Essays on the U.S. Labor Market

By

Matthew D. French

Dissertation
Submitted to the Faculty of the
Graduate School of Vanderbilt University
in partial fulfillment of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in

Economics

May, 2017

Nashville, Tennessee

Approved:

Kevin Huang, Ph.D.

Robert Driskill, Ph.D.

Gregory Huffman, Ph.D.

Hyunseung Oh, Ph.D.

Miguel Palacios, Ph.D.

Dedication

To Charlotte

Acknowledgments

I thank the Vanderbilt Ph.D. program in Economics for the financial and educational support they provided over the years.

I thank my cohort, classmates and friends. Jonah Yuen, Jason Campbell, James Harrison, Mike Mathes, and Sebastian Tello-Trillo all helped to clarify and focus my thinking at various times. Their time and willingness to act as a sound board are much appreciated.

My wife, Charlotte, and my parents who have been a source of unwavering support in all possible ways. Without them, I would have had neither the inspiration to have started nor the perseverance to have completed this project.

I also thank the members of my committee. Professor Driskill provided perspective and advice in understanding the larger importance of my work. Professor Huffman taught me the tools of macroeconomics and sparked my interest all the way back in my first year. Professor Oh for sharing his experience and helping me prepare for the job market. And thanks to Professor Palacious for honest and actionable feedback, which was always delivered in a thoughtful manner.

Finally, I would like to thank my advisor, Kevin Huang. This work has benefited from many long and fruitful conversations with him. His guidance helped to stretch my abilities and to make me ask larger questions. I appreciate the balance he provided between guidance and allowing me to explore ideas on my own. This balance taught me to become an independent researcher.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

The recovery after the great recession has renewed interest in the functioning of labor markets in the US. For example, the 2016 Economic Report of the President mentions 'jobs' 100 times. The 2011 and 2006 reports mention 'jobs' only 86 and 17 times respectively. Major financial media outlets now hold monthly contests urging viewers to 'guess the number' prior to the release of BLS non-farm payrolls data. Further, labor income has become a widely discussed issue. Given the increased focus on employment, any model used for policy analysis needs to carefully consider the assumptions underlying the labor market.

All labor market models face tradeoffs between accurate prediction, plausible assumptions, and complexity. When applying these models for policy analysis, one needs to carefully consider the implications of the assumptions underlying the labor market. For example, a model that incorporates a textbook search model of unemployment will not generate sufficient fluctuations in unemployment. Some large scale DSGE models make predictions about hours worked, however they are unable to make predictions about unemployment. Few models consider the asymmetry of unemployment over the business cycle. In the US, unemployment rises further above trend during recessions than it falls below trend during expansions. Models producing symmetric data systematically under-predict the depth of recessions for shocks of a given size.

This dissertation fills a gap in the literature by emphasizing the importance of higher order moments in labor markets. The current literature largely ignores the importance of higher order moments in labor markets. I consider these moments in three ways. First, I use the observed higher order moments of the US data as an empirical target. Second, I use models with heterogeneous agents. In Chapter 2, endogenous variables vary by wealth.

Therefore the higher order moments of the wealth distribution are important in determining aggregate variables. In Chapter 3, identical agents may receive different wages. The endogenous distribution of wages determines the dynamics of the aggregate wage. Finally I use solution methods that preserve higher order moments of simulated data. Standard practice is to use log linearization, however this approximation loses any higher order characteristics of the model in question. In Chapter 2, I develop a nonlinear solution method that preserves the higher order moments of simulated data.

In Chapter 2 I consider the relationship between inequality and income. This relationship has been the subject of much research since Kuznets [1955] hypothesized that inequality first rises with income and then falls. This relationship was originally the result of migration between agricultural and industrial sectors of the economy; however this explanation has since been discredited. This chapter adds to the literature by developing a new mechanism linking inequality and income. The mechanism works via labor markets. The incentive for households to supply labor varies with the initial level of wealth. Therefore, the second moment of the wealth distribution becomes important for determining aggregate variables. The model implies that there is a non-monotone relationship between income and inequality. However, it is rotated relative to the Kuznets curve. Finally, I find empirical support by examining cross country data to evaluate this new mechanism relative to other solutions proposed in the literature.

I propose solutions to several issues regarding the predictions of frictional labor market models in Chapter 3. I develop a new wage setting mechanism that generates realistic moments of the aggregate wage. Secondly, I show that an alternate calibration of the model generates realistic volatility in the job finding rate as well as unemployment skewness. Finally, I link unemployment asymmetry to investment asymmetry though a household savings problem. This allows me to produce realistic second and third moments of both unemployment, the average wage, and investment. This realistic replication of the labor market is important because unemployment and the average wage are important components of aggregate household income. Therefore, producing realistic moments and co-movements of unemployment and the average wage is crucial to generating accurate dynamics in household income.

Finally I document a new empirical fact regarding recent changes in the cyclicality of the real wage in the U.S. I show that the wage has changed from procyclical to countercyclical in recent years. There are currently no papers that acknowledge or address this change. I hypothesize that this change may be due to rising wage inequality. If low wage jobs are more sensitive to the business cycle, then rising wage inequality strengthens compositional effects on the average wage. Using a structural VAR approach, I construct a counterfactual in which inequality is no longer rising and is not subject to further shocks. In this counterfactual, the average real wage remains procyclical. This indicates that rising wage inequality is able to account for a large portion of the change in real wage cyclicality.

Chapter 2

Inequality and Income When Information is Costly

Dating back to Kuznets' seminal 1955 work, the relationship between inequality and growth has been the subject of significant scrutiny. The original work hypothesized that inequality initially rises with per capita income, reaches a maximum, and then eventually begins to fall. Much empirical work has been devoted to examining the existence of such a relationship. The research on the Kuznets curve, and more broadly the relationship between income and inequality, often reaches very different conclusions. Some papers find evidence of a monotonic relationship between income and inequality, both positive and negative. Others support the existence of Kuznets curve by finding evidence of a non-monotone relationship. Still others find evidence that does not support the existence of such a relationship.

This chapter develops a new theoretical mechanism that links inequality and macroeconomic performance, and also presents supporting empirical evidence. The mechanism in this chapter arises from relaxing the standard assumption of full, costless information. This costly information regarding aggregate uncertainty results in households dividing their time between forming expectations about the future and supplying labor. This represents a significant departure from the usual framework. The standard economic framework for analyzing inter-temporal decision-making in the face of uncertainty assumes that agents fully understand the uncertainty. Agents are fully informed about the possible future realizations and the likelihood associated with these realizations. This is a difficult assumption to justify empirically. Under the rational expectations hypothesis, consumption should be a random walk as shown in Hall [1978]. However, subsequent empirical studies have shown that aggregate consumption data does not behave as predicted [Flavin, 1981, Campbell and Deaton, 1989]. Specifically, consumption is excessively sensitive to past information and surprisingly smooth with respect to changes in current income. These puzzles have motivated a literature in which agents are rationally inattentive or subject to informational frictions. This chapter develops a model that incorporates the idea of rational inattention into a two period, general equilibrium savings model. The result is that the wealth distribution and inequality become relevant in determining macroeconomic aggregates. Using empirical specifications from the literature, I then test to determine if the data is consistent with the model's prediction.

The chapter proceeds by examining several strands of related literature in Section 2. Section 3 lays out the theoretical model. Finally, Section 4 tests the implication of the proposed model.

2.1 Related Literature

2.1.1 Growth and Inequality

The original Kuznets hypothesis postulated that inequality first increases with development but eventually begins to fall. The hypothesis reasoned that migration from a low productivity, low inequality agricultural sector to a high productivity, high inequality industrial sector caused the inverted 'U' curve. Lacking reliable, wide-spread data on inequality, Kuznets instead examined the development experiences of the United States, the United Kingdom and Germany. Subsequently, inter-sector migration has largely been discredited as the driving force behind such a relationship [Anand and Kanbur, 1993b]; however, others have developed additional theories to generate a Kuznets curve. These mechanisms broadly fall into three categories, political economy, credit market imperfections, or demographics and fertility.

It is important to note that the hypothesis originally related inequality and the level of income. Much of the empirical work that followed examined the relationship between growth and inequality; however this chapter will focus on levels of income. The original hypothesis is an inter-temporal prediction that applies within a country. A lack of longitudinal data caused much empirical work to focus on cross sectional relationships. For comparability this chapter will employ both panel and cross-sectional methods. The existing literature that analyzes the relationship between inequality and macroeconomic performance is expansive, and frequently arrives at differing conclusions[1]. Even among papers that agree on the nature of the empirical relationship, there is disagreement about the underlying mechanism.

One set of models use a political economy mechanism to link growth and inequality. Political economy models that introduce a social tradeoff between redistributive policy and growth promoting policy conclude that inequality is harmful for growth [Persson and Tabellini, 1994]. Policy, specifically taxation on investment, is determined by majority rule. A larger distance between the median voter and the median person in the income distribution results in more support for redistributive policy which is financed through taxation on investment. This generates a predicted negative relationship between inequality and growth. The authors then determine if this prediction is supported in the data. Persson and Tabellini [1994] estimates the following pooled regression:

$$Growth_{it} = \alpha_0 + \alpha_1 IncSh + \alpha_2 NoFran + \alpha_3 School + \alpha_4 GDPGap + u_{it}$$

The first independent variable is a measure of income inequality, in this case the income share of the top 20% of the population. The second independent variable is a measure of political participation. Average level of schooling is included in the regression to control for human capital, and it is expected to have a positive sign. The final explanatory variable is a measure of relative level of development. It is included to control for possible convergence, another long-standing issue in the growth literature. A major concern in cross country growth regressions is potential endogeneity of inequality related co-variates. To

[1] See Benabou [1996] and Ehrhart [2009] for surveys

mitigate this issue, the authors use a beginning of period notion of inequality; that is, inequality is predetermined relative to the growth rate. Persson and Tabellini [1994] finds that income share is negatively and significantly related to subsequent growth. Voting is key to the theoretical mechanism; therefore, they also split the sample into democracies and non-democracies. The same regression on the two sub-samples reveals that the income distribution variable is significant in democracies but not significant in non-democracies. In subsequent empirical work, I control for an index of political rights for this reason.

Croix and Doepke [2003]also found evidence of a negative relationship between inequality and growth, but for different reasons. They claim inequality is harmful for growth through a fertility channel. The mechanism relies on differential fertility rates between the rich and the poor. They model fertility and education choices as endogenous. Parents face a trade off between the number of children and the amount of education they can afford. This choice affects the accumulation of human capital, thereby influencing the growth rate. To test this mechanism, differential fertility rate is added to to a standard growth regression. Differential fertility is significant and has the expected sign. Croix and Doepke [2003] also includes nonlinear effects to examine the Kuznets hypothesis. However, these terms are insignificant, indicating a monotonic relationship.

Other papers found evidence of a positive, monotonic relationship between inequality and growth [Forbes, 2000, Li and Zou, 1998, Albuquerque, 2004]. Forbes [2000] highlights and seeks to remedy several problems with the existing literature. First, regressions that show a negative relationship between inequality and growth are not robust to specification. Second, prior work could be influenced by measurement error and omitted variable bias. Forbes [2000] uses Arellano-Bond estimation with the more consistent data of Deininger and Squire [1996] to help reduce omitted variable bias and measurement error. Using a specification similar to Perotti [1996], the main result is to estimate the following equation:

$$
\begin{aligned}
Growth_{it} &= \alpha_1 Inequality_{i,t-1} + \alpha_2 Income_{i,t-1} + \alpha_3 MaleEducation_{i,t-1} + \\
&\quad \alpha_4 FemaleEducation_{i,t-1} + \alpha_5 InvestmentPrice_{i,t-1} + \beta_i + \eta_t + u_{it}
\end{aligned}
$$

The main term of interest is the coefficient on inequality. Income is included to control for possible convergence, and education is included to control for human capital. The final control is the price level of investment from the Penn World Tables. This measures the cost of investment across countries, which should reflect taxation, regulation, and other distortions. Arellano-Bond estimation suggests that an increase in inequality is associated with a subsequent increase in the growth rate. However, this result does not directly conflict with the existing literature, as it applies in the short run. Albuquerque [2004] develops a theoretical model that is consistent with the findings in Forbes [2000].

Li and Zou [1998] used a political economy mechanism to generate a positive relationship between inequality and growth. The paper's theoretical model relies on including consumption of public goods in the consumers' utility function. As in other political economy mechanisms, taxation is determined by majority rule. However, in contrast to prior work, government expenditure is used for consumption. To optimize, agents balance the marginal utility from private consumption with that of public consumption. As inequality changes, so does the median voter. The median voter's preferences determine the tax rate, and so inequality influences future growth. The paper then tests the theoretical mechanism by estimating panel data models with specifications similar to prior work that found evidence of a positive relationship [Alesina and Rodrik, 1994]. The main contribution of the paper is to use panel methods and the expanded income inequality data-set developed in Deininger and Squire [1996].

Yet other papers predict or find evidence supporting a non-monotone relationship consistent with the original hypothesis (Barro [1999, 2008], Banerjee and Duflo [2000], Acemoglu and Robinson [2002]). Williamson (1985) attributes the relationship to technological change. If technological change causes wages to rise more quickly than returns to capital, then technological change can generate the Kuznets curve. Aghion and Bolton [1997] relies on capital market imperfections to create a trickle down effect.

Acemoglu and Robinson [2002] develops a model of political economy related to that of Persson and Tabellini [1994]. In the model, a group of ruling elites determine policy initially. However, they are subject to the threat of revolution. As inequality grows so do the benefits to such a potential revolution. Extending enfranchisement is the only credible action available to the wealthy ruling class that can prevent a revolution. This extended suffrage implies the median voter becomes less wealthy, and therefore policy increasingly supports redistribution. This generates the inverted U relationship of the Kuznets curve. However, the eventual decrease in inequality associated with advanced development is not an unavoidable outcome. Rather, the political institutions in western countries drove the reduction. The authors identify two additional cases to consider. The first is the "autocratic disaster" which is characterized by high inequality and low output. The second case is the "East Asian miracle" which is the opposite. These cases arise from countries that begin with different political and social institutions; however in all three scenarios, the relationship between inequality and growth depends on political institutions. Therefore in this framework democracy is a necessary condition for the existence of a Kuznets curve. For this reason this chapter will attempt to control for political institutions in any empirical work.

Barro [2000] and Barro [2008] find that after controlling for several additional sources of variation, the Kuznets curve is an empirical regularity. However it has little power in explaining cross country variation in inequality. The regression model estimates the following specification using a random effects estimator. The set of controls is larger that in most other empirical work, which may help to support the random effects hypothesis. However no Hausman test is performed.

Banerjee and Duflo [2000] cautions against imposing unwarranted linearity on the data. In the absence of a guiding theory, assuming monotonicity or linearity is a strong assumption. The paper finds that growth shows an inverted U pattern relative to changes in inequality. Changes in inequality in either direction are associated with lower subsequent growth. However they find little evidence of a short run link between the level of inequality and growth. They perform non-parametric estimation of the relationship using the control variables from Barro [2000] and Perotti [1996]. They then test the assumption of linearity in the regression model. They include a quartic polynomial in the change in inequality to the Barro control variables. A test for joint significance of the nonlinear terms rejects a linear relationship. Similarly, including a piece-wise term reveals the same inverted U shape.

While some papers confirm the Kuznets hypothesis, many also reject it [Deininger and Squire, 1998, Anand and Kanbur, 1993a]. [Deininger and Squire, 1998] rejects evidence of the Kuznets hypothesis and also is not supportive of the political economy mechanism developed in Persson and Tabellini [1994]. The authors use a data-set of higher quality assembled in prior work [Deininger and Squire, 1996]. The data-set contains information about the method and coverage used in collecting the inequality data, and has become the basis for nearly all subsequent work. The authors impose a quality standard to filter the observations. At a minimum, the data should be based on nationally representative surveys and should cover all sources of income or expenditure. Additionally, the authors have collected data on both income inequality and asset inequality as proxied by land holdings. Using this data-set the authors examine two possible relationships between inequality and growth.

First they focus on the effect of initial inequality on subsequent growth. This closely follows the prior literature. They estimate the following traditional regression:

$$
\begin{aligned}
Growth_{it} &= \alpha_0 + \alpha_1 InitGDP_{it} + \alpha_2 InitGini_{it} + \alpha_3 Investment_{it} + \\
&\quad \alpha_4 BlackMarketPremium_{it} + \alpha_5 Educ_{it} + u_{it}
\end{aligned}
$$

Using income inequality data from the higher quality data-set finds that higher initial inequality is associated with lower future growth. However this is not robust to specification, and the relationship loses significance once regional dummies are included in the regression. Interestingly, the use of asset inequality, as proxied by land holdings, is also significant and robust to the inclusion of regional dummies. Also this work questions the political economy mechanism because there does not exist a significant relationship between inequality and growth for democracies. Instead, they find support for a credit channel mechanism. An implication of the credit channel mechanism is that initial inequality affects developing and developed countries differently. Deininger and Squire [1998] finds that initial inequality is important for subsequent growth in a sample of developing countries; however it is not significant in the sample of OECD countries.

Next, Deininger and Squire [1998] shifts from examining the relationship between growth and inequality to considering the existence of the Kuznets curve. Due to the availability of data on inequality, prior attempts to identify the Kuznets curve use cross sectional data to draw conclusions about an inter-temporal relationship. Countries at different levels of development mimic the development process in a single country. The comprehensive data compiled in their earlier work allows The paper examine the existence of any contemporaneous relationship between levels of income and inequality. This is highly consistent with the original Kuznets hypothesis in that they are not dealing with first differences in income or in inequality.

$$
Gini_{it} = A_i + B_i(Y_{it}) + C_i(1/Y_{it}) + DS + error
$$

Here S is a dummy variable indicating socialist countries. Allowing for country specific intercepts and slopes allows the authors to test for a 'universal' Kuznets curve versus a country specific one. Once this is estimated in decadal differences, the data does not support the existence of a cross country Kuznets curve. Stronger still, they also reject the possibility of an intra-country Kuznets curve. That is a Kuznets curve with country specific slope parameters.

There are many varieties of theoretical basis for a relationship between inequality and growth. These can largely be grouped into political economy, credit channel, or fertility. This chapter adds a new mechanism to this existing group. Any new theoretical relationship between inequality and growth needs to be evaluated in relation to these existing theories. Therefore I will draw on them heavily in my empirical work in order to test the relative validity of my mechanism. Also it is evident from examining the literature that choice of estimator greatly influences the results, to the point of contradictory conclusions. Early cross section work is highly supportive of the Kuznets hypothesis, while fixed effects estimation and dynamic panel methods frequently reject it. However, random effects estimation in Barro [2000] and Barro [2008] offers support for the Kuznets curve.

2.1.2 Rational Inattention Literature

The mechanism developed in this paper relies on costly information to link inequality and income. Therefore, it is also related to the rational inattention literature. This literature is motivated by empirical puzzles regarding consumption. The excess sensitivity and smoothness puzzles make it natural to hypothesize that perhaps consumers are unable or unwilling to use all relevant information in forming expectations. Reis [2006] and Sims [2006] take this as a starting point and then propose models of consumption behavior under different informational frictions.

Reis [2006] models inattentive consumers as choosing to update expectations infrequently rather than instantaneously and continuously to shocks. Agents must pay a fixed cost each time they wish to update information, and thereby re-optimize the consumption decision. Sims [2006] models consumers as being constrained in the amount of information that they are able to consume and process. The motivation for modeling agents as having finite capacity for attention is that ‘information that is freely available to an individual may not be used, because of the individuals limited information processing capacity’ (Sims [2006], p. 160). He then applies the idea of Shannon capacity for measuring information flow in the context of agent decision making. Agents optimize over what information to use in the formation of expectations. This finite capacity for attention results in agents exhibiting inertia in decision making.

I propose a method that has some similarities to both the Reis and Sims approaches. I assume information to be costly for agents to process and utilize in forming expectations. This is similar to the Reis approach. However, the cost of acquiring information is not a fixed cost, but rather a variable cost that depends on the amount of information consumed. Also, rather than choosing the frequency of updating expectations, in my model agents optimize over the consumption of information. This is more similar to the Sims approach. However in my model inattention will arise from the cost of consuming information rather than a finite capacity for attention. This allows me to use the traditional tools of maximization to the formation of expectations. The main contribution of my approach is to endogenize the effort put into forming expectations. Rather than the discrete choice in Reis [2006], agents will choose from a continuum. As opposed to the information constraint in Sims [2006], agents in my model may consume all information but they may choose not to do so.

2.1.3 Financial Literacy Literature

If time spent forming expectations is interpreted as financial literacy then this paper is related to the literature linking financial literacy, planning, and wealth accumulation. With respect to the trade off between productive activity and acquiring information, my paper is most closely related to Lusardi et al. [2013]. In this model, agents acquire financial knowledge in order to gain access to a risky savings technology. The expected return of this savings device is increasing in the amount of financial knowledge acquired. My framework differs in that it is a general equilibrium approach. This general equilibrium aspect is necessary in my framework to make predictions regarding inequality and output.

Lusardi and Mitchell [2007] shows there are large portions of the US population that lack even basic financial literacy. This lack of literacy is associated with lack of a financial plan for the future and also reduced wealth accumulation. This correlation will be consistent with my model, in that initially wealthy agents will choose to spend more time acquiring information. However more recent work in Lusardi and Mitchell [2011] uses specific questions in the Health and Retirement Survey to address the direction of causality in this relationship. Using an instrumental variable approach, they conclude that financial literacy causes more wealth accumulation. However, they do not address they question of why financial illiteracy exists. If literacy causes a significant increase in long term wealth, then why do people not acquire these skills? In my model, this will be the result of optimizing behavior rather than an innate inability to perform the type of necessary calculations.

2.2 A Simple Model Economy

The theoretical model is presented in the simplest form possible in order to facilitate understanding and clarity. The model can be extended in various ways, however even in the simplest form it does not admit a closed form solution. Time is finite, and there are 2 periods. There are two types of agents, firms and households.

Firms are identical and have access to a constant returns to scale, Cobb-Douglas technology to produce a consumption good. The production technology is subject to an aggregate TFP shock. The productivity shock can take on two possible values, $z_g$or z_b where $z_g > z_b$. Conditional on the current state, transitions are governed by an exogenous 2 by 2 transition matrix. The realization of the shock is perfectly observed by firms prior to hiring and rental decisions. Upon completion of production, firms pay the wage and capital bills. Firms act as price takers in the labor and capital markets.

There exists a unit mass of households. The households have identical logarithmic preferences over consumption, but they differ in the initial endowment of wealth. Household wealth, denoted k, indexes the households. Let G(k) denote the cumulative distribution of initial wealth with associated density function g(k). The distribution has support $\mathcal{K}$ and and mean $\bar{K}$. Given the unit mass of agents, the mean also coincides with the aggregate amount of capital in the economy. Each household is also endowed with one unit of time per period. Households may spend their time working for a firm in exchange for labor income, or they may use a portion of their time to learn about the transition probabilities governing the TFP shock. The manner in which this learning occurs will be made specific later. These probabilities are used by the household in forming expectations about the future. Households do not value leisure, and so the entire time endowment is allocated between these activities.

Households engage in several activities. They lend capital and supply labor to the firms. Therefore they must make two decisions. They must determine how much labor to supply, which in turn determines their labor earnings. Also, they must determine how much of their total income to save for consumption in the second period. When making this decision, the agents' objective is to maximize ex-ante expected utility.

2.2.1 Sequence of Events

Understanding the timing of the model is important to forming the agents' maximization problems. When the consumers are 'born', they perfectly observe the state of the world. The state variables at this point include the distribution of capital and an initial realization of the productivity shock. In contrast to standard models, they do not know the transition probabilities without expending time to learn about them.

Households receive a wage for the time spent supplying labor. If the agent chooses to spend time forming expectations, they forgo labor income. Agents do not value leisure, so the entirety of the time endowment will be spend on working or collecting information for expectation formation. After observing the initial state, agents must decide how to allocate the time endowment. It is important to note that the agents make the labor supply choice prior to the realization of the random shock. This arises directly from the necessity of forming expectations prior to the shock realization. Since in this framework forming expectations requires time, and time is in limited supply, agents are forced into implicitly making the labor supply choice prior to the realization of the technological shock.

After the labor supply decision and expectations are formed, the productivity shock is realized. Consumers then choose how much of their income to devote to consumption and how much to save for the second period. The savings decision will be optimal given the income, as the uncertainty has been resolved prior to the savings decision. In the second period, there is no production and agents consume their savings. This could easily be relaxed, as in the final period there would be no incentive to form accurate expectations about the future. Thus the agents would spend their full time endowment working. Firms hire labor and capital after the realization of the productivity shock. Thus they are not subjected to the same ex-ante uncertainty as agents.

2.2.2 Household Problem

The objective of each household is to maximize its ex-ante expected utility. The transition probabilities the household assigns to the potential future states depend on the amount of time spent forming expectations. Let $f(\tau) : [0,1] \to [0,1]$ be an increasing function that reflects the usefulness of information utilized in forming expectations given τ. I will term this function the learning function, and I will use it to introduce a wedge between the true probability governing the economy and agent's perception of this probability. The agent's perception of the probability that the future state will match the currently observed state is given by $f(\tau) * \pi_s$, where $\pi_s$is the true probability governing the shock process. All agents are capable of attaining the true probabilities; however this would require them to spend the entire time endowment on forming expectations. It is most natural to think of this function as having the same properties as a typical production function. However, the only restrictions are that the function be monotonically increasing, $f(0) = 0$, and $f(1) = 1$.

Each type household of type i solves the following two stage maximization problem:

$$\max_{\tau^i} E\left[\max_{k_2^i}\left[log(c_1^i) + \beta log(c_2^i)\right] | f(\tau^i)\right]$$

$$\text{s.t.} c_1^i + k_2^i \le k_1^i \times R + (1 - \tau^i) \times W$$

$$c_2^i \le k_2^i$$

$$\tau^i + l^i \le 1$$

The conditional expectation notation is meant to suggest that the agent's assessment of the transition probabilities is influenced by the amount of time spent on collecting and processing information. The nested structure of the maximization problems arises from the sequence of events. To solve the model, I proceed backwards chronologically. I begin with the household savings problem, the firms' profit maximization problem, and finally the consumers labor supply choice.

2.2.3 Household Savings Problem

To begin solving the model, consider the inner maximization problem and the two associated budget constraints. Taking wages and capital rental rates as fixed, agents solve the inter-temporal utility maximization problem subject to the budget constraint. The solution to the problem is standard, and it yields the familiar value function:

$$V^i(k_1, z_1, l^i) = (1+\beta)log(Rk_1^i + W(1-\tau^i)) + \beta log\beta - (1+\beta)log(1+\beta))$$

Note that the time spent forming expectations by households of type i, τ^i, is included in the state space. It is predetermined relative to the inner maximization problem.

2.2.4 Firm Profit Maximization Problem

Firms have access to constant returns to scale, Cobb-Douglas technology. The sole purpose of firms in this setting is to convert capital and labor into the consumable good. Using capital letters to denote aggregate capital and labor supplies, firms operate according to the familiar profit maximization conditions:

$$R = \alpha z_1 \left(\frac{\bar{L}}{\bar{K}}\right)^{1-\alpha}$$

$$W = (1-\alpha) z_1 \left(\frac{\bar{K}}{\bar{L}}\right)^{\alpha}$$

Here $\bar{K}$ and $\bar{L}$ are used to indicate the aggregate capital and labor supplies respectively. Given the continuum of agents, the aggregate capital and labor supply are not influenced by the decision of any individual agent. Denoting the support of the capital distribution with $\mathcal{K}$ and integrating gives the aggregate capital and labor supply:

$$\bar{K} = \int_{\mathcal{K}} k\, g(k)\, dk$$

$$\bar{L} = \int_{\mathcal{K}} l\, g(k)\, dk$$

2.2.5 Labor Supply Decision

Given the value function that resulted from the savings decision, that is the inner maximization problem, it is possible to reformulate the outer maximization problem. The objective becomes a weighted combination of the savings problem value function in the two possible future states. The weights are the agents assessment of the transition probabilities, which depend on the amount of time spent forming expectations.

To complete the solution characterization, consider the outer maximization problem faced by households. Given the solution to the prior two portions of the model, agents of type i solve the following maximization problem:

$$\begin{aligned} \max_{l^i} E[V^i(z_s,k_1)|f(1-l^i)] = \quad & (1+\beta)(1-l^i)\pi_s log\left(\tfrac{z_s}{z_j}\right) + (1+\beta)log\left(Wl^i + Rk^i\right) \quad + \\ & (1+\beta)log(z_j) + \beta log\beta - (1+\beta)log(1+\beta) \end{aligned}$$

$$\text{s.t.} 0 \leq l^i \leq 1$$

The objective makes the tradeoff faced by households explicit. Increasing labor supply increases income from labor, but at the expense of facing greater uncertainty about the potential states of the economy. The first term in the expression captures this reduction in uncertainty, while the second term reflects the utility from income. To begin, consider the first order necessary condition for an interior optimum:

$$\frac{\partial E[V^i(k_1^i)|f(1-l^i)]}{\partial l^i} = \frac{(1+\beta)W}{Wl^i + Rk_1^i} - (1+\beta)\pi_s log\left(\frac{z_s}{z_j}\right) = 0$$

The first term is the marginal income gained from supplying labor in utility terms. The second term is the cost in terms of expected utility from spending time in labor instead of forming expectations; so it is the utility cost of increased uncertainty. Substituting the firm's profit maximization conditions and rearranging yields the optimal labor supply for agents of type i if the time endowment constraints do not bind:

$$l^*(k^i) = \frac{1}{\pi_s log\left(\frac{z_s}{z_j}\right)} - \left(\frac{\alpha}{1-\alpha}\right)\left(\frac{\bar{L}}{\bar{K}}\right)k^i$$

The optimal labor supply therefore is decreasing and linear in the agents' level of capital while the solution is in the interior of the constraint space. The marginal utility received from working an additional unit of time is different between households with differing levels of wealth. Therefore the initial wealth distribution will become relevant for determining macroeconomic aggregates. The monotonicity of the necessary first order condition for an interior optimum implies that there are two important threshold values in the wealth distribution. Below the lower threshold value the time constraint always binds upwards, while above the upper threshold value the time constraint always binds downwards. Therefore the full characterization of the the optimal labor supply, including the potentially binding time constraints, is a piecewise function:

$$l^*(k^i) = \begin{cases} 1 & : k^i \leq \hat{k} \\ \frac{1}{\pi_s log\left(\frac{z_s}{z_j}\right)} - \left(\frac{\alpha}{1-\alpha}\right)\left(\frac{\bar{L}}{\bar{K}}\right)k^i & : k^i \in (\hat{k}, \tilde{k}) \\ 0 & : k^i \geq \tilde{k} \end{cases}$$

Where $\tilde{k}$ and $\hat{k}$ denote the threshold levels of capital at which the two time endowment constraints become binding. These thresholds must be characterized to complete the optimal labor supply function. I begin by solving for the aggregate labor supply $\bar{L}$. Integrating the piecewise optimal labor supply function over the distribution of capital yields the following expression for the aggregate labor supply:

$$\bar{L} = \int\limits_{\underline{k}}^{\hat{k}} g(k^i)\, dk^i + \int\limits_{\hat{k}}^{\tilde{k}} \left(\frac{1}{\pi_s log\left(\frac{z_s}{z_j}\right)} - \left(\frac{\alpha}{1-\alpha}\right)\left(\frac{\bar{L}}{\bar{K}}\right) k^i \right) g(k^i)\, dk^i$$

The first term's lower bound of integration is the lower bound of the support of the capital distribution, denoted $\underline{k}$. The upper bound is the threshold level of capital at which agents no longer spend the entire time endowment supplying labor. The second term is the aggregate labor supplied by agents whose optimal labor supply is in the interior of the time constraint space. It is important to note that aggregate labor and capital are constant with respect to the variable of integration. Evaluating the first integral and splitting the second integral yields:

$$\bar{L} = G(\hat{k}) + \frac{1}{\pi_s log\left(\frac{z_s}{z_j}\right)} \int\limits_{\hat{k}}^{\tilde{k}} g(k^i) dk^i + \left(\frac{\alpha}{1-\alpha}\right)\left(\frac{\bar{L}}{\bar{K}}\right) \int\limits_{\underline{k}}^{\tilde{k}} k^i g(k^i)\, dk^i$$

Again, evaluating the first remaining integral and then solving yields the aggregate labor supply in terms of the two cutoff values in the capital distribution.

$$\bar{L} = \frac{G(\hat{k})\left(1 - \frac{1}{\pi_s log\left(\frac{z_s}{z_j}\right)}\right) + \frac{G(\tilde{k})}{\pi_s log\left(\frac{z_s}{z_j}\right)}}{1 + \left(\frac{\alpha}{1-\alpha}\right)\left(\frac{\int_{\hat{k}}^{\tilde{k}} k^i g(k^i)\, dk^i}{\bar{K}}\right)}$$

The final term in the denominator is of particular interest and warrants further consideration. It represents the proportion of capital held by households that are not at a corner solution for the labor supply choice. Therefore it can be related to the Lorenz curve associated with the capital distribution. To see this, begin by defining the Lorenz curve in the usual way:

$$\mathcal{L}(G(k)) \equiv \frac{\int_{\underline{k}}^{k} k^i g(k^i)\,dk^i}{\bar{K}}$$

The term of interest can be re-expressed as the difference between two integrals, which in turn is the difference of two values of the Lorenz curve:

$$\frac{\int_{\hat{k}}^{\tilde{k}} k^i g(k^i)\,dk^i}{\bar{K}} = \frac{\int_{\underline{k}}^{\tilde{k}} k^i g(k^i)\,dk^i - \int_{\underline{k}}^{\hat{k}} k^i g(k^i)\,dk^i}{\bar{K}} = \mathcal{L}(G(\tilde{k})) - \mathcal{L}(G(\hat{k}))$$

These are the values of the Lorenz curve associated with the two points in the capital distribution at which the time endowment constraints become binding. Therefore this term represents the amount of capital held by the households that lie withing the two threshold values of capital. The level of inequality clearly is relevant in determining the labor supply, which in turns determines aggregate output. Using the expression for aggregate labor, the optimal choice of individual labor supply becomes:

$$l^*(k^i) = \begin{cases} 1 & : k^i \geq \tilde{k} \\ \frac{1}{\pi_s log\left(\frac{z_s}{z_j}\right)}\left(1 - \frac{\left[\pi_s log\left(\frac{z_s}{z_j}\right)G(\hat{k}) + G(\tilde{k}) - G(\hat{k})\right]\alpha k^i}{(1-\alpha)\bar{K}\left[1+\frac{\alpha}{1-\alpha}[\mathcal{L}(G(\tilde{k}))-\mathcal{L}(G(\hat{k}))]\right]}\right) & : k^i \in [\hat{k}, \tilde{k}] \\ 0 & : k^i \leq \hat{k} \end{cases}$$

Next I examine the existence and identification of the threshold values in the capital distribution.

The lower capital threshold, at which households supply a full unit of labor, is defined by:

$$l^*(\hat{k}) = 1 \Longleftrightarrow \pi_s log\left(\frac{z_s}{z_j}\right) = 1 - \frac{\left[\pi_s log\left(\frac{z_s}{z_j}\right)G(\hat{k}) + G(\tilde{k}) - G(\hat{k})\right]\alpha\hat{k}}{(1-\alpha)\bar{K}\left[1+\frac{\alpha}{1-\alpha}[\mathcal{L}(G(\tilde{k}))-\mathcal{L}(G(\hat{k}))]\right]}$$

Similarly the upper capital threshold is defined by:

$$l^*(\tilde{k}) = 0 \Longleftrightarrow 1 = \frac{\alpha}{(1-\alpha)\bar{K}} \frac{\left[\pi_s log\left(\frac{z_s}{z_j}\right) G(\hat{k}) + G(\tilde{k}) - G(\hat{k})\right] \tilde{k}}{\left[1 + \frac{\alpha}{1-\alpha}[\mathcal{L}(G(\tilde{k})) - \mathcal{L}(G(\hat{k}))]\right]}$$

Begin by assuming that the distribution of capital follows a Pareto distribution. This implies that both the CDF and the Lorenz curve are continuous functions of capital. This in turn implies that the left hand side is a continuous function of capital as well. At the lower bound of the capital distribution support, the left hand side of the above equation is zero. For a fixed distribution G, the left hand side becomes arbitrarily large as k goes to infinity. Since the support of the Pareto distribution has no upper bound, it is always possible to choose some value of capital sufficiently large such that the intermediate value theorem can be applied. Thus in the case of a Pareto distribution, we can conclude there is always a value of $\tilde{k}$ such that the above holds.

After specifying a functional form for the wealth distribution I am unable to find a closed form solution to this system of equations. Instead, I must rely on numeric methods to arrive at a solution.

2.2.6 A Simple Calibration and Testable Implications

Despite the difficulties of obtaining a closed form analytic solution, it is possible to obtain numerical solutions. I perform comparative static exercises to determine how the aggregates in a sample economy respond to changes in inequality. Following the work of Blaum [2012], I assume that the wealth distribution in the US follows a Pareto distribution parametrized by a minimum wealth level of $10,000 and a shape parameter of 1.7. This implies an average wealth of $24,250. I take this as a starting point and perform a series of mean preserving spreads to the capital distribution. To perform the mean preserving spread, I form a vector of potential shape parameters to include the target of 1.7. I then calculate the minimum wealth level associated with each value of the shape parameter and a mean of \$24,250.

Parameter	Assumed Value/Range
α	0.35
π_s	0.8
z_s/z_j	1.66
θ	[1.05, 5]

Table 2.1: Calibrated Values

The Pareto distribution has several important advantages. First, as shown in Blaum [2012], it fits the tail of the wealth distribution well. Second, the Pareto distribution simplifies the relationship between the Lorenz curve and the Gini coefficient. For a Pareto distribution with shape parameter θ, the Gini coefficient is simply $\frac{1}{2\theta-1}$. Values for the other parameters are listed in Table 2.1. As a robustness check I have examined other values of the uncertainty structure, and the results remain directionally unchanged.

To solve the model, I use a grid search method over the support of the capital distribution. For each of the sequence of Pareto distributions, I evaluate the CDF and the Lorenz function at each point on the grid. Finally, I evaluate the expression that defines the two thresholds of capital for each point on the grid. I am then able to select the appropriate levels of capital that solve the equations defining the extensive margins at a given level of precision.

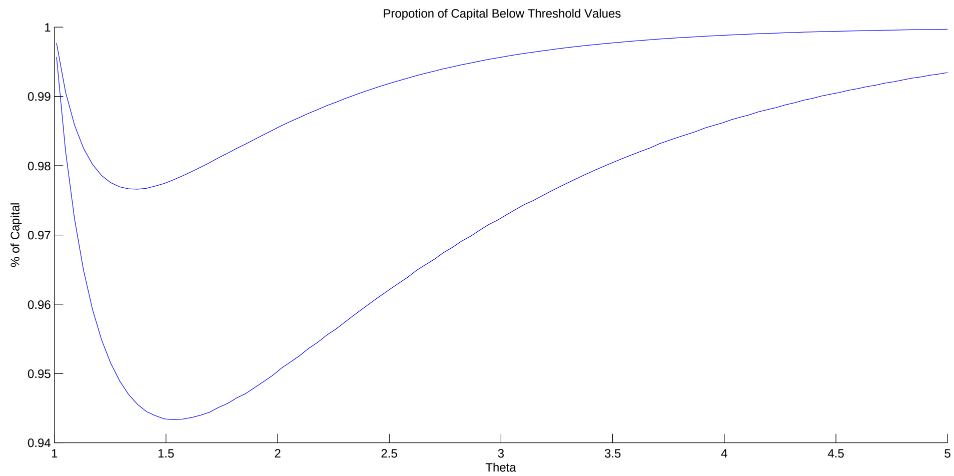

Figure 2.1: Percent of Capital Below Cutoff Value of Capital

Figure 2.1 shows the proportion of capital held by households below each of the extensive margins. The two extensive margins are nonlinear functions of the shape parameter of the wealth distribution. Therefore, as the wealth distribution moves toward equality, the percent of capital held by agents supplying positive labor also increases.

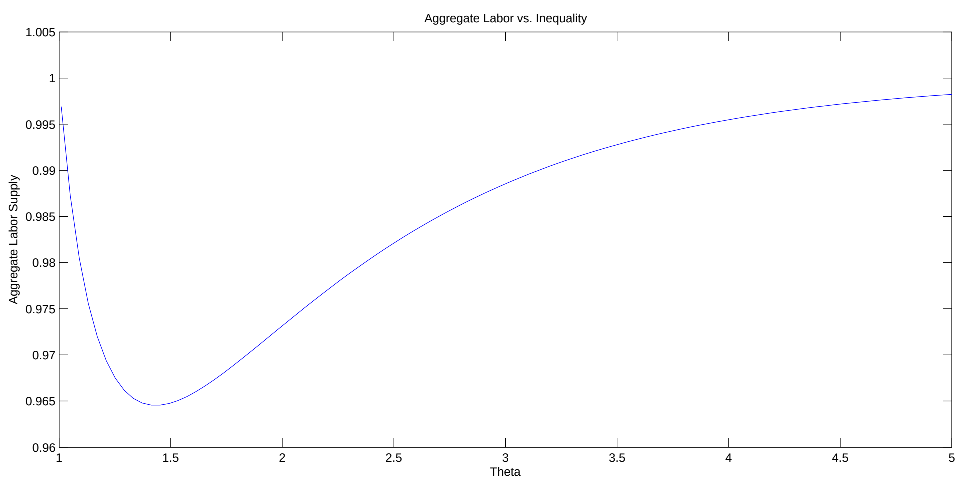

Figure 2.2: Aggregate Labor Supply vs. Shape Parameter

Figure 2.2 shows that the relationship between aggregate labor and inequality is non-monotone. Given the specification of technology and holding all other inputs fixed, the relationship between output and inequality will mirror the relationship between labor and inequality. Therefore, this non-monotone relationship between inequality and labor is closely related to the Kuznets hypothesis. However the relationship in this chapter is in a different space than the original Kuznets hypothesis. Kuznets predicts an inverted U shape in the income/inequality space. The theoretical model predicts a U shape in the inequality/output space. Whereas the original hypothesis predicts an inverted U, my model predicts this is reflected around the 45 degree line and therefore would appear to be a C shape.

2.3 Empirical Evidence

The main implication of the theoretical model is that aggregate hours worked is related to inequality in a non linear way. To determine if this is consistent with the data, I estimate several models that are closely related to traditional empirical specifications in the literature on growth and inequality. Directly testing this is complicated by the lack of a closed form solution for the cutoff levels of capital and the corresponding complexity of functional form for the aggregate labor supply. Therefore to test the implication of the model, I add a square term reflecting the non-linearity to cross country growth regressions.

$$Hours_{it} = \alpha_0 + \alpha_1 Theta_{it-1} + \alpha_2 ThetaSq_{it-1} + \alpha_3 Controls_{it-1} + u_{it}$$

In selecting the set of control variables, I use Barro [2000] and Perotti [1996] for guidance. One advantage is that both of these specifications are used in subsequent papers [Banerjee and Duflo, 2000, Forbes, 2000]. Additionally, the Barro [2000] set of controls is the most extensive in the literature, while [Perotti, 1996] is the most concise. The Perotti specification controls for per capital GDP, its square, the price level of investment, and average years of secondary education. In addition to these, the Barro specification controls for government expenditure as a percent of GDP, investment as a percent of GDP, the total fertility rate, an index of political rights, and terms of trade. This chapter tries to use the same data sources whenever possible, while still updating to include more recent data points.

To help reduce potential endogeniety problems, all co-variates are lagged one period so as to be predetermined. The majority of the literature utilizes averages over distinct five year periods for the dependent variable. This is regressed onto the set of controls from the most recent year prior to the relevant five year period. I deviate slightly here in that I use yearly observations for the dependent variable. I also require that all independent variables come from the same year. This reduces sample size due to intermittent data coverage, but it is more rigorous.

Most papers that examine growth and inequality are most tightly constrained by inequality data, and also most papers in the literature are concerned with growth and inequality. However, the simplest and most direct implication of the theoretical model uses labor supply as the dependent variable. However when testing the direct implications of my model, the hours worked data has lower coverage than inequality. For this reason, future work could use output as a dependent variable while adding capital stock, labor share, and productivity to the set of controls.

Much of the literature has come to focus on the relationship between growth and rather than income and inequality. The limited dynamics of my theoretical model will constrain me to testing hypotheses regarding income instead of growth. However this is consistent with the original Kuznets hypothesis.

2.3.1 Data

The dependent variable is the average annual hours worked per worker. This measures the intensity of employment, rather than an extensive margin. Therefore I do not control for country level demographics, such as dependency ratios.[2] This data is collected by the

[2]Initial work indicates that dependency ratios were not significant in either the Barro or Perotti specification

OECD.The data-set is unbalanced with coverage from 1974 to 2012 for all OECD member countries. However from 2000 onward, the data-set contains observations for each country. Data on inequality is typically the most sparse in the literature. However, the data on the average number of hours worked is the most binding constraint in my data-set. The reliance on OECD data implies limited ability to speak to the experience of developing countries.

Gini coefficients come from the June 2014 release of the World Income Inequality Database. To make the data directly comparable to the theoretical model, I transform the Gini coefficient to the Pareto shape parameter. This database builds on the work of Deininger and Squire [1996] to collect measures of inequality with high coverage across countries and time. The cost of this coverage however is comparability and complexity. As documented in Atkinson and Brandolini [2000], there is great variation across countries and time in the population coverage, survey method, and definition of income. For many country year pairs there are multiple observations based on different methodologies. Simply averaging across multiple observations is inappropriate and can potentially add unintended bias.[3] Incompatibility or observations across countries or within country and across time can result in spurious results. Therefore it is important to have a sensible, transparent method for selecting amongst multiple observations. The data-set provides the means to do this. It includes a details on survey method, source, and population coverage. Additionally, the data-set contains information on the notion of income and the unit of measure, for example household or individual.

I consider two selection methods as a robustness check. In the pooled regressions, my greatest concern is introducing the appearance of a nonlinear relationship by selecting differing measures of inequality. For example Deininger and Squire [1996] finds an average difference of 6.6 percentage points between income based and expenditure based Gini coefficients. If the type of measurement is correlated with other regressors, then there could be a spurious non-linearity[4]. Therefore, the first method of selecting Gini coefficient

[3]Incidentally, averaging across multiple sources increases the significance of the following results.

[4]I do not perform this adjustment for two reasons. First, Atkinson and Brandolini [2000] gives reason

seeks to maximize the comparability of the data across countries. To increase comparability of the data, I prioritize amongst multiple observations according to criteria laid out by Kuznets. First, the data should measure inequality amongst households rather than individuals. Second, measurement should be representative of the national population, while possibly noting households including young adults and retirees. Finally, the measure of income should reflect "income received by individuals, including income in kind, before and after direct taxes, excluding capital gains." While complete adherence to these properties would be a "statistical economist's pipe dream" (Kuznets 1955), I will follow these principles whenever possible in selecting data among multiple observations. This ensures high comparability across countries, but means there may be excess variation within country.

The primary method is to select all observations that come from a single source for each country. Switching between sources or survey methodology within a country can cause unintended breaks in the temporal dimension. This results in significant problems with fixed effects estimation. When multiple sources are available, I select the source that is of the highest quality according the above criteria. The use of a single source reduces the number of observations available, but greatly increases comparability through time. Additionally, I utilize the same source across countries whenever possible.

For the additional controls, I try to match the data sources used by the existing literature as closely as possible. I draw GDP per capita, the price level of investment, government expenditure, and investment expenditure from the Penn World Tables 8.0. As standard in the literature, I use the Barro and Lee education data-set to control for human capital. Specifically, I use the average years of secondary education. This data is collected every five years, and so inclusion of it will greatly reduce sample size. Therefore, I show all regressions with and without this variable included. As in Barro [1999], the subjective index of political rights is taken from Freedom House, an independent organization. Total fertility and terms of trade are from the World Bank.

to doubt the validity of the adjustment. Secondly, after prioritizing according to my criteria, only 3 Gini measurements remain that are based on consumption or income.

2.3.2 Pooled Regressions

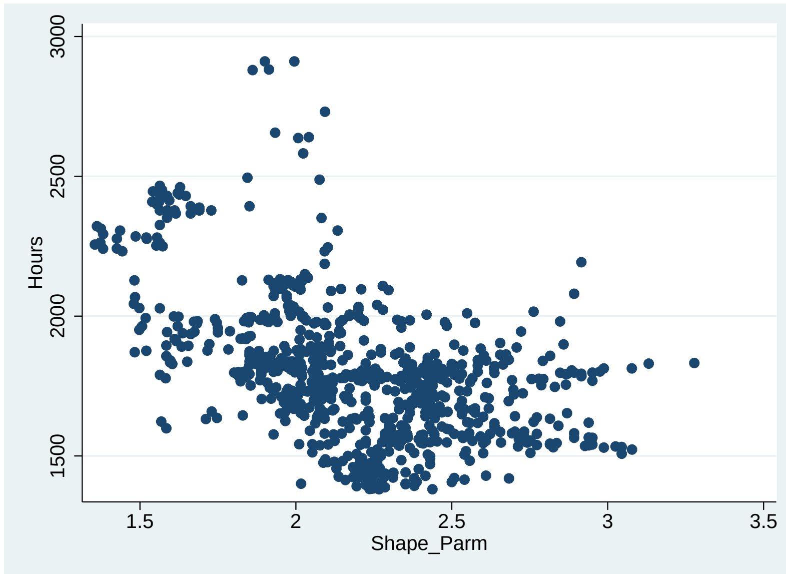

Figure 2.3: Average Annual Hours Worked vs Implied Shape Parameter

Figure 2.3 and an initial regression of hours worked onto the shape parameter and its square suggests the data appears consistent with the testable hypothesis of the model[5]. However of this may be misleading, as there is great variation across countries on other dimensions.

The results of the theoretical model are essentially a counterfactual. It examines subsequent changes in aggregate labor supply and output under the conditions that only the initial wealth distribution is changed. This insight will inform the method of estimation. I begin with cross sectional regressions as it is consistent with the early literature examining the Kuznets curve and my theoretical model. A more complex theoretical model with full dynamics would require dynamic panel data estimation. The pooled regression uses natural variation in inequality, while controlling for other factors, to mimic the experience of a single entity.

[5]Results not shown.

	Specification			
	Perotti (1996)	Perotti and Education	Barro (2000)	Barro and Education
Pareto Shape Parameter	-2615.6***	-2746.2***	-1819.5***	-2149.7***
	(0.000)	(0.000)	(0.000)	(0.001)
Pareto Shape Parameter Squared	514.0***	527.7***	322.7***	386.1***
	(0.000)	(0.000)	(0.003)	(0.009)
N	470	113	269	73
Adj R Squared	0.51	0.53	0.71	0.70

Table 2.2: Pooled Regressions

Table 2.2 contains pooled OLS estimates on the entire sample for which data is available[6]. The first two columns reflect the Perotti specification, while the second two columns are based on the Barro specification. The original specification in Persson and Tabellini [1994] uses education attainment by gender, however I have used overall education attainment to increase comparability with the Barro specification. The first and third columns do not contain the average years of secondary education as a co-variate to examine the effect of sample size reduction. The nonlinear inequality term is highly significant in both specifications, with and without the education data[7]. All standard errors reported are clustered by country and therefore robust to heteroscedasticity and auto-correlation. The non linear terms remain positive and highly significant, suggesting a nonlinear relationship consistent with the model's prediction. However there may be country specific variation that is not being accounted for correctly in pooled regression. Therefore I also perform fixed effects estimations.

[6]The appendix contains the estimates for all variables

[7]The appendix contains the same regressions performed on the sub-sample for which all variables are available

2.3.3 Fixed Effects Regressions

	Specification			
	Perotti (1996)	Perotti and Education	Barro (2000)	Barro and Education
Pareto Shape Parameter	-427.8	-246.9	-838.7**	669.1
	(0.12)	(0.39)	(0.02)	(0.22)
Pareto Shape Parameter Squared	95.2*	54.7	191.2**	-133.3
	(0.07)	(0.34)	(0.01)	(0.24)
N	470	113	269	73
Adj R Squared	0.26	0.37	0.38	0.69

Table 2.3: Country Fixed Effects Regressions

Using fixed effects to to estimate the same equation provides a more rigorous test of the model's prediction. Table 2.3 contains baseline fixed effects estimates. The nonlinear term is significant and of the expected sign in both the Barro and Perotti specifications without the education data. However once education is controlled for, the term loses significance and in fact switches signs in the Barro specification. This large difference in the estimate could be due to systematic differences between countries with and without education data, or it could be attributed to the inclusion of the additional co-variate. However, the average years of secondary education is not significant in the Barro specification with fixed effects.

	Specification			
	Perotti (1996)	Perotti and Education	Barro (2000)	Barro and Education
Pareto Shape Parameter	-278.2	-246.9	675.3	669.1
	(0.39)	(0.39)	(0.25)	(0.22)
Pareto Shape Parameter Squared	64.1	54.7	-134.0	-133.3
	(0.3)	(0.34)	(0.25)	(0.24)
N	113	113	73	73
Adj R Squared	0.30	0.37	0.69	0.69

Table 2.4: Country Fixed Effects Regressions-Full Data Sample

To answer this, Table 2.4 contains the same regression on the sub-sample for which education data is available.[8] The similarity of the coefficient on the squared inequality

[8]Therefore, columns 2 and 4 are the same between tables 2.3 and 2.4

term between columns 3 and 4 indicate that the sign reversal is due to systematic differences between the observations with education data and those missing the education data. This combined with the insignificance of the education data indicate that it can be safely be ignored in favor of the larger sample.

2.4 Conclusions

This chapter develops a model that relaxes the traditional assumption of costless and free information. Instead, agents must spend time analyzing information to form accurate expectations about the future. This creates a non-monotone relationship between inequality and the aggregate labor supply. However this relationship is rotated when compared to the original Kuznets hypothesis. Therefore when compared in the same space as the Kuznets curve, the model predicts a 'C' shape rather than an inverted 'U'. Using specifications from the literature, I test to see if the data exhibits this pattern. Pooled estimation is highly supportive, while fixed effects estimation is more mixed but generally supportive.

Chapter 3

Wage Progression and U.S. Labor Market Asymmetry

Search models have become the preferred method for studying equilibrium unemployment. However, numerous studies have shown that these models produce unrealistic second and third moments for key labor market variables. Existing models that attempt to account for volatility or skewness are incomplete. Some papers focus on the volatility issue [Gertler and Trigari, 2009, Hagedorn and Manovskii, 2008]. Others focus on volatility and skewness for a selected subset of variables in a partial equilibrium setting [Ferraro, 2013]. No existing model successfully accounts for the observed second and third moments simultaneously in a general equilibrium framework.

This paper addresses three specific problems with the higher order moments generated by existing theory relative to the observed data. First, common parameterizations of the textbook search model of unemployment, originally proposed by Diamond, Mortensen, and Pissarides (subsequently referred to as DMP) result in very little unemployment volatility compared to observed data [Shimer, 2005, Costain and Reiter, 2008, Hall, 2005]. This issue is commonly referred to as the unemployment volatility puzzle. Second, U.S. labor market variables and investment show a distinct pattern of skewness over the business cycle. Unemployment and job-finding probability are positively skewed and negatively skewed respectively; consequently, troughs associated with recessions are on average larger in magnitude than peaks. Additionally, investment shows significant negative skewness. I refer to these observations as the asymmetry puzzle. The third issue is the relative volatility and co-movement between wages and productivity. In the DMP model, wages are determined by Nash bargaining. The wage is perfectly flexible in the sense that bargaining occurs in every period. The resulting average wage is too volatile relative to the U.S. data. Also, in the DMP model the wage moves one-for-one with productivity, which is at odds with observation. I refer to this problem as the wage-productivity puzzle. No current theory is able to address these three issues simultaneously. The model put forth in this paper, referred to as the Wage Progression model, will address each of the puzzles.

This paper addresses these three problems in two ways. First, I develop a new wage setting mechanism that links on-the-job search and wage rigidity. This method of wage determination results in workers progressing along a wage ladder. The flows along this ladder vary endogenously with labor market conditions. This helps to address the wage-productivity puzzle. Second, I adopt a calibration that is consistent with not only the second moments of the job-finding rate but also with the third moments. This allows for more realistic dynamics of aggregate household income. Consequently the proposed model is able to more accurately model the volatility and skewness of consumption and investment without relying on higher order shocks.

This paper contributes to the literature by proposing a new wage determination mechanism. Bargaining over wages only occurs when an employed worker forms a match with another firm. This is similar to the steady state model of Cahuc et al. [2006] however in a dynamic setting. The number of matches between employed workers and firms is driven by labor market conditions. This gives rise to wage rigidity for employed workers that endogenously varies over the business cycle. During expansions, wage rigidity is low because a relatively large number of wage renegotiations occur. Therefore, the number of renegotiations occurring each period will have cyclical properties which will effect the dynamics of the average wage. The Wage Progression model is able to capture the correlation and relative volatility of the average wage and productivity.

The new method of wage determination proposed in this paper also contributes to the literature by ensuring that a worker and a firm renegotiate only when the worker can credibly threaten to move to another employer. Under Calvo or Taylor pricing, wage setting occurs without regard to the implication for the firm and employee's surplus. In contrast, I change the outside option in the bargaining game played by employed workers. An entering firm that matches with an employed worker recognizes the value of the worker's existing wage contract, rather than using the value of unemployment as the outside option. This approach ensures that renegotiating the wage delivers positive incremental surplus to the worker. The incumbent firm may then either choose to lose the employee and get zero surplus or to match the offer and receive positive surplus.

I propose a calibration of the model that generates both realistic volatility in the job-finding rate and skewness of unemployment. This calibration features a higher replacement ratio than is standard in the literature. While controversial, I show that a higher replacement ratio is consistent with the observed skewness of labor market variables as well as with aggregate consumption and investment. The responsiveness of firms' vacancy creation to changes in the state is key to this result. Calibrations utilizing a high replacement ratio to generate volatility have been criticized for failing to address the wage-productivity relationship [Gertler and Trigari, 2009]. The new method of wage determination proposed in this paper addresses this criticism.

I begin by describing previous work related each of these three puzzles in Section 2. Next, I describe the Wage Progression model in detail in Section 3. Section 4 details my calibration of the model and describes the data used to evaluate the model. Section 5 compares the moments of the U.S. data to the simulated Wage Progression model, and to simulated moments of several models in the literature. Section 6 sets forth my conclusions.

3.1 Proposed Solutions in the Literature

Each of the three puzzles has been addressed to varying degrees in the previous literature, although they have not been fully resolved. The calibrated model I propose differs from other solutions described in the literature because the Wage Progression model is able to jointly address these three puzzles. Further, these puzzles have not been addressed in a unified general equilibrium framework.

Broadly, there are two categories of existing solutions to the unemployment volatility puzzle; however, both proposed solutions have shortcomings. Furthermore, existing solutions to the unemployment volatility puzzle are unable to address the other puzzles previously discussed. The first category involves a natural solution to the unemployment volatility puzzle, which is to assume that wages are rigid. To generate additional variation in unemployment, workers flowing from the pool of unemployed into the work force are subject to wage rigidity. Wage rigidity in existing employer-employee matches is insufficient to generate unemployment volatility because it does not change the vacancy creation decision of firms. However, if wages for newly hired workers are rigid, then employment adjusts along the extensive margin. The result is lower volatility of the wage and higher unemployment volatility. This solution is pursued by Gertler and Trigari [2009] through a Calvo mechanism in wage setting to generate wage rigidity. They assume that newly hired workers are subject to the same wage rigidity as workers in an ongoing employment relationship. This assumption creates realistic volatility; however, the model generates symmetric simulated data. Further, there is a large body of evidence suggesting that the wages of new hires are more responsive to changes in productivity than the wage of existing workers [Haefke et al., 2013, Pissarides, 2009, Kudlyak, 2014]. This evidence casts doubts on the plausibility of role of wage rigidity as an amplification mechanism.

The second category of existing solutions to the unemployment volatility puzzle involves an alternate calibration of the DMP model. Hagedorn and Manovskii [2008] (subsequently referred to as HM) employs an alternate calibration that features worker bargaining power near zero and a very high replacement ratio. The HM calibration has the effect of creating very high labor supply elasticity, which generates large movements in employment as a result of small changes in the wage. Further the very low worker bargaining power results in low elasticity of the wage with respect to labor productivity, which is consistent with observation. However, this set of parameters faces three significant problems. First, the HM calibration is controversial because the assumption that worker bargaining power near zero is at odds with empirical findings. Point estimates of worker bargaining power find a value significantly higher than what is used in the HM calibration [Flinn, 2006]. Second while the HM calibration matches the elasticity of wages with respect to productivity, it is unable to resolve the wage-productivity puzzle by matching the correlation and relative volatility of wages and productivity. Finally, this parametrization overstates cyclical variation in the job-finding rate in order to match the variation in unemployment. This overstatement is present in any model that assumes a constant, exogenous separation rate while matching the variation in unemployment. The constant separation rate assumption is supported by Hall [2005] and Shimer [2012] and is widely utilized throughout the literature. However, Fujita and Ramey [2012] finds that variation in the separation rate accounts for a non-trivial proportion of the variation in unemployment. Therefore, the total variation in unemployment could be considered the wrong empirical target for a model that assumes a constant separation rate.

While this paper features a high replacement ratio as does the HM calibration, I address each of the three difficulties with the HM calibration described above. First, I use a realistic value for worker bargaining power that is consistent with both the literature and empirical observation. Also, the new wage determination mechanism proposed in this paper generates realistic volatility and co-movement of wages and productivity. Finally, this paper will focus on the volatility in the job-finding rate rather than the volatility of unemployment.

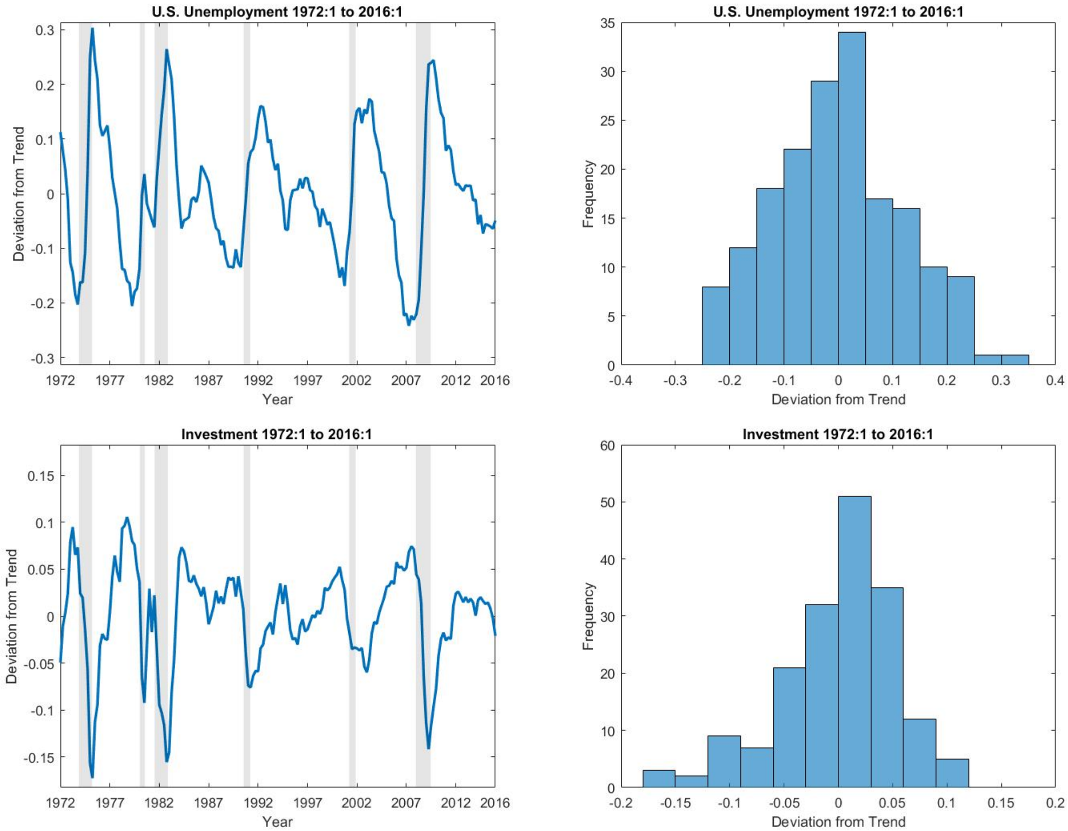

Figure 3.1: U.S. Unemployment and Investment

Note: Investment is the sum of Real Gross Private Domestic Investment and Real Personal Consumption Expenditures on Durables. The unemployment rate is the Civilian Unemployment rate. All data is at the quarterly frequency, measured in logs, and HP filtered with a smoothing parameter of 1600.

Regarding the asymmetry puzzle, this paper is concerned with two main facts that are not addressed by the standard DMP model. First, unemployment shows significant positive skew. Second, investment shows significant negative skew. Figure 3.1 depicts the pattern of skewness for unemployment and investment in the U.S. I use the third central, standardized moment to measure asymmetry. Tests of normality based on the third and fourth moments strongly reject normality of the unemployment, employment, and investment time series. When viewed as a time series, the unemployment peak deviations from trend are larger in magnitude than the troughs. The reverse is true of the investment time series. Alternatively, when viewed as a cross section, the unemployment distribution shows a long right tail, while investment displays a long left tail. These findings are consistent with a long body of literature[1]. The findings in the literature regarding the asymmetry of output are less consistent. McKay and Reis [2008] and Belaire-Franch and Peiro [2003] find evidence that employment is strongly asymmetric while output is not. Others find evidence supporting that output is asymmetric [Ferraro, 2013, Van Nieuwerburgh and Veldkamp, 2006].

The bulk of the literature on business cycle asymmetry have focused on either investment or output. While a few papers have addressed the asymmetry of labor markets [Andolfatto, 1997, McKay and Reis, 2008, Ferraro, 2013], this paper is the first to link the asymmetry of unemployment with that of investment. Ferraro [2013] develops a partial equilibrium search model that delivers realistic unemployment skewness while also accounting for selected second moments. That model delivers realistic volatility for unemployment, job-finding rate, and vacancies. However, it does not address the volatility of the wage relative to the unemployment time series as highlighted in Gertler and Trigari [2009]. Van Nieuwerburgh and Veldkamp [2006] develops asymmetry over the business cycle that relies on informational differences. Agents are more easily able to detect the beginning of a downturn than the start of a recovery. This generates asymmetry in output and hours worked; however, the model is agnostic to the unemployment rate. McKay and Reis [2008] develops a model that delivers labor market asymmetry via three modeling elements. The most striking of these elements is that jobs are assumed to be more easily destroyed than created. Similar to the unemployment volatility puzzle, the asymmetry puzzle has not been adequately resolved particularly in the class of models with a frictional labor market.

The wage-productivity puzzle is the least addressed in the literature. When considering the relationship between wages and productivity many papers consider the elasticity between the two. This elasticity is the regression coefficient of log wages on log productivity. Therefore, it is the product of the correlation and the relative volatility of wages and

[1]See Rothman [1991], Bai and Ng [2005], Belaire-Franch and Peiro [2003], Van Nieuwerburgh and Veldkamp [2006], McKay and Reis [2008], Jovanovic [2006] for evidence supporting the asymmetry of labor markets and investment.

productivity. The standard DMP model fails to match the elasticity, relative volatility, and correlation between wages and productivity. Wages and productivity are perfectly correlated, and they are of nearly the same volatility. The resulting elasticity between wages and productivity is near unity in the DMP model. Gertler and Trigari [2009] addresses this issue by introducing wage rigidity into the wage determination. This results in correlation and volatility consistent with observation. Further, Gertler et al. [2008] shows that while the HM calibration matches the elasticity of wages to productivity, it fails to match the correlation and relative volatility between the two. Thus the wage determination mechanism of the Hagedorn and Manovskii [2008] model does not properly capture the relationship between productivity and the wage outcome.

3.2 The Wage Progression Model

I make several modifications to a frictional model of the labor market. First, I introduce on-the-job search and link it to wage rigidity. This results in endogenously counter-cyclical wage rigidity for employed workers as well as wage dispersion among workers. Therefore, the average wage will have different dynamics relative to the wage setting mechanisms in other models. Additionally, I describe a new computational algorithm for solving the model.

3.2.1 General Environment

Time is discrete and infinite. There exists a unit mass of agents and a positive mass of firms and potential entrants. Workers are identical in ability and firms are identical in the technology they operate. Firms employ a single worker and combine capital and labor to produce output according to a constant returns, Cobb-Douglas production function. Firms either produce output or post vacancies to attract workers. Productivity follows an AR(1) process, which I approximate using the method of Tauchen [1986]. I assume aggregate productivity can take on one of nine possible values. Let $\pi_{s,s'}$ denote the probability that next periods productivity is $z_{s'}$, conditional on the current state being z_s.

I assume matches are dissolved with constant and exogenous probability $(1-\rho)$. Empirically, approximately 75% of unemployment fluctuations are caused by variation in the job-finding rate [Shimer, 2012]. I follow Shimer [2005] and Shimer [2012] by focusing on the job-finding probability. This will mean the model cannot explain 100% of the variation in unemployment present in the data, and in fact should not explain all of the variation. Further, I abstract from the labor participation margin; this is another source of variation in unemployment. The model should however explain a large portion of it through variation in the job-finding rate. Making the separation rate endogenous to the model would no doubt be preferable and will be the focus of future work. Other works have stressed the contribution of the job separation rate to the unemployment volatility puzzle [Fujita and Ramey, 2009, 2012]. Still others stress the importance of the job separation rate in generating asymmetric dynamics [Ferraro, 2013].

Workers have linear preferences over the wage and maximize the expected discounted sum of utility. Employment is subject to matching frictions, and so matches result in a positive surplus. As is standard, there exists a range of wages at which both the firm and the employee would prefer to remain matched. When a firm and an unemployed worker form a match, the two play a Nash bargaining game to determine the wage and the split of the expected future surplus created by the match. Wages are subject to rigidity, and therefore may last for multiple periods. The bargaining game takes into account this possibility.

I also allow for a simple version of on-the-job search, which results in matches between firms and currently employed workers. Employed workers costlessly and effortlessly sample job postings. Upon matching with a potential entering firm, the timing of events is as follows. The employed worker and the entrant play a two player Nash bargaining game that recognizes the value of the worker's existing wage contract as the outside option. The outcome of the bargaining game is a wage contract offered to the worker by the entrant firm. Then the incumbent firm has the right to match the newly determined wage contract. The alternative to matching the new wage contract is to lose the worker and get the value of posting a vacancy. Therefore I require an assumption to break indifference between the offers. I assume that employed workers prefer to remain at the incumbent firm if all else is equal.

Given the structure of the problem, the value to the worker from the incumbent's wage offer and the entrant's wage are are the same. The incumbent always chooses to match the entrant's offer, as the alternative is to post a vacancy and receive 0. Also, the incumbent has no incentive to offer more surplus to the worker than the entrant's offer, because I assume that ties are broken in favor of the incumbent. Therefore when given the opportunity, an employed worker that matches with an entrant continues to work at the incumbent firm. However, the incumbent must match the offer of the entrant to retain the worker.

I assume that matches with unemployed and with employed workers each are governed by separate matching technologies. The probability that an unemployed worker matches with a firm is a function of the labor market tightness, which is defined as the ratio of vacancies to job seekers. Let $\theta_u = {}^{v}/_{u}$ and $\theta_e = {}^{v}/_{e}$ denote the tightness with respect to unemployed and employed workers respectively. A vacancy can result in a match with either an employed worker or an unemployed worker. The number of matches formed with unemployed and employed workers are respectively:

$$m_u = \mu_u v^{\gamma_u} u^{1-\gamma_u}$$

$$m_e = \mu_e v^{\gamma_e} e^{1-\gamma_e}$$

The probability that a vacancy is filled by an unemployed worker or an employed worker is respectively given by:

$$q_u = \frac{m_u}{v} \quad q_e = \frac{m_e}{v}$$

The probability that an unemployed or employed workers is matched is given by:

$$p_u = \frac{m_u}{u} \quad p_e = \frac{m_e}{e}$$

3.2.2 Workers and Firm Value Functions

Given the probabilities, I now define the firm and worker value functions. The vector of relevant state variables includes the current wage, the level of employment, and the level of aggregate productivity. The aggregate level of employment enters the state space as it will be important in determining the transition of the wage. For compactness, let $s = (e, z, k)$ denote the aggregate state vector. The present value of employment at wage ω in state z to the worker is:

$$W(\omega, s) = \omega + \beta E\left\{\rho W(\omega', s') + (1-\rho)U(s')\right\}$$

Let b denote the flow value of unemployment to the worker. The present value of unemployment to the worker is:

$$U(s) = b + \beta E\left\{p_u W(\omega', s') + (1-p_u)U(s')\right\}$$

Given the assumption of constant returns in production, it is simplest to consider per worker quantities. All firms will employ the same amount of capital, and all workers operate at the same marginal and average products. Letting k denote capital per worker, and $J(\omega, e, z, k)$ denote the value to the firm of employing a worker at wage ω in state z with capital level k.

$$J(\omega,s) = \max_{k} zk^{\alpha} - \omega - rk + \beta E\left\{\rho J(\omega',s')\right\}$$

Capital markets are perfectly competitive and frictionless. Therefore the value of the firm can be re-expressed by using the usual capital first order condition:

$$J(\omega,s) = z(1-\alpha)k^{\alpha} - \omega + \beta E\left\{\rho J(\omega',s')\right\}$$

Also let $V(z)$ denote the value of posting a vacancy in state z. I assume free entry, and the value of posting a vacancy is 0 in all periods.

$$V(s) = -\kappa + \beta E\left\{q_u J(\omega',s') + q_e 0\right\}$$

Each of these value functions does not involve a max operator. It is possible to consider the problem as a dynamic discrete choice with trivial policy functions. The policy functions dictate that workers always accept employment offers and that firms always match competitor's wage offers. Given this, it is possible to think of the problem as a standard dynamic programming problem with a singleton choice set, and thereby apply standard results.

It is important to note that the decisions of the individual worker and firm do not depend on the aggregate distribution of workers over wages. Instead, the decisions depend only on labor market conditions through the ratio of vacancies to employment or unemployment. This is similar to block recursion of Menzio and Shi [2010], and it ensures a finite state space.

The value functions described are not fully characterized in the sense that they did not define the state transition. I now turn to the evolution of the state variables. Let $\omega_e^*(\omega,s)$ denote the outcome of an employed worker bargaining a new wage given that the current wage is ω and the aggregate state is (e,z,k). The Nash bargaining process that determines this wage will be made specific later. Similarly, let $k*$ denote the outcome of the representative household's savings decision; this decision will be discussed later. Following Tauchen [1986], I allow the productivity shock to take on nine possible values. Let $\pi_{i,j}$ denote the probability that the productivity shock takes on the j-th value next period conditional on currently being in the i-th state. Conditional on the current aggregate state vector being (ω, e, z_s), the aggregate state vector evolves according to:

$$(\omega', e', z', k') = \begin{cases} (\omega, \rho e + m_u, z_1, k*) & pr = (1 - p_e)\pi_{s,1} \\ \vdots & \vdots \\ (\omega, \rho e + m_u, z_N, k*) & pr = (1 - p_e)\pi_{s,N} \\ (\omega_e^*(\omega, e', z_1, k*), \rho e + m_u, z_1, k*) & pr = p_e \pi_{s,1} \\ \vdots & \vdots \\ (\omega_e^*(\omega, e', z_N, k*), \rho e + m_u, z_N, k*) & pr = p_e \pi_{s,N} \end{cases}$$

This imposes the law of motion of employment:

$$e' = \rho e + m_u$$

The number of workers transitioning into employment, m_u, is governed by labor market tightness and the matching technology. Next period's level of employment is predetermined, as firms' choice of vacancies today determines the number of matches that flow into employment next period. With probability $(1 - p_e)$ the existing wage contract is remains in place next period. With probability p_e the wage is renegotiated next period, and the outcome of that negotiation is denoted $\omega_e^*(\omega, s')$. The level of employment is relevant to the state space as it determines the renegotiation probability, p_e. This is due to the specification of matching technologies and the fact that $\theta_e = \frac{u}{e}\theta_u$.

I assume that all workers flowing from unemployment into employment optimally negotiate an entry wage. There is varied evidence on the cyclicality of the wage of new hires. Gertler and Trigari [2009] assumes that newly hired workers are brought in at the existing wage. This is key to the mechanism in their model. Wage rigidity effects the present value of employing a newly hired worker. Therefore overall wage rigidity influences the vacancy creation and the hiring decision of firms. However, Haefke et al. [2013] presents compelling evidence from CPS micro-data that the wages of newly hired workers adjust one for one with labor productivity. Similarly, Pissarides [2009] casts doubt on extent of wage rigidity for workers transitioning from unemployment into employment. If the wages of newly hired workers are flexible, then wage rigidity cannot generate realistic unemployment fluctuations. Haefke et al. [2013] demonstrates this via a search and matching model with wage rigidity in existing employment relationships but not in new ones. Let $\omega_u^*(s)$ denote the outcome of a wage negotiation with a newly hired employee. Given this and the transition of the aggregate state, the value functions can be rewritten as:

$$W(\omega,s) = \omega + \beta E\left\{\rho(1-p_e)W(\omega,s') + \rho p_e W\left(\omega_e^*(\omega,s'),s'\right) + (1-\rho)U(s')\right\}$$

$$U(s) = b + \beta E\left\{p_u W(\omega_u^*(s'),s') + (1-p_u)U(s')\right\}$$

$$J(\omega,s) = z(1-\alpha)k^{\alpha} - \omega + \beta E\left\{\rho(1-p_e)J(\omega,s') + \rho p_e J\left(\omega_e^*(\omega,s'),s'\right)\right\}$$

$$V(s) = -\kappa + \beta E\left\{q_u J(\omega_u^*(s'),s') + q_e 0\right\}$$

It is important to note that the probability an employed worker forming a match and thereby renegotiating the wage contract is time variant. However it is predetermined with respect to the expectation operator above. It is determined by the free entry condition, the current level of employment, and the employed worker matching technology. Therefore the expectation is over the possible future values of productivity conditional on the current value.

Since Nash bargaining occurs over surpluses, it is often helpful to define value functions that deal in surpluses directly. For unemployed workers, the surplus generated by employment is given by:

$$H(\omega,s) = W(\omega,s) - U(s)$$

For workers that are currently employed at wage $\check{\omega}$, the surplus generated by negotiating a new wage contract is given by:

$$G(\omega,\check{\omega},s) = W(\omega,s) - W(\check{\omega},s)$$

Substituting for the worker's value functions yields:

$$H(\omega,s) = \omega - b + \beta E\left\{\rho H(\omega,s') + \rho p_e G(\omega_e^*(\omega,s'),\omega,s') - p_u H(\omega_u^*(s'),s')\right\}$$

$$G(\omega,\check{\omega},s) = \omega - \check{\omega} + \beta E\left\{\rho(1-p_e)G\left(\omega,\check{\omega},s'\right) + \rho p_e G\left(\omega_e^*(\omega,s'),\omega_e^*(\check{\omega},s'),s'\right)\right\}$$

It is important to note that the value functions are differentiable with respect to the state variables, and further the partials of the firm and worker value functions with respect to the wage are equal in absolute value. The equivalence of the partials in absolute value is important for two reasons. First, it ensures that the frontier of the bargaining set is linear, and therefore the bargaining set is convex. This ensures that the maximization problems that will define the wage bargains have unique solutions. Second, it allows for simplification of the first order conditions from these maximization problems that implicitly define the optimal wages.

In addition, it is often convenient to work in terms of total surplus generate by a match with an worker. For a match with an unemployed worker the total surplus is given by:

$$S^u(\omega,s) \equiv W(\omega,s) + J(\omega,s) - U(s)$$

Inserting the earlier definitions for the agent surplus value functions yields:

$$\begin{aligned} S^u(\omega,s) = {} & z(1-\alpha)k^\alpha - b + \beta\rho E\left\{S^u\left(\omega,s'\right)\right\} + \beta\rho p_e E\left\{S^u\left(\omega_e^*(\omega,s'),s'\right) - S^u\left(\omega,s'\right)\right\} \\ & - \beta p_u E\left\{H\left(\omega_u^*(s'),s'\right)\right\} \end{aligned}$$

Similarly, the total surplus generated by a match with a worker that is currently employed at $\check{\omega}$ is given by:

$$S^e(\omega,\check{\omega},s) \equiv G(\omega,\check{\omega},s) + J(\omega,s)$$

$$S^e(\omega,\check{\omega},s) = z(1-\alpha)k^\alpha - \check{\omega} + \beta\rho(1-p_e)E\left\{S^e\left(\omega,\check{\omega},s'\right)\right\} + \beta\rho p_e E\left\{S^e\left(\omega_e^*(\omega,s'),\omega_e^*(\check{\omega},s'),s'\right)\right\}$$

It is important to note that the total surplus created by the match is independent of the wage determined by the Nash bargaining game. The wage determines the split of the total surplus, but the total surplus does not change with the wage. An alternative interpretation of the linear bargaining set is that the total surplus is independent of the wage. This is in contrast to Gertler and Trigari [2009] where the bargaining set was non-convex. Therefore the total surplus actually depended on the outcome of the bargaining game.

Using the above result, it is possible to simplify the total surplus for a match with an unemployed worker.

$$S^u(s) = z(1-\alpha)k^\alpha - b + \beta\rho E\left\{S^u\left(s'\right)\right\} - \beta p_u E\left\{H\left(\omega_u^*(s'), s'\right)\right\}$$

$$S^e(\check{\omega}, s) = z(1-\alpha)k^\alpha - \check{\omega} + \beta\rho(1-p_e)E\left\{S^e\left(\check{\omega}, s'\right)\right\} + \beta\rho p_e E\left\{S^e\left(\omega_e^*(\check{\omega}, s'), s'\right)\right\}$$

Whenever the wage is being determined optimally in the next period, it is possible to relate the agent value functions to the germane total match surplus value function. The first order condition that implicitly determines the result of wage bargaining will imply this relationship.

3.2.3 Wage Determination

The wage for workers transitioning from unemployment to employment is determined through Nash bargaining. The solution to the Nash bargaining game corresponds to the SPNE of the two player bargaining game due to Rubinstein as the time between bargaining rounds goes to zero. The negotiated wage remains in place during the next period with probability $(1-p_e)$. With probability p_e the wage is renegotiated in the subsequent period. In the case of a renegotiation, the new wage incorporates the outside option of the employed worker in a manner to be formalized later. For now, denote the outcome of a renegotiation with an employed worker by $\omega_e^*(\omega, e, z, k)$. Upon forming a match with an unemployed worker, the firm and the worker negotiate the wage to solve:

$$\omega_u^*(s) \equiv \arg\max_{\omega} \left[W(\omega, s) - U(s)\right]^\eta \left[J(\omega, s) - V(s)\right]^{1-\eta}$$

As discussed previously, the bargaining set is convex. Therefore there exists a unique solution to the maximization problem. The first order condition of the maximization problem implicitly determines the wage for new workers:

$$\eta J\left(\omega_u^*(s), s\right) \frac{\partial W(\omega, s)}{\partial \omega} + (1-\eta)\left[W\left(\omega_u^*(s), s\right) - U(s)\right] \frac{\partial J(\omega, s)}{\partial \omega} = 0$$

where

$$\frac{\partial W(\omega, s)}{\partial \omega} = 1 + \beta E_t \left\{ \rho(1-p_{e,t}) \frac{\partial W(\omega, s')}{\partial \omega} + \rho p_{e,t} \frac{\partial W\left(\omega_e^*(\omega, s'), s'\right)}{\partial \omega} \frac{\partial \omega_e^*(\omega, s')}{\partial \omega} \right\}$$

$$\frac{\partial J(\omega, s)}{\partial \omega} = -1 + \beta E_t \left\{ \rho(1-p_{e,t}) \frac{\partial J(\omega, s')}{\partial \omega} + \rho p_{e,t} \frac{\partial J\left(\omega_e^*(\omega, s'), s'\right)}{\partial \omega} \frac{\partial \omega_e^*(\omega, s')}{\partial \omega} \right\}$$

With period by period Nash bargaining, the partials of the value function with respect to the wage would simply be unity. With wage rigidity, the partials form a first order difference equation. In Gertler and Trigari [2009], the partial of the worker's value function and the partial of the firm's value function were different. The mismatch between partials of the value function arose because it was assumed that firms hired new workers at the existing wage. They termed this mismatch the 'horizon effect', and it resulted in time variation in the effective bargaining power of the worker. However they showed that the horizon effect was of little importance to their proposed solution to the volatility puzzle. In addition, the horizon effect also resulted in a non-convex bargaining set. Or equivalently, the present value of the total surplus created by a match with an unemployed worker depended on the level of the wage. So the wage was formed by solving an non-convex optimization problem. Therefore numerical methods were required to ensure the solution was a global optimum. Removing the horizon effect from the model ensures that the optimization problem that determines the wage is well behaved.

In my model, workers that transition to employment always negotiate a wage optimally, rather than accept an existing wage. The cyclicality of the wage for newly hired workers versus existing workers has been the subject of some contention.

One result of my assumption that newly hired workers negotiate a wage is that the partials of the value functions sum to 0. This allows for simplification of the first order necessary condition, and also ensures that the bargaining set is convex. Additionally, assuming new hires are brought in at the optimal wage ensures that the wage distribution does not enter the individual's problem. In the model of Gertler and Trigari [2009], the average wage was introduced into the agents' problems because workers transitioning from unemployment to employment could expect to receive the average wage. This introduced spillover effects that were a source of additional rigidity.

Given the equality of the partials in absolute value, rearranging the first order condition yields a variation of the familiar 'split the difference' rule:

$$\eta J\left(\omega_u^*(s), s\right) = (1-\eta)\left[W\left(\omega_u^*(s), s\right) - U(s)\right] = (1-\eta) H\left(\omega_u^*(s), s\right)$$

Or equivalently:

$$H\left(\omega_u^*(s), s\right) = \eta S^u(s)$$

$$J\left(\omega_u^*(s), s\right) = (1-\eta) S^u(s)$$

Substituting the value function definitions and grouping like terms show that the optimal wage for a worker entering employment is:

$$\begin{aligned}
\omega_u^*(s) \quad = \quad & \eta z(1-\alpha)k^{\alpha} + (1-\eta)b + \beta E\left\{\eta\rho(1-p_e)J\left(\omega_u^*(s), s'\right) - (1-\eta)\rho(1-p_e)H\left(\omega_u^*(s), s'\right) + (1-\eta)p_u H\left(\omega_u^*(s'), s'\right)\right\} + \\
& \beta E\left\{\eta\rho p_e J\left(w_e^*(\omega_u^*(s), s'), s'\right) - (1-\eta)\rho p_e G\left(\omega_u^*(s), s'\right)\right\}
\end{aligned}$$

This can be simplified further. In the case of a renegotiation next period, which takes place with probability p_e, the wage is determined optimally. So, the final term in the above equation will be shown to be zero through the first order condition of the employed worker bargaining problem from the subsequent period. Therefore the wage can be expressed as:

$$
\begin{aligned}
\omega_u^*(s) &= \eta z(1-\alpha)k^{\alpha} + (1-\eta)b + \beta\rho(1-p_e)E\left\{\eta J\left(\omega_u^*(s),s'\right) - (1-\eta)H\left(\omega_u^*(s),s'\right)\right\} + \\
&\quad \beta(1-\eta)E\left\{p_u H\left(\omega_u^*(s'),s'\right) - \rho p_e H\left(\omega_u^*(s),s'\right)\right\}
\end{aligned}
$$

The first two terms are related to wage that arises without rigidity. The second term is unique to a setting with wage rigidity. With probability $\rho(1-p_e)$ the wage that is negotiated in the current period will survive to the next period. So the optimal wage this period is forward looking and reflects the rigidity. The final term reflects the difference in next period's surplus evaluated at today's optimal wage and tomorrow's optimal wage conditional on being employed.

Employed workers that match with potential entrants determine wages through a similar two player bargaining game. As discussed previously, incumbent firms then match the new offer and the worker chooses to remain at the incumbent firm. Given an existing wage $\check{\omega}$, the new wage is determined as follows:

$$\omega_e^*(\check{\omega},s) \equiv \arg\max_{\omega} \left[W(\omega,s) - W(\check{\omega},s)\right]^{\eta} \left[J(\omega,s)\right]^{1-\eta}$$

The wage bargain for an employed worker recognizes the value of the current wage contract as the worker's outside option. This ensures that renegotiation results in higher surplus for the worker. From the incumbent firm's perspective, matching the new wage contract is also optimal. Again using the equality of the partials in absolute value, the corresponding first order condition implicitly defines the renegotiated wage:

$$\eta J(\omega_e^*(\check{\omega},s),s) - (1-\eta)\left[W(\omega_e^*(\check{\omega},s),s) - W(\check{\omega},s)\right] = 0$$

Substituting for the firm and worker's surplus functions yields an expression for the optimal wage:

$$\begin{aligned}\omega_e^*(\check{\omega},s) &= \eta z(1-\alpha)k^{\alpha} + (1-\eta)\check{\omega} + \beta\rho(1-p_e)E\left\{\eta J\left(\omega_e^*(\check{\omega},s),s'\right) - (1-\eta)G(\check{\omega},s')\right\} \\ &\quad + \beta\rho p_e E\left\{\eta J\left(\omega_e^*\left(\omega_e^*(\check{\omega},s),s'\right),s'\right) - (1-\eta)G\left(\omega_e^*(\check{\omega},s'),s'\right)\right\}\end{aligned}$$

Iterating the first order condition forward shows that the final term is zero. That is, in the case of a renegotiation next period the wage is determined optimally.

3.2.4 Consumption Choice

I use the representative family framework to model the intertemporal savings decision. This follows the work of Merz [1995] and Gertler and Trigari [2009]. Let Π denote firm profits which are returned to the households and T denote government transfers used to finance payments to the unemployed. I assume the government runs a balanced budget in all periods, and therefore $T + (1-e)b = 0$.

$$\Omega(\bar{\omega},s) = \max_{k'} log(c) + \beta E\left\{\Omega(\bar{\omega}',s')\right\}$$

subject to

$$c + k' = \bar{\omega}e + b(1-e) + (1-\delta+r)k + T + \Pi$$

Using value or policy function iteration to solve for the savings decision would require characterizing the transition of the average wage. This is a nontrivial task, as the wage distribution evolves according to the sequence of shocks and is non-parametric in they dynamic setting. Additionally, the evolution of the average wage depends on vacancy creation through the renegotiation probability. This creates a simultaneity problem. Vacancy creation incentives vary with the future level of capital while the savings decision depends on the mass of vacancies. For these reasons, it is convenient to solve for the policy function using the total resource constraint. The total resource constraint splits the total resources in the economy between consumption, savings and vacancy creation costs.

$$c + k' = zk^{\alpha} + (1-\delta)k - \kappa v$$

Considering the social planner's problem instead of the individual problem is valid in the representative family framework. All household in come is pooled and then distributed, therefore all individuals have perfect consumption insurance. Additionally, this approach highlights the link between the labor market and investment. Correctly modeling the fluctuations in firm's vacancy creation incentives is important to accurately modeling fluctuations in aggregate household income.

3.2.5 Steady State

Considering the steady state of the model gives good intuition for how the model operates in the stochastic environment. Additionally the model generates a non-degenerate wage distribution even in the steady state. To my knowledge, this is an unique feature worth exploration.

Let a tilde denote the steady state value of a variable. In the steady state z=z', which I normalize to unity; let $\tilde{s}$ denote the aggregate state vector in the steady state. The model has the feature of a non-degenerate distribution of workers across wages in the steady state. This is a direct result of worker renegotiations occurring each period with probability $\tilde{p}_e$. However the distribution is time invariant as worker flows across the different wages are balanced.

The steady state is determined by a system of nonlinear equations. The following equations are definitional:

$$\tilde{u} = 1 - \tilde{e}$$

$$\tilde{\theta}_u = \frac{\tilde{v}}{\tilde{u}}$$

$$\tilde{\theta}_e = \frac{\tilde{v}}{\tilde{e}}$$

For $i \in \{e, u\}$, the matching technologies specify the job-finding probability, vacancy filling probability and mass of matches formed:

$$\tilde{p}_i = \mu_i \tilde{\theta}_i^{\gamma_i}$$

$$\tilde{q}_i = \frac{\tilde{p}_i}{\tilde{\theta}_i}$$

$$\tilde{m}_i = \tilde{p}_i i$$

The steady state 'law of motion' governing employment is:

$$(1 - \rho)\tilde{e} = \tilde{m}_u$$

That is the mass of separations is equal to the mass of workers flowing into employment. Without any time variation in total factor productivity, the wage of workers entering employment is the same every period. Similarly, the first renegotiation of the wage, $\omega_e^*(\omega_u^*(\tilde{s}), \tilde{s})$, does not depend on the timing of the renegotiation relative to the entry into employment. This is true for all renegotiations as well. Therefore, in the steady state, the wage depends only on the total number of times it has been negotiated. Let v_i denote the mass of workers that have renegotiated the wage i times for $i \in \{0,1,2,..\}$. Let v_i' denote the mass of workers at that same wage in the 'next period' in the case that there is no variation in the productivity shock. The law of motion governing the mass of workers that enter employment from unemployment therefore is given by:

$$v_0' = v_0 + \tilde{m}_u - \rho \tilde{p}_e v_0 - (1-\rho) v_0$$

The mass next period is given by the mass this period, plus workers flowing into employment less renegotiating workers and dissolved matches. The mass of workers at any given wage is constant in the steady state.

$$\tilde{v}_0 = \frac{\tilde{m}_u}{1-\rho(1-\tilde{p}_e)}$$

For $i \geq 1$, the mass of workers next period that have renegotiated i times can be expressed as:

$$v_i' = v_i + \rho \tilde{p}_e v_{i-1} - (1-\rho) v_i - \rho \tilde{p}_e v_i$$

Again applying the time invariance of the mass of workers at each wage yields a recursion, which can also be expressed in terms of $\tilde{v}_0$:

$$\tilde{v}_i = \frac{(\rho \tilde{p}_e)\, \tilde{v}_{i-1}}{(1-\rho(1-\tilde{p}_e))} = \frac{(\rho \tilde{p}_e)^i\, \tilde{m}_u}{(1-\rho(1-\tilde{p}_e))^{i+1}}$$

Each of these is an absolute notion of the mass of workers employed at each wage in the steady state wage sequence. Therefore to convert it to a density I normalize by the steady state level of employment, $\tilde{e} = \frac{\tilde{m}_u}{1-\rho}$. The percentage of employed workers that have renegotiated i times is given by:

$$\tilde{v}_i = \frac{(\rho \tilde{p}_e)\,\tilde{v}_{i-1}}{(1-\rho(1-\tilde{p}_e))} \frac{(1-\rho)}{\tilde{m}_u} = \frac{(\rho \tilde{p}_e)^i (1-\rho)}{(1-\rho(1-\tilde{p}_e))^{i+1}}$$

It is worth noting that the mass of employed workers follows a geometric distribution over the sequence of wages. The probability parameter of the distribution is $\frac{1-\rho}{1-\rho(1-\tilde{p}_e)}$.

Free entry implies that the value of entering production or posting a vacancy is 0. The free entry condition relates the value of firm surplus when employing a worker at $\omega_u^*(\tilde{s})$ to the vacancy filling probability:

$$J(\omega_u^*(\tilde{s}), \tilde{s}) = \frac{\kappa}{\beta \tilde{q}_u}$$

To pin $\tilde{q}_u$ I use the first order condition of the unemployed wage bargain and the steady state value of the total surplus generated by a match with an unemployed worker. Solving for the steady state value of total surplus yields:

$$S^u(\omega_u^*(\tilde{s}), \tilde{s}) = \frac{(\tilde{k}^\alpha - b)}{(1-\beta\rho + \beta \tilde{p}_u \eta)}$$

Therefore, applying the first order condition yields the equation that pins the key variable of the model, θ_u:

$$\frac{(1-\eta)(\tilde{k}^\alpha - b)}{(1-\beta\rho + \beta \tilde{p}_u \eta)} = \frac{\kappa}{\beta \tilde{q}_u}$$

In the steady state, the unemployed wage equation becomes:

$$\omega_u^*(\tilde{s}) = \eta \tilde{k}^\alpha + (1-\eta)b + \beta\rho(1-p_e)\{\eta J(\omega_u^*(\tilde{s}),\tilde{s}) - (1-\eta)H(\omega_u^*(\tilde{s}),\tilde{s})\} + \beta(1-\eta)\{p_u H(\omega_u^*(\tilde{s}),\tilde{s}) - \rho p_e H(\omega_u^*(\tilde{s}),\tilde{s})\}$$

The first order condition for the unemployed wage problem always holds in the steady state, therefore the third term is zero. The first order condition can be used to express the final term in units of firm surplus, instead of worker surplus. Then using the free entry condition yields a reduced form for outcome of the unemployed bargaining problem.

$$\omega_u^*(\tilde{s}) = \eta\tilde{k}^{\alpha} + (1-\eta)b + \frac{\eta\kappa(\tilde{p}_u - \rho\tilde{p}_e)}{\tilde{q}_u}$$

This value is the beginning of an infinite sequence of steady state wages. Each renegotiated wage is related to the prior wage through a recursion that is derived from the steady state version of the employed wage equation. In the steady state, the employed wage equation simplifies to:

$$\omega_e^*(\check{\omega}, \tilde{s}) = \eta\tilde{k}^{\alpha} + (1-\eta)\check{\omega}$$

Rearranging shows that the renegotiated wage is greater than the existing wage if the flow profits to the firm are positive at the existing wage contract.

$$\omega_e^*(\check{\omega}, \tilde{s}) - \omega = \eta\left(\tilde{k}^{\alpha} - \check{\omega}\right)$$

Since the number of renegotiations is sufficient to determine the wage in the steady state, a minor change in notation will simplify the characterization of the average wage. Let $\omega_e^*(i)$ denote the wage that results from the i-th renegotiation for $i \in \{0, 1, 2, ..\}$, with $\omega_e^*(0) = \omega_u^*(\tilde{z})$. Then we can express the sequence of steady state wages as through the following recursion:

$$\omega_e^*(i) = \eta\tilde{k}^{\alpha} + (1-\eta)\omega_e^*(i-1)$$

Again using substitution, any wage $\omega_e^*(i)$ can be expressed in terms of the originally negotiated wage $\omega_e^*(0)$:

$$\omega_e^*(i) = \sum_{j=1}^{i}(\eta\tilde{k}^{\alpha})(1-\eta)^{i-1} + (1-\eta)^i\omega_e^*(0)$$

The characterization of the sequence of steady state wages and the distribution of employed worker at each wage allows me to express the average wage, $\tilde{\tilde{\omega}}$. The average wage is the inner product of the two infinite sequences:

$$\tilde{\tilde{\omega}} = \sum_{i=0}^{\infty} v_i \omega_e^*(i)$$

Substituting the expressions derived previously yields a double summation for the average wage.

$$\tilde{\tilde{\omega}} = \tilde{v}_0 \omega_e^*(0) + \sum_{k=1}^{\infty} \left[\frac{(\rho \tilde{p}_e)^k (1-\rho)}{(1-\rho(1-\tilde{p}_e))^{k+1}} \sum_{j=1}^{k} \left((\eta \tilde{k}^{\alpha})(1-\eta)^{j-1} + (1-\eta)^j \omega_e^*(0) \right) \right]$$

This expression is well defined. Both the sequence of wages and masses are monotone and convergent. Therefore Abel's test implies that the average wage converges. While analytically cumbersome, this expression is easily approximated computationally.

3.2.6 Solving the Stochastic Model

I develop a new computational algorithm to solve the stochastic model. This is necessary for two reasons. First, there exists a simultaneity problem between the wage equations and the value functions. The worker and firm value functions depend on the transition of the state, and specifically the outcome of the renegotiated wage. At the same time, the outcome of the renegotiated wage depends on the firm and worker value functions, as made clear in the wage equations. Second, log linearization is not suitable for analyzing higher order moments of the endogenous variables. Log linearization could be used to break the simultaneity, however all higher order moments would be lost.

I solve the model by value function iteration over a state space grid. The unemployed total surplus value function is straightforward as it is independent of the the wage. For an initial guess of $S_0^u(e,z)$, I use the free entry condition to determine an associated vacancy filling probability q_{u1}. Given the vacancy filling probability, I solve for the remaining transition probabilities as well as the mass of matches formed. Using the mass of matches formed with unemployed workers and the law of motion for employment, I solve for the subsequent period's level of employment. Finally, using the values of the transition probabilities and employment found previously, I update the value function according to the following:

$$S_i^u(s) = z(1-\alpha)k^\alpha - b + (\beta\rho - \beta p_{u,i}\eta)E\left\{S_{i-1}^u(s')\right\}$$

The expectation is with respect to the possible realizations of the productivity shock. This is repeated until the value function converges to a given tolerance level. This procedure pins the total surplus for a match with an unemployed worker as well as the employment transition probabilities, given the current state.

Solving for the individual agents' value functions is more challenging, because it requires simultaneously dealing with two unknown functions. The transition of the aggregate state, and therefore the value functions, depends on the optimal wage in the case of renegotiation. Also, the optimal wage equations show that the outcome of wage renegotiations depends on the unknown value functions. Gertler and Trigari [2009] resolves a similar issue by log linearizing the model. However, this solution method is not preferred for my purposes, as I am also considering higher order moments of the data. Instead, I use a particular form of grid search that considers all possible combinations of ω_e^*. First I form a grid over the wage space. Using the wage equation, I approximate $\omega_e^*(\omega, s)$ with the steady state wage renegotiation equation:

$$\omega_e^*(\omega, s) \approx \eta z'(1-\alpha)k^\alpha + (1-\eta)\omega$$

Given this approximation, I then consider nearby points in the wage grid. I define a search region that consists of a specified number of adjacent points on either side of the approximate wage. Next, I consider all possible combinations of points in the search region for the future possible realizations of z. The relatively high persistence of the technology shock is key to making this computationally feasible. I find that in any given current state, there are exactly three possible future realizations of the technology shock that occur with positive probability. Therefore I am able to restrict my attention to the future states that occur with positive probability. For concreteness, consider a search region consisting of three points around the approximate value of $\omega_e^*(\omega, s')$. Given the current state, there are three possible realizations of z'. For each of these I consider 3 possible values of $\omega_e^*(\omega, s')$. In total for each point in the aggregate state vector, I would evaluate 3^3 possible values of the renegotiated wage.

Given the candidate values of the renegotiated wage, I then perform value function iteration on the worker and firm surplus functions. To arrive at the final wage transition from among the candidates, I take advantage of two important things. First, the value of the total match surplus with an unemployed worker, which is determined via value function iteration as described above. The first order conditions from the wage bargains are the second important factor in determining the wage transition. From the candidates for the renegotiated wage, I select the one that makes the employed wage bargain first order condition hold.

This approach introduces a trade-off between the granularity of the wage grid and the region under consideration for the wage transition. A finer grid over the wage space would imply a 'smaller' search region for a given level of computational complexity. Adaptive grid methods are particularly useful in this setting. Determining appropriate grid bounds for each a given set of parameters ensures maximum efficiency for a given number of points. This helps to ensure that the employed wage first order condition holds with greater precision.

3.3 Data and Calibration Strategy

In this section I describe the data used to evaluate model performance as well as the calibration strategy. All data is quarterly from Q1 1972 to Q1 2016. All series are measured in logs and HP filtered with a smoothing parameter of 1600. I use a monthly calibration and then aggregate the data to make it comparable to the quarterly U.S. data.

All data sources and definitions are standard and comparable with the literature. The job-finding probability is calculated in the same manner as Shimer [2012], but I extend the data through 2016. Short term unemployment is measured as the number unemployed less than 5 weeks.[2] The unemployment rate is Civilian Unemployment Rate. I measure the average real wage by Average Hourly Earnings of Production and Non-supervisory Employees. The wage data is then deflated using the CPI. Consumption is measured by real personal consumption expenditures per capital on non-durable goods. Investment is the sum of Real Gross Private Investment and Real Personal Consumption Expenditures on durables. Productivity is the average output per worker. Output is Real Output from the Nonfarm Business Sector.

With the exception of the flow value of unemployment, my calibration of the Wage Progression model uses conservative and well agreed upon parameter values to target observable long term averages in the steady state. In total there are eleven parameters to be calibrated for the Wage Progression model. The discount rate is calibrated to be consistent with a 5% interest rate. I set the separation rate by matching the average job separation rate as reflected in the JOLTS data. This value is in line with other values in the literature. I set the Cobb Douglas parameter to be 1/3, and the capital depreciation rate to be 0.025/3. There are four parameters related to the matching technologies. I use 0.5 for the elasticity parameters. This value is well within the range found in the literature. The scale parameter for the unemployed matching technology is calibrated to match the long term

[2] I perform the adjustment for data from 1994 and beyond suggested by Shimer [2012] in order to correct for changes in BLS survey methodology.

US unemployment rate of 5.9% in the steady state. The scale parameter for the employed matching technology is set to imply an average wage contract duration of 9 months in the steady state. This duration is similar to the duration implied by the static Calvo parameter in Gertler and Trigari [2009]. The vacancy creation cost is calibrated to match a steady state value of unemployed labor market tightness of 0.72 found by Pissarides [2009]. The literature in this area utilizes a variety of steady state values for labor market tightness, ranging from values of 0.64 to 1. The value I select to target is well within the range of the literature.

Parameter	Calibrated Value	Empirical Target or Source
Discount Rate	0.9959	Real Interest Rate
Job Destruction Rate	0.028	JOLTs Separation Rate
Production Function Parameter	0.33	Standard
Capital Depreciation Rate	0.008	Literature
Unemployed Matching Elasticity	0.5	Literature Midpoint
Employed Matching Elasticity	0.5	Literature Midpoint
Unemployed Worker Scale	0.53	Steady State Unemployment Rate
Employed Worker Scale	0.52	Average Wage Duration
Vacancy Cost	0.38	Steady State Labor Market Tightness
Worker Bargaining Power	0.5	Literature
Replacement Ratio	0.9	Midpoint of Literature and HM
Technology Autoregressive	0.923	U.S. Data
Technology Standard Deviation	0.0044	U.S. Data

Table 3.1: Calibration of the Wage Progression Model

I use a worker bargaining share of 50%. This is consistent with most values used in the literature as well as direct estimates of worker bargaining power [Flinn, 2006]. Hagedorn and Manovskii [2008] stands in contrast to the bulk of the literature by using a value of 5%. The low value of the worker bargaining power is required by Hagedorn and Manovskii [2008] in order to match the elasticity of wages to productivity. However as previously noted, this fails to capture the correlation and relative volatility of wages and productivity. The new wage determination mechanism derived in this paper remedies these issues. With it, I am able to address the wage-productivity puzzle and avoid this departure from a well established parameter value.

The final parameter, the flow value of unemployment, is somewhat controversial in the literature e.g. Hagedorn and Manovskii [2008] and Hall and Milgrom [2008]. Values that result in a replacement ratio near unity result in a highly elastic labor supply, since workers are nearly indifferent between working and unemployment. This is the solution to the unemployment volatility puzzle implemented by Hagedorn and Manovskii [2008]. I reduce the flow value of unemployment relative to the HM calibration. For the value of unemployment, I select a value between the Hagedorn and Manovskii [2008] calibration and values used elsewhere in the literature, e.g. Ferraro [2013], Gertler and Trigari [2009], and Hall and Milgrom [2008]. This results in realistic volatility of the job-finding rate but understates the volatility of unemployment due to the assumption of a constant separation rate.

In addition to the calibration of the Wage Progression model, I consider three additional models from the literature. First, I calibrate a textbook DMP model. This uses typical values for worker bargaining power, 0.5, and a replacement ratio of 60%. The value of the replacement ratio is below the value of 72% used by Hall and Milgrom [2008] but above the 42% used in Gertler and Trigari [2009]. The matching parameters, the Cobb-Douglas parameter, and the capital depreciation parameter are taken to be the same as in the Wage Progression calibration. Second, I calibrate a partial equilibrium search model that corresponds to that of Hagedorn and Manovskii [2008]. The HM calibration uses a worker bargaining power near zero, 0.05, and a high replacement ratio, 0.95. In both the DMP and HM calibrations, the cost of vacancy creation is set to maintain a steady state labor market tightness of 0.72 and unemployment rate of 5.9%. Finally, I replicate the staggered Nash Bargaining model of Gertler and Trigari [2009]. In my replication, the Calvo parameter is set to 8/9, implying an average wage contract duration of 9 months. I use the calibration of the original paper, Gertler and Trigari [2009], for the remaining parameters.

3.4 Results

The Wage Progression model is able to simultaneously address the three puzzles that I previously described. Table 3.2 contains moments of the U.S. data, the corresponding moments of the Wage Progression model, and the moments from several related search models.[3] With the calibration described, the Wage Progression model generates more than 75% of the observed volatility in the job-finding rate. It also results in directionally correct skewness for the job-finding rate, unemployment, and investment. With regards to the wage-productivity puzzle, the Wage Progression model correctly captures the correlation and relative volatility between productivity and the average wage as shown in Table 3.3. Further, the earlier models are unable to deliver the observed skewness in investment and consumption.

[3] All reported moments are averages from 4000 replications of 2765 months. The first 2000 observations are then dropped. The draws of the stochastic driving process are identical between the DMP, HM, and Wage Progression model. The Staggered Nash Bargaining model uses the stochastic process of the original paper.

	Selected Variables					
	P_u	U	W	C	I	Y
A. US Economy 1972:Q1 to 2016:Q1						
Volatility Relative to Y	4.27	5.76	0.51	0.59	2.57	1.00
Skewness	-0.28	0.20	0.16	-0.18	-0.85	-0.55
Correlation with Y	0.83	-0.88	0.28	0.83	0.95	1.00
B. DMP Model: Standard Calibration						
Volatility Relative to Y	1.12	1.02	0.90	0.64	3.37	1.00
Skewness	-0.02	0.02	0.01	-0.07	-0.04	0.00
Correlation with Y	1.00	-0.95	1.00	0.64	0.88	1.00
C. DMP Model: HM Calibration						
Volatility Relative to Y	5.01	4.53	0.24	NA	NA	1.00
Skewness	-0.21	0.17	0.14	NA	NA	-0.09
Correlation with Y	1.00	-0.99	1.00	NA	NA	1.00
D. Gertler and Trigari						
Volatility Relative to Y	5.07	4.44	0.56	0.41	2.71	1.00
Skewness	0.02	-0.01	-0.06	-0.03	0.01	0.00
Correlation with Y	0.94	-0.77	0.66	0.90	0.99	1.00
E. Wage Progression						
Volatility Relative to Y	3.31	3.07	0.28	1.09	2.47	1.00
Skewness	-0.17	0.15	0.11	-0.09	-0.67	-0.03
Correlation with Y	0.99	-0.96	0.81	0.75	0.78	1.00

Table 3.2: Moments and Correlation of US and Model Data

Note: The job-finding probability (P_u) is calculated as in Shimer [2012] but updated to current. Unemployment (U) is the Civilian Unemployment Rate. The average real wage (W) is Average Hourly Earnings of Production and Nonsupervisory Employees which is then deflated using the CPI. Consumption (C) is Real Personal Consumption Expenditures. Investment (I)is the sum of Real Gross Private Domestic Investment and Real Personal Consumption Expenditures on Durables. Output (Y) is Output in the Nonfarm Business Sector. All simulated model moments are averages from 4000 replications of 2765 months. The first 2000 observations are then dropped. The draws of the stochastic driving process are identical between the DMP, HM, and Wage Progression model. The Staggered Nash Bargaining model uses the stochastic process of the original paper [Gertler and Trigari, 2009]. For both the U.S. and all simulated models, all data is measured in logs and HP filtered with a smoothing parameter of 1600.

The standard calibration of a general equilibrium version of the DMP model is unable to address any of the three puzzles I address in this paper. Panel B of Table 3.2 confirms the unemployment volatility and asymmetry puzzles. There is very little amplification of productivity shocks, so the job-finding probability and unemployment are approximately as volatile as output. This is expected since it does not implement new hire wage rigidity as in Gertler and Trigari [2009] or a parameter calibration that results in an elastic labor supply. Additionally, all the variables are nearly symmetric. Finally, Table 3.3 shows the moments regarding the wage-productivity puzzle. The wage and productivity are nearly perfectly correlated and the relative volatility is near unity. This results in an elasticity well above the observed value.

The model of Hagedorn and Manovskii [2013] performs well with regard to the unemployment volatility puzzle and the asymmetry puzzle. However, it is unable to address the wage-productivity puzzle. Panel C of Table 3.2 shows the simulated moments of the labor market variables. For comparability with the original paper, I employ a partial equilibrium version of the HM model.[4] This calibration generates realistic volatility in unemployment; however, it also results in counterfactually large variation in the job-finding rate. This is due to the assumption of a constant job separation rate. Additionally, this calibration produces realistic third moments in unemployment and the job-finding rate. However, this solution to the unemployment volatility puzzle poses problems for the wage-productivity puzzle, as argued in Gertler and Trigari [2009] and shown in Table 3.3. This calibration is unable to fully capture the relationship between wages and productivity. Wage rigidity in existing employment relationships is critical to capturing these dynamics. Although Table 3.3 shows that the HM model matches the elasticity of wages to productivity, it does so by reducing the relative volatility. The correlation between the average wage and productivity

[4] In a general equilibrium extension of the Hagedorn and Manovskii [2008] model, the addition of perfectly mobile capital acts to dampen the volatility in the job-finding rate and unemployment. Additionally, the asymmetry of the labor market variables is reduced relative to the partial equilibrium setting. However the skewness is still directionally correct. Similarly, a partial equilibrium version of the Wage Progression model shows significantly more volatility and asymmetry in the labor market variables at the same calibration.

is still unity.

	Wage, Productivity Moments		
	σ_w/σ_a	$corr(W,A)$	$el(W,A)$
US Economy 1972:Q1 to 2016:Q1	0.83	0.43	0.35
DMP Baseline	0.96	1.00	0.96
Hagedorn and Manovskii	0.33	1.00	0.33
Gertler and Trigari	0.79	0.63	0.50
Wage Progression	0.76	0.41	0.31

Table 3.3: Wage and Productivity Moments

Note: The real wage is Average Hourly Earnings of Production and Nonsupervisory Employees which is then deflated using the CPI. Productivity is measured as Output in the Nonfarm Business Sector per Nonfarm Employee. All data is in logs and HP filtered with a smoothing parameter of 1600.

The Staggered Nash Bargaining framework of Gertler and Trigari [2009] addresses the unemployment volatility puzzle and the wage-productivity puzzle; however, the simulated data is entirely symmetric. Panel D of Table 3.2 shows the moments produced by a replication of the Staggered Nash Bargaining framework of Gertler and Trigari [2009]. The Staggered Nash Bargaining model assumes wage rigidity for newly hired workers and thereby produces unemployment volatility in line with observation. But, it overstates volatility in the job-finding rate in order to produce the observed unemployment volatility. Additionally, this solution to the unemployment volatility puzzle generates entirely symmetric simulated data. As seen in Table 3.3, the Staggered Nash Bargaining model is able to capture the dynamics of the average wage and productivity. Wage rigidity for employed workers is critical to matching the correlation and relative volatility of the average wage and productivity.

The skewness of the Wage Progression model is caused by curvature of the endogenous variables when the variables are viewed as a surface over the state space. If these surfaces are curved, then the direction of the shock affects the magnitude of the response. In contrast when these surfaces are very nearly linear, the direction of the shock has no effect on the magnitude of the response of the endogenous variable. To illustrate, I compare the stochastic equilibrium of the standard DMP model and the Wage Progression model. The stochastic equilibrium is defined as an equilibrium in which the shock repeats itself in perpetuity. The stochastic equilibrium simplifies the endogenous variables from a mapping from the four dimensional state space into a single dimension. Figure 3.2 shows the endogenous variables over the various values of the productivity shock. In both the DMP model and the Wage Progression model, the labor market tightness is nearly linear. The job-finding probability is a concave function of the labor market tightness. The relatively flat response of labor market tightness also means there is very limited skewness.[5] Conversely, the steeper response in the Wage Progression model implies that the job-finding probability, vacancies, and unemployment all feature curvature.

In a general equilibrium setting with a multidimensional state space, curvature in the productivity shock direction is not the only determinant of higher order moments. Curvature in the other dimensions of the state space also contributes to these moments. Therefore, impulse responses more completely characterize the response of the endogenous variables relative to the stochastic equilibrium. Comparing the impulse responses for a positive and negative shock make the asymmetry explicit, as the magnitude of the response will differ based on the direction of the shock. Figures 3.3, 3.4, 3.5, and 3.6 depict impulse responses for the DMP model, the HM model, the Staggered Nash Bargaining model, and the Wage Progression model, respectively. To form the impulse response, I begin with the endogenous variables in the steady state. I then compare the paths of the variables resulting from beginning at two different points in the shock space. For the positive shock, I assume the technology shock is two points above the median or steady state value. This implies a shock of 1.5 standard deviations.[6] Finally, I take the difference between the two resulting

[5]This relationship between volatility and skewness is pointed out in a partial equilibrium setting by Ferraro [2013].

[6]This is slightly different for the Gertler and Trigari [2009] impulse responses. Since I recreate the model, I also adopt their calibration of the shock process. This slightly differs from my calibrated values, as I have extended the productivity through Q1 2016. In addition, the impulse responses shown are for a shock of one standard deviation.

paths and average across 50,000 replications. All responses to a positive shock are reflected around the horizontal axis for ease of comparison.

Figure 3.6 confirms that the Wage Progression model is able to resolve the three puzzles identified earlier. First, the magnitude of the responses of the job-finding rate and unemployment is relatively large when compared to the DMP model. The source of this amplification is the same as in Hagedorn and Manovskii [2008]. Specifically, firm profits are relatively small and quite sensitive to changes in the productivity shock; therefore, the vacancy creation incentive is also sensitive. The labor market variables exhibit significant asymmetry in the impulse responses. Both unemployment and investment show a stronger response to a negative shock than to a positive shock. This implies negative skewness for investment and positive skewness for unemployment. As discussed above, this skewness is closely related to the amplification of shocks. Additionally, the increased curvature of the vacancy creation surface spills over into the savings decision. This results in significant asymmetry in the investment response. Finally, the response of the average wage is consistent with an explanation of the wage-productivity puzzle. The introduction of wage rigidity for employed workers results in hump-shaped dynamics that stand in contrast to the DMP and HM models. The hump-shaped dynamics of the average wage stands in contrast to the behavior of the productivity shock. Therefore, it is consistent with the moderate correlation between wages and productivity. Further, the endogenously varying level of this wage rigidity introduces slight asymmetry that results in a positively skewed average wage. The response of the average wage shows a trade-off between amplification and persistence. The response to a positive shock is larger in magnitude while the response to a negative shock is more persistent.

An examination of the impulse responses from the other models discussed in this paper shows that each of them fails to resolve at least one of the three puzzles. Figure 3.3 shows that the standard DMP model is unable to address any of the three puzzles. The response of the labor market variables is limited in magnitude which is consistent with the low amplification of shocks and the low volatility in the job-finding rate. Also, these responses are symmetric to a positive and negative shock. Finally, the average wage response closely follows the path of the shock. Therefore, contrary to observation, the correlation of the wage and productivity is near one. Figure 3.4 shows the impulse responses for the HM model. The response of the wage in the HM model is too strongly correlated with the shock; therefore, it is unable to address the wage-productivity puzzle. Finally, Figure 3.5 shows the impulse responses for the Staggered Nash Bargaining model. The impulse responses of the Staggered Nash Bargaining model are identical regardless of shock direction. The Wage Progression model is able to address the three issues simultaneously.

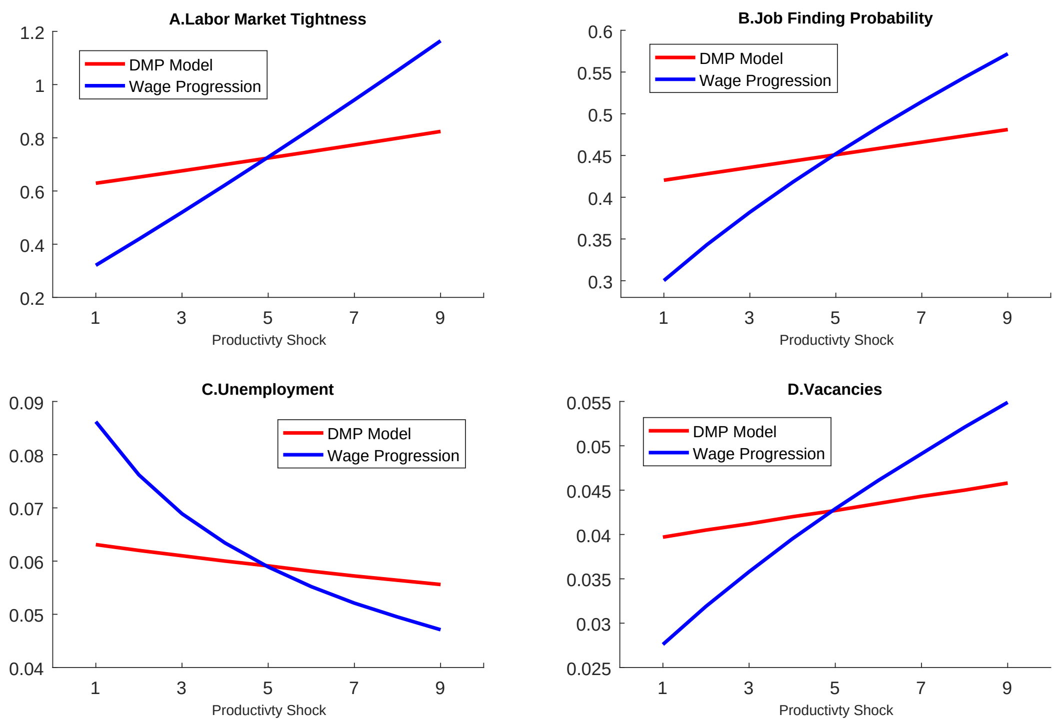

Figure 3.2: Stochastic Equilibrium of the DMP and Wage Progression Models

Note: The stochastic equilibrium is defined as an equilibrium in which the shock repeats itself in perpetuity. The horizontal axis indexes the possible values of the productivity shock.

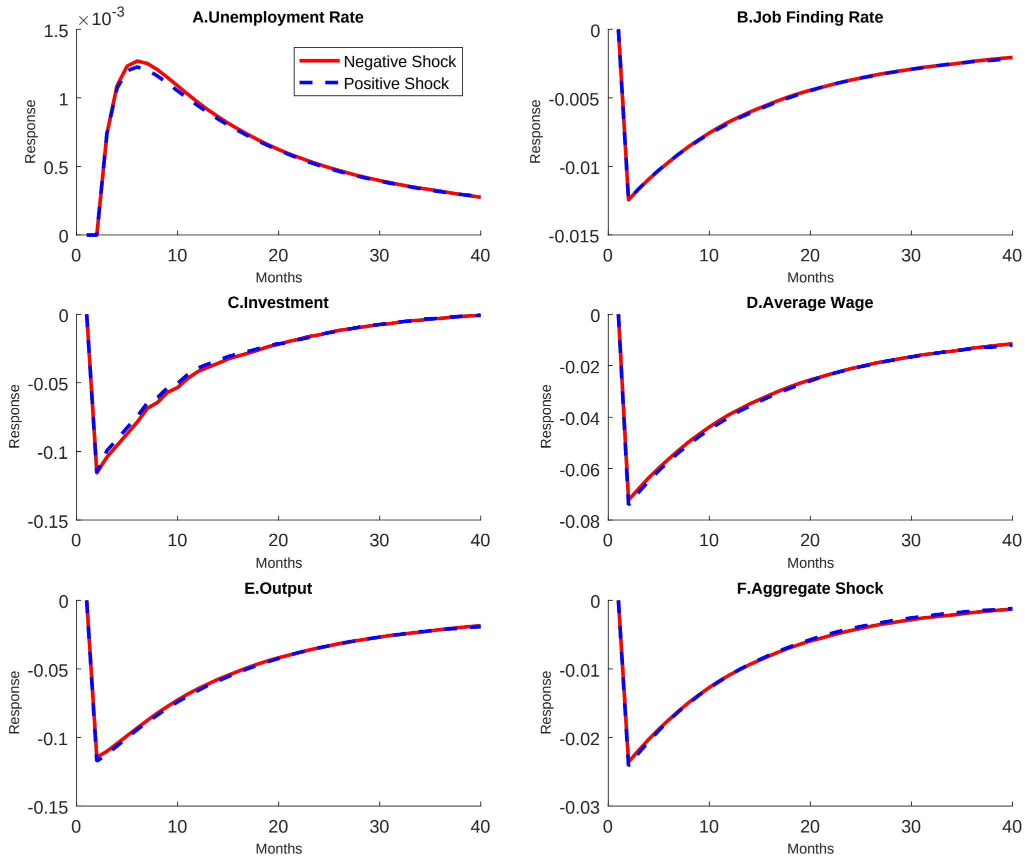

Figure 3.3: Impulse Responses of the DMP Model

Note: Impulse responses are formed by beginning with the endogenous variables at the steady state. I then compare the paths of the variables resulting from beginning at two different points in the shock space. Finally, I take the difference between the two resulting paths and average across 50,000 replications. Responses shown are generated by a shock of 1.5 standard deviations. All responses to a positive shock are reflected around the horizontal axis for ease of comparison.

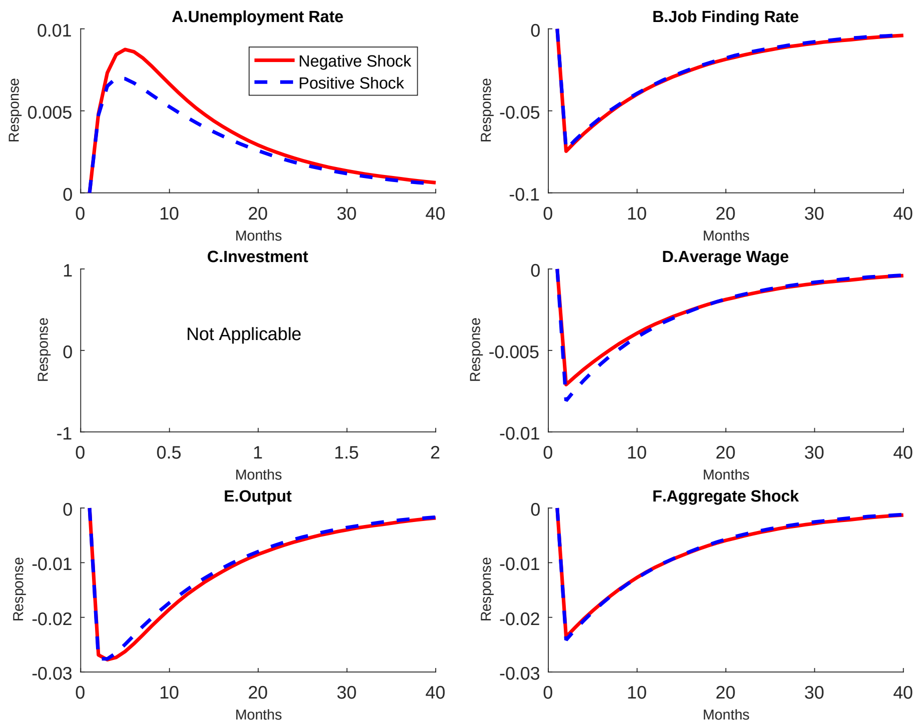

Figure 3.4: Impulse Responses of the HM Model

Note: Impulse responses are formed by beginning with the endogenous variables at the steady state. I then compare the paths of the variables resulting from beginning at two different points in the shock space. Finally, I take the difference between the two resulting paths and average across 50,000 replications. Responses shown are generated by a shock of 1.5 standard deviations. All responses to a positive shock are reflected around the horizontal axis for ease of comparison.

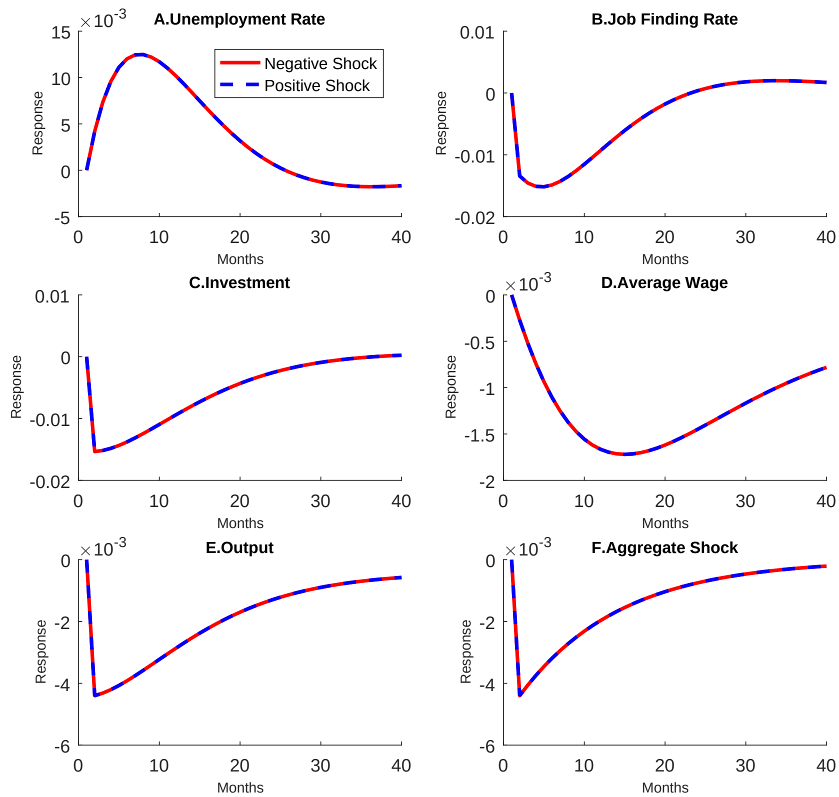

Figure 3.5: Impulse Responses of the GT Model

Note: Impulse responses are formed by beginning with the endogenous variables at the steady state. I then compare the paths of the variables resulting from beginning at two different points in the shock space. Finally, I take the difference between the two resulting paths and average across 50,000 replications. Responses shown are generated by a shock of 1.5 standard deviations. All responses to a positive shock are reflected around the horizontal axis for ease of comparison.

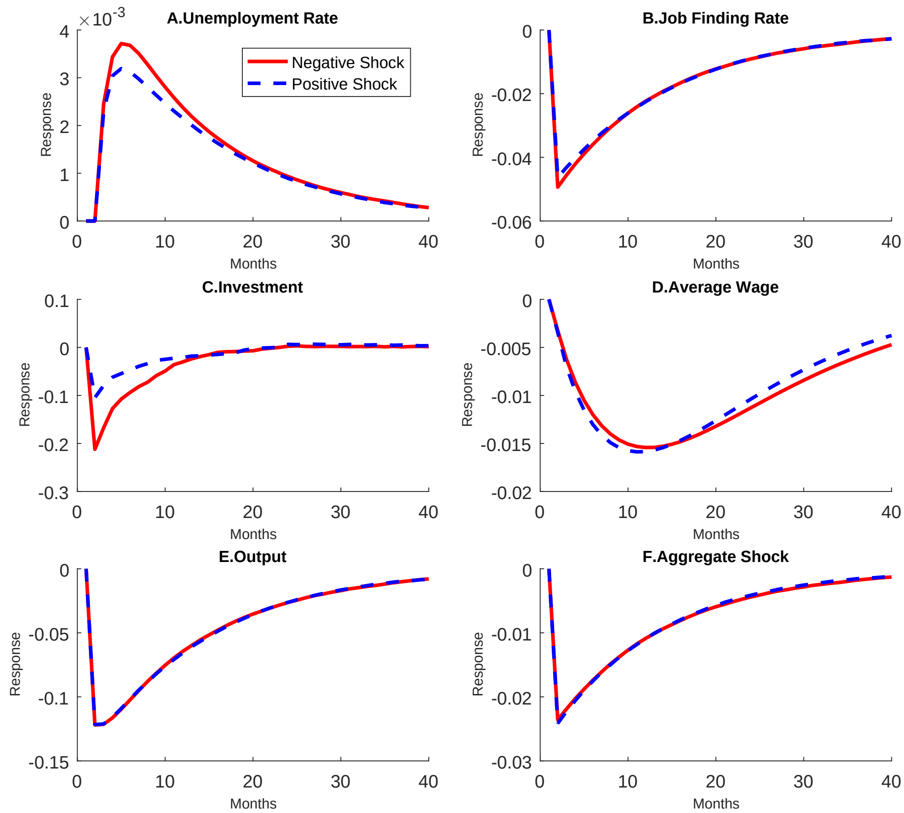

Figure 3.6: Impulse Responses of the Wage Progression Model

Note: Impulse responses are formed by beginning with the endogenous variables at the steady state. I then compare the paths of the variables resulting from beginning at two different points in the shock space. Finally, I take the difference between the two resulting paths and average across 50,000 replications. Responses shown are generated by a shock of 1.5 standard deviations. All responses to a positive shock are reflected around the horizontal axis for ease of comparison.

3.5 Conclusions

This paper develops a general equilibrium search model of the labor market that incorporates wage rigidity and on-the-job search. The calibrated Wage Progression model developed in this paper is able to simultaneously solve the unemployment volatility puzzle, the asymmetry puzzle, and the wage-productivity puzzle.

The Wage Progression model simultaneously addresses the three puzzles via two channels. First, I develop a new wage determination mechanism that links wage rigidity and on-the-job search. This new mechanism results in time-varying wage rigidity for employed workers, because in the Wage Progression model employed workers renegotiate the wage upon forming a match with a new firm. I assume that there is no wage rigidity for workers flowing from unemployment to employment. This is consistent with the bulk of empirical findings about wage rigidity for newly hired workers. Second, I show that when calibrated with a high replacement ratio, the simulated moments of the model are consistent with the volatility and asymmetry of the labor market. The slope and curvature of the vacancy creation surface over the state space is key to this finding.

The Wage Progression model is also able to account for more than 75% of the observed negative skewness of investment. Unemployment and vacancy creation costs are key components of aggregate household income. Therefore, the dynamics of the labor market are critical to correctly modeling fluctuations in household income and the subsequent savings decision.

Future work based on this paper could progress in several directions. First, inserting the Wage Progression framework into a monetary DSGE framework would allow for Bayesian estimation. This would add additional rigor to the comparison with other models. Second, the framework developed in this paper could be expanded to endogenize the job destruction margin. Endogenous separation models have the property that employees and firms mutually agree to dissolve matches. The Wage Progression framework could break this mutual agreement and thereby make a distinction between quits and fires. The Wage Progression model developed in this paper advances the literature on search models by producing more realistic labor market moments and investment dynamics.

Chapter 4

Inequality and the Changing Cyclicality of the U.S. Real Wage

In this chapter I document a significant change regarding the cyclicality of the average real wage in the U.S. Then I consider rising wage inequality as a possible cause for this change. I show that the cyclicality of the real wage in the US has recently reversed. Before 2000, the wage is procyclical; while after 2000 the average wage is countercyclical. The magnitude of the reversal is large, and the reversal is robust to detrending method. There is a substantial literature regarding the cyclical behavior of the real wage, however no existing work has documented this particular change in the cyclical behavior of the real wage. [1] This chapter contributes to this literature by showing the the real wage changed from procyclical to countercyclical around the year 2000 and then considering rising wage inequality as a potential cause.

The reversal that occurred in 2000 is not the first such reversal to be observed in the U.S. Hanes [1996] find that the wage changed from countercyclical in the interwar period to procyclical post war. Basu and Taylor [1999]confirms a similar finding in a set of 15 countries, including the U.S. They divide the time frame of 1870 to 1990 into four monetary regimes. Under the first two regimes, lasting from 1870 to 1939 excluding World War I, they find the real wage was acyclical. Under the final two regimes, Bretton Woods and floating exchange rates, they find that the real wage was moderately procyclical. Having observed this change, Huang et al. [2004] develops a DSGE model that model endogenously results in changes to the real wage cyclicality. These changes are caused by nominal wage rigidity, price rigidity, and the increasingly complex input-output structure of the economy. However, no existing work notes the reversal occurring in 2000.

Having established that a reversal in the cyclical behavior in the wage is an empirical

[1] See Abraham and Haltiwanger [1995] for an excellent survey

regularity, I then turn to possible causes. I find that increasing wage inequality is able to explain a large proportion of the observed change in correlation. Wage cyclicality is linked to inequality through a composition effect. Employment of low wage workers is significantly more sensitive to the business cycle. Therefore, the composition of the stock of employed workers effects the average wage. Stockman [1983] is the first work to conjecture that cyclical variation in the composition of the employed worker pool could give rise to a composition effect. Stockman [1983] attributes the observed low cyclicality of the average wage to aggregation bias. The author thereby concludes that this correlation should not be considered an empirical target for theoretical models. Solon et al. [1994] also makes note of this compositional effect. The paper shows disaggregated real wages are more procyclical than the average wage once worker attributes are controlled for. However, other papers conclude that the compositional changes have only a modest effect on the cyclicality of the real wage [Bils, 1985, Keane et al., 1988]. This chapter extends this idea by considering that the compositional effect has been strengthening as wage inequality has risen. Further, I show the composition effect is responsible for a countercyclical average wage in recent years. I use a structural vector autoregression approach to construct counterfactuals. These counterfactuals simulate the endogenous variables under different inequality scenarios. I find that increasing wage inequality can explain approximately 70% of the observed change in wage cyclicality.

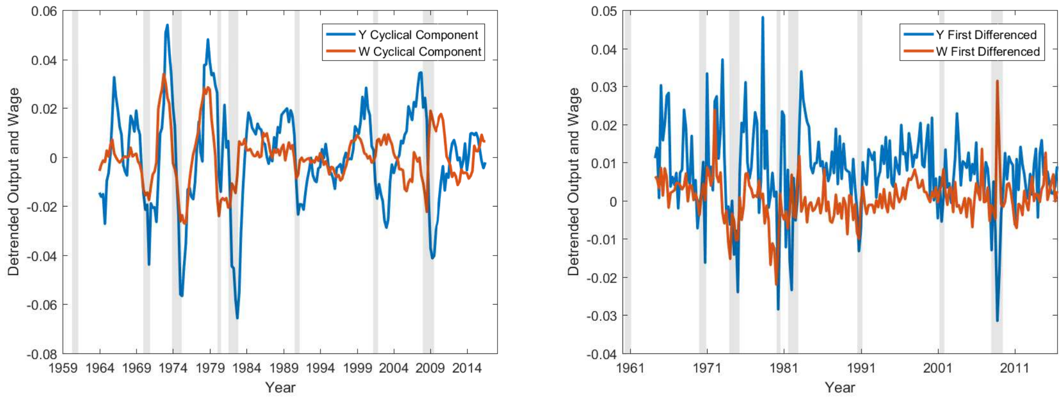

Figure 4.1: US Real Average Hourly Earnings and Real Output

Figure 4.1 shows the detrended real wage and output series from Q1:1964 to Q3:2016. Quarterly real output is from the Nonfarm Business Sector. The wage is the average hourly earnings of production and non-supervisory employees.[2] The wage is deflated using the CPI. All variables are measured in logs and detrended either with the HP filter or by first differencing. Visually, the correlation appears to switch sometime in the late 1990's or early 2000's. Table 4.1 shows the magnitude of the correlation reversal assuming a break point of Q1 2000. With both detrending methods, the wage is moderately procyclical prior to 2000 and moderately countercyclical after 2000. This is significant as Abraham and Haltiwanger [1995] shows that the cyclical nature of the real wage can be sensitive to the detrending method selected. The magnitude of the reversal is striking, as the pre-2000 and post-2000 correlations are roughly equivalent but opposite in sign.

Corr(Y,W)	HP Filtered	First Differenced
Q1 1963 to Q3 2016	0.34	0.15
Pre 2000 Sample	0.62	0.37
Post 2000 Sample	-0.66	-0.41

Table 4.1: Correlation of Wage and Output

[2]I also consider other measures of labor compensation. Average Real Total Compensation (BEA series A576RC1) per employee shows a similar pattern of reduced correlation. However, the change in cyclicality is not as large and total compensation remains procyclical after the year 2000.

Next, I fit a bivariate vector autoregression to the HP filtered series. This will allow for further formalization of the empirical observation. Estimation is by least squares. I use various information criteria and top down sequential testing to select the lag order. Both methods suggest that four lags are sufficient to capture the data generating process. In addition to the full data sample, I fit the model to the data for the period before the year 2000 and post 2000. Table 4.2 contains the estimated coefficients and test statistics. It is worth while to compare the estimated coefficients from the two sub-samples. The most important difference between the two is the coefficient on the once-lagged real wage in the Output equation. In the pre-2000 sample, this coefficient is positive and significant. However in the post-2000 sample, the coefficient is negative and no longer significant. Cross variable coefficients such as this are particularly of interest since they directly effect cross variable moments, such as correlations.

	Covariates							
	Y_{-1}	W_{-1}	Y_{-2}	W_{-2}	Y_{-3}	W_{-3}	Y_{-4}	W_{-4}
A. Complete Sample								
Y_t	1.00	0.33	-0.08	-0.09	-0.08	-0.08	-0.08	0.07
$T-stat$	14.30	2.39	-0.79	-0.41	-0.80	-0.40	-1.19	0.49
W_t	-0.06	1.17	0.11	-0.38	-0.07	0.35	-0.01	-0.33
$T-stat$	-1.85	17.38	2.33	-3.71	-1.43	3.42	-0.35	-4.77
B. Pre 2000								
Y_t	0.86	0.67	-0.05	-0.29	-0.11	-0.12	-0.07	0.28
$T-stat$	9.87	3.49	-0.43	-0.97	-0.94	-0.43	-0.79	1.34
W_t	-0.07	1.21	0.08	-0.40	-0.05	0.48	-0.01	-0.44
$T-stat$	-1.97	15.18	1.74	-3.19	-1.05	3.86	-0.13	-5.13
C. Post 2000								
Y_t	1.11	-0.15	-0.16	0.25	0.03	-0.05	-0.21	-0.17
$T-stat$	7.71	-0.72	-0.79	0.92	0.14	-0.18	-1.48	-0.84
W_t	-0.20	0.89	0.18	-0.32	-0.15	0.15	0.06	-0.11
$T-stat$	-2.01	6.13	1.27	-1.69	-1.06	0.83	0.63	-0.80

Table 4.2: Bivariate VAR(4) Coefficients

After fitting the empirical model to the sub-samples, I am able to more rigorously document that there has been a change between the pre and post-2000 periods. Using the Chow test, I am able to confirm a structural break in the data. Table 4.3 contains the test statistics and critical values. The Chow test rejects the null hypothesis of constant parameters cross the two time periods. This suggests a structural break in both the output and wage equation, but most prominently in the output equation. This is consistent with the observation above regarding the lagged wage variable in the output equation.

Chow Tests	Test Stat	Critical Value
Y equation	3.57	1.98
W equation	2.67	1.98

Table 4.3: Chow Test Results

To further strengthen this empirical observation, I also consider tests of multiple structural changes at unknown dates as in Bai and Perron [1998] and Bai and Perron [2003]. I find strong evidence of a structural break occurring around the year 2000. I consider a test of partial structural change, where I assume that the coefficients for lags of the dependent variable are unchanged across regimes. As I am particularly interested in the correlation between wages and output, the coefficients on the lags of wages in the output equation and vice versa are of interest. Testing for partial structural change is consistent with the idea that it is the cross terms in the reduced form regressions that are changing. The estimated coefficients for lags of the dependent variable are assumed to remain the same. Also, comparing the reduced form estimates for the pre-2000 sample and the post-2000 sample reveals that the most significant change is the one period lag of wages in the output equation. The Bai and Perron [2003] test on the output equation indicates a structural break occurring in the Q2:2000. The 95% confidence interval for this break is Q3:1996 to Q2:2004. SupF tests for a fixed number of breaks and UD max tests against an unknown number of breaks both indicate that there is a structural break at the 1% level. However, information criteria approaches are less supportive. Both the BIC and LWZ criteria select a model with no breaks. The test for partial structural change on the wage equation is less conclusive. SupF and UD max tests indicate the possibility of a single break; however, the test values are only marginally significant at the 10% level. The 95% confidence interval for this first break is Q1:1974 to Q3: 1986. However, a structural change in either equation is sufficient to alter the correlation between output and wages. Therefore the structural break in the output equation is sufficient to suggest a change in the relationship between average real wage and output.

4.2 Inequality as a Cause?

Having established that a significant change occurred in the relationship between wages and output around 2000, I now turn to possible explanations. One possible explanation for the change in the cyclicality of the real wage is increasing wage inequality. The hours worked and employment of low wage workers are substantially more cyclical than that of high wage workers. Therefore the wage distribution of wages effect the correlation of the average wage with output. If sufficiently many low wage workers are laid off during a recession, then the average wage may rise. If sufficiently strong, this composition effect could create a countercyclical aggregate wage. I hypothesize that rising wage inequality has strengthened this composition effect over time. This continued increase in the level of wage inequality thereby reversed the cyclicality of the real wage wage. I use a structural vector autoregression and historical decompositions to test this hypothesis.

4.2.1 Tri-variate VAR

To assess the effect of the strengthening compositional effect, I fit a vector autoregression to output, wage, and inequality data. I use the Gini coefficient to measure income inequality. Wage inequality is a low frequency variable and measured annually. Rather than use mixed frequency methods, I use cubic spline interpolation to increase the frequency of inequality data to match the frequency of the other time series.

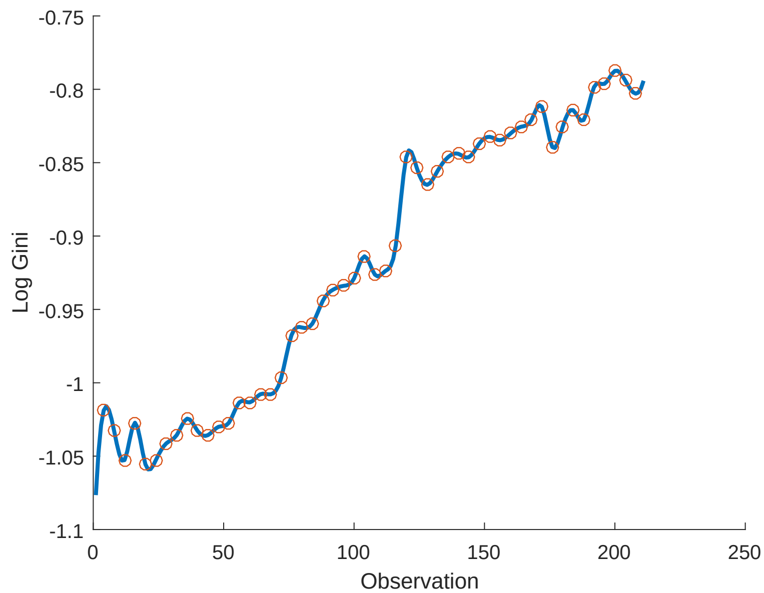

Figure 4.2: Annual and Quarterly Interpolated Inequality Data

The detrended output and wage series are stationary by construction, however the inequality data is not[3]. The augmented Dickey-Fuller test indicates that the inequality data is trend stationary. Therefore whenever including inequality data, I will include a trend term as well as a constant. I rely on various information criterion as well as 'top down' sequential testing to select the lag order. The parameter estimates and t-statistics of the reduced form model are in Table 5.5.

The estimated models fit the data well. In particular, the implied unconditional moments of the fitted model match the observed moments. Critically, the VARs are able to reproduce the unconditional correlation of wages and output observed in the data. This is true for a model fit to the entire sample as well as models fit to the pre-2000 and post-2000 periods individually. The actual correlations and the correlations implied by the companion matrices of the VARs are listed in table 4.4.

[3]I also considered leaving all variables in levels. However, the Johansen test for cointegration indicated that there was only a single cointegrating relationship between the three variables.

Corr(Y,W)	U.S. Data	Implied by VAR
Q1 1963 to Q3 2016	0.344	0.336
Pre 2000 Sample	0.620	0.623
Post 2000 Sample	-0.660	-0.666

Table 4.4: Actual Wage, Output Correlations vs. Model Implied Correlations

4.2.2 Identification and Historical Decompositions

I use short run restrictions to identify the vector autoregression. The recursive ordering I employ is the average real wage, followed by Gini coefficient, and finally output. This ordering is supported by three assumptions. First, if wages are rigid within quarters then the real wage should come before output. Second, I assume that total income does not contemporaneously effect the spread of the income distribution. This implies that the Gini coefficient should be before output. Finally, I assume that the current average wage is not effected by inequality. This implies that wages should be ordered prior to inequality. For completeness, impulse responses are shown in Appendix tables 5.1 through 5.3. However, the focus of my casual analysis is not on impulse responses but rather on counterfactual historical decompositions.

Having identified the structural shocks, I can now construct counterfactual time series for the system. In the first counterfactual, I examine the correlation between wages and output as if there were no inequality shocks. This corresponds to a textbook historical decomposition. In the second counterfactual, I remove the disturbances to inequality and also its trend at various points in time. What follows is based on setting the inequality trend coefficient to zero from the year 2000 onward. The results are robust to the choice of date at which the trend is flattened, and the results are slightly strengthened if an earlier date is chosen.

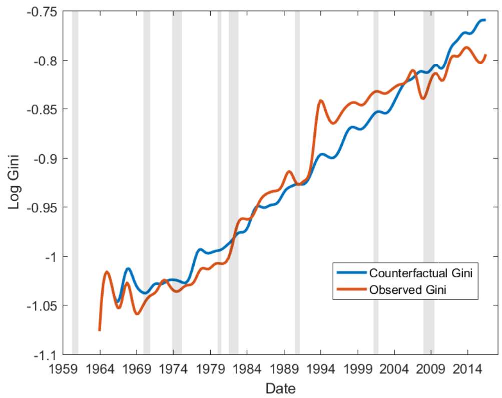

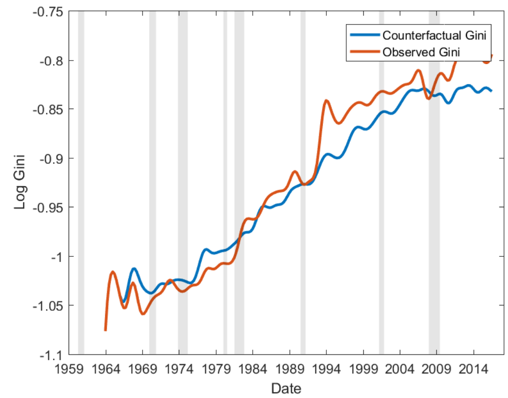

Figure 4.3: Counterfactual Inequality Time Series

Figure 4.3 depicts the two counterfactual inequality time series as well as the observed data. In the first counterfactual, inequality continues to increase throughout the time frame. This necessitates the second counterfactual, as the composition effect I wish to examine is driven by the level of wage inequality. Therefore, I include a trend break in the second counterfactual. This has the effect of flattening the level of wage inequality, which is consistent with the original hypothesis that the rising level of wage inequality drives the correlation reversal. Regarding the other two time series, the majority of the difference between the observed data and the counterfactual data occurs in the output series. This is consistent with the reduced form estimates, as they primarily show inequality to be significant in the output equation.

Corr(Y,W)	U.S. Data	No Inequality Shocks	No Shocks and Trend Break
Q1 1963 to Q3 2016	0.344	0.37	0.55
Pre 2000 Sample	0.620	0.61	0.61
Post 2000 Sample	-0.660	-0.72	0.25

Table 4.5: Actual and Counterfactual Correlations

Table 4.5 contains the correlations between the average real wage and output observed in the actual data and in the two counterfactual scenarios. Under the first counterfactual scenario, the correlations have a similar pattern to the observed data. The real wage is procyclical pre-2000 and then counter cyclical post-2000. Since the first counterfactual is not sufficient to generate a procyclical wage, shocks to inequality are not the cause of the correlation reversal. However, the correlations are significantly different in the second counterfactual scenario. With a trend break, the wage remains procyclical in the post 2000 period. If inequality had stopped rising in the year 2000 then the real wage would have remained procyclical. The post-2000 counterfactual correlation is able to account for 70% of the absolute value of the observed change in correlation. While this is certainly a significant proportion of the observed change, the decline in the strength of procyclicality suggests additional factors may be at work as well.

4.3 Conclusion

This chapter shows that there has been a significant change in the cyclical behavior of the real wage. Visual inspection indicates that the real wage switches from being procyclical to counter cyclical around the year 2000. Tests of structural change verify this conclusion. This change has not been noted or explained in the existing literature. Having established that this observation is robust, this chapter then considers increasing wage inequality as a potential cause. I use short run restrictions to identify a structural vector autoregression. With the structural shocks identified, I construct two counterfactuals. The first counterfactual examines the cyclicality of the wage in the absence of inequality shocks. This counterfactual is unable to explain the correlation reversal. The second counterfactual considers the case of no inequality shocks as well as flattening the persistent positive trend in inequality. If inequality shocks stopped and the positive trend in inequality broke in the year 2000, then the real wage would have remained procyclical.

Chapter 5

Conclusion

Labor markets have become increasingly an area of focus in macroeconomics since the Great Recession. This dissertation extends the literature on labor markets by emphasizing the importance of higher order moments. I consider higher order moments in several ways. First, I develop theoretical models in which second order moments of aggregate distributions influence the macroeconomic outcomes. I also seek to explain observed higher order moments of empirical data.

In the second chapter, I develop a model with heterogeneous agents in which the second moment of the wealth distribution determines macroeconomic aggregates. This represents a move beyond representative agent models, as agents make differential labor supply decisions based on the initial level of wealth. Having developed a theoretical model, I then use cross country data to find supporting evidence.

In the third chapter, I develop a theoretical model that simultaneously addresses three gaps between observed data and simulated data from existing models in the literature. The first gap is the well known unemployment volatility puzzle. Second, theoretical models generate symmetric unemployment and investment time series. In reality these are significantly skewed. Finally, labor search models have struggled to capture the relationship between wages and productivity. To address these issues, I develop and calibrate a model featuring non-linearity that results in realistic volatility and asymmetry of the labor markets. The improved labor market fluctuations result in more realistic dynamics for investment.

Finally, I document a significant change in the cyclicality of the average real wage in the U.S. From 1964 to 2000 the real wage is strongly procyclical. From 2000 to 2016 the real wage is strongly countercyclical. This empirical observation has not yet been documented or explored by other work. Tests that endogenously determine the dates of structural breaks identify the year 2000. I hypothesize that increasing wage inequality may cause this change. Using a structural vector autoregression approach, I then construct counterfactuals that show increases in wage inequality account for around 70% of the reversal in the correlation between wages and output.

References

Katharine G. Abraham and John C. Haltiwanger. Real Wages and the Business Cycle. *Journal of Economic Literature*, 33(3):1215–1264, 1995.

Daron Acemoglu and James A. Robinson. The political economy of the kuznets curve. *Review of Development Economics*, 6(2):183–203, 2002.

Aghion and P. Bolton. A theory of trickle-down growth and development. *lead article in the Review of Economic Studies*, 64(2):151–172, 1997.

Pedro H. Albuquerque. Inequality-driven growth: Unveiling aggregation effects in growth equations. Econometric Society 2004 Far Eastern Meetings 769, Econometric Society, 2004.

Alberto Alesina and Dani Rodrik. Distributive politics and economic growth. *The Quarterly Journal of Economics*, 109(2):465–490, May 1994.

Sudhir Anand and S. M. R. Kanbur. Inequality and development a critique. *Journal of Development Economics*, 41(1):19–43, 1993a.

Sudhir Anand and S. M. R. Kanbur. The kuznets process and the inequality–development relationship. *Journal of Development Economics*, 40(1):25–52, 1993b.

David Andolfatto. Evidence and Theory on the Cyclical Asymmetry in Unemployment Rate Fluctuations. *Canadian Journal of Economics*, 30(3):709–21, 1997.

A. B. Atkinson and A. Brandolini. Promise and pitfalls in the use of 'secondary' data -sets: Income inequality in OECD countries. Paper 379, Banca Italia - Servizio di Studi, 2000.

Jushan Bai and Serena Ng. Tests for Skewness, Kurtosis, and Normality for Time Series Data. *Journal of Business & Economic Statistics*, 23:49–60, 2005.

Jushan Bai and Pierre Perron. Estimating and Testing Linear Models with Multiple Structural Changes. *Econometrica*, 66(1):47–78, 1998.

Jushan Bai and Pierre Perron. Computation and analysis of multiple structural change models. *Journal of Applied Econometrics*, 18(1):1–22, January 2003.

Abhijit V. Banerjee and Esther Duflo. Inequality and growth: What can the data say? NBER Working Paper 7793, National Bureau of Economic Research, Inc, 2000.

Robert J. Barro. Inequality, growth, and investment. NBER Working Paper 7038, National Bureau of Economic Research, Inc, 1999.

Robert J. Barro. Inequality and growth in a panel of countries. *Journal of Economic Growth*, 5(1):5–32, 2000.

Robert J. Barro. Inequality and growth revisited. Working Papers on Regional Economic Integration 11, Asian Development Bank, 2008.

Susanto Basu and Alan M. Taylor. Business Cycles in International Historical Perspective. *Journal of Economic Perspectives*, 13(2):45–68, June 1999.

Jorge Belaire-Franch and Amado Peiro. Conditional and Unconditional Asymmetry in U.S. Macroeconomic Time Series. *Studies in Nonlinear Dynamics & Econometrics*, 7 (1), 2003. ISSN 1558-3708. doi: 10.2202/1558-3708.1108.

Roland Benabou. Inequality and growth. NBER chapters, National Bureau of Economic Research, Inc, 1996.

Mark J. Bils. Real Wages over the Business Cycle: Evidence from Panel Data. *Journal of Political Economy*, 93(4):666–689, 1985.

Joaquin Blaum. Wealth inequality and the losses from financial frictions. 2012 Meeting Paper 1077, Society for Economic Dynamics, 2012.

Pierre Cahuc, Fabien Postel-Vinay, and Jean-Marc Robin. Wage Bargaining with On-the-Job Search: Theory and Evidence. *Econometrica*, 74(2):323–364, 2006. ISSN 0012-9682.

John Y. Campbell and Angus Deaton. Why is consumption so smooth? *Review of Economic Studies*, 56(3):357–73, 1989.

James S. Costain and Michael Reiter. Business cycles, unemployment insurance, and the calibration of matching models. *Journal of Economic Dynamics and Control*, 32(4): 1120–1155, 2008.

David de la Croix and Matthias Doepke. Inequality and growth: Why differential fertility matters. *American Economic Review*, 93(4):1091–1113, 2003.

Klaus Deininger and Lyn Squire. A new data set measuring income inequality. CEMA Working Paper 512, China Economics and Management Academy, Central University of Finance and Economics, 1996.

Klaus Deininger and Lyn Squire. New ways of looking at old issues: inequality and growth. *Journal of Development Economics*, 57(2):259–287, 1998.

Christophe Ehrhart. The effects of inequality on growth: a survey of the theoretical and empirical literature. Working Paper 107, ECINEQ, Society for the Study of Economic Inequality, 2009.

Domenico Ferraro. The Asymmetic Cyclical Behavior of the U.S. Labor Market. November 2013.

Marjorie A. Flavin. The adjustment of consumption to changing expectations about future income. *Journal of Political Economy*, 89(5):974–1009, October 1981.

Christopher J. Flinn. Minimum Wage Effects on Labor Market Outcomes under Search, Matching, and Endogenous Contact Rates. *Econometrica*, 74(4):1013–1062, 2006.

Kristin J. Forbes. A reassessment of the relationship between inequality and growth. *The American Economic Review*, 90(4):869–887, September 2000.

Shigeru Fujita and Garey Ramey. The Cyclicality Of Separation And Job Finding Rates. *International Economic Review*, 50(2):415–430, 2009.

Shigeru Fujita and Garey Ramey. Exogenous versus Endogenous Separation. *American Economic Journal: Macroeconomics*, 4(4):68–93, 2012.

Mark Gertler and Antonella Trigari. Unemployment Fluctuations with Staggered Nash Wage Bargaining. *Journal of Political Economy*, 117(1):38–86, 2009.

Mark Gertler, Luca Sala, and Antonella Trigari. An Estimated Monetary DSGE Model with Unemployment and Staggered Nominal Wage Bargaining. *Journal of Money, Credit and Banking*, 40(8):1713–1764, 2008.

Christian Haefke, Marcus Sonntag, and Thijs van Rens. Wage rigidity and job creation. *Journal of Monetary Economics*, 60(8):887–899, 2013.

Marcus Hagedorn and Iourii Manovskii. The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited. *American Economic Review*, 98(4): 1692–1706, 2008.

Marcus Hagedorn and Iourii Manovskii. Job Selection and Wages over the Business Cycle. *American Economic Review*, 103(2):771–803, 2013.

Robert E. Hall. Employment Efficiency and Sticky Wages: Evidence from Flows in the Labor Market. *The Review of Economics and Statistics*, 87(3):397–407, 2005.

Robert E. Hall and Paul R. Milgrom. The Limited Influence of Unemployment on the Wage Bargain. *American Economic Review*, 98(4):1653–74, 2008. doi: 10.1257/aer.98.4.1653.

Christopher Hanes. Changes in the Cyclical Behavior of Real Wage Rates, 1870â1990. *The Journal of Economic History*, 56(4):837–861, December 1996.

Kevin X. D. Huang, Zheng Liu, and Louis Phaneuf. Why Does the Cyclical Behavior of Real Wages Change Over Time? *American Economic Review*, 94(4):836–856, 2004.

Boyan Jovanovic. Asymmetric Cycles. *Review of Economic Studies*, 73(1):145–162, 2006.

Michael Keane, Robert Moffitt, and David Runkle. Real Wages over the Business Cycle: Estimating the Impact of Heterogeneity with Micro Data. *Journal of Political Economy*, 96(6):1232–66, 1988.

Marianna Kudlyak. The cyclicality of the user cost of labor. *Journal of Monetary Economics*, 68:53–67, November 2014. ISSN 0304-3932. doi: 10.1016/j.jmoneco.2014.07.007.

Simon Kuznets. Economic growth and income inequality. *The American Economic Review*, pages 1â–28, 1955.

Hongyi Li and Heng-fu Zou. Income inequality is not harmful for growth: Theory and evidence. *Review of Development Economics*, 2(3):318–334, October 1998. ISSN 1467-9361.

Annamaria Lusardi and Olivia Mitchell. Financial literacy and retirement planning: New evidence from the rand american life panel. Working Paper wp157, University of Michigan, Michigan Retirement Research Center, 2007.

Annamaria Lusardi and Olivia S. Mitchell. Financial literacy and retirement planning in the united states. *Journal of Pension Economics and Finance*, 10(04):509–525, 2011.

Annamaria Lusardi, Pierre-Carl Michaud, and Olivia S. Mitchell. Optimal financial knowledge and wealth inequality. NBER Working Paper 18669, National Bureau of Economic Research, Inc, 2013.

Alisdair McKay and Ricardo Reis. The brevity and violence of contractions and expansions. *Journal of Monetary Economics*, 55(4):738–751, 2008.

Guido Menzio and Shouyong Shi. Block recursive equilibria for stochastic models of search on the job. *Journal of Economic Theory*, 145(4):1453–1494, 2010.

Monika Merz. Search in the labor market and the real business cycle. *Journal of Monetary Economics*, 36(2):269–300, 1995.

Roberto Perotti. Growth, income distribution, and democracy: What the data say. *Journal of Economic Growth*, 1(2):149–187, June 1996. ISSN 1381-4338, 1573-7020.

Torsten Persson and Guido Tabellini. Is inequality harmful for growth? *American Economic Review*, 84(3):600–621, 1994.

Christopher A. Pissarides. The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer? *Econometrica*, 77(5):1339–1369, September 2009. ISSN 1468-0262. doi: 10.3982/ECTA7562.

Ricardo Reis. Inattentive consumers. *Journal of Monetary Economics*, 53(8):1761–1800, 2006.

Philip Rothman. Further evidence on the asymmetric behavior of unemployment rates over the business cycle. *Journal of Macroeconomics*, 13(2):291–298, 1991.

Robert Shimer. The Cyclical Behavior of Equilibrium Unemployment and Vacancies. *American Economic Review*, 95(1):25–49, 2005.

Robert Shimer. Reassessing the Ins and Outs of Unemployment. *Review of Economic Dynamics*, 15(2):127–148, 2012.

Christopher A. Sims. Rational inattention: Beyond the linear-quadratic case. *American Economic Review*, 96(2):158–163, 2006.

Gary Solon, Robert Barsky, and Jonathan A. Parker. Measuring the Cyclicality of Real Wages: How Important is Composition Bias. *The Quarterly Journal of Economics*, 109 (1):1–25, 1994.

Alan Stockman. Aggregation Bias and the Cyclical Behavior of Real Wages. *unpublished*, July 1983.

George Tauchen. Finite state markov-chain approximations to univariate and vector autoregressions. *Economics Letters*, 20(2):177–181, 1986. ISSN 0165-1765. doi: 10.1016/0165-1765(86)90168-0.

Stijn Van Nieuwerburgh and Laura Veldkamp. Learning asymmetries in real business cycles. *Journal of Monetary Economics*, 53(4):753–772, 2006.

	(1)	(2)	(3)	(4)
Shape Parameter	-2615.6***	-2746.2***	-1819.5***	-2149.7***
	(0.000)	(0.000)	(0.000)	(0.001)
Shape Parameter Squared	514.0***	527.7***	322.7***	386.1***
	(0.000)	(0.000)	(0.003)	(0.009)
GDP	-0.0000441***	-0.0000336*	-0.0000255***	-0.0000470**
	(0.000)	(0.063)	(0.004)	(0.023)
GDP Squared	2.57e-12***	1.86e-12	1.69e-12**	3.00e-12**
	(0.000)	(0.148)	(0.011)	(0.042)
Price of Investment	-2.992***	-1.778***	-2.970***	-2.126*
	(0.000)	(0.005)	(0.000)	(0.057)
Secondary Education		-41.27**		5.229
		(0.026)		(0.841)
Government Expenditure			28.39***	33.23***
			(0.000)	(0.001)
Investment Expenditure			15.29***	13.35***
			(0.000)	(0.003)
Total Fertility Rate			-45.82*	-53.00
			(0.055)	(0.198)
Political Rights			37.41*	14.12
			(0.060)	(0.670)
Terms of Trade			2.390*	5.891
			(0.098)	(0.200)
Constant	5281.7***	5560.7***	3643.8***	3685.1***
	(0.000)	(0.000)	(0.000)	(0.000)

p-values in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5.1: Full Sample Pooled Regressions

	(1)	(2)	(3)	(4)
Shape Parameter	-2972.8***	-2746.2***	-2140.5***	-2149.7***
	(0.000)	(0.000)	(0.001)	(0.001)
Shape Parameter Squared	574.9***	527.7***	385.1***	386.1***
	(0.000)	(0.000)	(0.008)	(0.009)
GDP	-0.0000463**	-0.0000336*	-0.0000457**	-0.0000470**
	(0.013)	(0.063)	(0.021)	(0.023)
GDP Squared	2.33e-12*	1.86e-12	2.95e-12**	3.00e-12**
	(0.093)	(0.148)	(0.041)	(0.042)
Price of Investment	-2.047***	-1.778***	-2.029*	-2.126*
	(0.002)	(0.005)	(0.052)	(0.057)
Secondary Education		-41.27**		5.229
		(0.026)		(0.841)
Government Expenditure			32.97***	33.23***
			(0.001)	(0.001)
Investment Expenditure			13.13***	13.35***
			(0.002)	(0.003)
Fertility Rate			-51.87	-53.00
			(0.200)	(0.198)
Political Rights			13.73	14.12
			(0.676)	(0.670)
Terms of Trade			5.982	5.891
			(0.203)	(0.200)
Constant	5695.0***	5560.7***	3679.8***	3685.1***
	(0.000)	(0.000)	(0.000)	(0.000)
Observations	113	113	73	73
Adjusted R^2	0.516	0.531	0.709	0.704
F	.	.	.	.

p-values in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5.2: Education Subsample Pooled Regressions

	(1)	(2)	(3)	(4)
Shape Parmameter	-427.8	-246.9	-838.7**	669.1
	(0.119)	(0.389)	(0.016)	(0.224)
Shape Parmameter Squared	95.17*	54.73	191.2**	-133.3
	(0.069)	(0.338)	(0.010)	(0.244)
GDP	-0.0000843**	-0.0000561	-0.0000138	0.00000134
	(0.050)	(0.155)	(0.214)	(0.928)
GDP Squared	4.13e-12*	2.70e-12	2.17e-13	-1.72e-13
	(0.068)	(0.177)	(0.715)	(0.798)
Price of Investment	-0.757***	-0.790***	-0.964***	-1.130***
	(0.008)	(0.003)	(0.000)	(0.004)
Secondary Education		-34.05**		-9.283
		(0.017)		(0.250)
Government Expenditure			-1.142	5.881
			(0.815)	(0.357)
Investment Expenditure			4.589***	5.047***
			(0.000)	(0.001)
Total Fertility Rate			-54.34**	-36.94
			(0.035)	(0.454)
Political Rights			14.88	72.75***
			(0.137)	(0.000)
Terms of Trade			-0.630	0.208
			(0.392)	(0.861)
Constant	2361.4***	2292.2***	2806.9***	881.3
	(0.000)	(0.000)	(0.000)	(0.175)
Observations	470	113	269	73
Adjusted R^2	0.259	0.374	0.381	0.686
F	.	.	.	.

p-values in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5.3: Full Sample Fixed Effect Regressions

	(1)	(2)	(3)	(4)
Shape Parameter	-278.2	-246.9	675.3	669.1
	(0.390)	(0.389)	(0.227)	(0.224)
Shape Parameter Squared	64.05	54.73	-134.0	-133.3
	(0.302)	(0.338)	(0.249)	(0.244)
GDP	-0.0000936**	-0.0000561	-0.00000403	0.00000134
	(0.031)	(0.155)	(0.794)	(0.928)
GDP Squared	4.57e-12**	2.70e-12	7.19e-14	-1.72e-13
	(0.039)	(0.177)	(0.918)	(0.798)
Price of Investment	-0.814***	-0.790***	-1.253***	-1.130***
	(0.004)	(0.003)	(0.002)	(0.004)
Secondary Education		-34.05**		-9.283
		(0.017)		(0.250)
Government Expenditure			5.298	5.881
			(0.420)	(0.357)
Investment Expenditure			5.266***	5.047***
			(0.000)	(0.001)
Fertility Rate			-34.87	-36.94
			(0.454)	(0.454)
Political Rights			66.76***	72.75***
			(0.000)	(0.000)
Terms of Trade			0.263	0.208
			(0.830)	(0.861)
Constant	2207.9***	2292.2***	843.2	881.3
	(0.000)	(0.000)	(0.196)	(0.175)
Observations	113	113	73	73
Adjusted R^2	0.301	0.374	0.686	0.686
F	.	.	.	.

p-values in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5.4: Education Subsample Fixed Effect Regressions

	Covariates					
Equation	Constant	Trend	W_{-1}	G_{-1}	Y_{-1}	
W_t	-0.00	0.00	1.12	-0.86	-0.06	
T-stat	-0.05	0.05	15.1	-1.15	-1.56	
G_t	-0.01	0.01	-0.00	3.53	0.00	
T-stat	-2.78	2.92	-0.27	58.5	0.96	
Y_t	0.01	-0.00	0.26	-4.50	0.96	
T-stat	0.44	-0.44	1.70	-2.85	12.8	

Equation	W_{-2}	G_{-2}	Y_{-2}	W_{-3}	G_{-3}	Y_{-3}
W_t	-0.38	1.82	0.11	0.30	1.22	-0.05
T-stat	-3.42	0.69	2.17	2.59	0.28	-0.93
G_t	-0.00	-5.00	-0.00	0.00	3.52	-0.00
T-stat	-0.52	-23.5	-0.25	0.17	10.0	-0.04
Y_t	-0.05	13.7	-0.10	-0.17	-16.0	-0.08
T-stat	-0.19	2.46	-0.90	-0.69	-1.75	-0.74

Equation	W_{-4}	G_{-4}	Y_{-4}	W_{-5}	G_{-5}	Y_{-5}
W_t	-0.26	-7.05	-0.03	0.05	8.57	0.03
T-stat	-2.24	-1.43	-0.56	0.47	1.76	0.66
G_t	-0.01	-1.96	0.00	0.01	2.68	-0.01
T-stat	-0.51	-4.91	0.10	1.44	6.80	-1.61
Y_t	0.34	7.88	0.09	-0.43	-0.21	-0.11
T-stat	1.38	0.76	0.87	-1.77	-0.02	-1.03

Equation	W_{-6}	G_{-6}	Y_{-6}	W_{-7}	G_{-7}	Y_{-7}
W_t	0.05	-4.90	-0.01	-0.06	1.26	-0.08
T-stat	0.46	-1.18	-0.16	-0.54	0.53	-1.79
G_t	-0.00	-3.36	-0.00	-0.01	2.09	0.00
T-stat	-0.07	-9.95	-0.06	-1.00	10.9	0.28
Y_t	0.24	-0.05	-0.03	-0.05	-1.50	-0.05
T-stat	1.01	-0.01	-0.28	-0.21	-0.30	-0.48

	W_{-8}	G_{-8}	Y_{-8}
W_t	-0.11	-0.06	0.08
T-stat	-1.56	-0.09	2.30
G_t	0.00	-0.51	0.00
T-stat	0.66	-10.0	0.57
Y_t	0.05	0.79	-0.01
T-stat	0.30	0.60	-0.13

Table 5.5: VAR(8) Estimates

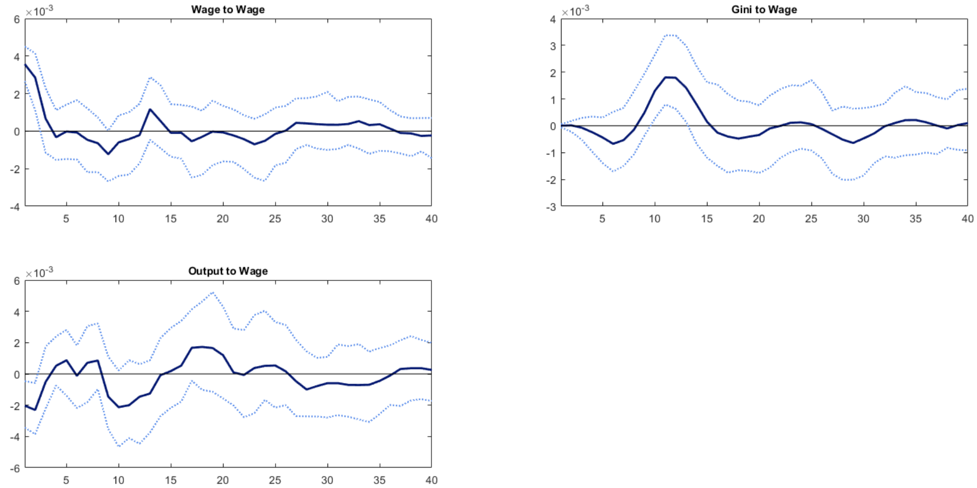

Figure 5.1: Structural Impulse Responses to Wage Shock

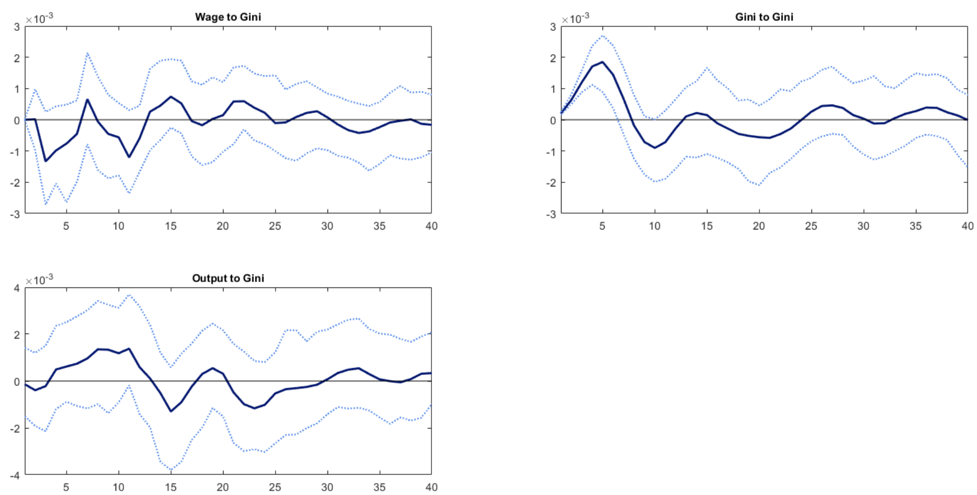

Figure 5.2: Structural Impulse Responses to Inequality Shock

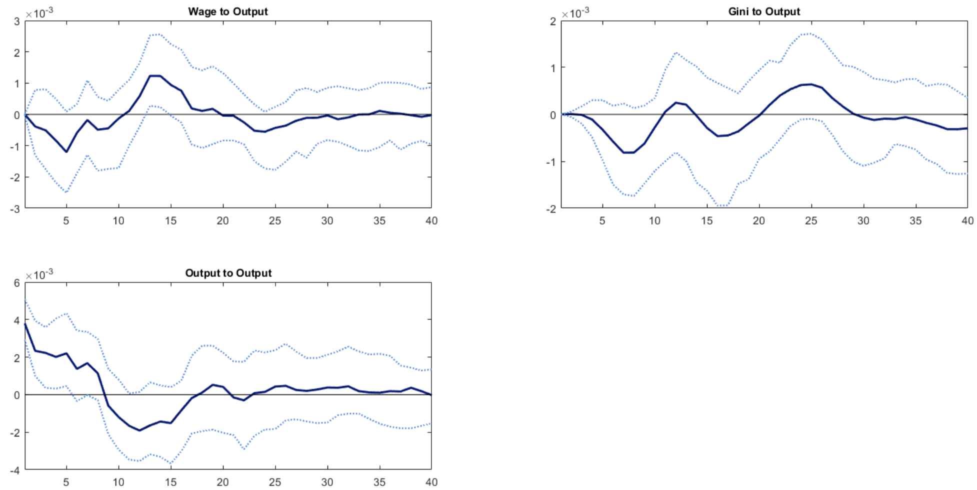

Figure 5.3: Structural Impulse Responses to Output Shock