Title
Three essays on macroeconomics

Author
Pruitt, Seth James

Publication Date
2008

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Three Essays on Macroeconomics

A Dissertation submitted in partial satisfaction of the
Requirements for the degree Doctor of Philosophy

in

Economics

by

Seth James Pruitt

Committee in charge:

Professor James D. Hamilton, Chair
Professor Valerie A. Ramey
Professor Garey Ramey
Professor Allan Timmermann
Professor Patrick J. Fitzsimmons
Assistant Professor Nir Jaimovich

2008

The dissertation of Seth James Pruitt is approved, and it is acceptable in quality and form for publication on microfilm:

Chair

University of California, San Diego

2008

Dedicated to Courtney.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS

I would like to acknowledge Jim Hamilton for his guidance and insightful advice. Additionally, Nir Jaimovich gave generously of his valuable time and support, and Valerie Ramey gave wonderful guidance to the formulation of this analysis.

Chapter 2, in full, is joint work with Nir Jaimovich and Henry E. Siu. The dissertation author was a primary author of this paper.

Chapter 3, in full, is joint work with Max Floetotto and Nir Jaimovich. The dissertation author was a primary author of this paper.

VITA

2002	Bachelor of Arts, University of Arizona
2003-2008	Teaching Assistant, Department of Economics, University of California, San Diego
2005	Master of Arts, University of California, San Diego
2005-2007	Research Assistant, Department of Economics, University of California, San Diego
2008	Doctor of Philosophy, University of California, San Diego

FIELDS OF STUDY

Major field: Macroeconomics

Secondary fields: Econometrics, forecasting

ABSTRACT OF THE DISSERTATION

Three Essays on Macroeconomics

by

Seth James Pruitt

Doctor of Philosophy

University of California, San Diego

2008

Professor James D. Hamilton, Chair

The dissertation is comprised of three chapters, each a free-standing paper.

The first chapter addresses the issue of data revision: the fact that many major macroeconomic data series are revised over time. I argue that decision-makers take this reality, and the resulting data uncertainty, into account. Therefore models of these agents must take account of the measurable data uncertainty surrounding their choices.

The second chapter notes that the United States' labor force can be demographically partitioned into groups with different cyclical quantity and price characteristics. We argue that the main determinant of these differences comes from the supply side of the economy, i.e. the production function.

The third chapter argues that investment-specific technological change, as measured in recent literature, can be contaminated by endogenous firm entry and exit. We describe a two-sector model of the economy, calibrate it to the United States, and analyze its performance in explaining the dynamics of consumption, investment, and labor input used for the production of both.

Chapter 1:
When Data Revisions Matter to Macroeconomics

Abstract

This paper finds important effects in two different models from assuming that optimizing agents have "data uncertainty." Data uncertainty results from the fact that many macroeconomic time series undergo revision and therefore data releases should be treated as predictions that may contain error. In the first model, the paper builds on the framework of Sargent, Williams, and Zha (2006). I show that ignoring data uncertainty leads to an overestimate of the Federal Reserve's model uncertainty and predicts Fed unemployment rate forecasts that are very different from Greenbook forecast evidence. Using a Bayesian Markov Chain Monte Carlo algorithm to estimate the nonlinear model's Extended Kalman filter approximation, I explicitly model the Fed's data uncertainty and show how this approach significantly mitigates those problems. In the second model, abstracting from data uncertainty leads to a monotone impulse response function that is at odds with VAR estimates I obtain from U.S. data. I show that a hump-shaped labor response matching those VAR estimates is the effect of assuming firms (1) can only observe real-time data and (2) recognize this informational limitation.

Contents

List of Tables

List of Figures

1.1 Introduction

This paper finds important effects in two different models from assuming that optimizing agents have "data uncertainty." Data uncertainty results from the fact that many macroeconomic time series undergo revision and therefore data releases should be treated as predictions that may contain error. The two models are chosen in order to show that the effects of data uncertainty are not merely artifacts of stylized assumptions. To this end, their differences are an asset. With the Federal Reserve model I show that ignoring data uncertainty necessarily amplifies estimates of the volatility of time-varying parameters, and can do so with striking results. On the other hand, with the production model I show that data uncertainty and real-time data form an informational friction that creates propagation in otherwise frictionless models. From these results I argue that data revisions matter to macroeconomics when agents hold model uncertainty or when the data display a hump-shaped response to real shocks.

The first model builds on Sargent, Williams, and Zha (2006) which depicts the Federal Reserve as optimally setting inflation in light of time-invariant unemployment and inflation targets. The Fed is uncertain of its economic model and thus lets it evolve in response to new data.[1.1] The Sargent, Williams, and Zha (2006) model very accurately explains the history of American inflation as optimal policy given changing estimates of the Philips curve. However, it suffers from a key problematic implication:

[1.1] Model uncertainty in this case is represented by the variance of shocks to time-varying parameters.

that the Fed's unemployment rate forecasts, the basis for setting inflation, are very inaccurate and significantly different than Greenbook forecasts. Accompanying these is an estimate of the Fed's model uncertainty that is biased upward. I change the model by allowing the Fed to account for multiple vintages of real-time data and to recognize that data is revised. By accounting for data uncertainty, the model explains inflation without the fore mentioned problems. The Fed is less uncertain of its model and makes accurate unemployment forecasts that resemble Greenbook forecasts. This supports the main story that inflation went high during the 70s and 80s because the Fed believed in and acted upon an evolving Philips curve trade-off.

The second application models output as a function of productivity and labor. The representative firm's new data vintages represent increasingly precise estimates of past productivity. I use U.S. output and labor data series to calibrate different versions of the model. When the firm sees true productivity, the impulse response function of labor is monotone, as is typical in benchmark RBC models. However, when the firm recognizes data uncertainty and observes real-time data, the model predicts a hump-shape similar to that estimated from U.S. data.

This paper joins previous literature highlighting the impact of real-time data on economic behavior. Oh and Waldman (1990) used data revisions to identify the effects of real-time macro announcements on future economic activity, highlighting how data mismeasurement itself can influence agents' behavior. Orphanides (2001) demonstrated that policy analysis using real-time preliminary data delivers significantly different conclusions than anachronistically using the current data vintage.[1.2] Ghysels, Swanson, and Callan (2002) extended this idea by using multiple data vintages to fit policy rules, some of which are adaptively estimated, forecasting the Fed Funds rate. These papers show that agents' behavior is connected to the real-time data they actually saw when making decisions, and Aruoba (2004) found welfare consequences associated with a one-period signal extraction problem motivated by the ensuing data uncertainty. Multiple vintages of real-time data are publicly available in the St. Louis Fed's ArchivaL Federal Reserve Economic Data (ALFRED) collection and the Philadelphia Fed's Real-Time Data Set for Macroeconomists pioneered by Croushore and Stark (2001). My paper's contribution is to show the impact of assuming that agents not only use real-time data, but take the existence of their revisions into account when optimizing.

Howrey (1978) was among the first to suggest a linear state space framework for forecasting in the presence of data revisions. Springing from this idea, I model agents as using a state space model to account for data uncertainty by observing multiple vintages of real-time data. They are then able to predict revisions, as Aruoba (2005) finds is possible for numerous macroeconomic data series. Since as far back as Burns and Mitchell, who revised their business cycle indicators as underlying data were revised, or Zellner (1958), who admonished readers to be "careful" with "provisional" data, the existence of revisions has been known. A series of recent papers (see Cunningham, Jeffery, Kapetanios, and Labhard (2007) and references therein)

[1.2] Runkle (1998) made a similar point around the same time.

by central bank researchers are evidence that macroeconomic data revisions remain significant to policy decisions even today.

Explicitly modeling data uncertainty is similar in spirit to the macroeconomic learning literature where the main idea is that agents are themselves econometricians. Whereas much of that literature focuses on the problem of learning parameters – for instance see Evans and Honkapohja (2001) or Orphanides and Williams (2006) – I draw attention to the continual problem of learning the true economic state as measured by real-time data. I show that data uncertainty (learning the true economic state) remains a key to understanding agent behavior even in the presence of traditional macroeconomic learning (learning about parameters).

The paper is organized as follows. Section 1.2 briefly describes how I model data uncertainty. Section 1.3 presents the Federal Reserve model and highlights the important differences created by accounting for data uncertainty. Section 1.4 presents the production model calibrated to U.S. data and conducts several simulation experiments highlighting the impact of real-time data and data uncertainty. Section 1.5 concludes. Technical discussion is put in Appendix 1.A.

1.2 Including Data Uncertainty

Consider the civilian unemployment rate u_t. A portion of its actual vintage data array is shown in Table 1.1.[1.3]

[1.3] For expositional clarity, I have shifted the vintage dates (column labels) from their actual dates, which are usually around the first week of the month following the observation date (row labels). For instance, the unemployment rate for November 1964 was actually announced on December 4th,

Table 1.1: Example Vintage Data Array

	Nov-1964	Dec-1964	Jan-1965
Nov-1964	5	5	4.9
Dec-1964		4.9	5
Jan-1965			4.8

Note: Civilian unemployment rate, from ALFRED. Vintage dates are shifted for clarity.

The preliminary measurement of the unemployment rate for Nov-1964 is known 0 periods after the fact; I call this value $u^0_{Nov-1964}$ and it equals 5. I denote the measurement of the Nov-1964 unemployment rate *one month later* as $u^1_{Dec-1964}$ since this datum is observed in Dec-1964 and pertains to 1 period previous.[1.4] The measurement $u^2_{Jan-1965}$ is 0.1 less than $u^1_{Dec-1964}$, meaning that a revision of –0.1 occurred: I label this $\nu^2_{Jan-1965}$. Similarly there is a data revision for the Dec-1964 unemployment rate revealed in Jan-1965

$$u^1_{Jan-1965} - u^0_{Dec-1964} = 5 - 4.9 = 0.1 = \nu^1_{Jan-1965}$$

Looking at Table 1.1, the rightmost column $(4.9, 5, 4.8)'$ is the Jan-1965 *vintage* of data. The main diagonal $(5, 4.9, 4.8)'$ is the *preliminary* data, or the *0 period horizon* real-time data.

In general,

$$u^0_t + \sum_{k=1}^{f} \nu^k_{t+k} = u_t$$

where $f \leq \infty$. This says that there are f revisions to data for the unemployment

1964 – I call the vintage period of this announcement November 1964.

[1.4] An alternative notation has been suggested to me wherein the Dec-1964 report of the Nov-1964 unemployment rate is denoted $u^1_{Nov-1964}$. I find it easier to keep track of data measurements using the notation here.

rate at time t. In practice, it is often necessary to suppose $f < \infty$. We see from the above example that these revisions can take the value 0.

Using the identity

$$u_{t+1}^1 = u_t^0 + \nu_{t+1}^1$$

it follows that

$$u_{t+j}^j = u_t - \sum_{k=j+1}^{f} \nu_{t+k}^k$$

In words, the j-period horizon real-time data for time t is equal to the true unemployment rate at time t minus revisions made more than j periods after the fact. This means that a measurement vector observed in time t is

$$\begin{pmatrix} u_t^0 \\ u_t^1 \\ \vdots \\ u_t^{f-1} \\ u_t^f \end{pmatrix} = \begin{pmatrix} u_t - \nu_{t+1}^1 - \nu_{t+2}^2 - \nu_{t+3}^3 - \cdots - \nu_{t+f}^f \\ u_{t-1} - \nu_{t+1}^2 - \nu_{t+2}^3 - \cdots - \nu_{t+f-1}^f \\ \vdots \\ u_{t-f+1} - \nu_t^f \\ u_{t-f} \end{pmatrix}$$

It might be that the time t unemployment rate has revisions that relate to one another. That is, there may be a relationship between the revisions to the time t unemployment rate that are revealed at $t+1, t+2, \ldots$. This is accomplished through the covariance matrix of the time t revision shock vector $\tilde{\boldsymbol{\nu}}_t$, described below. On the other hand, Aruoba (2005) found statistically significant dynamics between revisions of the same horizon. This is accomplished by autoregression.

To be clear about these forms of dependence, assume there are only two possible revisions. The equations for these revisions are

$$\begin{aligned} \nu_{t+1}^1 &= \mu_{\nu^1} + \rho_{\nu^1} \nu_t^1 + \tilde{\nu}_t^1 \\ \nu_{t+2}^2 &= \mu_{\nu^2} + \rho_{\nu^2} \nu_{t+1}^2 + \tilde{\nu}_t^2 \end{aligned}$$

The parameters $\rho_{\nu^1}, \rho_{\nu^2}$ govern the dependence between revisions of the same horizon. The covariance matrix of the time t revision shock vector $\tilde{\boldsymbol{\nu}}_t \equiv \begin{pmatrix} \tilde{\nu}_t^1 \\ \tilde{\nu}_t^2 \end{pmatrix}$ determines the relationship between revisions to the t period variable.[1.5] These structures are embedded into the models in Sections 1.3 and 1.4.

Looking ahead, I will be modeling revisions to the civilian unemployment rate, nominal GNP, and non-farm payrolls. In order to get a rough idea of the relative magnitude of revisions to these series, I calculate summary statistics of the *relative final revision*, expressed as a percentage, which I now define.

Take the final vintage value to be the data as it is reported at T=Mar-2007 or T=2007:I depending on whether the frequency of the data is monthly or quarterly. The relative final revision, as a percentage, for time t (a monthly or quarterly frequency index, depending on the series) is defined as the difference between the final data vintage value and the preliminary data value divided by the final data vintage value.

$$100 \times \frac{x_T^{T-t} - x_t^0}{x_T^{T-t}}$$

Summary statistics for relative final revisions, expressed as a percentage, are in Table 1.2. These values tell us how real-time preliminary data differs from what we now believe the data to be, as a percentage of the latter. Preliminary data on nominal GNP is biased downwards by 3.8%, and has been as much as 9.6% too low or 2.8% too high. The preliminary data on the unemployment rate is essentially unbiased, but is usually off by about 2.1% and has been as much as 8.1% too high or 5.4% too

[1.5] This can be generalized, as Appendix 1.A.1 explains.

Table 1.2: Summary Statistics: Relative Final Revisions

Variable	Mean	StDev	RMSE	Min	Max
Civilian Unemployment Rate	0.0222	2.0548	2.0527	–8.1081	5.4054
Nominal GNP	3.8042	2.4961	4.5471	–2.8514	9.5806
Non-farm Payrolls	0.5255	0.9472	1.0826	–1.0814	3.9851

Note: Relative final revision as a percentage; for data revisions series in this paper. RMSE is Root Mean Squared Error, viewing the relative final revision as a prediction error of the final vintage value.

low. Non-farm payroll preliminary data is biased downwards by 0.5% but is never off by more than 4% in either direction.

1.3 Federal Reserve Model

Policy-makers base decisions on macroeconomic data. Data revisions alter the statistics forming policy-makers' model of the economy. Hence, the existence of revisions implies that a savvy policy-maker associates some uncertainty to the latest observations of the most recent data vintage. I propose that the policy-maker optimizes accordingly.

Clarida, Gali, and Gertler (2000) provided evidence of different U.S. monetary policy responses over different parts of the postwar era. Boivin (2005) elaborated on this finding by using a time-varying Taylor rule estimated on preliminary data to characterize the response, while not giving a structure behind its evolution. Meanwhile, Orphanides and Williams (2006) and Sims (2006) have drawn attention to the role of model uncertainty in driving policy.

I follow the synthesis of these ideas presented in Sargent, Williams, and Zha (2006) that assumes that the Federal Reserve's economic model can be usefully approximated by a Philips curve with time-varying parameters. By specifying that the Fed believes the parameters follow a random walk we introduce persistent model uncertainty, as discussed in Cooley and Prescott (1976) and Primiceri (2005). Sargent, Williams, and Zha (2006) then explains inflation as the solution to an optimal control problem with a law of motion that changes according to the evolution of filtered parameter estimates. The Federal Reserve's objective remains the same while new data alter its best estimate of the effects of its actions (summarized by the Philips curve trade-off).[1.6] The Sargent, Williams, and Zha (2006) model explains inflation data quite well. However, it also implies that the Fed is extremely uncertain of its structural model and makes very inaccurate unemployment forecasts, at odds with Greenbook evidence, when setting inflation policy.

I emphasize the importance of conditioning an agent's behavior on realistic information. By incorporating multiple vintages of real-time data, the extended model I develop predicts that the Fed is much less uncertain of its model and that the Fed's unemployment forecasts are far more accurate, in line with Greenbook forecasts. The new model's predictions now support the idea that inflation was set high because the Fed believed its forecasted effect on unemployment would be beneficial.

[1.6] Alternatively, Romer and Romer (1990) and Owyang and Ramey (2004) suggest that the time series of American inflation could be explained by changing Fed objectives. This model instead focuses on data-driven changes in Fed beliefs (about how the economy works) and so eliminates the possibility of changes to the Fed's objective function.

1.3.1 Model

To be consistent with the main idea, the Philips curve should always be responsive to new data. If the Philips curve were assumed to be constant, new data would influence parameter estimates less and less over time. An example would be if the curve were estimated by OLS, in which case new data would be weighted by the reciprocal of the increasing number of observations. So instead the Philips curve is assumed to have time-varying parameters. In this way, the Fed is never certain of its model and changes it in the face of new data.

The direct effect of Fed activity is the rate of inflation in the economy – it sets inflation up to some exogenous shock beyond its control. This control shock could be thought of as unpredictable market reaction to Fed policy. The model, to be reasonable, should predict that this exogenous shock is relatively small since it proposes that American inflation is predominantly the result of Federal Reserve policy.

The Fed's inflation control is

$$\pi_t = x_{t-1} + \frac{1}{\zeta_1}\omega_{1t} \tag{1.1}$$

where π_t is the annual inflation rate in time t, x_{t-1} is the part of inflation controllable by the Federal Reserve using information through time $t-1$, and $\omega_{1t} \sim iid(0,1)$ is the exogenous control shock.

The Fed uses a Philips curve to understand the relationship between unemployment and inflation. However, the Fed is always uncertain of its estimated model

– a way of accomplishing this is by assuming the parameters follow a random walk:

$$u_t = \boldsymbol{\alpha}'_{t-1} \begin{pmatrix} \pi_t \\ \pi_{t-1} \\ u_{t-1} \\ \pi_{t-2} \\ u_{t-2} \\ 1 \end{pmatrix} + \frac{1}{\zeta_2}\omega_t \equiv \boldsymbol{\alpha}'_{t-1}\boldsymbol{\Phi}_t + \frac{1}{\zeta_2}\omega_{2t} \quad (1.2)$$

$$\boldsymbol{\alpha}_t = \boldsymbol{\alpha}_{t-1} + \boldsymbol{\Lambda}_t \quad (1.3)$$

where $\omega_{2t} \sim iid(0,1)$ and $\boldsymbol{\Lambda}_t$ is a vector with $\mathbb{E}(\boldsymbol{\Lambda}_t) = 0$, $\mathbb{E}(\boldsymbol{\Lambda}_t\boldsymbol{\Lambda}'_t) = \boldsymbol{V}$, and $\mathbb{E}(\boldsymbol{\Lambda}_t\omega_{2t}) = \mathbf{0}$.

The Fed has inflation and unemployment goals and experiences loss when the actual variables are away from their respective targets. The objective function, which Sargent (1999) calls the Phelps problem, is written

$$\min_{\{x_{t-1+j}\}_{j=0}^{\infty}} \hat{\mathbb{E}}_t \sum_{j=0}^{\infty} \delta^j \Big((\pi_{t+j} - \pi^*)^2 + \lambda(u_{t+j} - u^*)^2 \Big) \quad (1.4)$$

where $\delta \in (0,1)$ is a time discount factor, $\lambda > 0$ gives the Fed's relative weighting of its two objectives, and $\pi^*, u^* \geq 0$ are inflation and unemployment targets. $\hat{\mathbb{E}}$ is expectation with respect to the probability model formed by equations (1.1), (1.2), and (1.3). Because the parameters follow a random walk whose steps are independent of everything else, the Fed's estimate of $\boldsymbol{\alpha}_{t-1}$ is also its estimate of $\boldsymbol{\alpha}_{t+j}$, $\forall j \geq 0$. Hence, the time t solution to the dynamic programming problem {(1.4) s.t. (1.1), (1.2), (1.3)} is found after plugging the time t estimate of $\boldsymbol{\alpha}_{t-1}$ into the law of motion (1.2) for all $j \geq 0$. I set parameters in line with Sargent, Williams, and Zha (2006): $\delta = .9936$, $\lambda = 1$, $i^* = 2$, $u^* = 1$. They note that the results are unaffected by letting u^* be closer to typical "natural unemployment" rates and I have confirmed that this is indeed the case for both the model with and without data uncertainty.

1.3.1.1 Without Data Uncertainty

The Fed estimates the relationships (1.2) and (1.3): given the model without data uncertainty, estimation is a linear filtering problem whose solution is given by the Kalman filter. The notation $\sigma(\cdot)$ denotes the information set (σ-algebra) formed by random variables within the parentheses. Let $\mathbb{E}(\boldsymbol{\alpha}_t|\mathcal{I}_s) \equiv \boldsymbol{a}_{t|s}$ and $\mathbb{V}ar\,(\boldsymbol{\alpha}_t|\mathcal{I}_s) \equiv \boldsymbol{P}_{t|s}$ for $\mathcal{I}_s \equiv \sigma(u_1, \pi_1, \ldots, u_s, \pi_s)$. Given initial conditions $\boldsymbol{a}_{1|0}$ and $\boldsymbol{P}_{1|0}$, the Kalman updating occurs using the formulae:

$$\boldsymbol{a}_{t+1|t} = \boldsymbol{a}_{t|t-1} + \frac{\boldsymbol{P}_{t|t-1}\boldsymbol{\Phi}_t(u_t - \boldsymbol{\Phi}_t'\boldsymbol{a}_{t|t-1})}{(\frac{1}{\zeta_2})^2 + \boldsymbol{\Phi}_t'\boldsymbol{P}_{t|t-1}\boldsymbol{\Phi}_t} \tag{1.5}$$

$$\boldsymbol{P}_{t+1|t} = \boldsymbol{P}_{t|t-1} - \frac{\boldsymbol{P}_{t|t-1}\boldsymbol{\Phi}_t\boldsymbol{\Phi}_t'\boldsymbol{P}_{t|t-1}}{(\frac{1}{\zeta_2})^2 + \boldsymbol{\Phi}_t'\boldsymbol{P}_{t|t-1}\boldsymbol{\Phi}_t} + \boldsymbol{V} \tag{1.6}$$

For the model without data uncertainty, I follow Sargent, Williams, and Zha (2006) in setting $\boldsymbol{a}_{1|0}$ to the value given by a regression on data before 1960. Moreover, the parameter ζ_2 is unidentified in both that paper and in the present model without data uncertainty. Sargent, Williams, and Zha (2006) normalize ζ_2 such that $\frac{1}{\zeta_2}$ is one-tenth the standard deviation of the shock in an additional equation for the "true DGP" of the unemployment rate. Apart from this choice, the "true DGP" equation does not influence the belief-generating mechanism I have described, and thus it is not part of my model. In the interest of making a clear comparison to its results, I take this value from Sargent, Williams, and Zha (2006) and do not estimate it. This does not drive any of the key results below.[1.7]

If we do not allow the Fed to recognize data revisions, then the model is

[1.7] Appendix 1.A.3 explains the reasons for this.

completed by assuming that the Fed observes the true values of both inflation and unemployment each period.

1.3.1.2 Including Data Uncertainty

Graphs of the data show that while the unemployment rate experiences revisions, the inflation rate actually does not.[1.8] Therefore I only model revisions on unemployment and assume 12-month-ended inflation rates are not subject to revision. Additionally, I choose $f = 72$ months in order to capture the idea that the Fed does not want to miss many revisions (only some of which occur 6 years after the fact).[1.9]

To reduce the problem's dimension, every revision does not explicitly enter the model. Instead, for each u_t I assume that revisions are possible one month, two months, three months, and 72 months later. The first revision allows for inference on the preliminary report of unemployment. The middle two revisions allow for inference on the latent value of the two lags of unemployment that enter the policy rule emerging from the dynamic programming problem. The final possible revision acknowledges that data revision happens even after two months. By grouping all revisions past the third into the final revision, we are relabeling the sum $\sum_{k=4}^{f} \nu_{t+k}^{k}$ as simply ν_{t+f}^{f}.[1.10] This assumption greatly reduces the dimension of the state vector from around 2700 to around 160.

[1.8] See Appendix 1.A.2 for these figures.

[1.9] The results presented below are not very sensitive to this. In fact, changing f to 36, which accords with the horizon of historical values included in the Greenbook, delivers very similar results while increasing f does very little. It is only a very small f where most revisions are ignored that affects the estimates.

[1.10] I have experimented with explicitly accounting for each of these revisions and the results do not change, while the computational burden increases substantially.

Therefore the unemployment measurement vector is

$$\begin{bmatrix} u_t^0 \\ u_t^1 \\ u_t^2 \\ u_t^3 \\ u_t^f \end{bmatrix} = \begin{bmatrix} u_t - \nu_{t+f}^f - \nu_{t+3}^3 - \nu_{t+2}^2 - \nu_{t+1}^1 \\ u_{t-1} - \nu_{t-1+f}^f - \nu_{t-1+3}^3 - \nu_{t-1+2}^2 \\ u_{t-2} - \nu_{t-2+f}^f - \nu_{t-2+3}^3 \\ u_{t-3} - \nu_{t-3+f}^f \\ u_{t-f} \end{bmatrix} \tag{1.7}$$

In any modification that accounts for both data revision and time-varying parameters, both the latent economic quantities (u_t) and the time-varying parameters ($\boldsymbol{\alpha}_t$) are state variables. Since the latter govern the dynamics of the former, the state equation is always nonlinear in the state variables. In this application, the state shock also enters nonlinearly, though this is not generally the case. The transition equation is:

$$\boldsymbol{\beta}_t = \boldsymbol{g}_t(\boldsymbol{\beta}_{t-1}, \boldsymbol{\eta}_t) \tag{1.8}$$

Ignoring trivial equations that merely shift the position of elements from one state vector to the next, (1.8) is the compact notation for the system of equations

$$\pi_t = x_{t-1}(\cdot) + \frac{1}{\zeta_1}\omega_{1t} \tag{1.9}$$

$$u_t = \boldsymbol{\alpha}_{t-1}'\boldsymbol{\Phi}_t + \frac{1}{\zeta_2}\omega_{2t} \tag{1.10}$$

$$\boldsymbol{\alpha}_t = \boldsymbol{\alpha}_{t-1} + \boldsymbol{\Lambda}_t \tag{1.11}$$

$$\nu_{t+2}^2 = \tilde{\nu}_t^2 \tag{1.12}$$

$$\nu_{t+3}^3 = \tilde{\nu}_t^3 \tag{1.13}$$

$$\nu_{t+f}^f = \breve{g}\nu_{t+f-1}^f + \tilde{\nu}_t^f \tag{1.14}$$

$$\tag{1.15}$$

where

$$\begin{pmatrix} \tilde{\nu}_t^2 \\ \tilde{\nu}_t^3 \\ \tilde{\nu}_t^f \end{pmatrix} \sim iid(\mathbf{0}, \tilde{\boldsymbol{V}}) \tag{1.16}$$

(1.9), (1.10), and (1.11) repeat (1.1), (1.2), and (1.3), respectively. Note that x has been written as a function in (1.9) in order to point out that it will be a policy rule depending on the best estimate of $\boldsymbol{\beta}_t$. Furthermore, (1.9) is an element of the vector $\boldsymbol{\Phi}_t$ in (1.10). Hence, (1.9) and (1.10) demonstrate the nonlinear parts of the transition equation.

Similar to how Sargent, Williams, and Zha (2006) set $\boldsymbol{a}_{1|0}$ to equal a regression estimate from data before 1960, I set $\boldsymbol{a}_{1|0}$ to equal the time-varying parameter estimate at the appropriate time from the model without data uncertainty.[1.11] Additionally, now ζ_2 can be estimated.[1.12]

Preliminary analysis of the data shows that there is no predictive relationship between first, second, or third revisions, but that there is some dependence between final revisions. Therefore I only estimate an autoregressive structure in (1.14).

Turning to the measurement equation, the Fed observes

$$\boldsymbol{y}_t = (\pi_t, u_t^0, u_t^1, u_t^2, u_t^3, u_t^f)'$$

so that the state is measured

$$\boldsymbol{y}_t = \boldsymbol{H}\boldsymbol{\beta}_t + \boldsymbol{\varepsilon}_t = \boldsymbol{H}\boldsymbol{\beta}_t + \begin{pmatrix} 0 \\ \tilde{\nu}_t^1 \\ \mathbf{0} \end{pmatrix} \tag{1.17}$$

[1.11] The model including data uncertainty must be initialized, so *its* period 1 is actually period 73 to the point of view of the model without data uncertainty.

[1.12] This is because the variance of the time-varying parameters is not only known in relation to ζ_2, but now also affects the derivative of the policy rule with respect to the time-varying parameters, which allows $\boldsymbol{V}$ to be pinned down apart from ζ_2.

where $\tilde{\nu}_t^1 \sim i.i.d.\,(0,(1/\zeta_\varepsilon)^2)$. $\boldsymbol{H}$ forms the linear combinations written in (1.7) along with the inflation measurement.

1.3.2 Estimation

While (1.17) is linear, (1.8) is nonlinear. The Extended Kalman filter provides a first-order approximation of (1.8), as suggested by Anderson and Moore (1979) and following Tanizaki (1996). Having a method of approximating the optimal predictions and updates that relies only on matrix multiplication makes the entire estimation procedure computationally reasonable. In the interest of exposition, discussion of the Extended Kalman filter and the likelihood is put in Appendix 1.A.4.

The parameter estimated is

$$\boldsymbol{\Psi} \equiv \left(\zeta_1, \zeta_2, \mathrm{vech}\left(\mathrm{Chol}\left(\boldsymbol{V}\right)\right)', \mathrm{vech}\left(\mathrm{Chol}\left(\boldsymbol{P}_{1|0}\right)\right)', \breve{g}, \mathrm{vech}\left(\mathrm{Chol}\left(\tilde{\boldsymbol{V}}\right)\right)', \zeta_\varepsilon\right)'$$

where Chol$(\cdot)$ is the Cholesky factor of positive definite matrix $\boldsymbol{A}$ such that

$$\mathrm{Chol}(\boldsymbol{A})\mathrm{Chol}(\boldsymbol{A})' = \boldsymbol{A}.$$

Note that for the model without data uncertainty, the only part of $\boldsymbol{\Psi}$ estimated is

$$\left(\zeta_1, \mathrm{vech}\left(\mathrm{Chol}\left(\boldsymbol{V}\right)\right)', \mathrm{vech}\left(\mathrm{Chol}\left(\boldsymbol{P}_{1|0}\right)\right)'\right)'$$

Because of $\boldsymbol{\Psi}$'s large dimension, I follow Sargent, Williams, and Zha (2006) and use a Bayesian empirical method discussed in Appendix 1.A.6. From the simulated posterior distribution I report medians as my point estimates and quantiles as probability intervals for the parameters. One could instead numerically maximize the

Table 1.3: Without Data Uncertainty: Parameter Estimates

Value of log likelihood at estimate: 7.8469					
ζ_1:	4.2414	(4.0849, 4.3864)		(3.9690, 4.4781)	
Estimate of $\boldsymbol{P}_{1\|0}$					
0.1071	0.1426	0.0225	–0.2522	–0.0092	–0.1014
0.1426	0.1947	0.0305	–0.3409	–0.0124	–0.1369
0.0225	0.0304	0.0056	–0.0532	–0.0029	–0.0214
–0.2522	–0.3409	–0.0532	0.5995	0.0215	0.2408
–0.0092	–0.0124	–0.0028	0.0215	0.0018	0.0086
–0.1014	–0.1369	–0.0214	0.2409	0.0086	0.0968
Estimate of $\boldsymbol{V}$					
0.0824	–0.0777	0.0092	0.0493	–0.0082	–0.4135
–0.0777	0.0809	0.0001	–0.0502	0.0191	0.6794
0.0092	0.0001	0.0297	0.0013	0.0370	0.7170
0.0493	–0.0502	0.0013	0.0313	–0.0103	–0.3907
–0.0082	0.0191	0.0370	–0.0103	0.0517	1.0066
–0.4135	0.6794	0.7170	–0.3907	1.0066	25.8238

Note: Parameter estimates at the median. ζ_1^2 is the precision of ω_{1t}, the additive shock to the Fed's inflation control; with 68% and 90% probability intervals in parentheses. $\boldsymbol{P}_{1|0}$ is the Fed's initial step-ahead uncertainty over the initial Philips curve parameter estimate $\boldsymbol{a}_{1|0}$. $\boldsymbol{V}$ is the covariance matrix of the $\boldsymbol{\Lambda}_t$ shock to the time-varying parameters $\boldsymbol{\alpha}_t$.

likelihood under the classical paradigm, but I have found that this is computationally burdensome.

Given a prior $p(\boldsymbol{\Psi})$ and the likelihood $\mathcal{L}(\mathcal{Y}_T|\boldsymbol{\Psi})$ (in Appendix 1.A.4), the posterior distribution is

$$p(\boldsymbol{\Psi}|\mathcal{Y}_T) \propto \mathcal{L}(\mathcal{Y}_T|\boldsymbol{\Psi})p(\boldsymbol{\Psi}) \tag{1.18}$$

I sample from (1.18) using a Metropolis algorithm with random walk proposals (cf. Robert and Casella (2004)).[1.13]

1.3.3 Without Data Uncertainty: Results

Data on inflation and the civilian unemployment rate for ages 16 and older comes from the BEA and BLS, respectively, as reported in December 2003. [1.14] The data begins in 1960 and I end the sample at December 1995.[1.15] Table 1.4 shows estimates from 75,000 draws derived from 100,000 MCMC iterations where the first 25% are burned in hopes that the Markov Chain has, for practical purposes, converged to its ergodic distribution.[1.16] The estimates in Table 1.4 are virtually identical to those of Sargent, Williams, and Zha (2006).

Figure 1.1 shows the predicted inflation control choices. Figure 1.1 shows the Fed *choosing* to set inflation high in the two high-inflation episodes of the mid 1970s and early 1980s. With ζ_1 estimated at about 4.24, the standard deviation of the ω_{1t} is around 0.24, reflecting the Fed's belief that it has rather tight control of inflation.

The Philips curve beliefs $a_{t-1|t-1}$ are used to forecast the next month's unemployment rate for any inflation control setting. According to the model, the Fed sets inflation with this forecast in mind. Therefore an important aspect of the model-

[1.13] I must use a accept/reject simulation technique because, due to the effects of $\mathbf{\Psi}$ on the whole sequence of forecasts, the form of (1.18) is not known. Further details are in Appendix 1.A.6.

[1.14] The inflation rate data, following Sargent, Williams, and Zha (2006) is the annual rate of change of the seasonally-adjusted Personal Consumption Expenditure chain price index from the BEA, as reported in December 2003. I use PCE inflation here in the interest of making a direct comparison although I must use real-time CPI inflation for the model including data uncertainty. Estimates of the model without data uncertainty using current vintage non-seasonally-adjusted or seasonally-adjusted CPI inflation are very similar and so this does not drive the difference between the models' estimates; see Appendix 1.A.6.1.

[1.15] The results are the same if the data runs through December 2003; I end the sample at 1996 in order to ensure comparability between all the models I have estimated for robustness (one of which assumes final data is seen 10 years after the fact) and because I consider the model to be descriptive of Fed beliefs only through the 1980s.

[1.16] Sargent, Williams, and Zha (2006)'s results come from a sequence of 50,000 draws with an unspecified burn-in interval.

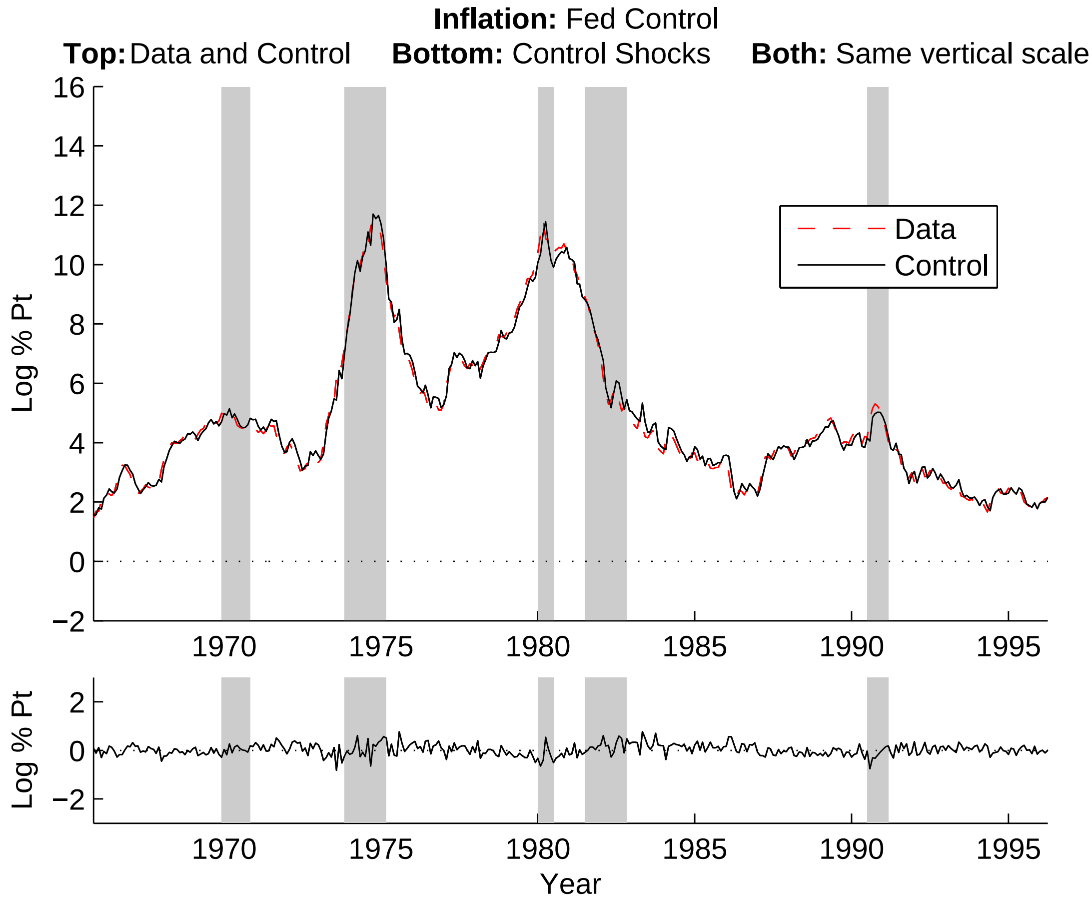

Figure 1.1: Without Data Uncertainty: Inflation

Note: Actual inflation versus the Fed's control, using the model without data uncertainty. NBER recessions shaded. Figures 1.1 and 1.3.4 are on the same scale.

predicted Fed beliefs are what they deliver in terms of unemployment forecasts: these are plotted in Figure 1.3.3.[1.17] These forecasts have some negative bias and a Root Mean Squared Error of 3.3 percentage points. Roughly speaking, every four months the Fed expects its month-ahead unemployment rate forecasts to be off by 3.3.

An important aspect of these unemployment rate forecasts are how well they

[1.17]The information displayed in Figure 1.3.3 is implied by Sargent, Williams, and Zha (2006)'s Figure 7. The difference is that in their Figure 7 the inflation control is replaced by the "Ramsey" inflation rate policy of 2, the graph is shifted downward by the estimate of the true natural rate of unemployment, and the parameter estimate used for time t is the updated estimate $\boldsymbol{a}_{t|t}$ (which is identically $\boldsymbol{a}_{t+1|t}$).

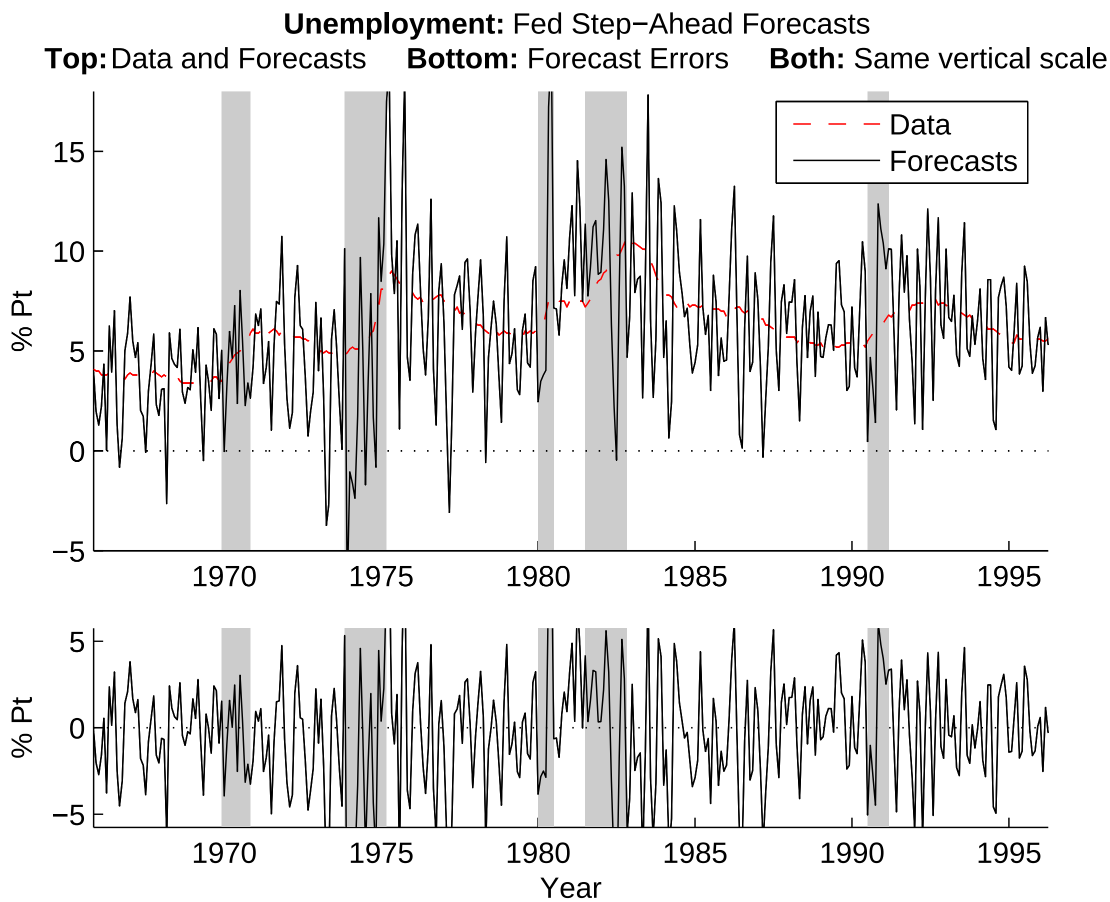

Figure 1.2: Without Data Uncertainty: Unemployment

Note: Actual unemployment versus the Fed's step-ahead unemployment forecasts, using the model without data uncertainty. These step-ahead unemployment forecasts come from the Philips curve (1.2) using the Fed's inflation setting and actual unemployment and inflation data. NBER recessions shaded. Figures 1.3.3 and 1.3.4 are on the same scale.

explain actual Federal Reserve unemployment rate forecasts. There is evidence of the latter from the Greenbook forecasts published at irregular intervals over this time span.[1.18] An appropriate statistical test of the similarity between the model-predicted forecasts and the Greenbook forecasts is that of Diebold and Mariano (1995). The statistic S_1 is a two-sided test of the null hypothesis that the model's unemployment forecasts have accuracy equal to the Greenbook forecasts; S_1 has an asymptotic stan-

[1.18] Appendix 1.A.5 discusses these forecasts and provides more details on the test statistic.

dard normal distribution. Here $|\hat{S}_1| = 2.5044$, so we can reject the null of equal forecasting accuracy at the 99% level. This means that in terms of accuracy the model-predicted forecasts are statistically different than actual Greenbook forecasts.

Largely volatile and inaccurate unemployment forecasts are evidence of either greatly varying parameters or misspecification of the forecast rule: since the model imposes that the forecast rule (1.2) is never discarded, the estimate of the parameter shock's covariance matrix is driven up. In particular, notice that the estimated $V^{(6,6)}$ implies that the Fed believes that every month the Philips curve's intercept experiences an i.i.d. shock with a standard deviation of 5.14 unemployment rate points.[1.19]

The results suggest a problem to the inflation-as-optimal-policy story. On the one hand, the estimates imply the Fed regarded its unemployment forecasting tool as erratic and is able to see its persistent inaccuracy. On the other hand, the model says that the estimated Philips curve relationship motivated inflationary policy in the hopes of lowering the unemployment rate. Given the estimates, it is difficult to agree with the idea that the Fed believed its filtered model enough to set inflation so far from target.

1.3.4 Including Data Uncertainty: Results

Multiple vintages of real-time data on inflation and the civilian unemployment rate for ages 16 and older comes from the ALFRED archive maintained by the Federal

[1.19] One might suppose that an Orphanides-type critique is at play here: this says that the model should use preliminary inflation and unemployment rate data. In fact, this does not alleviate the problems I discuss below and using preliminary data leaves the results qualitatively the same – see Appendix 1.A.6.2.

Table 1.4: Including Data Uncertainty: Parameter Estimates

Value of log likelihood at estimate: 1055.2			
ζ_1:	4.6573	(4.4617, 4.9476)	(4.2607, 5.0900)
ζ_2:	5.1624	(4.4097, 5.6742)	(4.0352 , 6.1434)
$\breve{g}$:	–0.0048	(–0.0185, 0.0012)	(–0.0244, 0.0028)
ζ_ε:	3.8759	(3.8700, 3.8774)	(3.8669,3.8823)

Estimate of $\tilde{\boldsymbol{V}}$

0.0100	0.0001	0.0001
0.0000	0.0100	0.0001
0.0001	0.0001	0.1002

Estimate of $\boldsymbol{P}_{1|0}$

0.1799	0.1130	0.3503	–0.4104	0.1806	0.3657
0.1130	0.9961	1.0329	–1.1465	–0.0811	0.3275
0.3503	1.0329	2.1429	–1.5729	–0.0729	0.5853
–0.4104	–1.1465	–1.5729	1.7901	–0.2264	–0.9296
0.1806	–0.0811	–0.0729	–0.2264	0.3437	0.5246
0.3657	0.3275	0.5853	–0.9296	0.5246	1.3844

Estimate of $\boldsymbol{V}$

0.0001	0.0004	–0.0003	–0.0003	0.0002	0.0005
0.0004	0.0020	–0.0015	–0.0015	0.0007	0.0055
–0.0003	–0.0015	0.0016	0.0011	–0.0011	–0.0011
–0.0003	–0.0015	0.0011	0.0013	–0.0004	–0.0071
0.0002	0.0007	–0.0011	–0.0004	0.0010	–0.0044
0.0005	0.0055	–0.0011	–0.0071	–0.0044	0.1149

Note: Parameter estimates at the median. ζ_1^2 is the precision of ω_{1t}, the additive shock to the Fed's inflation control. ζ_2^2 is the precision of ω_{2t}, the additive shock in the Philips curve. $\breve{g}$ governs the relationship between ν_t^f and ν_{t-1}^f. ζ_ε^2 is the precision of the preliminary data revision ν_t^1. $\tilde{\boldsymbol{V}}$ is the covariance matrix of the revision shock $\tilde{\boldsymbol{\nu}}_t$. $\boldsymbol{P}_{1|0}$ is the Fed's initial step-ahead uncertainty over the initial Philips curve parameter estimate $\boldsymbol{a}_{1|0}$. $\boldsymbol{V}$ is the covariance matrix of the $\boldsymbol{\Lambda}_t$ shock to the time-varying parameters $\boldsymbol{\alpha}_t$.

Reserve Bank of St. Louis, downloaded in 2007.[1.20] Table **??** reports estimates from 700,000 MCMC iterations from two separate runs of 400,000 with different initial conditions where the first 50,000 of each run is burned.

The estimated standard deviation of the inflation control shock is 0.22, re-

[1.20] I use non-seasonally-adjusted CPI inflation because real-time PCE inflation is not available for the time span under consideration. If we assume current vintage PCE inflation data was known in real-time and underwent no revisions, the results presented here are virtually identical.

In earlier versions I used the real-time non-seasonally-adjusted CPI data collected by Ghysels, Swanson, and Callan (2002) which is identical to that in ALFRED which became available after this paper's first version.

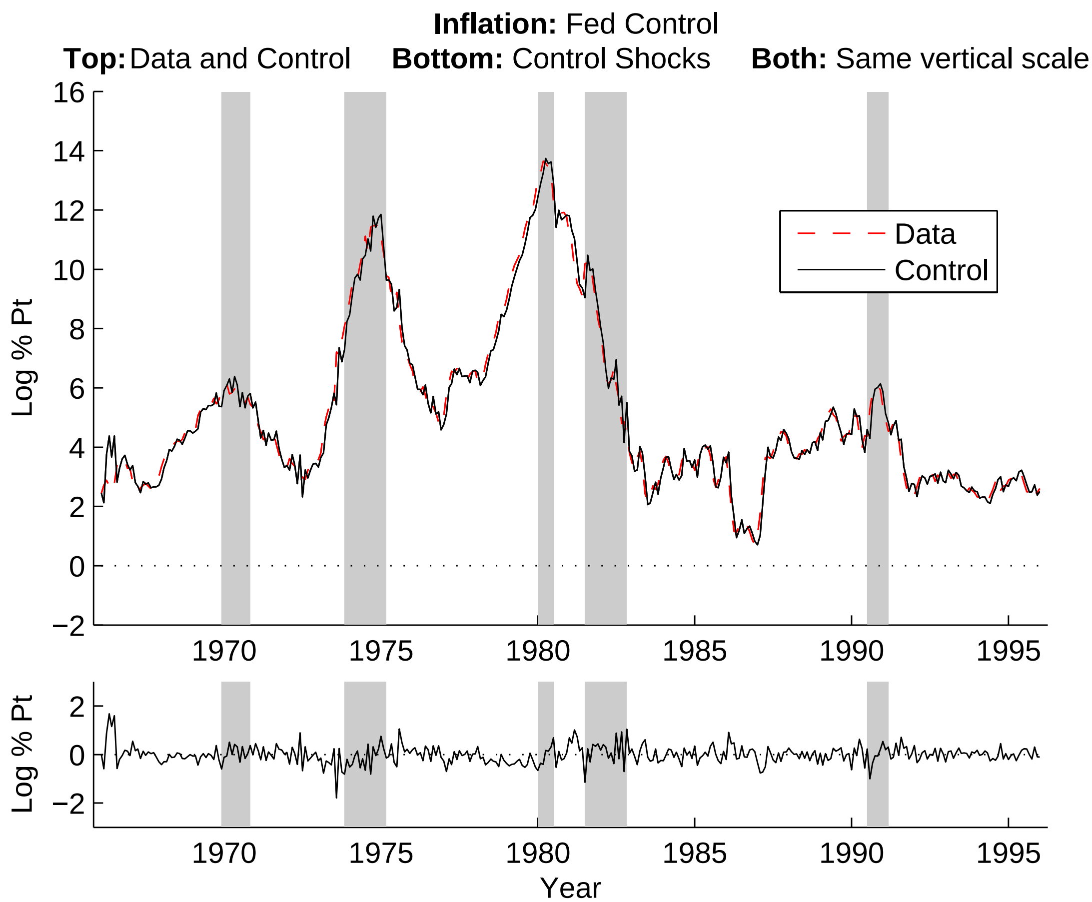

Figure 1.3: Including Data Uncertainty: Inflation

Note: Actual inflation versus the Fed's control, using the model including data uncertainty. NBER recessions shaded. Figures 1.1 and 1.3.4 are on the same scale.

flecting the Fed's belief that it has tight control of inflation. The estimate for ζ_2 says that the Fed expects its Philips curve forecast rule, if it knew the correct parameters, to deliver forecasts with a RMSE of 0.19, which is a little less than the RMSE of a random-walk forecast. The estimate of ζ_ε implies a standard deviation around 0.2536: the Fed overestimates the preliminary revision variance, having apparent difficulty in untangling the effect of ν_t^1 from the other shocks and time-varying parameters. The estimates of $\breve{g}$ and $\tilde{V}^{(3,3)}$ indicate that the Fed acts as though there is essentially no predictability between final revisions. Figure 1.3.4 shows that the Fed's inflation

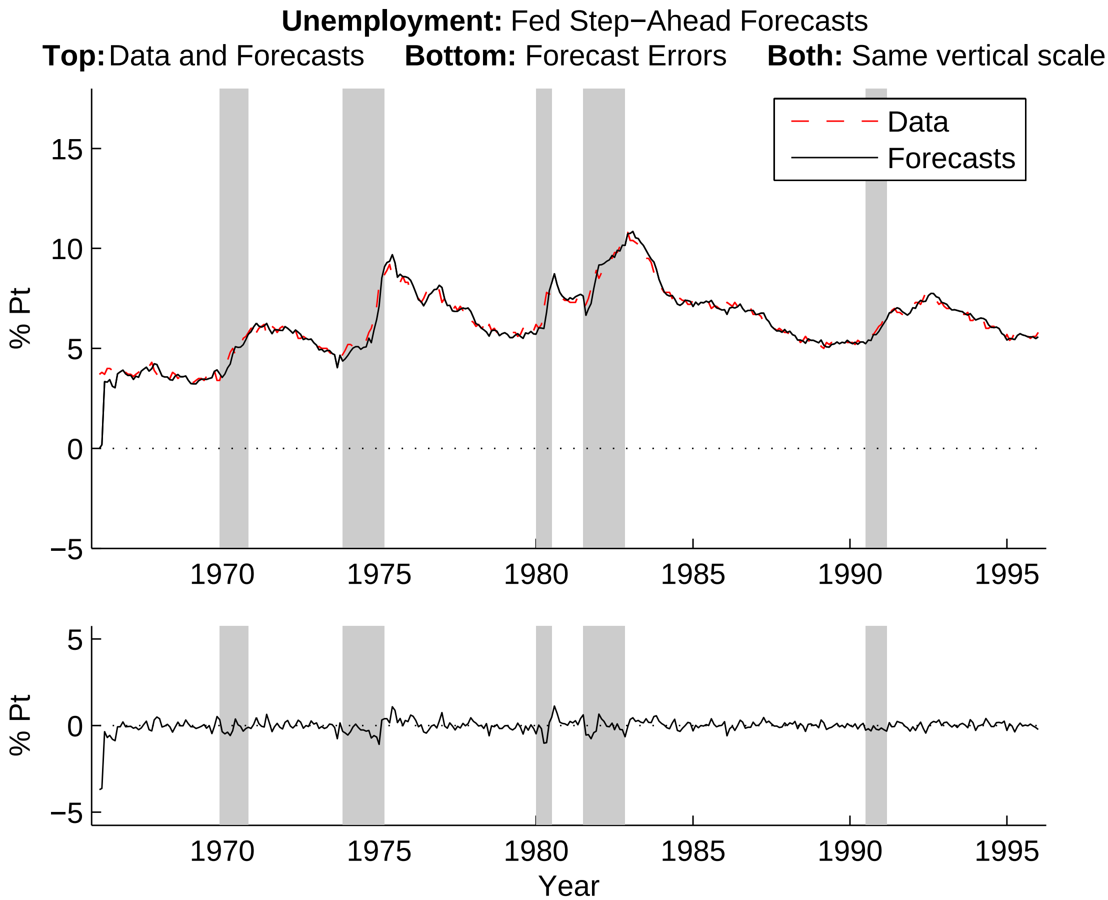

Figure 1.4: Including Data Uncertainty: Unemployment

Note: Actual unemployment versus the Fed's step-ahead unemployment forecasts, using the model including data uncertainty. These step-ahead unemployment forecasts come from the Philips curve (1.10). NBER recessions shaded. Figures 1.3.3 and 1.3.4 are on the same scale.

control explains the rise and fall of American inflation.

Turning now to the Fed's unemployment rate forecasts in Figure 1.3.4, we find a far different picture than in the model without data uncertainty. The Fed's forecasts are considerably more accurate than before, with a RMSE of 0.4007 percentage points and no significant bias.[1.21]

Again, seeing as the model forecasts are intended to predict the Fed's actual

[1.21] The other measurement equation forecasts are pictured in Appendix 1.A.6.

expectations about the future of unemployment, we can directly compare them to Greenbook forecasts. Using the Diebold and Mariano (1995) test of equal accuracy between model forecasts and Greenbook forecasts, $|\hat{S}_1| = 0.4858$ and we accept the hypothesis that the model forecasts and Greenbook forecasts are equally-accurate. Therefore in terms of accuracy the model-predicted unemployment rate forecasts are statistically the same as Fed unemployment rate forecasts.

Since the unemployment rate forecast errors are much smaller on average, the adapted model does not require a greatly-varying Philips curve to account for the maintained assumption that the Fed always retains (1.2) as its forecast rule. Accordingly, the estimate of $\boldsymbol{V}$ in Table **??** is much smaller than before. In particular, the estimated $V^{(6,6)}$ implies that the Philips curve's intercept has a monthly shock with a standard deviation of about $\frac{1}{3}$ as opposed to the 5 points estimated with data uncertainty left out.[1.22]

It is worth noting that the estimate of the Fed's initial Philips curve uncertainty, $\boldsymbol{P}_{1|0}$, is somewhat larger here. This is due to the choice of $\boldsymbol{a}_{1|0}$ as the estimate $\boldsymbol{a}_{t|t-1}$ for January 1966 from the model without data uncertainty. This Philips curve forecasts unemployment poorly, explaining why the Fed is estimated to be rather uncertain of its value. If $\boldsymbol{a}_{1|0}$ is chosen so as to forecast unemployment reasonably well in the first period, the estimate of $\boldsymbol{P}_{1|0}$ drops with very little change to the other

[1.22] It has been suggested that one could adopt part of the model without data uncertainty *ad hoc* and include unemployment forecast errors as part of a modified version of an estimation procedure resembling the model with data uncertainty. Doing so might by itself lead to smoother TVP estimates and better fitting inflation and unemployment forecasts. Looking into this, I found that while the unemployment forecasts errors do improve, the inflation forecasts greatly deteriorate; see Appendix 1.A.6.3.

parameter estimates.

By including data uncertainty, the inflation-as-optimal-policy story is supported by all the data brought to the model. The Fed is confident in its rule, as measured by the small estimates of the main diagonal of $\boldsymbol{V}$: the estimated Philips curve serves as a good predictor of unemployment. It is now plausible that the Fed sets inflation in hopes of affecting unemployment in the way (1.2) forecasts.

1.3.5 Discussion

By introducing data uncertainty, I make the Philips curve (1.2) more flexible. Clearly, data uncertainty is not the only possible form of misspecification leading (1.2) to forecast much differently than is evidenced by Greenbook forecasts. On the other hand, the flexibility I introduce is an observable form of measurement error and naturally follows from allowing the modeled Fed to handle multiple real-time data vintages in a realistic fashion. My results suggest that nonlinearity is part Fed's forecasting rule as it relates to the Fed's chosen inflation level. Accompanying this are estimates that imply the Fed perceives shocks to its parameters to be much smaller on average.

This is actually indicative of a more general issue: when parameter shocks and data revisions are independent, parameter volatility is always overestimated when data revisions exist but are ignored. For a simple demonstration of this fact, consider the data generating process $y = (b + c)(x + e)$ where b and x are known to an agent named F. The real numbers b and x are F's predictions of the true parameter and data values, respectively; c and e are the mean zero, uncorrelated prediction errors. F makes the forecast $\hat{y} = bx$ and then observes the forecast error $y - \hat{y}$.

Suppose that F ignores data uncertainty, assuming $e = 0$ always. Then F can calculate $\mathbb{V}ar(y - \hat{y})$ as equal to $x^2 \mathbb{V}ar(c)$ – therefore F perceives the variance of the parameter shock as $\frac{\mathbb{V}ar(y-\hat{y})}{x^2}$. However, this quantity is actually equal to

$$\mathbb{V}ar(c) + \frac{b^2}{x^2}\mathbb{V}ar(e) + \frac{1}{x^2}\mathbb{V}ar(e)\mathbb{V}ar(c)$$

The second and third terms are nonnegative. Moreover, they are nonlinearly related to the levels of the data and parameter predictions and the variances of the parameter and data prediction errors.[1.23]

A lesson from this is that time-varying parameters are sensitive to whether or not we account for data revisions – we always overestimate the volatility of parameter shocks by ignoring existing data revisions when the two are independent. The nonlinear relationship between perceived parameter volatility and revision volatility means that this bias depends on the level of data and parameter values. Therefore, for some applications this volatility bias may be insignificant, while for others – as seen here – it is important.

According to the preceding sections, the model including data uncertainty best describes the data. In light of the simple demonstration, it's no surprise that the model without data uncertainty overestimates the variance of the parameter shocks.

[1.23] A similar bias, involving more symbols, is present for DGPs including more data and more parameters. In this case, a likelihood-based estimation method would deliver the specific decomposition of a function of the forecast error variance into different parameter shocks' variance (in this simple example here, this function is $\mathbb{V}ar(y - \hat{y})/x^2$). The point made in this simple example still holds – the function of the forecast error would overestimate the volatility of the parameter shocks in a nonlinear fashion that may amplify the effect of revision shocks' small variances.

This biased estimate then feeds into the policy rule: as the parameters on unemployment are thought to be subject to larger shocks, they are predicted to substantially shift from month to month. But actual inflation never experiences an enormous monthly jump, which means that next month's optimal policy must be similar to this month's. This leads the model without data uncertainty to use a volatile constant coefficient to attenuate the change in the policy rule brought about by shifts in the other Philips curve parameters. On the other hand, the model including data uncertainty has the means to decompose forecast errors into not only parameter shocks but also data prediction errors and estimates a much smoother evolution for the constant coefficient.[1.24]

By ignoring data uncertainty, the maintained assumption – that the Fed makes forecasts with (1.2) that it believes – appears implausible. The implausibility follows from high estimates of the Fed's uncertainty over the rule (represented by large main diagonal elements of $\boldsymbol{V}$) and persuasive evidence of its misspecification (represented by persistently poor forecasting performance). However, by recognizing that the Fed is aware of data revisions and took into consideration multiple vintages of data, these problems with the basic story subside. Estimates support the idea that the inflationary episodes of the 1970s and 1980s can be modeled as optimal policy stemming from evolving estimates of a Philips curve trade-off.

[1.24]See Appendix 1.A.6.4 for plots of the evolution of the constant coefficient estimates.

1.4 Production Model

Firms are uncertain about macroeconomic and firm-specific conditions, but the latter are usually assumed to be the major influence on firm behavior – see Comin and Philippon (2005). In a representative agent model firm-specific information is aggregated across identical firms to calculate macroeconomic data. In this case macroeconomic data revisions are equivalent to firm-specific data revisions.

In this model I assume that firms cannot perfectly measure the real marginal product of their workers in real-time. The model is otherwise without friction, and if the firm is able to perfectly observe productivity there is no propagation in labor's response to productivity shocks (no hump-shaped impulse response function). By incorporating firm-level data uncertainty, on the other hand, aggregate labor responds to productivity with the hump-shape estimated in the data.

The general equilibrium model of Aruoba (2004) is similar in spirit to the partial equilibrium model presented here. That model uses a single period signal extraction problem to imply increased welfare from increasingly precise data releases. In contrast, I use the fact that there are multiple revisions to any period's report of productivity. This allows me to show that data uncertainty is an informational friction that can mean a real shock affects optimal decisions both on impact and in later periods.

1.4.1 Model

Labor and productivity are the only inputs of the representative firm's production function

$$Y_t = Z_t N_t$$

which outputs a homogenous consumption good. The firm is unable to perfectly see the productivity of its workers. Instead, the firm observes a year's worth of quarterly measures of their past productivity. That is, at time t the firm sees the measurements $\{z_t^0, z_t^1, z_t^2, z_t^3, z_t^4\}$ of what z_t is and was.

The log of labor productivity follows an AR(1). Letting $z_t \equiv \log(Z_t)$,

$$z_t = \mu_z + \rho_z z_{t-1} + \varepsilon_t \tag{1.19}$$

where $|\rho_z| < 1$ and $\varepsilon_t \sim iid(0, \sigma_z^2)$.

The labor supply decision is exogenously given. Assuming the detrended economy experiences small deviations about the deterministic steady state, a linear log labor supply function is a good approximate specification requiring one parameter: the Frisch labor supply elasticity λ_{Frisch}. Hence the log labor supply is

$$n_t = \lambda_{\text{Frisch}} w_t \tag{1.20}$$

for w the log real wage.

The consumption good and labor markets are perfectly competitive. Labor supply can meet any labor demand. Given the environment, the firms bid down the log real wage rate until it reaches their (mutual) best guess of productivity, $\hat{z}$. The log real wage bill is therefore $(\lambda_{\text{Frisch}} + 1)\hat{z}$, as is predicted log output.

Note the following complication: actual log output is going to be $\lambda_{\text{Frisch}}\hat{z} + z$ which in the conventional setup would allow the representative firm to immediately back out true log productivity from the equilibrium price of the good. But this eliminates any role for data revisions because in this case productivity is known at the end of the period. To avoid this, I assume that claims to log final good are sold as soon as $\hat{z}$ is announced, but before production happens. This means that the households do not actually buy part of the consumption good, but instead buy a claim to a proportion of future output whose *ex post* real price is

$$\frac{\lambda_{\text{Frisch}}\hat{z} + z}{(\lambda_{\text{Frisch}} + 1)\hat{z}}$$

Regardless of what the actual value of z turns out to be, the *ex ante* price of this claim is unity because $\frac{\lambda_{\text{Frisch}}\hat{z}+z}{(\lambda_{\text{Frisch}}+1)\hat{z}}$ is rationally predicted to be unity by all agents. This way, the equilibrium price seen by the firm (for a claim to part of overall output) does not reveal z. An anonymous market matches claims with actual output, which the firms do not observe. Since household income equals expenditure, the goods market clears. Moreover, I assume the household and firm do not share information, as the household's actual consumption would also reveal z to the firm.[1.25]

From this point on, in the text (not in figures or tables) I refer to n_t as *labor*, to z_t as *productivity*, and to ε_t as a *productivity shock*.

[1.25] It is due to the complication discussed here that I do not pursue a more standard representative agent model including capital, where it would necessitate even more complicated market clearing and information-sharing assumptions. A heterogenous agent model may get around these uncommon assumptions while necessitating more sophisticated considerations – see for instance Lorenzoni (2006).

1.4.2 Data

The data were downloaded in March 2007 from the ALFRED archive maintained by the Federal Reserve Bank of St. Louis. Aggregate labor supply is calculated monthly as the Average Weekly Hours in Total Private Industries (AWHNONAG) times Total Nonfarm Payrolls (PAYEMS) times 4: the three months of the quarter are then added up to get the quarterly aggregate labor input. Output is the seasonally-adjusted nominal GNP series and the deflator is GNPDEF.[1.26] I avoid revisions due to changes in base year, as are experienced by Real GNP: for this reason I use Nominal GNP and deflate by the 2007 March 1 vintage GNP deflator. This abstracts from revisions to the GNP deflator over time. The only remaining source of data uncertainty has to do with revisions to nominal output and aggregate labor.[1.27] The time span of the data is limited by AWHNONAG which begins in 1964: therefore the quarterly data runs from 1964:I to 2006:III. The firm detrends productivity data using the HP-filter (with $\lambda = 1600$): the productivity value is directly calculated from output and aggregate labor data.

Using the current data vintage, the plot of the orthogonalized impulse response of labor to a productivity shock ε is in Figure 1.5. This impulse response function

[1.26] Real-time data on GNP is available for a much longer period than for GDP because the former was the masthead output value prior to the 1990s. Real-time data on quarterly GNP is available only in seasonally-adjusted form.

[1.27] In this detrended economy, the agents find productivity values by detrending the newest data vintage. I want to highlight the data uncertainty stemming only from revisions to the data itself, not due to trend estimation. Therefore for each time t I calculate a data revision only as a change in the underlying log-output relative to the trend value that was estimated when there was a preliminary measurement of time t log-output. See Appendix 1.A.8 for details. Edge, Laubach, and Williams (2004) address the disentangling of productivity deviations from trend changes in real-time, an interesting issue that is beyond the scope of this paper.

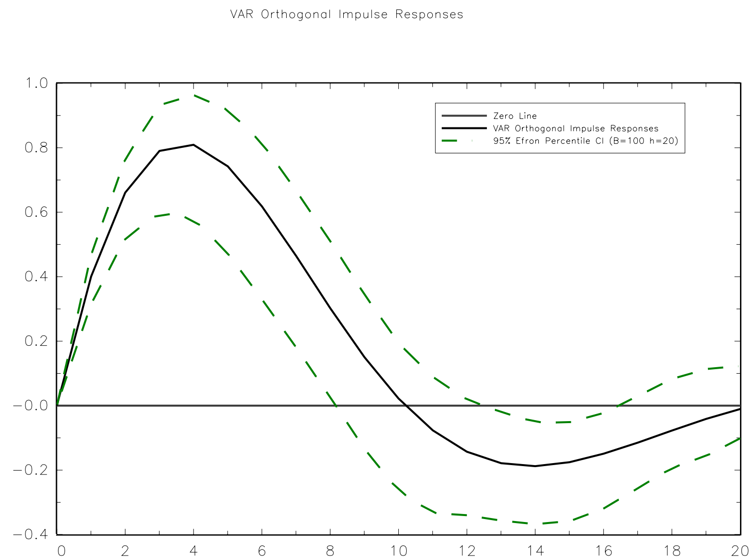

Figure 1.5: Impulse Response Function: Data

Note: Impulse response function of log hours fluctuations to log productivity fluctuation shocks; as a proportion of a one-standard-deviation shock. Both hours and productivity are current (2007) vintage data. From a bivariate VAR with 1 lag (chosen by BIC). The horizontal axis is in quarters.

is estimated by a bivariate VAR with lag-length chosen by the Bayesian Information Criterion.[1.28] This is the feature of the data that I attempt to match.[1.29] The impulse response functions and bootstrapped error bands follow Efron and Tibshirani (1993) and Lutkepohl and Kratzig (2004). Key features of this impulse response are its hump-shape, peak response at three to four quarters after the shock, and peak magnitude being about 80% of the shock to productivity.

Figure 1.6 plots productivity and productivity shocks (the residuals of the regression of (1.19)). We see that productivity shocks from current vintage data

[1.28]In fact, for the data and the full model the BIC chooses one lag.

[1.29]This exercise is not meant to join the discussion on labor impulse responses to productivity shocks coming from more sophisticated DSGE models: the discussion is found, for example, in Burnside, Eichenbaum, and Rebelo (1995), Gali (1999), Francis and Ramey (2002), Christiano, Eichenbaum, and Evans (2005), and Basu, Fernald, and Kimball (2006). Instead, this exercise demonstrates that a by incorporating multiple vintages of real-time data a simple model delivers propagation, which has been a weakness of RBC models (see King and Rebelo (1999)).

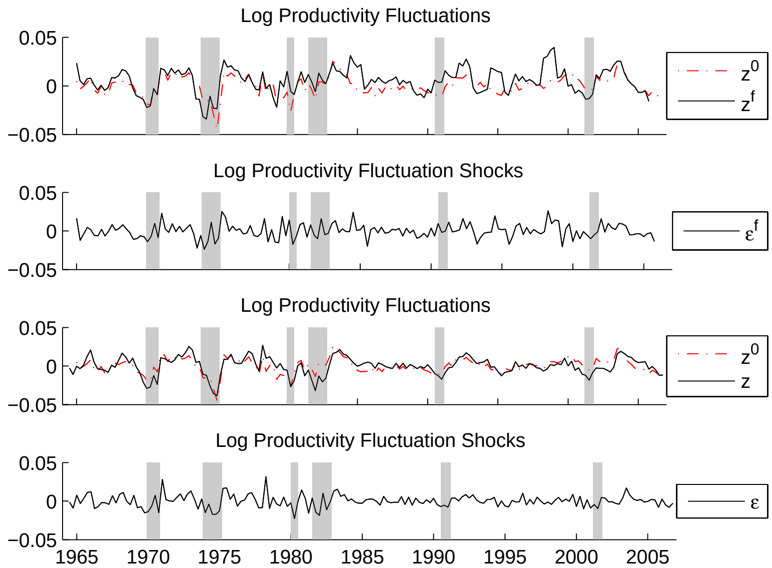

Figure 1.6: Productivity and Productivity Shocks
First panel: log productivity; the preliminary values are z^0 and the final vintage values (data as reported after 4 quarters) are z^f. **Second panel:** residual from regression $z_t^f|z_{t-1}^f, 1$; these are the estimated log productivity shocks for final vintage data. **third panel:** log productivity; the preliminary values are z^0 and the current (2007) vintage values are z. **Fourth panel:** residual from regression $z_t|z_{t-1}, 1$; these are the estimated log productivity shocks for current (2007) vintage data. NBER recessions are shaded.

exhibit the Great Moderation starting around 1984. Interestingly, a drop in volatility is not clearly evident in productivity shocks from one-year horizon real-time data.

For both current vintage productivity and final vintage productivity (data as reported after 4 quarters), autoregression estimates are reported in Table 1.5. Table 1.6 shows summary statistics for current vintage productivity shocks, year-horizon productivity shocks, and the first year's revisions. The standard deviation of each revision is at least 60% that of the perceived productivity shock.

Table 1.5: Productivity Autoregressions

Log Productivity Autoregression Estimates

$$z_t = \mu_z + \rho_z z_{t-1} + \varepsilon_t$$

μ_z	ρ_z	N	$\bar{R}^2$
–0.0000 (0.0006)	0.7217 (0.0706)	170	0.51

$$z_t^f = \mu_z^f + \rho_z^f z_{t-1}^f + \varepsilon_t^f$$

μ_z^f	ρ_z^f	N	$\bar{R}^2$
0.0015 (0.0009)	0.6859 (0.0720)	164	0.46

Top panel: estimates using current vintage data. **Bottom panel:** estimates using final vintage data (data as reported after 4 quarters). Robust standard errors are in parentheses.

Table 1.6: Summary Statistics: Productivity Shocks and Revisions

Log Productivity Shock and Revision Statistics

Univariate Autoregressions					
Var	μ	ρ	N	$\bar{R}^2$	StDev
ε	0.0000 (0.0006)	0.0571 (0.0952)	169	–0.00	0.0080
ε^f	–0.0001 (0.0007)	0.0187 (0.0798)	163	–0.01	0.0094
ν^1	0.0013 (0.0005)	–0.0710 (0.0320)	167	–0.00	0.0057
ν^2	0.0013 (0.0005)	–0.0790 (0.0305)	166	0.00	0.0056
ν^3	0.0013 (0.0004)	–0.0754 (0.0275)	165	–0.00	0.0055
ν^4	0.0012 (0.0004)	–0.0704 (0.0267)	164	–0.00	0.0054

Contemporaneous Correlations						
	ε	ε^f	ν^1	ν^2	ν^3	ν^4
ν^1	0.02	0.19	·			
ν^2	–0.04	0.18	–0.08	·		
ν^3	–0.07	0.17	–0.06	–0.08	·	
ν^4	–0.01	0.48	–0.03	–0.11	–0.07	·

Top Panel: the second and third columns are constant and AR(1) parameter estimates, the fourth is the number of observations, and the fifth is the adjusted-R^2 from a univariate autoregression of the variable in the first column; the sixth column is the standard deviation of the variable; robust standard errors in parentheses. **Bottom Panel:** contemporaneous correlation matrix between variables.

1.4.3 Simulation Results

The statistics above calibrate the model. Given the evidence, I assume that productivity shocks are mean zero and i.i.d.. I have found that the key calibration values are the standard deviations of revisions and productivity shocks and the autoregressive parameter for productivity – the other covariance and autoregressive parameters are inconsequential. I set $\lambda_{\text{Frisch}} = 1$ following the evidence surveyed in Blundell and Macurdy (1999), Rotemberg and Woodford (1999), and Christiano, Eichenbaum, and Evans (2005). This parameter attenuates the magnitude of the impulse response while leaving the shape unaffected.

For each of the four experimental designs described below, I simulate a labor series by feeding the actual productivity shocks and revisions into the model. The simulated labor series is then put into a bivariate VAR with the current vintage values of productivity, as was done with the actual labor data, to produce the impulse response function plotted in Figure 1.5. The main result of each experiment is this estimated impulse response function.

I build up to presenting the impulse response function predicted by the full model by highlighting its two key elements: (1) the firm recognizes the actual data uncertainty coming from four possible data revisions to each period's productivity value; (2) the final productivity measurement that the firm sees is actually subject to subsequent revisions the firm ignores. I show how both elements contribute by eliminating both (the benchmark Neoclassical case), eliminating only the first, and then eliminating only the second. The fourth impulse response function includes both elements and demonstrates that together they allow the model to match the data.

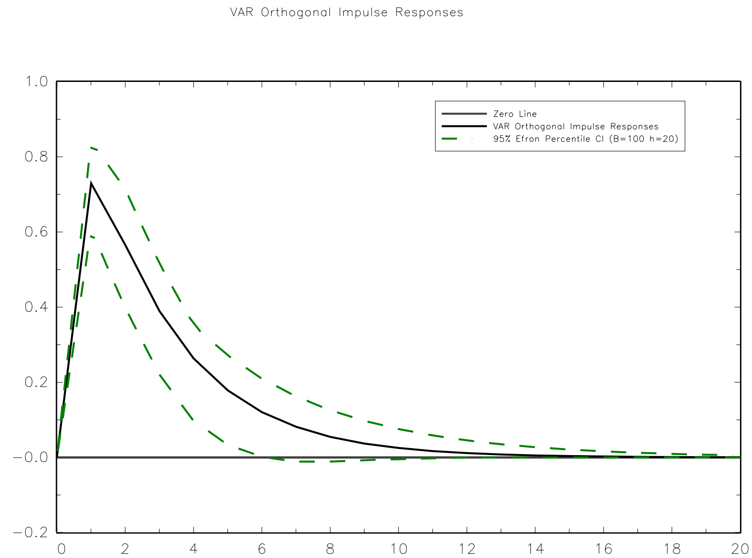

Figure 1.7: Impulse Response Function: No Data Uncertainty, No Ignored Revisions
Note: Impulse response function of simulated log hours fluctuations to log productivity fluctuation shocks; as a proportion of a one-standard-deviation shock. Productivity is current (2007) vintage data. Hours are simulated assuming (1) the firm *has no* data uncertainty, and (2) there *are no* data revisions the firm ignores. From a bivariate VAR with 1 lag (chosen by BIC). The horizontal axis is in quarters.

1.4.3.1 Benchmark

It is standard to assume the firm sees the true value of the data, which is the data as the researcher sees it. The experimental design is:

> *Suppose the firm sees true productivity for time t after labor is committed and believes the structure estimated in Table 1.5 governs productivity.*

The firm enters period t knowing the true value of productivity last period. The optimal forecast of time t productivity is $\rho_z z_{t-1}$. The firm uses this forecast to set the log real wage, which elicits labor according to (1.20). The estimated impulse response function is plotted in Figure 1.7.

The productivity shock directly affects next period's labor due to the autoregressive dynamics of productivity. The response geometrically declines after the first quarter according to the autoregressive parameter ρ_z. The magnitude of the peak is about that seen in the data, but there is no hump-shape. Figure 1.7 demonstrates the weak internal propagation mechanism found in benchmark Neoclassical models, a feature highlighted by King and Rebelo (1999).

1.4.3.2 Only Ignored Data Revisions

This is the model as subject to an Orphanides-type critique: agents actually react to preliminary real-time data. The firm is not sophisticated enough to recognize that data undergoes revision. The experimental design is:

> *Suppose the firm sees only the preliminary productivity value for time t after labor is committed and believes the structure estimated in Table 1.5 governs productivity.*

The firm enters period t knowing only the preliminary report of productivity last period. Since the firm does not hold any uncertainty over its data, the optimal forecast of time t productivity is $\rho_z z^0_{t-1}$. However, the firm ignores all the data revisions that imply, in general, $z^0_{t-1} \neq z_{t-1}$. The firm uses this forecast to set the log real wage, which elicits labor according to (1.20). The estimated impulse response function is plotted in Figure 1.8.

When the firm thinks that no data revisions exist, the impulse response of labor to true productivity has a hump-shape that peaks two quarters after the shock,

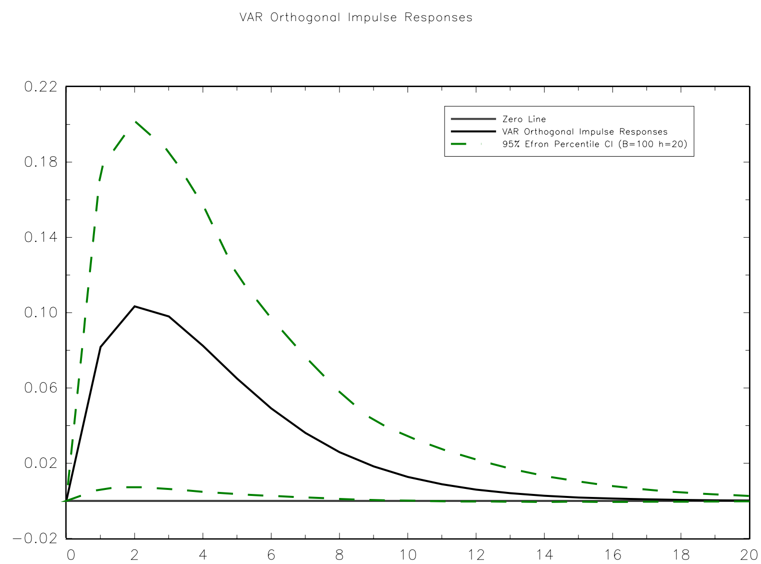

Figure 1.8: Impulse Response Function: No Data Uncertainty, Ignored Revisions
Note: Impulse response function of simulated log hours fluctuations to log productivity fluctuation shocks; as a proportion of a one-standard-deviation shock. Productivity is current (2007) vintage data. Hours are simulated assuming (1) the firm *has no* data uncertainty, and (2) there *are* data revisions the firm ignores. From a bivariate VAR with 1 lag (chosen by BIC). The horizontal axis is in quarters.

about 1-2 quarters too early when compared to the data. This peak response is about 12% of what is seen in the data.

This hump-shape is coming from the difference between the preliminary report of productivity and true productivity, both plotted in the third panel of Figure 1.6. Since the firm does not recognize any revision, there is no economic reason why this hump-shape should exist. Instead, it is due to the empirical observation that many of the productivity shocks in preliminary data lag the true productivity shocks by one or two quarters. This does not always happen, which leads to the large confidence interval about the impulse response function. However, it happens enough for the VAR to pick up a hump-shape.

1.4.3.3 Only Data Uncertainty

This simulation demonstrates the effects of data uncertainty alone. There are no data revisions that the firm does not observe. The experimental design is:

> *Suppose the firm sees a year's worth of productivity measures – at the end of the year, it sees the true productivity values. The firm believes the structure estimated in Tables 1.5 and 1.6 governs productivity and revisions.*

The firm enters period t knowing multiple vintages of real time data on productivity for the past year. Moreover, the productivity report one-year after the fact is actually the truth. Altogether, this means that each period t the firm gets the new observation

$$(z^0_{t-1}, z^1_{t-1}, z^2_{t-1}, z^3_{t-1}, z_{t-5})$$

When we assume agents face data uncertainty a simple filtering framework is required. The optimal linear predictor $z_{t|t-1} \equiv \hat{\mathbb{E}}(z_t|\mathcal{Y}_{t-1})$ is the solution to this filtering problem, which is put in Appendix 1.A.8 in the interest of space. The firm uses the forecast $z_{t|t-1}$ of time t productivity to set the log real wage, which elicits labor according to (1.20). The estimated impulse response function is plotted in Figure 1.9.

When the firm recognizes data revisions within the first year but sees true productivity after that, the impulse response of labor to true productivity is hump-shaped. The gradual response to a productivity shock is tied to the dynamic nature of the model.

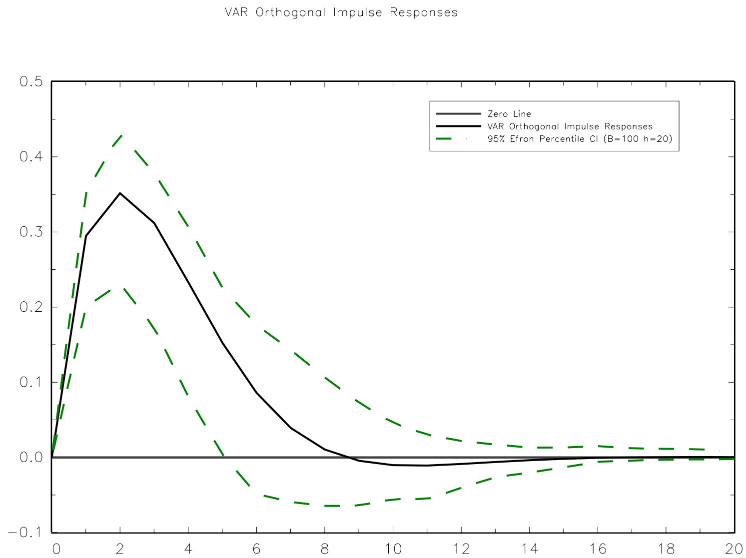

Figure 1.9: Impulse Response Function: Data Uncertainty, No Ignored Revisions
Note: Impulse response function of simulated log hours fluctuations to log productivity fluctuation shocks; as a proportion of a one-standard-deviation shock. Productivity is current (2007) vintage data. Hours are simulated assuming (1) the firm *has* data uncertainty, and (2) there *are no* data revisions the firm ignores. From a bivariate VAR with 2 lags (chosen by BIC). The horizontal axis is in quarters.

To see this, assume for example that $\rho_z \in (0,1)$, $\mu_z = 0$, and there is only one mean zero revision one quarter after the fact. At the end of period $t-1$ the report z^0_{t-1} arrives and the firm knows that some of the report is revision error. The firm decomposes the difference

$$\begin{aligned} z^0_{t-1} - \rho_z z_{t-2|t-1} &= z_{t-1} - \nu^1_t + \rho_z z_{t-2} \\ &= \rho_z z_{t-2} + \epsilon_{t-1} - \nu^1_t + \rho_z z_{t-2} \\ &= \epsilon_{t-1} - \nu^1_t \end{aligned}$$

So some of this difference is due to an actual productivity shock and some is due to the yet-to-be-seen revision. The optimal linear predictor assigns portions of this difference to a shock prediction $\epsilon_{t-1|t-1}$ and a revision prediction $\nu^1_{t|t-1}$ based on the proportion of each random variable's variance to the the sum of their variances.[1.30] If the variance of the revision error was σ^2 and the variance of the productivity shock was $\kappa\sigma^2$, then the update would give us that

$$\begin{aligned}
\epsilon_{t-1|t-1} &= \frac{\kappa}{1+\kappa}(z^0_{t-1} - \rho_z z_{t-2|t-1}) \\
\nu^1_{t|t-1} &= -\frac{1}{1+\kappa}(z^0_{t-1} - \rho_z z_{t-2|t-1})
\end{aligned}$$

That is, if the difference is positive some of it is predicted to come from a positive shock and some is predicted to come from a negative revision (since the revision enters negatively). Note that $z_{t-1|t-1} = z^0_{t-1} + \nu^1_{t|t-1}$ and the time t labor demand is $\rho_z z_{t-1|t-1}$. The labor demand for time t is increasing in ϵ_{t-1}.

At the end of period t the firm gets z^1_t, which in our example is z_{t-1}. The firm's prediction error turns out to be

$$z_{t-1} - z_{t-1|t-1} = \frac{\kappa}{1+\kappa}\nu^1_t + \frac{1}{1+\kappa}\epsilon_{t-1}$$

As written above, the prediction error is how far the firm under-predicted the true time $t-1$ productivity, and we can see that this under-prediction is increasing in ϵ_{t-1}. Thus a positive ϵ_{t-1} will lead the firm to under-predict time t productivity. Once the firm sees this, it will increase its prediction $z_{t|t}$ which will lead to a larger $z_{t+1|t}$ and, thus, an increased labor demand for time $t+1$. Thus, labor demand for time $t+1$ is increasing in ϵ_{t-1}. This means that in this example the impulse response does not follow the monotonic decay of the benchmark model.

[1.30] The following can be calculated as part of a Kalman updating step – see equation (1.37).

The full-blown linear prediction problem works in the same way. With more revisions, the productivity shock affects productivity predictions further into the future (in this example, its effect stops after two periods because the truth is seen after two periods). However, its impact on future productivity predictions diminishes geometrically according to the autoregressive parameter ρ_z. The second-quarter peak is a result of these offsetting effects.

1.4.3.4 Data Uncertainty and Ignored Data Revisions

This simulation combines the previous two experimental designs' main features: restricting firms to see only a year's worth of real-time data, and allowing them to recognize that data will be revised. The experimental design is:

> *Suppose the firm sees a year's worth of productivity measures – at the end of the year, it sees the productivity value as reported at the year horizon, which can be different from true productivity. The firm believes the structure estimated in Tables 1.5 and 1.6 governs productivity and revisions.*

The firm enters period t knowing multiple vintages of real time data on productivity for the past year. The longest horizon productivity report is not necessarily the truth – the firm ignores revisions to the data at horizons greater than one year. Altogether, this means that each period t the firm gets the new observation

$$(z^0_{t-1}, z^1_{t-1}, z^2_{t-1}, z^3_{t-1}, z^4_{t-1})$$

and in general $z^4_{t-1} \neq z_{t-5}$.

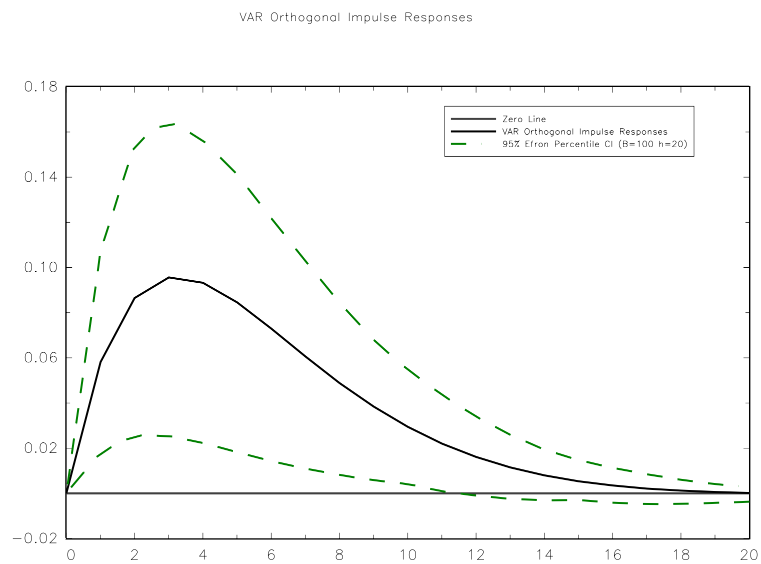

Figure 1.10: Impulse Response Function: Data Uncertainty, Ignored Revisions
Note: Impulse response function of simulated log hours fluctuations to log productivity fluctuation shocks; as a proportion of a one-standard-deviation shock. Productivity is current (2007) vintage data. Hours are simulated assuming (1) the firm *has* data uncertainty, and (2) there *are* data revisions the firm ignores. From a bivariate VAR with 1 lag (chosen by BIC). The horizontal axis is in quarters.

The firm uses the forecast $z_{t|t-1}$ of time t productivity to set the log real wage, which elicits labor according to (1.20). The estimated impulse response function is plotted in Figure 1.10. The third- and fourth-quarter responses are very similar, as in the impulse response function in Figure 1.5.

1.4.4 Discussion

The full model's impulse response function shows a promising place for data uncertainty and real-time data as part of the explanation of hump-shaped responses that are typically estimated in the data. There are several attractive features to the model. First, the additional equations it requires – relative to the benchmark model – are immediately suggested by the data already under consideration. Second, the parameters governing this informational friction are estimable directly from the data. Third, the effect of revisions to past data is not a result of a stylized state vector where firm behavior depends, for some reason, on past conditions; the effect of revisions comes as a result of their effect on today's best guess of the current economic state. This last feature suggests that the results of this simple model could be found in more complicated frameworks.

This model is obviously not the first to incorporate informational frictions in order to deliver labor response propagation. The sticky information models of Mankiw and Reis (2002) and Mankiw, Reis, and Wolfers (2003) are recent papers that have also done so. In Mankiw and Reis (2002) the "inattentiveness" of firms is assumed such that they update their information once a year. In Mankiw, Reis, and Wolfers (2003) the stickiness is set by matching the model predictions to observed interquartile ranges in various notable surveys of forecasts.

The similarity between my model and those papers' is that, roughly speaking, the agents are unable to perfectly see the relevant state variables. Because of this uncertainty agents' response to real changes in the economy is prolonged. A key difference is that in my model the firm's information set is constantly updating. The informational friction is introduced directly from the data series that the frictionless model already incorporated, so secondary data sources or exogenous calibration is unnecessary.

This experiment shows the promise of explaining propagation effects found in the data by modeling the information limitation represented by real-time data and data uncertainty.

1.5 Conclusion

This paper shows that data uncertainty warrants explicit modeling. In two different models this uncertainty comes from the actual revision processes apparent in the relevant data series. I extend the Federal Reserve model of Sargent, Williams, and Zha (2006) and show that accounting for unemployment data uncertainty in the Fed's optimization leads to model-predicted unemployment rate forecasts that closely resemble actual Fed forecasts found in the Greenbook. Additionally, the model highlights the fact that model uncertainty can be largely overestimated when even small amounts of data uncertainty are present but ignored. In a separate model, the predicted labor response to productivity shocks is hump-shaped when the firm observes real-time data on labor productivity and accounts for data uncertainty when optimizing.

Here I take the data seriously insofar that changes to historical data constitute events of interest. The points made here, and the framework suggested, carry over to situations where agents purchase forecasts constructed from information outside the agent's information set. There are many such transactions. Forward-looking agents obtain these outside forecasts well in advance of the event-date, allowing the possibility that these forecasts be revised as the event-horizon diminishes. The agent might have information that would allow prediction of these revisions, as is possible with some macroeconomic data revisions. Future work can thus explore responses to purchased or published forecasts that change over time.

1.A Appendix

1.A.1 Including Data Revisions

Generally, linear dependence between revisions, collected in a vector, can be expressed by the finite order matrix lag polynomial $\boldsymbol{B}(L)$. The lag operator is understood to operate on the subscripts, as is usual. In this case, the general linear dependence structure can be written

$$\boldsymbol{B}(L)\begin{pmatrix} \nu_{t+1}^{1} \\ \nu_{t+2}^{2} \\ \vdots \\ \nu_{t+f}^{f} \end{pmatrix} \equiv \boldsymbol{B}(L)\boldsymbol{\nu}_t = \tilde{\boldsymbol{\nu}}_t$$

$$\mathbb{E}(\tilde{\boldsymbol{\nu}}_t) = \mathbf{0}, \quad \mathbb{E}(\tilde{\boldsymbol{\nu}}_t\tilde{\boldsymbol{\nu}}_t') = \boldsymbol{\Sigma}_{\tilde{\nu}}$$

The above allows for dependence in the revisions of the same horizon (ν_t^j, ν_{t+1}^{j+1}, etc.) for different period variables – this is the kind of dependence found by Aruoba (2005). It also allows for dependence between revisions to the same period variable (ν_t^j, ν_{t+1}^{j+1}, etc.) in two ways. The first way is through the off-diagonal elements of the $\boldsymbol{B}$ matrices. This kind of dependence structure is assumed to not exist in the above models, which means that the $\boldsymbol{B}$ matrices considered are always diagonal (moreover, I always assume an AR(1) structure to revisions, so there is only one $\boldsymbol{B}$ matrix considered for each model).

The second way to account for dependence between revisions to the same period variable is through the covariance matrix $\boldsymbol{\Sigma}_{\tilde{\nu}}$. If this matrix is not diagonal, there is comovement between the shocks to revisions to the same period variable, which translates into comovement between the revisions themselves. I allow for this

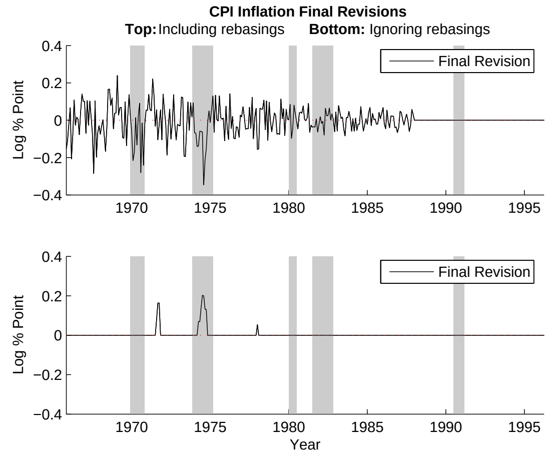

Figure 1.11: Impact of Rebasing on Revisions

Note: NBER recessions are shaded.

above, in Section 1.3 by the matrix $\tilde{\boldsymbol{V}}$, and in Section 1.4 by the nonzero part of the matrix $\boldsymbol{Q}$.

1.A.2 Data Revisions

This section analyzes data revisions for inflation and unemployment.[1.31] I begin with inflation.

Consider the difference between the current vintage and the preliminary vintage of CPI inflation data. Calling this our final revision, we have the top panel of Figure 1.11. Looking only at this panel would lead us to believe that inflation data is heavily revised. However, this exercise is a misleading look at the inflation revision process because it does not account for the periodic rebasing of the index: our clue that this is the case is found in near absence of revisions from 1988 onwards.

[1.31] The following statements apply to non-seasonally-adjusted CPI numbers. Since I am measuring inflation as an annual growth rate of the index, seasonal adjustment in theory should not bear upon the inflation measure; in practice, however, seasonal adjustment affects even annual growth rates, and so might experience revision characteristics different than what I report. I use the non-seasonally-adjusted numbers as suggested by Ghysels, Swanson, and Callan (2002).

As an index, the CPI is conventionally reset to 100; in fact, this was done in 1988. Values of the index *prior to* the rebase period are scaled. Logarithmic rates are invariant to such a procedure alone; however, index values are then rounded after scaling, which breaks the invariance.

If look for data revision only within each base period, we get the bottom panel of Figure 1.11. This illustrates that data revision, properly ignoring changes due to arbitrary rebasings, is not an issue for CPI inflation.

A larger issue that this figure brings to light is that economic data coming from indices are changed over time due solely to the (rather arbitrary) choice to change the base year of the index. What are the accurate observations of these historic economic phenomena? I here ignore the revisions due to changing the base period of the index: this leads me to model inflation's preliminarily reported value as the true value of inflation and hence to say inflation is not subject to a revision structure like what is developed below for unemployment.

The civilian unemployment rate is not an index, hence is immune to basing issues. Thus, I assume changes to the data represent true data revision. Figure 1.12 displays the final revision as a percentage of the final value, giving an idea of the relative importance of this revision error to the underlying value. Apparently data revision is an issue for the unemployment rate: its Root Mean Squared Error is 2.1% of the final vintage value.[1.32]

[1.32]Figure 1.12 shows the revision as a percentage of the final vintage value: for example, if the final vintage unemployment rate was 4%, a 20 basis-point revision would represent 5% relative revision. I show relative revisions to account for the fact that though a 20 basis-point revision is sizeable when the unemployment rate is 4%, it might be of less importance when the unemployment rate level is

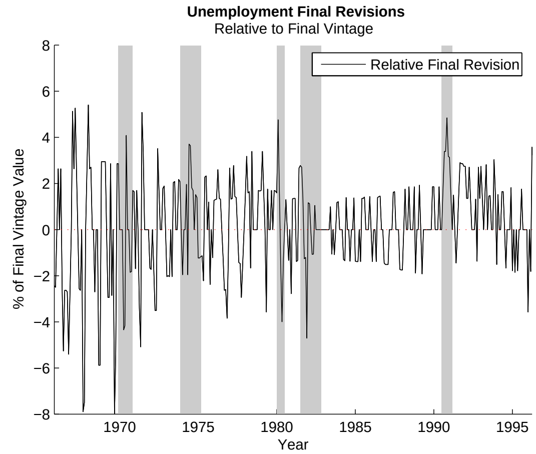

Figure 1.12: Relative Final Revisions: Unemployment
Note: RMSE = 2.099. NBER recessions shaded.

As a practical matter, I assume that data revision ends at some finite time so that the state vector is finite-dimensional. In reality, unemployment's revisions occur even after prohibitively (empirically-speaking) long intervals. Figure 1.13 shows the revisions missed by assuming the final revision occurs 24, 48, 72, or 120 months after the fact. The data shows that it is rather safe to assume the final revision to occur 10 years after the fact, and not very problematic to assume that revisions end after 72 months.

There is no evidence of a time-series relationship between first, second, and third revisions, but some evidence of a relationship between final revisions. To analyze this, I propose the autoregression

$$\nu_t^{72} = \gamma_0 + \gamma_1 \nu_{t-1}^{72} + error_t \tag{1.21}$$

at 12%.

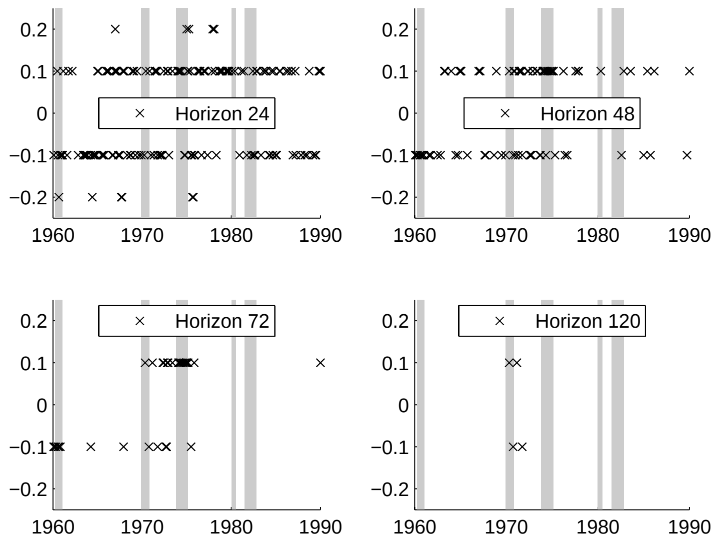

Figure 1.13: Ignored Unemployment Rate Data Revisions by assuming Final Vintage Horizon

Note: Data revisions that are ignored by assuming the final vintage of data exists. An ignored data revision is the difference between the current data vintage and the final data vintage for varying final vintage horizons.

Over the period March 1960 to December 1995:

$$\nu_t^{72} = \underset{(0.0053)}{-0.0031} + \underset{(0.0555)}{0.3873}\ \nu_{t-1}^{72}, \qquad R^2 = 0.15,\ N = 430$$

Robust standard errors in parentheses

Lags past the first (results not reported) are statistically insignificant at the 10% level. These results show that there is some predictability of the final revisions over the entire time span. How has this predictability appeared in real-time?

I estimate (1.21) on expanding windows starting in January 1961 (when there are 10 observations). The resulting time paths of $\hat{\gamma}_0$, $\hat{\gamma}_1$ with two-standard-error bands and the associated R^2 of each regression are plotted in Figure 1.14. Note that the coefficients are aligned such that these regression results, on final data, would have been available at the time on the x-axis.

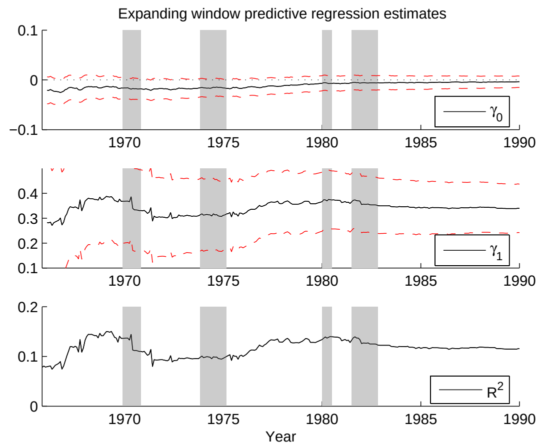

Figure 1.14: Rolling window regression estimates

Note: Estimates from regression (1.21) run over expanding windows. NBER recessions are shaded.

These figures indicate that the relationship (1.21) is rather stable: $\hat{\gamma}_0 \approx 0$ and $\hat{\gamma}_1 \approx 0.35$ with an $R^2 \approx 0.12$ over the time span. This observation guides my choice of which parameters to estimate (and which to impose are zero to reduce the number of estimated parameters) in the transition equation specified in Section 1.3.

In summary, I find inflation to be accurately reported by its preliminary values and so do not model its revision process. The unemployment rate, however, does see revisions. There is a moderately stable predictive relationship between final revisions but none between the others, leading me to estimate in Section 1.2 the autocorrelation (governed by $\breve{g}$) between final revisions and assume no autocorrelation between revisions of other horizons.

1.A.3 Federal Reserve Model Comparison

I have dropped Sargent, Williams, and Zha (2006)'s Lucas natural-rate Philips curve from my model presentation. The estimation of the Sargent, Williams, and

Zha (2006) Lucas natural-rate Philips curve does not affect the estimates of the belief-formation parameters apart from ζ_2 ($\boldsymbol{V}, \boldsymbol{P}_{1|0}, \zeta_1$) because the relevant priors are independent and, otherwise, the likelihood does not tie the two together. However, ζ_2 is unidentified by the model without data uncertainty.

Sargent, Williams, and Zha (2006) deals with this problem by normalizing it such that $\frac{1}{\zeta_2}$ is one-tenth the standard deviation of the shock in the Lucas natural-rate equation. In practice, that paper's assumption implies that $\zeta_2 = 59.7108$ which means that the Fed thinks that, if it knew the parameters, the Philips curve would forecast unemployment up to an exogenous error with standard deviation $\frac{1}{59.7108} = 0.0167$. That paper describes the assumption as implying that the Fed believes "that the standard deviation of the [its] regression error is smaller by a factor of ten than the standard deviation exogenous unemployment shocks."

Since the Philips curve is used by the Fed to forecast its affects on unemployment, consider a very common forecasting rule: the random-walk unemployment forecast. I can make an assumption relative to random-walk unemployment forecast errors that is similar to Sargent, Williams, and Zha (2006)'s assumption relative to Lucas natural-rate Philips curve shocks. This assumption is: "the Fed believes that the standard deviation of its Philips curve forecast error is smaller by a factor of ten than the standard deviation of random-walk unemployment forecast errors." This assumption in practice implies that $\zeta_2 = 54.9753$ since the standard deviation of random-walk unemployment forecast errors is 0.1819. Hence I arrive at choice for ζ_2 that is similar to Sargent, Williams, and Zha (2006)'s, but without the Lucas natural-rate Philips curve.

To ensure that my results are robust to these factors, I have done the following. I have reproduced the results of Sargent, Williams, and Zha (2006)'s entire model. I have estimated the Section 1.3.1.1 model while setting $\zeta_2 = 59.7108$ (calculated from Sargent, Williams, and Zha (2006)'s estimate of 35.6538 as the precision of the shock to the Lucas natural-rate Philips curve). And I have estimated the Section 1.3.1.1 model by using $\zeta_2 = 54.9753$ based on assumed Fed beliefs as to its forecasting ability relative to random-walk forecasts. The results (parameter estimates, beliefs, inflation choices, and unemployment forecasts) are virtually the same.[1.33] In order to provide a clear comparison to previous literature, I report the results from simply setting $\zeta_2 = 59.7108$ in Section 1.3.3.

It might be supposed that the estimate of $\boldsymbol{V}$ is large in Section 1.3.3 and small in Section 1.3.4 because ζ_2 is set to about 50 in Section 1.3.3 and estimated at around 5 in Section 1.3.4. However, this is not the case. As Sargent, Williams, and Zha (2006) note, in the Section 1.3.3 model $\boldsymbol{V}$ and $\boldsymbol{P}_{1|0}$ are related to $\frac{1}{\zeta_2^2}$ such that both can be scaled by $\kappa > 0$ with no effect on the likelihood. Let $\kappa = 10$ and notice that scaling ζ_2 by $\frac{1}{\kappa}$ is the same as scaling $\frac{1}{\zeta_2^2}$ by κ^2. Therefore, scaling ζ_2 by $\frac{1}{10}$ (making ζ_2 smaller) means that $\boldsymbol{V}$ and $\boldsymbol{P}_{1|0}$ could be scaled by 10^2 with no effect on the likelihood. That is, in the model without data uncertainty the decrease in ζ_2 would actually work in the opposite direction from what is being supposed – a smaller

[1.33] The value of log-likelihood (multiplied by the prior) for my reproduction of Sargent, Williams, and Zha (2006) is 555.0586, comparable to the 564.92 they find. The log-likelihood value in Section 1.3.3 comes as a result of the likelihood of that model being only a function of the inflation control equation.

ζ_2 in the model without data uncertainty implies a *larger* $\boldsymbol{V}$ at the same likelihood value. This means that the difference in ζ_2 is not driving the different estimates of $\boldsymbol{V}$.

1.A.4 Extended Kalman Filter and Likelihood

Extended Kalman Filter To implement, we first approximate the state space model by first-order expansions.[1.34] Let $\mathbb{E}(\boldsymbol{\beta}_t|\mathcal{Y}_s) \equiv \boldsymbol{b}_{t|s}$ and $\mathbb{V}ar\,(\boldsymbol{\beta}_t|\mathcal{Y}_s) \equiv \Sigma_{t|s}$ for $\mathcal{Y}_s \equiv \sigma(\boldsymbol{y}_s, \boldsymbol{y}_{s-1}, \ldots)$. The expansion of (1.17) about $(\boldsymbol{\beta}_t, \boldsymbol{\varepsilon}_t) = (\boldsymbol{b}_{t|t-1}, 0)$ is exact:

$$\boldsymbol{h}_t(\boldsymbol{\beta}_t, \boldsymbol{\varepsilon}_t) = \boldsymbol{H}\boldsymbol{b}_{t|t-1} + \boldsymbol{H}(\boldsymbol{\beta}_t - \boldsymbol{b}_{t|t-1}) + \boldsymbol{\varepsilon}_t \tag{1.22}$$

The expansion of (1.8) about $(\boldsymbol{\beta}_{t-1}, \boldsymbol{\eta}_t) = (\boldsymbol{b}_{t-1|t-1}, 0)$ is approximate:

$$\boldsymbol{g}_t(\boldsymbol{\beta}_{t-1}, \boldsymbol{\eta}_t) \approx \boldsymbol{g}_{t|t-1} + \boldsymbol{T}_{t|t-1}(\boldsymbol{\beta}_t - \boldsymbol{b}_{t-1|t-1}) + \boldsymbol{R}_{t|t-1}\boldsymbol{\eta}_t \tag{1.23}$$

where

$$\begin{aligned}
\boldsymbol{g}_{t|t-1} &= \boldsymbol{g}_t(\boldsymbol{b}_{t-1|t-1}, 0) \\
\boldsymbol{T}_{t|t-1} &= \left.\frac{\partial \boldsymbol{g}_t(\boldsymbol{\beta}_{t-1}, \boldsymbol{\eta}_t)}{\partial \boldsymbol{\beta}'_{t-1}}\right|_{(\boldsymbol{b}_{t-1|t-1}, 0)} \\
\boldsymbol{R}_{t|t-1} &= \left.\frac{\partial \boldsymbol{g}_t(\boldsymbol{\beta}_{t-1}, \boldsymbol{\eta}_t)}{\partial \boldsymbol{\eta}'_t}\right|_{(\boldsymbol{b}_{t-1|t-1}, 0)}
\end{aligned}$$

I motivate the derivation of optimal prediction and updating for the approximating system (1.22) and (1.23) by assuming Gaussian shocks, as in Howrey (1978), Watson and Engle (1983), and Harvey (1989).[1.35] In this case, the relevant conditional

[1.34] I experimented with the Second Order Extended Kalman filter, but found that the numerical second derivatives of the optimal policy rule were both inaccurate and computationally burdensome.

[1.35] Technically, I must assume that the vector $\boldsymbol{\eta}_t$ appearing in (1.23) is Gaussian; assuming a Gaussian $\boldsymbol{\eta}_t$ for the general nonlinear case (1.8) does not assure that the shock in the first-order expansion would be Gaussian.

expectations have the known forms given below. In particular, I assume

$$\boldsymbol{\eta}_t \sim iid\,\mathcal{N}(\mathbf{0}, \boldsymbol{Q}), \qquad \boldsymbol{\varepsilon}_t \sim iid\,\mathcal{N}(\mathbf{0}, \boldsymbol{N}), \qquad \boldsymbol{\eta}_t \perp \boldsymbol{\varepsilon}_\tau, \forall t, \tau$$

The Extended Kalman Filtering equations are

$$\boldsymbol{b}_{t|t-1} = \boldsymbol{g}_{t|t-1} \tag{1.24}$$

$$\boldsymbol{\Sigma}_{t|t-1} = \boldsymbol{T}_{t|t-1}\boldsymbol{\Sigma}_{t-1|t-1}\boldsymbol{T}'_{t|t-1} + \boldsymbol{R}_{t|t-1}\boldsymbol{Q}\boldsymbol{R}'_{t|t-1} \tag{1.25}$$

$$\boldsymbol{y}_{t|t-1} = \boldsymbol{H}\boldsymbol{b}_{t|t-1} \tag{1.26}$$

$$\boldsymbol{F}_{t|t-1} = \boldsymbol{H}\boldsymbol{\Sigma}_{t|t-1}\boldsymbol{H}' + \boldsymbol{N} \tag{1.27}$$

$$\boldsymbol{M}_{t|t-1} = \boldsymbol{H}\boldsymbol{\Sigma}_{t|t-1} \tag{1.28}$$

$$\boldsymbol{K}_t = \boldsymbol{M}'_{t|t-1}\boldsymbol{F}^{-1}_{t|t-1} \tag{1.29}$$

$$\boldsymbol{\Sigma}_{t|t} = \boldsymbol{\Sigma}_{t|t-1} - \boldsymbol{K}_t\boldsymbol{F}_{t|t-1}\boldsymbol{K}'_t \tag{1.30}$$

$$\boldsymbol{b}_{t|t} = \boldsymbol{b}_{t|t-1} + \boldsymbol{K}_t(\boldsymbol{y}_t - \boldsymbol{y}_{t|t-1}) \tag{1.31}$$

Conditional on the data $\mathcal{Y}_T$, parameters $\{\boldsymbol{G}, \boldsymbol{H}, \boldsymbol{Q}, \boldsymbol{N}\}$, and initial conditions $\{\boldsymbol{b}_{1|0}, \boldsymbol{\Sigma}_{1|0}\}$, the sequences of left hand side variables (1.24)-(1.31) are found by matrix multiplication.

$\boldsymbol{Q}$ is a block diagonal matrix composed of $\boldsymbol{V}, (1/\zeta_1)^2, (1/\zeta_2)^2, \tilde{\boldsymbol{V}}$, and zeros otherwise. $\boldsymbol{N}$ only has one nonzero entry: $N^{(2,2)} = (1/\zeta_\varepsilon)^2$. $\boldsymbol{b}_{1|0}$ is comprised of the true values of the revisions and economic variables, along with the estimate of the time-varying parameters from the model without data uncertainty for the starting time period of the model including data uncertainty.

$\boldsymbol{\Sigma}_{1|0}$ is a block diagonal matrix composed of $\boldsymbol{P}_{1|0}$ and the remaining elements specifying the initial uncertainty over the revisions and unemployment values. Three of these elements specify the initial government one-step-ahead uncertainty over the yet-to-be-seen data revisions (assumed to be a diagonal matrix), declaring that the uncertainty be the same for revisions of same type (for instance, ν_t^1 and ν_s^1) but possibly different for revisions of different type (i.e. ν_t^1, ν_t^2, and ν_t^f). The remaining three specify the initial one-step-ahead uncertainty about latent unemployment values.

Likelihood The likelihood is

$$\mathcal{L}(\mathcal{Y}_T|\mathbf{\Psi}) \propto \prod_{t=1}^{T} \left|\boldsymbol{F}_{t|t-1}\right|^{-1/2} \exp\left\{-\frac{1}{2}(\boldsymbol{y}_t - \boldsymbol{y}_{t|t-1})'\boldsymbol{F}_{t|t-1}^{-1}(\boldsymbol{y}_t - \boldsymbol{y}_{t|t-1})\right\}$$

where the $\boldsymbol{F}_{t|t-1}$ come from (1.27).

1.A.5 Greenbook Forecasts

The main issue with comparing the model predicted unemployment forecasts to Greenbook forecasts is the difference in the frequency of observation. The model forecasts the monthly unemployment rate one month into the future. The Greenbook forecasts quarterly unemployment rates and are released without rigid frequency. For example, there are Greenbook forecasts published monthly through the 1970s, but into the 1980s and 1990s these forecasts are published almost at a bimonthly frequency. I take the following steps to make the comparison.

First, I form a quarterly unemployment rate series as the average of unemployment rate for the three underlying months. It is against these series that the forecasts produce forecast errors.

Second, I form a quarterly model-forecast series as the average of the step-ahead forecasts for the three underlying months. That is, the model's quarterly unemployment forecast for quarter q composed of months m_1, m_2, m_3 is

$$\frac{1}{3}(u_{m_1|m_1-1} + u_{m_2|m_1} + u_{m_3|m_2})$$

where $u_{j|j-1}$ is the forecast made at time $j-1$ pertaining to time j

Third, I form the quarterly Greenbook-forecast series as an average of all the forecasts made the month before or anytime during a quarter. That is, the Greenbook quarterly forecast for q composed of months m_1, m_2, m_3 and immediately preceded by month m_0 is

$$\frac{1}{n_{obs}}(gb_{m_0} + gb_{m_1} + gb_{m_2} + gb_{m_3})$$

where gb_j is the Greenbook forecast for quarter q published in month j. It should be noted that all four of these forecasts do not exist for every quarter, in which case only those observed are summed and n_{obs} adjusts to however many forecasts *are* observed.

The Diebold and Mariano (1995) statistic S_1 takes forecast error series $\{e_{it}\}$ and $\{e_{jt}\}$

$$S_1 = \frac{\bar{d}}{\sqrt{\frac{2\pi \widehat{f_d(0)}}{T}}}$$

where

$$\bar{d} = \frac{1}{T}\sum_{t=1}^{T}(e_{it} - e_{jt})$$

and I take $\widehat{f_d(0)}$ to be Andrews (1991) quadratic-spectral HAC estimator. The errors under consideration run from 1970 through 1995 so that $T = 104$. The forecast errors

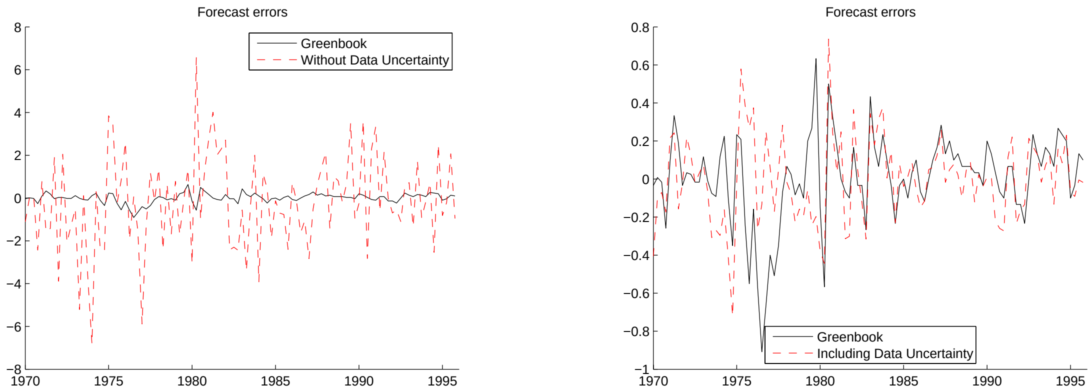

Figure 1.15: Forecast errors entering equal-accuracy test

from the Greenbook and the two models are graphed in Figure 1.15 (note different vertical scales).

1.A.6 MCMC Implementation and Robustness

Priors The prior for $\mathbf{\Psi}$ is multivariate normal with a non-zero mean and a diagonal covariance matrix – so equivalently, the priors for each parameter are independent normals. The exact specifications are listed below where

$$\boldsymbol{\varphi} \equiv \left(\mathrm{vech}\left(\mathrm{Chol}\left(\boldsymbol{V}\right)\right)', \mathrm{vech}\left(\mathrm{Chol}\left(\boldsymbol{P}_{1|0}\right)\right)\right)'$$

following the notation of Sargent, Williams, and Zha (2006), and $\boldsymbol{\psi} \equiv \mathrm{vech}\left(\mathrm{Chol}\left(\tilde{\boldsymbol{V}}\right)\right)'$:

$\boldsymbol{\zeta}$ $\mathcal{N}\left(\begin{bmatrix} 5 \\ 5 \end{bmatrix}, \begin{bmatrix} \frac{5^2}{3^2} & 0 \\ 0 & \frac{5^2}{3^2} \end{bmatrix}\right)$

$\boldsymbol{\varphi}$ Follows Sargent, Williams, and Zha (2006). For each element on the diagonal of $\mathrm{Chol}\left(\boldsymbol{V}\right)$ or $\mathrm{Chol}\left(\boldsymbol{P}_{1|0}\right)$ the prior is $\mathcal{N}\left(0, 5^2 \times 0.5\right)$; for those elements off the diagonal, it is $\mathcal{N}\left(0, 2.5^2 \times 0.5\right)$

$\breve{g}$ $\mathcal{N}(0, 0.3)$

$\boldsymbol{\psi}$ The diagonal elements have mean $(\frac{1}{10}, \frac{1}{10}, \frac{1}{3})'$ and the off-diagonal elements have mean $\frac{1}{1000}$. $\mathbb{V}ar(\boldsymbol{\psi}) = \frac{1}{9} \times diag\left(\mathbb{E}\left(\boldsymbol{\psi}\right)\right)$ where the $diag(\cdot)$ makes a diagonal matrix out of its argument.

ζ_ε $\mathcal{N}(5, \frac{5^2}{3^2})$

$\boldsymbol{\Sigma}_{1|0}$ The prior on the diagonal elements associated with past unemployment values is $\mathcal{N}(0.5, \frac{1}{9} \times 0.5)$. The prior on the diagonal elements associated with yet-to-be-seen revisions have a mean depending on the type of revision (second, third, or final) corresponding to the diagonal element of $\tilde{\boldsymbol{V}}$ parameterizing the variance of that revision shock; the standard deviation of prior for these elements is one-third of the mean.

Convergence of the MCMC To address the convergence of the MCMC algorithm to its posterior distribution, I computed the number of iterations required to estimate the 0.025 quantile with a precision of 0.02 and probability level of 0.950 using the method of Raftery and Lewis (1992). For each chain (with different initial conditions) the max of these across $\boldsymbol{\Psi}$ was below the $3.5E5$ iterations taken from each chain, suggesting that mixing the two chains (after burn-in) yields satisfactory precision.

Metropolis Algorithm An important part of the MCMC algorithm sampling from the posterior in a reasonable number of iterations is the covariance matrix of the proposal random step in the Metropolis algorithm. The Metropolis algorithm is

1. Given $\boldsymbol{\Psi}^{\text{previous}}$, propose a new value

$$\boldsymbol{\Psi}^{\text{proposal}} = \boldsymbol{\Psi}^{\text{previous}} + \boldsymbol{\xi}$$

   where $\boldsymbol{\xi}$ is normal with mean zero and covariance matrix $c\boldsymbol{\Sigma}_{\xi}$

2. Compute

$$q = \min\left\{\frac{p\left(\boldsymbol{\Psi}^{\text{proposal}}|\mathcal{Y}_T\right)}{p\left(\boldsymbol{\Psi}^{\text{previous}}|\mathcal{Y}_T\right)}, 1\right\}$$

3. Randomly draw $w \sim U(0,1)$

4. If $w \leq q$, accept $\boldsymbol{\Psi}^{\text{proposal}}$ as current draw; otherwise, set $\boldsymbol{\Psi}^{\text{previous}}$ as the current draw

Given the manner in which all parameters affect the optimal policy, I arrived at this proposal covariance matrix $\boldsymbol{\Sigma}_{\xi}$ by doing the following. Using the covariance matrix for $\boldsymbol{\varphi}$ numerically solved for as described in Sargent, Williams, and Zha (2006)'s Appendix D and the prior covariance terms for all other elements of $\boldsymbol{\Psi}$ given above, the MCMC was started. For tens of thousands of iterations based on one initial condition, I considered only elements of the MCMC chain where a proposal had been accepted. From these chain elements I calculated the sample covariance matrix of the successful proposal shocks and set $\boldsymbol{\Sigma}_{\xi}$ equal to this. I tried different initial conditions and took the weighted average of the Cholesky factors of these sample covariance matrices. The tuning parameter c was adjusted to achieve an acceptance rate of around 25-35% during the first 20,000 iterations: after this, it was unadjusted, as continual

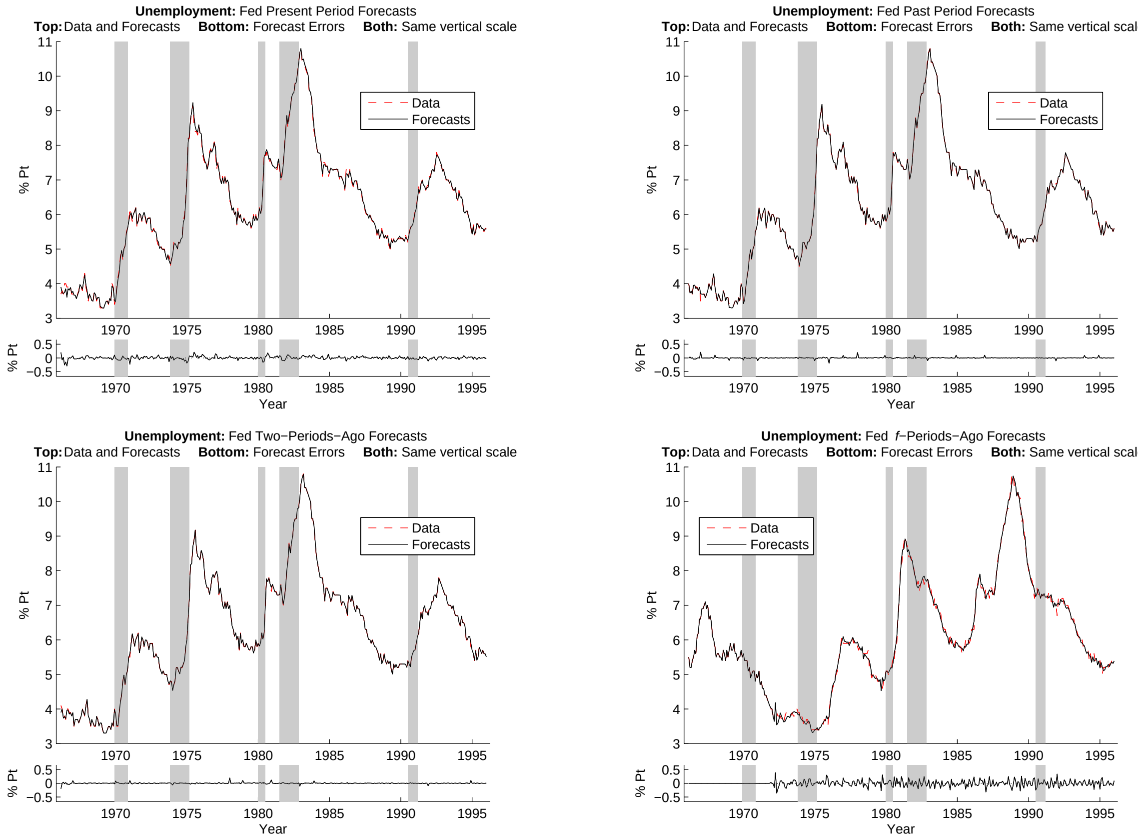

Figure 1.16: Relevant measurement vector predictions

Note: Fed unemployment now, one-step past-, two-step past-, and 72-step past-casts. NBER recessions are shaded. These figures are all on the same scale, but not on the scale of Figures 1.3.3 and 1.3.4.

chain-dependent adjustment of Metropolis step-size can negate the ergodicity upon which MCMC methods are based (see Robert and Casella (2004)).

Remaining $y_{t|t-1}$ forecasts Figure 1.A.6 shows the remaining measurement equation predictions: current, 1-lag, 2-lag, and 72-lag unemployment report forecasts.

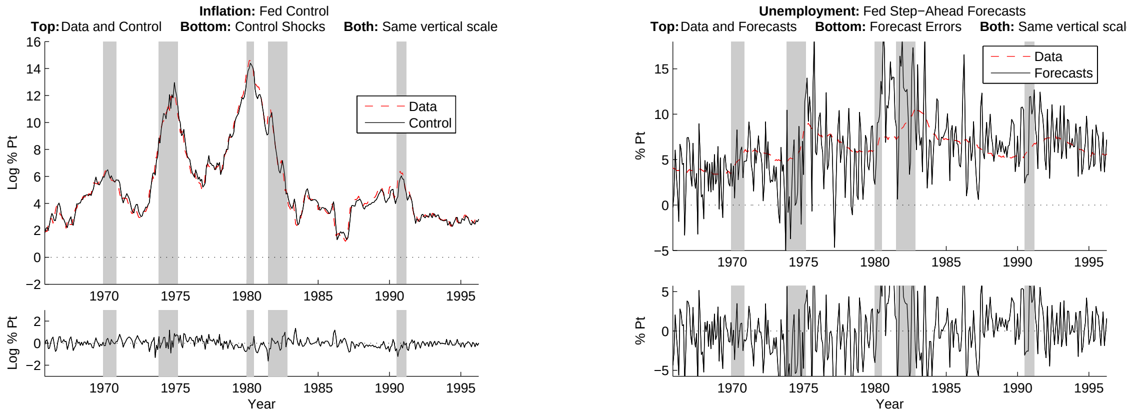

Figure 1.17: Inflation and Unemployment step-ahead forecasts: Current CPI
Note: Using current vintage seasonally-adjusted CPI inflation data. NBER recessions are shaded.

1.A.6.1 Using Current Vintage CPI Inflation and Unemployment Rate Data

The plots in Figure 1.A.6.1 are produced with inflation data as current vintage of CPI inflation (made available by those authors). The $\boldsymbol{V}$ and $\boldsymbol{P}_{1|0}$ estimates are close to those reported in Table 1.4 for the model without data uncertainty while the estimate of ζ_1 is about one-half as large – they are available upon request.

1.A.6.2 Using Preliminary Inflation and Unemployment Rate Data

The plots of Figure 1.A.6.2 are produced by instead using preliminary CPI inflation and unemployment rate data. The $\boldsymbol{V}$ and $\boldsymbol{P}_{1|0}$ estimates are close to those reported in Table 1.4 for the model without data uncertainty while the estimate of ζ_1 is about one-third as large – they are available upon request.

1.A.6.3 Modified Likelihood

It has been suggested on earlier drafts of this paper that the results may stem not so much from the data revisions as much as only the modified likelihood function

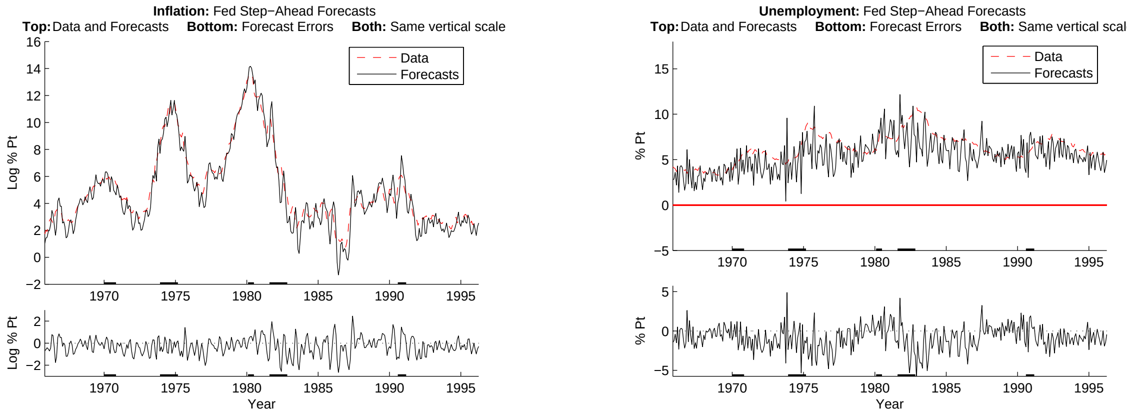

Figure 1.18: Inflation and Unemployment step-ahead forecasts: Preliminary CPI
Note: Using preliminary CPI inflation data (non-seasonally-adjusted). NBER recessions are shaded.

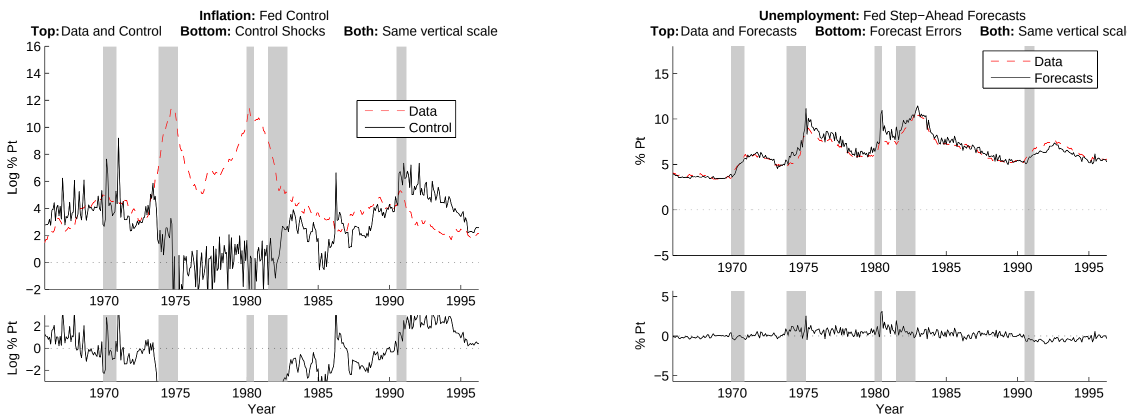

Figure 1.19: Inflation and Unemployment step-ahead forecasts: Modified Likelihood
Note: Using the model without data uncertainty and a modified likelihood placing more weight on the unemployment rate forecasts. NBER recessions are shaded.

taking account of revisions. This modification involves having the unemployment forecasts enter the likelihood, which might do most of the "smoothing" that is evident.

Doing this, we indeed see smoother Fed unemployment forecasts in the right panel of Figure 1.A.6.3. As expected, the likelihood penalizes unemployment forecast errors and delivers far better ones. However, in order to accomplish this the time-varying Philips curve estimates are such that inflation is far from target far too often, as the left panel shows.

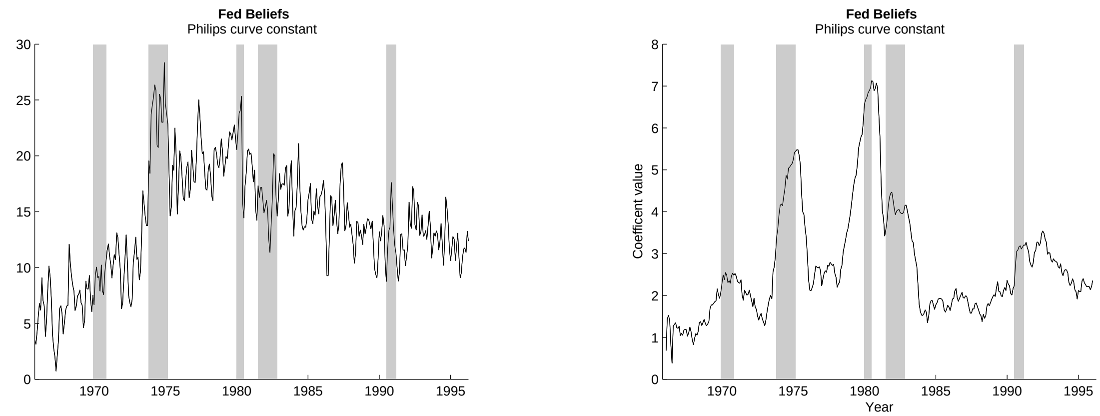

Figure 1.20: Philips curve constant estimates
Left: From model without data uncertainty. **Right:** From model including data uncertainty. NBER recessions are shaded.

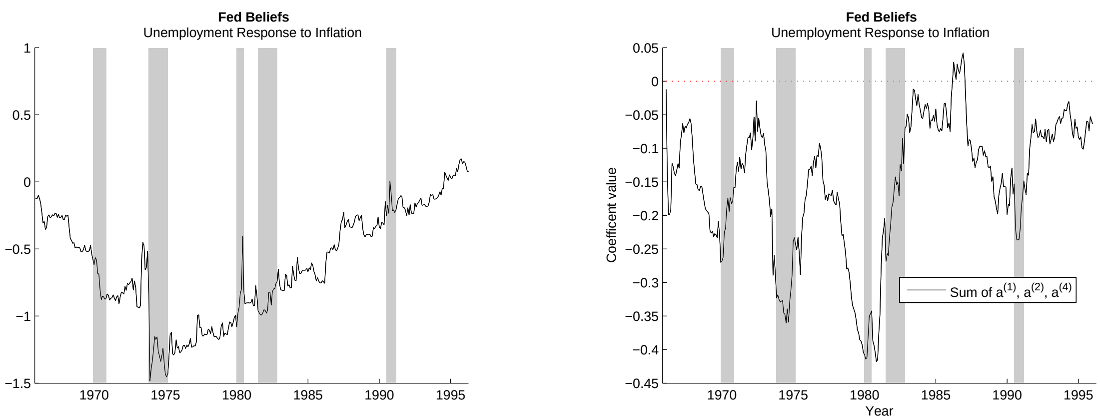

Figure 1.21: Philips curve inflation response estimates
Left: From model without data uncertainty. **Right:** From model including data uncertainty. NBER recessions are shaded.

1.A.6.4 TVP Estimates

Figure 1.A.6.4 shows the evolution of the Fed's beliefs about the constant term in the Philips curve.

Figure 1.A.6.4 shows the evolution of the Fed's beliefs about the response of unemployment to inflation. This is the sum of the coefficients on inflation in the Philips curve.

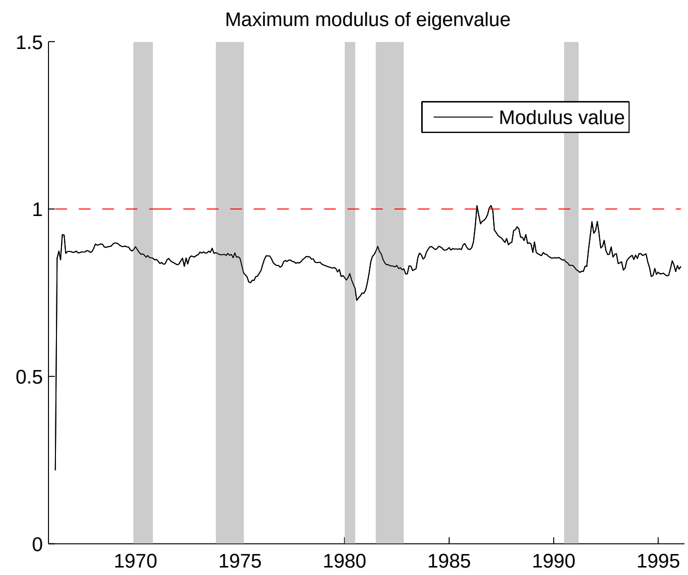

Figure 1.22: Unemployment Dynamic Instability: Model

Note: Maximum modulus of eigenvalue of characteristic polynomial of (1.2) coming from filtered time-varying parameter estimates in Section 1.2. NBER recessions are shaded.

1.A.7 Autoregressive Stability

One issue that has been raised with the estimated beliefs coming from Sargent, Williams, and Zha (2006) is that they imply autoregressive instability for more than a dozen months around 1973. The present adapted model seems to "fix" this problem during the 1970s and through the mid 1980s, but has a period of perceived instability coming at the end of 1986 – see Figure 1.22. This section explains what exactly this issue is and also argues that it is not problematic to the purposes of using the Philips curve as a state transition equation.

Consider the autoregressive structure of the Philips curve the government is estimating over time. The estimates $\boldsymbol{a}_{t-1|t-1}$ determine the perceived transition law for the Fed's optimal control problem.

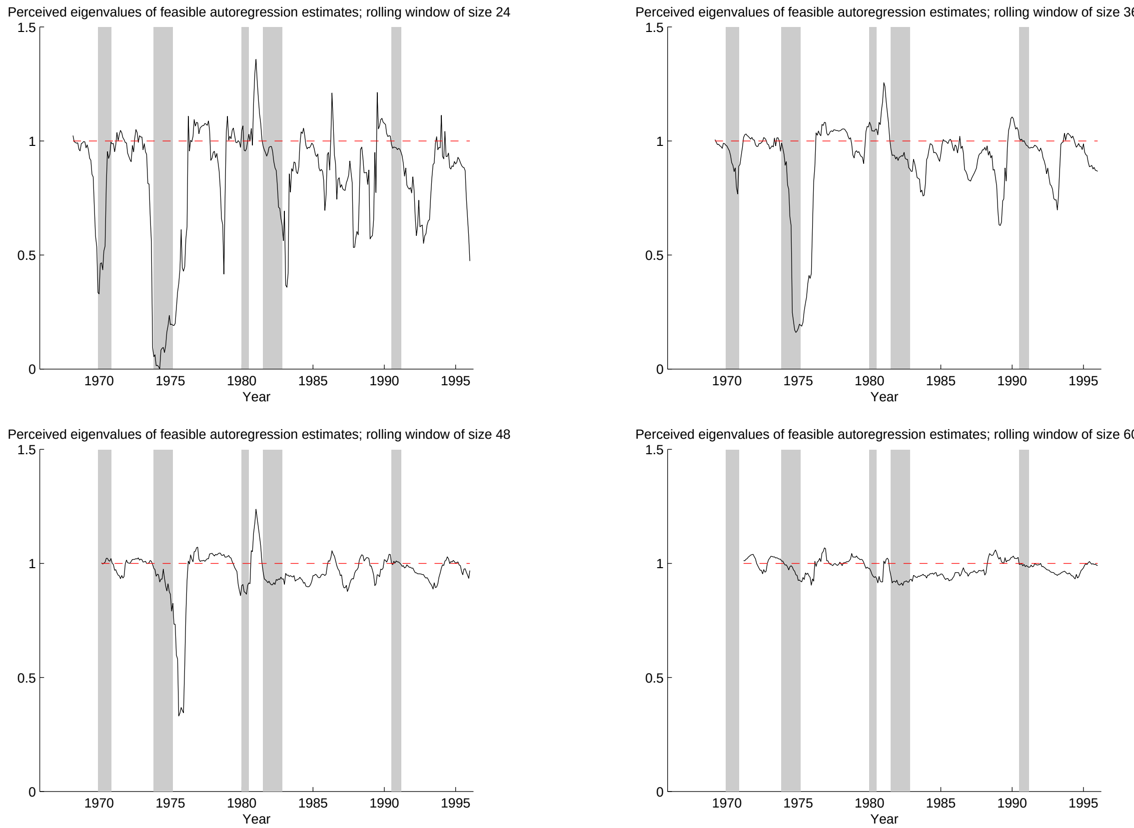

Figure 1.23: Unemployment Dynamic Instability: Rolling Autoregressions

Note: Maximum modulus of eigenvalue of characteristic polynomial of (1.2) coming from bivariate VAR; aligned so that the maximum modulus is at the time when it could be estimated with data at least f months ago. NBER recessions are shaded.

Strictly-speaking, unemployment is not an explosive process: it must certainly take on values in [0,100]. Nonetheless, its high persistence, especially during the 1970s and 1980s, certainly makes it appear not "very" covariance-stationary. Consider estimating a VAR on the final data (f periods after the preliminary data) over rolling windows of varying sizes; Figure 1.23 shows this for windows of 2, 3, 4, or 5 years.[1.36] If the Fed was concerned with breaks, a rolling window estimation procedure would be

[1.36]The window sizes are chosen with Friedman (1968) page 11 in mind: "[T]here is always a temporary trade-off between inflation and unemployment...I can at most venture a personal judgment, based on some examination of the historical evidence, that the initial effects of a higher and unanticipated rate of inflation last for something like two to five years."

a straightforward way of picking this up. Notice that this eigenvalue condition is close to the unstable region most of the time. Hence, without positing a more sophisticated mechanism of Fed belief formation, we see that simple estimation methods would also have given evidence of unemployment instability.

Nonetheless, it is not clear that autoregressive instability would have led the Fed to reject (1.2) as unemployment's dynamic structure for the purposes of setting optimal policy, given the accurate unemployment one-step-ahead predictions it yields (in the case where data uncertainty is acknowledged). To the optimal controller, what is important from forecasts are accuracy, not the description of the world they engender. The Fed had little reason to adjust its Philips curve just on the basis of these forecast errors. So the Fed may have done well to forecast unemployment using a rule which implied unemployment was not stationary, if the rule performed better.[1.37]

[1.37] For a related point, see the modeling motivation given in Primiceri (2005).

1.A.8 Production Model

Avoiding revisions stemming from trend recalculation The revised predictions of yet-to-be-seen productivity are transformed to deviations from the trend that was calculated when each of them, respectively, was the preliminary vintage value. This ensures that revisions to productivity fluctuations come only from new data and not from changes in the estimate of the trend. My procedure ensures that revisions to the productivity data exist only if revisions to the underlying output and payrolls data exist.[1.38] Here I describe the procedure.

For time t the most recent productivity data vintage is detrended using the Hodrick-Prescott filter. This delivers a productivity fluctuation z_t^0 that is inherently related to a log-productivity trend value at time t that I call z_t. By definition, for the preliminary measurement of log-productivity Z_t^0 it is the case that $\mathsf{Z}_t^0 = z_t^0 + \mathsf{z}_t$.

At time $t+1$, we get a new measurement of time t log-productivity, Z_{t+1}^1. I then define the new measurement of the time t log-productivity fluctuation as

$$z_{t+1}^1 = \mathsf{Z}_{t+1}^1 - \mathsf{z}_t$$

The following three horizons are likewise calculated. This assures that the firm only faces data uncertainty created by revisions to the underlying data, not due to trend re-estimation.[1.39]

[1.38] To the extent that these output and payrolls revisions are translated to productivity by assuming a production model, the exercise is not entirely model-free. However, I am attempting to avoid revisions due to trend calculations coming from some sort of model, which could exist as data accrues even if the preliminary data values were inviolate.

[1.39] Allowing for trend re-estimation leads to similar results as those I present because it increases the size and frequency of "revisions" stemming from differences in the underlying data *and* differences in the trend. However, it gets away from my point, which is that there are shocks directly in the data that may affect agents' behavior.

True productivity is detrended using the 2007 vintage of data. Therefore there is some difference between these measurements and the truth due to different trends. This is fine since it comprises one effect to firm behavior (with respect to true shocks) that is distinct from the effect due solely to data uncertainty, as discussed in Section 1.4.3.3.

Filtering Framework When we assume agents face data uncertainty a simple filtering framework is required. For each simulation, the period-by-period labor choice is $\lambda_{\text{Frisch}} z_{t|t-1}$ such that $z_{t|t-1}$ is given by the optimal linear predictor $\hat{\mathbb{E}}(z_t|\mathcal{Y}_{t-1})$. The firm faces a filtering problem described by the following state space model.

$$\boldsymbol{\beta}_t = \boldsymbol{g} + \boldsymbol{G}\boldsymbol{\beta}_{t-1} + \boldsymbol{\eta}_t \tag{1.32}$$

$$\boldsymbol{y}_t = \boldsymbol{H}\boldsymbol{\beta}_t \tag{1.33}$$

The state vector $\boldsymbol{\beta}_t$ holds present and past productivity values and yet-to-be-seen revisions. The constant vector $\boldsymbol{g}$ has 0s and μ_xs. The shock $\boldsymbol{\eta}_t$ is such that

$$\mathbb{E}\left(\boldsymbol{\eta}_t\boldsymbol{\eta}_t'\right) = \boldsymbol{Q}$$

and $\boldsymbol{Q}$ is block-diagonal with zeros and $\mathbb{E}\left((\varepsilon_t^f, \tilde{\nu}_t^1, \tilde{\nu}_t^2, \tilde{\nu}_t^3, \tilde{\nu}_t^4) \cdot (\varepsilon_t^f, \tilde{\nu}_t^1, \tilde{\nu}_t^2, \tilde{\nu}_t^3, \tilde{\nu}_t^4)'\right)$. The measurement vector is

$$\boldsymbol{y}_t = (n_t, z_t^0, z_t^1, z_t^2, z_t^3, z_t^4)'$$

The matrix $\boldsymbol{G}$ holds AR(1) coefficients and 1s. The matrix $\boldsymbol{H}$ forms linear forecasts of productivity observations from the predictions of past productivity and yet-to-be-seen revisions. There are no shocks to the observation equation (1.33) because

I assume there are no measurement shocks apart from the revisions, and these are part of the state. The standard Kalman filter equations give the sequence of linear predictions that solve this system:

$$\boldsymbol{b}_{t|t-1} = \boldsymbol{g} + \boldsymbol{G}\boldsymbol{b}_{t-1|t-1} \tag{1.34}$$

$$\boldsymbol{\Sigma}_{t|t-1} = \boldsymbol{G}\boldsymbol{\Sigma}_{t-1|t-1}\boldsymbol{G} + \boldsymbol{Q} \tag{1.35}$$

$$\boldsymbol{y}_{t|t-1} = \boldsymbol{H}\boldsymbol{b}_{t|t-1} \tag{1.36}$$

$$\boldsymbol{b}_{t|t} = \boldsymbol{b}_{t|t-1} + \left(\boldsymbol{H}\boldsymbol{\Sigma}_{t|t-1}\right)' \left(\boldsymbol{H}\boldsymbol{\Sigma}_{t|t-1}\boldsymbol{H}'\right)^{-1} (\boldsymbol{y}_t - \boldsymbol{y}_{t|t-1}) \tag{1.37}$$

$$\boldsymbol{\Sigma}_{t|t} = \boldsymbol{\Sigma}_{t|t-1} \tag{1.38}$$

and I initialize $\boldsymbol{\Sigma}_{0|0} = \boldsymbol{Q}$ and fill $\boldsymbol{b}_{0|0}$ with the true values of productivity, labor, and revisions.

References

Anderson, B. D. O., and J. B. Moore (1979): *Optimal Filtering*. Prentice-Hall.

Andrews, D. W. K. (1991): "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation," *Econometrica*, 59(3), 817–58.

Aruoba, S. B. (2004): "Data Uncertainty in General Equilibrium," Computing in Economics and Finance 2004 131, Society for Computational Economics.

——— (2005): "Data Revisions are not Well-Behaved," Discussion paper, University of Maryland.

Basu, S., J. G. Fernald, and M. S. Kimball (2006): "Are Technology Improvements Contractionary?," *American Economic Review*, 96, 1418–1448.

Blundell, R., and T. Macurdy (1999): "Labor supply: A review of alternative approaches," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card.

Boivin, J. (2005): "Has US Monetary Policy Changed? Evidence from Drifting Coefficients and Real-Time Data," NBER Working Papers 11314, National Bureau of Economic Research, Inc.

Burnside, C., M. Eichenbaum, and S. Rebelo (1995): "Capital Utilization and Returns to Scale," vol. NBER Macroeconomics Annual, pp. 67–110. NBER.

Christiano, L., M. Eichenbaum, and C. Evans (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113(1), 1–45.

Clarida, R., J. Gali, and M. Gertler (2000): "Monetary Policy Rules And Macroeconomic Stability: Evidence And Some Theory," *The Quarterly Journal of Economics*, 115(1), 147–180.

Comin, D., and T. Philippon (2005): "The Rise in Firm-Level Volatility: Causes and Consequences," NBER Working Papers 11388, National Bureau of Economic Research, Inc.

Cooley, T. F., and E. C. Prescott (1976): "Estimation in the Presence of Stochastic Parameter Variation," *Econometrica*, 44(1), 167–84.

Croushore, D., and T. Stark (2001): "A real-time data set for macroeconomists," *Journal of Econometrics*, 105(1), 111–130.

Cunningham, A., C. Jeffery, G. Kapetanios, and V. Labhard (2007): "A State Space Approach To The Policymaker's Data Uncertainty Problem," Money Macro and Finance (MMF) Research Group Conference 2006 168, Money Macro and Finance Research Group.

DIEBOLD, F. X., AND R. S. MARIANO (1995): "Comparing Predictive Accuracy," *Journal of Business & Economic Statistics*, 13(3), 253–63.

EDGE, R. M., T. LAUBACH, AND J. C. WILLIAMS (2004): "Learning and Shifts in Long-Run Productivity Growth," Discussion paper.

EFRON, B., AND R. J. TIBSHIRANI (1993): *An Introduction to the Bootstrap*. Chapman & Hall.

EVANS, G. W., AND S. HONKAPOHJA (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press.

FRANCIS, N., AND V. A. RAMEY (2002): "Is the Technology-Driven Real Business Cycle Hypothesis Dead?," NBER Working Papers 8726, National Bureau of Economic Research, Inc.

FRIEDMAN, M. (1968): "The Role of Monetary Policy," *The American Economic Review*, 58(1), 1–17.

GALI, J. (1999): "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?," *American Economic Review*, 89(1), 249–271.

GHYSELS, E., N. R. SWANSON, AND M. CALLAN (2002): "Monetary Policy Rules with Model and Data Uncertainty," *Southern Economic Journal*, (69), 239–265.

HARVEY, A. C. (1989): *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.

HOWREY, E. P. (1978): "The Use of Preliminary Data in Econometric Forecasting," *The Review of Economics and Statistics*, 60(2), 193–200.

KING, R. G., AND S. T. REBELO (1999): "Resuscitating real business cycles," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1 of *Handbook of Macroeconomics*, chap. 14, pp. 927–1007. Elsevier.

LORENZONI, G. (2006): "A Theory of Demand Shocks," NBER Working Papers 12477, National Bureau of Economic Research, Inc.

LUTKEPOHL, H., AND M. KRATZIG (eds.) (2004): *Applied Time Series Econometrics*. Cambridge University Press, Cambridge.

MANKIW, N. G., AND R. REIS (2002): "Sticky Information Versus Sticky Prices: A Proposal To Replace The New Keynesian Phillips Curve," *The Quarterly Journal of Economics*, 117(4), 1295–1328.

MANKIW, N. G., R. REIS, AND J. WOLFERS (2003): "Disagreement about Inflation Expectations," NBER Working Papers 9796, National Bureau of Economic Research, Inc.

OH, S., AND M. WALDMAN (1990): "The Macroeconomic Effects of False Announcements," *The Quarterly Journal of Economics*, 105(4), 1017–34.

ORPHANIDES, A. (2001): "Monetary Policy Rules Based on Real-Time Data," *American Economic Review*, 91(4), 964–985.

ORPHANIDES, A., AND J. C. WILLIAMS (2006): "Monetary Policy with Imperfect Knowledge," *Journal of the European Economic Association*, 4(2-3), 366–375.

OWYANG, M. T., AND G. RAMEY (2004): "Regime switching and monetary policy measurement," *Journal of Monetary Economics*, 51(8), 1577–1597.

PRIMICERI, G. E. (2005): "Time Varying Structural Vector Autoregressions and Monetary Policy," *Review of Economic Studies*, 72(3), 821–852.

RAFTERY, A., AND S. LEWIS (1992): "How Many Iterations in the Gibbs Sampler," .

ROBERT, C. P., AND G. CASELLA (2004): *Monte Carlo Statistical Methods*. Springer, 2 edn.

ROMER, C. D., AND D. H. ROMER (1990): "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz," NBER Working Papers 2966, National Bureau of Economic Research, Inc.

ROTEMBERG, J. J., AND M. WOODFORD (1999): "The cyclical behavior of prices and costs," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford.

RUNKLE, D. E. (1998): "Revisionist history: how data revisions distort economic policy research," *Quarterly Review*, (Fall), 3–12.

SARGENT, T. J. (1999): *The Conquest of American Inflation*. Princeton University Press.

SARGENT, T. J., N. WILLIAMS, AND T. ZHA (2006): "Shocks and Government Beliefs: The Rise and Fall of American Inflation," *American Economic Review*, 96(4), 1193–1224.

SIMS, C. A. (2006): "Improving Monetary Models," Discussion paper.

TANIZAKI, H. (1996): *Nonlinear Filters: Estimation and Applications*. Springer-Verlag, 2 edn.

WATSON, M. W., AND R. F. ENGLE (1983): "Alternative Algorithms for the Estimation of Dynamic Factor, MIMIC, and Varying Coefficient Regression Models," *Journal of Econometrics*, 23, 385–400.

Zellner, A. (1958): "A Statistical Analysis of Provisional Estimates of Gross National Product and its Components, of Selected National Income Components, and of Personal Saving," *Journal of the American Statistical Association*, 53(281), 54–65.

Chapter 2: Capital-Experience Complementarity and Demographics: Implications for the Volatility of Hours and the Great Moderation

Abstract

In the United States, the employment and hours worked of young individuals fluctuate much more over the business cycle than prime-aged individuals' employment and hours. The hypothesis in this paper is that understanding the mechanisms underlying this observation is key to explaining the volatility of aggregate hours over the cycle. We argue that the joint behavior of U.S. age-specific hours and wages point to the importance of differences in labor demand over the cycle. These considerations lead us to consider a production environment characterized by capital-experience complementarity. We structurally estimate the key parameters governing the degree of capital-experience complementarity and show that the model can account for the joint behavior of age-specific hours and wages while generating a series of aggregate hours that is as volatile as aggregate output. Finally, we utilize the model to quantify the role of the labor force's age composition in explaining the Great Moderation in the U.S., finding that it contributes about one fifth and one quarter of the decline in output and hours volatility, respectively.

Contents

List of Tables

List of Figures

2.1 Introduction

In the United States, young individuals' employment and hours worked fluctuate much more over the business cycle than do prime-aged individuals'. The hypothesis in this paper is that understanding the mechanisms underlying this observation, while interesting in its own right, has the potential to shed light on a long standing puzzle in the business cycle literature: why aggregate hours are as volatile as output over the cycle.

Our hypothesis is based on the observation that cyclical fluctuations in aggregate hours are disproportionately accounted for by young workers. For example, in the postwar era 15-29 year olds account for about one quarter of total hours worked in the U.S.; however, this same group accounts for almost one half of the the volatility of aggregate hours at the business cycle frequency. Thus, developing a quantitative theory that can account for the relative volatility of this age group is crucial to understanding the volatility of aggregate hours and, ultimately, the mechanisms that amplify and propagate business cycle fluctuations.

To maintain comparability with the real business cycle (RBC) literature, we investigate a model economy that represents a minimal deviation from the standard RBC model. One can argue that within the RBC framework, differences across age groups arise from differences in preferences (or succinctly, differences in labor supply), factors relating to technology (labor demand), or both.[2.1]

[2.1] See Rios-Rull (1996) and Gomme, Rogerson, Rupert, and Wright (2004) for models highlighting differences in labor supply owing to life-cycle considerations. They show that life-cycle mechanisms are successful at explaining volatility differences between prime-aged and old workers; however, such

How does one distinguish between these two potential channels? We suggest that the joint behavior of the cyclicality of age-specific hours and wages has the potential to shed light on this question. As we document in Section 2.2, not only is the hours volatility of young individuals greater then that of prime-age individuals in the U.S., but so too is the volatility of their wages. Any modification to the RBC framework incorporating age-specific labor supply differences alone would not be able explain this observation.[2.2] Jointly matching the behavior of hours and wages requires an important role for differences in labor demand over the cycle.

These considerations lead us to consider in Section 2.3 an environment characterized by labor demand differences due to capital-experience complementarity. The large body of literature studying capital-skill complementarity has concentrated on education as a proxy for skill (see Krusell, Ohanian, Rios-Rull, and Violante (2000), and the references therein). We concentrate on the other significant observable dimension of skill emphasized in Mincerian wage regressions, namely labor market experience. To highlight the potential in this approach, we assume that there are only 2 groups of workers, young an old; we posit that an individual's age directly determines his or her labor market experience and that production exhibits capital-experience complementarity. With this technology, differences in the cyclical demand for experienced and inexperienced labor arise naturally.

The intuition for this is straightforward. As an extreme case, suppose that

considerations cannot fully account for the volatility of young workers. See Nagypál (2004) for an alternative approach highlighting the interaction between age and worker-occupation match.

[2.2] Appendix 2.A.2 discusses this in depth.

capital and old, or experienced, labor are perfect complements, while capital and young, or inexperienced, labor display some substitutability. If capital services are a state variable and firms are profit maximizing and price-taking, any shock generating a response in inputs results in variation in only the quantity of young labor hired.

The challenge in our framework is to quantitatively account for the observed differences in hours volatility. In order to discipline our analysis, in Section 2.4 we estimate the key parameters governing the degree of capital-experience complementarity in a manner that does *not* target differences in the cyclical volatility of hours. Rather, our strategy entails estimating the parameters governing this complementarity from the model's factor demand equations and to explore the identification that emerges from the relationship between aggregate prices and quantities observed in the data.

Based on this structural estimation we simulate the model economy in Section 2.5. We find that the model generates volatilities of hours by age groups that are similar to the ones observed in the U.S. data. As a by product of this success, the model generates volatility of aggregate hours that is very close to the volatility of aggregate output. We then show that the model can account for the joint behavior of age-specific hours and relative wages. Specifically, we show in our preferred calibration that the model matches the volatility ratios of aggregate hours to output, young and old hours to output, young hours to old hours, and young wages to old wages.

As is well known, there has been significant change in the relative shares of young and old individuals in the U.S. population over the past 50 years. Combining the exogenous shift in labor force age demographics with the age group differences in the cyclical volatility of hours sheds some light the role of demographic change in the U.S. Great Moderation. We address this issue within our quantitative model in Section 2.6 where we demonstrate that variation in the age composition of the labor force manifests itself in variation of macroeconomic volatility. Specifically we show that within our quantitative model the demographic change accounts for about a quarter of the moderation in U.S. business cycle volatility observed in the past 25 years.[2.3]

Concluding remarks are provided in Section 2.7. Appendices 2.A.2 and 2.A.3 discuss the importance of a labor demand channel and capital-experience complementarity in explaining hours and wage volatility, and 2.A.4 describes the data construction.

2.2 The Cyclicality of Age-Specific Hours and Wages

2.2.1 Age-Specific Hours

In this section, we analyze the responsiveness of market work to the U.S. business cycle for data disaggregated by age. We consider both the behavior of hours worked and employment by age. The analysis of the latter is "episodic" is the sense that we estimate the unemployment rate response for various age groups to postwar U.S. recessions.

[2.3]In related work, DellaVigna and Pollet (2007) recently found that shifts in the age-composition of the consumer base predictably affect profits across industries to an economically-significant degree.

Table 2.1: Volatility of Hours Worked by Age Group, U.S.

	15 - 19	20 - 24	25 - 29	30 - 39	40 - 49	50 - 59	60+
filtered volatility	5.596	2.290	1.786	1.687	1.462	1.731	2.408
R^2	0.87	0.72	0.54	0.61	0.40	0.30	0.11
cyclical volatility	5.223	1.938	1.315	1.316	1.017	1.094	0.821
share of hours (%)	3.83	10.87	12.94	25.46	23.25	16.79	6.86
share of volatility (%)	14.38	15.13	12.22	24.07	16.98	13.19	4.04

Note: HP filtered data from the March CPS, 1968-2005. Log hours.

Our approach to studying differences in business cycle volatility by age is similar to that of Gomme, Rogerson, Rupert, and Wright (2004). We use data from the March supplement of the CPS to construct annual series for per capita hours worked from 1963 to 2005 for individuals within specific age groups. We also construct an aggregate series for all individuals 15 years and older. See Appendix 2.A.4 for detailed information on data sources used throughout the paper.

To extract the high frequency component of hours worked, we remove the trend from each series using the Hodrick-Prescott (HP) filter. Since we are interested in fluctuations at business cycle frequencies (those higher than 8 years), we use a smoothing parameter of 6.25 for annual data.[2.4]

Table 2.1 presents results on the volatility of hours worked in the U.S. for the 15-19, 20-24, 25-29, 30-39, 40-49, 50-59, and 60+ year-old age groups. The first

[2.4]Through analysis of the transfer function of the HP filter, Ravn and Uhlig (2002) find this to be the optimal value for annual data. Using a similar approach, Burnside (2000) recommends a smoothing parameter value of 6.65. Finally, see Baxter and King (1999), who recommend a value of 10, through visual inspection of the transfer function. Throughout this paper, we have repeated our analysis of annual data using the band-pass filter proposed by Baxter and King (1999), removing fluctuations less frequent than 8 years. The results are essentially identical in all cases.

row presents the percent standard deviation of the detrended age-specific series. We see a decreasing relationship between the volatility of hours worked fluctuations and age, with an upturn at the end around retirement. We are not interested in the high frequency fluctuations in these time series per se, but rather those that are correlated with the business cycle. For each age-specific series, we identify the business cycle component as the projection on a constant, current detrended output, and on current and lagged detrended aggregate hours; we refer to these as the *cyclical* hours worked series. The second row of Table 1 reports the R^2 from these regressions. This is high for most age groups, indicating that the preponderance of high frequency fluctuations are attributable to the business cycle. The exception is the 60+ age group: here a larger fraction of fluctuations are due to age-specific, non-cyclical shocks.

The third row indicates the percent standard deviation of the cyclical age-specific series. Compared to row one, the largest differences between *filtered* and *cyclical* volatilities are for those aged 60 years and up, reflecting the point made immediately above. Now we see a pattern of decreasing with age. The young experience much greater cyclical volatility in hours than the prime-aged. Moreover, the differences across age groups are large. The standard deviation of cyclical hours fluctuations for 15-19 and 20-24 year old workers is at least 5.2 and 2 times that of 50-59 year olds, respectively.[2.5]

[2.5] These results corroborate the findings of Gomme, Rogerson, Rupert, and Wright (2004), and extend them to include data from the 2001 recession. See also Clark and Summers (1981), Rios-Rull (1996), and Nagypál (2004) who document differences in cyclical sensitivity across age groups. More broadly, the literature documents differences as a function of skill; see for instance, Kydland and Prescott (1993) and Hoynes (2000), and the references therein. Note that those studies are confined to the analysis of U.S. data.

Table 2.2: Volatility of Hours Worked by Age and Gender, U.S.

	15 - 19	20 - 24	25 - 29	30 - 39	40 - 49	50 - 59	60+
filtered volatility							
female	5.196	2.181	2.030	1.855	1.608	1.918	2.885
male	6.284	2.849	1.942	1.821	1.611	2.024	3.172
cyclical volatility							
female	4.44	1.538	1.186	1.122	0.929	0.993	0.704
male	5.893	2.339	1.444	1.472	1.134	1.181	1.146

Note: HP filtered data from the March CPS, 1968-2005. Log hours.

The fourth row indicates the average share of aggregate hours worked during the sample period by each age group. The fifth row indicates the share of "aggregate hours volatility" attributable to each age group. Here, aggregate hours volatility is represented by the share-of-hours-weighted average of age-specific cyclical volatilities. What is striking is the extent to which fluctuations in aggregate hours are disproportionately accounted for by young workers. Although those aged 15-29 make up only 26% of aggregate hours worked, they account for 43% of aggregate hours volatility. By contrast, prime-aged workers in their 40s and 50s account for 41% of hours but only 30% of hours volatility.

These age patterns remain when we undertake further demographic breakdowns. We summarize these results, found in Tables 2.2 and 2.3. Firstly, we disaggregate the U.S. workforce by age and gender. Again, the decreasing pattern exists for both men and women. Moreover, the magnitude of volatility differences by age is roughly similar. Importantly, the differences across age groups within gender are more pronounced than the differences across genders within age group. A simple

Table 2.3: Volatility of Hours Worked by Age and Education, U.S.

	15 - 19	20 - 24	25 - 29	30 - 39	40 - 49	50 - 59	60+
filtered volatility							
high school and less	5.950	2.984	2.669	2.403	2.006	2.286	3.120
more than high school	5.269	2.492	1.888	1.817	1.889	2.273	3.929
cyclical volatility							
high school and less	5.518	2.330	2.029	1.622	1.353	1.399	1.118
more than high school	3.587	1.476	0.626	1.071	0.739	0.831	1.378

Note: HP filtered data from the March CPS, 1968-2005. Log hours.

average across age groups indicates that males have about 1.3 times cyclical hours volatility of females. On the other hand, 15-24 year olds hours are roughly 3.1 and 3.5 times more cyclically volatile than 40-59 year olds hours, for females and males respectively. Gomme, Rogerson, Rupert, and Wright (2004) discuss age differences with further demographic breakdowns (e.g., marital status, industry of occupation) for the U.S. Their results corroborate those presented here, indicating large and important differences in the volatility of hours worked by age.

For disaggregation by age and educational attainment, the results remain. For brevity, we present results only for two education groups: those with high school diplomas or less (labeled high school and less), and those with at least some postsecondary education (more than high school). Again, the differences across age groups within education groups are more pronounced than the differences across education within age groups. The less educated have about 1.75 times the cyclical hours volatility of those with more than a high school education. On the other hand, 15-24 year olds hours are roughly 2.9 and 3.2 times more cyclically volatile than the 40-59 year olds hours, for less- and more-education respectively.

It is possible that we are mistaking labor market experience, as measured by age, for the last-in/first-out character of work situations that are negotiated by labor unions.[2.6] A greater degree of unionization in an industry, then, could associate with the age pattern we have found. Investigating this, we find that the age pattern remains the most robust feature of cyclical hours volatility, and that industry-types or industry unionization rates have much weaker relationships with cyclical hours volatility

We disaggregate hours by age and nine BLS-defined nonfarm supersectors, and use unionization rate data that is available from the BLS. For brevity, we focus on our preferred measure of cyclical volatility. In the the top panel of Table 2.4 the cyclical hours volatility for the 40-49 year olds for each supersector category is given the first column; columns 2-7 report the relative cyclical hours volatility of all the age groups in comparison to the 40-49 year olds. The eighth column reports the average cyclical hours of each supersector across age groups, and column 9 displays the supersectors' average unionization rates. The final two columns of the top panel show that there is a weak relationship between supersector unionization and supersector cyclical hours volatility; the first row of the bottom panel quantifies this correlation as 0.05.

The differences across age groups within supersectors or unionization groups are more pronounced than the differences across supersector/unionization groups within age groups. Rows 2 and 3 show that the standard deviation of volatilities

[2.6]We thank Valerie Ramey for bringing this idea to our attention.

Table 2.4: Volatility of Hours Worked by Age and Supersector, U.S.

	Cyclical volatility								
	Actual	*Relative*						*Actual*	Union
Supersector	$40-49$	$15-19$	$20-29$	$30-39$	$40-49$	$50-59$	60+	Avg	%
1	1.96	6.26	2.24	2.30	1	1.66	2.65	5.26	13.7
2	1.92	7.11	2.74	1.43	1	1.07	0.73	4.51	19.1
3	1.58	5.05	1.43	1.09	1	0.94	1.03	2.77	17.3
4	1.20	7.16	1.14	1.42	1	1.20	1.34	2.66	6.2
5	2.34	4.36	1.07	0.55	1	0.30	0.96	3.22	35.8
6	1.27	4.61	0.95	0.24	1	0.67	1.37	1.87	8.4
7	2.94	3.15	0.78	0.85	1	0.86	0.62	3.56	2.8
8	0.54	5.53	1.02	1.22	1	0.29	1.01	0.91	22.9
9	1.12	3.22	1.40	0.63	1	1.69	1.10	1.68	6.6

Summary statistics	
Corr between supersector unionization and supersector cyclical hours volatility	0.05
Std of cycl hours vol. within age group, across supersector	1.58
Std of cycl. hours vol. within supersector, across age	2.69
Avg ratio of cycl. hours vol. of three highest-unionized supersectors to three lowest	1.18
Avg ratio of cycl. hours vol. of 15-29 to 40-59	3.50

Note: HP filtered data from the March CPS, 1968-2005. Log hours. Unionization data from BLS, 1983-2006. Supersectors are the following: 1 – Forestry, Mining (SIC 0190-0280,0370 - 0490); 2 – Durable Manufacturing (2470 - 2590,2670 - 2990,3070 - 3290,3360 - 3390/3470, 3490,3570 - 3690,3770 - 3870,3890,3960 - 3990); 3 – Nondurable Manufacturing (1070 - 1290,1370,13901470 - 1790,1870 - 1990/6470 - 6490,2070, 2090,2170 - 2290,2370 - 2390); 4 – Trade (4070 - 4590,4670 - 5790); 5 – Transportation and Utilities (6070 - 6390,0570 - 0690/7790); 6 – Information Services (6675/6692, 6695,6680, 6690/6770, 67807270 - 7490/7570/7580 - 7780); 7 – FIRE (6870 - 6970/6990/7070/7080 - 7190); 8 – Health and Education Services (7860 - 7890, 8190,7970 - 8180, 8270, 8290,8370 - 8470); 9 – Leisure, Hospitality, and Other Services (8560 - 8590/6570, 6590/6670,8660, 8670,8680, 8690,8770 - 8890,8970 - 9090,9160 - 9190). The first column $40-49$ gives the cyclical hours volatility of the $40-49$ year olds; columns 2 through 7 give the relative cyclical hours volatility of all age groups; Avg gives the within-supersector average cyclical hours volatility; and, Union % give the supersector unionization rate.

across age groups is 1.7 times the standard deviation of volatilities across supersectors, showing greater variation by age in cyclical hours volatility than by supersector. Ordering supersectors by unionization, we find that the three most-unionized supersectors have, on average, 1.2 times the cyclical hours volatility of the three least-unionized supersectors. On the other hand, the 15-29 year olds have, on average, 3.5

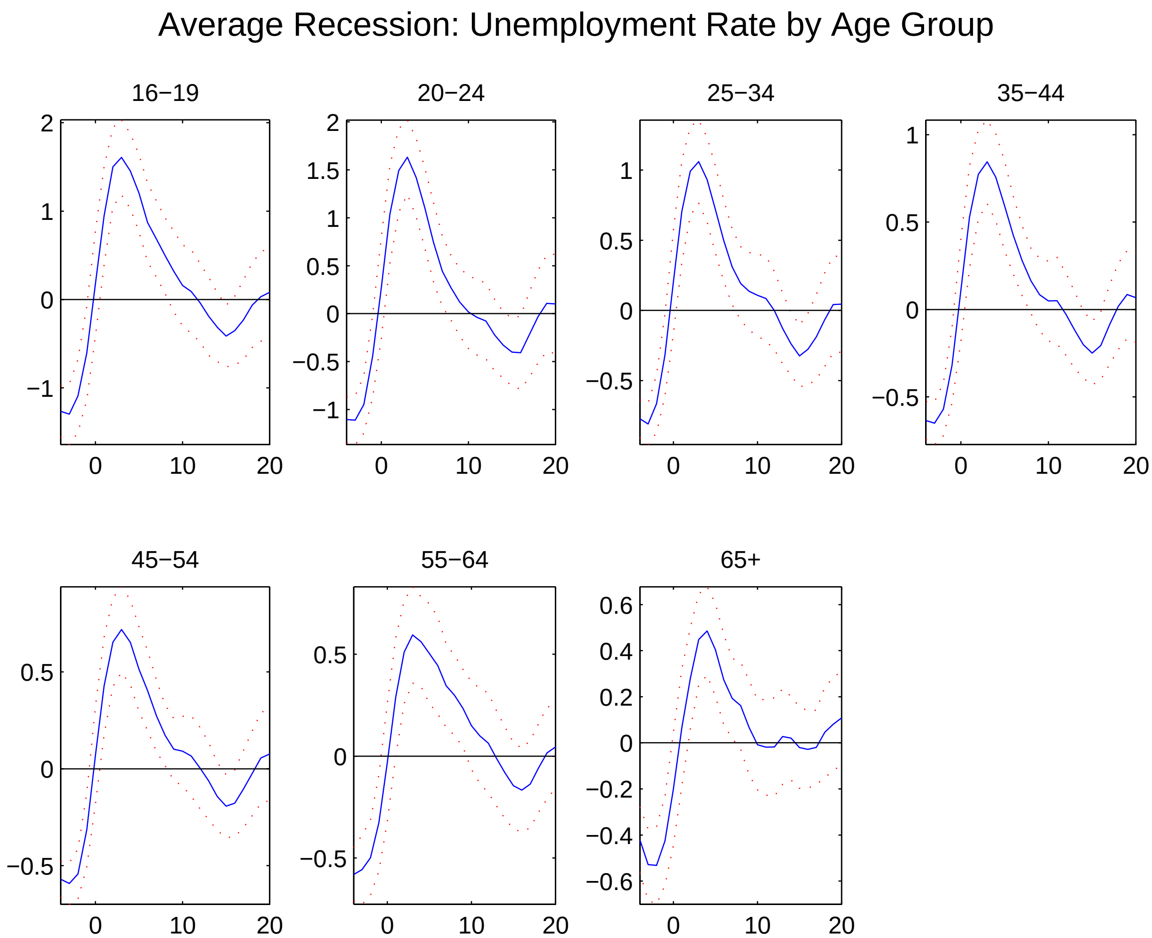

Figure 2.1: Unemployment Response to Recession
Note: Unemployment data from BLS, 1948:I-2004:II. HP filtered. Recession dates from NBER.

times the cyclical hours volatility of 40-59 year olds.[2.7]

Age-Specific Unemployment Additional evidence on the differences in business cycle sensitivity across age groups is presented in Figure 2.2.1 where we present the average response of unemployment to a postwar U.S. recession.[2.8] The

[2.7]In fact, the ratio of most-unionized to least-unionized cyclical hours volatilities is less than 1.8 for any age group; in contrast, the ratio of 15-29 to 40-59 cyclical hours volatilities is greater than 3 except in the case of Information Services (2.1) Leisure, Hospitality, and Other Services (1.7).

[2.8]See also Nagypál (2004) who provides an analysis of age group differences during recessionary episodes.

unemployment rate data come from the BLS, cover the period 1948:I-2004:II, and are available for the age groups presented. Along the horizontal axis, date 0 represents the first quarter of the recession as the NBER defines it. The figure tracks the filtered age-specific unemployment rates for 20 quarters beyond this date. The solid blue line represents the recessionary response averaged across episodes, while the dashed red lines represent 2 standard deviation bands. Unemployment rates for all age groups rise quickly in response to a recession, crossing above trend at date 0, then peaking at date 4 or 5 before slowly returning to trend.

Magnitudes of the recessionary response, however, differ across age groups. The peak response of unemployment is much stronger for younger aged individuals. While the unemployment rate of 16-19 and 20-24 year olds increases by 1.5% above trend, the increase is only about 0.6% for prime-aged workers. Moreover, the 16-19 and 20-24 age groups experience average trough-to-peak responses of approximately 2.4% about the trend. This compares with a trough-to-peak response of only 1.2% for prime-aged individuals. In summary, young workers' unemployment rate responds to recessions roughly twice as much as prime-aged workers' unemployment rate.

Age-Specific Labor Force Participation Labor fluctuations are due to changes in labor force participants' hours or changes in the number of the labor force participants. The latter we refer to as the "participation margin." The relative contribution of each of these two margins to the cyclical volatility of the labor input is important in guiding us to the relevant modeling approach. If the participation

Table 2.5: Hours Decomposition, Participation Margin

	15 - 64	15 - 19	20 - 24	25 - 29	30 - 39	40 - 49	50 - 59	60+
filtered volatility								
covariance not included (%)	9.05	28.82	8.73	10.49	10.48	16.24	25.38	57.23
covariance included (%)	19.78	47.01	20.46	11.80	17.19	16.23	29.15	60.41
cyclical volatility								
covariance not included (%)	5.35	23.64	4.06	2.01	1.46	2.78	4.26	30.08
covariance included (%)	16.80	45.68	17.17	6.81	7.81	2.66	10.09	35.94

Note: HP filtered data from the March CPS, 1968-2005. Log hours and log labor force shares. Shown are shares of total hours variation attributed to the participation margin. "Covariance not included" means no covariance terms are added to arrive at total variation. "Covariance included" means all covariance terms are added *and* all covariance terms including labor force are attributed to the participation margin.

margin is the main driver of hours variation, then it is reasonable to argue the practical necessity of explicitly modeling a labor force participation decision. However, if the main driver of hours variation is fluctuation in labor force participants' hours, then the decision of hours worked is the most important channel to be modeled. Table 2.5 shows the portion of hours variation by age group that can be attributed to the participation margin.

The first row, using the filtered volatility measure, tells us that the participation margin accounts for less than 30% of hours variation for all age groups except those over the age of 60. For those near retirement the labor force decision appears to be an important source of variance in the hours these 60+ year olds contribute; however, these workers provide less than 7% of the total hours worked. The second row makes an optimistic assumption on the explanatory power of labor force participation, adding all possible terms of the hours variance decomposition to this margin (see Appendix 2.A.1 for more details). In this case, still only the 60+ have more than 50% of their hours variation explained by participation choices. For all but the 15-19 year olds this margin still explains less than 30% of the variation.

The third and fourth rows give us the same information as rows one and two, respectively, now using the business cycle component of the age-specific hours and labor-force participation series.[2.9] Cyclical volatility tells an even less favorable story for the participation margin: with covariance terms not included, it explains less than 31% of the variation of any age group, most notably less than 5% of 20-59 year olds variation. Even with the inclusion of covariance terms, participation explains less than 50% of the variation of the 15-19 and 60+ year olds, and less than 20% of the 20-59 year olds.[2.10] We argue this evidence suggests that the participation margin is of secondary importance to the matter at hand. Instead, the behavior of hours-per-worker constitutes the major driver of hours volatility.

Summary The data presented in this Section suggests that the business cycle volatility of hours worked is decreasing with age. This finding is robust to further disaggregation to gender, educational attainment, major industry type, or unionization rate. The decrease of hours volatility with age is much greater than decreases along the gender, education, or unionization dimension. Business-cycle-

[2.9] Again, calculated as a projection on current detrended output, and current and lagged detrended hours.

[2.10] The portion of cyclical volatility for the 60+ year olds is noticeably lower than for filtered volatility; this comes from the low R^2 of these business cycle projections for these age groups' labor force participation decisions.

conditional volatility of the unemployment rate amongst the young is about twice that of the prime-aged, concurrent with a recessionary response that is, again, twice as large. The participation margin (movements into and out of the labor force) accounts for less than 31%, 25%, and 5% of the cyclical variation in hours among 60+ year olds, 15-19 year olds, and 20-59 year olds, respectively. Altogether, these data suggest that modeling an age-sensitive hours-per-worker mechanism is of first importance.

2.2.2 Age-Specific Wages

The previous section highlighted the fact that the hours worked by younger workers are more volatile than the hours worked by older workers. Within the RBC framework, this could potentially be driven by differences in labor supply or differences in factors related to technology that operate through the labor demand channel. The premise of this section is that an analysis of the cyclicality of real wages can be used to differentiate between these two mechanisms. If differences in labor supply are responsible for the higher volatility of the young workers' hours, then their wages should be less volatile over the business cycle than the wages of prime aged workers.[2.11] By contrast, if the higher volatility of the hours of young workers is due to differences in labor demand, then their wages should be more volatile than the wages of prime aged workers.

We thus look at the behavior of age-specific real wages. From the March CPS, we use information on labor income and hours worked to construct annual time series

[2.11] We make this argument in more detail in Appendix 2.A.2.

Table 2.6: Volatility of Real Hourly Wages by Age Group, U.S.

	15 - 19	20 - 24	25 - 29	30 - 39	40 - 49	50 - 59	60+
filtered volatility	2.865	1.588	1.268	1.126	1.355	1.338	2.215
R^2	0.26	0.25	0.22	0.29	0.18	0.24	0.09
cyclical volatility	1.363	0.816	0.663	0.580	0.656	0.580	0.659

Note: HP filtered data from the March CPS, 1963-2005. Log wages.

for wage rates for the period 1963 through 2005. We are careful to estimate wage rates per year as an average for each of 220 highly disaggregated demographic groups, properly weighted into a measure per age group, following Katz and Murphy (1992) and Krusell, Ohanian, Rios-Rull, and Violante (2000).[2.12] We then HP-filter these series to isolate fluctuations with frequency greater than 8 years.

The first row in Table 2.6 reports the percent standard deviation of the HP-filtered hourly real wage rates. We see a moderately decreasing pattern with an upturn for 60+ year olds. The second row reports the R^2 from our business cycle projection, showing a moderate cyclical association for age-specific wages in all but the 65+ group. Row three shows the now familiar decreasing pattern of cyclical volatility. The standard deviation of cyclical volatility for 15-19 and 20-24 year olds is 130% and 40% more than that of 50-59 year olds, respectively.

This pattern of wage volatility decreasing with age is robust to further disaggregation. Turning to Table 2.7 find that male age-specific wages follow the decreasing pattern until an upturn near retirement. For females, this same pattern is present in

[2.12] See Appendix 2.A.4 for more details.

Table 2.7: Volatility of Hourly Real Wages Worked by Age and Gender, U.S.

	15 - 19	20 - 24	25 - 29	30 - 39	40 - 49	50 - 59	60+
filtered volatility							
female	3.552	1.600	1.597	1.539	1.526	1.723	2.945
male	3.384	1.858	1.353	1.383	1.653	1.534	2.971
cyclical volatility							
female	1.334	0.381	0.338	0.384	0.409	0.407	0.753
male	1.909	1.131	0.737	0.713	0.579	0.752	1.089

Note: HP filtered data from the March CPS, 1963-2005. Log real wage.

Table 2.8: Volatility of Hourly Real Wages Worked by Age and Education, U.S.

	15 - 19	20 - 24	25 - 29	30 - 39	40 - 49	50 - 59	60+
filtered volatility							
high school and less	3.046	1.872	1.453	1.253	1.502	1.465	2.573
more than high school	3.727	1.923	1.307	1.369	1.533	1.751	2.875
cyclical volatility							
high school and less	1.485	0.684	0.366	0.268	0.489	0.484	0.499
more than high school	1.897	0.892	0.589	0.569	0.384	0.313	0.735

Note: HP filtered data from the March CPS, 1963-2005. Log real wage.

the filtered measure, but less so in our preferred cyclical measure due to a drop in the 25-29 year olds. On average across age groups, male wages have 1.75 times the cyclical volatility of female wages. In contrast, the cyclical wage volatility of 15-24 year olds is 2.1 and 2.3 times the cyclical wage volatility of 40-59 year olds, for females and males respectively.

When disaggregated by education in Table 2.8 the decreasing pattern of age-specific wages is still visible for both education groups, particular so for those with more than high school, with an upturn near retirement. On average across age groups, workers with more than high school education experience cyclical wage volatility that is 1.1 times the cyclical wage volatility of workers with only high school or less. Notably, we actually see that the cyclical volatility of the more-educated is greater than the less-educated *except* for the prime-aged 40-59 year olds where that pattern reverses. The cyclical wage volatility of 15-24 year olds is 4 and 2.2 times the cyclical wage volatility of 40-59 year olds, for less- and more-educated respectively.

The analysis of the cyclicality of age-specific wages in the U.S. data reveals that while all age-specific wages are procyclical, wages of young workers are more volatile over the cycle than those of others.

2.2.3 Summary

The empirical evidence that has been presented describes the cyclical behavior of both age-specific hours and age-specific wages.[2.13] The cyclical volatility of age-specific hours and wages decrease with age. The fact that the volatilities of young workers' hours *and* young workers' wages are higher than those of prime age workers suggests that differences in the the labor demand for each age groups has the potential of being an important mechanism in explaining the data.

[2.13] We have focused on hourly wages since the concept is both common and the price unit we explicitly model. Katz and Murphy (1992) and following labor literature use the March CPS for data on weekly wages. It is worth noting that the patterns we describe for hourly wages are intact when talking about weekly wages instead; the relevant table is available from the authors upon request.

2.3 Benchmark Model

2.3.1 Households

The economy is populated by a large number of identical, infinitely-lived households. Each household is composed of a unit mass of family members. For simplicity, we assume there are only two types of family members, Young and Old. Let s_Y denote the share of family members that are Young. Young family members derive utility from consumption C_Y and disutility from hours spent working N_Y according to U_Y. Old family members have similar preferences U_O defined over consumption C_O and working hours N_O.

The representative household's date t problem is to maximize

$$E_t \sum_{j=t}^{\infty} \beta^{t-j} (s_Y U_{Yt} + (1 - s_Y) U_{Ot}) \tag{2.1}$$

In our benchmark case, we specify U_Y and U_O such that this problem is

$$E_t \sum_{j=t}^{\infty} \beta^{j-t} \quad \left(s_Y \left[\log C_{Yj} - \psi_Y N_{Yj}^{1+\theta_Y} / \left(1 + \theta_Y \right) \right] \right.$$
$$\left. +(1 - s_Y) \left[\log C_{Oj} - \psi_O N_{Oj}^{1+\theta_O} / (1 + \theta_O) \right] \right)$$

subject to

$$s_Y C_{Yj} + (1 - s_Y) C_{Oj} + \tilde{K}_{j+1} = (1 - \delta)\tilde{K}_j + r_j \tilde{K}_j + s_Y W_{Yj} N_{Yj} + (1 - s_Y) W_{Oj} N_{Oj}, \ \forall j \geq t.$$

Here $\tilde{K}_t$ denotes capital holdings at date t, $0 \leq \delta \leq 1$ is the depreciation rate on capital, r_t is the rental rate, W_{Yt} is the wage rate of Young workers, and W_{Ot} is the wage rate of Old workers. $0 < \beta < 1$ is the time discount rate and $\theta_Y, \theta_O \geq 0$ govern the Frisch labor supply elasticity. $\psi_Y, \psi_O > 0$ are parameters used to calibrate the steady state of N_Y and N_O. The household takes all prices as given. We normalize the time endowment of all family members to unity, so that $0 \leq N_{Yt} \leq 1$ and $0 \leq N_{Ot} \leq 1$.

Because of additive separability in preferences, optimality entails equating consumption across all family members:

$$C_{Yt} = C_{Ot} = C_t. \tag{2.2}$$

The first-order condition for capital holdings is given by:

$$C_t^{-1} = \beta E_t \left[C_{t+1}^{-1}(r_{t+1} + 1 - \delta) \right].$$

The first-order conditions for hours worked are given by:

$$W_{Yt} = \psi_Y C_t N_{Yt}^{\theta_Y},$$

$$W_{Ot} = \psi_O C_t N_{Ot}^{\theta_Y}.$$

Condition (2.2) implies that the income effect of a consumption change on labor supply is equal across Young and Old workers. In our benchmark calibration, we set $\theta_Y = \theta_O$ so that the substitution effect of wage changes on labor supply is equated across workers. Adopting identical income and substitution effects allows us to isolate the role of capital-experience complementarity in generating volatility differences across Young and Old workers.

Finally, we assume that an individual's age directly determines his or her labor market experience, so that all Young workers are "inexperienced" while all Old workers are "experienced."

2.3.2 Firms

Production exhibits capital-experience complementarity. Formally, we assume that final goods are produced by perfectly competitive firms according to the CES production function:

$$Y_t = \left[\mu \left(A_t H_{Yt}\right)^{\sigma} + (1-\mu) \left[\lambda K_t^{\rho} + (1-\lambda) \left(A_t H_{Ot}\right)^{\rho}\right]^{\frac{\sigma}{\rho}}\right]^{\frac{1}{\sigma}}, \quad \sigma, \rho < 1.$$

Here H_{Yt} is labor input of Young or inexperienced workers, H_{Ot} is labor input of Old or experienced workers, and K_t is capital services hired at date t. Technology follows a deterministic growth path with persistent transitory shocks:

$$A_t = \exp\left(gt + z_t\right),$$

$$z_t = \phi z_{t-1} + \varepsilon_t, \quad 0 < \phi < 1,$$

where $E\left(\varepsilon\right) = 0$, $0 \leq var\left(\varepsilon\right) = \sigma_{\varepsilon}^2 < \infty$, and $g > 0$ is the trend growth rate of technology.

The elasticity of substitution between experienced workers and capital is given by $(1-\rho)^{-1}$, while the elasticity of substitution between inexperienced workers and the H_O-K composite is $(1-\sigma)^{-1}$. Following Krusell, Ohanian, Rios-Rull, and Violante (2000), we define production as exhibiting capital-experience complementarity when $\sigma > \rho$.

Firms rent capital, and Young and Old workers' time from perfectly competitive factor markets to maximize profits:

$$\Pi_t \equiv Y_t - r_t K_t - W_{Yt} H_{Yt} - W_{Ot} H_{Ot}.$$

Optimality entails equating factor prices with marginal revenue products:

$$r_t = Y_t^{1-\sigma}(1-\mu)\left[\lambda K_t^{\rho} + (1-\lambda)\left(A_t H_{Ot}\right)^{\rho}\right]^{\frac{\sigma-\rho}{\rho}} \lambda K_t^{\rho-1},$$

$$W_{Ot} = Y_t^{1-\sigma}(1-\mu)\left[\lambda K_t^{\rho} + (1-\lambda)\left(A_t H_{Ot}\right)^{\rho}\right]^{\frac{\sigma-\rho}{\rho}} (1-\lambda) A_t^{\rho} H_{Ot}^{\rho-1},$$

$$W_{Yt} = Y_t^{1-\sigma}\mu A_t^{\sigma} H_{Yt}^{\sigma-1}.$$

2.3.3 Equilibrium

Equilibrium is defined as follows. Given $\tilde{K}_0 > 0$ and the stochastic process, $\{z_t\}$, a *competitive equilibrium* is an allocation, $\{C_t, N_{Yt}, N_{Ot}, \tilde{K}_{t+1}, Y_t, H_{Yt}, H_{Ot}, K_t\}$, and a price system, $\{W_{Yt}, W_{Ot}, r_t\}$, such that: given prices, the allocation solves both the representative household's problem and the representative firm's problem; and factor markets clear for all t:

$$K_t = \tilde{K}_t; \quad H_{Yt} = s_Y N_{Yt}; \quad H_{Ot} = (1-s_Y)\, N_{Ot}.$$

Walras' law ensures clearing in the final goods market:

$$C_t + K_{t+1} = Y_t + (1-\delta)\, K_t, \quad \forall t.$$

Finally, for the purposes of model evaluation, we define aggregate hours worked as $H_t = s_Y H_{Yt} + (1-s_Y)\, H_{Ot}$.

2.3.4 The Effects of Capital-Experience Complementarity

In this subsection, we provide analytical results regarding the relative cyclicality of hours worked and real wages for Young and Old agents. To begin, we show that when production displays capital-experience complementarity, the response of

Young hours to a technology shock is greater than that of Old hours; this result holds even when there are no differences in labor supply characteristics.

Proposition 1. *Let $\theta_Y = \theta_O \geq 0$ and $\sigma > \rho$. The response of labor input of Young workers to a business cycle shock is greater than the response of labor input of Old workers.*

The proof is contained in Appendix 2.A.3. Here, we demonstrate this result for the special case in which $\rho = 0$.

When $\rho = 0$, the H_O-K composite becomes Cobb-Douglas, and the firm's FONCs become:

$$W_{Yt} = \mu Y_t^{1-\sigma} A_t^{\sigma} H_{Yt}^{\sigma-1},$$

$$W_{Ot} = (1-\mu)(1-\lambda) K_t^{\lambda\sigma} Y_t^{1-\sigma} A_t^{(1-\lambda)\sigma} H_{Ot}^{(1-\lambda)\sigma-1}.$$

In $\log W$ - $\log H$ space, these define linear labor demand curves, with slope $(\sigma - 1)$ for Young labor and slope $[(1-\lambda)\sigma - 1]$ for Old labor. Since $0 < \lambda < 1$, and $0 < \sigma < 1$ (recall that capital-experience complementarity is defined as $\sigma > \rho$, and we have assumed $\rho = 0$), the demand curve for Young labor is flatter than that of Old labor. Moreover, a shock to technology (a change in $\log A$) generates a vertical shift in the Young labor demand curve of σ, which is larger than the shift in the Old labor demand curve of $(1-\lambda)\sigma$. These two factors combine to generate the result of Proposition 1.

This can be seen diagrammatically in Figure 2.3.4. The left panel depicts the demand curve for Young labor, the right panel for Old labor. In each panel, the

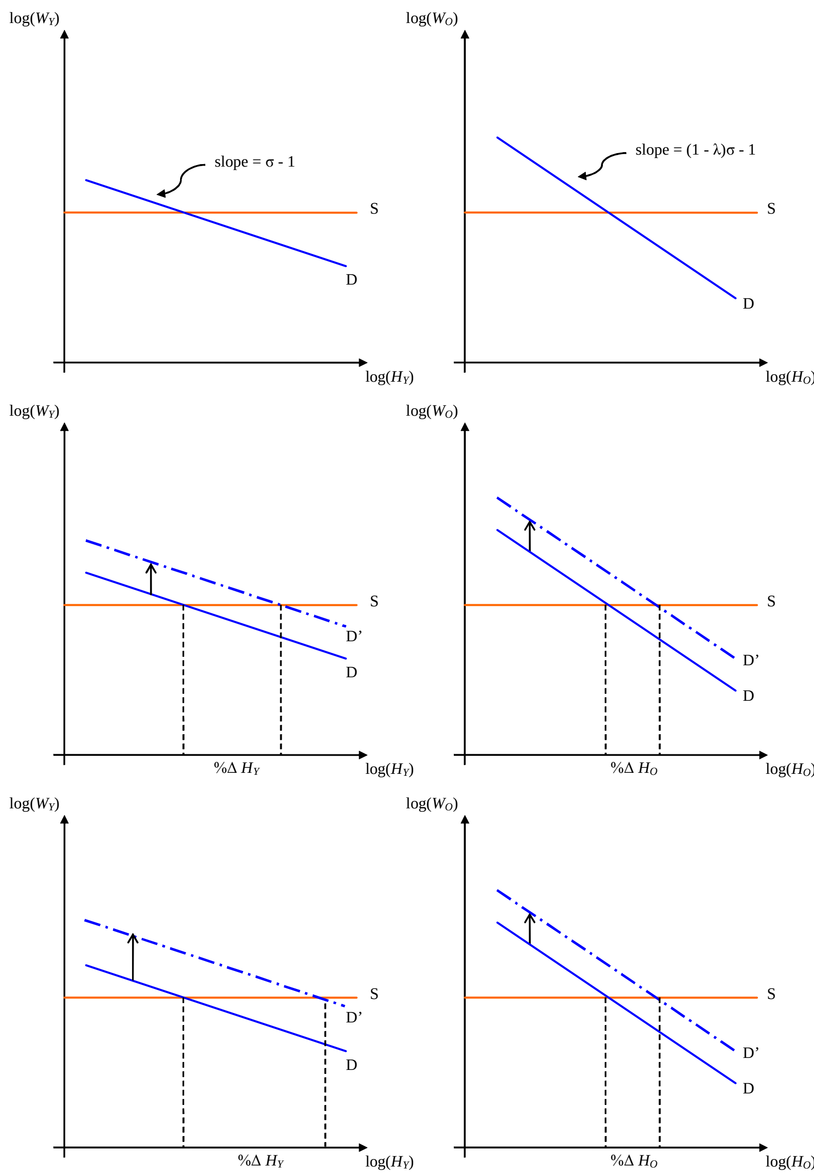

Figure 2.2: Diagrams

All panels: Horizontal lines depict the labor supply curves derived from the household's FONCs with Rogerson-Hansen preferences in log-log space with common slope $\theta_Y = \theta_O = 0$. **Top panel:** slope of demand curve for H_Y is flatter than the demand curve for H_O. **Middle panel:** we abstract from the wealth effects of a productivity shock since they are identical across Young and Old; the shock causes both demand curves to shift up; the "relative slope" effect is evident in $\%\Delta H_Y > \%\Delta H_O$. **Bottom panel:** the "demand response" effect is evident from the labor demand for H_Y shifting up by more than for H_O, increasing $\%\Delta H_Y$ even more.

horizontal line depicts the labor supply curves derived from the household's FONCs with Rogerson-Hansen lotteries; in log-log space, both are linear with common slope, $\theta_Y = \theta_O = 0$. Using a circumflex to denote log deviations, consider a positive shock to technology, $\hat{A}$. In equilibrium, the positive wealth effect of this shock generates an upward shift in the labor supply curves; since our model assumes identical wealth effects across agents, we abstract from these shifts in the diagram for the sake of clarity. The technology shock also results in an equilibrium output response, $\hat{Y}$; since the effect of this response is identical across labor demand curves, we abstract from these in the diagram as well. Finally, note that capital is a state variable, so that the response of capital to the shock is $\hat{K} = 0$.

Hence, the only effect that requires diagrammatic consideration is the direct effects of the shock on the labor demand curves, and we plot these in the middle and bottom rows of Figure 2.3.4. Suppose, momentarily, that the technology shock results in identical vertical shifts in the demand curves of Young and Old agents: this is illustrated as the dotted lines in the middle row. As is geometrically obvious, this results in a larger equilibrium response of Young labor input relative to Old labor input, i.e. $\hat{H}_Y > \hat{H}_O$. This is due to the relative complementarity of Old labor to capital, implying that the marginal revenue product of labor is more sensitive to changes in labor for H_O relative to H_Y. After a positive shift in labor demand, a smaller change in Old labor is required to achieve the same change in its marginal revenue product, and we call this the "relative slope" effect.

But note that the positive technology shock actually generates a larger shift in the demand for H_Y than for H_O: $\sigma\hat{A} > (1-\lambda)\,\sigma\hat{A}$. That is, the shock has a larger direct effect on the marginal revenue product of Young labor relative to Old labor. This is depicted by the dash-dot line in the left panel. This additional "demand response" effect reinforces "relative slope" effect due to capital-experience complementarity. Hence, in equilibrium, $\hat{H}_Y > \hat{H}_O$. The intuition for this is the following: when agents have identical labor supply curves, the only way to induce a greater hours response for Young workers is through a larger wage response.

Analytically, this can be derived from the household's FONCs with respect to labor supply. Using the fact that consumption is equated across agents:

$$W_{Yt}/\psi_Y N_{Yt}^{\theta_Y} = W_{Ot}/\psi_O N_{Ot}^{\theta_O}.$$

Substituting in the labor market clearing conditions, this can be rewritten in terms of log deviations as:

$$\hat{W}_Y - \hat{W}_O = \theta_Y \hat{H}_Y - \theta_O \hat{H}_O.$$

When $\theta_Y = \theta_O > 0$, $\hat{W}_Y > \hat{W}_O$ follows directly from the fact that $\hat{H}_Y > \hat{H}_O$. Note, however, that this condition implies a stronger result. With capital-experience complementarity, the wage response of Young workers is greater than that of Old workers even when Young labor supply is more elastic (i.e. when $\theta_Y < \theta_O$).

On the other hand, assume there is no capital-experience complementarity. As we discuss formally in Appendix 2.A.2, then one must assume that the Frisch labor supply elasticity of the Young is higher than that of the Old in order to match the fact that $\hat{H}_Y$ is more volatile than $\hat{H}_O$. However, in this case for identical labor demand shifts such a model cannot match the fact that $\hat{W}_Y$ is more volatile than $\hat{W}_O$. Hence, matching *both* the higher relative volatility of Young hours *and* Young wages requires a model where labor demand shocks are not age neutral.

2.4 Structural Estimation

In this section, we describe the quantitative specification used for evaluation of the model. To maintain comparability with the RBC literature, we adopt a standard calibration when possible. However, the model's parameters governing elasticities of substitution in production, σ and ρ, cannot be calibrated to match standard first moments in the U.S. data. Instead, we adopt a structural estimation procedure to identify these values using data from the NIPA and CPS. After describing the procedure, we discuss calibration of the remaining parameter values. Given the results in Section 2.2, we classify 15-29 year olds as Young and 30-64 year olds as Old.

Our strategy entails estimating σ and ρ from the model's labor demand equation.[2.14] Consider the firm's FONC with respect to the demand for H_{Yt} rewritten in logged, first-differenced form

$$\Delta \log W_{Yt} = a_0 + (\sigma - 1)\Delta \log \left(H_{Yt}/Y_t\right) + \sigma u_t, \tag{2.3}$$

where a_0 is a constant, and u_t is a function of current and lagged shocks

$$u_t = \varepsilon_t - (1-\phi)\left(\varepsilon_{t-1} + \phi\varepsilon_{t-2} + \phi^2\varepsilon_{t-3} + ...\right).$$

Hence, σ is determined from the response of W_Y to exogenous changes in H_Y and Y.

[2.14] A similar approach is used in Burnside, Eichenbaum, and Rebelo (1995) and the references therein.

The age-specific wage measures analyzed in Section 2.2.2 are constructed using hours data in order to translate direct information on labor income into information on hourly wages, similar to Krusell, Ohanian, Rios-Rull, and Violante (2000)'s hourly wage and Katz and Murphy (1992)'s weekly wage constructions. To ensure that we sidestep any potential problem stemming from this fact, we estimate a variant of (2.3) for which direct data on the left-hand side variable is available without having to be constructed using a right-hand side variable.[2.15] This is obtained by multiplying both sides of the first-order condition by H_{Yt}

$$\Delta \log LI_{Yt} = a_1 + \sigma \Delta \log H_{Yt} + (1-\sigma)\Delta \log Y_t + \sigma u_t, \tag{2.4}$$

where $LI_{Yt} \equiv W_{Yt} H_{Yt}$ denotes labor income earned by Young workers. If there were no endogeneity issues (see below), σ could be estimated from a simple restricted least-squares regression.

To estimate ρ, we proceed in a similar manner. Combining the firm's first-order conditions with respect to H_{Ot} and K_t and performing similar manipulations obtains

$$\Delta \log (Q_{Ot}/Q_{Kt}) = a_2 + \rho \Delta \log (H_{Ot}/K_t) + \rho u_t, \tag{2.5}$$

where Q_{Ot} denotes the share of national income earned by Old labor, and Q_{Kt} the share of national income earned by capital.

Importantly, this procedure does not require imposing any restrictions from

[2.15]The data used in estimation come from standard sources. Briefly, Y_t, K_t, and Q_{Kt} come from the BEA's NIPA and Fixed Asset Tables. H_{Yt}, H_{Ot}, LI_{Yt}, and Q_{Ot} are constructed using March CPS data. Because of this, our data comprise annual observations for the period 1968 - 2005. See Appendix 2.A.4 for a detailed discussion of the data construction.

the model's specification of household behavior.[2.16] The only assumptions required to pin down σ and ρ are: (*i*) profit maximization on the part of firms, and (*ii*) that factor prices reflect marginal revenue products. No aspect of our approach imposes $\sigma > \rho$. Whether this condition is satisfied depends on the relationship between aggregate prices and quantities observed in the data.

Instruments Since both of our estimating equations (2.4) and (2.5) are based on the estimation of factor demand equations, we need to address the endogeneity of the regressors to the error term. The structural equations identify the error term as due to shocks to productivity. It follows that in order to obtain unbiased estimates, we must identify variation in our regressors that is unrelated to shocks shifting firms' input demand, be they technology shocks or other omitted factors.

Specifically, we use two instruments: the Ramey-Shapiro dates (Ramey and Shapiro (1998); Ramey (2006)) and lagged birth rates. The Ramey-Shapiro dates correspond to dummy variables indicating the onset of government spending increases due to war and military build-ups. The validity of this instrument is easy to justify. In a standard RBC model like the one we consider, the introduction of exogenous government spending shocks introduces exogenous shifts in labor supply due to the wealth effect of such shocks (see Christiano and Eichenbaum (1992)). This results in changes in H_Y, H_O, and Y that are unrelated to shifts in factor demand. Given Ramey and Shapiro's narrative approach in identification, these government spending

[2.16]We see this as a virtue since our goal is to abstract from labor supply differences, and isolate the quantitative role of differences in the cyclical demand for Young and Old labor.

dates are exogenous to shocks to technology.

Our second instrument is lagged birth rates. This instrument allows us again to identify changes in current labor supply, this time due to changes in past fertility that are uncorrelated to shifts in factor demand. Recall that

$$u_t = \varepsilon_t - (1-\phi)\left(\varepsilon_{t-1} + \phi\varepsilon_{t-2} + \phi^2\varepsilon_{t-3} + ...\right).$$

In the case $\phi < 1$ lagged birth rates remain valid if we assume that birth rates are exogenous to all past technology shocks, $\{\varepsilon_{t-j}\}_{\forall j>0}$. If one believes that fertility decisions, say, 15 years ago might be endogenous to shocks at least 15 years ago, then some bias might be induced with these instruments. However, note that in the case of the 15-year lagged birth rate, the concern is its correlation with the sum $(1-\phi)\sum_{j=14}^{\infty}\phi^j\varepsilon_{t-j-1}$ in u_t. For standard values of shock persistence, ϕ, relevant for business cycle analysis, the impact of this is almost negligible. Obviously, for birthrates of larger lag, this is even smaller. We thus conclude that, from an empirical standpoint, lagged birth rates are valid instruments. We obtain IV estimates of $\hat{\sigma} = 0.591$ and $\hat{\rho} = 0.003$ with standard errors of 0.263 and 0.219, respectively.

2.5 Quantitative Evaluation

Given the estimated values for σ and ρ, the remaining parameters are calibrated in the standard way. $\beta = 0.99$ and $\delta = 0.025$. We set $s_Y = 0.31$ to match the average population share of Young individuals in $1968-2005$. N_{Yss} and N_{Oss} are set to jointly match the observed ratio of Young-to-Old hours worked in $1968-2005$, and a steady-state value for aggregate hours of $H_{ss} = s_Y N_{Yss} + (1-s_Y)N_{Oss} = 0.3$.

Table 2.9: Data and Model Moments

		Relative Volatilities				
Num.	*Denom.*	U.S. Data	Model: Separable Preferences	Model: Separable Preferences	Model: GHH	Model: GHH
Aggregate Hours	*Output*	0.97	0.91	0.83	0.97	0.96
"Y" Hours	*Output*	1.61	1.50	1.16	1.39	1.41
"O" Hours	*Output*	0.87	0.68	0.70	0.77	0.75
"Y" Hours	*"O" Hours*	1.85	2.22	1.66	1.81	1.88
"Y" Wages	*Output*	0.57	0.26	0.41	0.20	0.21
"O" Wages	*Output*	0.37	0.26	0.27	0.11	0.13
"Y" Wages	*"O" Wages*	1.54	1.00	1.54	1.82	1.56

Column 1: Data HP filtered from March CPS, 1968-2005. **Column 2:** Rogerson-Hansen preferences, $\theta_Y = \theta_O = 0$. **Column 3:** $\theta_O = 0$, $\theta_Y = 0.15$ chosen to match data on relative age-specific wages (seventh row). **Column 4:** $\theta_Y = \theta_O = 0.15$ chosen to match data on relative aggregate hours (first row). **Column 5:** $\theta_Y = 0.15$ and $\theta_O = 0.18$ chosen to match data on relative aggregate hours (first row) and relative age-specific wages (seventh row).

The only two "new" parameters are μ and λ. Following Krusell, Ohanian, Rios-Rull, and Violante (2000) we calibrate these to match national income shares. Specifically, in our model we set μ and λ to match the $1968 - 2005$ national income shares of $Q_K = 0.373$ and $Q_O = 0.494$. With values for $\{\sigma, \rho, \mu, \lambda\}$ and data on output and factor inputs, we back out the implied technology series $\{A_t\}$.[2.17]

Column 1 in Table 2.9 presents several statistics for the U.S. economy. As it is well known, the volatility of hours is almost identical to the volatility of output

[2.17]We are calibrating a quarterly model, however up to now we have dealt with annual data measures. The reason for this is that quarterly data on age-specific hours do not begin until 1976. We do have *semiannual* data on age-specific hours from 1968-2005 (constructed by the authors from the March CPS and the October CPS surveys held by NBER). There we see that the important series (for our purposes) display the same volatilities relative to output: the relative volatilities of Young and Old hours are the same in both the annual and semiannual time series. We conclude that for these relationships the frequency of observation does not much matter.

(the ratio is 0.97 in the U.S. data.) The following rows in column 1 report that the volatility of hours and wages of the Young and the Old relative to output and each other. Note that the hours of the Young are about 60% more volatile than output, while the volatility of the hours worked by the Old is less than that of output. This implies that the volatility of hours of the Young to Old is almost double (1.85). Similarly, the volatility of the Young wages is about 50% higher than that of the Old.

We begin the quantitative evaluation of the model in column 2 where we set $\theta_Y = \theta_O = 0$. I.e., we set the disutility from work to be linear as in the "indivisible labor" literature.[2.18] This is a useful benchmark since, as it is well known, the benchmark RBC with one type of labor and a Cobb Douglas production function requires this high labor supply elasticity in order to generate a significant volatility of hours. As column 2 reports, the capital-experience complementarity model generates a volatility of total hours that is almost identical to the one observed in the data. For comparison, a calibration of the benchmark RBC model with the same parameter values generates a ratio of the volatility of hours to the volatility of output of 0.79.

The next two rows show that the key to this success lays in the model's ability to generate a series of hours worked by the Young that fluctuates much more than output and Old hours over the business cycle. The model generates a 1.5 ratio of Young hours volatility to output volatility, which is slightly below the 1.6 observed in the data. Similarly to the data, the model generates a series of Old hours that is less volatile than output. Overall the model generates a 2.2 ratio of relative volatilities of

[2.18]See for example, Hansen (1985), Rogerson (1988), King and Rebelo (1999).

Young and Old hours, which is somewhat higher than the one observed in the data (1.85). However, while the benchmark model is successful with respect to the hours dimension, it cannot account for the behavior of relative wages, failing to replicate the relative volatility of Young to Old wages. This is not surprising since our the calibration in column 2 uses identical Frisch labor supply elasticities.

Hence, in column 3 we study the following modification; we change one of Frisch labor supply elasticities such that we match the volatility of wages of the Young relative to that of the Old. As column 3 reports, the cost of such a procedure is mainly along the hours behavior. Not surprisingly, a less elastic supply of hours induces a fall in the volatility of hours relative to that of output (0.83) and reduces the volatility of Young hours relative to output.[2.19]

To conclude, the benchmark model with separable preferences and capital experience complementarity can easily match the behavior of hours at the cost of not matching the relative volatility of wages (column 2). When targeted to match the relative volatility of wages, the model underperforms somewhat with respect to the volatility of age specific hours and overall aggregate hours (column 3).

Recall that neither the idea of capital-experience complementarity nor our structural estimation depend on the specification of the household side of the economy. Therefore, we investigate the model's performance when using preferences as suggested in Greenwood, Hercowitz, and Huffman (1988) (GHH) instead of the bench-

[2.19] Again, for comparison, the benchmark RBC model generates a ratio of 0.67 with the same $\theta_Y = \theta_O = \theta$ as used in this example. This shows again the importance of the capital-experience channel for the volatility of hours.

mark preferences. As shown below, this alternative utility function enables us to both match the relative behavior of Young and Old hours and the relative behavior of Young and Old wages.

The GHH momentary utility function is given by

$$U_j = \log\left(C_j - \psi_j N_j^{1+\theta j} / (1+\theta_j)\right) \quad \text{for } j \in \{Y, O\} \tag{2.6}$$

and we plug these U_Y and U_O into the representative household's problem (2.1). As is well known, this utility function exhibits no wealth effects on the supply of labor and it generates a higher volatility of hours. We begin the analysis in column 4 of Table 2.9 by considering the case of $\theta_Y = \theta_O$. Under this assumption and given our previous calibration we search for the θ that generates a volatility of aggregate hours relative to output that matches this ratio in the data (0.97).[2.20] As this column suggests, the model matches the relative volatility of hours worked by the Young and Old (1.81 in the model relative to 1.85 in the data). However, the model generates a wage series for the Young that is about 80% more volatile than that of the Old, while in the data this ratio is about 50%. Hence, in column 5 we repeat the same exercise we conducted in column 3 where we alter one of the labor supply elasticities to match the volatility of the Young wages relative to that of the Old. As column 5 reports, the model almost identically matches the behavior of aggregate hours relative to that of output. Similarly, the model continues to match the relative volatility of hours worked by the Young and the Old, while also being consistent with the behavior of the age specific

[2.20] We do not set $\theta_Y = \theta_O = 0$ as in column 2 since in this case the model generates an aggregate series of hours that is more volatile than output.

hours relative to output (in the model we get 1.41 for the Young and 0.75 for the Old, while in the data the ratios are 1.61 for the Young and 0.87 for the Old). Note that while the model in column 5 matches the volatility of the Young wages relative to that of the Old, it underperforms with respect to the ratio of the volatility of each of the wage groups hours relative to output. The lackluster performance of the model with respect to this last statistic characterizes the other versions of the model as well.

To summarize this section, the capital-experience complementarity model, when calibrated in the standard way and using independently-estimated[2.21] elasticity of substitution parameters, can easily match the high volatility of aggregate hours which has been a long standing puzzle in the business cycle literature. Moreover, not only does the model generate a series of aggregate hours that is almost as volatile as output, it also generates age-specific hours that are consistent with the pattern of age-specific hours observed in U.S. data. This points to the importance of explicitly accounting for the hours of Young inexperienced workers because the cyclical fluctuations of aggregate hours are disproportionately affected by this labor type. Assuming the utility function suggested in Greenwood, Hercowitz, and Huffman (1988) allows the model to match the behavior of relative wages moderately well.

2.6 The Great Moderation

Recently, a great deal of attention has been devoted to studying the moderation in business cycle volatility in the U.S since the mid-1980s. Since the mid-1980s

[2.21] We mean that the estimation is independent of, and so cannot target, (*i*) the volatility of aggregate hours and (*ii*) the relative volatility of Young and Old hours.

the U.S. has undergone a substantial decline in business cycle volatility. Indeed, determining the causes of "the Great Moderation" is the objective of a growing body of literature. Potential explanations include: a reduction in inflation volatility that is potentially related to improved monetary policy (see Clarida, Gali, and Gertler (2000), Blanchard and Simon (2001), Stock and Watson (2002)); regulatory changes and financial market innovation related to household borrowing (Campbell and Hercowitz (2005), Fisher and Gervais (2006), Justiniano and Primiceri (2006)); changes that have reduced the volatility of production relative to sales (McConnell and Perez-Quiros (2000), Ramey and Vine (2005)); and good luck, in the form of a reduction in the variance of business cycle shocks (Stock and Watson (2002) and Stock and Watson (2003), Justiniano and Primiceri (2006), Arias, Hansen, and Ohanian (2007)).

Jaimovich and Siu (2007) propose demographic change as a force that can rationalize the evolution of U.S. macroeconomic volatility over the last four decades. Moreover, demographic change is there shown to be relevant to understanding the evolution of cyclical volatility observed in industrialized economies other than the U.S. during the postwar period. Using panel data methods, Jaimovich and Siu (2007) exploit cross country differences in demographic change to show that the age composition of the workforce has a large and statistically significant effect on cyclical volatility. Using simple accounting exercises, demographic change is found to account for about a quarter of the moderation in the volatility of output and hours worked, respectively. Those results indicate that demographic composition plays an important role in the propagation of business cycle fluctuations and they serve as our motivation to articulate this notion within our quantitative macroeconomic framework.

Hence, having shown in Section 2.5 that the benchmark model is capable of accounting for differences in the volatility of hours worked across age groups, we study how changing age composition manifests itself in changing macroeconomic volatility. Our maintained hypothesis within the model is that the Great Moderation is due to two factors: a fall in the volatility of technology shocks, and a fall in the share of aggregate hours worked by young agents.[2.22] Therefore, we proceed as follows. We use the estimated $\hat{\sigma}$ and $\hat{\rho}$, as in the previous section, and we back out the implied technology series for the pre- and post-Moderation periods. Specifically we find a drop of 55% in the standard deviation of the innovation across the two periods. In the pre-Moderation period, $s_Y = 0.35$ is set to match the average population share of Young individuals in $1968 - 1984$. N_{Yss} and N_{Oss} are set to jointly match the observed ratio of Young-to-Old hours worked in $1968 - 1984$, and a steady-state value for aggregate hours of $H_{ss} = s_Y N_{Yss} + (1 - s_Y) N_{Oss} = 0.3$. To match the increasing share of Old hours, we drop $s_Y = 0.27$ and increase N_{Oss} by 12% in the post-Moderation period to match the average values observed in $1985 - 2005$.

2.6.1 Results

To evaluate the model's predictions, we separately simulate data for the pre- and post-Moderation periods according to the calibration just described. Aside from the changes to the shock process and demographics across periods, all other pa-

[2.22]See Arias, Hansen, and Ohanian (2007)'s investigation of the role of decreased Solow residual volatility in the standard RBC model.

Table 2.10: Great Moderation, Benchmark Case

	A. Data			*B. Model*			*C. Counterfactual*	
	PRE	POST	% CHG	PRE	POST	% CHG	POST	% CHG
std(Y)	2.11	0.94	−80.9	1.65	0.89	−61.7	0.95	10.2
std(H)	1.99	0.99	−69.8	1.60	0.82	−66.8	0.91	16.9
std(H)/std(Y)	0.94	1.05	11.0	0.97	0.92	− 5.1		
std(H_Y)/std(Y)	1.50	2.02	29.8	1.41	1.45	2.8		
std(H_O)/std(Y)	0.89	0.85	− 4.6	0.76	0.77	1.3		
std(H_Y)/std(H_O)	1.69	2.38	34.4	1.86	1.88	1.5		

Note: PRE period is 1968-1984, POST period is 1985-2005. **Panel A:** Data HP filtered from March CPS, 1968-2005. **Panel B and C** are described in the text.

rameters are held fixed – we call this our benchmark case.[2.23] Table 2.10 presents second-moment statistics for HP-filtered output and hours worked for the U.S.: the first column covers the $1968-1984$ period, the second column covers $1985-2004$, and the third column presents the log-percentage difference. The volatility of output and aggregate hours both exhibit drastic moderation, on the order of a more than 70 log-percentage fall across the two periods. Interestingly, the fall in the volatility of hours worked by Young individuals has been smaller (only 30 log-percentage points) so that, relative to output, the standard deviation of the Young hours has actually risen by 30 log-percentage points. Panel B of Table 2.10 presents the same statistics for model simulated data. For the benchmark calibration, the model generates volatility of Young and Old hours relative to output that matches the average values found in the U.S. for the $1968-2005$ period. Note that the model does a good job of replicating the Great Moderation driven solely by changes in the volatility of shocks

[2.23]The model uses GHH preferences, for which we found the best model performance in Section 2.5.

and the share of young and old; specifically, the model generates moderation in the volatility of aggregate output and hours that is about 80% as large as those found in the data.

To assess the role of demographic change in accounting for the model-generated moderation, we perform a counterfactual experiment where we re-simulate data for the post-Moderation period holding demographic factors fixed at their pre-Moderation values, allowing only the shock volatility to fall. The results are reported in panel C of Table 2.10. Had demographics stayed constant across periods but the variance of productivity shocks had fallen, aggregate volatility would have fallen by only 55 log-percentage points. Hence, demographic change accounts for about 10% of the moderation in output and 17% of the moderation in aggregate hours.[2.24]

Note that the benchmark case does not capture the increase in the relative volatility of Young workers' hours since 1984. As a result, the benchmark counterfactual likely understates the role of demographic change. In the post-Moderation period, we not only see a fall in the share of volatile Young workers, we also see those workers become more volatile. In this case, holding shares constant at pre-Moderation values would entail larger demographic effects.[2.25] Because the model cannot account for this, we propose simple reduced-form modifications to gauge its quantitative importance – we call this our alternative case. Specifically, we set the

[2.24]We also performed the counterfactual in which the post-1985 period is re-simulated with the shock process of the pre-1984 period, allowing only demographics to change. In this experiment, demographic change accounts for virtually the same fraction of the moderation in hours and output as discussed here.

[2.25]The same is true for Old workers, as their hours became slightly more stable: see panel A, row 5, Table 2.10 or 2.11. Quantitatively, this effect is likely much weaker.

Table 2.11: Great Moderation, Alternative Case

	A. Data			*B. Model*			*C. Counterfactual*	
	Pre	Post	% Chg	Pre	Post	% Chg	Post	% Chg
std(Y)	2.11	0.94	−80.9	1.94	0.99	−67.3	1.11	18.6
std(H)	1.99	0.99	−69.8	2.06	1.01	−71.3	1.18	24.6
std(H)/std(Y)	0.94	1.05	11.0	1.06	1.02	− 4.0		
std(H_Y)/std(Y)	1.50	2.02	29.8	1.51	1.85	20.3		
std(H_O)/std(Y)	0.89	0.85	− 4.6	0.85	0.77	− 9.9		
std(H_Y)/std(H_O)	1.69	2.38	34.4	1.78	2.40	30.2		

Note: Pre period is 1968-1984, Post period is 1985-2005. **Panel A:** Data HP filtered from March CPS, 1968-2005. **Panel B and C** are described in the text.

pre-Moderation values of θ_Y and θ_O to match the relative volatilities of Young and Old hours to output in $1968 - 1984$; we set the post-Moderation values of θ_Y and θ_O to match the relative volatilities in $1985 - 2005$.[2.26] We perform the same counterfactual as described above and the results are shown in panel C of Table 2.11. We find that demographic change now accounts for about 19% of the moderation in output volatility and 25% of the moderation in aggregate hours volatility.

In summary, this simple variant of the RBC model with capital-experience complementarity attributes a similar role to demographic change in the moderation of macroeconomic volatility to what is predicted in the reduced form results of Jaimovich and Siu (2007). We view these results as suggesting the importance of a structural model that is capable of replicating the observed changes in the relative volatility of

[2.26] As panel B in Table 2.11 reports, in the post-Moderation calibration the model does not match the relative volatilities of Young hours to output. We find that if we match this ratio, then the relative volatility of the volatility of Old hours to output is too low in light of the data. This is the reason we opt to generate "sufficient" volatility in Young hours while not underperforming with respect to the volatility of Old hours.

hours worked by young inexperienced and old experienced agents. Our results suggest that such a model would potentially attribute a significant role to demographics in the great moderation.[2.27]

2.7 Conclusion

We have constructed an RBC model where capital and experienced old labor are greater complements than capital and inexperienced young labor. We do this after investigating the behavior of both age-specific hours input and age-specific wages over the business cycle. We find that young individuals' hours and wages are more volatile than the hours and wages of old individuals. We argue that, within an RBC framework, differences in age groups' labor demand curves can explain this fact while differences in labor supply curves cannot. A straightforward and precedented manner of introducing this differential labor demand is assuming capital-experience complementarity in the production side of the economy. In a procedure that does not impose any ordering on the parameters, we estimate two production function elasticities of substitution that bear out the idea that old labor and capital are greater complements. In our subsequent simulation exercises, we maintain these estimated elasticities in order to discipline our analysis.

Our calibrated model is able to match well the relative volatility of young

[2.27] Note that the current model has only two groups of workers. Thus, our counterfactuals ignore important composition changes within the $15-29$ and $30-64$ year old age groups. Specifically, the counterfactuals understate the fall in the share of $15-19$ year olds and concurrent rise in the share of $40-49$ year olds observed in the post-moderation period. Because these are the most volatile and most stable age groups, respectively, a more disaggregated treatment of the age composition would suggest an even greater role.

and old hours with respect to output and with respect to output. Importantly, it is also able to replicate the relative volatility of aggregate hours with respect to output, which had previously been a lackluster aspect of RBC models. Given our model's success, we explore its usefulness in understanding structural forces behind the widespread diminishment of business cycle variation around 1984 referred to as the "Great Moderation." In particular, we ask if a large shift in the age-composition of the labor force might explain part of the the decline in output and aggregate hours volatility. Calibrating the model to match key business cycle facts for the pre-1984 and post-1984 periods, our counterfactual exercises suggest that demographic change accounts for between 10-18% of the moderation in output volatility and 19-25% of the moderation in aggregate hours volatility.

ACKNOWLEDGEMENT: *this chapter is joint work with Nir Jaimovich and Henry E. Siu. The dissertation author was a primary author of this paper.*

2.A Appendix

2.A.1 Hours Variance Decomposition

In this subsection, we provide a decomposition of the volatility of hours worked into its primary components. Our view is that such a decomposition provides preliminary evidence on the relative importance of labor supply and labor demand factors in understanding differences in volatility across age groups. This type of analysis is informative with respect to the structural model we pursue. For example, differences owing to life-cycle considerations would plausibly attribute a greater proportion of the volatility of hours to volatility in labor force participation among some age groups as compared to others. Schooling decisions might be a margin of adjustment which is more relevant for young agents than to those in their prime-age. Similarly, the decision to re-enter or drop out of the labor force over the business cycle might be more relevant for those above the retirement age than others.

Note that hours worked per member of an age group (H) can be written as the following product:

$$H = h \times lf,$$

where h is hours per labor force participant and lf is labor force participant per age-group member. Hence, changes in H can be due to changes in any of these two terms, which we refer to as the "hours margin," and the "participation margin," respectively. As such we decompose the variance of hours worked in the following way:

$$\text{var}\left(\hat{H}\right) = \text{var}\left(\hat{h}\right) + +\text{var}\left(l\hat{f}\right) + 2 \times \text{cov}\left(\hat{h}, l\hat{f}\right) \tag{2.7}$$

where $\hat{H}$ denotes deviations of the log of H from its HP-filtered trend, and similarly for the other variables.[2.28] Our point in decomposing var $\left(\hat{H}\right)$ is to see how influential is the participation margin, which we do not model, relative to the hours margin, which we do.

Again, our primary interest is in fluctuations of hours worked that are correlated with the business cycle. We thus regress each HP-filtered series on the same measures of the business cycle as in Sections 2.2-2.2.2, and perform the decomposition on the projected series.

The goal of simply accounting for the participation margin does not tell us how to attribute the covariance term in (2.7) and so we do the following. For one set of calculations, we ignore the covariance term in (2.7). Then the participation margin is simply

$$\frac{\text{var}\left(l\hat{f}\right)}{\text{var}\left(\hat{h}\right)+\text{var}\left(l\hat{f}\right)}$$

We denote these results "covariance not included."

For the other set of calculations, we include the covariance term. Moreover, we take a very favorable stance towards the participation margin and attribute to it the covariance. Here the participation margin is

$$\frac{\text{var}\left(l\hat{f}\right)+2\times\text{cov}\left(\hat{h},l\hat{f}\right)}{\text{var}\left(\hat{H}\right)}$$

using (2.7). We denote these results "covariance included."

As discussed in the text, the picture that emerges here is that the participation margin is less influential than the hours margin we model for all age groups.

[2.28] See Hansen (1985) for a similar decomposition.

2.A.2 Labor Supply Models

As discussed in the introduction, in order to maintain comparability with the literature, we are interested in a model that represents a minimal deviation from the standard RBC model. We begin by analyzing two simple models based on labor supply differences. As expected from the previous discussion, while these models can account for the differences in the cyclicality of age-specific hours, they have counter-factual implications regarding the cyclicality of age-specific wages. We then conclude that within the RBC framework, labor demand differences are crucial for matching differences in the cyclicality of age-specific wages and analyze such a model in the following Section.

Differences in Labor Supply: Model I In the first model we consider we assume that final goods are produced by perfectly competitive firms according to the Cobb-Douglas production function:

$$Y_t = A_t K_t^{\alpha} H_t^{1-\alpha}$$

Here H_t is the aggregate labor input in the economy and it satisfies

$$H_t = H_{Yt} + EH_{Ot}$$

That is, the aggregate labor input is the sum of the hours of the Young, H_{Yt}, and the hours of the Old, H_{Ot}. The parameter E allows for a difference in the efficiency of hours supplied between the Young and the Old. The rest of the model is identical to the one we consider in the paper.

In this enviorment we have

$$W_{Ot} = (1-\alpha)E\frac{Y_t}{H_t}, \tag{2.8}$$

$$W_{Yt} = (1-\alpha)\frac{Y_t}{H_t}. \tag{2.9}$$

From the FOC of the household we get

$$\hat{N}_{Yt} = \frac{\theta_O}{\theta_Y}\hat{N}_{Ot} \tag{2.10}$$

Thus, if $\theta_Y < \theta_O$, i.e., the Frisch labor supply elasticity of the Young is higher then that of the Old (recall that the Frisch labor supply elasticity equals $\frac{1}{\theta}$), then the model can match $Var(\hat{N}_{Yt}) > Var(\hat{N}_{Ot})$. However, from (2.8) and (2.9) it follows that by construction the relative volatility of wages equals one in this model. Thus, this model cannot account for the joint behavior of age specific hours and age specific relative wages that is observed in the U.S. data.

Differences in Labor Supply: Model II We maintain the same assumptions regarding the households as in the previous. We only vary the production function by postulating the following production function

$$Y_t = A_t K_t^{\alpha}\left(H_{Yt}^{\mu}\left(EH_{Ot}\right)^{1-\mu}\right)^{1-\alpha}$$

Note that in this version we allow for the labor input of the Young and the Old to differ. However, we assume that both of these inputs have the same elasticity of substitution with capital. Given this production function we get that

$$W_{Ot} = (1-\gamma)(1-\alpha)\frac{Y_t}{H_{Ot}}, \tag{2.11}$$

$$W_{Yt} = \gamma(1-\alpha)\frac{Y_t}{H_{Yt}}. \tag{2.12}$$

Using the prices determined in (2.11)-(2.12) and the fact that $H_{Yt} = s_Y N_{Yt}$ and $H_{Ot} = (1 - s_Y) N_{Ot}$ we get

$$\frac{(1-\gamma)(1-\alpha)}{\psi_O (1-s_Y)} = \frac{C_t N_{Ot}^{1+\theta_O}}{Y_t}, \tag{2.13}$$

$$\frac{\gamma(1-\alpha)}{\psi_Y} = \frac{C_t N_{Yt}^{1+\theta_Y}}{Y_t}. \tag{2.14}$$

Log-linearizing the ratio of these two equations it follows that

$$\hat{N}_{Yt} = \frac{(1+\theta_O)}{(1+\theta_Y)} \hat{N}_{Ot} \tag{2.15}$$

Similarly to the previous model we have considered, the case of $\theta_Y = \theta_O$ implies the volatility of hours worked is identical for the Young and the Old. The only case in which $Var(\hat{N}_{Yt}) > Var(\hat{N}_{Ot})$ is if $\frac{(1+\theta_O)}{(1+\theta_Y)} > 1$ – i.e. $\theta_Y < \theta_O$. However, from the labor supply equations it follows

$$\hat{W}_{Yt} - \hat{W}_{Ot} = \theta_Y \hat{N}_{Yt} - \theta_O \hat{N}_{Ot},$$

and thus

$$\hat{W}_{Yt} - \hat{W}_{Ot} = \left(\frac{\theta_Y - \theta_O}{1+\theta_Y} \right) \hat{N}_{Ot}.$$

Since we are interested in calibrations where $\theta_Y < \theta_O$, it follows that $\hat{W}_{Yt} < \hat{W}_{Ot}$. The immediate implication is that the fluctuations in the wage of the Young are smaller then the fluctuations in the wage of the Old.[2.29]

To conclude this section: the two models based on "labor supply" differences can easily match the relative volatility of age specific hours. The first model can also match the relative volatilities of Young and Old hours to output. However, it

[2.29] This is true as long as the two wages correlate positively.

is inherent to the models' mechanisms that they have counterfactual implications regarding the volatility of age specific wages. These results lead us to consider a "Labor-Demand" channel.

2.A.3 Proofs

The method of proof follows the arguments made in the text. Assume $\sigma > \rho$, so that production exhibits capital-experience complementarity. The firm's FONCs written in log deviation form are:

$$\hat{W}_Y = (1-\sigma)\hat{Y} + \sigma\hat{A} + (\sigma-1)\hat{H}_Y,$$

$$\hat{W}_O = (1-\sigma)\hat{Y} + \left(\frac{\sigma-\rho}{\rho}\right)\hat{X} + \rho\hat{A} + (\rho-1)\hat{H}_O.$$

Here, $X = \lambda K^\rho + (1-\lambda)(AH_O)^\rho$, so that:

$$\hat{X} = \frac{(1-\lambda)(AH_O)^\rho}{X}\rho(\hat{A}+\hat{H}_O) \equiv X_2\rho(\hat{A}+\hat{H}_O).$$

We have used the fact that $\hat{K} = 0$ in the impact period of a shock. Note that $0 < X_2 < 1$. Hence:

$$\hat{W}_O = (1-\sigma)\hat{Y} + [(\sigma-\rho)X_2 + \rho]\hat{A} + [(\sigma-\rho)X_2 + \rho - 1]\hat{H}_O.$$

Assuming $\theta_Y = \theta_O = \theta$, the household's FONCs in log deviation form are:

$$\begin{aligned}\theta\hat{H}_Y &= \hat{W}_Y - \hat{C},\\ \theta\hat{H}_O &= \hat{W}_O - \hat{C},\end{aligned}$$

so that:

$$\theta\hat{H}_Y - \hat{W}_Y = \theta\hat{H}_O - \hat{W}_O.$$

Substituting in the firm's FONCs and simplifying, we obtain:

$$\frac{\hat{H}_Y}{\hat{H}_O} = \frac{\theta + 1 - \rho - (\sigma - \rho) X_2}{\theta + 1 - \sigma} + \frac{(\sigma - \rho)(1 - X_2)}{\theta + 1 - \sigma} \frac{\hat{A}}{\hat{H}_O}.$$

The first term on the right-hand side of the equality is greater than one since $\sigma > \rho$. Moreover, since $0 < X_2 < 1$, the second term on the right-hand side is greater than zero. Hence, capital-experience complementarity implies that $\hat{H}_Y > \hat{H}_O$ in response to a positive technology shock, $\hat{A} > 0$.

2.A.4 Data

Data on hours, employment shares, and wages come from the Current Population Survey (CPS) conducted by the Census Bureau. To obtain wage data, we use questions in the March CPS about income obtained in the previous (last) year.[2.30] In order to turn this income data into wage data, we must know how many hours the individual worked last year. The hours for the previous year are constructed as the number of weeks worked last year multiplied by some measure of how many hours-per-week were worked by the individual last year. We follow Krusell, Ohanian, Rios-Rull, and Violante (2000) in imputing the hours-per-week from the data on how many hours the individual worked *in the previous (last) week*.

Our measure of hours-per-week is different than Krusell, Ohanian, Rios-Rull, and Violante (2000) in the following. We note whether the worker described her work last year as either full-time (FT) or part-time (PT). Her last week's hours are imputed as the hours-per-week only if the value falls within believable values, given

[2.30] As noted below, a specific question reporting wages only appears in the CPS survey starting in 1982.

that her work last year was either FT or PT. If her previous week's hours are not consistent with FT or PT work, we impute a "disaggregated" group average as the hours-per-week; by contrast, Krusell, Ohanian, Rios-Rull, and Violante (2000) impute a "disaggregated" group average only if the worker reported that she worked last year but worked zero hours last week.

Our "disaggregated" groups are formed by dividing respondents by age, education, gender, and last year's FT/PT status. Given that there are eleven 5-year age bins (15-19,20-24,...,60-64,65+), 5 eduction bins (below HS, HS, some college, college graduate, postgraduate work), 2 genders, and a FT or PT status, there are 220 possible groups. Our "disaggregated" groups combine education bins for some age-gender-FT/PT groups to ensure that for every year in 1964-2006 our 114 "disaggregated" groups each have at least fifty members.[2.31] This is done so that the "disaggregated" group average is not overly reliant on only a few observations.

Let g be a "disaggregated" group of workers: say g is a group of individuals of a certain age-education-gender category that claimed to have been working PT last year. If $i \in g$ says that they were working 24 hours last week, we impute that as person i's hours-per-week last year: this is because the CPS defines PT work as between 0-34 hours. On the other hand, suppose $j \in g$ says they were working 50 hours last week – this number is not *FT-status-fitting* as an imputed number for hours-per-week last year since j described their work last year as PT. Therefore, we take a CPS-weighted average over the people in g who, like i, had a FT-status-fitting

[2.31]Additionally cutting by race (white/nonwhite) does not change matters much.

value for their hours last week: we impute this average to people like j whose hours last week aren't FT-status-fitting for imputation because they aren't appropriately as PT work.[2.32] This means

- If a person claims to be PT last year and works between 1 and 34 hours last week, we impute their hours-per-week last year as their hours last week; otherwise (they worked 0 or more than 34 hours last week) they are imputed the group average

- If a person claims to be FT last year and works 35 or more hours last week, we impute their hours-per-week last year as their hours last week; otherwise (they worked 0-34 hours last week) they are imputed the group average

Let g' be the part of g with hours last week that are FT-status-fitting for imputation purposes (given the FT/PT nature of g), and g'' be those whose hours last week are not FT-status-fitting. Let h_i, m_i, y_i, and μ_i be worker i's hours last week, number of weeks worked last year, wage and salary income last year, and CPS Person weight, respectively. Then the measures of group g's "disaggregated" group

[2.32] In the March supplement, we have both a CPS Basic Person weight, and a CPS Supplemental Person weight. Personnel at the Census Bureau have advised us to use the latter for all the data questions we are addressing, even though some of these data are not part of the March Annual Supplement.

average, weight, hours worked last year, and income last year are

$$h_{g'} = \frac{1}{\sum_{i\in g'} \mu_i} \left(\sum_{i\in g'} h_i \mu_i \right) \tag{2.16}$$

$$\mu_g = \sum_{k\in g} \mu_k \tag{2.17}$$

$$h_g = \frac{1}{\mu_g} \left(\sum_{i\in g'} h_i \mu_i + \sum_{j\in g''} h_{g'} \mu_j \right) \tag{2.18}$$

$$y_g = \frac{1}{\mu_g} \left(\sum_{k\in g} y_k \mu_k \right) \tag{2.19}$$

Let γ be a set of gs: this is a larger group, such as all workers in the 15-19 age category, comprised of smaller "disaggregated" groups. Our construction of an efficiency wage measure for γ is similar to that of Krusell, Ohanian, Rios-Rull, and Violante (2000): our efficiency measurement f for each g is the average of their wage (y_g/h_g) for the years 1985-1989.[2.33]

$$W_\gamma = \frac{\sum_{g\in\gamma} y_g \mu_g}{\sum_{g\in\gamma} h_g f_g \mu_g} \tag{2.20}$$

It is worth mentioning that the March CPS has a specific question "On average, how many hours per week did you work last year, when you worked?" starting in 1976. We find that making sure the hours imputation is FT-status-fitting leads to hours measures that are close to the post-1976 question when both are available. By ignoring the FT-status, one underreports the groups' hours.

Our data on hours come directly from the hours last week question. Likewise, our labor force share data comes from a labor force status question pertaining to last week.

[2.33]Krusell, Ohanian, Rios-Rull, and Violante (2000) use the wage in 1980 as the efficiency measurement, but we wish to derive the efficiency from wages across both a recessionary and expansionary period.

We have found that these last week hours have level shifts between the 1967 and 1968 survey years and therefore start our hours series at 1968. The last year information used in the wage series appears unaffected during this time, so we use data going back to the 1964 survey year (data about 1963). The statistics on wages remain virtually identical if we start the wage series at survey year 1968.

References

Arias, A., G. Hansen, and L. Ohanian (2007): "Why have business cycle fluctuations become less volatile?," *Economic Theory*, 32(1), 43–58.

Baxter, M., and R. G. King (1999): "Measuring Business Cycles: Approximate Band-Pass Filters For Economic Time Series," *The Review of Economics and Statistics*, 81(4), 575–593.

Blanchard, O., and J. Simon (2001): "The Long and Large Decline in U.S. Output Volatility," *Brookings Papers on Economic Activity*, 2001(1), 135–164.

Burnside, C., M. Eichenbaum, and S. Rebelo (1995): "Capital Utilization and Returns to Scale," NBER Working Papers 5125, National Bureau of Economic Research, Inc.

Campbell, J. R., and Z. Hercowitz (2005): "The Role of Collateralized Household Debt in Macroeconomic Stabilization," NBER Working Papers 11330, National Bureau of Economic Research, Inc.

Christiano, L. J., and M. Eichenbaum (1992): "Current Real-Business-Cycle Theories and Aggregate Labor-Market Fluctuations," *American Economic Review*, 82(3), 430–50.

Clarida, R., J. Gali, and M. Gertler (2000): "Monetary Policy Rules And Macroeconomic Stability: Evidence And Some Theory," *The Quarterly Journal of Economics*, 115(1), 147–180.

Clark, K. B., and L. H. Summers (1981): "Demographic Differences in Cyclical Employment Variation," *The Journal of Human Resources*, XVI, 61–79.

DellaVigna, S., and J. M. Pollet (2007): "Demographics and Industry Returns," *American Economic Review*, 97(5), 1667–1702.

Fisher, J., and M. Gervais (2006): "First-Time Home Buyers," 2006 Meeting Papers 432, Society for Economic Dynamics.

Gomme, P., R. Rogerson, P. Rupert, and R. Wright (2004): "The business cycle and the life cycle," Working Paper 0404, Federal Reserve Bank of Cleveland.

Greenwood, J., Z. Hercowitz, and G. W. Huffman (1988): "Investment, Capacity Utilization, and the Real Business Cycle," *American Economic Review*, 78(3), 402–17.

Hansen, G. D. (1985): "Indivisible labor and the business cycle," *Journal of Monetary Economics*, 16(3), 309–327.

Hoynes, H. W. (2000): "Local Labor Markets and Welfare Spells: Do Demand Conditions Matter?," *The Review of Economics and Statistics*, (3), 351–368.

Jaimovich, N., and H. E. Siu (2007): "The young, the old, and the restless: demographics and business cycle volatility," Staff Report 387, Federal Reserve Bank of Minneapolis.

Justiniano, A., and G. E. Primiceri (2006): "The Time Varying Volatility of Macroeconomic Fluctuations," Nber working papers, National Bureau of Economic Research, Inc.

Katz, L. F., and K. M. Murphy (1992): "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *The Quarterly Journal of Economics*, 107(1), 35–78.

King, R. G., and S. T. Rebelo (1999): "Resuscitating real business cycles," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1 of *Handbook of Macroeconomics*, chap. 14, pp. 927–1007. Elsevier.

Krusell, P., L. E. Ohanian, J.-V. Rios-Rull, and G. L. Violante (2000): "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis," *Econometrica*, 68(5), 1029–1054.

Kydland, F. E., and E. C. Prescott (1993): "Cyclical movements of the labor input and its implicit real wage," *Economic Review*, (Q II), 12–23.

McConnell, M. M., and G. Perez-Quiros (2000): "Output Fluctuations in the United States: What Has Changed since the Early 1980's?," *American Economic Review*, 90(5), 1464–1476.

Nagypál, E. (2004): "Learning-by-Doing Versus Learning About Match Quality: Can We Tell Them Apart?," Discussion paper, Northwestern University.

Ramey, V. A. (2006): "Identifying Government Spending Shocks: Its All in the Timing," Discussion paper, UCSD.

Ramey, V. A., and M. D. Shapiro (1998): "Costly Capital Reallocation and the Effects of Government Spending," *Carnegie-Rochester Conference Series on Public Policy*, 48, 145–194.

Ramey, V. A., and D. J. Vine (2005): "Tracking the source of the decline in GDP volatility: an analysis of the automobile industry," Finance and economics discussion series, Board of Governors of the Federal Reserve System (U.S.).

Ravn, M. O., and H. Uhlig (2002): "On Adjusting the Hodrick-Prescott Filter for the Frequency of Observations," *The Review of Economics and Statistics*, 84, 371–380.

Rios-Rull, J.-V. (1996): "Life-Cycle Economies and Aggregate Fluctuations," *Review of Economic Studies*, 63(3), 465–89.

Rogerson, R. (1988): "Indivisible labor, lotteries and equilibrium," *Journal of Monetary Economics*, 21(1), 3–16.

Stock, J. H., and M. W. Watson (2002): "Has the Business Cycle Changed and Why?," NBER Working Papers 9127, National Bureau of Economic Research, Inc.

——— (2003): "Has the Business Cycle Changed and Why? Evidence and Explanations," Discussion paper.

Chapter 3:
Markup Variation and Endogenous Fluctuations in the Price of Investment Goods

Abstract

We present a two sector model where the evident countercyclicality of the price of investment arises endogenously due to fluctuations in the markups charged in the consumption and investment sectors. The mechanism that generates countercyclical markups is based on a model where net business formation is endogenously procyclical, as is evident in U.S. data. Based on this result, the paper suggests a simple structural method for decomposing investment price variation into exogenous and endogenous sources. The endogenous channel described here is an interaction between firms' entry and exit decisions and the degree of competition in the two sectors. This decomposition suggests that around a quarter of the movements in the price of investment goods can be attributed to this interaction. Finally, we show that the model accounts for the countercyclicality of the price of investment goods and that endogenous fluctuations in the markups are necessary to match this feature of the data.

Contents

List of Tables

3.1 Introduction

It is well known that the price of investment goods has been trending down over the last century in the U.S. (see Greenwood, Hercowitz, and Krusell (2000) and Cummins and Violante (2002)). It has been argued that this decline accounts for a significant fraction of economic growth during that period. For instance, Greenwood and Krusell (2007) concludes that more than half of postwar growth can be attributed to investment-specific technological progress.

Recent work has been focusing on the cyclical properties of this price series and has emphasized that in the U.S. (*i*) the price of investment goods is countercyclical, and that (*ii*) fluctuations in investment-specific technological progress (the inverse of the price of investment goods) contribute significantly to postwar U.S. business cycles. For example, Greenwood, Hercowitz, and Krusell (2000) suggests that this form of technological change is the source of about 30% of output fluctuations. Similarly, Fisher (2006) and Justiniano and Primiceri (2006) argue that investment-specific technological progress is the most important determinant of output variability.

Motivated by this evidence we ask if the cyclical fluctuations in the price of investment goods are entirely exogenous, or if instead a fraction of them can be attributed to endogenous movements. For example, Fisher (2006) points out that while "investment-specific technology shocks could play a key role in short-run fluctuations, the short-run correlations might be driven at least partly by factors other than technological change, such as time-varying markups." To address this issue we build a two-sector dynamic general equilibrium model, modified from Jaimovich and Floetotto (2008), where variations in the level of technology lead to changes in the number of competitors in both the consumption goods and investment goods sectors.[3.1]

Based on this model, which we present in some detail in Section 3.2, we derive a simple structural method for decomposing variation in the price of investment goods into exogenous sources and sources originating endogenously from the interaction of entry and exit decisions and markup variations. Specifically, in our model, the price of investment goods is positively correlated with the ratio between the investment sector markup and the consumption sector markup. Therefore, endogenous movements in these markups induce endogenous fluctuations in the price of investment goods. Based on our preferred calibration we show in Section 3.3 that around a quarter of the movement in the price of investment goods can be attributed to these endogenous fluctuations. It is important to note that this result depends only on the assumptions we make regarding the technology and market structure. It is independent of the preference specification we later assume.

In Sections 3.4 and 3.5 we model the household side and calibrate and simulate the model, showing its consistency with the main empirical regularities characterizing the postwar U.S. business cycle. In particular, we show that the model, when simulated, succeeds in accounting for the countercyclicality of the price of investment goods. We show that endogenous fluctuations in sectoral markups are necessary to

[3.1] Previous work on two-sector neoclassical models includes Baxter (1996), Hornstein and Praschnik (1997), Huffman and Wynne (1999), and Harrison (2003). Of course, the seminal paper by Long and Plosser (1983) was a multi-sector model.

match this feature of the data.

The approach of this paper is related to a literature following Hall (1986) which suggests that measured Total Factor Productivity (TFP) has important endogenous components (e.g. Hall (1988), Hall (1990), Basu and Fernald (2002)). As this literature has emphasized, the presence of these endogenous components leads researchers to overestimate the variance of TFP shocks.[3.2] In the same spirit, we propose a mechanism that leads to endogenous movements in the price of investment goods, and then quantify its contribution to the overall fluctuations in the data.

Before presenting the model it is worth emphasizing that the model represents a minimal perturbation of the prototype perfect competition two sector real business cycle (RBC) model. This greatly simplifies comparison with existing work and allows for a simple structural decomposition of the price of investment goods. However, this simplicity is purchased at the cost of descriptive realism, and two empirical caveats should be highlighted. First, the model here is symmetric. This implies that the number of firms moves together in all industries within one sector. One might be worried that the procyclicality in the number of firms in the data is really driven by only a few industries. Second, in our model entrants have the same size as existing firms. It is well known, however, that smaller firms make up the majority of entrants and exits. This may imply that variations in their number are potentially less important and that entry rates should be weighted by the size of entrants. We refer the reader to Jaimovich and Floetotto (2008) for a detailed discussion of this.

[3.2]Similarly, Kim (2006) found that investment-specific technology shocks are Granger-caused by variables used in Evans (1992)'s analogous finding for Solow residuals.

3.2 Technology and Market Structure

There are two sectors of production: consumption and investment. Production of each the consumption and the investment good follow the setup of Jaimovich and Floetotto (2008). Within each period, capital can be costlessly reallocated from one sector to the other.[3.3] We thus do not have to track sector-specific capital but only aggregate capital K as a state variable. The remainder of this section presents the consumption sector in some detail. The investment sector is exactly analogous.

3.2.1 The Consumption Sector

The consumption good is produced in a large number of differentiated industries. In each industry, there is a finite number of intermediate firms each producing differentiated goods which are imperfect substitutes in the production of a industrial good. Similarly, the industrial goods are imperfect substitutes in the production of the consumption sector good. Entry into the existing industries is costless for intermediate producers. Hence, a zero-profit condition is satisfied in each period and every industry.

The sectoral good is produced with a constant-returns-to-scale production function, which aggregates a measure one continuum of industrial goods

$$C_t = \left[\int_0^1 Q_t^c(j)^{\omega} dj \right]^{\frac{1}{\omega}}, \ \omega \in (0,1)$$

where $Q_t^c(j)$ denotes output of industry j. The elasticity of substitution between any

[3.3] See Huffman and Wynne (1999) for a two-sector model with adjustment costs in the allocation of capital between sectors.

two industrial goods is constant and equals $\frac{1}{1-\omega}$. The consumption good producers behave competitively.

In each of the consumption industries, there are N_t^c firms producing differentiated intermediate goods. A CES function aggregates those to yield the output of industry j

$$Q_t^c(j) = (N_t^c)^{1-\frac{1}{\tau}} \left[\sum_{i=1}^{N_t^c} x_t^c(j,i)^\tau \right]^{\frac{1}{\tau}}, \ \tau \in (0,1) \tag{3.1}$$

where $x_t(j,i)$ is the output of firm i in industry j.[3.4] The elasticity of substitution between any two goods within a industry is constant and equals $\frac{1}{1-\tau}$. The market structure of each industry exhibits monopolistic competition; each differentiated $x_t^c(j,i)$ is produced by one firm that sets the price for its good in order to maximize profits. Finally, it is assumed that the elasticity of substitution between any two goods within an industry is higher than the elasticity of substitution across industries, $\frac{1}{1-\omega} < \frac{1}{1-\tau}$.

Each intermediate good, $x_t^c(j,i)$, is produced using capital, $k_t^c(j,i)$, and labor, $h_t^c(j,i)$. The parameter $\phi^c > 0$ represents an *overhead cost*. In each period, an amount ϕ^c of the intermediate good is immediately used up, independent of how much output is produced. As in Rotemberg and Woodford (1996) the role of this parameter is to allow the model to reproduce the apparent absence of pure profits in actual U.S. industries despite the presence of market power. Therefore, production of the intermediate good follows the function

$$x_t^c(j,i) = z_t^c k_t^c(j,i)^\alpha h_t^c(j,i)^{1-\alpha} - \phi^c, \ \alpha \in [0,1] \tag{3.2}$$

[3.4] The term $N^{1-\frac{1}{\tau}}$ in (3.1) implies that there is no *variety effect* in the model.

As is standard in the literature, log technology shocks follow a stationary first order auto-regressive process. It is assumed that $|\zeta^c| < 1$, and that ε_t^c is a normally distributed random variable, with a mean of zero and a standard deviation σ_ε^c.

$$\ln z_t^c = \zeta^c \ln z_{t-1}^c + \varepsilon_t^c \tag{3.3}$$

The sectoral good producer solves a static optimization problem that results in the usual conditional demand for each industrial good, $Q_t^c(j)$, where $p_t^c(j)$ is the price index of industry j at period t and P_t^c is the price of the consumption good at period t,

$$Q_t^c(j) = \left[\frac{p_t^c(j)}{P_t^c}\right]^{\frac{1}{\omega-1}} C_t \tag{3.4}$$

$$P_t^c = \left[\int_0^1 p_t^c(j)^{\frac{\omega}{\omega-1}} dj\right]^{\frac{\omega-1}{\omega}}. \tag{3.5}$$

Denoting the price of good i in industry j in period t by $p_t^c(j,i)$, the conditional demand faced by the producer of each $x_t^c(j,i)$ variant is similarly defined as

$$x_t^c(j,i) = \left[\frac{p_t^c(j,i)}{p_t^c(j)}\right]^{\frac{1}{\tau-1}} \frac{Q_t^c(j)}{N_t^c} \tag{3.6}$$

$$p_t^c(j) = (N_t^c)^{1-\frac{1}{\tau}} \left[\sum_{i=1}^{N_t^c} p_t^c(j,i)^{\frac{\tau}{\tau-1}}\right]^{\frac{\tau-1}{\tau}}.$$

Using (3.4) and (3.6), the conditional demand for good $x_t^c(j,i)$ at period t can then be expressed in terms of the consumption good as

$$x_t^c(j,i) = \left[\frac{p_t^c(j,i)}{p_t^c(j)}\right]^{\frac{1}{\tau-1}} \left[\frac{p_t^c(j)}{P_t^c}\right]^{\frac{1}{\omega-1}} \frac{C_t}{N_t^c}.$$

3.2.2 The Elasticity of Demand

In a standard setup following Dixit and Stiglitz (1977) each firm is small relative to the economy. When setting its price, the firm therefore does not take the effect of its own behavior on the remaining firms into account. In that case, the $x_t^c(j,i)$ producer has no effect on the sectoral price level, $p_t^c(j)$, or on the aggregate price level, P_t^c. The elasticity of demand is thus constant which implies the well known constant markup rule

$$\begin{aligned} \eta_{x^c(j,i)p^c(j,i)} &= \frac{1}{\tau - 1} < 0 \\ \frac{p_t^c(j,i)}{MC_t^c(j,i)} &= \mu^c = \frac{1}{\tau}. \end{aligned} \tag{3.7}$$

In our setup, within each sector there is a continuum of industries, but within each industry the number of operating firms is finite. While an individual firm's decisions have no affect on the general price level P_t^c, they do affect the industrial price level $p_t^c(j)$. The resulting price elasticity of demand is then a function of the number of firms within a industry, N_t^c. In a symmetric equilibrium, the elasticity becomes

$$\eta_{x^c(j,i)p^c(j,i)}(N_t^c) = \frac{1}{\tau - 1} + \left[\frac{1}{\omega - 1} - \frac{1}{\tau - 1}\right]\frac{1}{N_t^c} \tag{3.8}$$

implying that an increase in N_t^c leads to a more elastic demand curve. At the solution to the monopolistic firm's problem, marginal revenue equals marginal cost

$$\frac{p_t^c(j,i)}{MC_t^c(j,i)} = \mu^c(N_t^c) = \frac{(1-\omega)N_t^c - (\tau - \omega)}{\tau(1-\omega)N_t^c - (\tau - \omega)} > 1. \tag{3.9}$$

Note that the markup function is monotonically decreasing in the number of firms, i.e. $\frac{d\mu^c}{dN^c} < 0$, and that $\tau\mu^c(N^c) > 1$.

We assume that the economy's technology is symmetric with respect to all intermediate inputs and hence we focus on symmetric equilibria.

$$\forall(j,i) \in [0,1]\text{x}[1,N_t]:$$

$$x_t^c(j,i) = x_t^c,\ k_t^c(j,i) = k_t^c,\ h_t^c(j,i) = h_t^c,\ p_t^c(j,i) = p_t^c,\ N_t^c(j) = N_t^c$$

Total capital and hours in the consumption sector are then given by $K_t^c = N_t^c k_t^c$ and $H_t^c = N_t^c h_t^c$ respectively. Finally, in the symmetric equilibrium, a zero-profit condition is imposed in every sector in every period.

$$\left(\mu_t^c - 1\right) x_t^c = \phi^c. \tag{3.10}$$

The number of firms per industry and aggregate final output can now be found by using (3.2) and the zero-profit condition (3.10).[3.5]

$$\begin{aligned} N_t^c &= z_t^c \left(\frac{K_t^c}{H_t^c}\right)^\alpha H_t^c \left[\frac{\mu_t^c - 1}{\mu_t^c \phi^c}\right] && (3.11) \\ C_t &= \frac{z_t^c}{\mu_t^c} \left(\frac{K_t^c}{H_t^c}\right)^\alpha H_t^c. && (3.12) \end{aligned}$$

We can see that N_t^c is procyclical, implying that μ^c is countercyclical, by rewriting (3.11) and (3.12) as

$$N_t^c = \left[\frac{\mu^c(N_t^c) - 1}{\phi^c}\right] C_t. \tag{3.13}$$

We use P_t^c as the numeraire and set it to 1. This implies that the price charged by an intermediate producer in the consumption sector is also 1 in a symmetric equilibrium.

[3.5] To see this, multiply (3.2) with N_t^c and use the zero-profit condition to plug in for x_t^c. In order to find C_t, multiply (3.11) by x_t^c and use the zero-profit condition again.

3.2.3 The Investment Sector

The setup in the investment good sector is exactly analogous. There is a continuum of industries, each with a finite number N_t^i of intermediate producers. The investment good, the industrial goods, and the differentiated goods are thus produced according to

$$\begin{aligned} I_t &= \left[\int_0^1 Q_t^i(j)^\omega dj\right]^{\frac{1}{\omega}}, \ \omega \in (0,1) \\ Q_t^i(j) &= \left(N_t^i\right)^{1-\frac{1}{\tau}} \left[\sum_{i=1}^{N_t^i} x_t^i(j,i)^\tau\right]^{\frac{1}{\tau}}, \ \tau \in (0,1) \\ x_t^i(j,i) &= z_t k_t^i(j,i)^\alpha h_t^i(j,i)^{1-\alpha} - \phi^i, \ \alpha \in [0,1]. \end{aligned}$$

The solution to the optimization problem leads to an analogous expression for the markup charged in the investment sector

$$\frac{p_t^i(j,i)}{MC_t^i(j,i)} = \mu^i(N_t^i) = \frac{(1-\omega)N_t^i - (\tau-\omega)}{\tau(1-\omega)N_t^i - (\tau-\omega)} > 1.$$

Total capital and hours in the investment sector are given by $K_t^i = N_t^i k_t^i$ and $H_t^i = N_t^i h_t^i$ respectively. The zero-profit condition holds in every sector in every period,

$$\left(\mu_t^i - 1\right) x_t^i = \phi^i,$$

which as above allows us to derive the number of firms per sector and aggregate investment

$$\begin{aligned} N_t^i &= z_t^i \left(\frac{K_t^i}{H_t^i}\right)^\alpha H_t^i \left[\frac{\mu_t^i - 1}{\mu_t^i \phi^i}\right] \\ I_t &= \frac{z_t^i}{\mu_t^i} \left(\frac{K_t^i}{H_t^i}\right)^\alpha H_t^i. \end{aligned} \tag{3.14}$$

3.2.4 The Price of Investment

In this economy capital and labor are mobile across sectors and industries. In equilibrium, factor prices have to be equalized in the consumption and investment sector. This allows us to derive a simple expression for the price of investment.

$$\frac{z_t^c}{\mu_t^c}(1-\alpha)\left(\frac{k_t^c}{h_t^c}\right)^{\alpha} = W_t \quad = \quad P_t^i \frac{z_t^i}{\mu_t^i}(1-\alpha)\left(\frac{k_t^i}{h_t^i}\right)^{\alpha} \tag{3.15}$$

$$\frac{z_t^c}{\mu_t^c}\alpha\left(\frac{k_t^c}{h_t^c}\right)^{\alpha-1} = R_t \quad = \quad P_t^i \frac{z_t^i}{\mu_t^i}\alpha\left(\frac{k_t^i}{h_t^i}\right)^{\alpha-1} \tag{3.16}$$

Solving (3.15) or (3.16) for P_t^i then leads to the following expression for the price of investment

$$P_t^i = \left(\frac{z_t^c}{z_t^i}\right)\left(\frac{\mu_t^i}{\mu_t^c}\right)\left(\frac{k_t^c}{h_t^c}\right)^{\alpha}\left(\frac{k_t^i}{h_t^i}\right)^{-\alpha}.$$

This, however, can be simplified further by combining (3.15) and (3.16) which yields the well known result that the capital labor ratio is the same in both industries, $\frac{k_t^c}{k_t^c} = \frac{k_t^i}{k_t^i}$. The price of investment then becomes

$$P_t^i = \left(\frac{z_t^c}{z_t^i}\right)\left(\frac{\mu_t^i}{\mu_t^c}\right). \tag{3.17}$$

The first term is standard: when productivity in the investment sector increases relative to the consumption sector, investment goods become cheaper in terms of consumption goods. The second term, however, is a result of the particular sectoral structure that we have assumed. When the investment sector becomes more competitive relative to the consumption sector, i.e. $\frac{\mu_t^i}{\mu_t^c}$ falls, the price of investment falls. This equation is the basis of our quantitative exercise in the next section.

Note that aggregate output can be expressed in terms of consumption goods as

$$
\begin{aligned}
Y_t &= C_t + P_t^i I_t \\
&= \frac{z_t^c}{\mu_t^c}\left(\frac{K_t^c}{H_t^c}\right)^{\alpha} H_t^c + P_t^i \frac{z_t^i}{\mu_t^i}\left(\frac{K_t^i}{H_t^i}\right)^{\alpha} H_t^i \\
&= \frac{z_t^c}{\mu_t^c}\left(\frac{K_t}{H_t}\right)^{\alpha} H_t
\end{aligned}
$$

which implies that

$$
\begin{aligned}
W_t &= (1-\alpha)\frac{Y_t}{H_t} \\
R_t &= \alpha\frac{Y_t}{K_t}.
\end{aligned}
$$

3.3 Decomposing the Price of Investment

The simple expression for the price of investment gives rise to an interesting decomposition exercise. Using a hat to denote log deviations from the steady state, (3.17) can be expressed as

$$\hat{p}_t = \hat{\mu}_t^i - \hat{\mu}_t^c + \hat{z}_t^c - \hat{z}_t^i. \tag{3.18}$$

While the price of investment is in principal observable, all terms on the right hand side of the equation are latent. However, the model's equilibrium conditions imply that we can express $\hat{\mu}_t^i$ and $\hat{\mu}_t^c$ as functions of observable data. Using (3.9) and (3.13) one can show that $\frac{1}{\tau^c\mu^c} = \frac{C_t}{\phi^c}\left(\frac{1-\omega^c}{\tau^c-\omega^c}\right)$. Log linearization then yields the following expressions

$$\hat{\mu}_t^c = \left(\frac{1-\tau^c\mu^c}{\tau^c\mu^c}\right)\hat{c}_t = A\hat{c}_t \tag{3.19}$$

$$\hat{\mu}_t^i = \left(\frac{1-\tau^i\mu^i}{\tau^i\mu^i}\right)\hat{\imath}_t = B\hat{\imath}_t, \tag{3.20}$$

where A and B are constants. Using these two equation we can restate (3.18) as

$$\hat{p}_t = B\hat{\imath}_t - A\hat{c}_t + \hat{z}_t^c - \hat{z}_t^i.$$

Note that four parameters are required to assign values to A and B. First, there are the steady state markups in the consumption and investment sectors. To the best of our knowledge, there is no clear evidence on the average size of markups in each of these sectors. Similar to Jaimovich and Floetotto (2008), we thus calibrate the steady state value of the markup to 1.3 in both sectors.

Second, we have to calibrate τ^c and τ^i, the parameters that determine the elasticity of substitution between differentiated goods in the consumption and investment sector. Given the steady state markup, it is straightforward to derive these. Taking the consumption sector as an example, use (3.9) to find $N_t^c = \left(\frac{\tau^c - \omega^c}{1-\omega^c}\right)\frac{\mu_t^c - 1}{\tau^c \mu_t^c - 1}$. Log linearization then leads to an expression for the elasticity of the number of firms with respect to that sector's output,

$$\hat{n}_t^c = \left(\frac{1-\tau^c}{\tau^c(\mu^c - 1)}\right)\hat{c}_t, \tag{3.21}$$

with a similar formula for the investment sector

$$\hat{n}_t^i = \left(\frac{1-\tau^i}{\tau^i(\mu^i - 1)}\right)\hat{\imath}_t. \tag{3.22}$$

Using these expressions, it is clear that we can easily back out the values for τ^c and τ^i from the elasticity of the number of firms in a given sector with respect to that sector's output. Estimation of τ^c thus requires data on the number of firms in the consumption sector while estimation of τ^i requires data on the number of firms in the investment sector. Jaimovich and Floetotto (2008) argue that changes in the number of establishments might be a better measure of changes in the number of competitors in the economy. We therefore use data on the number of establishments in the two sectors, constructed as follows.

Data on the number of establishments in thirteen nonagricultural supersectors come from the Bureau of Labor Statistics' Business Employment Dynamics database. At present, these data run from 1992:3-2007:2. To arrive at the total number of establishments in a supersector, we add Expansions (businesses that were already in existence and added employees) to Contractions (businesses that were already in existence and shed employees) to Openings (businesses that came into existence), and then subtract Closings (businesses that closed).

To normalize the number of establishments across supersectors, we turn to data from the Small Business Administration. They record the value of estimated receipts, annual payrolls, and employment for twenty large nonagricultural industry groups on an annual basis for 1988-2005. These twenty groups are a subpartition of the partition of thirteen above and can be added up to get values for these three variables for the supersectors. We rank the relative size of the supersector by each variable, and average this weighting across the time span. In our paper we prefer to use the normalization based on annual payrolls because it most resembles a weighting by output share (if the capital-output ratio across supersectors is identical); nonetheless, our results are robust to using any of them.

With normalized establishment counts for the supersectors in hand, it remains to divide them into the Consumption and Investment sectors. To do this, we modify the procedure of Harrison (2003). Looking at the Bureau of Economic Analysis' Input-Output Use table, we are able to see the amount of output used for Personal Consumption Expenditure or for Fixed Private Investment for eighty-four nonagricultural industries similar to 2-digit SIC industries. Adding up the industries within each supersector, we arrive at a value of Personal Consumption Expenditure and Fixed Private Investment for each supersector. We then define a supersector's consumption sector share as (Personal Consumption Expenditure)/(Personal Consumption Expenditure plus Fixed Private Investment), while the investment sector share is Fixed Private Investment over the same denominator.[3.6] This procedure provides us with a quarterly series of establishments for the consumption and the investment sector.[3.7]

The resulting two series are our measures of N_t^c and N_t^i. We use the Hodrick-Prescott (HP) filter (smoothing parameter set at 1600) to calculate the deviations $\hat{n}_t^c$ and $\hat{n}_t^i$ and estimate the elasticity in (3.21) by regressing $\hat{n}_t^c$ on $\hat{c}_t$ and a constant.[3.8] We find an elasticity of 0.64 which implies a value of 0.84 for τ^c. Using (3.22) for the investment sector, we find an elasticity of 0.28. The implied value of τ^i is 0.92.[3.9]

With a calibration of τ^c, τ^i, μ^c, and μ^i we find that the constant $A = -0.1055$ and $B = -0.1817$. We can now use a simple variance covariance decomposition to

[3.6] These shares are virtually identical regardless of the Use table year.

[3.7] For example, take the Transportation & Warehousing supersector. We see that 92% of its output goes to consumption and 8% goes to investment, according to the Use table. This supersector accounts for 3.38% of aggregate payrolls and has 76,000 establishments in 1992:3. Therefore, it accounts for $0.92 \times 0.0338 \times 76000 = 2363.296$ consumption sector establishments in that quarter.

[3.8] Data on real consumption and investment are taken from NIPA via FRED, acronyms PCECC96 and FPIC96. We have that $\sigma(\hat{n}^c) = 0.0082, \sigma(\hat{n}^i) = 0.0130, \sigma(\hat{c}) = 0.0060, \sigma(\hat{\imath}) = .0343, \rho(\hat{n}^c, \hat{n}^i) = 0.92, \rho(\hat{n}^c, \hat{c}) = 0.48, \rho(\hat{n}^i, \hat{\imath}) = 0.75, \rho(\hat{c}, \hat{\imath}) = 0.67$.

[3.9] The implied values of $\tau^\cdot$ are robust to using log differences instead of the HP deviations.

calculate the share of the investment price variation that can be attributed to the mechanism emphasized in this paper.

$$\mathrm{Var}(\hat{p}_t) = \mathrm{Var}(B\hat{\imath}_t - A\hat{c}_t) + \mathrm{Var}(\hat{z}_t^c - \hat{z}_t^i) + 2 \times \mathrm{Cov}(B\hat{\imath}_t - A\hat{c}_t, \hat{z}_t^c - \hat{z}_t^i)$$

Using the parameters implied by our baseline calibration we find that 24% of the variation in the price of investment are due to the time-variation in markups. It is important to note that none of the above results require imposing any restrictions on the model's specification of household behavior. The simple expression for the price of investment goods can de derived from the assumptions on technology alone.

3.4 Household and Preferences

In order to simulate the economy we need to close the model by specifying the household side. It is well known that comovement of hours worked does not arise in the benchmark two sector model with separable preferences as in King, Plosser, and Rebelo (1988); for a discussion of this issue see Christiano and Fitzgerald (1998). The same holds true for our extension of the benchmark model. We use the results in Jaimovich and Rebelo (2005) who show that this failure can be remedied by assuming a utility functions with a weak short-run wealth effects on the labor supply. We thus use the utility function of Greenwood, Hercowitz, and Huffman (1988), which implies the absence of a wealth effect on labor supply.

At each point in time the economy is inhabited by a continuum of identical households. The mass of households is normalized to one. It is assumed that the representative agent has preferences over random streams of consumption and leisure.

The representative agent chooses a sequence of consumption, hours and investments in capital to maximize

$$\max_{\{H_t, C_t, K_{t+1}\}} E_0 \sum_{t=0}^{\infty} \beta^t \log \left(C_t - \frac{\theta}{1+\frac{1}{\chi}} H_t^{1+\frac{1}{\chi}} \right)$$

subject to the sequential budget constraint and the law of motion for capital

$$\begin{aligned} C_t + P_t^i I_t &= R_t K_t + W_t H_t + \Pi_t \\ K_{t+1} &= (1-\delta) K_t + I_t \end{aligned}$$

where the initial capital stock is given and equal to K_0. C_t and H_t denote consumption and hours worked by the household in period t. $\beta \in (0,1)$ and $\delta \in (0,1)$ denote the subjective time discount factor and the depreciation rate of capital. $\chi \geq 0$ is the Frisch labor supply elasticity, and $\theta > 0$. Households own the capital stock and take the equilibrium rental rate, R_t, and the equilibrium wage, W_t, as given. Finally, the households own the firms and receive their profits, Π_t. Combining the two constraints and denoting the Lagrange multiplier by λ we can find the first order conditions

$$\begin{aligned} \theta H_t^{\frac{1}{\chi}} &= w_t \\ \lambda_t &= \left(C_t - \frac{\theta}{1+\frac{1}{\chi}} H_t^{1+\frac{1}{\chi}} \right)^{-1} \\ \lambda_t &= \lambda_{t+1} \beta \left((1-\delta) \frac{P_{t+1}^i}{P_t^i} + \frac{R_{t+1}}{P_t^i} \right). \end{aligned}$$

3.5 Calibration and Simulation

We now continue by simulating the model economy. We adopt a standard calibration of the parameters in the model – see Table 3.1. In order to simulate the

Table 3.1: Calibration

	Parameter	
$\mu^* - 1$	Markup in steady state	30%
τ^c	Elasticity within industry (consumption sector)	0.84
τ^i	Elasticity within industry (investment sector)	0.92
α	Capital share	0.30
H^*	Time spent working	0.33
β	Time discount factor	0.99
δ	Depreciation rate	0.025

Note: The calibration of μ, τ^c and τ^i is explained in the main text. The value of ω^c and ω^i do not matter for the results. The AR(1) parameter on productivity in the two sectors are estimated to be $\zeta^c = 0.83$ and $\zeta^i = 0.84$ while the shocks have standard deviations $\sigma(\varepsilon_t^c) = 0.0078, \sigma(\varepsilon_t^i) = 0.0068$. The correlation between the shocks is $\rho(\varepsilon_t^c, \varepsilon_t^i) = 0.61$. The remaining parameters are standard.

model we also need to specify the parameters governing the stochastic process of z_t^c and z_t^i. We use the model's equilibrium conditions to identify the technology shocks in the two sectors. Log-linearizing equations (3.12) and (3.14) and using the exact same substitutions as in (3.19) and (3.20) leads to

$$\hat{z}_t^c = (1+A)\hat{c}_t - \alpha\left(\hat{k} - \hat{h}\right) - \hat{h}_t^c \tag{3.23}$$

$$\hat{z}_t^i = (1+B)\hat{\imath}_t - \alpha\left(\hat{k} - \hat{h}\right) - \hat{h}_t^i. \tag{3.24}$$

There are two ways in which a series of H_t^c and H_t^i could be constructed. First, we could use a similar approach to the one we used for the construction of the number of establishments in the two sectors discussed above. However, we can also use the equilibrium conditions of the model to infer the two series. To see this, note that the price of investment goods equals

$$P_t^i = \frac{C_t}{H_t^c}\frac{H_t^i}{I_t} = \frac{C_t}{I_t}\frac{(H_t - H_t^c)}{H_t^c},$$

which follows from (3.12) and (3.14). Hence, using data on P_t^i, C_t, I_t, and H_t we can construct a series of H_t^c and H_t^i.

We then use H_t^c and H_t^i in (3.23) and (3.24) to estimate $\hat{z}_t^c$ and $\hat{z}_t^i$. For $\hat{z}_t^c$ we estimate the AR1 coefficient ζ^c to equal 0.83 and a standard deviation of $\sigma(\varepsilon_t^c) = 0.0078$. Similarly, for $\hat{z}_t^i$ we estimate the AR1 coefficient ζ^i to equal 0.84 and $\sigma(\varepsilon_t^i) = 0.0068$. Finally, we estimate the correlation $\rho(\varepsilon_t^c, \varepsilon_t^i) = 0.605$.[3.10]

3.5.1 Results of Simulation

Panel I in Table 3.2 reports moments for the U.S.[3.11] As previously documented the price of investment is as volatile as output and negatively correlated with output at the business cycle frequency. Our benchmark model (Panel II) produces a price of investment series that is countercyclical with a similar magnitude to the one observed in the data. The model underperforms with respect to the volatility of the series: the ratio of standard deviations of the price of investment to output is 1.07 in the data while this ratio equals 0.75 in the model.

In order to assess the role of the endogenous markups in generating this neg-

[3.10]The estimation is done as follows. We treat $\hat{z}_t^c$ and $\hat{z}_t^i$ as first differences and from which we can build a level series of the two shocks. We then follow the approach in King and Rebelo (1999). They assume that $\log z_t^c$ and $\log z_t^i$ exhibit a linear trend which they use to construct deviations. Using these approach we then estimate $\zeta^c = 0.96$ and $\sigma(\varepsilon_t^c) = 0.008$, $\zeta^c = 0.99$, and $\sigma(\varepsilon_t^i) = 0.008$ and $\rho(\varepsilon_t^c, \varepsilon_t^i) = 0.56$. When calibrated with these parameters, the model generates a countercyclical price of investment too. However, the resulting series of $\log z_t^c$ and $\log z_t^i$ exhibit a non linear trend. Hence, our preferred calibration is based on an estimate that uses a more flexible specification of the trend, i.e. an HP trend.

[3.11]We use data from 1955:1-2000:4. We thank Jonas Fisher for the price of investment series.

Table 3.2: Data and Model Moments

	I – Data			II – Benchmark			III – Constant Markups		
	$\sigma(x)$	$\frac{\sigma(x)}{\sigma(y)}$	$\rho(x,y)$	$\sigma(x)$	$\frac{\sigma(x)}{\sigma(y)}$	$\rho(x,y)$	$\sigma(x)$	$\frac{\sigma(x)}{\sigma(y)}$	$\rho(x,y)$
Output (y)	0.015	1	1	0.018	1	1	0.017	1	1
Consumption	0.012	0.80	0.85	0.011	0.59	0.85	0.012	0.72	0.97
Investment	0.047	3.16	0.87	0.065	3.57	0.86	0.035	2.11	0.87
Hours	0.018	1.17	0.78	0.011	0.60	1.00	0.010	0.60	1.00
Hours (C Sector)	0.012	0.80	0.49	0.006	0.33	0.32	0.006	0.35	0.87
Hours (I Sector)	0.037	2.47	0.86	0.049	2.69	0.91	0.030	1.82	0.94
Price of Investment	0.016	1.07	**−0.39**	0.014	0.75	**−0.32**	0.008	0.48	**0.56**
Markups (C Sector)				0.001	0.06	−0.85			
Markups (I Sector)				0.012	0.65	−0.86			

Note: Second moments of data, benchmark model, and model with constant markups. Benchmark model has endogenous markups. Sectoral hours constructed as described in Section 3.5. See text for details.

ative correlation, Panel III reports the results of the same model with the same technology shocks where the markup is a constant. Hence, the only difference between Panels II and III is along the mechanism emphasized in this paper, i.e. the endogeneity of the markup. Note from Panel III that the model generates price of investment time series that is both (*i*) less volatile than in the benchmark model, and more importantly (*ii*) *positively* correlated with output. Hence, endogenous movements in the markup are necessary for the model to generate a countercyclical process of the price of investment goods.[3.12]

With respect to other variables of interest, the performance of the model is rather standard. Investment is more volatile than output, consumption is less volatile

[3.12] When we simulate the model with the stochastic process based on a linear trend in the shocks as discussed in footnote 3.5 the model generates a negative correlation between the price of investment goods and output of −0.17 and a relative standard deviations equal to 0.68.

than output, and the model underestimates the volatility of hours worked. Interestingly, the benchmark model (Panel II) generates a correlation between the hours in the two sectors and output that resembles the estimates we obtain in the data.[3.13]

3.6 Conclusion

This paper formulates a simple structural two sector model in a general equilibrium framework in which technology shocks induce the entry and exit of competitors. Endogenous variation in the number of operating firms in the two sectors lead to endogenous variation in the degree of competition over the business cycle. We show that this model economy implies that the price of investment goods can be decomposed into an exogenous component as well as an endogenous component resulting from the entry and exit of firms. Based on this decomposition, the paper suggests that about a quarter of the variation in the price of investment in the U.S. are due to this interaction. Moreover, the model, when simulated, accounts for the countercyclicality of the price of investment goods. We show that, within our model, endogenous fluctuations in the markups are necessary to match this feature of the data.

Acknowledgement: *this chapter is joint work with Max Floetotto and Nir Jaimovich. The dissertation author was a primary author of this paper.*

[3.13]The moments we report in Panel I with respect to H^c and H^i are based on series we construct as discussed in Section 3.5.

References

Basu, S., and J. G. Fernald (2002): "Aggregate productivity and aggregate technology," *European Economic Review*, 46(6), 963–991.

Baxter, M. (1996): "Are Consumer Durables Important for Business Cycles?," *The Review of Economics and Statistics*, 78(1), 147–55.

Christiano, L. J., and T. J. Fitzgerald (1998): "The Business Cycle: It's Still a Puzzle," *Economic Perspectives*, 22.

Cummins, J. G., and G. L. Violante (2002): "Investment-Specific Technical Change in the US (1947-2000): Measurement and Macroeconomic Consequences," *Review of Economic Dynamics*, 5(2), 243–284.

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

Evans, C. L. (1992): "Productivity shocks and real business cycles," *Journal of Monetary Economics*, 29(2), 191–208.

Fisher, J. D. M. (2006): "The Dynamic Effects of Neutral and Investment-Specific Technology Shocks," *Journal of Political Economy*, 114(3), 413–451.

Greenwood, J., Z. Hercowitz, and G. W. Huffman (1988): "Investment, Capacity Utilization, and the Real Business Cycle," *American Economic Review*, 78(3), 402–17.

Greenwood, J., Z. Hercowitz, and P. Krusell (2000): "The role of investment-specific technological change in the business cycle," *European Economic Review*, 44(1), 91–115.

Greenwood, J., and P. Krusell (2007): "Growth accounting with investment-specific technological progress: A discussion of two approaches," *Journal of Monetary Economics*, 54(4), 1300–1310.

Hall, R. E. (1986): "Market Structure and Macroeconomic Fluctuations," *Brookings Papers on Economic Activity*, 2, 285–338.

——— (1988): "The Relation between Price and Marginal Cost in U.S. Industry," *Journal of Political Economy*, 96(5), 921–47.

——— (1990): "The Invariance Propereties of Solow's Productivity Residual," in *Growth, Productivity, and Unemployment*, ed. by P. Diamond. MIT Press.

Harrison, S. G. (2003): "Returns to Scale and Externalities in the Consumption and Investment Sectors," *Review of Economic Dynamics*, 6(4), 963–976.

Hornstein, A., and J. Praschnik (1997): "Intermediate inputs and sectoral comovement in the business cycle," *Journal of Monetary Economics*, 40(3), 573–595.

Huffman, G. W., and M. A. Wynne (1999): "The role of intratemporal adjustment costs in a multisector economy," *Journal of Monetary Economics*, 43(2), 317–350.

Jaimovich, N., and M. Floetotto (2008): "Firm Dynamics, Markup Variations, and the Business Cycle," Discussion paper, Stanford University.

Jaimovich, N., and S. Rebelo (2005): "Can News About the Future Drive the Business Cycle?," Discussion paper.

Justiniano, A., and G. E. Primiceri (2006): "The Time Varying Volatility of Macroeconomic Fluctuations," Nber working papers, National Bureau of Economic Research, Inc.

Kim, K. H. (2006): "Is Investment-Specific Technological Change Really Important for Business Cycles?," Discussion paper, UCSD.

King, R. G., C. I. Plosser, and S. T. Rebelo (1988): "Production, growth and business cycles : I. The basic neoclassical model," *Journal of Monetary Economics*, 21(2-3), 195–232.

King, R. G., and S. T. Rebelo (1999): "Resuscitating real business cycles," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1 of *Handbook of Macroeconomics*, chap. 14, pp. 927–1007. Elsevier.

Long, Jr., J. B., and C. I. Plosser (1983): "Real Business Cycles," *The Journal of Political Economy*, 91(1), 39–69.

Rotemberg, J. J., and M. Woodford (1996): "Imperfect Competition and the Effects of Energy Price Increases on Economic Activity," *Journal of Money, Credit and Banking*, 28(4), 550–77.