approximating the mode) is on the low end of the plausible levels from the New Keynesian models, and well beyond those from the flexible price model; the L_p-risks of the BVAR correlations assuming flexible prices are all essentially one, and assuming the New Keynesian model all above 0.99 (calling for the BVAR model's "rejection at the 99% level"). The BVAR(4) correlation distribution can be interpreted as the range of correlations between output and inflation compatible with the data; the contribution of the restrictions on dynamics given by the DSGE model is to improve the precision of the correlation estimates. If we assign a substantial posterior probability to the hypothesis that the DSGE models are the "true" models, as the Bayes factors suggest we should, then the lower precision of the BVAR estimates makes for a much greater expected loss from using the results of the BVAR over the DSGE.

By the same token, if we assign a high prior probability (and so a high posterior probability) to the idea that prices are perfectly flexible, the New Keynesian model turns out to be substantially riskier than the flexible price model, given the wider range of correlations it allows without a substantially higher posterior probability (unless the prior odds in favor of price stickiness are extremely higher). The New Keynesian models' distribution of correlations nests the flexible price correlations fairly comfortably; in no case, given the distribution from the New Keynesian model, are the L_p-risks of the flexible price mean correlations higher than 0.95, and most are below 0.9. However, given the distribution from the flexible price model, which essentially restricts the correlations of output and lagged inflation to zero (which the New Keynesian model does not), the L_p-risks of the New Keynesian model's mean correlations of output and lagged inflation are essentially one, though the risks for the correlations with current and leads of inflation are well below 0.9.

The cash-in-advance model suggests a slightly too negative correlation between current output and future inflation; the New Keynesian model compensates for that somewhat, as indicated by L_p-losses closer to 0.5 which would indicate a "perfect fit." (The New Keynesian model, as it were, includes the Phillips curve relation allowing high output to spur an eventual rise in inflation.) Under the cash in advance model, there is essentially no correlation between current output and past inflation; the New Keynesian model allows for a more realistic negative correlation between output and inflation lagged one period, but for essentially none before then.

The fit for the correlations individually is considerably better than that reported in Schorfheide (2000), especially for the relation between output and current inflation, which he found to be too sharply negative, and the relation between output and future inflation (which he found not negative enough). The reason for the better fit, both by the “eyeball metric” and by L_p-loss measure, has, however, as much to do with the more general preferences in my model as opposed to Schorfheide's. In particular, the (implausibly) low risk aversion parameter makes consumption (and hence real money balances) more closely related to the capital stock and less sensitive to technology shocks, reducing fluctuations in consumption (and so inflation) over the cycle.

To get a sense of how risky overall is each model, I evaluated the L_p-losses of all five correlations jointly. The probability $p(\hat{\varphi}_{p,i} \mid Y_T)$ of the posterior modes of the correlations from the DSGE models is approximately $\sum_{j=0}^{2} \pi_{j,T}\, p(\hat{\omega}_{q,i} \mid Y_T, M_j)$, the posterior mean of the probabilities of the posterior means of the DSGE correlations. In Table 2.3b are reported the probabilities of the posterior mean of each model, given each model's distribution, where the distributions are assumed to be multivariate normal distributions with the mean and covariance matrices being the posterior mean and covariance matrices of the correlation draws from each model taken from the Metropolis algorithm. If the New Keynesian model is taken to be the true model (so that $\pi_{2,T} = 1$ and $\pi_{0,T}, \pi_{1,T} = 0$, the probability of the BVAR(4) mean correlations (model 0) are essentially zero, while those from the flexible price model (model 1) are only about 8% as likely as the mean from the New Keynesian model (model 2). By that formulation, clearly not assuming price stickiness is riskier than assuming it. However, this result is sensitive to the assumption that the chance of price stickiness is negligible; if we put take the posterior odds for model 1 versus model 2 given in Table 2.2, the implied probability of the flexible price mean correlations is $\dfrac{(0.224)e^{27.11} + (0.760)e^{19.90}}{(0.224)e^{-392.4} + (0.760)e^{22.47}} \approx 30.6$ times that of the New Keynesian mean correlations. In effect, unless one assumes the probability of perfect price flexibility is very low a priori, the least risky hypothesis a posteriori is that of perfect price flexibility—or in other words, that inflation-related demand shocks have no impact on output fluctuations.

A fairly similar conclusion can be drawn from the calculated L_{χ^2}-losses for the competing models. The L_{χ^2}-losses of a given model *i* given another model *j* can be approximated by the *p*-value of a test of the null hypothesis $H_0 : C_{\chi^2}(\hat{\varphi}_{q,i} \mid Y_T, M_j) \sim \chi^2$ with five degrees of freedom. Taking the BVAR(4) first as the benchmark, the weighted distance $C_{\chi^2}(\hat{\varphi}_{q,i} \mid Y_T, M_0)$ between the mean correlations from the two DSGE models and the means from the BVAR(4) give the edge to the New Keynesian model. However, for both DSGE models the L_{χ^2}-risks are low enough to make "rejection" impossible at any reasonable confidence level. If we take instead the flexible-price model as the benchmark, the values for $C_{\chi^2}(\hat{\varphi}_{q,i} \mid Y_T, M_1)$ reported here imply L_{χ^2}-losses of near one for both the BVAR(4) and the New Keynesian model (and rejection of both models, as it were). Taking finally the New Keynesian model as the benchmark, the value for $C_{\chi^2}(\hat{\varphi}_{q,1} \mid Y_T, M_2)$ reported for the flexible price model has an implied L_{χ^2}-loss that cannot allow rejection at any reasonable confidence level; however the value $C_{\chi^2}(\hat{\varphi}_{q,0} \mid Y_T, M_2)$ reported for the BVAR(4) implies an L_{χ^2}-loss of about one (and clear rejection). Clearly, unless a priori perfect price flexibility is ruled out, inference using a model assuming perfect price flexibility is a much less risky proposition, in part because inflation effects on output are small enough, given the size and persistence of inflation disturbances that they can be for many purposes be ignored in analyzing output fluctuations.

Impulse-response functions

Looking (in Figure 2.1) at the posterior means of the impulse responses of output growth and inflation to one-standard-deviation shocks in technology and in money (i.e. in inflation disturbances unrelated to technology shocks) will help make this clear. Looking first at the impact of a technology shock on output, the impulse-response function for the flexible price and for the New Keynesian model are visually indistinguishable. On the other hand, that from the BVAR(4) using Blanchard-Quah ordering suggests a much lower impact of technology on output—and by implication a much larger role for monetary/aggregate demand shocks in output fluctuations. The posterior mean response from the BVAR(4) of a monetary shock on output growth is more than twice that of the New Keynesian model. (By assumption, the impact of a monetary shock on output in the flexible price model is zero.)

The mean impact of a technology shock on inflation implied by the BVAR(4) is also more than twice that of the New Keynesian model, while that for the flexible price model is only about a third that of the New Keynesian model. These results match those found with the correlations, with the highest correlations between output and inflation coming from the BVAR and the lowest from the flexible price model. The BVAR suggests a larger role for technology/supply shocks in inflation determination than do the DSGEs, and a correspondingly smaller role for monetary/demand shocks; the difference between the mean responses of inflation to a monetary shock for the two DSGE models is small, and both are considerably higher than that implied by the BVAR(4).

Simply looking at the posterior means is a little misleading, however, for in fact, again, the DSGE provides a good deal more precision in disaggregating output and inflation into technology and monetary shocks than the BVAR can. Figure 2.2 reports 90% confidence intervals for the IRFs from the New Keynesian DSGE model along with those from the BVAR(4). Except for the money-on-inflation IRF, the BVAR(4) confidence interval for each IRF comfortably nests the New Keynesian IRF confidence interval. The BVAR(4) confidence intervals for technology and money shocks on outpupt are particularly wide; for both shocks the 90% confidence interval allows for an impact of essentially zero to up to 0.8 percent, near the average size of an output disturbance in quarterly US GDP growth data. In other words the data is compatible with either all output disturbances being attributed to permanent (technology) shocks, all being attributed to temporary (monetary/demand) shocks, or any case in between. The restrictions imposed by the Blanchard-Quah identification scheme alone are not strict enough to shed light on which shock, temporary or permanent, plays a larger role in output fluctuations. The restrictions imposed by the DSGE model go far enough beyond those of the Blanchard-Quah scheme (though inasmuch as technology shocks are permanent and inflation shocks temporary, the Blanchard-Quah restrictions are satisfied as well) to greatly increase the precision of the IRFs and reveal, at last, that permanent technology shocks are by far the more important, the impact of a technology shock being near the maximum the Blanchard-Quah identification scheme would allow as plausible.

Table 2.4 reports the loss function calculations for the IRFs (taken at lags 0 through 11) of the various competing models, given as benchmarks either the New Keynesian model or the BVAR(4). (As the distribution of $d\Delta \ln GDP / d\varepsilon_m$ is degenerate for the flexible price model, being zero at all lags, L_{χ^2}-risk for any other model implying non-zero responses; hence, L_{χ^2}-risks taking flexible-price as the benchmark are not reported.) For the most part, they confirm what the "eyeball metric" suggests for the precision, and the compatibility of the various models. Taking the New Keynesian model as benchmark, the calculated $C_{\chi^2}(\hat{\varphi}_{q,0} \mid Y_T, M_2)$ values for all four BVAR(4) mean IRFs are enormous and L_{χ^2}-risk essentially one. However, the calculated $C_{\chi^2}(\hat{\varphi}_{q,1} \mid Y_T, M_2)$ values for the four flexible-price IRFs are extremely low. For all but $d\Delta \ln GDP / d\varepsilon_m$, L_{χ^2}-risk is close to zero, and for $d\Delta \ln GDP / d\varepsilon_m$ itself L_{χ^2}-risk is only about .31, suggesting that a hypothesis of zero (or rather negligible) impact of inflation on output cannot be rejected at any reasonable confidence level. Meanwhile, if we take the BVAR(4) as a benchmark, the L_{χ^2}-risks are very low for the mean IRFs of both DSGE models, indicating they are well within the range of possibilities suggested by the VAR. For all but the response of inflation to a temporary (monetary) shock, the L_{χ^2} risks are near zero; only for $d\Delta \ln P / d\varepsilon_m$ are they even above .1, and even then none are less than .4.

I also report the L_q-losses, given either the New Keynesian model or BVAR(4) as benchmark, for the competing models, using as a weighting matrix a 12 x 12 identitiy matrix times 1/12 (giving the average square deviation of the benchmark IRF from the competing model). Using the New Keynesian model as benchmark, the BVAR(4) has consistently larger L_q-losses than the flexible price model; for $d\Delta \ln GDP / d\varepsilon_z$, in particular, the L_q-losses are over 25 times greater, and for $d\Delta \ln P / d\varepsilon_m$ nearly ten times greater. Using the BVAR(4) on the other hand, the margins between the two DSGE models are not nearly as large; however, the New Keynesian model's losses are consistently lower. If one takes, as many would, the BVAR to be the true model, the New Keynesian DSGE would appear to be the least risky. However, when a substantial or even overwhelming probability exists that the DSGE is the true one, the BVAR(4) becomes much riskier by several orders of magnitude.

Conclusions

The loss-function based evaluation of a New Keynesian DSGE model estimated using Bayesian methods leads to several conclusions. Contrary to Schorfheide's (2000) findings for a cash-in-advance Bayesian DSGE model, it is possible to construct and estimate Bayesian DSGE models with New Keynesian features that are competitive, if not superior, in fit to a Bayesian vector autoregression in output and inflation. The advantage of the Bayesian DSGE model is its ability to greatly improve the precision of estimates of such moments as correlations between output and inflation as well as impulse response functions, without an undue cost in terms of fit.

Another conclusion may be more surprising. The improved performance of the New Keynesian DSGE model used here comes from elements such as investment adjustment costs and a monetary rule in terms of inflation (not, as in Schorfheide, a monetary growth rule without feedback) that have little to do with the New Keynesian features of the model. Disappointingly, perhaps, for a researcher devoted to a New Keynesian explanation of business cycles, the output and inflation data bear surprisingly little information about the average price stickiness in the US economy. The problem is that inflation shocks have contributed such a small amount to US business cycles (enough that they can easily be regarded as negligible) that the amount of information one could reasonably expect the US GDP series alone to bear regarding the impact of inflation on output is quite small. One conclusion that could easily be drawn is that studies that mine output data, or marginal cost data derived from output data, for information on the slope of the New Keynesian Phillips curve may well be best abandoned in favor of micro surveys of price stickiness, and plausible levels of price stickiness imposed on the data a priori for the purpose of analyzing the impact of monetary shocks.

It should be added that there are, naturally, other sources of information besides improved priors, that is, other sources of data such as nominal interest rates. If anything, the monetary policy model in this paper is unrealistically simple, not least because central banks set nominal interest rates, not inflation as such. Also, it is rather doubtful all the variation in output attributed to temporary shocks in the Blanchard-Quah identification scheme comes from monetary/inflation shocks; much could easily come from non-monetary sources of demand shocks, such as government spending and preference shocks. In the next chapter, we will take a closer look at how to disaggregate data using a DSGE model into technology, money and non-monetary demand shocks, and explore each one's roles in business cycle fluctuations.

Table 2.1a: Prior distributions for structural parameters of DSGE models

Parameter	Range	Density	Mean	Std. deviation
$\ln\gamma$	$\Re$	Gaussian	0.005	0.005
$\ln\bar{\pi}$	$\Re$	Gaussian	0.010	0.005
β	[0,1]	Beta	0.999	0.0005
δ	[0,1]	Beta	0.015	0.0025
ϕ	[0,1]	Beta	0.700	0.025
$\varphi/(1-\varphi)$	[0,1]	Uniform	0.5	0.289
ρ	[0,1]	Uniform	0.5	0.289
$corr(\varepsilon_z,\varepsilon_\pi)$	[-1,0]	Uniform	-0.5	0.289
$\alpha/(1-\alpha)$	$\Re^+$	Gamma	2	0.5
			Scale	d.f.
σ_z	$\Re^+$	Inverse gamma	0.008	2
σ_π	$\Re^+$	Inverse gamma	0.002	2

Table 2.1b: Posterior distributions for structural parameters of DSGE models

	Model			
	Flexible price		New Keynesian	
Parameter	Mean	Std. deviation	Mean	Std. deviation
$\ln\gamma$	0.0053	0.0009	0.00512	0.0009
$\ln\bar{\pi}$	0.0102	0.0021	0.0103	0.0018
β	0.999	0.0005	0.999	0.0005
δ	0.0150	0.0025	0.0149	0.0024
ϕ	0.700	0.0251	0.701	0.0237
$\varphi/(1-\varphi)$	0.361	0.161	0.301	0.154
ρ	0.891	0.0378	0.878	0.0327
$corr(\varepsilon_z,\varepsilon_\pi)$	-0.0932	0.0577	-0.312	0.130
$1/(1-\alpha)$			3.099	0.467
σ_z	0.0116	0.0012	0.0118	0.0012
σ_π	0.0028	0.0002	0.0028	0.0002

Table 2.2: Posterior odds for competing models

Model	Flex. price DSGE	N-K DSGE	BVAR(1)	BVAR(2)	BVAR(3)	BVAR(4)
Prior prob.	1/3	1/3	1/12	1/12	1/12	1/12
Log Bayes factor	1188.36	1189.58	1184.70	1184.94	1185.95	1186.32
Post. prob	0.224	0.760	0.001	0.002	0.005	0.007
Post. odds, $\pi_{i,T}/\pi_{1,T}$	1	3.39	0.006	0.008	0.022	0.032

Table 2.3a: Distributions and loss function values for output-inflation correlations of competing models, taken individually

			Correlations, $corr(\Delta \ln GDP_t, \Delta \ln P_{t+h})$				
	Model		$h=-2$	$h=-1$	$h=0$	$h=1$	$h=2$
Distributions	Flex. price	Mean	-0.009	-0.014	-0.054	-0.047	-0.042
		Std. err.	0.007	0.010	0.037	0.031	0.027
	N-K	Mean	-0.090	-0.083	-0.092	-0.080	-0.070
		Std. err.	0.045	0.043	0.059	0.051	0.043
	BVAR(4)	Mean	-0.258	-0.253	-0.247	-0.219	-0.198
		Std. err.	0.128	0.128	0.129	0.130	0.132
	Benchmark	Model	$h=-2$	$h=-1$	$h=0$	$h=1$	$h=2$
L_p-risk	Flex-price	N-K	~1	~1	0.702	0.698	0.690
		BVAR(4)	~1	~1	~1	~1	~1
	N-K	Flex-price	0.930	0.891	0.478	0.477	0.475
		BVAR(4)	~1	~1	0.991	0.994	0.997
	BVAR(4)	Flex-price	0.948	0.939	0.866	0.809	0.763
		N-K	0.810	0.817	0.771	0.711	0.669

Table 2.3b: Loss function values for output-inflation correlations of competing models, taken jointly

		$\ln \tilde{p}(\hat{\omega}_{q,i} \mid M_j)\|$		
		$j=0$	$j=1$	$j=2$
L_p-loss	$i=0$ (BVAR(4))	12.16	-5876	-453.45
	$i=1$ (flex-price)	8.92	27.11	19.90
	$i=2$ (N-K)	10.97	-392.4	22.47
		$C_{\chi^2}(\omega_{q,i} \mid Y_T, M_j)$		
		$j=0$	$j=1$	$j=2$
	$i=0$ (BVAR(4))	0	11807	951.79
	$i=1$ (flex-price)	6.48	0	5.10
L_{χ^2}-risk	$i=2$ (N-K)	2.38	839.09	0
		L_{χ^2}-risk		
		$j=0$	$j=1$	$j=2$
	$i=0$ (BVAR(4))	0	~1	~1
	$i=1$ (flex-price)	0.737	0	0.596
	$i=2$ (N-K)	0.205	~1	0

Table 2.4: Loss function values for impulse response functions of competing models, taken jointly

			Impulse response functions			
			$\frac{d\Delta \ln GDP}{d\varepsilon_z}$	$\frac{d\Delta \ln P}{d\varepsilon_z}$	$\frac{d\Delta \ln GDP}{d\varepsilon_m}$	$\frac{d\Delta \ln P}{d\varepsilon_m}$
Benchmark						
	L_q-loss	$i=0$	0.0086	0.0033	0.0086	0.0033
		$i=1$	0.00023	0.0012	0.0046	0.00034
New Keynesian	$C_{\chi^2}(\omega_{q,i} \mid Y_T)$	$i=0$	1.12E12	8.09E11	3.42E12	3.78E12
		$i=1$	1.67	2.78	9.14	0.92
	L_{χ^2}-risk	$i=0$	~1	~1	~1	~1
		$i=1$	0.00023	0.0031	0.309	9.01E-6
BVAR(4)	L_q-loss	$i=1$	0.0116	0.0084	0.0260	0.0050
		$i=2$	0.0086	0.0033	0.0086	0.0033
	$C_{\chi^2}(\omega_{q,i} \mid Y_T)$	$i=1$	5.38	6.05	4.86	10.03
		$i=2$	4.36	3.76	1.87	9.77
	L_{χ^2}-risk	$i=1$	0.056	0.087	0.038	0.386
		$i=2$	0.024	0.012	0.00042	0.363

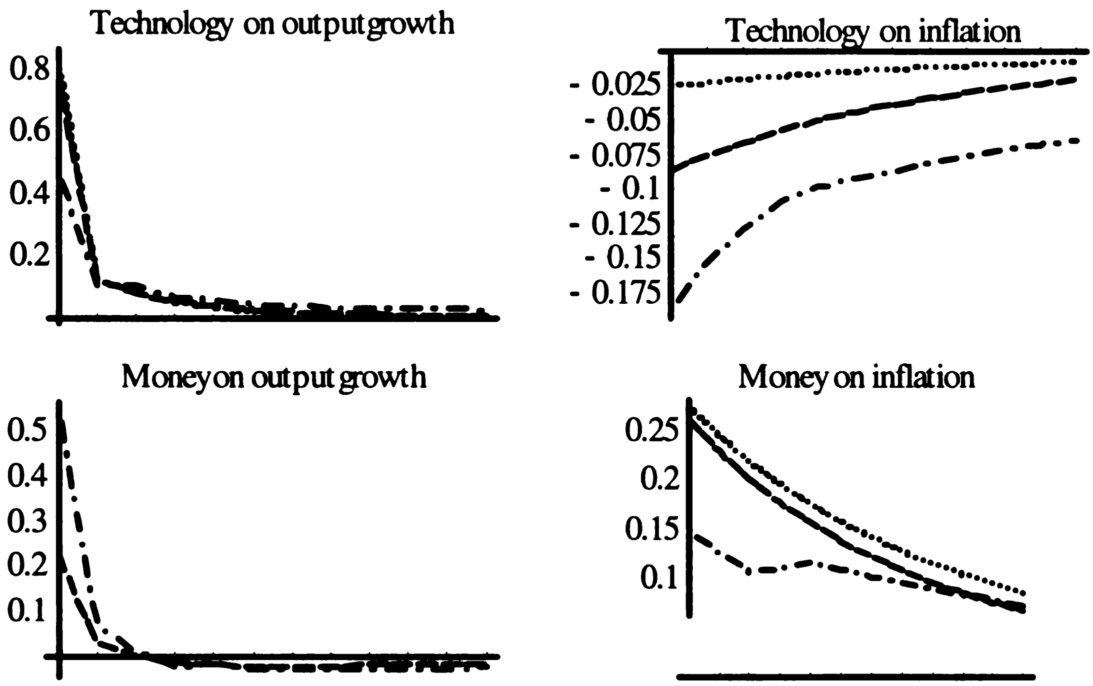

Figure 2.1: Posterior means of impulse response functions from estimated cash-credit DSGE model (dotted lines), New Keynesian DSGE model (dashed lines) and BVAR(4) model (dotted-dashed lines)

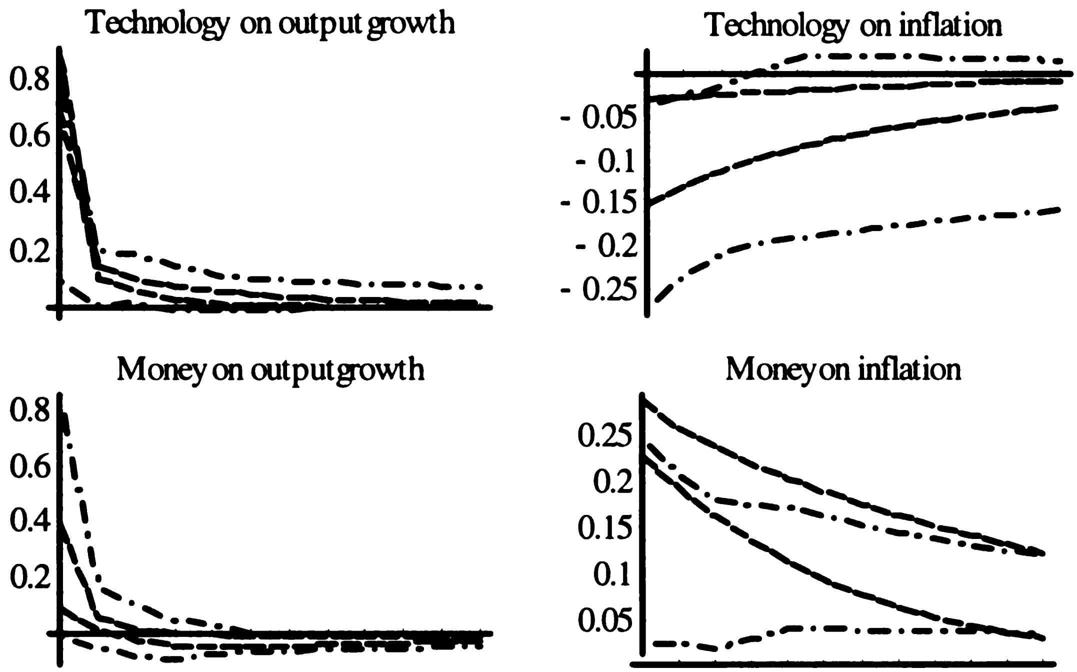

Figure 2.2: 90% confidence bands of impulse response functions from New Keynesian DSGE model (dashed lines) and BVAR(4) model (dotted-dashed lines).

CHAPTER 3

TECHNOLOGY SHOCKS VERSUS MONETARY SHOCKS: IDENTIFYING SOURCES OF BUSINESS CYCLE FLUCTUATIONS WITH A NEW KEYNESIAN DSGE MODEL

Introduction

To answer the questions of "What does monetary policy do?" and "Does monetary policy generate recessions?" we need to be sure, keeping the critique of Rudebusch (1998) in mind, that the measure of monetary policy shocks used by our models make sense. Most of the methods used to identify monetary shocks in the vector autoregression (VAR) literature, which still comprises the bulk of empirical literature on the effects of monetary policy, are not wholly satisfactory in that regard. The most widely used method for identifying monetary shocks are zero restrictions on contemporaneous effects (e.g. Bernanke and Mihov 1998; Christiano, Eichenbaum and Evans 1996, 1998; Leeper et al. 1996), but the institutional rigidities used to defend these restrictions are at best debatable.

Other restrictions that have sometimes been used are long-run restrictions on the effects of shocks (e.g. Gali 1992) and sign restrictions on their effects (e.g. Uhlig 1999); Identification schemes based on the heteroskedasticity of structural shocks have also been proposed, for example by Rigobon (2003). Identification schemes restricting signs and long-run effects of temporary monetary shocks can be defended with a number of models, among them so-called "New Keynesian" models, which have become quite popular in recent work on monetary policy and its effects (e.g. Clarida, Gali and Gertler 2000; Leeper and Zha 2001; McCallum and Nelson 1999). New Keynesian models are dynamic stochastic general equilibrium models (DSGE models) of various levels of elaboration, descended from the "real business cycle" (RBC) models of Kydland and Prescott (1982). New Keynesian models add such features as imperfect competition and price and wage stickiness to a basic RBC model to permit monetary and other demand disturbances to play a larger role in economic fluctuations. The long-run effect restrictions imposed by Gali (1992) were inspired by restrictions that would be imposed by an IS-LM model, or a modern New Keynesian DSGE equivalent.

Arguably more appealing from a theoretical (or an aesthetic) standpoint would be to have done with ad-hoc identification schemes and use an estimated DSGE model to identify sets of monetary shocks. In practice, however, little attention has been paid to the possibility of identifying monetary shocks with DSGE models. The purpose of this paper is to fill that gap.

The earliest papers in the DSGE literature devoted to estimating series of shocks (e.g. Hansen and Prescott 1993) did so using calibrated models. The implicit assumption was that DSGE models were too highly stylized to be worth estimating. However, more recent, richer DSGE models have done a better job of fitting US macro time series than their predecessors. A large literature now exists using classical maximum likelihood methods and/or generalized method of moments (e.g. Christiano and Eichenbaum 1992; McGrattan 1994; McGrattan et al. 1997) to fit RBC models containing several shocks to the data. More recent work has used Bayesian methods to estimate and assess non-monetary DSGE models. Dejong et al. (2000a, 2000b) studied a non-monetary DSGE model containing total factor productivity and marginal efficiency of investment shocks. Among the advantages of a Bayesian approach is the ability to supplement the data with prior information on structural parameters, which is helpful when parameter estimates conflict markedly with other evidence.

Several recent papers answer the common criticism that DSGE models, monetary or non-monetary, are not competitive with classical VAR models in terms of fit. Kim (2000) finds a monetary DSGE model with a realistic monetary rule fits the output, inflation, money supply and interest rate data about as well as a classical VAR. del Negro and Schorfheide (2003) found that a stylized monetary DSGE model could be used to construct a prior for a Bayesian VAR that performed as well, or better, than the random-walk "Minnesota" prior (Litterman 1986) in forecasting US output and inflation. Smets and Wouters (2002) estimated a DSGE model of the euro area using seven data series and ten types of shocks, including such real rigidities as habit formation as well as nominal wage and price rigidity. They assessed the fit of their model by calculating Jeffreys-Bayes posterior odds (Zellner 1971), and found it comparable to that of a Bayesian VAR using a Minnesota prior for each time series.

The present paper tries to use a monetary DSGE model to answer the question: Do monetary policy shocks generate recessions? For that matter, have they been as important as supply shocks in US business cycles? The approach used here to answer the question is to estimate a series of technology/supply shocks, monetary shocks and non-monetary demand shocks using a well-fitted DSGE model, and examine what information content these shocks—in particular monetary shocks—can provide in post-war US business cycles.

The technique of the paper is similar to Dejong et al. (2000a, 2000b). The DSGE model used is estimated with output, inflation and interest rates only, similar to the del

Negro and Schorfheide (2002) study. Following Dejong et al., the likelihood function implied by the monetary DSGE model is combined with a prior distribution for the model parameters to derive a posterior distribution for model parameters, impulse responses and series of shocks.

Dejong et al. found that their model forecast output and investment about as well as a Bayesian VAR using a "Minnesota" prior, but did not directly compare in-sample fit. To compare the fit of this paper's DSGE model to that of competing models, and assess how trustworthy the model is for business cycle analysis, this paper follows Schorfheide (2000) and Smets and Wouters (2002) in calculating Bayesian posterior odds for the DSGE model versus a competing model, specifically a Bayesian VAR model of output, inflation and the interest rate. The DSGE model is competitive with the Bayesian VAR, suggesting that it describes the dynamics of output, inflation and interest rates about as well as the BVAR. Among the advantages of using the DSGE model (as opposed to a BVAR of roughly equal fit) lies in its identifying the shocks with direct appeal to the theoretical model, giving the shock series firmer basis in theory than the ad hoc methods used to construct monetary shocks from a VAR. Another is that the restrictions from the DSGE allow a good deal more precision in describing the relationships of the various time series used to estimate the model.

With the various shock series in hand, it is possible to study the behavior and interactions of the shocks around NBER recessions. As a rule, monetary shocks lead US recessions, while supply shocks coincide with recessions. However, supply shocks are a more reliable indicator of NBER recessions. Logit models of the probability of recessions, estimated using current and lagged shock series as explanatory variables, yield accurate predictions of the turning points.

This paper's results differ from those of Smets and Wouters in their implications for the relative importance of shocks. Smets and Wouters found technology shocks to be of secondary importance in European business cycles compared to preference and labor supply shocks, as well as monetary shocks. Most of the explanatory power from the logit model for NBER recessions estimated here comes from supply shocks. Models omitting supply shocks are inadequate, and by standard model selection criteria suggest the evidence that monetary shocks have had much predictive content for recessions, particular post-1982 recessions, is quite weak.

The main source of the difference between the results below versus those of Smets and Wouters, this paper will argue, is the different prior assumptions about the levels of nominal price rigidity and intertemporal substitution of consumption. The prior distributions for nominal price rigidity used here restricts it to levels that micro data suggest are most plausible for the US economy. The prior used for intertemporal substitution of consumption reflects the very low estimates found in the consumption literature (e.g. Hall 1988). Given the fairly steep IS curves and Phillips curves that result, demand-side shocks generally, and monetary shocks in particular, are given less role in business cycles as compared to supply or real technology shocks than Smets and Wouters implied for Europe. The results below come closer to the conclusions of the VAR literature that has found for the most part "responses of real variables to monetary policy shifts are estimated as modest or nil, depending on specification" (Sims 1998, p. 933).

A Small Scale New Keynesian Macro Model

The model used in what follows is similar to models used in del Negro and Schorfheide (2002) and Smets and Wouters (2002), the biggest difference being a modification to make preferences compatible with balanced growth. Consumers' preferences also allow for habit formation in consumption and leisure, following recent research such as Fuhrer (2000). Nominal rigidity is added by following Calvo (1983) and allowing imperfectly competitive firms only a positive probability (less than one) of being able to adjust their prices in a given period. The systematic part of monetary policy is taken to be a modified Taylor rule (Taylor 1993), which specifies the interest rate target as a function of output and inflation.

Preferences

A large number of consumers have preferences that can be expressed by the lifetime utility function

$$E_t \sum_{t=0}^{\infty} \beta^t \left(\varepsilon_t \frac{1}{1-\sigma} \left((x_t - \phi X_{t-1})^{1-\sigma} + (m_t / P_t)^{1-\sigma} \right) \right),$$

m_t / P_t is individual holdings of real money balances at time t, with m_t being holdings of nominal balances and P_t being the economy-wide price level. x_t is a composite of consumption and leisure so:

$$x_t = c_t^{1-\eta} l_t^{\eta},$$

where c_t is individual consumption and $l_t = 1 - h_t$, where h_t is the fraction of period t engaged in labor. (In this section small letters will denote individual decision variables, and capitals will denote economy-wide aggregates.)

Economy wide average consumption and leisure (denoted by X_{t-1}), which individuals take as given, enters individual utility functions to generate habit formation in consumption and leisure; the parameter ϕ determines the amount of habit formation. Lifetime expected utility has a finite value if and only if $\bar{\beta} = \beta\gamma^{(1-\eta)(1-\sigma)} < 1$.

The term ε_t measures shifts in preferences that follow the AR(1) process

$$\ln \varepsilon_t = (1-\rho_\varepsilon)\ln \varepsilon_{t-1} + v_{\varepsilon t}, v_{\varepsilon t}\,\text{iid}, E_{t-1}v_{\varepsilon t} = 0, E_{t-1}v_{\varepsilon t}^2 = \sigma_\varepsilon^2.$$

It will become clear that the effect of these shocks is to raise output, employment and inflation, suggesting an aggregate demand shock. With that in mind, preference shocks will be referred to as "autonomous demand shocks" or just "demand shocks" from now on.

While labor/leisure is homogeneous in this economy, consumer goods are not. c_t is a composite of a large number of imperfectly substitutable consumer goods c_{it}, each produced by an individual firm *i*. Specifically, aggregate consumption as a function of the individual goods is

$$c_t = \left(\int_0^1 c_{it}^{(\delta-1)/\delta} di \right)^{\delta/(\delta-1)},$$

where $\delta > 1$. Cost minimization yields individual demand functions for each good c_{it} of the form

$$c_{it} = (p_{it} / P_t)^{-\delta} c_t,$$

where p_{it} is the price of good *i*. δ then is the elasticity of demand for each good *i*. P_t can be expressed in terms of the individual p_{it}'s as

$$P_t = \left(\int_0^1 P_{it}^{1-\delta} di \right)^{1/(1-\delta)}.$$

At the beginning of each period t each individual carries forward bonds with a nominal value of b_{t-1} dollars, and money balances with a nominal value of m_{t-1} dollars. Individuals sell h_t units of labor in the competitive labor market at the real wage rate w_t per hour. They also receive dividends from ownership of monopolistically competitive firms in the amount of r_t. Finally, individuals receive a helicopter drop of x_t dollars each on top of the receipts from maturing bonds of b_{t-1} dollars, leaving them with total cash balances of $b_{t-1} + m_{t-1} + x_t$ dollars. Individuals then purchase consumer goods, privately issued bonds and money balances subject to the budget constraint

$$c_t + \frac{1}{1+R_t}\frac{b_t}{P_t} - \frac{b_{t-1}}{P_t} + \frac{m_t}{P_t} - \frac{m_{t-1}}{P_t} - \frac{x_t}{P_t} - w_t h_t - r_t \leq 0,$$

R_t being the prevailing nominal interest rate.

The first order conditions of the consumer's problem are

(C1: consumption demand) $$(1-\eta)c_t^{-\eta}(1-h)_t^{\eta}(x_t - \phi X_{t-1})^{-\sigma} = \lambda_t,$$

(C2: labor supply) $$\eta c_t^{1-\eta}(1-h)_t^{\eta-1}(x_t - \phi X_{t-1})^{-\sigma} = w_t \lambda_t,$$

(C3: demand for bonds) $$\frac{1}{1+R_t}\lambda_t = E_t\left(\frac{P_t}{P_{t+1}}\lambda_{t+1}\right),$$

(C4: money demand) $$(m_t / P_t)^{-\sigma} = \lambda_t - E_t\left(\frac{P_t}{P_{t+1}}\lambda_{t+1}\right) = \frac{R_t}{1+R_t}\lambda_t.$$

(C1) and (C2) easily yield

(C5) $$\frac{\eta}{1-\eta}c_t = w_t(1-h_t),$$

while (C1) and (C4) yield

(C4')
$$\frac{R_t}{1+R_t}(1-\eta)c_t^{-\eta}(1-h)_t^{\eta}(x_t-\phi X_{t-1})^{-\sigma}=(m_t/P_t)^{-\sigma},$$

the money demand function for this economy.

The assumed preferences for consumption and leisure fix the marginal rate of substitution between consumption and leisure at one. This ensures balanced growth, with individual consumption and the wage rate following the same trend, and labor effort being a constant $\bar{h}$ at the steady state.

Production

Consumers own, and supply labor to, monopolistically competitive firms that produce according to a technology expressible as

(F1: Technology)
$$y_{it}=(z_t(h_{it}-h_0)),$$

where y_{it} is the output of firm i, h_{it} is the amount of labor used by that firm and z_t is economy-wide level of technology, the evolution process for which follows a random walk with drift process

$$\Delta\ln z_t=\ln\gamma+v_{zt}, v_{zt}\text{iid}, E_{t-1}v_{zt}=0, E_{t-1}v_{zt}^2=\sigma_z^2,$$

the technology shock v_{zt} being white noise and γ being the long-run rate of economic growth. h_0 is a fixed cost in terms of labor that ensures a nontrivial barrier to entry. Technology can be broadly interpreted here as an aggregate of non-labor inputs to production, including institutional factors (for instance, regulatory regimes) as well as technology. Supplies of natural resources are also a "supply-side" constraint on productive capacity; the measure of technology shocks reported below will certainly pick up the disruptions in oil supplies that so commonly caused economic disruptions in recent

US history. Given this, "supply shocks" and "technology shocks" will be used interchangeably in what follows.

The cost minimization problem of firm i is

$$\min tc_{it} = w_t h_{it} + r_t \text{ such that } y_{it} = \left(z_t (h_{it} - h_0)\right).$$

The solution of that problem can be expressed in the form $tc_{it} = w_t h_0 + mc_t y_{it}$, where mc_t is the real marginal cost of producing an additional unit of y_i, given as

(F2: Labor demand) $$mc_t = w_t / z_t .$$

The instantaneous real profit of firm i at time t is then

$$\begin{aligned} \Pi_{it} &= (p_{it} / P_t) y_{it} - tc_{it} = (p_{it} / P_t) y_{it} - w_t h_0 - mc_t y_{it} \\ &= \left(p_{it} / P_t\right)^{1-\delta} y_t - w_t h_0 - mc_t \left(p_{it} / P_t\right)^{-\delta} Y_t \end{aligned}$$

where Y_t is economy-wide aggregate supply. Under perfect price flexibility, maximizing profits would imply $\left(p_{it} / P_t\right) = \dfrac{\delta}{\delta - 1} mc_t$ for all firms; as it would be the case, all firms being identical, that $p_{it} = P_t$, it would follow that $mc_t = \dfrac{\delta - 1}{\delta}$ for all t. The average economy wide markup would be $\mu \equiv \dfrac{Y_t}{w_t h_t} = \dfrac{1}{mc_t} = \dfrac{\delta}{\delta - 1}$. While this is indeed the long-run trend of the markup, it will not hold in general for all firms if prices are not perfectly flexible.

The precise mechanics of price-stickiness, as in the previous sections, follow that in Calvo (1983). Each firm faces a probability $\alpha > 0$ of getting a "stop signal" and being unable to adjust their prices as it wishes, being only able to raise it by the long-run average rate of inflation $\bar{\pi}$. The remaining fraction 1- α of firms receives a "go signal" to adjust their prices freely. If a firm is allowed adjust its price, its choice for its product's price is the solution to the maximization problem

$$\max_{p_{it}} E_t \sum_0^\infty (\alpha\beta)^j (\lambda_{t+j} / \lambda_t)(((\pi^j p_{it} / P_{t+j})^{1-\delta} - mc_t (\pi^j p_{it} / P_{t+j})^{-\delta})Y_{t+j} - w_{t+j} h_0),$$

which has the solution

(F3: Price-setting equation)

$$p_t^* = \frac{\delta}{\delta - 1} \frac{E_t \sum_0^\infty (\alpha\beta)^j (\lambda_{t+j} / \lambda_t)(Y_{t+j} / Y_t) mc_{t+j} P_{t+j}^{\delta} \pi^{-j\delta}}{E_t \sum_0^\infty (\alpha\beta)^j (\lambda_{t+j} / \lambda_t)(Y_{t+j} / Y_t) P_{t+j}^{\delta-1} \pi^{-j(\delta-1)}}.$$

Meanwhile, firms whose prices are "stopped" choose the price $p_{it} = \bar{\pi} p_{it-1}$. As all firms, given their status as "stop" or "go" firms, are otherwise identical, these first order conditions and the price level aggregation function $P_t = \left(\int_0^1 p_{it}^{1-\delta} di \right)^{1/(1-\delta)}$ imply a process for the price level obeying

(F4: Price-level evolution process) $\quad P_t^{1-\delta} = (1-\alpha) p_t^{*1-\delta} + \alpha(\pi P_{t-1})^{1-\delta}$.

Monetary policy

To close the system completely requires specifying a monetary rule for the central bank. Here it is assumed the central bank follows a monetary rule that can be expressed in the form

(MP1: Monetary policy)

$$(1 + R_t) = v_{rt} R((1 + R_{t-1}), Y_t / (z_t (\bar{h} - h_0)), \pi_t, Y_t / Y_{t-1}, \pi_t / \pi_{t-1}).$$

That is, the Fed sets the interest rate with regard to the output gap (the deviation of aggregate output Y_t from its long-run trend $z_t(\bar{h} - h_0)$ where $\bar{h}$ is the long-run average for h_t, the current inflation rate π_t, the economic growth rate Y_t / Y_{t-1}, and the rate of change of the inflation rate, π_t / π_{t-1}. The central bank also smoothes changes in the interest rate, so that the lagged interest rate $(1 + R_{t-1})$ enters the reaction function. Also influencing the target interest rate is a multiplicative, non-systematic and iid monetary shock v_{rt} which has the properties $E_{t-1} \ln v_{rt} = 0, E_{t-1}(\ln v_{rt})^2 = \sigma_r^2$.

Market clearing

The aggregate demand and supply for consumer goods are equal, so that

(MK1: Goods market clearing) $$C_t = Y_t;$$

there being no inventories, this is a result of all firms selling their production at the going price for their good, so that $c_{it} = y_{it}$ for all i. Labor demand and supply also equalize at the competitive wage rate w_t, so that

(MK2: Labor market clearing) $$h_t = H_t = \int_0^1 h_{it} di.$$

Aggregate demand and supply for money and bonds also equalize. In the case of bonds, as all individuals are identical, net demand for bonds is zero at the market interest rate. We have

(MK3: Money market clearing) $$m_t = M_t.$$

(MK4: Bond market clearing) $$b_t = 0.$$

How To Estimate The Model

The model above must be log-linearized and put in state-space form to make estimation feasible. Log-linearizing (C1) through (C4), imposing the goods market clearing condition (MK1) (which log-linearizes as $\hat{C}_t = \hat{Y}_t$) and a little manipulation, gives me at last

(C5') $$\hat{Y}_t = \hat{w}_t - \frac{\bar{h}}{1-\bar{h}}\hat{H}_t,$$

where the deviation of a variable from its long-run trend is denoted by marking it with a circumflex. Log-linearizing the production function (F1) yields

(F1') $$\hat{Y}_t = \mu\hat{H}_t;$$

The fact that the output gap $\hat{Y}_t$ is proportional to employment $\hat{H}_t$ can be used to write (C3) in the form of an expectations-augmented Phillips curve. (C5) and the fact that in the long run $Y_t = \mu w_t h_t$ can be used to write the leisure preference parameter in terms of the average markup and the average labor effort as $\eta = \frac{1-\bar{h}}{1-\bar{h}+\mu\bar{h}}$. With those facts in hand, and the fact that the demand shock process log-linearizes as $\hat{\varepsilon}_t = \rho_\varepsilon \hat{\varepsilon}_{t-1} + \hat{v}_{\varepsilon t}$ implying $E_t\hat{\varepsilon}_{t+1} = \rho_\varepsilon \hat{\varepsilon}_t$, the log-linearized form of (C3) can be written in the form of an expectations-augmented IS curve, so:

(IS) $$\begin{aligned}&(1\hat{+}R)_t - E_t\hat{\pi}_{t+1} \\ &= \sigma\left(1-\frac{(1-\bar{h})\mu+\bar{h}}{\mu(1-\bar{h}+\mu\bar{h})}\right)\left(\frac{1}{1-\phi}E_tY_{t+1} - \frac{1+\phi}{1-\phi}Y_t + \frac{\phi}{1-\phi}Y_{t-1} - \frac{\phi}{1-\phi}v_{zt}\right) \\ &+\left(\frac{(1-\bar{h})\mu+\bar{h}}{\mu(1-\bar{h}+\mu\bar{h})}\right)\left(E_tY_{t+1}-Y_t\right)+(1-\rho_\varepsilon)\hat{\varepsilon}_t.\end{aligned}$$

Log-linearizing (F3) yields

$$\hat{p}_t^* = (1-\alpha\beta)\left(\sum_0^\infty (\alpha\beta)^j \left(E_t \hat{mc}_{t+j} + E_t\hat{\lambda}_{t+j} + E_t\hat{Y}_{t+j} + \delta E_t \hat{P}_{t+j}\right)\right)$$

$$-(1-\alpha\beta)\left(\sum_0^\infty (\alpha\beta)^j \left(E_t \hat{mc}_{t+j} + E_t\hat{\lambda}_{t+j} + E_t\hat{Y}_{t+j} + (\delta-1) E_t \hat{P}_{t+j}\right)\right)$$

$$= (1-\alpha\beta)\left(\sum_0^\infty (\alpha\beta)^j \left(E_t \hat{mc}_{t+j} + E_t \hat{P}_{t+j}\right)\right),$$

implying $$\hat{p}_t^* - \alpha\beta E_t \hat{p}_{t+1}^* = (1-\alpha\beta)(\hat{mc}_t + \hat{P}_t).$$

Log-linearizing the price index process (F4) gives us $\hat{P}_t = (1-\alpha)P_t^* + \alpha\hat{P}_{t-1}$, given that in the long run $P_t^* = P_t$. The fact that $\hat{\pi}_t = \hat{P}_t - \hat{P}_{t-1}$ implies $\hat{\pi}_t = (1-\alpha)(\hat{P}_t^* - \hat{P}_t)$. Substituting into (F3) (and a little rearrangement finally gets us our Calvo-type Phillips curve:

(F3') $$\hat{\pi}_t = \beta E_t \hat{\pi}_{t+1} + \frac{(1-\alpha)(1-\alpha\beta)}{\alpha}\hat{mc}_t.$$

(F2) implies

(F2') $$\hat{mc}_t = \hat{w}_t,$$

so that combining (F2'), (F3') and (C5') gives a Phillips curve of the form:

(Phillips curve) $$\frac{\mu(1-\bar{h})+\bar{h}}{\mu(1-\bar{h})}\frac{(1-\alpha)(1-\alpha\beta)}{\alpha}\hat{Y}_t = \hat{\pi}_t - \beta E_t \hat{\pi}_{t+1}.$$

Finally, taking the elasticity of the target interest rate with regard to each of the endogenous variables to be constant, log-linearizing the monetary policy reaction function gives a rule of the form

(Monetary policy) $$\begin{aligned}(1\hat{+}R_t) &= \rho_r(1\hat{+}R_{t-1}) + (1-\rho_r)(r_y\hat{Y}_t + r_\pi\hat{\pi}_t) \\ &+ r_{\Delta y}(\hat{Y}_t - \hat{Y}_{t-1} + v_{zt}) + r_{\Delta\pi}(\hat{\pi}_t - \hat{\pi}_{t-1}) + \hat{v}_{rt}.\end{aligned}$$

The three equations of the model can be written as

(IS curve)

$$
\begin{aligned}
&(1\hat{+}R)_t - E_t\hat{\pi}_{t+1} \\
&= \sigma\left(1-\frac{(1-\bar{h})\mu+\bar{h}}{\mu(1-\bar{h}+\mu\bar{h})}\right)\left(\frac{1}{1-\phi}E_t\hat{Y}_{t+1}-\frac{1+\phi}{1-\phi}\hat{Y}_t+\frac{\phi}{1-\phi}\hat{Y}_{t-1}-\frac{\phi}{1-\phi}v_{zt}\right) \\
&+\left(\frac{(1-\bar{h})\mu+\bar{h}}{\mu(1-\bar{h}+\mu\bar{h})}\right)\left(E_t\hat{Y}_{t+1}-\hat{Y}_t\right)+(1-\rho_\varepsilon)\hat{\varepsilon}_t,
\end{aligned}
$$

(Monetary policy)

$$
\begin{aligned}
&(1\hat{+}R_t) = \rho_r(1\hat{+}R_{t-1})+(1-\rho_r)(r_y\hat{Y}_t+r_\pi\hat{\pi}_t) \\
&+r_{\Delta y}(\hat{Y}_t-\hat{Y}_{t-1}+v_{zt})+r_{\Delta\pi}(\hat{\pi}_t-\hat{\pi}_{t-1})+\hat{v}_{rt},
\end{aligned}
$$

where ρ_r is the interest rate smoothing parameter, and

(Phillips curve)

$$
\frac{\mu(1-\bar{h})+\bar{h}}{\mu(1-\bar{h})}\frac{(1-\alpha)(1-\alpha\beta)}{\alpha}\hat{Y}_t = \hat{\pi}_t - \beta E_t\hat{\pi}_{t+1}.
$$

The log-linearized system can be interpreted as an expectations-augmented IS-LM model. Positive demand shocks influence aggregate demand by shifting the IS curve, while monetary shocks do so by shifting the "LM curve" formed by targeting of the interest rate. However, the full impact of the shocks on aggregate demand is not felt in the period of the shock because of habit formation. The Phillips curve describes the deviation of aggregate supply from trend $\hat{Y}_t$ as proportional to the current inflation rate minus expected future inflation.

The log-linearized version of the model can be written (following Hamilton 1994) in the state space form:

$$
\begin{aligned}
?_t &= \mathbf{F}?_{t-1}+\mathbf{v}_t \\
\mathbf{y}_t &= \mathbf{A}'\mathbf{x}_t+\mathbf{H}'?_t
\end{aligned}
$$

The vector of observable variables is $\mathbf{y}_t = \{\Delta \ln GDP_t, \Delta \ln P_t, \ln(1+R)_t\}'$. The vector of deterministic variables is $\mathbf{x}_t = \{\ln \gamma, \ln \bar{\pi}, \ln \bar{\beta}\}$ for all t, mapped into the constant terms for the dependent variable equations by $\mathbf{A} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 1 & -1 \end{pmatrix}$. The state variables are $?_t = \{\hat{v}_{zt}, \hat{\varepsilon}_t, \hat{v}_{rt}, (1 \hat{+} R_{t-1}), \hat{Y}_{t-1}, \hat{\pi}_{t-1}\}$, and the mapping matrix $\mathbf{H}$ and the transition matrix $\mathbf{F}$ are

$$\mathbf{H}' = \begin{pmatrix} 1+\varphi_{yz} & \varphi_{ye} & \varphi_{yv} & \varphi_{yr} & \varphi_{yy}-1 & \varphi_{y\pi} \\ \varphi_{\pi z} & \varphi_{\pi e} & \varphi_{\pi v} & \varphi_{\pi r} & \varphi_{\pi y} & \varphi_{\pi\pi} \\ \varphi_{rz} & \varphi_{re} & \varphi_{rv} & \varphi_{rr} & \varphi_{ry} & \varphi_{r\pi} \end{pmatrix} \text{ and}$$

$$\mathbf{F} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \rho_{\varepsilon} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \varphi_{rz} & \varphi_{re} & \varphi_{rv} & \varphi_{rr} & \varphi_{ry} & \varphi_{r\pi} \\ \varphi_{yz} & \varphi_{ye} & \varphi_{yv} & \varphi_{yr} & \varphi_{yy} & \varphi_{y\pi} \\ \varphi_{\pi z} & \varphi_{\pi e} & \varphi_{\pi v} & \varphi_{\pi r} & \varphi_{\pi y} & \varphi_{\pi\pi} \end{pmatrix},$$

where the φ's are complicated functions of the model parameters. The mapping matrix uses the fact that the observables can be approximated by

$$\begin{aligned} \Delta \ln GDP_t &= \hat{v}_{zt} + \hat{Y}_t - \hat{Y}_{t-1} \\ \Delta \ln P_t &= \ln \bar{\pi} + \hat{\pi}_t \\ \ln(1+R)_t &= \ln \gamma + \ln \pi - \ln \bar{\beta} + (1 \hat{+} R)_t. \end{aligned}$$

The shock vector $\mathbf{v}_t = \{v_{zt}, v_{\varepsilon t}, 0, v_{rt}, 0, 0\}$ has the covariance matrix

$$\mathbf{Q}=\begin{pmatrix}\sigma_z^2 & 0 & 0 & 0 & 0 & 0\\ 0 & \sigma_\varepsilon^2 & 0 & 0 & 0 & 0\\ 0 & 0 & 0 & 0 & 0 & 0\\ 0 & 0 & 0 & \sigma_r^2 & 0 & 0\\ 0 & 0 & 0 & 0 & 0 & 0\\ 0 & 0 & 0 & 0 & 0 & 0\end{pmatrix}.$$

The Kalman filter can easily be used to calculate the likelihood $L(\mathbf{Y}\mid\mathbf{B})$, of the system, where $\mathbf{Y}=\{\mathbf{y}_t\}_{t=1}^{T}$ are the 3xT matrix observations of $\mathbf{y}_t$ over the full period from time 1 to time T, and $\mathbf{B}=\{\gamma,\bar{\beta},\mu,\bar{h},\sigma,\pi,\alpha,\rho_\varepsilon,\rho_r,r_y,r_\pi,r_{\Delta y},r_{\Delta\pi},\sigma_z,\sigma_\varepsilon,\sigma_r\}$ is the vector of the model's structural parameters. Combining the likelihood with a suitable prior distribution $f_{prior}(\mathbf{B})$ will imply a posterior distribution with a PDF satisfying $f_{post}(\mathbf{B})\propto L(\mathbf{Y}\mid\mathbf{B})f_{prior}(\mathbf{B})$.

By taking draws from the posterior, it is straightforward to calculate posterior distributions for the moments of interest. These include the structural parameters, impulse response functions, and the values of the structural shocks $\{v_{zt},v_{\varepsilon t},v_{rt}\}$ at time *t*. In a Kalman filter framework (see Hamilton 1994 for details) it is straightforward to calculate a historical series of the values of the state variables, given the information we have at time *T*, $?_{t|T}$, with the values for the expectation of the states at each time *t* given information up to *t*-1, $?_{t|t-1}$. With those in hand, the historical series of shocks, given the data up to time *T* and a certain value of $\mathbf{B}$, are just $\mathbf{v}_{t|T,B}=?_{t|T,B}-\mathbf{F}?_{t-1|T,B}$.

On the data and the priors

The data used to calculate the likelihood were quarterly data from 1965:I to 2003:I on economic growth (logged growth rate of per capita real GDP), inflation (logged growth rate of the chain type GDP deflator) and the federal funds rate. These were obtained from the Federal Reserve Bank of St. Louis FRED database. Data from 1959:I to 1964:IV were used to initialize the Kalman filter, but not to actually estimate the likelihood.

The prior used to estimate the Bayesian VARs were Normal-Wishart priors of the Kadiyala-Karllson type, similar to those used in chapter 2; the hyperparameter values used for the experiments in chapter 2 were also used here, and so it is pointless discussing the prior for the VARs in detail here. Of more interest is the prior distributions for the 17 model parameters for the DSGE model are given in table 3.1. All the prior distributions must be proper (i.e. integrate to one) for estimated Bayes factors to be meaningful (Kass and Raftery 1995). However, the results were for the most part not sensitive to the priors (with a few exceptions), so most were picked to be in line with values common in the literature, where the data were likely to be uninformative, or as "gently uniform" as possible when the data were likely to be informative. The prior distributions for the economic growth rate γ and the long-run inflation rate $\bar{\pi}$ are centered on their sample means for the 1965-2003 period. Large prior standard deviations were assumed to prevent either prior from being too informative.

The standard deviations of the technology shock σ_z, the demand shock σ_ε and the monetary shock σ_r were assumed to have inverse gamma distributions with 2 degrees of freedom, implying an infinite prior variance. The scale parameters s were picked after some experimentation, guided by results from estimation with a truly diffuse prior. Uniform distributions were used on the persistence parameter for demand shocks ρ_ε and on the monetary policy smoothing parameter ρ_r. The prior distributions on the output gap term r_y and the inflation term r_π were picked so as to center them at values consistent with Taylor's original rule, giving values of 0.5 for the output gap and 1.5 for inflation. There was less guidance regarding the output growth term $r_{\Delta y}$ and the inflation growth term $r_{\Delta\pi}$; some experimentation suggested prior means of 0.125 for $r_{\Delta y}$ and 0.5 for $r_{\Delta\pi}$, with wide standard deviations. All the terms in the monetary rule were assigned gamma distributions, to constrain them above zero. The data were also fairly informative regarding the discount factor β, suggesting a value in the ballpark of 0.99875 for quarterly data (essentially one).

The prior mean for the average markup μ was 1.2, in line with the fairly low estimates for the parameter found in recent studies (e.g. Basu and Fernald 1995). The distribution was a gamma distribution in terms of $\mu-1$, to ensure a value for μ above one. The prior mean of average amount of free time devoted to labor $\bar{h}$, was 0.3, a common value in the RBC literature. The prior distribution of habit formation ϕ, was assumed uniform between 0 and 1.

More critical are the priors for the price fixity parameter, α, and the curvature parameter for the utility function, σ. Many older studies (Roberts 1995; Gali and Gertler 1999) suggested price fixity in the range of five to six quarters; however, this is at odds with survey data (e.g. Blinder 1994) that suggests price fixity for a typical firm is closer to three quarters. Some more recent macro studies (e.g. Sbordone 2002) concur. The results below assume a prior gamma distribution on $\alpha/(1-\alpha)=1/(1-\alpha)-1$ centered at three with a standard deviation of one; this implies values for $1/(1-\alpha)$ from two to six, a range most researchers would find plausible, and centers average price fixity a priori at one year.

The value of the reciprocal of the curvature parameter for the utility function, $1/\sigma$, is related to the intertemporal elasticity of substitution of consumption. The value of this parameter remains somewhat controversial; what little has been determined regarding the intertemporal elasticity of substitution of consumption in the US economy suggests that it is probably not large. A large literature starting with Hall (1988) has been unable to show that intertemporal elasticity of consumption is much above zero with aggregate consumption data. By contrast, research with DSGE models generally assumes values for the reciprocal of the intertemporal elasticity of consumption from one to five; in the Bayesian New Keynesian DSGE literature, del Negro and Schorfheide use a prior implying a range for the reciprocal from one to three, while Smets and Wouters center their prior for the reciprocal at one.

The results reported below assume a gamma prior distribution for $1/\sigma$, with a prior mean of 0.01 and one degree of freedom, so that the prior variance of σ is infinite and the prior mode of $1/\sigma$ itself is zero. The implication of the prior for intertemporal elasticity can be best understood by recalling the IS equation

(IS curve)
$$
\begin{aligned}
&(1\hat{+}R)_t - E_t\hat{\pi}_{t+1} \\
&= \sigma\left(1 - \frac{(1-\bar{h})\mu + \bar{h}}{\mu(1-\bar{h}+\mu\bar{h})}\right)\left(\frac{1}{1-\phi}E_t\hat{Y}_{t+1} - \frac{1+\phi}{1-\phi}\hat{Y}_t + \frac{\phi}{1-\phi}\hat{Y}_{t-1} - \frac{\phi}{1-\phi}v_{zt}\right) \\
&+ \left(\frac{(1-\bar{h})\mu + \bar{h}}{\mu(1-\bar{h}+\mu\bar{h})}\right)\left(E_t\hat{Y}_{t+1} - \hat{Y}_t\right) + (1-\rho_\varepsilon)\hat{\varepsilon}_t,
\end{aligned}
$$

Ignoring habit formation, the response of the output gap to a unit change in the interest rate is not $1/\sigma$ but rather $1/\left(\sigma\left(1-\frac{(1-\bar{h})\mu+\bar{h}}{\mu(1-\bar{h}+\mu\bar{h})}\right)+\left(\frac{(1-\bar{h})\mu+\bar{h}}{\mu(1-\bar{h}+\mu\bar{h})}\right)\right)$; the prior mean for this term has a value of about 0.1, Hall's (1988) point estimate for intertemporal elasticity of substitution of consumption.

Estimation results

Posterior distributions of structural parameters

Following Schorfheide (2000) and Smets and Wouters (2002), the well-known Metropolis-Hastings algorithm was used to get the DSGE model's posterior distribution and the posterior odds. (See, for example, Geweke (1999) for a discussion of the Metropolis-Hastings algorithm.) Table 2a gives the posterior means and standard deviations of the structural parameter values drawn from the Metropolis-Hastings algorithm. Most values are not far from their prior means, and in most cases standard deviations are smaller, suggesting that the data is informative and that the priors are consistent with the data. The demand shock persistence parameter is high, approximately in the 0.9 range, as is the interest rate smoothing parameter. The amount of habit formation is substantial, with a posterior estimate of ϕ of around 0.3.

The posterior mean of $1/(1-\alpha)$ is slightly higher than the prior mean, suggesting price fixity of about four quarters (a bit longer than a year). $1/\sigma$ is also not too far from its prior mean; its value along with the value of habit formation suggests intertemporal substitution of around 0.07. Both have posterior variances not far from the prior, suggesting little updating. The lack of updating suggests though that the data are actually fairly uninformative about $1/(1-\alpha)$ and $1/\sigma$, and so of the IS and Phillips curves.

Compare the results in Table 2a to those of Smets and Wouters (2002), who, using a rather different prior, get different results for Europe. As noted above, they use a prior suggesting a prior mean of $1/\sigma$ of around one, along with a loose prior on α restricting it to be between 0.5 and 1. Their model also fits about as well as a Bayesian VAR. However, their posterior estimate for α is about 0.9, suggesting average price fixity $1/(1-\alpha)$ of about ten quarters. Other evidence suggests price stickiness is much less than this, in Europe as well as in the US. Gali, Gertler and Lopez-Salido (2001) find a more realistic estimate of $1/(1-\alpha)$ for the euro area would be closer to three or four. Like Gali and Gertler (1999), Smets and Wouters do not trust their high point estimate for α, attributing it to misspecification of the marginal cost process.

The finding of counterfactually high α, however, is not robust to the prior distribution, particularly a more informative one on α constraining it away from such values. Part of Smets and Wouters' problem derives from their relatively uninformative prior for α, which places a non-trivial weight on values well above plausible levels. The real problem, however, is subtler than this. Study of the posterior correlation matrix (not reported in full here) suggests that draws of α and $1/\sigma$ are positively correlated; the correlation coefficient is about 0.3. The data, then, are somewhat informative for a combination of the two parameters, but not for each individually. The data are compatible with both high price fixity and high intertemporal elasticity of consumption/output (and a large role for monetary and other demand shocks), or for both low price fixity and low intertemporal elasticity of consumption (and a small role for monetary and demand shocks).

In essence we have an identification problem, in that the data alone do not give the model much guidance in disaggregating disturbances into supply, demand or monetary shocks. To be able to say anything useful about which shocks are important in the business cycle, then, we need the best prior information we have available regarding parameters for price fixity and intertemporal elasticity of consumption. The marginal value of Bayesian estimation, where outside information on these parameters can be taken into account, is clearly positive here.

What about regime shifts in monetary policy?

The time period of the data set used here is similar to that of data sets commonly used in VAR studies, but it spans a period usually considered to contain several policy regimes. Rudebusch (1998) strongly objected to use of such data sets to estimate VARs. Rudebusch claimed that one reason monetary policy shocks from VAR models did not make sense was their counterfactual assumption of a time-invariant reaction function.

It turns out that splitting the sample to allow for the policy to vary over different time periods (especially the 1979-1982 period) did not alter the qualitative nature of the results. Among the more noticeable effects was a reduction in the estimated long-run rate of inflation—allowing for a higher volatility of shocks in the 1979:III-1982:II period (the period of the Volcker "monetarist" experiment) down-weights these observations somewhat, and this period was marked by high inflation. The policy rules for the pre-Volcker experiment period (1965:I-1979:II) and post-Volcker experiment period (1982:III-2003:I) were also more accommodating; during the Volcker experiment, monetary shocks were less persistent than during the Martin-Burns-Miller and late Volcker-Greenspan periods. However, the differences were not large enough to affect overall conclusions. The models with multiple regimes did soundly beat the model without regime shifts, on a posterior odds basis, but most of the improvement came from allowing for the variance of monetary shocks (and to a lesser extent, supply shocks) to vary over periods.

Posterior odds-based evaluation of the model

Following the method described in chapter 2, to compare the DSGE model to the competition, posterior odds were calculated for the DSGE model as well as a number of Bayesian VARs, estimated with up to four lags, by calculating Bayes factors for each model. As in Chapter 2, the approximate Bayes factor reported for each model is the modified harmonic mean estimator suggested in Geweke (1999). Table 3 gives the estimated Bayes factors and posterior odds for the DSGE model along with those estimated for BVARs with the Normal-Wishart prior. For the DSGE model, the results are based on 50,000 draws from the Metropolis-Hastings algorithm; for the BVARs, they are based on 10,000 draws from the Normal-Wishart posterior.

Of the BVARs, the model with four lags is best from a posterior odds viewpoint. However, the posterior odds of the BVAR (4) are lower than those for the DSGE model. Using the prior distribution in Table 1, the posterior odds are more than 5 to 1 in favor of the DSGE model versus the BVAR (4). For the DSGE versus all the BVARs, it is slightly less than 4 to 1. The margin is nowhere large enough to suggest abandoning VAR models entirely in favor of DSGE models. However, it does suggest that the DSGE model is competitive in fit compared to a four lag BVAR, and so is well specified enough to be useful in business cycle analysis. The Bayes factor results parallel the results of the forecasting experiments of Dejong et al. (2000a). Geweke (1999) argues that Bayes factors can be interpreted as a measure of forecasting performance over the sample period and so the results here suggest that the forecasting performance of the DSGE model is roughly that of the BVAR.

The impulse responses: What do monetary policy shocks do?

The 95% confidence bands for estimated impulse responses of output, inflation and interest rates to a one-standard deviation shock of each type are illustrated in Figure 1. The distributions are taken from 10,000 draws from the posterior distribution of the DSGE model.

All three sets of impulse responses are of reasonable shape. A technology shock raises output, and also lowers inflation slightly. The response of interest rates to a technology shock is ambiguous; the median is slightly above zero, but the confidence band nests zero easily. Less ambiguous is the response of economic variables to a demand shock. A demand shock raises inflation and output, and also raises interest rates, partly from higher expected future inflation, partly from anti-inflationary monetary policy. A monetary policy shock raises the interest rate and lowers output and inflation. However, notice that the scale of output responses to "average" demand and monetary shocks is considerably smaller in magnitude than the responses of output to the technology shock. Visual inspection of the impulse responses alone suggests that the contribution of typical monetary shocks to the movement of output will be minor compared to that of technology.

Because of habit formation, the full effects of shocks on output only come with a slight lag. The lag, however, is not nearly as long for monetary shocks as VAR studies often imply; the output effect peaks here at two quarters, whereas in many VAR studies output does not “hit bottom” in response to a monetary shock until up to three years after the shock, which seems too long for a transitory shock. Also encouraging is the unambiguously negative response of inflation to the monetary shock; the “price puzzle” that arises in small-scale VARs, which postulate a positive response of inflation to a monetary shock, is absent here.

Moving from real to nominal variables, study of the impulse responses of inflation and interest rates to the various shocks suggests that the demand shock is most important for determining inflation, with monetary shocks of at most secondary importance. Finally, the systematic portion of monetary policy seems to be mostly made up of responses to demand shocks, not so much technology shocks. This is as it should be, if we presume countercyclical monetary policy is designed to moderate deviations of output from the long-run trend determined by technology, not to counteract expansions driven by genuine growth in the economy’s productive capacity. This fact adds to the plausibility of the model’s measure of monetary shocks (which are deviations from the demand shock-driven systematic policy). Thanks to interest rate smoothing, the full adjustment of the target only comes with a lag of up to two quarters.

What might not be obvious from the impulse responses is that a positive technology shock results in output rising by less than the full proportion of the technology shock. Price stickiness results in firms being unable to adjust their prices quickly enough to be able to sell a proportionately large number of goods. The response of firms whose prices are "stuck" to a technology shock that lowers marginal cost is to lay off workers, rather than increase output. Hence labor demand, real wages and employment fall in response to a technology shock.

This is a common phenomenon in New Keynesian models, interesting here mostly for its implications for the importance of demand versus supply shocks. The data favor (in both VAR and DSGE models) a mildly countercyclical inflation rate. If prices are assumed to be stickier than they actually are, the response of output to a technology shock will be measured as much smaller than is the case. As a result inflation will be too countercyclical; to patch up the problem demand side shocks that imply a procyclical output-inflation relationship will be assigned a excessive role in recessions by the model.

Cross-correlations of time series: How do the BVARs and the DSGE differ?

An "identification-free" method of comparing the implications of the competing models for the data is to compare the correlations at various leads and lags of the various series implied by the DSGE model versus a BVAR. These are reported in Tables 3.3a-c. Only the cross-correlations for the best-fitting BVAR (the BVAR (4)) are reported. The 95% confidence intervals of the cross-correlations for both the BVAR and DSGE models are calculated from 10,000 draws from the posterior distributions of each.

The DSGE and BVAR agree on a negative correlation between output and inflation, though the DSGE model's correlations are smaller in magnitude. They also agree on a positive correlation between inflation and nominal interest rates. Where they differ qualitatively is the correlations between output and nominal interest rates; while the correlation between present output and future (nominal) interest rates is negative in the BVAR, it is positive in the DSGE. This is consistent with the positive relation of economic growth and monetary policy assumed in the monetary rule. The correlation of lagged nominal interest rates and present output is positive in both, but those for the

DSGE are much smaller in magnitude, the median being about 0.03, essentially zero, versus the median from the BVAR of about 0.3.

Another qualitative difference in the DSGE versus the BVAR correlations is the smaller standard errors of the DSGE correlations. To quantify the greater precision of the DSGE correlations versus the BVAR correlations, I calculated the minimal L_{χ^2}-losses from using each model's set of correlations, given that the other was the true model (for details, see chapter 2). The L_{χ^2}-loss statistic penalizes imprecision, and it is obvious that the BVAR's estimates of the correlations are much less precise across the board; the BVAR's L_{χ^2}-losses are anywhere from 30 to 200 times larger than those of the DSGE.

Also reported is the CDF of the L_{χ^2}-loss given a chi-square distribution with 7 degrees of freedom; taking as a first approximation that the distributions of each set of correlations in each model are approximately multinormal, this will provide a test of the hypothesis that a given model could have produced the posterior mean of the correlations from the other model. The p-values of the posterior means of the BVAR correlations given the DSGE are all about one; clearly, the BVAR's estimates of the correlations are inconsistent with the restrictions of the DSGE. Meanwhile, the cross-correlations of output and inflation, and inflation and interest rates, reported by the DSGE cannot be rejected by the BVAR, as it were, at any reasonable significance level. However, the BVAR does reject the DSGE's estimates of the output-interest rate relationship at the 95% level.

The greater restrictions placed on the data by the DSGE result in much greater precision being put on the ranges in which the cross-correlations may fall; as a result, the

BVAR correlations have a much lower probability of being produced if the DSGE is the "true" model than vice versa, suggesting a much greater risk from choosing the BVAR estimates if the choice is wrong. Among the benefits of using the DSGE model is that it permits greater precision regarding the relationships of economic time series by adding more information about economic structure, without large costs in terms of fit.

The shocks themselves: Do monetary policy shocks generate recessions?

It is safe to conclude at this point that the DSGE model is well specified; its fit is competitive with a VAR, and it generates well-behaved impulse responses. It also allows more precision regarding the relationship of economic variables. Given this, some confidence is possible in assessing the effects of individual shocks as measured by the DSGE model in particular US business cycles.

The lowest-cost way of doing this is to visually inspect the shock series and their behavior around individual recessions. In figures 3.2a-c are reported, in standard deviation terms, the medians of the distributions of the technology (3.2a), demand (3.2b) and monetary (3.2c) shocks. By far the best indicator of NBER recessions (marked in the graphs by vertical lines) is the technology shock series; above-average and persistent technology shocks coincide with all the recessions of the post-1965 period (1970, 1974-5, 1980, 1981-2, 1990-1, 2001). All but the 1970 and 2001 recessions were associated with disruptions in oil supplies. More likely culprits for the 1970 and 2001 recessions, both fairly mild, were a decline in investment demand after the prolonged booms both followed, coupled with changes in tax policy just before the 1970 and the 1990-1 recessions.

The demand shock series that by rights should best capture shocks to investment demand, unrelated to long-run technology, has a much harder time matching the historical record of cycles. The 1974-5 recession is associated with a *positive* demand shock; while there was a negative shock in the 1970 recession, it was no larger than average for the pre-1982 period. There is no obvious demand shock associated with the 1990-1 recession; there is for the 2001 recession, however. A possible interpretation is that both the observed demand and technology shocks capture some elements of the decline in investment after the end of the 1990's boom.

Looking at monetary shocks, perhaps the best general conclusion is that a recession is generally associated with a lagged monetary shock; this is apparent for the 1974-1975, 1980 and 1981-2 recessions, as well as the 2001 recession, before which the Greenspan Fed was often taken to task for excessively tight monetary policy in response to the boom. However, shocks of the scale seen before the 2001 recession occurred without incident in the 1980s and 1990s, and no such shock is obvious before 1970 or 1990-1.

As a test of a DSGE model's ability to account for business cycles, Dejong et al. (2000b) suggested estimating a logit model of the probability of NBER recessions (explained variable equals one if a quarter was within an NBER recession, 0 otherwise), using as explanatory variables the current values and values lagged up to four quarters of the standardized levels of each shock series. The estimated probabilities of NBER recessions from the various logit models are reported in Figure 3.3. For each model, results are reported from the model with a lag length that maximized the Schwarz criterion; the motivation for using this criterion will become clearer below. The best model using all the shock series does a good job of picking out recessions; all the genuine recessions are predicted, and if we take a predicted probability of 0.5 as predicting a recession, only one "false recession" is predicted, in the mid 1990's.

More informative regarding the sources of recessions is examining how well each shock series does individually in predicting recessions. The performance of logit models using only demand shocks and/or monetary shocks is not promising. The pure demand shock model picks out the 1974-5, 1980 and 1981-2 recessions, all three of which were associated with disinflation; however, it cannot pick out the non-inflationary recessions of the 1990's. The pure monetary shock model is not much better; it picks out the 1979-82 recessions, associated with the large variations in interest rates during the Volcker period, as well as the 1970 and 2001 recessions, but misses the much larger 1974-5 recession, and is prone to predicting false recessions, most obviously in 1995. By contrast, a model using only technology shocks does much better; it picks all the post-1965 recessions (with only one possible "false alarm" in the late 1970's), and the false 1995 recession disappears.

Contrasting that "supply-side" model to a model with a demand side model (including only demand and monetary shocks) is instructive. During the 1970's, when volatile monetary policy and high and volatile inflation prevailed, a "demand-side" model of recessions does about as well as a supply-side one with technology shocks. However, in the 1980's and 1990's, the "demand-side" model's performance deteriorates, while the "supply-side" model does not. To the extent that aggregate demand shocks can generate business cycles, they clearly have taken a back seat to supply shocks in the 1980's and 1990's. This is not entirely surprising, of course; much of it is due to stricter adherence to systematic policy (and hence smaller variance of deviations from policy) during the tenure of Chairman Alan Greenspan than during the tenures of Chairmen Martin, Burns and Miller. Even for the 1970's, however, the effects of monetary shocks are hard to disentangle from those of supply shocks (i.e. oil shocks).

The last two logit models reported, one with supply and monetary shocks and one with supply and demand shocks, are hard to distinguish visually from the pure supply shock model. To figure out if accounting for monetary shocks adds anything to the pure supply shock model, it would help to calculate posterior odds for the various logit models. Complicating calculation of Bayes factors is a lack of prior information on the values of the logit model parameters. However, Kass and Raftery (1995) argue that the Schwarz criterion can be used as a rough "prior-free" approximation of the average likelihood. So, the Schwarz criteria are reported in Table 3.4 as approximations of the Bayes factors for each logit model.

The probabilities of models omitting supply shocks are near zero, confirming the inadequacy of purely "demand-side" theories of the cycle. As a general rule, models including non-monetary demand shocks are inferior to models without them; monetary shocks appear to be the best candidate for the source of the demand-side shocks that have buffeted the US economy. The "demand shock only" model is particularly bad, unable to beat a naïve model assuming recessions to be equally likely each period.

However, a pure supply shock model dominates all the models including monetary and other demand shocks; the next best model is a model including monetary shocks, but the posterior odds are about 7 to 1 against that model. While this is not a decisive rejection of the importance of monetary shocks in recessions, it is fair to conclude that they have at best been of secondary importance to supply or technology shocks. Such a result is not surprising in the context of much of the previous identified VAR literature on the effects of monetary shocks; Sims (1998) summarizes the literature's findings by stating that the effects of monetary shocks in business cycles is "modest to nil." In the context of the DSGE model, however, a structural interpretation of why this is so is much easier to propose; given the fairly steep IS curve the model assumes, only monetary shocks on the scale observed during the 1979-1982 period are likely to disrupt the real economy appreciably.

Correlations between shocks and time series

An alternative perspective on the role of the various shocks in the cycle is to look at the distributions of the cross-correlations of the shock series with output, inflation and interest rates. The results are reported in detail in Tables 3.5a-c. The signs of relationships between the shock series and the time series can, however, be summed up succinctly in the following table:

	Technology			Demand			Monetary		
	Lags	Current	Leads	Lags	Current	Leads	Lags	Current	Leads
GDP growth	+	+	0	-	+	+	-	-	+
Inflation	-	-	-	+	+	+	0	-	0
Fed. funds	-	-	-	+	+	+	+	+	+

For output growth these correlations largely confirm what eyeballing the shock series and the estimated logit models did. The contemporaneous correlation of the supply shock with output growth is very high (it is no lower than 0.8 and may be as high as 0.95). For the other two shocks it is much weaker. Correlations of output growth with demand shocks range between 0.17 and 0.36 for the demand shocks. The correlations for monetary shocks are if anything even weaker; they are no lower than –0.3 and may well be approximately zero. It is clear supply shocks are the best indicator for fluctuations in economic growth.

For nominal variables such as inflation and interest rates, the situation is more ambiguous, but some general tendencies can be noted. The highest median correlation for inflation is with demand shocks, between 0.3 and 0.5. Next highest in absolute value are supply shocks, fairly precisely estimated at around -0.22 to -0.27. Weakest, again, were monetary shocks, with correlations ranging from about zero to at most -0.25. The evidence is quite weak for a role in monetary shocks (as opposed to systematic policy) having had a great impact on inflation *or* output in the post-1965 period. Even for inflation the case for supply shocks (i.e. oil and other supply shocks) being more important for inflation fluctuations overall is more persuasive. Inflation was not obviously "always and everywhere a monetary phenomenon" even in the 1970s; monetary shocks were very large, but so were oil shocks.

The correlation of the rate with demand shocks and monetary shocks is unambiguously positive. However, the correlation with supply shocks is negative, suggesting that nominal interest rates fall in response to a supply shock. This result conflicts with the impulse response results reported above, which suggest interest rates should rise. One interpretation is that the positive effects on the target rate from a rise in output from the shock are swamped by the negative effect on inflation, with which the federal funds rate is strongly correlated.

Looking now at leads and lags of the shock series, lagged supply is positively correlated with output growth, but leads of supply are essentially uncorrelated with output. However, inflation and nominal interest rates are negatively correlated with supply at all leads and lags. Since supply shocks are assumed to be white noise, the correlations of the leads of the shocks with endogenous variables should all be zero. So, to the extent that the leads deviate from zero, they suggest potential sources of misspecification. The supply shock series is highly enough correlated with output growth to give the negative correlation with inflation and interest rates that we saw in the BVAR.

The correlations of lagged demand shocks are better behaved: lagged demand is negatively correlated with output growth, as output peaks in the period of the shock then falls back towards trend; correlations of lagged demand with inflation and interest rates are positive. However, while leads of demand shocks are uncorrelated with the federal funds rate, with confidence intervals all nesting zero, leads of demand shocks are positively correlated with current output growth, and, more weakly, with inflation. Lagged monetary shocks are negatively correlated with output growth, and essentially uncorrelated with inflation, and, unsurprisingly, positively correlated with interest rates. Leads of monetary shocks are essentially uncorrelated with inflation and interest rates; more problematic is the unambiguously positive correlation of leads of monetary shocks with output growth, that is, of current monetary shocks with lagged output growth. This suggests that the monetary shock series is still picking up policy responses to output growth not captured in the monetary rule. Steps should be taken in future work to minimize this problem and improve the accuracy of monetary policy shock estimates.

Conclusions

To sum up, then, there is little evidence, from the results reported above, that the DSGE model's estimated monetary policy shocks have had much information content for the real business cycle, at least since the 1980's. The estimated logit models for NBER recessions do not indicate that adding monetary shocks adds much to a model estimated using just supply shocks in explaining recessions. Also, the correlation between the model's estimate of monetary shocks with inflation as well as output growth is fairly weak. There is also little evidence that other, non-monetary demand-side shocks have much impact on the real business cycle either. Given the best evidence on price stickiness and intertemporal substitution of consumption, the chief culprits for US business cycles, contrary to Smets and Wouters' results for Europe, appear to be supply and technology shocks. What is apparent for now is that the degree of culpability for business cycles that a Bayesian DSGE model assigns to supply shocks versus monetary shocks is sensitive to its priors on the structural parameters, as these determine how our models disaggregate disturbances into various shocks. Our ability to learn the causes of recessions from Bayesian DSGE models depends on how well we specify our prior knowledge about the structure of the economy, and, by extension, how well that prior knowledge can be refined. Let that stand as the moral to this whole tangled tale.

The results are naturally conditional on the accuracy of the shock series, and their refinement is certainly possible. Information from other measures of output, inflation and interest rates (differentiated by measurement error) could be mined for information on the timing of monetary, supply, and other shocks. DSGE models could give a structural interpretation to the dynamic factors estimated from big data sets that have proven useful for forecasting business cycles (Stock and Watson 2002a, 2002b). Unfortunately, such a study must be left to future research, that is, to another day.

Slán agat agus beannacht Dé leat.

Table 3.1a: Prior distributions for structural parameters of DSGE model

Parameter	Distribution	Range	Mean	Standard deviation
	Long-run trend parameters			
$\ln \gamma$	Normal	$(-\infty,\infty)$	0.005	0.005
$\ln \pi$	Normal	$(-\infty,\infty)$	0.010	0.0075
	Preference parameters			
$\bar{\beta}$	Beta	[0,1]	0.99875	0.000625
ϕ	Uniform	[0,1]	0.500	0.289
$1/\sigma$	Gamma	$[0,\infty)$	0.010	0.010
$\bar{h}$	Beta	[0,1]	0.300	0.050
ρ_ε	Uniform	[0,1]	0.500	0.289
	Nominal/real rigidity parameters			
μ-1	Gamma	$[0,\infty)$	0.200	0.100
$\alpha/(1-\alpha)$	Gamma	$[0,\infty)$	3.000	1.000
	Monetary policy parameters			
ρ_r	Uniform	[0,1]	0.500	0.289
r_y	Gamma	$[0,\infty)$	0.125	0.050
r_π	Gamma	$[0,\infty)$	1.500	0.250
$r_{\Delta y}$	Gamma	$[0,\infty)$	0.125	0.050
$r_{\Delta \pi}$	Gamma	$[0,\infty)$	0.500	0.250
	Shock variance parameters			
			s	v
σ_z	Inverse gamma	$[0,\infty)$	1.00	2
σ_ε	Inverse gamma	$[0,\infty)$	3.00	2
σ_r	Inverse gamma	$[0,\infty)$	0.25	2

Table 3.1b: Posterior distributions for structural parameters of DSGE model

Parameter		Quantile	
	0.025	0.50	0.975
$\ln \gamma$	0.00324	0.00522	0.00716
$\ln \pi$	0.00656	0.0101	0.00140
$\bar{\beta}$	0.99752	0.99886	0.999649
ϕ	0.199	0.333	0.474
$1/\sigma$	0.00333	0.0105	0.0227
$\bar{h}$	0.195	0.290	0.387
ρ_{ε}	0.864	0.919	0.966
μ	1.051	1.163	1.420
$1/(1-\alpha)$	3.423	4.532	6.058
ρ_r	0.832	0.887	0.927
r_y	0.0511	0.126	0.271
r_{π}	1.194	1.585	2.073
$r_{\Delta y}$	0.0843	0.134	0.201
$r_{\Delta \pi}$	0.269	0.476	0.720
σ_z	1.025	1.237	1.536
σ_{ε}	3.089	4.908	8.853
σ_r	0.255	0.293	0.343

Table 3.2: Posterior odds for competing models

Model	DSGE	BVAR(1)	BVAR(2)	BVAR(3)	BVAR(4)
Prior probability	1/5	1/5	1/5	1/5	1/5
Approx. Bayes factor (Modified harmonic mean)	1873.01	1869.47	1868.65	1870.20	1871.26
Posterior probability	0.78	0.02	0.01	0.05	0.14

Table 3.3a: Cross-correlations of GDP growth and inflation

Model	Quantile	$i = -3$	$i = -2$	$i = -1$	$i = 0$	$i = 1$	$i = 2$	$i = 3$
	0.025	-0.165	-0.185	-0.197	-0.213	-0.136	-0.104	-0.086
DSGE	0.50	-0.095	-0.110	-0.122	-0.124	-0.058	-0.034	-0.024
	0.975	-0.047	-0.055	-0.065	-0.050	-0.002	0.015	0.019
		$L_{\chi^2} \mid \text{BVAR} = 5.60, p = 0.412$						
	0.025	-0.583	-0.584	-0.577	-0.571	-0.560	-0.548	-0.525
BVAR(4)	0.50	-0.210	-0.226	-0.218	-0.210	-0.180	-0.156	-0.117
	0.975	0.015	0.005	0.011	0.019	0.045	0.069	0.106
		$L_{\chi^2} \mid \text{DSGE} = 1077, p \approx 1$						

Table 3.3b: Cross-correlations of GDP growth and federal funds rate

Model	Quantile	$i = -3$	$i = -2$	$i = -1$	$i = 0$	$i = 1$	$i = 2$	$i = 3$
	0.025	-0.066	-0.066	-0.048	-0.010	0.041	0.047	0.044
DSGE	0.50	-0.029	-0.027	-0.014	0.064	0.098	0.104	0.098
	0.975	-0.007	-0.005	0.013	0.155	0.169	0.177	0.168
		$L_{\chi^2} \mid \text{BVAR} = 16.9, p = 0.982$						
	0.025	-0.606	-0.621	-0.623	-0.587	-0.560	-0.545	-0.533
BVAR(4)	0.50	-0.293	-0.329	-0.338	-0.255	-0.179	-0.139	-0.115
	0.975	-0.099	-0.133	-0.142	-0.060	0.020	0.059	0.086
		$L_{\chi^2} \mid \text{DSGE} = 1394, p \approx 1$						

Table 3.3c: Cross-correlations of inflation and the federal funds rate

Model	Quantile	$i = -3$	$i = -2$	$i = -1$	$i = 0$	$i = 1$	$i = 2$	$i = 3$
	0.025	0.123	0.158	0.197	0.238	0.237	0.232	0.222
DSGE	0.50	0.374	0.416	0.462	0.510	0.519	0.518	0.510
	0.975	0.686	0.713	0.744	0.776	0.785	0.788	0.785
		$L_{\chi^2} \mid \text{BVAR} = 2.63, p = 0.083$						
	0.025	0.084	0.135	0.178	0.216	0.220	0.227	0.217
BVAR(4)	0.50	0.505	0.552	0.598	0.639	0.653	0.660	0.658
	0.975	0.893	0.903	0.913	0.924	0.929	0.931	0.931
		$L_{\chi^2} \mid \text{DSGE} = 71.1, p \approx 1$						

Table 3.4: Schwarz criteria and posterior odds for estimated logit models.

Model	Naive	T	D	M	TD	TM	TDM	DM
Lag length	NA	2	3	4	2	2	2	3
Prior prob.	1/8	1/8	1/8	1/8	1/8	1/8	1/8	1/8
Schwarz crit.	-73.44	-36.63	-73.52	-68.82	-40.93	-38.45	-41.88	-69.57
Posterior prob.	9E-17	0.846	8E-17	9E-15	0.011	0.138	0.004	4E-15

Code in "Model" row denotes series included: T=technology, D=demand, M=monetary.

Table 3.5a: Cross-correlations from estimated DSGE of supply shocks with endogenous variables

Model	Quant.	i = -3	i = -2	i = -1	i = 0	i = 1	i = 2	i = 3
	0.025	0.085	0.196	0.185	0.797	-0.081	0.072	-0.020
Output growth	0.50	0.095	0.123	0.230	0.903	0.031	0.120	-0.073
	0.975	0.102	0.239	0.262	0.959	0.144	0.162	0.023
Inflation	0.025	-0.114	-0.178	-0.211	-0.278	-0.130	-0.173	-0.132
	0.50	-0.095	-0.162	-0.189	-0.249	-0.077	-0.132	-0.097
	0.975	-0.080	-0.145	-0.166	-0.220	-0.008	-0.077	-0.056
Federal funds	0.025	-0.072	-0.094	-0.134	-0.240	-0.302	-0.260	-0.165
	0.50	-0.048	-0.078	-0.116	-0.215	-0.255	-0.203	-0.109
	0.975	-0.031	-0.064	-0.093	-0.178	-0.192	-0.132	-0.044

Table 3.5b: Cross-correlations from estimated DSGE of demand shocks with endogenous variables

Model	Quantile	i = -3	i = -2	i = -1	i = 0	i = 1	i = 2	i = 3
	0.025	-0.139	-0.264	-0.085	0.172	-0.069	0.050	0.112
Output growth	0.50	-0.108	-0.219	-0.032	0.263	0.021	0.089	0.148
	0.975	-0.079	-0.169	0.016	0.361	0.117	0.124	0.176
	0.025	0.183	0.178	0.243	0.302	-0.087	-0.028	0.021
Inflation	0.50	0.272	0.270	0.337	0.398	0.047	0.088	0.126
	0.975	0.358	0.364	0.437	0.509	0.190	0.213	0.243
	0.025	0.212	0.169	0.143	0.110	-0.099	-0.116	-0.108
Federal funds	0.50	0.318	0.284	0.271	0.240	0.004	-0.016	-0.012
	0.975	0.429	0.408	0.411	0.386	0.132	0.104	0.105

Table 3.5c: Cross-correlations from estimated DSGE of money shocks with endogenous variables

Model	Quantile	i = -3	i = -2	i = -1	i = 0	i = 1	i = 2	i = 3
Output growth	0.025	-0.102	-0.290	-0.209	-0.319	0.104	0.050	-0.049
	0.50	-0.075	-0.239	-0.255	-0.151	0.182	0.104	0.001
	0.975	-0.047	-0.184	-0.158	0.007	0.245	0.151	0.049
	0.025	-0.133	-0.106	-0.156	-0.254	-0.073	-0.075	-0.148
Inflation	0.50	-0.023	0.006	-0.039	-0.126	0.057	0.037	-0.047
	0.975	0.055	0.085	0.047	-0.027	0.155	0.118	0.025
Federal funds	0.025	0.061	0.095	0.189	0.249	-0.007	-0.016	-0.006
	0.50	0.137	0.171	0.267	0.331	0.081	0.067	0.067
	0.975	0.214	0.249	0.345	0.417	0.173	0.152	0.140

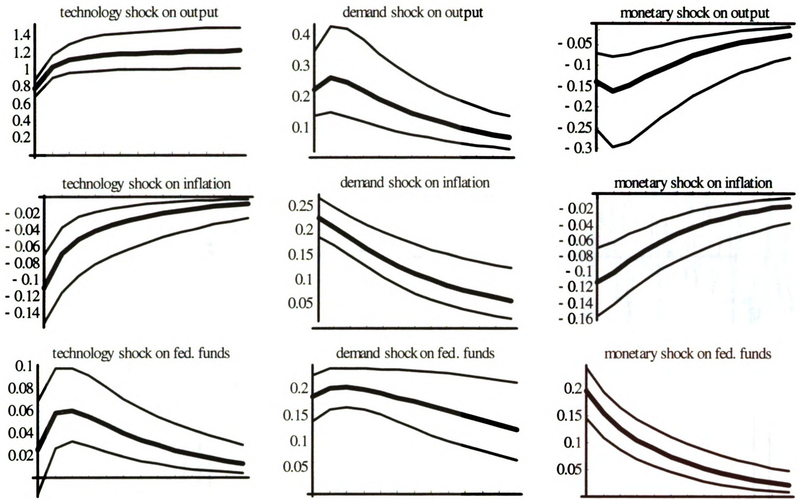

Figure 3.1: 95% confidence bands for impulse response functions from estimated DSGE; thick band in middle indicates 50th percentile

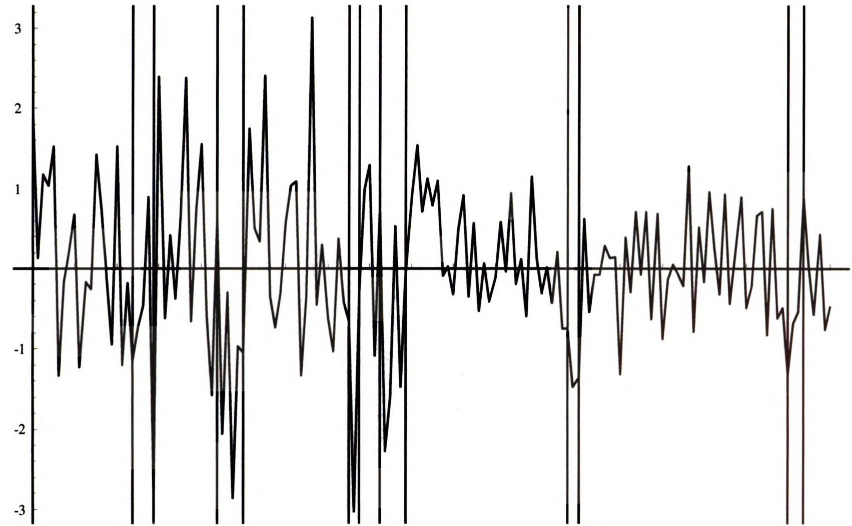

Figure 3.2a: 50th percentile of posterior distribution for technology/supply shocks from estimated DSGE model, 1965:I to 2003:I; vertical lines mark NBER recessions

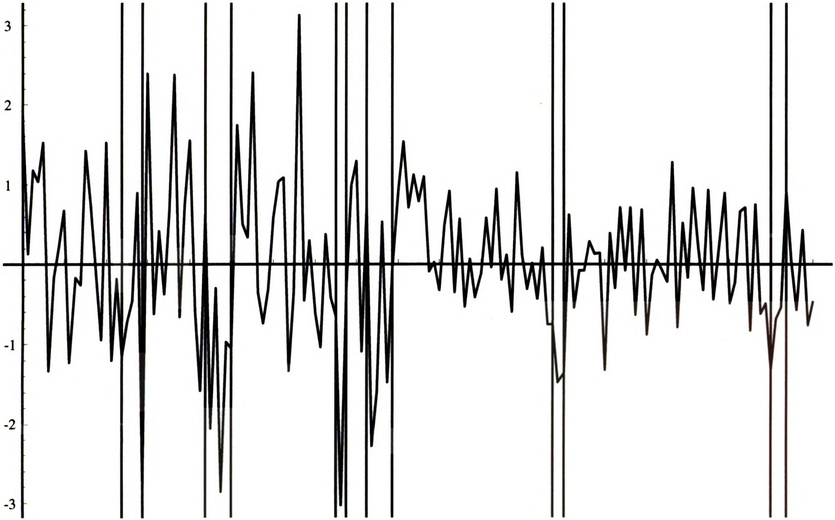

Figure 3.2b: 50th percentile of posterior distribution for demand shocks from estimated DSGE model, 1965:I to 2003:I; vertical lines mark NBER recessions

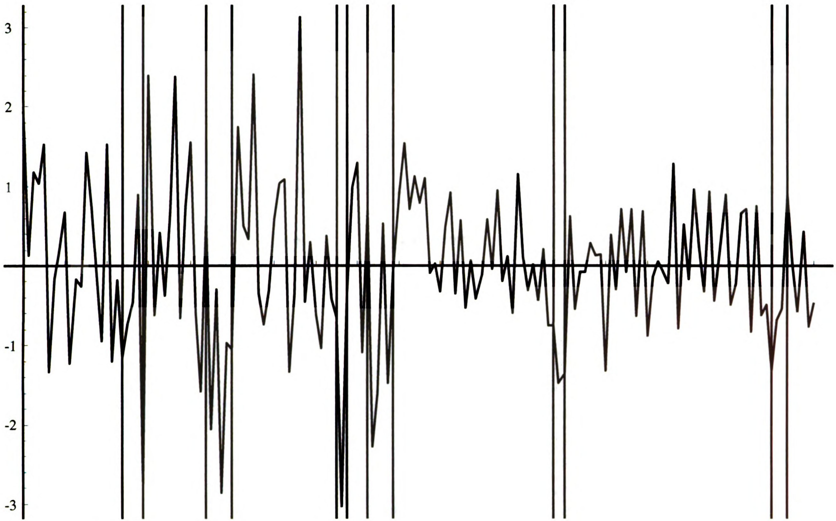

Figure 3.2c: 50th percentile of posterior distribution for monetary shocks from estimated DSGE model, 1965:I to 2003:I; vertical lines mark NBER recessions

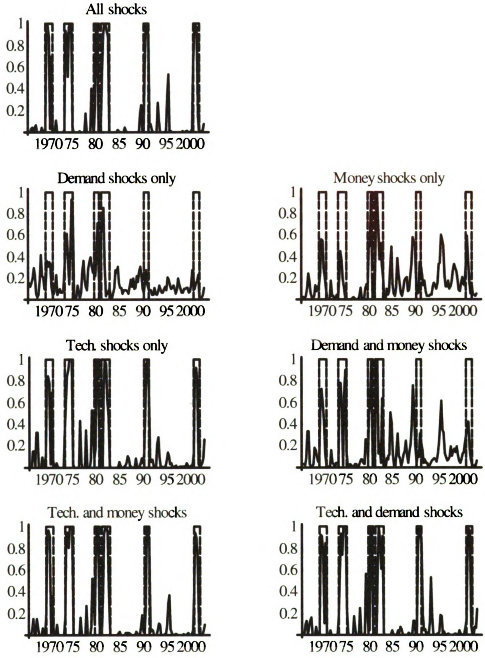

Figure 3.3: Estimated probabilities of NBER recessions from logit models using series of shocks from estimated DSGE model

REFERENCES

Bernanke, B., Mihov, I., 1998. Measuring monetary policy. *Quarterly Journal of Economics* 113, 869-902.

Blinder, A. 1994. On sticky prices: academic theories meet the real world. In *Monetary Policy*. Chicago and London: University of Chicago Press, pp. 117-150.

Calvo, G.A., 1983. Staggered prices in a utility maximizing framework. *Journal of Monetary Economics* 12, 383 –398.

Christiano, L., Eichenbaum, M., 1992. Current real-business-cycle theories and aggregate labor-market fluctuations. *American Economic Review* 82, 430-50.

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

Christiano, L., Eichenbaum, M, Evans, C. 1998. Monetary policy shocks: what have we learned and to what end? National Bureau of Economic Research Working Paper 6400.

Clarida, R. Gali, J. Gertler, M., 2000. Monetary policy rules and macroeconomic stability: evidence and some theory. *Quarterly Journal of Economics* 115, 147-80.

Dejong, D., Ingram, B., Whiteman, C., 2000a. A Bayesian approach to dynamic macroeconomics. *Journal of Econometrics* 98, 203-223.

Dejong, D., Ingram, B. Whiteman, C., 2000b. Keynesian impulses versus Solow residuals: identifying sources of business cycle fluctuations. *Journal of Applied Econometrics* 15, 311-329.

del Negro, M., Schorfheide, F., 2003. Take your model bowling: forecasting with general equilibrium models. Federal Reserve Bank of Atlanta Economic Review 88, 35-50.

Diebold, F.X., Mariano, R, 1995. Evaluating forecast performance. *Journal of Business and Economic Statistics* 13, 253-263.

Doan, T., Litterman, R., Sims, C., 1984. Forecasting and conditional projection using realistic prior distributions. *Econometric Reviews* 3, 1-100.

Fuhrer, J., 2000. Habit formation in consumption and its implications for monetary policy models. *American Economic Review* 90, 367-90.

Fuhrer, J., Moore, G. 1995. Inflation persistence. *Quarterly Journal of Economics* 110, 127-159.

Gali, J., 1992. How well does the IS-LM model fit postwar US data? *Quarterly Journal of Economics* 107, 709-738.

Gali, J. Gertler, M., 1999. Inflation dynamics: a structural econometric analysis. *Journal of Monetary Economics* 44, 195-222.

Gali, J., Gertler, M., Lopez-Salido, J., 2001. European inflation dynamics. *European Economic Review* 45, 1237-1270.

Geweke, J., 1999. Using simulation methods for Bayesian econometric models: inference, development and communication. *Econometric Reviews* 18, 1-73.

Greenwood, J., Hercowitz, Z., Huffman, G., 1988. Investment, capacity utilization and the real business cycle. *American Economic Review* 78, 402-17.

Hall, R., 1988. Intertemporal substitution in consumption. *Journal of Political Economy* 96, 339-357.

Hamilton, J., 1994. *Time Series Analysis*. Princeton University Press, Princeton, New Jersey.

Hansen, G., Prescott, E., 1993. Did technology shocks cause the 1990-1991 recession? *American Economic Review* 83, 280-86.

Ingram, B., Whiteman, C., 1994. Supplanting the 'Minnesota' prior: Forecasting macroeconomic time series using real business cycle priors. *Journal of Monetary Economics* 34, 497-510.

Judge, G., Griffiths, W., Hill, R., Lee, T., 1985. *The Theory and Practice of Econometrics*. 2nd ed. John Wiley & Sons, New York.

Kadiyala, K., Karlsson, S., 1997. Numerical methods for estimation and inference in Bayesian VAR models. *Journal of Applied Econometrics* 12, 99-132.

Kass, R., Raftery, A., 1995. Bayes factors. *Journal of the American Statistical Association* 90, 773-795.

Kim, J., 2000. Constructing and estimating a realistic optimizing model of monetary policy. *Journal of Monetary Economics* 45, 329-359

King, R. Wolman, A., 1996. Inflation targeting in a St. Louis model of the 21st century. Federal Reserve Bank of St. Louis *Review* 78, 83-107.

Kurmann, A., 2001. Can new Keynesian pricing explain US inflation dynamics? Mimeo, University of Virginia.

Kydland, F. Prescott, E., 1982. Time to build and aggregate fluctuations. *Econometrica* 50, 1345-1370.

Leeper, E., Sims, C., Zha, T., 1996. What does monetary policy do? *Brooking Papers on Economic Activity* 2, 1-63.

Leeper, E., Zha, T., 2001. Assessing simple policy rules: a view from a complete macroeconomic model. *Federal Reserve Bank of St Louis Review* 83, pp. 83-110.

Litterman, R. 1986. Forecasting with Bayesian vector autoregressions—five years of experience. *Journal of Business and Economic Statistics* 4, pp. 25-38.

McCallum, B., Nelson, E., 1999a. Performance of operational policy rules in an estimated semiclassical structural model. In: J.B. Taylor, ed., *Monetary Policy Rules*. University of Chicago Press, Chicago, pp. 15-45.

McCallum, B., Nelson, E., 1999b. An optimizing IS-LM specification for monetary policy and business cycle analysis. *Journal of Money, Credit, and Banking* 31, 296-316

McGrattan, E., 1994. The macroeconomic effects of distortionary taxation. *Journal of Monetary Economics* 33, 573-601.

McGrattan, E., Rogerson, R., Wright, R., 1997. An equilibrium model of the business cycle with household production and fiscal policy. *International Economic Review* 38, 267-290.

Mankiw, N. Rotemberg, J. Summers, L., 1985. Intertemporal substitution in macroeconomics. *Quarterly Journal of Economics* 100, 225-51.

Rigobon, R. 2003. Identification through heteroskedasticity. *Review of Economics and Statistics* 85, 777-92.

Roberts, J., 1995. New Keynesian economics and the Phillips curve. *Journal of Money, Credit and Banking* 27, 975-984.

Robertson, J, Tallman, E., 1999. Vector autoregressions: forecasting and reality. Federal Reserve Bank of Atlanta *Economic Review* 84, 4-18.

Rotemberg, J., 1982. Sticky prices in the United States. *Journal of Political Economy* 90. 1187-1211.

Rotemberg, J., 1987. The new Keynesian microfoundations. In: S. Fischer, ed., *NBER Macroeconomics Annual*: 1987. MIT Press, Cambridge, Mass., pp. 69-104.

Rotemberg, J., Woodford, M., 1995. Dynamic general equilibrium models with imperfectly competitive product markets. In T.F. Cooley, ed., *Frontiers of Business Cycle Research*. Princeton University Press, Princeton, New Jersey, pp. 243-293.

Rudebusch, G., 1998. Do measures of monetary policy in a VAR make sense? *International Economic Review* 39, 943-948.

Sbordone, A., 2002. Prices and unit labor costs: a new test of price stickiness. *Journal of Monetary Economics* 49, 265-292.

Schorfheide, F. 2000. Loss-function based evaluation of DSGE models. *Journal of Applied Econometrics* 15, 645-670.

Sims, C. 1998. Comment on Glenn Rudebusch's "Do measures of monetary policy in a VAR make sense?" *International Economic Review* 39, 933-941.

Smets, F, Wouters, R., 2002. An estimated stochastic general equilibrium model of the euro area. European Central Bank Working Paper 171.

Stock, J., Watson, M., 1999. Forecasting inflation. *Journal of Monetary Economics* 44, 293-335.

Stock, J. Watson, M., 2002a. Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association* 97, 1167-79.

Stock, J. Watson, M., 2002b. Macroeconomic forecasting using diffusion indexes. *Journal of Business and Economic Statistics* 20, 147-62.

Taylor, J., 1993. Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy* 39, 195-214.

Theil, H. Goldberger, A., 1961. On pure and mixed estimation in economics. *International Economic Review* 2, 65-78.

Uhlig, H., 1999. What are the effects of monetary policy on output? Results from an agnostic identification procedure. Centre for Economic Policy Research Discussion Paper 2137.

Yun, T., 1996. Nominal price rigidity, money supply endogeneity, and business cycles. *Journal of Monetary Economics* 37, 345-70.

Zellner, A., 1971. *An introduction to Bayesian inference in econometrics.* John Wiley and Sons, New York.