

UC Santa Cruz

UC Santa Cruz Electronic Theses and Dissertations

Title

Essays in Macroeconomic Dynamics

Permalink

<https://escholarship.org/uc/item/7qs0k3n7>

Author

Munro, David Ross

Publication Date

2016

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
SANTA CRUZ

ESSAYS IN MACROECONOMIC DYNAMICS

A dissertation submitted in partial satisfaction of the
requirements for the degree of

DOCTOR OF PHILOSOPHY

in

ECONOMICS

by

David R. Munro

June 2016

The Dissertation of David R. Munro
is approved:

Professor Carl Walsh, Chair

Professor Kenneth Kletzer

Professor Grace Gu

Dean Tyrus Miller
Vice Provost and Dean of Graduate Studies

Copyright © by

David R. Munro

2016

Table of Contents

| | |
|---|-----------|
| List of Figures | v |
| List of Tables | vi |
| Abstract | vii |
| Dedication | x |
| Acknowledgments | xi |
| 1 Consumer Behavior and Firm Volatility | 1 |
| 1.1 Introduction | 1 |
| 1.2 Demand Elasticity and Dispersion: An Example | 7 |
| 1.3 Empirics | 9 |
| 1.3.1 Countercyclical Dispersion | 9 |
| 1.3.2 Retail Data Analysis | 13 |
| 1.4 General Equilibrium Model | 31 |
| 1.4.1 Representative Household Problem | 31 |
| 1.4.2 Firms | 39 |
| 1.4.3 Model Solution | 41 |
| 1.4.4 Calibration | 42 |
| 1.4.5 Model Results | 45 |
| 1.4.6 Sensitivity Analysis | 48 |
| 1.5 Conclusion | 51 |
| A Additional empirical results | 55 |
| A.1 Additional analysis on countercyclical dispersion using SythLBD | 56 |
| A.2 Additional analysis on households' relative consumption price | 59 |
| A.3 Additional Store Analysis | 62 |

| | | |
|----------|---|------------|
| B | Comparison of IRF’s across calibrations | 64 |
| 2 | Experience, Skill Composition, and the Persistence of Unemployment Fluctuations (with Aspen Gorry) | 67 |
| 2.1 | Introduction | 67 |
| 2.2 | Model | 75 |
| 2.2.1 | Setup and Worker Flows | 76 |
| 2.2.2 | Value Functions and Equilibrium | 77 |
| 2.3 | Paramaterization | 81 |
| 2.4 | Numerical Results | 84 |
| 2.4.1 | Unemployment, Wages, and Worker Flows | 85 |
| 2.4.2 | Persistence | 87 |
| 2.4.3 | Thin Market Externality | 92 |
| 2.4.4 | Interpreting Business Cycles | 100 |
| 2.5 | Discussion | 105 |
| C | Steady State Model Solution | 109 |
| 3 | Unemployment Composition and Aggregate Dynamics | 113 |
| 3.1 | Introduction | 113 |
| 3.2 | Empirics | 120 |
| 3.2.1 | Unemployment by Reason | 122 |
| 3.2.2 | Job Finding Rates by Reason | 123 |
| 3.2.3 | Compositional effect | 126 |
| 3.2.4 | Elasticity of Job Finding Rate | 129 |
| 3.2.5 | Why compositional shifts have become stronger | 130 |
| 3.3 | Model | 132 |
| 3.3.1 | Model Solution | 137 |
| 3.3.2 | Steady-State | 139 |
| 3.4 | Calibration | 140 |
| 3.5 | Results | 142 |
| 3.6 | Discussion | 149 |
| D | Derivation of Steady-State | 154 |

List of Figures

| | | |
|-----|---|-----|
| 1.1 | Demand elasticity and dispersion example | 8 |
| 1.2 | Shopping trips | 15 |
| 1.3 | Household consumption price adjustment | 19 |
| 1.4 | Shopping behavior and demand elasticity | 39 |
| 1.5 | Impulse response functions | 46 |
| A.1 | Relative consumption price unweighted | 61 |
| B.1 | Output IRF in the two calibrations | 64 |
| B.2 | Shopping IRF in the two calibrations | 65 |
| B.3 | IQR IRF in the two calibrations | 65 |
| B.4 | Relative price IRF in the two calibrations | 66 |
| 2.1 | Model vs. Empirical wage growth | 86 |
| 2.2 | Job finding and job separation rates | 88 |
| 2.3 | Model unemployment persistence | 91 |
| 2.4 | Model unemployment persistence with thin market externality | 98 |
| 2.5 | Evaluating thin market externality | 99 |
| 2.6 | Recession simulation - unemployment rate | 102 |
| 2.7 | Recession simulation - job finding and separation rates | 104 |
| 3.1 | Unemployment persistence in post-war recessions | 115 |
| 3.2 | Unemployment shares by reason for unemployment | 122 |
| 3.3 | Job finding rates by reason of unemployment | 125 |
| 3.4 | Change in job finding rate relative to the start of the recession | 127 |
| 3.5 | Change in job finding rates due to the composition effect | 127 |
| 3.6 | Plots value of a vacancy as a function of θ^{ss} | 140 |

List of Tables

| | | |
|-----|---|-----|
| 1.1 | IQR results for retail industries | 12 |
| 1.2 | Firm market share regression | 22 |
| 1.3 | Market consumption price regression | 25 |
| 1.4 | Nationwide IQR regression | 27 |
| 1.5 | Local market IQR regression | 29 |
| 1.6 | Consumption price IQR regression | 30 |
| 1.7 | Model calibration | 44 |
| 1.8 | Empirical vs. Model responses | 51 |
| A.1 | IQR SynthLBD results for manufacturing | 56 |
| A.2 | IQR SynthLBD robustness | 58 |
| A.3 | Household consumption price regression | 60 |
| A.4 | Firm market share regression robustness test | 62 |
| A.5 | AR(1) regression on stores' relative price | 63 |
| 2.1 | Baseline values for model parameters along with targets. | 83 |
| 2.2 | State state results | 85 |
| 2.3 | Unemployment half-lives | 92 |
| 2.4 | Half-lives for different worker types | 100 |
| 3.1 | Compositional changes in unemployment during recessions | 124 |
| 3.2 | Elasticity of job finding rates | 130 |
| 3.3 | Calibration values for model with common elasticity. | 141 |
| 3.4 | Model half-lives | 145 |
| 3.5 | Empirical and model second moments with common elasticity | 147 |
| 3.6 | Empirical and model second moments with separate elasticity | 149 |

Abstract

Essays in Macroeconomic Dynamics

by

David R. Munro

The motivation of my dissertation research has been to develop a better understanding of the mechanisms behind business cycle fluctuations in employment and firm dynamics. I have an interest in these issues not only because I find business cycle phenomena interesting, but because it is crucial in designing economic policies that can help mitigate the severity of recessions. To answer these questions my dissertation research has focused on outcomes and behavior at the individual and firm level.

My first chapter is focused on firm dynamics over the business cycle. Growth rates of firms' employment and revenues becomes more disperse during recessions. Existing research on this business cycle phenomena has focused on information and frictions present on the firms' side of the economy. I argue that this increase in dispersion in firm-level growth rates can arise from changes in consumer behavior over the business cycle. The key link between consumer behavior and the dispersion of firm-level growth rates is demand elasticity: when demand elasticity is high, a cost shock has a larger impact on a firms' sales and employment, relative to when demand elasticity is low. Using a UPC-level data set of prices and quantities

at retail stores and a panel of household purchases, I find evidence that demand elasticity rises during recessions. Consistent with changing demand elasticity, the dispersion of stores' growth rates increases during recessions and this increase is larger in markets where the change in consumer behavior is the strongest. To assess the importance of this mechanism I construct a business cycle model with heterogeneous firms and frictions in product markets. In the model, it is costly for households to obtain the best prices in the market, and more shopping effort translates into lower consumption prices. With this margin of adjustment available, households increase shopping effort during recessions to obtain lower prices as a means to mitigate their fall in consumption. This behavior changes the demand elasticity faced by firms, leading to countercyclical dispersion. The model generates countercyclical shopping effort, procyclical relative consumption prices, and countercyclical dispersion, all of which are observed empirically. The baseline calibration of the model is able to generate countercyclical dispersion that is roughly one third of that observed empirically.

My second and third chapters focus on unemployment fluctuations over the business cycle. Traditional models of unemployment struggle to reproduce the persistence of unemployment fluctuations observed over the business cycle. In my second chapter, I extend a traditional search and matching model of unemployment to capture skill acquisition over a normal working life. The model shows

that recessions that are characterized by human capital loss lead to slow recoveries in unemployment. My third dissertation chapter examines the relationship between the characteristics of the pool of unemployed workers and the pace of unemployment recoveries. Using monthly U.S. census data I examine unemployed workers' reported reason for unemployment and how these workers differ in their subsequent job-finding rates. I find that workers who have been permanently displaced from their previous jobs with no expectation of recall have drastically lower job-finding rates than those on temporary layoffs and other reasons for unemployment. I show that recent recessions have been characterized by a stronger compositional shift towards these low job-finding rate permanent displacements, especially relative to earlier post-war recessions. These stronger compositional changes are an important factor behind slow unemployment recoveries in recent U.S. recessions. I find that this shift is due, in part, to the changing industrial landscape in the U.S., moving towards a professional and service-oriented economy and away from production-oriented industries.

To my family and my partner, Rita Jones,
for their continued support and encouragement.

Acknowledgments

I thank my advisor, Carl Walsh, for input and guidance throughout my PhD. I also thank Rudi Bachmann, Frank Caliendo, Johanna Francis, Yuriy Gorodnichenko, Aspen Gorry, Grace Gu, Ken Kletzer, Ezra Oberfield, Guillaume Rocheteau, Ajay Shenoy, Alan Spearot, and seminar participants at the 10th NP Quantitative Macroeconomics Workshop, 6th Southwest Search and Matching Workshop, the Fall 2013 Midwest Macro Conference, the 2014 AEA Meetings, Carleton University, Elon University, the Federal Reserve Bank of Atlanta, the Federal Reserve Board of Governors, the UC Conference on Social Science and Law, UCSC, Washington State University, the Washington University St. Louis Economics Grad Student Conference, and Wilfred Laurier University for helpful comments on the work contained in this dissertation.

Chapter 1

Consumer Behavior and Firm

Volatility

1.1 Introduction

Volatility in firm-level growth rates rises during recessions. This greater dispersion in firm outcomes is important to understand because it can have implications for the macroeconomy. Frictions present in labor, product, and credit markets can result in a misallocation of resources and lower productivity when firm outcomes become more disperse. Understanding the driving forces behind this increased dispersion is not only important for understanding business cycles but can help economists and policy makers design policy to mitigate the inefficiencies caused by

it. The main objective of this paper is to show that changes in consumer behavior over the business cycle is an important driver of countercyclical dispersion.

A number of authors have argued that countercyclical dispersion in firm-level growth rates arises from changes in “uncertainty” about future outcomes and that these changes are an important driver of business cycles (see Bloom (2009), Arellano et al. (2010), Schaal (2012), Bloom et al. (2012), Gilchrist et al. (2014), and Christiano et al. (2014)). However, others have argued that there is weak empirical support for these “uncertainty” shocks (see Bachmann et al. (2010) and Bachmann et al. (2011)). Other authors have argued that the causality runs in the opposite direction: business cycles themselves are the cause of countercyclical dispersion. For example, Ilut et al. (2014) argue that firm hiring/firing decisions can lead to countercyclical dispersion and Decker et al. (2014) argue that if firms reduce the number of markets they serve during recessions they become exposed to the idiosyncratic shocks of the individual markets which leads to countercyclical dispersion. Berger and Vavra (2013) find that firms become more responsive to changes in costs (exchange rates) during recessions. Their argument is very similar to the one made in this paper except that they focus on the responsiveness of firms to the prices they face, where I focus on the responsiveness of households.

A growing body of research has highlighted the importance of changes in consumer behavior over the business cycle. Using data from the American Time Use

Survey, Aguiar et al. (2013) documents how household allocation of time changes over the recession. Importantly, for this paper, they find that time spent shopping increased significantly during the great recession. Likewise, McKenzie and Schar-grotsky (2011) find that during the 2002 Argentina crisis the average household increased their shopping frequency even while the real value of goods purchased fell. In studying dispersion of prices, Kaplan and Menzio (2014) find that households with fewer employed members pay lower prices and that they achieve lower prices by visiting more stores rather than by shopping more frequently. Likewise, Coibion et al. (2012) find that prices actually paid by consumers decline significantly during recessions while little change occurs in the inflation rate of posted prices. They argue that this is a result of households reallocating expenditures towards lower priced retailers. Stroebel and Vavra (2014) document a causal response of local retail prices to changes in house prices. They argue that the change in retail prices is driven by changes in mark-ups rather than local costs, and that this mark-up variation arises because movements in housing wealth changes household demand elasticity. The common thread of these papers is that household consumption behavior changes systematically over the business cycle. In this paper I link the literatures on countercyclical dispersion and consumer behavior and show that changes in consumer behavior can have important implications for firm dynamics.

The link between consumer behavior and the dispersion in firm growth rates is through demand elasticity. The intuition behind this mechanism is simple: in an economy where firms experience idiosyncratic cost shocks, when demand elasticity is high, a shock has a strong impact on a firm’s sales relative to when demand elasticity is low.¹

I begin this paper by examining the empirical evidence on changes in the volatility of firm-level growth rates over the business cycle. The existing literature on countercyclical dispersion has typically focused on firm-level growth rates for manufacturing industries and rightly so, these firms are capital intensive and often face substantial labor frictions. Using firm-level synthetic census data I show that countercyclical dispersion is present in retail industries and that its magnitude is larger than that found in high-friction manufacturing industries.

I next employ UPC-level consumer panel data to investigate how consumer behavior changes over the business cycle to explore whether consumption side mechanisms are a plausible story behind countercyclical dispersion. Consistent with recent findings using the American Time Use Survey I find that households dedicate more effort to shopping, as measured by shopping trips, during recessions. With frictions present in product markets this increase in shopping effort

¹There is an important distinction in this literature between firms and establishments. In this paper I focus on the entities with which consumers interact with directly. As such, these business entities would be defined as establishments. For the sake of readability, I mostly use the term “firms” throughout the paper.

could plausibly increase demand elasticity because consumers may have better information about prices and/or are more willing to visit more stores to reduce the price of their consumption bundle. Using information on household purchases I next show that households shift their consumption towards lower priced stores when local unemployment is high, consistent with an increase in demand elasticity during bad economic times. Turning to firm outcomes, I find that lower priced stores gain a larger share of consumption when local unemployment rises, which also consistent with an increase in demand elasticity. Finally, with strong evidence of consumption-side mechanisms at play I explore whether countercyclical dispersion is present in the growth rates of store sales. Grouping at both nationwide and local market levels I find dispersion in firm growth rates increases with unemployment and that this increase in dispersion is larger in markets where the shift towards lower priced stores is the strongest.

To further explore the quantitative importance of this consumption-side mechanism I construct a simple business cycle model where frictions in product markets make it costly for households to obtain low priced goods. These product market frictions generate a distribution of prices in equilibrium and are the source of firms' market power, allowing them to set prices above marginal costs. In the model, households choose shopping effort and in expectation higher effort translates to lower priced consumption. Shopping effort becomes a margin of adjust-

ment through which households can mitigate their fall in consumption during a recession by obtaining lower prices. This increase in shopping effort during recessions increases the demand elasticity faced by firms, leading to countercyclical dispersion. When calibrated to the data, the model generates movements in shopping effort, relative consumption prices, and the dispersion of firm growth rates that are quantitatively similar to those observed empirically. The baseline calibration of the model generates countercyclical dispersion that is roughly one third of that observed empirically. The model's findings highlight the importance of consumption-side mechanisms in understanding countercyclical dispersion.

In this paper I focus on demand elasticity at the store level. Of course, consumers may also adjust along other dimensions, for example across products of different quality (name brand vs. generic) or even across product categories (steak vs. spaghetti).² However, obtaining an empirical metric for a product's quality is difficult and may be subjective. At the end of the paper I provide some discussion on how these different margins of consumer adjustment may work back through supply chains from retailers to ultimately impact the dispersion of firm growth rates in other industries besides retail.

The remainder of the paper is structured as follows: Section 1.2 presents a simple example to developed intuition on the mapping between demand elasticity and

²Jaimovich et al. (2015) argue that consumers adjust along quality margins over the business cycle.

countercyclical dispersion. Section 1.3 presents empirical results on countercyclical dispersion and consumer behavior. Section 1.4 develops a general equilibrium model with heterogeneous firms and frictions in product markets to explore the quantitative importance of consumption-side mechanisms in generating countercyclical dispersion. Section 1.5 provides some discussion and concluding remarks.

1.2 Demand Elasticity and Dispersion: An Example

In this section I discuss a simple example to build further intuition on the mapping between demand elasticity and the dispersion of firm growth rates. This example will help motivate the empirical analysis and highlight the important mechanisms driving the results in the general equilibrium model developed below. Assume there are a continuum of firms selling an identical good in a given market and that each of these firms is subject to idiosyncratic cost shocks. If these idiosyncratic shocks are of a typical stationary autoregressive form, one can solve for the ergodic distribution of costs which would map into a distribution of prices. Assume further that there are frictions in this product market such that each firm is able to sell a positive quantity but that there is an inverse relationship between a firm's price and quantity sold. The green distribution on the y-axis in 1.1 plots

an example of an ergodic density of prices charged by firms who are pricing at cost. Two examples of demand curves are shown by the red and blue lines. These demand curves differ only in their price elasticities. In a high price elasticity environment the density of prices can be used to solve for a density of quantities sold, given by the blue distribution under the x-axis.

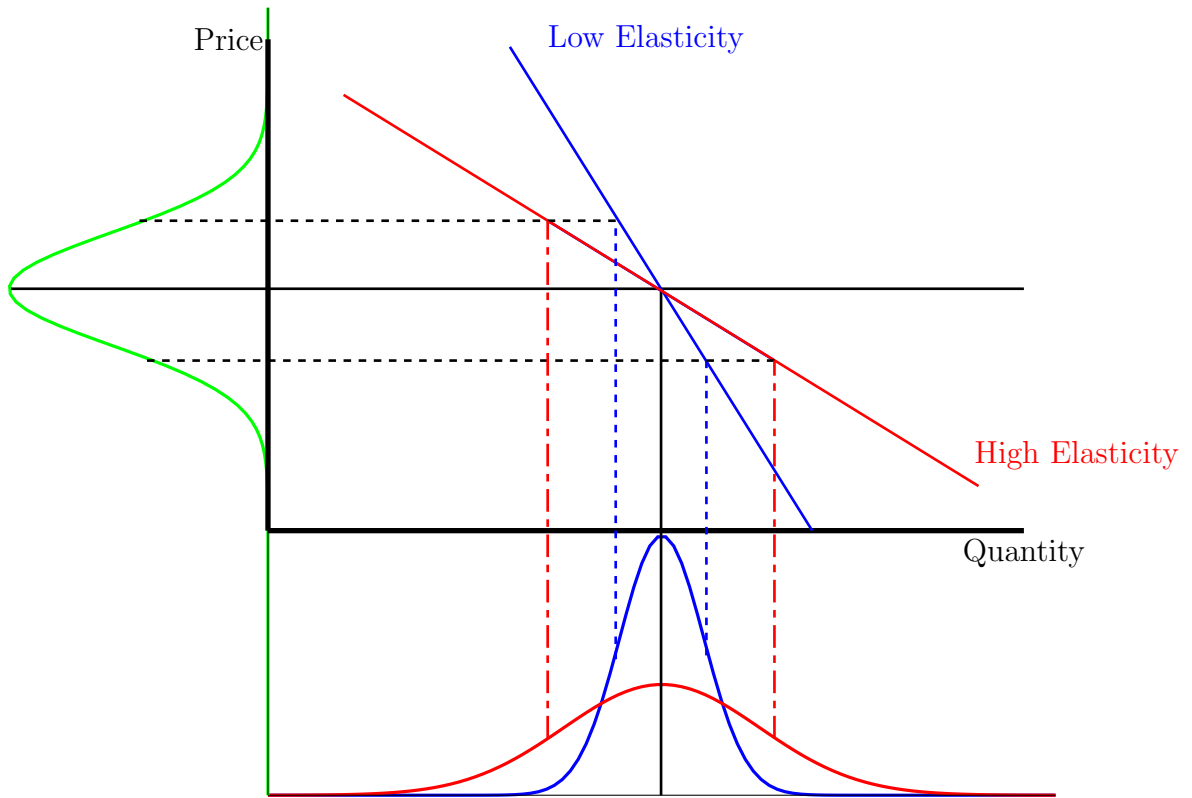


Figure 1.1: Demand elasticity and dispersion example

It is important to emphasize that while the aggregate cost distribution has converged to its ergodic form, the idiosyncratic cost shocks are still active, leading to a dispersion of growth rates across firms. The red line in Figure 1.1 is an

environment where price elasticity is much higher. Without changing the price density, this higher elasticity translates into a much more dispersed density of quantities as shown by the red distribution. This wider distribution is representative of the greater dispersion in growth rates across firms - the same idiosyncratic cost shock causes a firm's demand to change more dramatically in the high elasticity environment. For clarity of exposition I assume that only demand elasticity changes in these two examples. However, the intuition also holds if changes in demand elasticity are accompanied by shifts in demand since the focus is on second moments.³

1.3 Empirics

1.3.1 Countercyclical Dispersion

Countercyclical dispersion has been examined in various firm-level growth rates. In this section I focus on growth rates in firm employment. The magnitude of countercyclical dispersion in firm-level employment has been documented in Decker et al. (2013) and Ilut et al. (2014). Using confidential data from the Longitudinal Business Database (LBD) Ilut et al. (2014) report that on average

³This intuition is similar to the mechanism in Berger and Vavra (2013). They argue that firms become more responsive to costs (exchange rate fluctuations) during recessions and that this leads to an increase in the dispersion of price changes. The argument here is similar, except that it focuses on the increase in households' responsiveness to prices.

NBER recessions have an interquartile range (IQR) that is 4.2 percentage points larger in recessions versus expansions for manufacturing firms.

Manufacturing firms typically face high capital intensities and labor frictions. Focusing on industries with these operating frictions might produce large estimates of countercyclical dispersion if these frictions are the important driving factor behind the changes in firm growth rates. Since the main argument in this paper is that consumer behavior is an important driver of countercyclical dispersion, it should be the case that countercyclical dispersion is also present in industries outside of manufacturing. To explore this I use the Synthetic Longitudinal Business Database (SynthLBD) to investigate how the dispersion in firm-level growth rates changes over the business cycle.⁴ This data set spans 1976-2000 and therefore covers the early '80s recessions and the 1991 recession. Column 4 of Table 1.1 reports the difference between the average IQR of firm employment growth rates during NBER recessions and expansions for various 3-digit SIC retail industries.⁵ The level of disaggregation at the 3-digit SIC level is important because it helps control for differences in cyclical sensitivities of industries driving the dispersion results. These results show that for all retail industries the dispersion in firm-level

⁴See U.S. Census Bureau (2011) for further details on this dataset. I thank the creators of the Synthetic LBD for providing access to the dataset. Version 2 of the Synthetic LBD was funded by NSF Grant 0427889 and the Synthetic Data Server is funded through NSF grant SES-1042181.

⁵Following Davis and Haltiwanger (1990) I compute the growth rate of employment as $n_t^i = 2(L_t^i - L_{t-1}^i)/(L_t^i + L_{t-1}^i)$ where L_t^i is firm i 's employment at time t . This helps avoid cyclical entry and exit driving the dispersion results.

growth rates is larger during recessions than in expansions. Later in the paper I examine micro data from retail grocery stores and it is worth noting that this industry displays large countercyclical dispersion in the SynthLBD, see row three in Table 1.1.

Table A.1 in Appendix A presents the difference in IQRs for manufacturing industries in NBER recession and non-recession years using the SynthLBD. On average, manufacturing industries have an IQR that is 2.7 percentage points larger in recessions relative to expansions.⁶ As noted above, Ilut et al. (2014) report a difference in IQRs between recession and expansions of 4.2 percentage points when pooling firms across manufacturing industries. This estimate, while somewhat larger, is similar to the 2.7 percentage point difference that I find using the SynthLBD and thus suggests that the SynthLBD does a good job at capturing changes in the dispersion of firm-level growth rates observed during recessions.⁷ Comparing the SynthLBD results from manufacturing and retail industries shows that on average countercyclical dispersion is larger in retail industries. With low business frictions relative to manufacturing industries this comparison suggests

⁶It is important to note that the NBER recession years are not evenly spaced in the SynthLBD coverage and that early years in the data set are over represented by recessions. This is problematic given the Decker et al. (2013) finding that IQRs in employment growth rates have experienced a secular decline in many US industries over the past few decades. I describe how I control for this problem in Appendix A and show that the conclusions here are unchanged.

⁷It is important to emphasize that the 4.2 percentage point estimate reported in Ilut et al. (2014) pools across manufacturing industries which likely biased this result upwards due to the different cyclical sensitivities of various manufacturing industries. They are careful to control for this bias in other estimates in their paper.

that frictions on the firms' side of the economy may not be the only mechanisms driving countercyclical dispersion.

| Industry | SIC Code | Description | $\text{Avg}(IQR_{Rec.}) - \text{Avg}(IQR_{Boom})$ |
|---------------------------|----------|----------------------|---|
| Retail | | | |
| | 521-527 | Building Materials | 0.0395 |
| | 531-539 | General Merchandise | 0.0260 |
| | 541-549 | Food Stores | 0.0515 |
| | 551-559 | Auto Dealers & Serv. | 0.0365 |
| | 561-569 | Apparel & Access. | 0.0442 |
| | 571-573 | Home Furnishings | 0.0291 |
| | 581 | Eating & Drinking | 0.0625 |
| | 591-599 | Miscellaneous | 0.0494 |
| Unweighted Average | | | 0.042 |

Table 1.1: Reports the difference between the average IQR in firm employment growth rates for different industries in NBER recession and expansion years.

Firms in all industries are, of course, impacted by changes in consumer behavior to some degree. While I advocate for consumption-side mechanisms as an important driving force behind countercyclical dispersion, it is important to emphasize that this does not mean it is the *only* driving force. Understanding the relative importance of consumption mechanisms versus firm-side frictions and how those might vary across industries, while interesting, is beyond the scope of this paper. However, highlighting the importance of consumption-side mechanisms is an important step towards enriching our understanding of countercyclical dispersion and the data examined in this section provide evidence that countercyclical dispersion is large in consumer oriented industries.

1.3.2 Retail Data Analysis

1.3.2.1 Data Description

In this section I explore whether there is empirical support for consumer behavior as an important driver of countercyclical dispersion. Using data from IRI Marketing I examine both how household consumption behavior and firm outcomes change over the business cycle. The data set contains weekly UPC-level data from 2001 to 2011 on price and quantity for UPC codes in 31 product categories from stores in 47 retail chains in 50 geographic markets in the U.S.⁸ These product categories include such things as beer, milk, salty snacks, etc. In the 50 geographic markets there are a total of 505 drug stores and 1,588 grocery stores.

The data also contains a consumer panel in two markets. This panel tracks household purchases and demographic information for between 3,000 and 5,000 households in Pittsfield, MA and Eau Claire, WI. Just under 1,500 households are present over the entire life of the panel.

1.3.2.2 Consumer Behavior

Using the consumer panel I examine various elements of household shopping behavior and how this changes over the business cycle. Unfortunately, this panel

⁸See Bronnenberg et al. (2008) for a detailed description of the data set. I would like to thank IRI for making the data available. All estimates and analysis in this paper, based on data provided by IRI are by the author and not by IRI.

is only conducted in two geographic markets and does not have a time series of demographic information on households. As such, to get a sense of how consumer behavior changes over the business cycle I can only use aggregate economic conditions in those geographic areas as an explanatory variable. A regression analysis that took this approach would effectively reduce the number of observations to two, as all changes in household shopping behavior in a market would be explained by a common variable.⁹ Due to this issue, I present only descriptive data on household shopping behavior and leave statistical analysis to the richer store-level data set.

Shopping Trips:

As argued in the partial equilibrium example above, changing demand elasticity can result in changes in the dispersion of firm-level growth rates. In an environment where households need to acquire information about prices or spend time traveling to obtain the best prices, demand elasticity can change based on how much shopping effort is employed. When households have a large amount of information on prices or spend a lot of time visiting stores to obtain the best prices, demand elasticity will be high because it is more likely that stores with the best

⁹In this setting there is no appropriate way to control for any cross-sectional dependence in a statistical analysis that is surely to exist when all households behavior is being explained by a common variable. One approach would be to cluster standard errors at the market level, but with only two markets even bootstrapping along this dimension would be problematic.

prices will receive most of the demand. Conversely, in a low information or low shopping effort environment demand elasticity will be low because few households will have information on the location of the best prices or be less willing to travel to obtain them. Using the consumer panel data I examine the extent that shopping effort changes over the business cycle by examining the number of shopping trips a household takes in a given year. Figure 1.2 plots the average number of annual shopping trips taken by households in the panel across time.

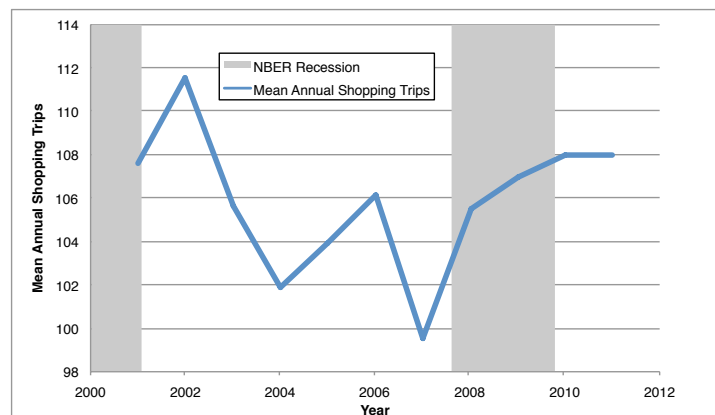


Figure 1.2: Displays the average number of annual shopping trips across all households in the consumer panel from 2001 to 2011. NBER recession years are denoted by the areas shaded in gray.

The average number of shopping trips from the consumer panel is generally consistent with countercyclical shopping effort. From the business cycle peak in 2007 to 2010 the average number of annual shopping trips increased by 8.5%. This increase is similar to findings reported in Aguiar et al. (2013). Using the

American Time Use Survey (ATUS) they find that the average time households reported shopping increased by 7% during the great recession. It is important to emphasize that shopping trips in the consumer panel does not capture visits to stores where nothing was purchased or the time dedicated towards shopping, and thus is a much different metric than the time spent shopping reported in the ATUS. However, both data sources are consistent in showing that shopping effort increases during recessions.

Store-Switching Behavior:

Evidence of increased shopping effort during recessions is only indirect evidence of a change in demand elasticity. A more direct test is to examine whether households shift consumption towards lower priced stores. To explore this prediction empirically, I follow the approach used in Coibion et al. (2012).¹⁰ In general, this involves obtaining a metric for the relative expensiveness of stores and seeing how households allocate their consumption bundles across these stores, and whether this varies over the business cycle. I begin by obtaining a metric of the relative expensiveness of a store then use this metric to calculate the relative expensiveness of a household's consumption bundle.

I begin by computing the median UPC-level price in each market in each

¹⁰Coibion et al. (2012) examine the implications of store-switching behavior for the effects of monetary policy. My results are consistent with their findings that households shift consumption to lower priced stores during recessions.

month. I then compute the relative price of that UPC code at each store in a given market as the log difference between a store’s price and the median local market price. With these UPC-level relative prices I then find the average relative price of a store across UPC codes. Because stores may be oriented on different quality products it is important that I restrict a store’s relative price to include only UPC codes that are sold in a wide range of stores. A “high end” store might, in general, sell products at higher prices than a “lower end” store simply because they are selling higher quality products. Since I would like to obtain a metric for the relative expensiveness of a store I need to control for these quality price effects. To do this I compare store prices across identical products. Following Coibion et al. (2012), I find a store’s average relative price in a product category by only using UPC codes that are sold in 75%, 90%, and 100% of stores in a market in a given month. These different sampling criteria are denoted S_{75} , S_{90} , and S_{100} . With these average relative prices of a store for each product category I find a store’s overall average price by taking a simple average and a consumption weighted average across the 31 product categories.¹¹ Denote the average relative price of store j in market m and month t using sampling criteria S as $\bar{R}_{j,m,t,S}$.

Using household purchase data, I next find the share of a household’s spending

¹¹To find the product category consumption weights I compute the total sales in each product category over all 50 geographic regions over all 11 years and divide this by the total sales over all product categories in all markets over all 11 years. This helps control for any shifts in the composition of consumption bundles that might be taking place over the business cycle.

done at each store j in each month t . I use this information to find the relative price of a household's consumption bundle by multiplying a store's relative price by the household's share of monthly consumption at that store. Summing across all the stores where a household shops gives a metric of the relative price of that household's consumption bundle: $\tilde{R}_{i,m,t,S} = \sum_{j \in m} \phi_{i,j,t} \bar{R}_{j,m,t,S}$

Where $\phi_{j,t}$ is the share that household i spends at store j in month t . To get a sense of how the relative price of a household's consumption bundle changes over the business cycle I compute the mean relative consumption price for all households in the consumer panel and plot its time series in Figure 1.3.¹²

¹²This relative consumption price is constructed using store-level mean relative prices that are weighted by product category consumption. Figure A.1 in the Appendix plots the same time series for the simple average mean relative consumption price (i.e. not the consumption weighted average). These results are very similar to the consumption weighted ones in Figure 1.3.

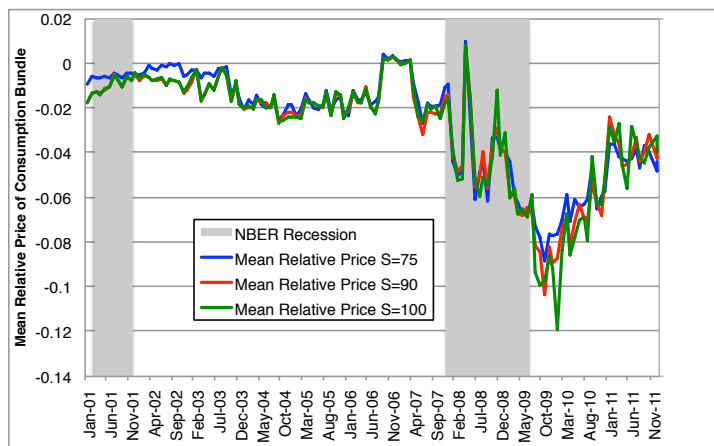


Figure 1.3: Displays the consumption weighted relative consumption price averaged across all households for the three UPC sampling criteria, $S = 75$, $S = 90$, and $S = 100$. See main text for a description of these criteria. NBER recession months are shaded in gray.

All UPC sampling criteria show a clear business cycle pattern to the relative price of a households' consumption bundles. During the great recession, households shifted consumption to lower priced stores.¹³ These descriptive statistics are consistent with household demand elasticity increasing during recessions.

¹³In Appendix A I also report regression results estimating the relationship between a household's relative consumption price and the local unemployment rate. These results are consistent with the patterns in Figure 1.3. The standard errors from these regressions are only robust to autocorrelation at the household level so statistical inference is not recommended. However, the coefficient estimates are useful in that they report the magnitude of the change in a household's relative consumption price that can be compared to the results from the firm analysis.

1.3.2.3 Store Outcomes

In the preceding sections I show descriptive data from the consumer panel that suggests households exert more shopping effort and shift consumption to lower priced stores during recessions. In this section I explore how this change in behavior impacts stores. Specifically, as households display a higher price sensitivity during recessions we should expect this change in behavior to a) result in firms being differentially impacted by a recession depending on their relative position in the price distribution, b) cause the relative consumption price in a given market to decline as households shift consumption to cheaper stores, and c) result in more disperse growth rates of firm sales. With the richer store-level data set (50 geographic markets) I explore the statistical support of these predictions.

Store Sales:

In this section I explore how the impact of a recession on a store's share of local market sales is different sales depending on that store's position in the price distribution. Due to the potential simultaneity between a store's market share and their relative price I simply examine how firms' market shares change over the great recession conditional on their relative price before the start of the recession. Specifically, I find a store's mean relative price during 2007 and explore their market share changes over the great recession (2008-2011) by running the following

specification:

$$\frac{Sales_{j,m,t}}{MarketSales_{m,t}} = \alpha_j + \alpha_1 UR_{m,t} + \alpha_2 UR_{m,t} * \bar{R}_{j,m,2007,S} + error \quad (1.1)$$

Where α_j is a store j fixed-effect, $UR_{m,t}$ is the seasonally adjusted unemployment rate in market m in month t , $\bar{R}_{j,m,2007,S}$ is the average relative price of store j in market m in 2007 with UPC sampling criteria S , and $\frac{Sales_{j,m,t}}{MarketSales_{m,t}}$ is the share of total local market sales for store j in market m in month t .¹⁴ In this specification, the coefficient of interest is on the interaction term. If the great recession was characterized by an increase in price sensitivity, we would expect the effect on a store's market share to differ conditional on their location in the price distribution.

¹⁴I leave out the relative price instrument as a regressor as it is subsumed in the store fixed effect

| Coefficient on $UR_{m,t} * \bar{R}_{j,m,2007}$ | | |
|--|------------------------------|----------------------------|
| UPC Sample | Unweighted Consumption Price | Weighted Consumption Price |
| S_{75} | -0.0022*** (0.0005) | -0.0025*** (0.0006) |
| S_{90} | -0.0021*** (0.0005) | -0.0024*** (0.0006) |
| S_{100} | -0.0019*** (0.0005) | -0.0024*** (0.0005) |

Table 1.2: Reports regression results for specification 1.1 for both unweighted and weighted store prices and for the three different UPC sampling criterion. The dependent variable is the market share of store j 's sales, in market m , in month t . The coefficients reported in the table are for the local unemployment rate which is seasonally adjusted using the X-12 ARIMA method. There are 79,048 observations. Robust standard errors are clustered at the store level. ***, **, * indicate significance at 1%, 5%, and 10% levels, respectively.

Table 1.2 reports the estimated coefficient on the interaction term $(UR_{m,t}x\bar{R}_{j,m,2007,S})$ in specification (1.1) for both weighted and simple average relative store prices.¹⁵

¹⁵The results in Table 1.2 cluster standard errors at the store level. One might also be concerned about correlation at the local market level. The same specification with standard errors clustered at the market level generates nearly identical levels of significance and are

These results show that the impact on a store’s market share is negative for stores above the median priced store in a market. This is consistent with consumption reallocating to lower priced stores during recessions. To give a sense of magnitude, results using consumption weighted prices and a sampling criteria of S_{100} predict that a store whose relative price is 10% above the median store’s price would have a decrease in market share of by 0.024% when the local unemployment rate increases by 10 percentage points.

Relative Market Prices:

An alternative approach to assessing the relative impact of a recession for stores with different prices is to examine how the market price of consumption changes over the business cycle. Descriptive data from the consumer panel showed that households shifted consumption to lower priced stores during the great recession. This was evident through the fact that relative prices of a household consumption bundles declined starting around the end of 2007 through to the end of 2009, at which point they began to recover. This behavior of the relative price of consumption can also be explored using the richer store-level data. To do this I begin by computing the share of total sales a store receives in month t , market m , denoted by $\epsilon_{j,m,t}$ and multiply this share by that store’s relative price in that month, $\bar{R}_{j,m,t,S}$. Summing across the stores in market m , month t , I obtain a metric for the

reported in Table A.4 in Appendix A.

mean relative price of consumption in that market: $\tilde{R}M_{m,t,S} = \sum_{j \in m} \epsilon_{j,m,t} \bar{R}_{j,m,t,S}$.

I explore how this relative market price changes over the business cycle with the following specification:

$$\tilde{R}M_{m,t,S} = \alpha_m + \beta_1 UR_{m,t} + error \quad (1.2)$$

Again, where $UR_{m,t}$ is the seasonally adjusted local unemployment rate in market m , month t and α_m is a market m fixed effect. Coefficient estimates on $UR_{m,t}$ for each UPC sampling criteria and for weighted and simple average store prices are presented in Table 1.3.

| UPC Sample | Simple Average Consumption Price | Weighted Consumption Price |
|------------|----------------------------------|----------------------------|
| S_{75} | -0.0021*** | -0.0015*** |
| S_{90} | -0.0022*** | -0.0017*** |
| S_{100} | -0.0024*** | -0.0019*** |

Table 1.3: Reports regression results for specification 1.2. The dependent variable is the expenditure weighted/simple average average price of a markets consumption in a given month. The coefficients reported in the table are for the local unemployment rate which is seasonally adjusted using the X-12 ARIMA method. There are 6,600 observations. Robust standard errors are clustered at the market level. ***, **, * indicate significance at 1%, 5%, and 10% levels, respectively.

These results show that consumption in a market shifts towards lower priced stores when the local unemployment rate rises. The results are statistically significant at the 1% level across the simple average and consumption weighted store prices and across the three different UPC sampling criteria. These point estimates can be compared to the household behavior observed from the consumer panel. Specification A.1 in Appendix A is used estimate how the relative price of households' consumption bundles changes with the local unemployment rate. The coefficient estimate on the local unemployment rate using weighted consumption

prices and the sampling criteria of $S = 100$ for households is -0.0044 which is similar in magnitude to the corresponding estimate of -0.0019 using store sales data.

Countercyclical Dispersion:

Unfortunately, the data set does not include information about employment at the participating retail stores. However, countercyclical dispersion has also been documented in the growth rates of firm sales, see Bloom et al. (2012). In this section I examine whether countercyclical dispersion is present in the growth rates of store-level sales in this data set. For each store I sum across total sales from each of the 31 product categories in each month to obtain store j 's total sales in month t , denoted $TotSales_{j,t}$.¹⁶ I then compute the percentage change in each store j 's total sales from one month to the next, $\% \Delta TotSales_{j,t-1 \rightarrow t}$.¹⁷ The majority (70%) of the geographic regions in the data set cover fewer than 50 stores (grocery and drug combined). The small number of store-month observations in these regions is problematic for computing IQRs as one store can change an estimate substantially. As such, I pool geographic regions and compute a nationwide IQR for each month. I seasonally adjust this IQR time series and explore how it changes

¹⁶It is important to note that the 31 product categories do not constitute the entire universe of store sales. The main product categories absent in this dataset are produce and meat, which are not standardized products and typically do not have UPC codes.

¹⁷As above, I compute the growth rate of sales as $\% \Delta TotSales_{j,t} = 2(TotSales_{j,t} - TotSales_{j,t-1}) / (TotSales_{j,t} + TotSales_{j,t-1})$.

over the business cycle using the specification in 1.3.¹⁸

$$IQR_t = \beta_0 + \beta_1 * UR_t + \beta_2 * month + error \quad (1.3)$$

Where UR_t is the seasonally adjusted national unemployment rate reported by the Federal Reserve Bank of St. Louis. The results from this regression are reported in Table 1.4.

| Indep. Var. | Coeff. | Std. Err. | p-value |
|-------------|------------------|---------------|--------------|
| UR | 0.0032*** | 0.0010 | 0.001 |
| month | -0.0004*** | 0.00007 | 0.000 |
| cons. | 0.1018*** | 0.00378 | 0.000 |

Table 1.4: Reports regression results for specification 1.3. The dependent variable is the nationwide IQR for the growth rates of total store sales. There are 131 observations and robust standard errors are reported. ***, **, * indicate significance at 1%, 5%, and 10% levels, respectively.

Table 1.4 shows that there is a positive and statistically significant relationship between the IQR of sales growth rates and the unemployment rate. The interpretation of the coefficient on UR_t is that a 1 percentage point increase in the

¹⁸All seasonal adjustment is done using the X-12 ARIMA method.

unemployment rate is associated with an increase in the IQR by 0.32 percentage points. To give a sense of magnitude of the change in the IQR over the business cycle, note that the national unemployment rate increased by 5 percentage points from December 2007 to October 2009. These regression results would predict an associated increase in the IQR by 1.6 percentage points over the great recession which is quantitatively similar to those reported above using the SynthLBD. It is also worth noting that there is a statistically significant downward trend in the IQR, which is consistent with the findings in Davis et al. (2007).

One might be concerned that the countercyclical dispersion reported in Table 1.4 is a result of pooling together stores from geographic regions that are differentially impacted by a national recession. As mentioned above, computing IQRs within a geographic region is problematic because many of the regions contain a small number of stores. This low number of stores makes the IQR sensitive to one store's outcome and thus may produce a noisy time series. Nevertheless, I explore whether countercyclical dispersion is still a feature of the data after controlling for geographic heterogeneity. To do this, I compute the IQR in each market and regress it on the local seasonally adjusted unemployment rate while allowing for market fixed effects. This specification is given below in (1.4) and results are reported in Table 1.5.

$$IQR_{m,t} = \beta_m + \beta_1 * UR_{m,t} + \beta_2 * month + error \quad (1.4)$$

| Indep. Var. | Coeff. | Std. Err. | p-value |
|-------------|-----------------|---------------|--------------|
| UR | 0.0021** | 0.0001 | 0.034 |
| month | -0.0002*** | 0.00006 | 0.000 |
| cons. | 0.1002*** | 0.0046 | 0.000 |

Table 1.5: Reports regression results for specification 1.4. The dependent variable is the IQR in each market for the growth rates of total store sales. Market level fixed effects are used and robust standard errors are clustered at the market level. There are 6550 observations. ***, **, * indicate significance at 1%, 5%, and 10% levels, respectively.

Allowing for these geographic differences generates similar estimates but, as expected, costs some precision: the coefficient on the local unemployment rate is now only significant at the 5% level. The data set supports the conclusion of statistically significant countercyclical dispersion even when examining IQRs within local markets. To the author’s knowledge this is the first estimate of countercyclical dispersion that controls for geographic heterogeneity.

Implicit in the above regression results is the fact that markets that experi-

enced larger shifts in consumption towards low priced stores are the markets that experienced larger increases in the dispersion of firm growth rates. This relationship can also be shown directly. For each market I compute the difference between the average relative price of consumption in 2007 and 2009 and regress this on the difference between the average IQR of the growth rates of stores' sale in 2007 and 2009. Results of this regression are presented in Table 1.6. They show a negative relationship between the change in a market's relative price of consumption and the dispersion in stores' growth rates which is significant at the 10% level.

| Indep. Var. | Coeff. | Std. Err. | p-value |
|---------------------|---------|-----------|---------|
| Δ Rel. Price | -0.734* | 0.408 | 0.078 |
| cons. | -0.011* | 0.0055 | 0.056 |

Table 1.6: The dependent variable is the change in a markets IQR of the growth rates of total store sales between 2007 and 2009. There are 50 observations. ***, **, * indicate significance at 1%, 5%, and 10% levels, respectively.

As a whole the results from this empirical analysis are consistent with demand elasticity being an important driver of countercyclical dispersion in firm growth rates. Households increase shopping effort and shift consumption to low priced stores during recessions and the impact of a recession is worse for firms with higher

prices. In addition, consistent with changing demand elasticity, the dispersion of stores' growth rates becomes larger during recessions and this change is stronger in markets where the change in consumer behavior is the largest.

1.4 General Equilibrium Model

In this section I formalize the partial equilibrium intuition outlined in section 1.2 into a general equilibrium business cycle model with heterogeneous firms where demand elasticity is determined by household shopping behavior. The model's purpose is twofold: first, the model serves as a tool to explore the quantitative importance of changing demand elasticities on firm-level growth rates; and second, the model's various predictions can be compared to the empirical results presented above.

1.4.1 Representative Household Problem

In standard macroeconomic models demand elasticity is assumed to be a time invariant preference parameter, commonly formalized in a Dixit and Stiglitz (1977) consumption aggregator. However, in a setting where it is costly for consumers to obtain the best prices available, demand elasticity will be determined, at least in part, by the extent of these product market frictions. Given the empirical

evidence of price dispersion across identical goods these product market frictions seem to be non-trivial.¹⁹ To capture the effect of changing consumer behavior on firm-level growth rates I depart from standard time invariant demand elasticities and instead allow demand elasticities to be determined by household shopping effort.

The representative household has period utility given by:

$$\frac{1}{1-\gamma}C_t^{1-\gamma} + \chi(1 - N_t - \mathcal{K}d_t) \tag{1.5}$$

C_t is aggregate consumption and is defined as $C_t \equiv \sum_{j=1}^N c_{j,t}$, where $c_{j,t}$ is the amount consumed from firm(s) with productivity j at time t . There are a continuum of firms offering the identical consumption good c . These firms are indexed by their level of productivity j . Since the firms are offering the same good, households prefer to consume the good from the firm with the lowest price. N_t is aggregate labor hours provided to the market by households, d_t is the amount of time spent shopping for goods, and \mathcal{K} measures the disutility of shopping.

When thinking about modeling product market frictions it is natural to think of a search problem: households need to search to gather information about prices. Indeed, a large literature in the 1970's and 80's used search models to generate equilibrium price dispersion. These models, however, often struggle to sustain

¹⁹For recent empirical evidence of price dispersion across identical goods see Kaplan and Menzio (2014).

price dispersion.²⁰ In addition, it is not entirely clear that all of shopping behavior is characterized by search. It is plausible that households have some information about which stores are generally less expensive and even which products are cheaper at specific stores.²¹ However, even with information about prices it still might be costly for households to obtain the best prices because doing so may involve substantial travel time. This is not to say that search (gathering information about prices) is not important, it is unreasonable to assume that households have perfect information about prices, but it is to say that the costs associated with obtaining the best prices are likely a combination of search and simply travel time to obtain the prices whose locations are already known. Because of this I do not take a stand on *why* it is costly for households to obtain the best prices in the market. I simply model product market frictions in a reduced-form fashion that captures the notion that households prefer to consume the lowest priced goods available but obtaining these best prices is costly.

In a local market there are a number of options (prices) available to a consumer for an identical good or bundle of goods. In each market there also exists a unique optimal (lowest priced) option. Frictions in product markets imply that not all demand is captured by the firm with the lowest price. To model these fric-

²⁰For an early discussion on these issues see Rothschild (1973).

²¹Using a store's monthly relative price I run AR(1) regressions to get a sense of the stability in a store's relative expensiveness. The coefficients in Table A.5 in Appendix A show that a store's monthly relative price is highly persistent with AR(1) coefficients ranging from 0.86 to 0.96. It seems plausible that households would have some information regarding the expensiveness of local stores through repeated shopping experience.

tions I borrow a technique commonly used in microeconomic decision modeling. In accounting for observed subject responses in various game settings, theorists often assume that there is a cognitive cost associated with arriving at the specific utility maximizing option and that modeling responses as a probability distribution around the optimal choice describes the data well. In this “control-cost” approach, the more a subject invests in decision making the tighter his or her distribution will be around the optimal choice.²² This approach has recently been used in macroeconomics to model firm pricing behavior, see Costain and Nakov (2013) and Costain and Nakov (2015). This way of modeling decisions maps well to the setting of product markets because there are a wide array of options (stores) each with potentially different payoffs (prices) for identical or near-identical goods. It captures the notion that the more effort invested in shopping (either through search or travel time) the more likely the consumer is to choose the optimal (lowest priced) option.

Specifically, I model this shopping cost by assuming that each household in the economy acts “as if” they are choosing a probability distribution (or mixed strategy) over the set of available prices $[\underline{p}_j, \bar{p}_j]$ in the market. In this sense, the more mass a probability distribution has over low prices the more favorable this “strategy” is for households because it gives them a lower expected price. As will be shown below, the cost associated with choosing a distribution is inversely

²²For more on control-costs see Stahl (1990) and Mattsson and Weibull (2002)

proportional to the expected consumption price. This captures the notion that while households can reduce the prices they pay for their consumption bundle by increasing their shopping effort, doing so comes with a cost.

It is assumed that the cost of choosing a probability distribution over the set of available prices is proportional to Kullbeck-Libler Divergence, or relative entropy, which is given by:

$$\mathcal{D}(\pi||\eta) = \sum_{j=1}^N \pi_j \ln \frac{\pi_j}{\eta_j}$$

There are N unique firm productivity levels each representing a different price. Here, household shopping costs are measured as the “distance” between the chosen probability distribution π and the underlying price distribution η . When π and η are the same distribution the shopping cost incurred is zero. In equilibrium, η will match the price distribution in the economy, therefore when no shopping effort is expended the household chooses its firm to consume from by randomly drawing from the distribution η . As will be shown below, when households increase their shopping effort the probability distribution becomes more concentrated on low priced firms.

I have not yet been specific as to what a “firm” represents. The natural interpretation given the empirical analysis above is that a firm represents a store and thus households are choosing which stores they would like to consume from. I will use firm and store interchangeably in what follows. However, this characteriza-

tion of consumer decisions could easily be used to model choices across products *within* a store, but in that setting it is less clear why there are costs associated with choosing the lowest priced option.

Assuming that there is a large number of households and that each household chooses the same probability distribution π , it follows that the representative household consumes $\pi_j C_t$ from stores with productivity j , again where C_t is aggregate consumption.

The representative household's time spent shopping is set equal to Kullbeck-Libler Divergence and is thus given by:

$$d_t = \sum_{j=1}^N \pi_j \ln \frac{\pi_j}{\eta_j}$$

The representative household solves the following period maximization problem:

$$\max_{\{C_t, N_t, \pi_{j,t}\}} \frac{1}{1-\gamma} C_t^{1-\gamma} + \chi(1 - N_t - \mathcal{K} \sum_{j=1}^N \pi_{j,t} \ln \frac{\pi_{j,t}}{\eta_{j,t}}) \quad (1.6)$$

s.t.

$$\sum_{j=1}^N p_{j,t} \pi_{j,t} C_t \leq W_t N_t + div_t \quad (\lambda_t) \quad (1.7)$$

$$\sum_{j=1}^N \pi_{j,t} = 1 \quad (\mu_t) \quad (1.8)$$

As can be seen by (1.7) I assume that households are rule-of-thumb consumers. The inclusion of bonds in this model would provide an standard looking Euler equation but would leave the aggregate dynamics of the model unchanged because the model is absent capital in the production process or an outside lender (the government). Since allowing for savings does not change the dynamics of the model I leave it out.²³

The first-order conditions for C_t and N_t yield:

$$\frac{\chi}{W_t} = \frac{C_t^{-\gamma}}{\sum_{j=i}^N p_{j,t} \pi_{j,t}} \quad (1.9)$$

The first-order condition for $\pi_{j,t}$ yields:

$$-\frac{\chi}{W_t} p_{j,t} C_t - \chi \mathcal{K} \left[\ln \frac{\pi_{j,t}}{\eta_{j,t}} + 1 \right] - \mu_t = 0$$

Some re-arrangement yields the following weighted multinomial logit expression:

$$\pi_{j,t} = \frac{\eta_{j,t} \exp \frac{-p_{j,t} C_t}{\mathcal{K} W_t}}{\sum_{j=1}^N \eta_{j,t} \exp \frac{-p_{j,t} C_t}{\mathcal{K} W_t}} \quad (1.10)$$

To further develop the intuition behind these probability distributions over

²³The inclusion of bonds would produce a time series for the interest rate, but that is not a variable of interest to this paper

prices I plot two examples constructed from (1.10) in Figure 1.4. These two different distributions vary only in the disutility of “shopping effort”, \mathcal{K} . With a lower disutility (the red line) it is less costly (in utility terms) to choose optimal (lower priced) stores and thus π is more concentrated on lower prices relative to the distribution with higher disutility (the dashed line). In the limit, as the disutility of shopping approaches zero, the lowest cost brand is consumed with probability one. As will be shown below, the market power of firms arises purely from this product market friction and in the limit where shopping disutility converges to zero the product market converges to one of pure competition.

Since the quantity demanded from stores with productivity j is given by: $c_{j,t} = \pi_{j,t}C_t$, π_t can be thought of as a demand curve: it determines the share of total consumption a firm receives. As is apparent from the different slopes in Figure 1.4, demand elasticity is larger in a low shopping disutility environment relative to a high shopping disutility environment.²⁴

One of the main contributions of this paper is to provide a tractable way to generate demand elasticities from frictions in product markets. With this connection, changes in shopping effort leads to changes in demand elasticities faced by firms.

²⁴This demand curve will be more familiar to mathematicians as quantity is on the y-axis.




Figure 1.4: Plots the steady state equilibrium values of π_j over $[\underline{p}_j, \bar{p}_j]$ for two different shopping disutilities. In this case η is assumed to be a uniform distribution. The red line and dashed line represent low and high shopping disutility environments, respectively.

1.4.2 Firms

There are a continuum of firms denoted by their productivity j . They produce output from a constant returns technology $Y_{j,t} = A_{j,t}N_{j,t}$. There is no capital involved in the production process. In the aggregate steady-state $A_{j,t}$ is a time-invariant idiosyncratic productivity process. Thus each firm's productivity $A_{j,t}$ is correlated with $A_{j,t-1}$ but uncorrelated with the productivity of other firms. When considering an optimal price firms must weigh the trade-off between profits and the probability of being selected by a household for consumption; higher priced

firms capture less market demand. With flexible prices each period firms solve:

$$\max_{p_{j,t}} p_{j,t} c_{j,t} - \phi_{j,t} c_{j,t} \quad (1.11)$$

Firms hire labor at the competitive wage rate W_t and thus $\phi_{j,t} = W_t/A_{j,t}$ is firm j 's nominal marginal cost of production. Substituting in for aggregate consumption and the households decision rule for $\pi_{j,t}$, the firms problem can be written as:

$$\max_{p_{j,t}} p_{j,t} \frac{\eta_j \exp \frac{-p_{j,t} C_t}{\mathcal{K} W_t}}{\sum_{j=1}^N \eta_j \exp \frac{-p_{j,t} C_t}{\mathcal{K} W_t}} C_t - \phi_{j,t} \frac{\eta_j \exp \frac{-p_{j,t} C_t}{\mathcal{K} W_t}}{\sum_{j=1}^N \eta_j \exp \frac{-p_{j,t} C_t}{\mathcal{K} W_t}} C_t \quad (1.12)$$

The first order condition yields:

$$\begin{aligned} \pi_{j,t} C_t + p_{j,t} \left[\frac{-C_t}{\mathcal{K} W_t} \pi_{j,t} \right] C_t - \phi_{j,t} \left[\frac{-C_t}{-\mathcal{K} W_t} \pi_{j,t} \right] C_t &= 0 \\ p_{j,t} &= \phi_{j,t} + \frac{\mathcal{K} W_t}{C_t} \end{aligned} \quad (1.13)$$

Some intuition can be gleaned from (1.13). Namely, the optimal price is equal to a firm's nominal marginal cost plus a mark-up that depends on the disutility of household shopping effort, \mathcal{K} . Thus profits in this model are derived from product market frictions and the model converges to perfect competition as the disutility of shopping approaches zero:

$$\lim_{\kappa \rightarrow 0^+} p_{j,t} = \phi_{j,t}$$

Definition 1 *An equilibrium is a set of household decisions $\{C_t, N_t, d_t, \pi_t\}$ and firm decisions $\{p_{j,t}, N_{j,t}\}$ that satisfy the following conditions:*

1. *Household decisions satisfy (1.9) and (1.10).*
2. *Firm decisions satisfy (1.13)*
3. *Aggregate consistency: Individual decisions are consistent with aggregate variables and the distribution of firm prices.*

1.4.3 Model Solution

Combining the household's first order conditions for consumption and labor gives (1.9). This can be combined with the households optimal choice of $\pi_{j,t}$, given by (1.10), and the firms' optimal pricing decision given by (1.13) to obtain an expression for C_t :

$$\frac{\chi}{W_t} \sum_{j=1}^N \left[\left[\phi_{j,t} + \frac{\kappa W_t}{C_t} \right] \frac{\eta_{j,t} \exp \left[-\frac{\phi_{j,t} C_t}{\kappa W_t} - 1 \right]}{\sum_{j=1}^N \eta_{j,t} \exp \left[-\frac{\phi_{j,t} C_t}{\kappa W_t} - 1 \right]} \right] - C_t^{-\gamma} = 0 \quad (1.14)$$

Normalizing the wage to one, (1.14) becomes a non-linear equation for C_t that

is only a function of the current productivity state. Given a productivity state, this equation can be solved numerically to obtain C_t .

1.4.4 Calibration

In the dynamic version of the model firm productivity is determined both by idiosyncratic shocks and aggregate shocks. It is important to ensure that the dispersion of idiosyncratic productivities is independent from aggregate productivity. This eliminates any possibility that changes in distribution of firm-level growth rates that are observed are a result of changes in the underlying productivity distribution. Specifically, I assume that $A_{i,t} = z_{i,t} + \psi_t - 1$ and that both evolve according to a mean zero AR(1) process in logs:

$$\log z_{i,t} = \rho_z \log z_{i,t-1} + \epsilon_{z,i,t}$$

$$\log \psi_t = \rho_\psi \log \psi_{t-1} + \epsilon_{\psi,t}$$

In the economy wide steady-state, aggregate shocks, ψ_t , are not active and only firm specific shocks, $z_{i,t}$, are active. Therefore, in the steady-state the productivity distribution of firms has converged to its ergodic distribution centered around an aggregate productivity $\psi_t = 1$. The additive separability of $A_{i,t}$ insures that, in response to aggregate shocks, the first moment of the ergodic productivity

distribution shifts while the higher moments remain unchanged.

Shutting down aggregate shocks, the variance of the firm productivity is given by:

$$\sigma_z^2 = \frac{\sigma_{\epsilon_z}^2}{1 - \rho_z^2} \quad (1.15)$$

The autocorrelation coefficient for both productivity processes is chosen to be 0.95. This matches the estimate from Blundell and Bond (2000) at a monthly frequency. The standard deviation of idiosyncratic firm productivity, σ_z , is set to match the standard deviation of “reference costs” found in Eichenbaum et al. (2011).²⁵ Using (1.15), this implies $\sigma_{\epsilon_z} = 0.000351$. Dividing the seasonally adjusted monthly real GDP time series constructed by Stock and Watson by monthly civilian employment reported by the BLS, I construct a time series of monthly labor productivity.²⁶ I compute the standard deviation of the percent deviation from trend of the HP-filtered data and use this to calibrate σ_{ϵ_ψ} .

²⁵Reference costs (prices) are defined as the most often quoted costs (prices) of a retail good at a specific store within a given time period. Eichenbaum et al. (2011) argue that stores set a reference price to maintain average markups over reference costs and as such, these reference costs are the important factor behind a firm’s pricing decisions.

²⁶Real monthly GDP from Stock and Watson was accessed at http://www.princeton.edu/~mwatson/mgdp_gdi.html

| Parameter | Value | Target |
|--------------------------|---------|---|
| \mathcal{K} | 0.062 | ATUS shopping time |
| γ | 0.7 | cov(GDP, Emp) |
| χ | 1.85 | Steady-state employment= 1/3 |
| σ_{ϵ_z} | 3.51e-4 | Eichenbaum et al. (2011) |
| σ_{ϵ_ψ} | 1.18e-5 | Std. Dev. of monthly agg. prod. |
| ρ_z, ρ_ψ | 0.95 | Blundell and Bond (2000) at monthly freq. |

Table 1.7: Calibration values of model parameters. Further explanations are located in Section 1.4.4

I choose γ so that the model generated covariance between consumption and employment matches that observed empirically. I compute this empirical covariance using seasonally adjusted Stock and Watson monthly real GDP data and civilian monthly employment from the BLS. The disutility of shopping, \mathcal{K} , is chosen so that steady-state shopping time, d_t , matches data from the ATUS that show households spend roughly 5% of their non-sleep hours obtaining goods and services. χ is chosen so that in the steady-state households spend roughly one third of their time working.

1.4.5 Model Results

I discretize the aggregate and firm productivity state spaces into 25 grid points using the method of Tauchen (1986). I shock the model with a 1% positive aggregate productivity shock and plot the response of consumption, shopping effort, the dispersion of firm-level growth rates, and relative consumption prices in Figure 1.5 below. I begin by examining how household shopping behavior changes over the business cycle. To offset a reduction in consumption households can choose to allocate more effort towards shopping as a means to obtain lower prices. This lowers the price of their consumption bundle and allows them to purchase more of the consumption good.

Figure 1.5 shows that shopping effort is countercyclical. Using ATUS data Aguiar et al. (2013) finds that households increase their time spent shopping by approximately 7% during the great recession. The great recessions constituted a roughly 5.5% drop in GDP relative to trend. To give a clearer sense of the magnitude of the response of shopping time I shock the model with a productivity shock that generates a decrease in consumption of 5.5%. This shock generates an increase in shopping effort of 3.5%, half of the change reported in Aguiar et al. (2013).

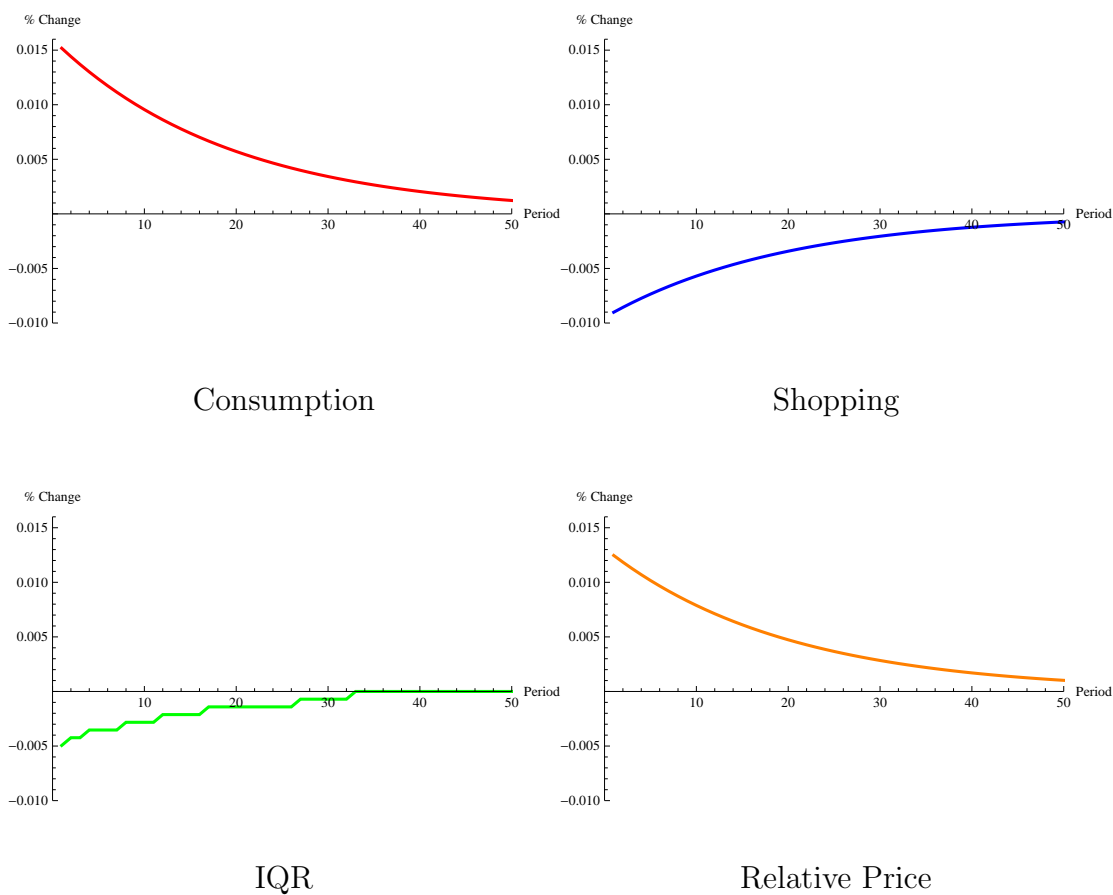


Figure 1.5: Impulse responses from a 1% positive aggregate productivity shock.

Given that model generates sizable countercyclical shopping effort, I now explore the impact this has for the dispersion of firm-level growth rates. Figure 1.5 shows the countercyclical behavior of the IQR of firm-level growth rates. To compare the model to my empirical findings I focus on the estimation result reported in Table 1.5 that controls for different geographical sensitivities of recessions. These results predict that over the great recession the IQR of the growth rates of store

sales would increase by 1.06 percentage points. Comparison of these predictions with the model's results is difficult as there is no unemployment in the model. However, feeding in a productivity shock that generates the 5.5% fall in consumption over the great recession results in an increase in the IQR of firm-level growth rates by 0.28 percentage points.²⁷ Therefore, the model generates a increase in the IQR of firm-level growth rates that is roughly one third of what is found empirically.

While the model predictions and empirical results are qualitatively similar it is worth understanding why the model under-predicts the empirical findings relating to countercyclical dispersion. The empirical evidence presented above shows that the relative price of a consumptions bundle falls during recessions, which is evidence of increasing demand elasticities. These results provide a measure of the magnitude of the change in demand elasticity observed over the business cycle that can be compared to the model results. As in the empirical section I construct the relative price of a firm in the model as the log difference between a firm's price and the median market price. I then compute the mean relative price of consumption by multiplying each firm's relative price by their share of total consumption (π) and sum across all firms to obtain the market relative price of consumption.

Figure 1.5 plots the IRF of the relative price of consumption. Consistent with

²⁷The mean IQR of firm-level sales growth rates is 14% in the model versus the empirical estimate of 10% reported in Table 1.5

the findings in Table 1.3, the model generates a procyclical relative consumption price: households shift purchases to lower priced stores during recessions.

However, the response of relative prices over the business cycle is small relative to what is reported in the empirical analysis above. Shocking the model with a productivity shock that generates a 5.5% decline in consumption results in a 4.4% (.08 percentage point) decline in the relative price of consumption. The empirical estimates using the sampling criteria of S_{100} and a consumption weighted price predict that over the great recession the relative price of consumption should have declined by 0.95 percentage points. As the change in the relative price of consumption is a proxy for the change in demand elasticity it is not surprising that degree of countercyclical dispersion is smaller in the model than what is found empirically.

1.4.6 Sensitivity Analysis

In this section I provide some insights regarding the sensitivity of the model's results to important parameters. Households can offset a fall in consumption by working or by reducing the price of their consumption bundle by exerting more shopping effort. The magnitude of the change in utility derived from consumption governs how labor and shopping will change in the model and this depends crucially, of course, on the curvature of the utility function with respect to con-

sumption.

In the baseline calibration, a relatively low curvature for the consumption utility function was chosen in order to match the empirical covariance between consumption and employment. This low value results in a relatively small increase in shopping effort during a recession. Increasing the curvature of the utility function amplifies the household's desire to work and shop as the fall in consumption during a recession becomes more costly in utility terms.

To illustrate the model's sensitivity to the curvature of the utility function I recalibrate the model with a larger γ . Specifically, I assign the representative household a utility function that is logarithmic in consumption ($\gamma = 1$).²⁸

This new calibration results in shopping time, the IQR of firm-level growth rates, and the relative price of consumption to respond more to a 1% productivity shock than the baseline calibration. Appendix B.2 plots the comparisons of the IRFs for consumption, shopping time, firm IQR's, and relative prices for the baseline calibration and this new calibration.

To give the reader a sense of how the dynamics of this new calibration compare to the empirical estimates reported above I again feed into the model a productivity shock that generates a 5.5% decline in consumption and observe the response of the key variables of interest. The model's responses from this new calibration

²⁸With this new value of γ I adjust χ so that steady-state employment remains at 1/3 and steady-state shopping time remains at 5%

vs. the baseline calibration and the empirical findings are summarized in Table 1.8 below.

The model predicts that shopping time would increase by 12.4% over the great recession, now slightly larger than the 7% increase reported in Aguiar et al. (2013). The IQR of firm-level sales growth would increase by 6.5% (.91 percentage points) which is now very close to the empirical estimate of an 1.06 percentage points increase. And finally, the model predicts a decrease in the relative price of consumption by 11.8% (.19 percentage points) which remains low in magnitude relative to the empirical findings (0.95) but is twice the decrease observed in the baseline calibration.²⁹

²⁹Some readers may also be interested in the behavior of mark-ups in the model. As can be seen from (1.13) the behavior of mark-ups over the business cycle depends on the change in productivity relative to the change in consumption. In the aggregate steady-state the mean mark-up in the model is 16.5%. The dynamics of mark-ups depend crucially on the model's parameterization of γ . For low levels of γ (below 1) the model generates countercyclical mark-ups, which is consistent with findings from Rotemberg and Woodford (1991). However, for $\gamma > 1$ the model generates pro-cyclical mark-ups, consistent with the recent findings of Stroebel and Vavra (2014).

| Variable | Responses | | |
|------------------|-------------|--------------------------|------------------------|
| | Empirical | Model ($\gamma = 0.7$) | Model ($\gamma = 1$) |
| Shopping Time | 7%-8.5% | 3.5% | 12.4% |
| IQR | 1.06% point | 0.28% point | 0.91% point |
| Rel. Cons. Price | 0.95% point | 0.08% point | 0.19% point |

Table 1.8: Empirical vs. Model responses to a recession the magnitude of the Great Recession.

While this new calibration improves the response of variables of key interest to this paper it is important to emphasize that it comes at the cost of matching the joint behavior of consumption and employment over the business cycle. Enriching the model to so that it is able to capture jointly the behavior of employment, consumption, shopping effort, dispersion of firm growth rates, and relative consumption prices is an important avenue of future research.

1.5 Conclusion

This paper provides a novel explanation for countercyclical dispersion in firm-level growth rates. To date, the literature on countercyclical dispersion had focused on information or frictions present on the firms' side of the economy as the

explanation behind this phenomenon. In this paper, I show that countercyclical dispersion can be generated from changes in consumer behavior over the business cycle. The important mechanism that maps consumer behavior to countercyclical dispersion is a change in the demand elasticity faced by firms.

I show that countercyclical dispersion is large in consumer oriented retail industries using synthetic census longitudinal business data. Using a UPC-level data set for retail firms and a consumer panel I find evidence for changing consumer behavior over the business cycle. I find that households take more shopping trips and shop more at low priced stores during recessions. In addition, I find that low priced stores gain a larger share of consumption when local unemployment rises. Consistent with this change in consumer behavior I find that the dispersion in the growth rates of firms' sales increases during recessions and that this increase in dispersion is larger in markets where the shift towards lower priced stores is the strongest.

To explore the implications of these changes in consumer behavior I develop a business cycle model with heterogeneous firms and frictions in product markets. In the model it is costly for households to obtain the best prices in the market: more time spent shopping reduces the price of consumption. These product market frictions are the source of firms' market power allowing them to set prices above nominal marginal costs. With product market frictions shopping effort becomes

an important margin of adjustment for households as it allows them to offset a decline in consumption by paying lower prices. The model captures well the movements of shopping effort, demand elasticity, and the dispersion of firm-level growth rates over the business cycle. In the baseline calibration, the model is able to capture roughly a third of the change in the IQR of firm growth rates observed empirically. The response of the IQR over the business cycle depends crucially on the curvature of the household's consumption utility function. An alternative calibration shows that a small increase in this curvature generates a stronger response in shopping effort over the business cycle leading to movements in the IQR of firm growth rates that is nearly identical to that found empirically.

While this paper highlights the importance of consumer behavior in understanding countercyclical dispersion it is important to emphasize that this is not the only mechanism at play. A large literature documents the importance of firm frictions in understanding countercyclical dispersion. An important and interesting avenue of future research will be to understand the relative importance of firm versus consumer mechanisms in accounting for the changes in the dispersion of firm growth rates and how this may vary across industries. This will involve a deeper understanding of how consumer behavior works back through supply chains to impact wholesale and manufacturing industries. The margin of consumer adjustment examined in this paper is across retail stores and thus may not

impact firms higher up the supply chain in a differential manner. However, there is also evidence that consumers adjust along quality margins over the business cycle, see Jaimovich et al. (2015). This adjustment would plausibly generate more dispersed firm outcomes higher up the supply chain if wholesale or manufacturing firms specialize in low or high quality products. Understanding these connections is an area worthy of further research.

Appendix A

Additional empirical results

A.1 Additional analysis on countercyclical dispersion using SythLBD

| Industry | SIC Code | Description | $Avg(IQR_{Rec.}) - Avg(IQR_{Boom})$ |
|---------------------------|----------|-----------------|-------------------------------------|
| Manufacturing | | | |
| | 201-214 | Food Kind. | 0.0124 |
| | 221-229 | Textile Mill. | 0.0338 |
| | 231-239 | Apparel | -0.0107 |
| | 241-249 | Lumber & Wood | 0.0661 |
| | 251-259 | Furniture Fix. | 0.0341 |
| | 261-267 | Paper Allied | 0.0000 |
| | 271-279 | Printing & Pub. | 0.0360 |
| | 281-289 | Chem. & Allied | 0.0186 |
| | 341-349 | Fabr. Metal | 0.0376 |
| | 361-369 | Elect. Equip. | 0.0143 |
| | 371-379 | Trans. Equip. | 0.0511 |
| Unweighted Average | | | 0.0267 |

Table A.1: Reports the difference between the average interquartile range in firm employment growth rates for different industries in NBER recession and expansion years.

As noted in above, Decker et al. (2013) document that IQR's in employment growth rates have experienced a secular decline in many US industries over the past few decades. Decker et al. (2013) argue that this is evidence of a decline in business dynamism. Since the recession years I examine are not evenly spaced in the SynthLBD data window, it is important to insure that the results reported in Table 1.1 and Table A.1 are not a product of oversampling early years from a downward trend. To control for this, I HP-filter the time series of IQR's and compute their average cyclical component during recession years. I employ a HP parameter of 6.25 for annual data as suggested by Ravn and Uhlig (2002). These results are reported in Table A.2. While the magnitude of countercyclical dispersion is smaller after controlling for this issue it is still present in all the industries I examine.

| Industry | SIC Code | Description | $Avg(IQR_{Rec.HP-Cycle})$ |
|---------------------------|----------|-----------------------|---------------------------|
| Retail | | | |
| | 521-527 | Building Materials | 0.0052 |
| | 531-539 | General Merchandise | 0.0079 |
| | 541-549 | Food Stores | 0.0176 |
| | 551-559 | Auto Dealers & Serv. | 0.0126 |
| | 561-569 | Apparel & Access. | 0.0196 |
| | 571-573 | Home Furnishings | 0.0190 |
| | 581 | Eat & Drink | 0.008 |
| | 591-599 | Miscellaneous | 0.0150 |
| Unweighted Average | | | 0.013 |
| Manufacturing | | | |
| | 201-214 | Food Kind. | 0.0112 |
| | 221-229 | Textile Mill. | 0.0137 |
| | 231-239 | Apparel | 0.0144 |
| | 241-249 | Lumber & Wood | 0.0136 |
| | 251-259 | Furniture Fix. | 0.0114 |
| | 261-267 | Paper and Allied | 0.0056 |
| | 271-279 | Printing & Publishing | 0.0157 |
| | 281-289 | Chemicals & Allied | 0.0104 |
| | 341-349 | Fabricated Metal | 0.007 |
| | 361-369 | Electrical Equip. | 0.0019 |
| | 371-379 | Transportation Equip. | 0.0102 |
| Unweighted Average | | | 0.0105 |

Table A.2: Report the average HP-filtered cyclical component of the interquartile range in firm employment growth rates during NBER recessions. The annual time series is filtered using a HP parameter of 6.25 as suggested by Ravn and Uhlig (2002)

A.2 Additional analysis on households' relative consumption price

I explore how the relative price of a households bundle changes over the business cycle using the following specification:

$$\tilde{R}_{i,m,t} = \alpha_{i,m} + \beta_1 UR_{m,t} + error \quad (A.1)$$

Where $\alpha_{i,m}$ is a household i market m fixed effect and $UR_{m,t}$ is the monthly seasonally adjusted unemployment rate in market m month t . Results are reported in Table A.3 below.

| UPC Sample | Unweighted Consumption Price | Weighted Consumption Price |
|------------|------------------------------|----------------------------|
| S_{75} | -0.0086*** | -0.0048*** |
| S_{90} | -0.0075*** | -0.0041*** |
| S_{100} | -0.0081*** | -0.0044*** |

Table A.3: Reports regression results for specification A.1 for both simple relative prices and standardized relative prices. The dependent variable is the expenditure weighted average price of a household’s consumption bundle in a given month. The coefficients reported in the table are for the local unemployment rate which is seasonally adjusted using the X-12 ARIMA method. There are 127,226 observations. Robust standard errors are clustered at the household level. ***, **, * indicate significance at 1%, 5%, and 10% levels, respectively.

As noted in the main text, these regressions do not control for an cross-sectional dependence that is likely present and are thus only useful for examining how the coefficient estimates compare to those found in the store analysis.

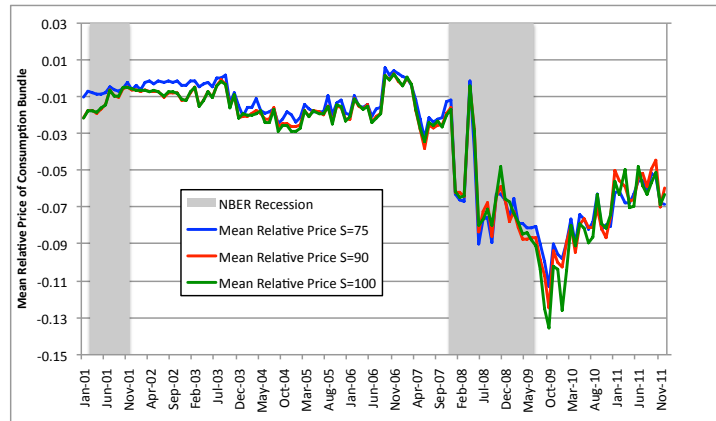


Figure A.1: Displays the simple average relative consumption price averaged across all households for the three UPC sampling criteria, $S = 75$, $S = 90$, and $S = 100$. See main text for a description of these criteria. NBER recession months are shaded in gray.

A.3 Additional Store Analysis

| Coefficient on $UR_{m,t} * \bar{R}_{j,m,2007}$ | | |
|--|------------------------------|----------------------------|
| UPC Sample | Unweighted Consumption Price | Weighted Consumption Price |
| S_{75} | -0.0022** (0.0009) | -0.0025*** (0.0009) |
| S_{90} | -0.0021** (0.0009) | -0.0024** (0.0009) |
| S_{100} | -0.0019** (0.0009) | -0.0024** (0.0009) |

Table A.4: Reports regression results for specification 1.1 for both unweighted and weighted store prices and for the three different UPC sampling criterion. The dependent variable is the market share of store j 's sales, in market m , in month t . The coefficients reported in the table are for the local unemployment rate which is seasonally adjusted using the X-12 ARIMA method. There are 79,048 observations. Robust standard errors are clustered at the market level. ***, **, * indicate significance at 1%, 5%, and 10% levels, respectively.

| AR(1) Coefficient | | |
|-------------------|------------------------------|----------------------------|
| UPC Sample | Unweighted Consumption Price | Weighted Consumption Price |
| S_{75} | 0.957*** | 0.927*** |
| S_{90} | 0.940*** | 0.899*** |
| S_{100} | 0.911*** | 0.855*** |

Table A.5: Reports coefficient estimates of AR(1) regressions of a stores monthly relative price. This regression sheds light on the stability of a stores relative price.

There are 231,445 observations. ***, **, * indicate significance at 1%, 5%, and 10% levels, respectively.

Appendix B

Comparison of IRF's across calibrations




Figure B.1: Plots the impulse response of consumption for the baseline calibration ($\gamma = 0.7$) in blue and an alternative calibration ($\gamma = 1$) in purple from a positive 1% shock to aggregate productivity.




Figure B.2: Plots the impulse response of shopping effort for the baseline calibration ($\gamma = 0.7$) in blue and an alternative calibration ($\gamma = 1$) in purple from a positive 1% shock to aggregate productivity.




Figure B.3: Plots the impulse response of the IQR of firm-level growth rates for the baseline calibration ($\gamma = 0.7$) in blue and an alternative calibration ($\gamma = 1$) in purple from a positive 1% shock to aggregate productivity.




Figure B.4: Plots the impulse response of the relative price of consumption for the baseline calibration ($\gamma = 0.7$) in blue and an alternative calibration ($\gamma = 1$) in purple from a positive 1% shock to aggregate productivity.

Chapter 2

Experience, Skill Composition,

and the Persistence of

Unemployment Fluctuations

(with Aspen Gorry)

2.1 Introduction

Fluctuations in monthly unemployment rates are highly persistent. The autocorrelation of monthly unemployment rates in the U.S. exceeds 0.95 and can be as high as 0.99 for prime age workers. This implies a half-life of shocks to unemploy-

ment ranging between 13 and 69 months.¹ Moreover, the level of persistence has become even higher in recent recessions. Identifying the mechanisms that generate these persistent fluctuations is important for understanding the propagation of shocks over the business cycle. While macroeconomic models account for many patterns in the data, the ability to provide an internal propagation mechanism for shocks remains a major challenge.

Neither standard real business cycle models nor the canonical Mortensen and Pissarides (1994) model successfully propagate shocks to unemployment. Standard real business cycle models generate time series of aggregate variables that closely follow the shock process, leaving the persistence problem largely unexplained.² While search frictions embodied in search and matching models provide an intuitive explanation for persistence, the ability of these models to explain the persistence of unemployment fluctuations depends on both the speed at which workers find jobs when unemployed and separate from their job when employed.

Observed levels of labor market flows imply a half-life of only one to two months.

¹The half-life of fluctuations in unemployment can be calculated from the empirical employment data by estimating an AR(1) process on monthly unemployment rates. Autocorrelations of U.S. monthly unemployment rates generate coefficients that exceed 0.95. Given coefficient ρ , the half-life of unemployment fluctuation is given by:

$$t_{hl} = \frac{\log .5}{\log \rho}$$

resulting in a lower bound estimate for the half-life of unemployment fluctuations from the data is 13.5 months. When $\rho = 0.99$, the half-life is 69.0 months.

²The ability of labor market frictions to provide a propagation mechanism in standard dynamic stochastic general equilibrium models has been explored by Merz (1995), Andolfatto (1996), and den Haan et al. (2000). See Pries (2004) for a discussion.

Even with a sizable decline in job finding rates after 2000, Tasci (2012) shows that the half-life of unemployment fluctuations is still less than two months.³ Existing search and matching models consistent with observed job finding and separation probabilities are therefore unable to generate persistent unemployment fluctuations.

The goal of this paper is to understand how shocks to the level of unemployment are propagated to generate persistent unemployment. We propose a model where changes in the composition of worker types away from their steady state distribution can generate persistence while maintaining the observed high levels of inflows and outflows of unemployment for all groups of workers. To show this, we extend a standard search and matching model to include two types of workers with different steady state unemployment rates. Having two groups is a deliberate simplification that keeps the model tractable and allows the quantitative implications of the mechanism to be easily assessed. While short-run unemployment

³To understand this point consider a standard search model where unemployed workers find jobs at rate f and are separated from their jobs at rate s . Using the continuous time formulation as discussed in Shimer (2012) and Elsby et al. (2009), if all workers start off unemployed then unemployment at time t is given by:

$$u(t) = u^* + (1 - u^*)e^{-(s+f)t}$$

where $u^* = \frac{s}{s+f}$ is the steady state level of unemployment. In this case the rate of convergence of the system is governed by $s + f$ since the half-life of any difference in unemployment from steady state is given by:

$$t_{hl} = \frac{-\log .5}{s + f}$$

Observed worker flows in the U.S. imply that $s + f \approx 0.5$. Therefore, the half-life is just over one month. Even with lower transitions rates of about 0.1 found in many European countries the half life is only about 6 months.

dynamics are governed by flows into and out of unemployment within each group as in previous models, our model generates new long-run unemployment dynamics that are governed by the speed of transitions between groups.

In particular, the two groups can represent experienced and inexperienced workers. They can differ in their productivity when employed, probability of finding a productive match, and their exogenous job separation probability so that they have different steady state rates of unemployment. Even with rapid within group worker flows, the model generates persistent unemployment if the transition rate between groups is slow. This is because compositional changes take a long time to return to their steady state distribution.

Next, we quantitatively assess the importance of these compositional changes. Since different groups of workers have different baseline unemployment rates, the model is calibrated to match life cycle patterns of employment outcomes. By targeting employment outcomes for experienced and inexperienced workers in the model to match those of young and old workers in the data, steady state outcomes replicate empirical age patterns of unemployment rates, job finding probabilities, job separation probabilities, and wages.⁴ Inexperienced workers have a higher

⁴Elsby et al. (2010) document that there are sizable differences in labor market flows by gender, age, race, and education. While a number of these characteristics are fixed, different outcomes by age and education could proxy for differences in skill or experience. Hence, one interpretation matching age patterns of employment outcomes for high school educated workers is that changes in the composition of workers after a shock could reflect the loss of skills during unemployment as in Ljungqvist and Sargent (1998). Moreover, parameterizing the model to match unemployment dynamics by age is appealing as Jaimovich and Siu (2009) show that accounting for the employment experiences of young workers is crucial to understand aggregate

baseline unemployment rate, which drives persistent unemployment fluctuations if there is a compositional change in workers from experienced to inexperienced even though inexperienced workers have higher job finding probabilities.

The calibrated model is simulated for different initial compositions of workers across types to assess the persistence of unemployment fluctuations. Increases in unemployment without changes in the composition of workers across type have the same rapid dynamics as the Mortensen and Pissarides (1994) model. When changes in unemployment include changes in the composition of workers the model generates long-run persistence of unemployment of similar magnitudes documented in the data. Compositional changes generate non-linear rates of convergence to the steady state. The model reproduces the same short-run dynamics as a regular search and matching model, but now has new long-run dynamics. The time to close half of the gap of the initial shock is rapid at 1.7 months. However, when closing the last 10% of the increase in unemployment the half-life rises to above 11 months and above 77 months when closing the final 5%. The non-linearity occurs because when workers are displaced and need to reacquire skills it takes them a long time to learn a new skill or regain skills in order to return to their previous lower average rate of unemployment. The calibration strategy of targeting employment outcomes of young and old workers may also understate the amount of persistence generated as workers who lose skills with job loss may

employment dynamics.

have lower job finding probabilities.⁵

An alternate explanation for the persistence of unemployment is the existence of a thin market externality as first proposed by Pissarides (1992). Such an externality arises in models with skill loss when the fraction of unskilled workers in the unemployment pool increases, causing firms to post fewer vacancies and hence reduce a worker's probability of finding a job. While Pissarides (1992) develops a simple theoretical model to highlight the possibility that these externalities can generate persistent unemployment, their quantitative importance has not been studied.⁶ After showing that compositional changes generate persistent unemployment fluctuations, we assess whether a thin market externality can generate similar levels of persistence. To study the role of such an externality, our baseline model where experienced and inexperienced workers have separate matching

⁵Inexperienced workers in the model have short unemployment durations corresponding to young workers who find jobs rapidly. In contrast, Valletta (1991) shows that high tenure workers have longer spells of unemployment following job displacement and Kletzer (1998) finds that displaced workers have an average unemployment duration of 17 weeks compared to just 7.2 for workers who are just laid-off. These high durations could arise due to hope by workers that they will regain their lost jobs, a buffer stock of assets, learning about future job quality as in Gorry (2012), or the combination of skill loss and unemployment insurance as in Ljungqvist and Sargent (1998). More generally, a large literature on displaced workers shows that such workers have worse outcomes than other unemployed workers. See for instance, Jacobson et al. (1993) who find long term wage losses for displaced workers and Stevens (1997) who shows that more frequent job loss explains an important part of the average wage loss experienced by displaced workers.

⁶Thin market externalities have been studied in Wasmer (2004). Such externalities can also give rise to multiple equilibria as noted by Diamond (1982), Howitt (1985), and Mortensen (1989). Despite the potential for multiple equilibria, there exists a unique steady state equilibrium for reasonable parameterizations of the model developed in this paper. This is consistent with estimates of the matching function that do not exhibit sufficient increasing returns to generate multiple equilibria as noted by Pissarides (1986), Blanchard and Diamond (1989), and Petrongolo and Pissarides (2001).

functions is modified so that both groups search for jobs in the same market. In this environment, fluctuations in labor market tightness arising from the thin market externality quantitatively generate only a small amount of persistence. The intuition for this result is that the aggregate match rate is bounded between the match probability of each type of worker when they have separate matching functions. Since each group of workers has a high job finding probability, a shock that increases the number of inexperienced workers in the unemployment pool has only a modest impact on persistence.

Finally, to assess the business cycle implications of the model, simulations are run where all experienced workers who lose their jobs unexpectedly become inexperienced for 18 months. While there are no new shocks added to the model, this exercise is consistent with interpreting business cycles as a time when job loss leads to skill loss among workers. The baseline model with separate matching functions and the model with a thin market externality generates both an increase in unemployment and a slight increase in job separation probabilities that decline slowly after the 18 month period. While the baseline model has the counterfactual implication that job finding probabilities increase as unemployment rises because inexperienced workers find jobs more rapidly, the model with a thin market externality generates lower job finding probabilities. These simulations suggest that while the thin market externality does not generate substantial persistence on

its own, it may be important in explaining observed cyclical patterns in worker flows.⁷

In related work, Pries (2004) shows that persistence can be generated by workers learning about the quality of a new job match. Learning implies that unemployed workers have rapid turnover on new jobs because with some probability they learn that they are unproductive soon after starting a new job. Additionally, explanations of persistent unemployment from heterogeneity have previously been discussed in Pries (2008) and Ravenna and Walsh (2012). This paper compliments previous explanations as the model generates both an increase in job separation probabilities and predictions about the cyclicalities of job finding probabilities.

While persistence has always been a feature of unemployment fluctuations, it has increased during the great recession. See Elsby et al. (2010) and Elsby et al. (2011) for a summary of labor market outcomes during the great recession and Coibion et al. (2013) for a discussion of the recent increase in persistence of unemployment fluctuations. The aim of this paper is to understand the mechanism that generates these persistent fluctuations rather than understanding how persistence has changed over time.

The remainder of this paper is structured as follows. Section 2 presents the

⁷Another way to reconcile cyclical patterns in worker flows is to have workers who lose their jobs experience lower job finding probabilities, rather than the higher ones assumed. The model could easily accommodate this by adding a third group of workers who have longer unemployment durations.

model. Section 3 describes the parameterization of the model. Section 4 presents the results on persistence, the effect of a thin market externality, and business cycle implications of the model. Section 5 concludes with a discussion of the results and their relation to explanations for unemployment during the great recession and recent jobless recoveries.

2.2 Model

This section presents the baseline model of heterogeneous workers who have different steady state unemployment rates. Experienced workers and inexperienced workers search for jobs in separate markets. The model is designed to match life-cycle patterns of unemployment as in Gorry (2013). The description of the model does not include any shocks. Alternatively, changes in the initial composition of workers across groups will be considered to measure the time to converge back to the steady state. Heterogeneity in worker types is the key feature that allows the model to generate persistent unemployment fluctuations through changes in the composition of workers. In the results section, the model will be modified to have a single matching function to understand the effects of a thin market externality on the persistence of unemployment fluctuations.

2.2.1 Setup and Worker Flows

Time is discrete. In any period there is a unit mass of workers who maximize the present discounted value of their consumption stream and discount the future at rate β . Workers can be either employed or unemployed and experienced or inexperienced. Inexperienced workers become experienced while employed with probability α and remain experienced until they exit the labor force. Workers leave the labor market at rate δ and are replaced by a new cohort of inexperienced, unemployed workers. This assumption prevents all workers from becoming experienced in the steady state model.

There is a continuum of infinitely lived firms that can search for workers by posting vacancies for either an experienced or an inexperienced worker at flow cost k per vacancy. Production occurs when a worker is paired with a firm. Workers of each type search for firms in a separate market characterized by a constant returns to scale matching function $m(v_i, u_i) = u_i^\eta v_i^{1-\eta}$ where $i \in \{e, n\}$. Let e denote experienced and n denote inexperienced workers. $\theta_i = \frac{v_i}{u_i}$ denotes the tightness of the labor market for workers of type i . With this matching function, an unemployed worker meets a job in a given period with probability $\lambda(\theta_i) = m(v_i, u_i)/u_i = \theta_i^{1-\eta}$ and open vacancies are matched with a worker with probability $q(\theta_i) = m(v_i, u_i)/v_i = \theta_i^{-\eta}$.

When a worker and a firm meet there is a probability that the match is produc-

tive. Experienced matches are productive with probability p_e and inexperienced matches are productive with probability p_n . These probabilities enable job finding probabilities to differ for experienced and inexperienced workers.⁸ When a worker and firm of type $i \in \{e, n\}$ form a productive match they produce y_i units of output. In general we assume that $y_e > y_n$. Nash bargaining determines wages for both types of workers.

With this setup, workers of type $i \in \{e, n\}$ find jobs with probability $(1 - \delta)\lambda(\theta_i)p_i$. Worker separations arise from labor force exit and exogenous employment separation shocks. Experienced workers separate from their jobs with probability $\delta + (1 - \delta)s_e$ and inexperienced workers separate with probability $\delta + (1 - \delta)(1 - \alpha)s_n$. Also, with probability $(1 - \delta)\alpha$ inexperienced workers become experienced, remaining employed.

2.2.2 Value Functions and Equilibrium

Value functions for unemployed and employed workers of each type are as follows:

$$U_n = b + \beta(1 - \delta) [\lambda(\theta_n)(p_n E_n + (1 - p_n)U_n) + (1 - \lambda(\theta_n))U_n] \quad (2.1)$$

⁸Alternately, differences in job finding rates could be generated by differences in the cost of posting vacancies k across different types of workers. While the results are identical for the baseline model, this setup simplifies the analysis when considering the model with a single matching function to understand the quantitative relevance of thin market externalities.

$$U_e = b + \beta(1 - \delta) [\lambda(\theta_e)(p_e E_e + (1 - p_e)U_e) + (1 - \lambda(\theta_e))U_e] \quad (2.2)$$

$$E_n = w_n + \beta(1 - \delta) [\alpha E_e + (1 - \alpha)(s_n U_n + (1 - s_n)E_n)] \quad (2.3)$$

$$E_e = w_e + \beta(1 - \delta) [s_e U_e + (1 - s_e)E_e] \quad (2.4)$$

Unemployed workers get flow value b and move to employment with probability $\lambda(\theta_i)p_i$ if they do not exit the labor market. b can be interpreted as some combination of unemployment benefits, the value of leisure, and the value of home production. When they become employed they get their employment value E_i .

Inexperienced employed workers receive their wage w_n and with probability α become experienced employed in the next period. When they do not become experienced they are separated from their job with probability s_n becoming unemployed inexperienced. Experienced employed workers receive wage w_e and are separated from their jobs with probability s_e when they do not exit the labor market.

Next, firms can choose to open vacancies to meet workers and search directly

for inexperienced or experienced workers. Their value functions are as follows:

$$V_n = -k + \beta q(\theta_n) p_n J_n \quad (2.5)$$

$$V_e = -k + \beta q(\theta_e) p_e J_e \quad (2.6)$$

$$J_n = y_n - w_n + \beta(1 - \delta) [\alpha J_e + (1 - \alpha)(s_n V_n + (1 - s_n) J_n)] \quad (2.7)$$

$$J_e = y_e - w_e + \beta(1 - \delta) [s_e V_e + (1 - s_e) J_e] \quad (2.8)$$

Firms post vacancies at period flow cost k . Jobs are then created if the workers and firms form a productive match. Inexperienced matches produce output y_n and the firm pays the worker wage w_n . In each period an inexperienced worker becomes experienced with probability α and when she does not become experienced the worker and firm separate with probability s_n . Likewise, experienced matches produce output y_e and earn w_e and workers are separated with probability s_e each period.

In this economy, a steady state equilibrium is defined as follows:

Definition 2 *A steady state equilibrium consists of the value functions for the worker, U_n , U_e , E_n , and E_e , the value functions of the firm, V_n , V_e , J_n , and J_e , the aggregate state variables, u_n , u_e , e_n , e_e , θ_n , and θ_e :*

1. *Value functions are satisfied: Given w_n , w_e , u_n , u_e , θ_n , and θ_e , then U_n , U_e , E_n , E_e , V_n , V_e , J_n , and J_e satisfy equations (2.1)–(2.8).*

2. *Match Formation:* Given $w_n, w_e, u_n, u_e, \theta_n$ and θ_e , it is optimal for workers to form productive matches.
3. *Free Entry:* The value of posting a vacancy for each type of worker is given by $V_n = V_e = 0$.
4. *Bargaining:* w_n and w_e are determined by Nash bargaining equations with weight γ given to workers:

$$E_n - U_n = \gamma[J_n + E_n - U_n]$$

$$E_e - U_e = \gamma[J_e + E_e - U_e]$$

5. *Steady State:* The following four worker flow equations hold:

$$\delta + (1 - \delta)(1 - \alpha)s_n e_n = (\delta + (1 - \delta)\lambda(\theta_n)p_n)u_n$$

$$(1 - \delta)\lambda(\theta_n)p_n u_n = (\delta + (1 - \delta)\alpha + (1 - \delta)(1 - \alpha)s_n)e_n$$

$$(1 - \delta)s_e e_e = (\delta + (1 - \delta)\lambda(\theta_e)p_e)u_e$$

$$(1 - \delta)\lambda(\theta_e)p_e u_e + (1 - \delta)\alpha e_n = (\delta + (1 - \delta)s_e)e_e$$

The steady state equilibrium can easily be solved. For details see the appendix.

2.3 Paramaterization

This section parameterizes the model to match key features of life-cycle patterns of unemployment rates in the United States. Matching life cycle employment outcomes provides discipline on model parameters. The approach is similar to the one used in Gorry (2013). The model period is assumed to be one month. Therefore, $\delta = \frac{1}{480}$ so that the expected length of time in the labor market for each worker is 40 years. The discount rate is set using $\beta(1 - \delta) = 0.9967$ to match an annual interest rate of 4%. As normalizations, $y_n = 1$ and $p_n = 1$ so that y_e is interpreted as the relative productivity of experienced workers and p_e is the relative probability that a match is productive for experienced workers compared to inexperienced.

The matching function takes the standard Cobb-Douglas form, $m(u, v) = u^\eta v^{1-\eta}$. η is set to 0.5. This value is at the lower end of the range of estimates found in Petrongolo and Pissarides (2001). The choice of $\gamma = \eta$ insures that the Hosios (1990) condition applies.

Next, observed job separation probabilities are used to set s_e . Micro-data from the Current Population Survey (CPS) is used to construct job finding and job separation probabilities for high school educated workers.⁹ Targeting only

⁹Throughout the paper job finding and job separation probabilities are measured in the same way as in Shimer (2012). Job finding probabilities are constructed from transitions between unemployment and employment (U to E) while job separation probabilities are constructed from transitions between employment and unemployment (E to U) in a three state model that includes employment, unemployment, and out of the labor force. The quarterly flows from

high school educated workers insures that the age patterns observed are due to experience rather than changes in composition as workers of different skills enter the labor force. The separation probability for experienced workers can be set directly from the measured job separation probability of 50-54 year old workers. The separation probability solves: $0.011 = \delta + (1 - \delta)s_e$. This gives $s_e = 0.009$.

The remaining parameters of the model are the productivity of experienced workers y_e , the separation probability for inexperienced workers s_n , the probability that a match is productive for experienced workers p_e , the probability with which experienced workers gain experience α , the value of unemployment b , and the cost of posting a vacancy k . These parameters are calibrated jointly to match targets about individual wage growth, job finding and job separation probabilities, and unemployment benefits.

The following targets are used. First, y_e is set to match the amount of wage growth observed in the data. Using data from the Merged Outgoing Rotation Groups (MORG) from the CPS the mean hourly wage for 18-year-old workers from 2002-2007 is \$8.44 and the mean hourly wage for 50-54 year-old workers is \$16.18 (both values are in 2009 dollars). Therefore, the wage for experienced workers is targeted to be 1.92 times the wage of inexperienced workers. The monthly

the procedure in Shimer (2012) are averaged over the period from 2002-2007 to get the values reported here. This period corresponds to an average for the expansion preceding the Great recession. Since the model only has employed and unemployed workers, focusing on U to E and E to U transitions makes the data more consistent with the model.

| Parameter | Value | Target |
|-----------|--------|---------------------------------------|
| δ | 1/480 | 40 year working life |
| β | 0.999 | Annual Interest rate of 4% |
| η | 0.5 | Petrongolo & Pissarides (2001) |
| p_n | 1 | Normalization |
| p_e | 0.766 | Ratio of job finding probabilities |
| y_n | 1 | Normalization |
| y_e | 1.76 | Wage growth from MORG |
| s_n | 0.037 | 20-24 year-old separation probability |
| s_e | 0.009 | 50-54 year-old separation probability |
| α | 0.0076 | Share of experienced workers is 0.78 |
| b | 0.45 | $b = 0.5w_n$ |
| k | 9.42 | 50-54 finding probability of 0.32 |

Table 2.1: Baseline values for model parameters along with targets.

job separation probability for 20-24 year old individuals is 3.9%. Using the flow equation for separations the following target is used: $0.039 = \delta + (1 - \delta)(1 - \alpha)s_n$. Third, $(1 - \delta)\theta_n^{1-\eta}p_n = 0.316$ is targeted to match the job finding probability for 20-24 year old workers of 31.6%. Next, since the mean hourly wage from 2002-2007 in the MORG for 18-64 year-old workers is \$14.46 this implies a target of the fraction of experienced workers in the population to be $\frac{e_e}{e_n + e_e} = 0.78$ so that the average wage in the model matches the average wage in the data. The flow value of unemployment is targeted to be half of the wage of inexperienced workers. Finally, $(1 - \delta)\theta_e^{1-\eta}p_e = 0.274$ is targeted to match the job finding probability for 50-54 year old workers.

These targets imply parameter values of $y_e = 1.76$, $s_n = 0.037$, $\alpha = 0.0076$, $b = 0.45$, and $k = 9.42$. Steady state wages generated by these parameters are

$w_e = 1.71$ and $w_n = 0.89$. The parameters are summarized along with their calibration targets in Table 3.3.

2.4 Numerical Results

This section reports the numerical results from the steady state model. First, the life-cycle outcomes of the model are simulated to demonstrate that the parameterized model matches observed patterns of unemployment and worker flows. Next, the ability of compositional changes in the distribution of workers across states to generate persistent unemployment is assessed. To do so, the dynamics of the model are solved for the case where 1% of workers employed in the steady state become unemployed and inexperienced. With this formulation, we compute the time required for the model to converge back to the steady state. After an initial period of quick convergence of workers finding new jobs, the model generates substantial persistence in unemployment rates. Because the baseline unemployment rates are higher for inexperienced compared to experienced workers, increases in the share of inexperienced workers in the economy leads to a persistent increase in unemployment. Next, the model is modified to have only one matching function to assess the ability of a thin market externality to generate persistence. Finally, to interpret the cyclical dynamics of worker flows, both models are simulated for an 18 month period where all workers who lose jobs unexpectedly become inex-

perienced. While the thin market externality alone does not generate meaningful persistence in unemployment, it does generate the lower job finding probabilities observed during recessions.

2.4.1 Unemployment, Wages, and Worker Flows

With the steady state values of θ_e and θ_n the flow equations can be solved for the steady state number of workers in each state $\{u_n, u_e, e_n, e_e\}$. Table 2.2 summarizes the number of workers in each state. Using these figures, the total steady state unemployment rate in the model is 5.6%, while the unemployment rates for inexperienced and experienced workers are 14.8% and 3.4% respectively. With these baseline unemployment rates the model matches average levels of unemployment among high school educated workers in the United States between 2002 and 2007.

| State | Quantity |
|-------|----------|
| u_n | 0.031 |
| u_e | 0.025 |
| e_n | 0.209 |
| e_e | 0.735 |

Table 2.2: Steady state results for share of population in each state in the economy.

The most important parameter in the model to determine the persistence of unemployment fluctuations is the rate at which workers become experienced, α , as it determines how long it takes workers to transition to the group with




Figure 2.1: Monthly wages simulated from the model and compared to the data in five year age bins.

lower steady state levels of unemployment. The calibration chooses α to match the average wages for high school workers by targeting the share of experienced workers in the economy. Figure 2.1 shows the simulated pattern of wage growth by age compared with mean wages for each 5 year age group from CPS MORG data. The parameterized model generates much of the observed wage growth in the data.

In order to assess the quality of this target, α also determines how quickly job finding and separation probabilities decline over the life-cycle.¹⁰ A strength of the model is that it is consistent with age patterns of job finding and job

¹⁰While the endpoints of both job finding and separation rates are targeted through other parameters of the model, α determines the life cycle patterns of job finding and separation probabilities between the endpoints.

separation probabilities. The left panel of Figure 2.2 shows average unemployment to employment transition probabilities for high school educated workers in the United States between 2002 and 2007 for each five year age group from 20-24 through 50-54. The figure also shows average job finding probabilities for each age for 10,000 worker outcomes simulated from the model where each worker enters the labor force unemployed and inexperienced at age 18. The model captures much of the observed decline in the job finding probability by age. Observed job separation probabilities by age are shown in the right panel of Figure 2.2. The simulated model closely replicates the job separation probabilities by age observed in the data, providing additional evidence that the calibrated value of α is reasonable.

2.4.2 Persistence

This section assesses the ability of compositional changes in the distribution of workers across states to generate persistent unemployment fluctuations. As discussed in the introduction, the half-life of convergence can be computed separately for inexperienced and experienced workers given their steady state worker flows. For inexperienced workers, the baseline calibration implies that $s = 0.04$ and $f = 0.38$.¹¹ This implies that the half-life for changes in the unemployment

¹¹This comes from converting the targeted monthly job finding and job separation probabilities into rates. The formula for the rate (given in lower case letters) is given by: $s = -\log(1 - S)$. Here S is the monthly probability. An analogous equation holds for f .




Figure 2.2: Monthly job finding probabilities (left panel) and job separation probabilities (right panel) for high school educated workers in the United States by five year age group and simulated from the model.

rate is 1.65 months. For experienced workers, $s = 0.011$ and $f = 0.32$ gives a half-life of 2.09 months. The short duration of deviations generated by each group is similar to the lack of persistence generated in standard search models that are calibrated to match the observed levels of worker flows. Individually, neither group of workers exhibits persistent deviations in their unemployment rates. Even with the observed decline in transition rates during the great recession standard search and matching models are unable to account for the observed levels of persistence.

To understand how much persistence in unemployment is generated by compositional changes, the model is simulated for monthly employment outcomes after

1% of employed workers (both experienced and inexperienced) from the steady state distribution start off unemployed. Two scenarios are considered. In the first scenario, workers remain in their original experience group. That is, 1% of experienced employed workers start experienced unemployed and 1% of inexperienced employed workers begin inexperienced unemployed. This scenario is referred to as no skill loss. In the second scenario, the 1% of experienced workers who start off unemployed also begin inexperienced. This scenario is referred to as skill loss in the results and documents the main mechanism for persistence in the model.

The 1% of employed workers who begin unemployed increases the unemployment rate by nearly one percentage point from the steady state rate of 5.59% to 6.53%. After computing the share of workers in each state, the unemployment dynamics of the model can be easily computed using the following first order difference equations that give the number of workers in each state in the next period:

$$u'_n = u_n + \delta + (1 - \delta)(1 - \alpha)s_n e_n - (\delta + (1 - \delta)\lambda(\theta_n)p_n)u_n$$

$$e'_n = e_n + (1 - \delta)\lambda(\theta_n)p_n u_n - (\delta + (1 - \delta)(1 - \alpha)s_n + (1 - \delta)\alpha)e_n$$

$$u'_e = u_e + (1 - \delta)s_e e_e - (\delta + (1 - \delta)\lambda(\theta_e)p_e)u_e$$

$$e'_e = e_e + (1 - \delta)\lambda(\theta_e)p_e u_e + (1 - \delta)\alpha e_n - (\delta + (1 - \delta)s_e)e_e$$

In the above equations, u'_n , e'_n , u'_e , and e'_e denote the values for the number of workers in each state in the next period. Because there are separate matching functions for each group, θ_e and θ_n do not depend on the composition of workers across states in the economy so the simulation is simple to execute.

Figure 2.3 plots the monthly unemployment rate for five years after the increase in unemployment for each scenario. While the simulations do not specify the shock that generates the increase in unemployment, the graphs can be interpreted as impulse response functions to changes in the composition of workers across states from their steady state distribution. The gray line depicts the steady state unemployment rate of 5.6%. In both simulations, the initial unemployment rate is 6.5%. The dashed line shows that when there is no skill loss the unemployment rate converges rapidly back to the steady state level of unemployment. The dotted line shows the monthly unemployment rates for the scenario with skill loss. The figure demonstrates two results. First, compositional shocks generate substantial persistence in unemployment, as the unemployment rate does not fully converge back to the steady state level after five years. Second, the convergence generated by the model is highly non-linear. In the first few months unemployment declines rapidly in both scenarios (in fact, it declines even more rapidly in the skill loss scenario as inexperienced workers have high job finding rates), but the rate of convergence slows dramatically in the scenario with skill loss as it takes workers




Figure 2.3: Monthly unemployment rate in response to 1% of employed workers starting out unemployed for the cases with and without skill loss.

a long time to become experienced. The larger fraction of inexperienced workers can generate a persistent increase in unemployment.

To get a better sense of how unemployment converges after compositional changes, Table 2.3 reports a number of measures of the speed of convergence for each scenario. Because of the non-linearity in convergence, the half-life is no longer a sufficient statistic for the speed of convergence in the case of skill loss. Therefore, the number of months it takes for the unemployment to close 50%, 75%, 90%, 95% and 97.5% of the initial shock are reported. In the case of no skill loss, the speed of convergence is rapid with a half-life of approximately 2 months. This rate of change remains nearly constant as the time to go from 90 to 95%

| Simulation | No Skill Loss | Skill Loss |
|----------------|---------------|------------|
| Half-life | 1.9 | 1.7 |
| Converge 75% | 3.9 | 3.6 |
| Converge 90% | 6.5 | 7.0 |
| Converge 95% | 8.5 | 18.5 |
| Converge 97.5% | 10.5 | 95.9 |

Table 2.3: Time to return to steady state unemployment rate from 1% of employed workers starting unemployed compared to the steady state distribution of workers for the cases with and without skill loss.

and 95 to 97.5% are each 2 months (each of these differences represents closing half of the remaining distance to the steady state). In contrast, the results for the case with skill loss are highly non-linear. For the half-life, there is less persistence than in the case with no skill loss as it takes 1.7 months to close half of the initial shock. This occurs as newly unemployed workers quickly converge to the baseline unemployment rate for inexperienced workers. This quick convergence continues through closing 75% of the gap, then slows down dramatically after closing 90%. It takes 11.5 months to go from 90-95% and over 77 months to go from 95-97.5%. A portion of the initial shock to unemployment remains highly persistent as it takes workers a long time to gain experience.

2.4.3 Thin Market Externality

An alternate explanation for persistent unemployment fluctuation is the presence of a thin market externality as proposed by Pissarides (1992). The intuition

is that as the composition of the pool of unemployed workers deteriorates firms have a lower incentive to post vacancies. Therefore, when unemployment pool has more low quality workers job finding probabilities are low and unemployment can remain higher than it otherwise would. Such externalities do not arise in the baseline model as experienced and inexperienced workers search for jobs in separate labor markets. While Pissarides (1992) develops the possibility of such an externality generating persistent unemployment fluctuations, the quantitative relevance of this channel has never been assessed.

To assess the role of a thin market externality, the model is modified so that there is a single matching function for both types of workers. The thin market externality arises as both experienced and inexperienced workers search for jobs in the same labor market. Since workers become experienced through employment, low rates of unemployment lead to a higher fraction of experienced workers in the unemployment pool. Experienced workers have higher work productivity and are more valuable to firms. Therefore, firms post more vacancies when the composition of the unemployment pool is better.

Specifically, it is assumed that both workers now match using the same constant returns to scale matching function $m(v, u) = u^\eta v^{1-\eta}$. Let $u = u_e + u_n$ be the aggregate number of unemployed workers where u_i is the number of unemployed workers of type $i \in \{e, n\}$. $\theta = \frac{v}{u}$ denotes the tightness of the labor market.

The fraction of experienced workers in the unemployment pool is denoted by $\mu \equiv \frac{u_e}{u_n + u_e}$. Given that firms post vacancies of a single type, their value of posting vacancies is given by:

$$V = -k + \beta q(\theta)[(1 - \mu)p_n J_n + \mu p_e J_e]$$

With probability $q(\theta)$ an open vacancy meets a worker who with probability μ is experienced and with probability $1 - \mu$ is inexperienced. For the model to generate an externality it is assumed that experienced workers are more productive than inexperienced ones so that $y_e > y_n$. Moreover, it must be the case that $J_e > J_n$ so that experience workers are more valuable to firms. A sufficient condition for this to be the case is that $s_n \geq s_e$ and $y_e - w_e \geq y_n - w_n$ with one strict inequality. The second inequality holds in a standard Nash bargaining solution.

By assumption firms cannot search separately for experienced workers. This assumption overstates the potential of the thin market externality to account for persistence as any ability to sort workers reduces the externality from changes in the quality of the pool of unemployed workers. To the extent that labor markets are able to sort workers, these externalities would be less important even though there is a fair amount of segmentation by education and experience.

In addition to having full segmentation, we assume that wages for each type of worker, w_e and w_n , are fixed at the steady state level of the baseline model.

If wages were allowed to adjust, they would lessen the impact of the thin market externality as workers in bad labor markets are willing to accept lower wages, partially offsetting the lower incentive for firms to post vacancies. While both of these assumptions may not be realistic, they generate an upper bound on the amount of persistence the externality can generate in this model.

With this setup, the model is reparameterized to match the same targets. Since there is only one matching function, the ratio of job finding rates for younger and older workers implies that $p_e = 0.866$. All of the remaining parameters are identical except for the cost of posting vacancies k . Since k must now account for the possibility of firms meeting different types of workers, targeting the job finding probability for young workers implies a value of $k = 10.6$.

To assess the role of the thin market externality, the model is simulated as follows. First, the proportion of experienced workers in the pool of unemployed workers, $\mu = \frac{u_e}{u_n + u_e}$ is computed. Second, the zero profit condition is used to find the labor market tightness θ associated with the current value of μ by finding the value of θ that solves:

$$0 = -k + q(\theta)\beta(1 - \delta)(\mu p_e J_e + (1 - \mu)p_n J_n)$$

Finally, the following first order difference equations are used to find the next

periods number of workers in each state:

$$u'_n = u_n + \delta + (1 - \delta)(1 - \alpha)s_n e_n - (\delta + (1 - \delta)\lambda(\theta)p_n)u_n$$

$$e'_n = e_n + (1 - \delta)\lambda(\theta)p_n u_n - (\delta + (1 - \delta)(1 - \alpha)s_n + (1 - \delta)\alpha)e_n$$

$$u'_e = u_e + (1 - \delta)s_e e_e - (\delta + (1 - \delta)\lambda(\theta)p_e)u_e$$

$$e'_e = e_e + (1 - \delta)\lambda(\theta)p_e u_e + (1 - \delta)\alpha e_n - (\delta + (1 - \delta)s_e)e_e$$

Where u'_n , e'_n , u'_e , and e'_e denote the number of workers in each state in the next period. Using the new values for u_n and u_e , the simulation method can be repeated to generate a monthly time series for θ and unemployment rates.

In order to compare the results of the models with and without an externality, two simulations are conducted. First, the model with an externality is simulated for the same skill loss scenario as presented above. This simulation explains how much more persistence the model with a thin market externality can generate than the baseline model. Second, to assess the ability of the externality alone to generate persistence, both models are simulated for the case where all of the 1% of workers who begin unemployed come from the pool of inexperienced workers. This scenario has an identical effect on the composition of the pool of unemployed workers as the simulation with skill loss, but does not change the composition

of experienced and inexperienced workers in the workforce. Therefore, this scenario assesses how much persistence in unemployment fluctuations a thin market externality generates on its own.

Figure 2.4 plots the results from the skill loss simulations for the baseline model and the model with a single matching function. The dotted line replicates the results from Figure 2.3 showing that the baseline model has quick convergence initially followed by persistent unemployment. The model with a single matching function shows a similar pattern. Unemployment begins at 6.5% before quickly dropping below 5.7%. After the initial decline, the model with the externality generates persistent unemployment. While the patterns are very similar, the figure shows that the thin market externality only slightly increases the persistence of unemployment fluctuations over the persistence generated by compositional changes alone.

Next, Figure 2.5 plots the monthly unemployment rate in response to a shock where all additional unemployment relative to the steady state comes from inexperienced workers for both the baseline model and the model with a single matching function. Even though the fraction of unemployed workers who are experienced, μ , is the same as in the skill loss simulation, the unemployment rate quickly converges back to the steady state level for both models. While this is expected for the baseline model, the fact that the model with a single matching function




Figure 2.4: Monthly unemployment rate in response to 1% of employed workers starting out unemployed and inexperienced for the baseline model and the model with a thin market externality.

converges rapidly implies that the externality alone does not generate persistent unemployment fluctuations.

Finally, the rates of convergence from each simulation are reported in Table 2.4. The first column replicates the rates of convergence from the baseline skill loss scenario reported in Table 2.3. Next, the table shows that the externality only modestly increases the half-life from 1.7 months to 2.0 months. However, when converging the final 10% the half-life increases from 11.5 to 31.3 months and even further from 77.4 to 78.8 months for the final 5%. While the model with an externality generates more persistence, it does not substantially alter the amount of persistence generated by the model. The final two columns show the




Figure 2.5: Monthly unemployment rate from the distribution where all additional workers who start unemployed come from inexperienced employed workers with and without the thin market externality.

rates of convergence for the case where all of the additional workers who start out unemployed come from the group of employed inexperienced workers. The baseline model converges very rapidly as inexperienced workers quickly regain their steady state level of unemployment. The half-life for convergence is 1.6 months at all durations. In the case with an externality it takes moderately longer to converge, but convergence is still rapid. Closing half the distance to the steady state occurs in 1.8 months and does not vary much as the model approaches the steady state. The thin market externality alone does not generate the level of persistence observed in the data.

| Simulation | Skill Loss | Skill Loss, Externality | n Job Loss | n Job Loss, Externality |
|----------------|------------|-------------------------|------------|-------------------------|
| Half-life | 1.7 | 2.0 | 1.6 | 1.8 |
| Converge 75% | 3.6 | 4.2 | 3.2 | 3.6 |
| Converge 90% | 7.0 | 8.5 | 5.3 | 5.9 |
| Converge 95% | 18.5 | 39.8 | 6.9 | 7.8 |
| Converge 97.5% | 95.9 | 118.6 | 8.5 | 9.6 |

Table 2.4: Time to return to steady state unemployment rate from 1% of employed workers starting out unemployed for each scenario. Note that “n Job Loss” refers to job loss for inexperienced workers.

2.4.4 Interpreting Business Cycles

This section uses the model to assess how unemployment recovers after a period of unanticipated skill loss. This exercise can be interpreted as a means to understand how unemployment recovers after a recession that generates skill loss among workers. Skill loss is a plausible outcome of recessions as the portion of workers who separate from their jobs due to layoffs increases while the portion of workers who quit declines. Davis et al. (2012) show that during the past recession quits dropped dramatically from their pre-recession high in 2006 of nearly 8 percent of employment to under 5.5 percent of employment by the end of the recession. At the same time, layoffs increased from about 6 percent to over 8 percent of employment.

To explore this feature of recessions in the model, these patterns are interpreted as experienced workers losing their skills when they are separated from a job. This section simulates a recessionary episode where experienced workers who are

separated from their jobs unexpectedly lose their skills by becoming inexperienced. The subsequent recovery is then simulated as in the previous sections. When thinking about the results of this exercise and recent recessions, a number of caveats are in order. First, the simulations assume that all workers lose their skills during the recession and no workers become inexperienced during the recovery. While in practice there are certainly workers who lose their skills and do not in any given period, these extreme assumptions clarify the mechanisms in the model. Second, the exercise does not change the magnitude of the shocks in the model. It is assumed that the probability of job separation remains constant for both types of workers over the business cycle. Hence, it will not attempt to generate the magnitude of fluctuations in unemployment observed during the recession. The benefit of this approach is that the value of filled matches of each type remains constant, which makes the model easier to solve. The simulation does not generate as much unemployment as observed during the past recession. Finally, the simulation assumes that the skill loss that moves experienced workers to inexperienced is unanticipated so that the value functions remain unchanged from those previously described. While modifying expectations modestly changes the dynamics of the system, the purpose of the simulation is to evaluate the persistence in the recovery rather than identify the shock that caused the recession.

We first present the time series of unemployment for the baseline model and the




Figure 2.6: Monthly unemployment rate simulated from 18 months of skill loss with job separation for each model. 18-month recession shaded in gray followed by five year recovery.

model with a single matching function. Figure 2.6 plots monthly unemployment rates for the period before the recession, an 18 month recession where workers who are separated from their jobs also lose their skills, and the following five years of recovery. The 18-month recessionary period is chosen to match the length of the great recession and is shaded in gray. In the baseline model the initial impact of the skill loss is for the unemployment rate to go *down*. This is the case because inexperienced workers have higher job finding rates than experienced ones, so the separations with skill loss lead to lower average unemployment durations. This effect dominates for a few months until the unemployment rate begins to increase due to the compositional effect of a now higher portion of workers who

are inexperienced and hence have higher unemployment rates. In the model with a single matching function where firms reduce the average number of vacancies per unemployed worker, the initial decline in unemployment is almost completely muted. The reduction in job finding probabilities for all workers implies that there is a larger increase in the unemployment rate from the change in the composition of the unemployment pool. In both cases, unemployment continues to rise after the recessionary period of skill loss ends as the job separation rate increases due to the compositional change in the workforce. Moreover, the recovery in each case is eventually characterized by persistence in that unemployment only slowly returns to its steady state level, as the composition of workers across experience groups is slow to recover.

Finally, we look at the cyclical patterns of job finding and job separation probabilities from each simulation. Aggregate job finding and job separation probabilities are computed using the composition of the pools of unemployed and employed workers in each period multiplied by the probability of each type of workers experiencing a change in their employment status. The left panel of Figure 2.7 plots the pattern of job finding probabilities for both simulations. Here, the baseline model generates the counterfactual result that job finding probabilities increase during the period of skill loss. This occurs due to the assumption that inexperienced workers have higher job finding probabilities than experienced ones. If

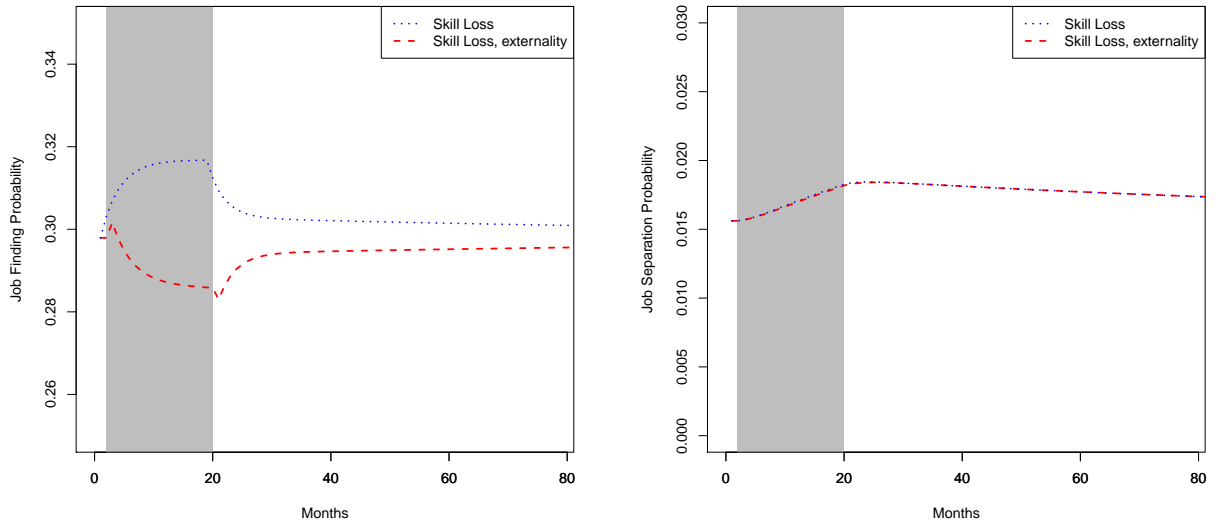


Figure 2.7: Monthly job finding probabilities (left panel) and job separation probabilities (right panel) for 18 months of skill loss from job separations simulated from each model. 18-month recession shaded in gray followed by five year recovery.

there were a third state for displaced workers with lower job finding probabilities this result could be reversed. With the thin market externality job finding probabilities move in the opposite direction. This is consistent with the empirical evidence that job finding probabilities are procyclical. Simulated results for job separations are depicted in the right panel of Figure 2.7. Both simulations generate nearly identical patterns of job separation probabilities as they slowly increase during the recession as the fraction of workers who are inexperienced (with high job separation probabilities) increases.

2.5 Discussion

The goal of this paper is to quantitatively assess potential channels to generate persistent unemployment fluctuations in search and matching models. The results suggest that compositional changes among heterogeneous groups of workers with different baseline unemployment rates generate persistent unemployment fluctuations. This explanation is related to the heterogeneity explanation explored in Ravenna and Walsh (2012) and the learning story in Pries (2004). This paper compliments previous explanations as it generates a theory of long-run unemployment fluctuations that also has predictions about the cyclicity of both job finding and job separation probabilities. Learning implies that periods of high unemployment are persistent due to higher than normal job separation probabilities. However, Shimer (2012) shows that variation in job finding rates are an important component of cyclical unemployment fluctuations. The cyclical properties of job finding and job separation probabilities in this paper depend on the separate probabilities for inexperienced and experienced workers. The baseline parameterization of the model where inexperienced workers have higher finding and lower separation probabilities implies changes in the composition of workers generates counterfactually high job finding rates. However, the model can generate higher job finding rates either with a thin market externality that amplifies persistence or by including additional worker types with different job finding probabilities.

While the focus of this paper is to understand the theoretical propagation mechanism that can generate persistent unemployment fluctuations rather than the shocks that cause unemployment to change, it relates to a number of papers that seek to understand changes in unemployment during the great recession. For a summary of the labor market with a focus on worker flows through the recession see Elsby et al. (2010) and Elsby et al. (2011). An increase in long-duration unemployment is a key feature in the recent US recession and has been a constant feature of higher rates of unemployment in Europe. While the baseline model does not account for increases in long-duration unemployment spells, a thin market externality can increase the duration of unemployment for all workers. Moreover, the modeling framework is flexible enough to accommodate more groups of workers that could generate long-duration unemployment. Another possible explanation for the deterioration of labor market conditions is mismatch as described in Shimer (2007). A large literature has attempted to assess the role of mismatch in increased unemployment after the recession, but has only found modest effects.¹²

When assessing the evidence for compositional changes proposed in this paper with respect to the thin market externality proposed by Pissarides (1992) there are a number of pieces of evidence to consider. First, the explanations are not exclusive in that both can play important roles in explaining cyclical patterns

¹²See papers by Barlevy (2011), Herz and van Rens (2011), and Sahin et al. (2012) among others.

of unemployment outcomes. Second, even with generous assumptions about the size of the externality including a single labor market, assuming that all workers become inexperienced when shocks hit the economy, and fixed wages to magnify their effect, the thin market externality alone only generates moderate amounts of persistence. In contrast, compositional changes can generate substantial persistence on their own that can be enhanced through an externality. Mueller (2012) provides further evidence between these mechanisms by showing that during recessions the pool of unemployed workers is composed of more workers who were separated from high wage jobs. While this evidence makes a thin market externality less likely as the composition of the unemployment pool is improving, such separations could still generate persistent unemployment fluctuations if they lose skills when they separate.

The mechanism of compositional changes in workers across skills is potentially related to a recent literature on job polarization.¹³ In particular, Jaimovich and Siu (2012a) show that the disappearance of jobs in occupations in the middle of the skill distribution has been concentrated during recessions. They argue that this factor contributes to jobless recoveries, but could also contribute to compositional changes where workers previously employed in middle skill occupations can no longer find jobs in that area. Therefore, job polarization could contribute to the compositional story proposed to account for the persistence of unemployment

¹³For a discussion see Acemoglu (1999), Autor et al. (2006) among many others.

fluctuations proposed in this paper.

Finally, in attempting to understand how high observed labor market flows can be reconciled with persistence in the unemployment rate, this paper has abstracted away from the influence of policies on labor market outcomes. Ljungqvist and Sargent (1998) and Pries and Rogerson (2005) show that policies can have important effects on the level of worker turnover. In relation to explanations that focus on the role of policy, this paper provides a complimentary explanation that emphasizes the compositional role of skill differences for unemployment outcomes in the absence of policy differences. Exploring how policy interacts with heterogeneity and labor market shocks is an intriguing avenue for future study.

Appendix C

Steady State Model Solution

This section shows the steps taken to solve for the steady state equilibrium of the model. Using the free entry condition for experienced worker firms, J_e can be solved for using (2.8):

$$J_e = \frac{y_e - w_e}{1 - \beta(1 - \delta)(1 - s_e)}$$

Subtracting (2.2) from (2.4) yields:

$$E_e - U_e = \frac{w_e - b}{1 - \beta(1 - \delta)(1 - s_e - \lambda(\theta_e)p_e)}$$

Substituting this into the Nash Bargaining solution yields:

$$w_e = \frac{\gamma y_e(1 - \beta(1 - \delta)(1 - s_e - \lambda(\theta_e)p_e))}{1 - \beta(1 - \delta)(1 - s_e - \gamma\lambda(\theta_e)p_e)} + \frac{(1 - \gamma)b(1 - \beta(1 - \delta)(1 - s_e))}{1 - \beta(1 - \delta)(1 - s_e - \gamma\lambda(\theta_e)p_e)}$$

$E_e - U_e$ can be substituted into (2.4) to solve for E_e :

$$E_e = \frac{1}{1 - \beta(1 - \delta)} \left[w_e - s_e\beta(1 - \delta) \frac{w_e - b}{1 - \beta(1 - \delta)(1 - \theta_e^{1-\eta}p_e - s_e)} \right]$$

Following the same approach, subtracting (2.1) from (2.3) gives:

$$(1 - \beta(1 - \delta)(1 - (1 - \alpha)s_n - \lambda(\theta_n)p_n))(E_n - U_n) = w_n - b + \alpha\beta(1 - \delta)(E_e - E_n)$$

Solving for $E_e - E_n$ and substituting into the above equation yields:

$$E_n - U_n = \frac{w_n - b + \alpha\beta(1 - \delta) \frac{(1 - \beta(1 - \delta))E_e - w_n}{1 - \beta(1 - \delta)(1 - \alpha)}}{A}$$

where:

$$A = 1 - \beta(1 - \delta)(1 - \lambda(\theta_n)p_n - (1 - \alpha)s_n) - \frac{\alpha(1 - \alpha)(\beta(1 - \delta))^2 s_n}{1 - \beta(1 - \delta)(1 - \alpha)}$$

Equation (2.7) and the zero profit condition combined with the solution for J_e

implies that J_n is given by:

$$J_n = \frac{1}{1 - \beta(1 - \delta)(1 - \alpha)(1 - s_n)} \left(y_n - w_n + \frac{\alpha\beta(1 - \delta)}{1 - \beta(1 - \delta)(1 - s_e)}(y_e - w_e) \right)$$

Finally, plugging these into the Nash bargaining equation and solving for w_n gives:

$$w_n = \frac{AC\gamma \left(y_n + \frac{\alpha\beta(1-\delta)(y_e-w_e)}{1-\beta(1-\delta)(1-s_e)} \right) + B(1-\gamma)(Cb - \alpha\beta(1-\delta)(1-\beta(1-\delta))E_e)}{B(1-\gamma)(C - \alpha\beta(1-\delta)) + AC\gamma}$$

where:

$$B = 1 - \beta(1 - \delta)(1 - \alpha)(1 - s_n)$$

and

$$C = 1 - \beta(1 - \delta)(1 - \alpha)$$

To solve for the steady state of the model, the above equations for J_n and J_e can be substituted into the value functions for vacancies with the zero profit condition imposed:

$$J_n = \frac{k}{\beta q(\theta_n) p_n} \tag{C.1}$$

$$J_e = \frac{k}{\beta q(\theta_e) p_e} \tag{C.2}$$

Solving the steady state flow equations as a function of θ_i provides an expression that can be substituted into the zero profit conditions. For any given set of parameters, these conditions determine the equilibrium number of workers in each state.

Chapter 3

Unemployment Composition and Aggregate Dynamics

3.1 Introduction

Unemployment has been slow to recover following recent recessions relative to prior post-war recessions. To highlight these differences, Figure 1 plots the gap between the unemployment rate and the pre-recession Congressional Budget Office measure of the long-run natural rate of unemployment for each post-war recession.¹ What is apparent in the figure is that unemployment in the past three recessions (plotted in red) peaks later and is slower to return to pre-recession levels.

¹These gaps are normalized to the peak gap for each post-war recession for ease of comparison. This normalization controls for differences in the size of unemployment increases across recessions and allows for a simple visual comparison of the pace of unemployment recoveries.

The pace of the unemployment recovery following the 2008 recession is the slowest recovery to date. In addition, the 1990 and 2001 unemployment recoveries are also slow relative to the average post-war recession. Following the 1990 recession the discussion of sluggish labor market recoveries began, and has since remained a topic of discussion and research. While various explanations for slow labor market recoveries have been examined, there has not yet emerged a comprehensive account for this increase in persistence.² In this paper I explore the extent to which compositional changes in the pool of unemployed workers, and how those differ across recessions, can account for slow unemployment recoveries.

In attempting to understand the observed increase in unemployment persistence it is important to examine the changes in the underlying flows between employment, unemployment, and non-participation in the labor market. Numerous papers have examined the driving forces behind the cyclical variation of the unemployment rate. Darby et al. (1986) argue that cyclical variation in U.S. unemployment is almost entirely driven by cyclical variation in the inflow rate (job separations). Shimer (2005) and Hall and Milgrom (2005) argue the opposite, that cyclical variation in unemployment is completely driven by changes in the outflow (job finding) rate. Recent papers by Elsby et al. (2009) and Fujita and Ramey (2009)

²See Berger (2012) and Coibion et al. (2013) for a discussion on slow labor market recoveries and various explanations. There is also a recent literature examining the decline in labor market fluidity over the past few decades, see e.g. Hyatt and Spletzer (2013) and Davis and Haltiwanger (2014).

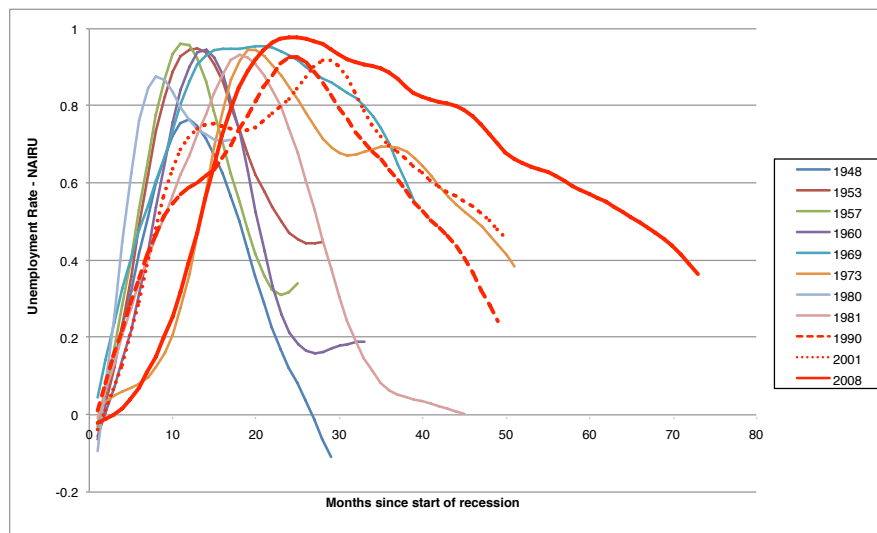


Figure 3.1: Plots the monthly unemployment gap (unemployment rate - NAIRU) from the start of the recession for all post-war recessions. All time series are normalized to the peak gap and are HP-filtered to eliminate some of the noisiness in the series. With month data, a low HP filter weight of 100 is used so not to mask the lower frequency trends in the data.

argue that while the pro-cyclical job finding rates account for the majority of the cyclicity in unemployment, countercyclical inflows are also quantitatively important. These recent studies find that outflows account for roughly two-thirds of the cyclical variation in the unemployment rate where inflows account for roughly one-third. Given these findings and the fact that inflows are quantitatively important early in recessions (see Elsby et al. (2009)), understanding the increased persistence of unemployment in recent recoveries involves accounting for the changes in the aggregate job finding rate during recoveries.

The aggregate job finding rate masks substantial heterogeneity in the job finding rates for different categories of unemployed workers. The distinction of un-

employed workers that I focus on in this paper is workers' reported reason for unemployment. I focus on this distinction because research in both micro and macro labor has found that "permanent displacements," workers who have been permanently dismissed from their previous employment, have adverse labor market outcomes.³

Disaggregating unemployment outflow rates by reason for unemployment in Census Population Survey (CPS) data confirms this general finding. Namely, that unemployment workers who report permanent displacement as their reason for unemployment have substantially lower job finding rates than other reported reasons for unemployment, see Figure 3.3 below. This heterogeneity in job finding rates across different groups leads to the possibility that compositional changes in the pool of unemployed can lead to fluctuations in the observed aggregate job finding rate.⁴

In this paper, I examine the extent to which compositional shifts towards

³Micro studies (e.g. Jacobson et al. (1993) and Schoeni and Dardia (2003)) find that permanently displaced workers face adverse earnings and re-employment probabilities for many years following a termination. Consistent with this, macro studies find that permanently displaced workers have substantially lower job finding rates than workers who are unemployed for other reasons (e.g. Elsby et al. (2009)).

⁴The extent to which this "heterogeneity hypothesis" accounts for the overall cyclicity in the job finding rate is explored in Baker (1992), Elsby et al. (2009), and Shimer (2012). Baker (1992) and Shimer (2012) find that changes in the composition of the pool of unemployed can impact the aggregate job finding rate but that it constitutes a modest fraction of the overall observed cyclicity. However, Elsby et al. (2009) finds that disaggregating flows by reason for unemployment is important for understanding the cyclicity of aggregate unemployment. I highlight that the strength of these compositional changes varies widely between recessions and I explore to what extent these compositional changes can account for the observed differences in the pace of unemployment recoveries.

permanent displacements varies across recessions, and the degree to which this compositional factor can account for the differences in the pace of unemployment recoveries. I find that the compositional shift in the unemployment pool towards permanent displacements is the largest in the 2008 recession relative to the prior four recessions. Due to the important re-design in the 1994 CPS I restrict my analysis to comparing the 1981 versus 1990 and 2001 versus 2008 recessions and corresponding recoveries. As the recovery following the 1990 recession was the first of the slow unemployment recoveries, it provides an interesting comparison to the 1981 recovery. I find that compositional changes are able to account for approximately 20% of the excess deterioration in the job finding rate observed in the 1990 recovery. While the 2000 and 2008 recessions can both be characterized as having slow labor market recoveries, what is apparent in Figure 3.1 is that the unemployment recovery following the 2008 recession was substantially slower than the 2000 recovery. I find that compositional changes are able to account for approximately 24% of this excess deterioration in the job finding rate observed in the 2008 recovery. In addition to these simple compositional shifts, I find evidence that the cyclical sensitivity of job finding rates is larger for permanent displacements which strengthens the quantitative importance of this compositional mechanism.

This finding that recent recessions are characterized by both stronger compositional shifts towards permanent displacements and slow unemployment recov-

eries relates to findings in Berger (2012). He finds that countercyclical business restructuring can generate slow employment recoveries and provides indirect evidence that this restructuring behavior has become more prevalent since the 1980s. I also explore why these compositional shifts towards permanent displacements have been stronger in recent recessions. I find that in 1976 the majority (73%) of “temporary layoffs” originated from “blue collar” occupations even though those occupations only accounted for 37% of all jobs. Temporary layoffs also enjoy high job finding rates relative to permanent displacements, see Figure 3.3 below. By 2013, however, the employment share of “blue collar” occupations had nearly halved, representing 21% of total employment. This occupational trend, often referred to as “job polarization,” sheds some light onto why compositional shifts towards permanent displacements have been stronger in recent recessions and why they also have been characterized by low job finding rates.⁵

The focus on heterogeneity also makes this paper closely related to Elsby et al. (2009) and Barnichon and Figura (2011). Elsby et al. (2009) is focused on understanding the relative importance of inflows versus outflows to/from unemployment in understanding movements in aggregate unemployment. Looking at recessions up to 2001, they find that “job losers” inflow and outflow rates are important to understand the overall cyclicity of unemployment. In this paper, I focus on

⁵See, among others, David et al. (2003) and Jaimovich and Siu (2012b), for more on job polarization.

the magnitude of compositional changes in the pool of unemployment towards “job losers” during recessions and to what degree differences in these compositional shifts can account for differences in aggregate job finding rates observed in recessions. Barnichon and Figura (2011) examines the causes behind fluctuations in aggregate matching efficiency. It relates to this paper in that they find that compositional changes in the pool of unemployment, especially in reasons for unemployment, are important drivers of the cyclical nature of matching efficiency.

To further assess the quantitative importance of these compositional changes, I construct a labor search model with heterogeneous workers. With the model I explore the significance of compositional shifts across reasons for unemployment on the pace of unemployment recoveries. In addition, the model allows for an examination of the relative importance of simple compositional shifts to groups with different steady-state job finding rates and the added importance of the differences in the cyclical sensitivity of job finding rates across these groups. I find that the model generates unemployment recoveries that are substantially slower than standard search models, and that allowing for differences in the job finding elasticities generates most of this increase. In addition, the model is able to generate substantial increases in the standard deviation of various labor market indicators. These findings highlight the importance of incorporating labor market heterogeneity in macroeconomic models.⁶

⁶Pries (2008) and Ravenna and Walsh (2012) also highlight the importance of labor market

The remainder of the paper is structured as follows: Section 2 discusses the data, constructs job finding rates across unemployment groups, and provides an analysis of the compositional changes across recoveries; Section 3 outlines the model, solution method, and steady-states; Section 4 discusses the model calibration; Section 5 discusses the results of the model; and Section 6 provides a discussion of the connection of these findings to existing literature and discusses policy implications.

3.2 Empirics

To construct job finding rates and explore the compositional changes in the pool of unemployed workers I use data from the Census Population Survey (CPS) from 1976 to 2013.⁷ The time structure of the CPS is such that households are surveyed for four sequential months, out of the sample for eight months, then surveyed again for another four sequential months. The CPS is a rolling panel, in that new households are added each month and households who have completed their time are dropped. In total, households are surveyed for eight months and are commonly distinguished by the number of months they have been in the panel,

heterogeneity, but focus on a different source of heterogeneity than the one highlighted in this paper.

⁷Unfortunately, data on reason for unemployment are, to my knowledge, unavailable prior to 1976. This prevents analysis of the quantitative importance of the composition mechanism in earlier recessions

known as “rotation groups.” Job finding rates are computed as in Shimer (2012):

$$F_t = 1 - \frac{u_{t+1} - u_{t+1}^s}{u_t} \quad (3.1)$$

Where u_t is the number of unemployed workers at time t and u_{t+1}^s is the number of short term unemployed workers.⁸ The 1994 CPS redesign presents issues for the consistency of the measure of short-term unemployed.⁹ One approach to correct for this issue is to calculate a short term unemployment series using only workers who are in the first or fifth rotation groups.¹⁰ As noted in Elsby et al. (2009) and Shimer (2012) this correction can generate substantial variation in this time series as one is effectively reducing the sample size to one quarter of that available. This susceptibility to sample variation is especially salient when one is grouping workers into subgroups, as I will be doing below. I follow the correction methodology used in Elsby et al. (2009). After calculating the monthly series for long term unemployment and short term unemployment I seasonally adjust the data using the U.S. Census Bureau X-12 ARIMA program. I construct job finding rates from these monthly series.

⁸More precisely, u^s are workers who have been unemployed for less than four weeks

⁹See Polivka and Miller (1998) and Abraham and Shimer (2001)

¹⁰The measurement of short term unemployed in these rotation groups was not changed by the 1994 redesign.

3.2.1 Unemployment by Reason

The 1994 CPS redesign also modified how the survey accounted for reasons of unemployment. Because of this discontinuity I split the data into pre and post 1994 eras. In the pre 1994 era I calculated unemployment for the following categories: job losers on layoff, other job losers (permanent), job leavers, re-entrants, and new-entrants. For the post 1994 era I compute the same five categories plus workers who are unemployed because temporary work ended. After tallying total unemployment for each reason, I compute the share of the unemployment pool represented by each type and plot the time series in Figure 3.2.¹¹




Figure 3.2: Shares of unemployment pool by reason for unemployment pre and post 1994 redesign

There are a few interesting features of the data worth discussing. The first is that since the twin recessions of the 1980s, the share of workers on temporary layoff shows very little pro-cyclicality. The importance of temporary layoffs in determining the speed of the unemployment recoveries is highlighted in Groshen

¹¹Before 1994 these unemployed workers would have been counted as either “other job losers” or “job losers on layoff.”

and Potter (2003). If permanent layoffs force workers to transition to employment in new sectors and if this transition takes time, a reduction in the use of temporary layoffs may result in slow unemployment recoveries. However, Aaronson et al. (2004) use an alternative empirical approach and show that sectoral reallocation does not set the 1990 and 2001 recessions apart from earlier ones. Though measured sectoral reallocation may not be quantitatively important in the 1990 and 2001 recessions, less use of temporary layoffs may still lead to slower unemployment recoveries if these unemployed workers suffer from low job finding rates for other reasons besides costly reallocation. The other prominent feature of the data is the magnitude of the rise of permanent job loss in the 2008 recession. Table 3.1 displays how the unemployment share of permanent job losers changed during recent recessions. The compositional shift in the unemployment pool towards these permanent job losers is nearly three times larger in the 2008 recession than in the 1980 recession. As will be documented in the next section, it is this class of unemployed workers who experience the worst job finding probabilities.

3.2.2 Job Finding Rates by Reason

As noted above, important consideration must be given to the method used to correct the short-term unemployment series after the 1994 CPS redesign. To reduce the sampling noise of job finding rates by reason for unemployment, I

| Recession | Share at start of recession | Share at Peak of recession | Peak change in unemployment share of permanent job losers |
|-----------|-----------------------------|----------------------------|---|
| 1980 | 28% | 35% | 7% point increase |
| 1982 | 33% | 41.8% | 8.8% point increase |
| 1980-1982 | 28% | 41.8% | 13.8% point increase |
| 1991 | 32% | 45% | 13% point increase |
| 2000 | 19% | 36.5% | 17.5% point increase |
| 2008 | 25% | 47% | 22% point increase |

Table 3.1: Reports pre and post-recession unemployment shares of permanent job losers along with the peak changes in these shares.

employ the discrete adjustment method used by Elsby et al. (2009). Specifically, I find the ratio of short-term unemployed to total unemployed for the first and fifth rotation groups and the ratio of short-term unemployed to total unemployed for all rotation groups in each period. I then compute the average of the ratio of those ratios:

$$x_i \equiv \frac{\frac{1}{T} \sum_{i=0}^T \frac{U_{1,5}^{is}}{U_{1,5}^i}}{\frac{1}{T} \sum_{i=0}^T \frac{U_{All}^{is}}{U_{All}^i}} \quad (3.2)$$

where $U_{1,5}^{is}$ and U_{All}^{is} are the number of short term unemployed in rotation groups 1 or 5 and all rotation groups i , respectively, and $U_{1,5}^i$ and U_{All}^i are the total number of unemployed in rotation groups 1 or 5 and all rotation groups i respectively. These discrete corrections are multiplied by the short-term unemployment time series for each group i to obtain a redesign adjusted time series. Averaging across all time periods when computing this adjustment to the short-term unemployment reduces sampling noise in the CPS data.

The left panel of Figure 3.3 displays the job finding rates pre 1994 for the five reasons for unemployment listed above, and the right panel of Figure 3.3 displays the job finding rates post 1994 for the six reasons for unemployment listed above. These time series show substantial heterogeneity in job finding rates across reasons for unemployment. In the pre 1994 period permanently displaced workers consistently have job finding rates that are roughly 40% lower than other groups. In the post 1994 redesign period this discrepancy ranges between 25% and can be as high as 60% depending upon the comparison group. Given the cyclical shifts in the composition of the pool of unemployed observed above, this heterogeneity in job finding rates may play an important role in understanding the differences in unemployment recoveries across recessions.




Figure 3.3: Job finding rates by reason of unemployment

3.2.3 Compositional effect

To compute the compositional effect for each recession I construct a time series of the aggregate job finding rate, holding job finding rates constant, and allow only the composition of the pool of unemployed to vary over time. Following the notation of Shimer (2012):

$$F_{comp,t} \equiv \sum_{i=1}^N \bar{F}_i \omega_{i,t} \quad (3.3)$$

$$F_{real,t} \equiv \sum_{i=1}^N F_{i,t} \omega_{i,t} \quad (3.4)$$

where F_i is the job finding rate for group i , N is the number of reasons for unemployment, \bar{F}_i is the average job finding rate for group i in their respective eras, and ω_i is the share of group i of the total number of unemployed. I compute the gap in F_{real} from the start of each recession and plot it in Figure 3.4. This is calculated as $F_{real,start} - F_{real,t}$. It is clear from this figure that the deterioration of the aggregate job finding rate in the 2008 recession is unlike those experienced in the prior four recessions.

For each recession I also compute the gap in F_{comp} since the start of each recession. This provides a measure of the change in the job finding rate due solely to the change in the composition of the pool of unemployed by reason for

unemployment. It is calculated as $F_{comp,start} - F_{comp,t}$ and plotted in Figure 3.5.




Figure 3.4: Change in job finding rate relative to the start of the recession




Figure 3.5: Change in job finding rates due to the composition effect

The change in the job finding rate due to the compositional changes is the largest for the 2008 recovery. This mechanism is also prominent for the 1990 recovery, and less so for the 1981 and 2001 recoveries.¹² Due to the important CPS redesign issues highlighted above I restrict my comparison to the 1981 vs. 1990

¹²Unfortunately, because the labor market had not fully recovered prior to the start of the 1981 recession, the 1980 recovery is of little use for the analysis conducted here.

recoveries and the 2001 vs. 2008 recoveries. I assess the importance this compositional mechanism by constructing the following time series for both comparisons:

$$y_t \equiv \frac{(F_{comp,start} - F_{comp,t})^A - (F_{comp,start} - F_{comp,t})^B}{(F_{real,start} - F_{real,t})^A - (F_{real,start} - F_{real,t})^B} \quad (3.5)$$

where A indicates the 1990 and 2008 time series and B indicates the 1981 and 2001 time series. In words, (3.5) computes the share of the excess decline in the recession A job finding rate relative to the decline in the recession B job finding rate that is explained by the excess compositional change in A relative to B. Since the decline in the job finding rates are similar in the onset of recessions, y_t suffers from small denominators early in the time series. I therefore compute y_t for each period past twenty months since the start of the recession.¹³ I find that relative to the 1981 recovery, the compositional channel accounts for approximately 20% of the excess decline in the job finding rate in the 1991 recovery, and that relative to the 2001 recovery the compositional channel accounts for approximately 24% of the excess decline in the job finding rate in the 2008 recovery. While much of the excess decline in the 1990 and 2008 recoveries remains unaccounted for by this decomposition it does highlight that accounting for compositional changes in the pool of unemployed by reason for unemployment is quantitatively important in understanding the pace of recent labor market recoveries.

¹³As can be seen in Figure 3.4, the 20-month target does a reasonable job at capturing the time to the peak decline in the job finding rates across recessions.

3.2.4 Elasticity of Job Finding Rate

The analysis in the preceding section took a naïve approach to estimating the importance of the compositional shifts in reason for unemployment. Namely, it assumes that the elasticity of the job finding rates with respect to the aggregate labor market tightness across reasons for unemployment is the same. It may be the case that the elasticity of job finding rates across reasons for unemployment is systematically different and therefore the importance of the compositional channel outlined above might be over or understated. Table 3.2 reports the estimated elasticity of job finding rates with respect to aggregate labor market tightness.¹⁴ I report elasticities for permanently displaced workers as one group and all other reasons for unemployment as another group from January 1994 through to December 2007, so as to leave out the Great Recession. Since JOLTS data on vacancies only begins in 2000 I use the monthly civilian unemployment rate as a proxy for labor market tightness.¹⁵

I find that the job finding rate for permanent layoffs is roughly twice as elastic to aggregate labor market tightness as is the job finding rate for the group of all other reasons of unemployment. The higher cyclical sensitivity of the job finding rate for permanent layoffs strengthens the compositional mechanism highlighted

¹⁴These estimates are obtained using a Prais-Winsten estimation to allow for first order autocorrelation in the residuals.

¹⁵This proxy is also used in Barnichon and Figura (2011). From 2000-2013 correlation between aggregate tightness and unemployment is -0.92

| Variable | Other reasons | | Permanent Layoff | |
|--------------------|---------------|----------|------------------|----------|
| | Coeff. | p-value | Coeff. | p-value |
| Constant | 0.139 | 0.151 | 0.133 | 0.4198 |
| $-\ln(Unemp.rate)$ | 0.413*** | 6.13e-11 | 0.805*** | 2.66e-19 |
| Durbin-Watson | 2.38 | | 2.06 | |

Table 3.2: Elasticity of job finding rates with respect to aggregate labor market tightness (unemployment rate) for permanent layoffs and all other reasons for unemployment

above. During recessions with large increases in permanent layoffs, not only is the economy loading on the group with the lower steady-state job finding rate but that is also the group that is more cyclically sensitive to the aggregate state of the labor market. To formalize this compositional mechanism and to gain a better sense of its quantitative importance, I construct a labor search model in Section 3.3 that captures the richness of the heterogeneity I find in the data.

3.2.5 Why compositional shifts have become stronger

Given the compositional effects documented above, a natural question to ask is why those effects have been stronger in recent recessions. In this section I investigate potential explanations behind these changes over time. Figure 3.2 shows that temporary layoffs were strongly countercyclical in the twin 1980s recessions, saw a modest increase in the 1990 recession and displayed very little cyclical behavior in the 2001 and 2008 recessions. This reduction in the use of temporary displacements has strengthened the compositional shift towards permanent dis-

placements.

To shed light on why the use of temporary displacements has declined, I explore the characteristics of workers on temporary layoffs in the CPS data. Importantly, I find that in 1976 the majority (73%) of temporary layoffs originate from “routine manual” (or “blue collar”) occupations, even though those occupations accounted for only 37% of all workers.¹⁶ One of the employment trends that has received considerable attention in recent years is “job polarization,” or the disappearance of middle skill jobs (see, for example, David et al. (2003) and Jaimovich and Siu (2012b)). In line with these papers I find that the share of employment in routine manual occupations declined from 37% in 1976 to 21% in 2013, a decline of 43%. Given that the majority of temporary layoffs originate from this class of occupations, this decline in employment share translates into a decline in the share of unemployed workers who are on temporary layoff.

In addition to a decline in the share of employment in routine manual occupations, another reason for a decline in the use of temporary layoffs may be a result of the decline in unionization. Temporary layoffs are heavily used in unionized occupations and there has been a sharp decline in unionization in the early 80s.¹⁷ In sum, examining the occupational/industrial trends in the U.S. over the past 40 years provides some insights into the dynamics of unemployment across recov-

¹⁶To categorize routine manual I use the same definitions as David et al. (2003).

¹⁷See Farber and Western (2000) for evidence on the decline in unionization.

eries and the general decline in the pace of unemployment recoveries in recent recessions.

3.3 Model

Over the past few decades macroeconomic models have placed more emphasis on the importance of frictions in labor markets (see, for example, Merz (1995), Andolfatto (1996), den Haan et al. (2000), and Walsh (2005)). In most cases, these DSGE models keep labor markets homogeneous for tractability. In attempting to understand business cycle phenomena this abstraction may be reasonable if heterogeneity is unimportant for aggregate dynamics or if it is unaligned with the business cycle. However, as documented above, the differences in employment outcomes across reasons for unemployment is both a source of substantial heterogeneity and is strongly aligned with the business cycle. In this section I formalize the compositional mechanisms highlighted above in an otherwise standard labor search and matching model and explore their quantitative implications.

The role of heterogeneity in labor search models in propagating unemployment fluctuations has been explored in Pries (2008) and Ravenna and Walsh (2012). The key mechanism in these models is that the composition of the unemployment pool shifts more toward low productivity workers during recessions, which lowers the expected value of a match to a firm and reduces their incentive to post va-

cancies. While this mechanism is found to amplify the unemployment response to shocks, the mechanism predicts a wage cyclical that is at odds with the data. Namely, if the unemployment pool shifts toward low skilled workers during recessions, the average previous wage in the unemployment pool should be procyclical. However, Mueller (2012) finds the opposite, that previous wages in the unemployment pool are countercyclical.

I propose a model that focuses on a different type of heterogeneity: differences in reason for unemployment. The model serves two purposes. The first is that it provides insights into how important these compositional shifts are for amplifying and propagating shocks relative to standard homogeneous worker labor search models. The second is that the model allows for a quantitative assessment of the importance of the two compositional channels highlighted in the empirics above. The first channel being simple compositional shifts in the pool of unemployed towards permanent displacements with lower steady-state job findings rates, and the second being the additional importance of the greater cyclical sensitivity of job finding rates for permanent displacements.

The key ingredients of the model are as follows. There are two groups of unemployed workers: permanently displaced and “other.” The “other” group represents workers who are unemployed for any reason besides permanent displacement. I maintain different groups of employed workers (“other” and permanent displace-

ment) and workers in both of these groups can transition into either unemployment pool.

The value to a firm of being matched with a displaced worker and a worker of type “other” is denoted $J_d(X)$ and $J_o(X)$, respectively, and are given by (3.6) and (3.7) below. The value to a firm of being matched with a worker is the net production value of the worker, $p + y_o - w_o$ or $p + y_d - w_l$ for the two types of workers, plus the continuation value the match. Where, p is the aggregate labor productivity in the economy, y_i is the labor productivity of the individual worker, and w_i is the wage for a worker of type i .¹⁸

$$J_d(X) = p + y - w_d(X) + \beta \mathbb{E}_{X'|X} [(1 - s_d(p) - s_o)J_d(X') + (s_d(p) + s_o)V(\mathfrak{U})] \quad (3.6)$$

$$J_o(X) = p + y - w_o(X) + \beta \mathbb{E}_{X'|X} [(1 - s_d(p) - s_o)J_o(X') + (s_d(p) + s_o)V(\mathfrak{U})] \quad (3.7)$$

There is regular churning in the labor market; workers flow into the “other” unemployed pool in an acyclical fashion. The value of being unemployed is given by (3.17) and (3.18) below, where z is the value of leisure/unemployment benefits, the probability of a entering into an employment match is given by $f_i(\theta(X)) = \theta^{1-\eta_i}$, and η_i is the elasticity of the job finding rate with respect to aggregate labor

¹⁸I do not allow productivity to vary across worker types. The literature on “scarring” suggests that permanently displaced workers may be lower productivity, or simply perceived that way, and thus have lower earnings trajectories after re-engagement (see Arulampalam et al. (2001)). Allowing for different productivities between displaced and “other” types would strengthen the results below as the average productivity of the pool of unemployed would deteriorate with a burst of permanent displacements, reducing firms’ incentive to post vacancies.

market tightness.¹⁹ I create two variants of the model in terms of job finding elasticities. In the first I allow permanent displacements to differ only in their steady-state job finding rates and parameterize both groups to have the same job finding elasticities, thus $\eta_d = \eta_o$. In the second version, I allow the two groups to differ both in their steady-state job finding rates but also in their job finding rate elasticity, η_i , to the aggregate state of the labor market.²⁰

$$U_d(X) = z + \beta \mathbb{E}_{X'|X}[f_d(\theta(X))E_d(X') + (1 - f_d(\theta(X)))U_d(X')] \quad (3.8)$$

$$U_o(X) = z + \beta \mathbb{E}_{X'|X}[f_o(\theta(X))E_o(X') + (1 - f_o(\theta(X)))U_o(X')] \quad (3.9)$$

Employed workers earn wages, $w_i(X)$, and separate from matches into either the “other” or permanently displaced unemployment pools with rates s_o and $s_d(p)$, respectively. I assume s_o to be constant and, as such, represent the regular churning in the labor force, and I assume $s_d(p)$ to be negatively correlated with aggregate labor productivity to capture the cyclical response of separations. This assumption will also generate compositional shifts in the pool of unemployment,

¹⁹One important difference between permanent displacements and “other” reasons for unemployment may be their access to unemployment benefits. I abstract from these differences here, but they are more likely to generate interesting labor market dynamics in a model that includes search intensity, as in Ljungqvist and Sargent (1998).

²⁰The potential reasons behind the differences in steady-state job finding rates and their sensitivities to the business cycle are many. I do not attempt to micro-found these reasons in the model for the sake of tractability. Instead I focus on highlighting the importance of these differences across groups of unemployed workers for the dynamics of labor market search models. In the discussion section of the paper I highlight some potential reasons for why these differences exist, which have important implications for labor market policies.

which will allow me to explore the quantitative importance of the compositional mechanisms emphasized in the empirical section above. $E_d(X)$ and $E_o(X)$ are the values of being employed for permanent displacements and “other” respectively:

$$E_d(X) = w_d(X) + \beta \mathbb{E}_{X'|X} [(1 - s_d(p) - s_o)E_d(X') + s_d(p)U_d(X') + s_o U_o(X')] \quad (3.10)$$

$$E_o(X) = w_o(X) + \beta \mathbb{E}_{X'|X} [(1 - s_d(p) - s_o)E_o(X') + s_d(p)U_d(X') + s_o U_o(X')] \quad (3.11)$$

The value to a firm for posting a vacancy is given by:

$$V(X) = -k + \beta \mathbb{E}_{X'|X} [q_d(\theta(X))\mu J_d(X') + q_o(\theta(X))(1 - \mu)J_o(X')] \quad (3.12)$$

where, $q_i(\theta(X)) = \theta^{-\eta}$ is the rate at which vacancies are matched to unemployed workers. In all the above value functions, $X = \{u_d, u_o, p\}$ is the state variable that describes the state of the labor market. These include the composition of the pool of unemployed ($\mu = u_d/(u_d + u_o)$) and aggregate labor productivity, p .

Finally, the surplus generated from a match is given by:

$$S_d(X) = J_d(X) + E_d(X) - U_d(X) - V(X) \quad (3.13)$$

$$S_o(X) = J_o(X) + E_o(X) - U_o(X) - V(X) \quad (3.14)$$

Workers and firms split this surplus via Nash Bargaining where workers receive $(1 - \gamma_i)(E_i(X) - U_i(X))$ and firms receive $\gamma_i J_i(X)$.

3.3.1 Model Solution

I solve the model computationally following the general steps used in Pries (2008). I discretize the state space (θ, u_d, u_o) into a 50x50x20 grid. Since workers are able to transition between groups in my model the solution becomes more slightly more complicated. As in Pries (2008) I attempt to solve for simple recursions in the surplus equations:

$$\begin{aligned}
S_d(X) &= J_d(X) + E_d(X) - U_d(X) - V(X) \\
&= p + y - z + \beta \mathbb{E}_{X'|X} [[1 - s_d(p) - so - f_d(\theta) p_d(\theta) \gamma_d] S_d(X') - s_o(U_d(X') - U_o(X'))]
\end{aligned} \tag{3.15}$$

$$\begin{aligned}
S_o(X) &= J_o(X) + E_o(X) - U_o(X) - V(X) \\
&= p + y - z + \beta \mathbb{E}_{X'|X} [[1 - s_d(p) - so - f_o(\theta) \gamma_o] S_o(X') - s_d(p)(U_o(X') - U_d(X'))]
\end{aligned} \tag{3.16}$$

Where I substitute in $E_i(X) = \gamma_i S_i(X) + U_i(X)$ from the Nash Bargaining conditions.

Because workers can transition between groups, my recursive surplus equations

also include $U_o(X')$ and $U_d(X')$. With these added complications I solve the model as follows:

Start with a guess of $U_o(X)$, $U_d(X)$, $S_o(X)$, and $S_d(X)$. With these guesses I solve jointly for the true values $U_o(X)$, $U_d(X)$, $S_o(X)$, and $S_d(X)$ by iterating on the system of equations that includes (3.15), (3.16) and:

$$U_d(X) = z + \beta \mathbb{E}_{X'|X} [f_d(\theta(X)) p_d \gamma_d S_d(X') + U_d(X')] \quad (3.17)$$

$$U_o(X) = z + \beta \mathbb{E}_{X'|X} [f_o(\theta(X)) \gamma_o S_o(X') + U_o(X')] \quad (3.18)$$

To solve this I need to track the evolution of the aggregate state X . Productivity will evolve via a standard AR(1) process.²¹ The following equations govern the evolution of unemployment:

$$u'_d = u_d(1 - f_d(\theta(X) p_d) + (1 - u_d - u_o) s_d \quad (3.19)$$

$$u'_o = u_o(1 - f_o(\theta(X))) + (1 - u_d - u_o) s_o \quad (3.20)$$

Because new unemployment and productivity levels will generally lie between grid points, I use linear interpolations of the value functions between the two nearest grid points to approximate their evolution.

²¹I discretize productivity into a grid of twenty states using the Rouwenhorst method. See Kopecky and Suen (2010) for details.

With the numerical solutions of the value functions I next to check if the free entry condition holds:

$$k = \beta \mathbb{E}_{X'|X} [q_d(\theta(X)) p_d \mu (1 - \gamma_d) S_d(X') + q_o(\theta(X)) (1 - \mu) (1 - \gamma_o) S_o(X')] \quad (3.21)$$

Given the numerical solutions of $S_d(X)$ and $S_o(X)$ I solve for a new matrix $\theta(X)$ from the non-linear equation (3.21) using the Newton-Rapson method. If the new values of $\theta(X)$ implied by free entry are sufficiently different than last periods iterate of $\theta(X)$ then the matrix is updated using a convex combination of last periods iterate and the new value. The above steps are repeated until the free entry condition is satisfied. With the solution in hand, the model will allow for various simulation exercises to investigate the quantitative importance of the mechanisms discussed above.

3.3.2 Steady-State

In the steady-state, $u'_d = u_d = u_d^{ss}$ and $u'_o = u_o = u_o^{ss}$ which yields:

$$u_d^{ss} = \frac{1}{1 - \frac{s_d^{ss} s_o}{(f_d(\theta^{ss}) p_d + s_d^{ss})(f_o(\theta^{ss}) + s_o)}} \left[\frac{s_d^{ss}}{f_d(\theta^{ss}) + s_d^{ss}} - \frac{s_d^{ss} s_o}{(f_d(\theta^{ss}) + s_d^{ss})(f_o(\theta^{ss}) + s_o)} \right]$$

$$u_o^{ss} = \frac{s_o}{s_o + f_o(\theta^{ss})} \left[1 - \frac{1}{1 - \frac{s_d^{ss} s_o}{(f_d(\theta^{ss}) p_d + s_d^{ss})(f_o(\theta^{ss}) + s_o)}} \left[\frac{s_d^{ss}}{f_d(\theta^{ss}) + s_d^{ss}} - \frac{s_d^{ss} s_o}{(f_d(\theta^{ss}) + s_d^{ss})(f_o(\theta^{ss}) + s_o)} \right] \right]$$

To solve for the steady-state labor market tightness (3.17) and (3.18) can be substituted into (3.15) and (3.16) where $U_d(X) = \mathbb{E}_{X'|X}U_d(X')$, $U_o(X) = \mathbb{E}_{X'|X}U_o(X')$, $S_d(X) = \mathbb{E}_{X'|X}S_d(X')$, and $S_o(X) = \mathbb{E}_{X'|X}S_o(X')$. These steady-state equations can be substituted into (3.21) to solve for θ^{ss} . See the Appendix for derivations. Figure 3.6 plots the value of a vacancy as a function of θ^{ss} for the model parameterization used below. It shows that the model emits a unique steady-state value of labor market tightness and that multiple equilibria do not exist.




Figure 3.6: Plots value of a vacancy as a function of θ^{ss} .

3.4 Calibration

This section provides a calibration of the model chosen to closely match important targets in the data, and where possible to match standard values in the

labor search literature. I explore two calibrations of the model. In the first, I allow permanently displaced and “other” unemployed workers to differ only in their steady-state job finding rates yet have the same job finding elasticity set to match the aggregate job finding elasticity in the data. In the second calibration, I relax this assumption and allow the two groups to differ in their job finding elasticities set to match the results from Table 3.2 above. Specifically, in the first calibration I set $\eta_d = \eta_o = \eta$ to 0.49 and in the second calibration I set $\eta_d = 0.2$ and $\eta_o = 0.59$. In all cases I choose γ_i so that the Hosios (1990) condition holds.

| Parameter | Value | Target |
|------------|--------|------------------------------------|
| β | 0.996 | Annual Interest rate of 4% |
| η | 0.49 | Estimated from CPS data |
| γ | 0.51 | Hosios Condition |
| k | 2.25 | Target from steady-state JF rates |
| p | | Discretized matrix |
| y | 1 | Normalization |
| s_d^{ss} | 0.0075 | Target from unemployment shares |
| s_o | 0.0275 | Target from unemployment shares |
| p_d | 0.7 | Steady-state difference in JF rate |
| z | 0.4 | See text |

Table 3.3: Calibration values for model with common elasticity.

I follow Pries (2008) and set $z = 0.4$. The labor productivity process, p , is a Markov process set to match the moments found in U.S. data. Worker productivity, y , is normalized to 1. Finally, an aggregate job separation rate of 0.035 found in CPS data is divided so that $s_o = 0.0275$ and $s_d^{ss} = 0.0275$, which matches

average unemployment composition, μ , of 0.3. To generate compositional shifts towards permanent displacements in the pool of unemployed I allow s_d to be negatively correlated with the aggregate state of the economy, p . I report the model's dynamics for two different covariances between s_d and p along with the baseline model where s_d is acyclical.

3.5 Results

I begin by assessing how the two versions of the model impact the speed of unemployment recoveries relative to a standard homogeneous worker labor search model. To do this I shock the model with a one period negative productivity shock and track the evolution of unemployment back to the steady-state. In this exercise I restrict productivity to its steady state level following the one period shock. Thus the path of unemployment back to the steady-state is purely a result of the propagation mechanisms.

To quantitatively assess the impact of these different calibrations on the evolution of unemployment it is worth comparing the half-lives of unemployment across calibrations to that of a standard model. The half-life of out of steady-state unemployment in a standard, continuous time, homogenous worker DMP model can be written as:

$$t_{hl} = \frac{-\log 0.5}{s + f}$$

With an aggregate job finding rate and separation rates of approximately 0.37 and 0.035 used in the calibration, this implies a half-life of 1.71 months for a standard homogeneous worker DMP model.²² Table 3.4 reports the unemployment half-lives for the model with common and separate job finding elasticities across the reasons for unemployment for the two parameterizations of separation cyclicity.

With common job finding elasticities the model generates a half-life of 2.37 and 2.45 months for the low and high of separation cyclicalities. Therefore, simple compositional changes towards permanent displacements who have lower steady-state job finding rates can increase the persistence of unemployment by roughly 30%. The reason for this increase very straightforward. Since labor market tightness fluctuates little (as will be shown below) most of the increase in persistence is from the fact that the increase in unemployment is coming from flows into permanent displacements. With a steady-state job finding rate of 0.27, $t_{hl} = \frac{-\log 0.5}{s+f}$ would predict a half-life of 2.27 months. Standard search models miss this because they are calibrated to aggregate job finding rates and do not distinguish between groups of unemployed workers who have different re-employment probabilities.

When the model is allowed to differ in both steady-state job finding rate differences and differences in the cyclicity of job finding rates across reasons for

²²The baseline DMP model is one where I simply set the two groups to have the same steady-state job finding rates and same job finding elasticities. These values are calibrated to match the aggregate data.

unemployment, the persistence of unemployment fluctuations increases dramatically. In this version of the model, half-lives of unemployment shocks are 4.61 and 5.14 for the low and high separation cyclicalities, respectively. This represents between a 170% to 200% increase from a baseline DMP model. This increase occurs because unemployment shocks are not only loading on the group with the lower steady-state job finding rate, but this is also the group who's job finding rate is more sensitive to fluctuations in the labor market. However, this mechanism only kicks in if there are sizable fluctuations in labor market tightness. As will be shown below, this version of the model generates more volatility in labor market tightness and therefore causes the differences in job finding rates between the two reasons for unemployment to become exaggerated during recessions.

The other interesting point worth noting is that larger compositional shifts towards permanent displacements increases the half-life of unemployment fluctuations, especially in the model with separate job finding elasticities. Unfortunately, there is no clear way to estimate an empirical counterpart to the model exercise in Table 3.4, but this result is consistent with the general empirical findings above: recessions with stronger shifts towards permanent displacements have slower unemployment recoveries.

| | Baseline | Common | Common | Separate | Separate |
|--------------------|----------|------------|------------|------------|------------|
| | DMP | Elasticity | Elasticity | Elasticity | Elasticity |
| | | Low Sep. | High Sep. | Low Sep. | High Sep. |
| Half-Life (months) | 1.71 | 2.37 | 2.45 | 4.61 | 5.14 |

Table 3.4: Half-life of aggregate unemployment in months following a one period productivity shock.

In addition to helping to understand the persistence of unemployment shocks, this model also has implications regarding the volatility of labor market variables.²³ Table 3.5 provides a comparison of the second moments of key labor market variables between U.S. data, the baseline DMP model, and the version of the model with common job finding elasticities. This exercise compares second moments of productivity, labor market tightness, job finding rates, unemployment, separation rates, and the composition of the pool of unemployed by reason for unemployment. I highlight a few important results. The first is the common result that the volatility in labor market tightness generated in the baseline DMP is far below that observed empirically and that the model with common job finding elasticities only generates modest increases. This is the reason that

²³The persistence of shocks and the volatility of labor market variables are, of course, are closely connected. A know well known literature has documented the inability of standard labor search models to replicate the degree of labor market volatility observed empirically. See, for example, Shimer (2005) and Costain and Reiter (2008).

the unemployment half-life in this version of the model is very close to the one simply calculated for permanent displacements using $t_{hl} = \frac{-\log 0.5}{s+f}$. There is only a modest increase in the volatility in labor market tightness in this version of the model because vacancies change little. Through equation (3.12), one can see that equilibrium tightness will not vary much through these simple compositional shifts (changes in μ). With the same job finding elasticity, η , changes in μ will only impact aggregate labor market tightness through the difference between J_o and J_d , which is small.

The other important model moments worth noting are those for unemployment, separation rates, and the composition of the unemployment pool. By construction, these moments increase with the degree of cyclicity of separation rates into permanent displacements but are worth comparing to their empirical counterparts to give a sense of how well the model is capturing the compositional shifts in the pool of unemployment over the business cycle.

| | U.S. data | Baseline (acyclical separations) | Cyclical Separations (Low) | Cyclical Separations (High) |
|----------------|-----------|-------------------------------------|-------------------------------|--------------------------------|
| Prod. | 0.02 | 0.02 | 0.02 | 0.02 |
| $\theta = v/u$ | 0.382 | 0.032 | 0.036 | 0.038 |
| $f(\theta)$ | 0.126 | 0.0166 | 0.0183 | 0.0193 |
| u | 0.190 | 0.0139 | 0.0373 | 0.0626 |
| s | 0.079 | 0 | 0.0210 | 0.0435 |
| μ | 0.1505 | 0.00149 | 0.0614 | 0.1248 |

Table 3.5: Empirical and model second moments. Model moments are constructed from the calibrated version of the model where worker types share a common job finding elasticity. All standard deviations are constructed after taking natural logs and removing the trend using an HP smoothing parameter of 10^5 .

The second version of the model, with separate job finding elasticities, generates more volatility in these labor market moments, reported in Table 3.6. The volatility of labor market tightness is 65% and 147% larger for the two separation cyclical parameterizations, respectively. Again, the amplification depends on equation (3.12). In the version of the model with different job finding elasticities, compositional changes impact labor market tightness through the difference be-

tween job finding elasticities, η_i . To see this, note that equation (3.12) depends on the rate that firms are matched with workers, $q_i(\theta(X)) = \theta^{-\eta_i}$. Calibrated to match the empirical findings above, η_d is much smaller than η_o , 0.2 versus 0.59, respectively. As the composition of the pool of unemployed shifts towards permanent displacements during recessions and μ increases, equation (3.12) loads more heavily on $\theta^{-\eta_d}$, which is much lower than $\theta^{-\eta_o}$. This reduces the value of posting a vacancy and lowers labor market tightness. This is also the effect behind the increase in the half-life of unemployment fluctuations highlighted above. Besides being relevant to the motivation of this paper, this is a point that is relevant to labor search models with heterogeneous workers: the different job finding elasticities across groups can have very important consequences for the dynamics of the model.

| | U.S. data | Cyclical Separations (Low) | Cyclical Separations (High) |
|----------------|-----------|-------------------------------|--------------------------------|
| Prod. | 0.02 | 0.02 | 0.02 |
| $\theta = v/u$ | 0.382 | 0.053 | 0.079 |
| $f_o(\theta)$ | 0.0799 | 0.0212 | 0.0322 |
| $f_d(\theta)$ | 0.1666 | 0.042 | 0.0628 |
| u | 0.190 | 0.049 | 0.084 |
| s | 0.0786 | 0.0210 | 0.0435 |
| μ | 0.1505 | 0.06879 | 0.13764 |

Table 3.6: Empirical and model second moments. Model moments are constructed from the calibrated version of the model where worker types have separate job finding elasticities estimated from the data. All standard deviations are constructed after taking natural logs and removing the trend using an HP smoothing parameter of 10^5 .

3.6 Discussion

Numerous micro and macro studies have documented the adverse impact of permanent displacements on subsequent earnings and employment outcomes. Us-

ing CPS microdata I confirm this by documenting that those reported as permanently displaced have substantially lower job finding rates as compared to other reported reasons for unemployment. In addition, I also document that recent recessions have been characterized by stronger compositional shifts towards permanent displacements relative to earlier post-war recessions. I investigate the extent to which these compositional changes in the unemployment pool by reason for unemployment contribute to the slow labor market recoveries observed in the past three recessions. This channel is able to account for a quantitatively important (20 to 24%) of the excess decline in the job finding rate observed in recent recessions. I also document that this channel is understated by the greater cyclical sensitivity in job finding rates for permanently displaced workers.

Given that this heterogeneity between reasons for unemployment is quantitatively important, and that compositional shifts across these groups vary strongly with the business cycle, I explore the quantitative importance of capturing this heterogeneity in a search and matching model. I find that it is able to substantially decrease the speed of aggregate unemployment recoveries and that it is able to amplify the response of key labor market variables in comparison to a model with homogenous workers. This exercise provides evidence that homogeneous worker labor search models widely used in macro miss much of the important business cycle dynamics of unemployment.

These findings highlight the importance of compositional shifts towards permanent displacements for the pace of unemployment recoveries. To mitigate this aspect of slow labor market recoveries, policy could be designed to either reduce the discrepancy in job finding rates between reasons for unemployment and/or reduce the size of the compositional shifts towards permanent displacements during recessions.

The causes behind differences in job finding rates across reasons for unemployment arise from numerous reasons. One key difference between permanent displacements and other reasons for unemployment is the access to unemployment insurance. Ljungqvist and Sargent (1998) argue that unemployment insurance can have perverse effects in turbulent labor markets. Specifically, it might reduce a workers incentive to undertake costly search for jobs and therefore limit the labor supplied to the market.²⁴ In addition to incentives related to unemployment insurance, a workers' access to credit may be an important determinant of their search intensity, see Herkenhoff (2013).

Permanently displaced workers may also face a difficult reallocation process following a termination. A large literature examines the role of sectoral reallocation (or mismatch) in labor markets. If a permanent displacement results from a plant or firm closure, workers may face a costly retraining or geographic relocation

²⁴The perverse incentives of unemployment insurance may be less important during recessions in a model where labor demand by firms is an important driver of unemployment, see Landais et al. (2010).

before regaining employment, leading to low observed job finding rates. Papers in this literature include Shimer (2007), Barlevy (2011), Herz and van Rens (2011), Sahin et al. (2012) among others.

Due to the various reasons driving the differences in job finding rates across reasons for unemployment, it may be easier to design policy that limits the compositional shifts toward permanent displacements during recessions. In my empirical results, I highlight that these compositional shifts have become larger in recent recessions and that these recessions are characterized by slow labor market recoveries. In related work, Berger (2012) asserts a connection between firm restructuring behavior and slow labor market recoveries. In consideration of the importance of these compositional shifts towards permanent displacements, policy could be designed to instead target firms' employment responses to recessions. For example, policies designed to disincentivize firms from laying off workers during recessions appeared to mitigate the use of layoffs in Germany during the Great Recession which may be partially responsible for the mild impact on their labor market, see Burda and Hunt (2011). These firm disincentives may increase the speed of unemployment recoveries in that they prevent a large compositional shift in the pool of unemployed, as was experienced by the U.S. in the Great Recession. Having a clearer understanding of how these incentives impact firms' employment decisions and aggregate labor market dynamics remains an important avenue for

further research.

Appendix D

Derivation of Steady-State

The steady-state equations are given by:

$$J_d^{ss} = \frac{p + y - w_d(X)}{1 - \beta(1 - s_d^{ss} - s_o)} \quad (\text{D.1})$$

$$J_o^{ss} = \frac{p + y - w_o(X)}{1 - \beta(1 - s_d^{ss} - s_o)} \quad (\text{D.2})$$

$$E_d^{ss} = \frac{w_d(X) + \beta s_d^{ss} U_d^{ss} + \beta s_o U_o^{ss}}{1 - \beta(1 - s_d^{ss} - s_o)} \quad (\text{D.3})$$

$$E_o^{ss} = \frac{w_o(X) + \beta s_d^{ss} U_d^{ss} + \beta s_o U_o^{ss}}{1 - \beta(1 - s_d^{ss} - s_o)} \quad (\text{D.4})$$

$$U_d^{ss} = \frac{z_d + \beta f_d(\theta^{ss}) p_d(\theta^{ss}) \gamma_d S_d^{ss}}{1 - \beta} \quad (\text{D.5})$$

$$U_o^{ss} = \frac{z_o + \beta f_o(\theta^{ss}) \gamma_o S_o^{ss}}{1 - \beta} \quad (\text{D.6})$$

Combining the above equations into:

$$S_i^{ss} = J_i^{ss} + E_i^{ss} - U_i^{ss}$$

yields,

$$S_o^{ss} = \frac{1}{A} \left[B + \frac{\beta^2 s_d f_d(\theta^{ss}) p_d \gamma_d}{(1-\beta)(1-\beta(1-s_d^{ss}-s_o))} S_d^{ss} \right] \quad (D.7)$$

$$S_d^{ss} = \frac{1}{C} \left[D + \frac{\beta^2 s_o f_o(\theta^{ss}) \gamma_o}{(1-\beta)(1-\beta(1-s_d^{ss}-s_o))} S_o^{ss} \right] \quad (D.8)$$

where,

$$\begin{aligned} A &= 1 + \frac{\beta f_o(\theta^{ss}) \gamma_o}{1-\beta} - \frac{\beta^2 s_o f_o(\theta^{ss}) \gamma_o}{(1-\beta)(1-\beta(1-s_d^{ss}-s_o))} \\ B &= \frac{p+y}{1-\beta(1-s_d^{ss}-s_o)} - \frac{z_o}{1-\beta} + \frac{\beta s_d^{ss} z_d + \beta s_o z_o}{(1-\beta)(1-\beta(1-s_d^{ss}-s_o))} \\ C &= 1 + \frac{\beta f_d(\theta^{ss}) p_d \gamma_d}{1-\beta} - \frac{\beta^2 s_d^{ss} f_d(\theta^{ss}) p_d \gamma_d}{(1-\beta)(1-\beta(1-s_d^{ss}-s_o))} \\ D &= \frac{p+y}{1-\beta(1-s_d^{ss}-s_o)} - \frac{z_d}{1-\beta} + \frac{\beta s_d^{ss} z_d + \beta s_o z_o}{(1-\beta)(1-\beta(1-s_d^{ss}-s_o))} \end{aligned}$$

Combining (D.8) and (D.7) yields a solution to the system in terms of θ^{ss} and other parameters which can be used to solve the steady-states of the other value functions and labor market tightness.

Bibliography

- Aaronson, D., Rissman, E., Sullivan, D. G., 2004. Can sectoral reallocation explain the jobless recovery? *Economic Perspectives* (Q II), 36–39.
- Abraham, K. G., Shimer, R., 2001. Changes in unemployment duration and labor force attachment. Tech. rep., National Bureau of Economic Research.
- Acemoglu, D., December 1999. Changes in unemployment and wage inequality: An alternative theory and some evidence. *American Economic Review* 89 (5), 1259–1278.
- Aguiar, M., Hurst, E., Karabarbounis, L., 2013. Time use during the great recession. *The American Economic Review* 103 (5), 1664–1696.
- Andolfatto, D., March 1996. Business cycles and labor-market search. *American Economic Review* 86 (1), 112–32.
- Arellano, C., Bai, Y., Kehoe, P., 2010. Financial markets and fluctuations in uncertainty. Federal Reserve Bank of Minneapolis Working Paper.

- Arulampalam, W., Gregg, P., Gregory, M., 2001. Unemployment scarring. *The Economic Journal* 111 (475), 577–584.
- Autor, D. H., Katz, L. F., Kearney, M. S., May 2006. The polarization of the u.s. labor market. *American Economic Review* 96 (2), 189–194.
- Bachmann, R., Elstner, S., Sims, E. R., 2010. Uncertainty and economic activity: Evidence from business survey data. Tech. rep., National Bureau of Economic Research.
- Bachmann, R., Moscarini, G., et al., 2011. Business cycles and endogenous uncertainty. In: 2011 Meeting Papers. Vol. 36. Society for Economic Dynamics.
- Baker, M., 1992. Unemployment duration: compositional effects and cyclical variability. *American Economic Review* 82 (1), 313–21.
- Barlevy, G., 2011. Evaluating the role of labor market mismatch in rising unemployment. *Economic Perspectives*.
- Barnichon, R., Figura, A., 2011. What Drives Matching Efficiency?: A Tale of Composition and Dispersion. Division of Research & Statistics and Monetary Affairs, Federal Reserve Board.
- Berger, D., 2012. Countercyclical restructuring and jobless recoveries. manuscript, Yale.

- Berger, D., Vavra, J. S., 2013. Volatility and pass-through. Tech. rep., National Bureau of Economic Research.
- Blanchard, O. J., Diamond, P., 1989. The beveridge curve. *Brookings Papers on Economic Activity* 20 (1), 1–76.
- Bloom, N., 2009. The impact of uncertainty shocks. *Econometrica* 77 (3), 623–685.
- Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., Terry, S. J., 2012. Really uncertain business cycles. Tech. rep., National Bureau of Economic Research.
- Blundell, R., Bond, S., 2000. Gmm estimation with persistent panel data: an application to production functions. *Econometric reviews* 19 (3), 321–340.
- Bronnenberg, B. J., Kruger, M. W., Mela, C. F., 2008. Database paper: The iri marketing data set. *Marketing Science* 27 (4), 745–748.
- Burda, M. C., Hunt, J., 2011. What explains the german labor market miracle in the great recession? Tech. rep., National bureau of economic research.
- Christiano, L. J., Motto, R., Rostagno, M., 2014. Risk shocks. *The American Economic Review* 104 (1), 27–65.
- Coibion, O., Gorodnichenko, Y., Hong, G. H., 2012. The cyclicity of sales, reg-

- ular and effective prices: Business cycle and policy implications. Tech. rep., National Bureau of Economic Research.
- Coibion, O., Gorodnichenko, Y., Koustas, D., October 2013. Amerisclerosis? the puzzle of rising u.s. unemployment persistence. Working Paper 19600, National Bureau of Economic Research.
- Costain, J., Nakov, A., 2015. Precautionary price stickiness. *Journal of Economic Dynamics and Control* 58, 218–234.
- Costain, J. S., Nakov, A., 2013. Logit price dynamics.
- Costain, J. S., Reiter, M., 2008. Business cycles, unemployment insurance, and the calibration of matching models. *Journal of Economic Dynamics and Control* 32 (4), 1120–1155.
- Darby, M. R., Haltiwanger, J. C., Plant, M. W., 1986. The ins and outs of unemployment: The ins win.
- David, H., FRANK, L., RICHARD, J. M., 2003. The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics* 118, 4.
- Davis, S. J., Faberman, R. J., Haltiwanger, J., 2012. Labor market flows in the cross section and over time. *Journal of Monetary Economics* 59 (1), 1–18.

- Davis, S. J., Haltiwanger, J., 1990. Gross job creation and destruction: Microeconomic evidence and macroeconomic implications. In: NBER Macroeconomics Annual 1990, Volume 5. MIT Press, pp. 123–186.
- Davis, S. J., Haltiwanger, J., 2014. Labor market fluidity and economic performance. Tech. rep., National Bureau of Economic Research.
- Davis, S. J., Haltiwanger, J., Jarmin, R., Miranda, J., 2007. Volatility and dispersion in business growth rates: Publicly traded versus privately held firms. In: NBER Macroeconomics Annual 2006, Volume 21. MIT Press, pp. 107–180.
- Decker, R., D’Erasmus, P., Moscoso Boedo, H. J., 2014. Market exposure and endogenous firm volatility over the business cycle.
- Decker, R., Haltiwanger, J., Jarmin, R., Miranda, J., 2013. The secular decline in business dynamism in the us. Manuscript, University of Maryland.
- den Haan, W. J., Ramey, G., Watson, J., June 2000. Job destruction and propagation of shocks. *American Economic Review* 90 (3), 482–498.
- Diamond, P. A., October 1982. Aggregate demand management in search equilibrium. *Journal of Political Economy* 90 (5), 881–94.
- Dixit, A. K., Stiglitz, J. E., March 1977. Monopolistic competition and optimum product diversity. *American Economic Review* 67 (3), 297–308.

- Eichenbaum, M., Jaimovich, N., Rebelo, S., 2011. Reference prices, costs, and nominal rigidities. *The American Economic Review*, 234–262.
- Elsby, M. W., Hobijn, B., Sahin, A., Valletta, R. G., 2011. The labor market in the great recession: an update. Working Paper Series 2011-29, Federal Reserve Bank of San Francisco.
- Elsby, M. W. L., Hobijn, B., Sahin, A., 2010. The labor market in the great recession. *Brookings Papers on Economic Activity* 41 (1 (Spring)), 1–69.
- Elsby, M. W. L., Michaels, R., Solon, G., January 2009. The ins and outs of cyclical unemployment. *American Economic Journal: Macroeconomics* 1 (1), 84–110.
- Farber, H. S., Western, B., 2000. Round up the usual suspects: the decline of unions in the private sector, 1973-1998.
- Fujita, S., Ramey, G., May 2009. The cyclicalities of separation and job finding rates. *International Economic Review* 50 (2), 415–430.
- Gilchrist, S., Sim, J. W., Zakrajšek, E., 2014. Uncertainty, financial frictions, and investment dynamics. Tech. rep., National Bureau of Economic Research.
- Gorry, A., 2012. Experience and worker flows. Working paper, University of California, Santa Cruz.

- Gorry, A., November 2013. Minimum wages and youth unemployment. *European Economic Review* 64, 57–75.
- Groshen, E. L., Potter, S., 2003. Has structural change contributed to a jobless recovery? *Current Issues in Economics and Finance* (Aug).
- Hall, R. E., Milgrom, P. R., March 2005. Employment fluctuations with equilibrium wage stickiness. *American Economic Review* 95 (1), 50–65.
- Herkenhoff, K. F., 2013. The impact of consumer credit access on unemployment.
- Herz, B., van Rens, T., Jul. 2011. Structural unemployment. *Economics Working Papers* 1276, Department of Economics and Business, Universitat Pompeu Fabra.
- Hosios, A. J., April 1990. On the efficiency of matching and related models of search and unemployment. *Review of Economic Studies* 57 (2), 279–98.
- Howitt, P., March 1985. Transaction costs in the theory of unemployment. *American Economic Review* 75 (1), 88–100.
- Hyatt, H. R., Spletzer, J. R., 2013. The recent decline in employment dynamics. *IZA Journal of Labor Economics* 2 (1), 1–21.
- Ilut, C., Kehrig, M., Schneider, M., 2014. Slow to hire, quick to fire: Employment

- dynamics with asymmetric responses to news. Tech. rep., National Bureau of Economic Research.
- Jacobson, L. S., LaLonde, R. J., Sullivan, D. G., September 1993. Earnings losses of displaced workers. *American Economic Review* 83 (4), 685–709.
- Jaimovich, N., Rebelo, S., Wong, A., 2015. Trading down and the business cycle. Tech. rep., National Bureau of Economic Research.
- Jaimovich, N., Siu, H. E., June 2009. The young, the old, and the restless: Demographics and business cycle volatility. *American Economic Review* 99 (3), 804–26.
- Jaimovich, N., Siu, H. E., Aug. 2012a. The trend is the cycle: Job polarization and jobless recoveries. NBER Working Papers 18334, National Bureau of Economic Research, Inc.
- URL <http://ideas.repec.org/p/nbr/nberwo/18334.html>
- Jaimovich, N., Siu, H. E., 2012b. The trend is the cycle: Job polarization and jobless recoveries. Tech. rep., National Bureau of Economic Research.
- Kaplan, G., Menzio, G., 2014. The morphology of price dispersion. Tech. rep., National Bureau of Economic Research.

- Kletzer, L. G., 1998. Job displacement. *The Journal of Economic Perspectives* 12 (1), 115–136.
- Kopecky, K. A., Suen, R. M., 2010. Finite state markov-chain approximations to highly persistent processes. *Review of Economic Dynamics* 13 (3), 701–714.
- Landais, C., Michaillat, P., Saez, E., 2010. Optimal unemployment insurance over the business cycle. Tech. rep., National Bureau of Economic Research.
- Ljungqvist, L., Sargent, T. J., June 1998. The european unemployment dilemma. *Journal of Political Economy* 106 (3), 514–550.
- Mattsson, L.-G., Weibull, J. W., 2002. Probabilistic choice and procedurally bounded rationality. *Games and Economic Behavior* 41 (1), 61–78.
- McKenzie, D., Schargrodsky, E., 2011. Buying less but shopping more: The use of non market labor during a crisis. *Economia*.
- Merz, M., November 1995. Search in the labor market and the real business cycle. *Journal of Monetary Economics* 36 (2), 269–300.
- Mortensen, D. T., 1989. The persistence and indeterminacy of unemployment in search equilibrium. *Scandinavian Journal of Economics* 91 (2), 347–70.
- Mortensen, D. T., Pissarides, C. A., July 1994. Job creation and job destruction in the theory of unemployment. *Review of Economic Studies* 61 (3), 397–415.

- Mueller, A. I., September 2012. Separations, sorting and cyclical unemployment. IZA Discussion Papers 6849, Institute for the Study of Labor (IZA).
- Petrongolo, B., Pissarides, C. A., June 2001. Looking into the black box: A survey of the matching function. *Journal of Economic Literature* 39 (2), 390–431.
- Pissarides, Christopher, A., October 1986. Unemployment and vacancies in Britain. *Economic Policy* III, 499–559.
- Pissarides, Christopher, A., November 1992. Loss of skill during unemployment and the persistence of employment shocks. *The Quarterly Journal of Economics* 107 (4), 1371–1391.
- Polivka, A. E., Miller, S. M., 1998. The cps after the redesign: Refocusing the economic lens. In: *Labor Statistics Measurement Issues*. University of Chicago Press, pp. 249–289.
- Pries, M., July 2008. Worker heterogeneity and labor market volatility in matching models. *Review of Economic Dynamics* 11 (3), 664–678.
- Pries, M., Rogerson, R., August 2005. Hiring policies, labor market institutions, and labor market flows. *Journal of Political Economy* 113 (4), 811–839.
- Pries, M. J., January 2004. Persistence of employment fluctuations: A model of recurring job loss. *Review of Economic Studies* 71 (1), 193–215.

- Ravenna, F., Walsh, C. E., December 2012. Screening and labor market flows in a model with heterogeneous workers. *Journal of Money, Credit, and Banking* 44 (2).
- Ravn, M. O., Uhlig, H., 2002. On adjusting the hodrick-prescott filter for the frequency of observations. *Review of economics and statistics* 84 (2), 371–376.
- Rotemberg, J. J., Woodford, M., 1991. Markups and the business cycle. In: *NBER Macroeconomics Annual 1991, Volume 6*. MIT Press, pp. 63–140.
- Rothschild, M., 1973. Models of market organization with imperfect information: A survey. *The Journal of Political Economy*, 1283–1308.
- Sahin, A., Song, J., Topa, G., Violante, G. L., 2012. Mismatch unemployment. Working paper, NYU.
- Schaal, E., 2012. Uncertainty, productivity and unemployment in the great recession. Federal Reserve Bank of Minneapolis, mimeo.
- Schoeni, R., Dardia, M., 2003. Estimates of earnings losses of displaced workers using california administrative data. Tech. Rep. 03-543, PSC Research.
- Shimer, R., March 2005. The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review* 95 (1), 25–49.

- Shimer, R., September 2007. Mismatch. *American Economic Review* 97 (4), 1074–1101.
- Shimer, R., April 2012. Reassessing the ins and outs of unemployment. *Review of Economic Dynamics* 15 (2), 127–148.
- Stahl, D. O., 1990. Entropy control costs and entropic equilibria. *International Journal of Game Theory* 19 (2), 129–138.
- Stevens, A. H., January 1997. Persistent effects of job displacement: The importance of multiple job losses. *Journal of Labor Economics* 15 (1), 165–88.
- Stroebel, J., Vavra, J., 2014. House prices, local demand, and retail prices. Tech. rep., National Bureau of Economic Research.
- Tasci, M., 2012. The ins and outs of unemployment in the long run: Unemployment flows and the natural rate. Working paper, Federal Reserve Bank of Cleveland.
- Tauchen, G., 1986. Finite state markov-chain approximations to univariate and vector autoregressions. *Economics letters* 20 (2), 177–181.
- U.S. Census Bureau, 2011. Synthetic LBD Beta version 2.0. [computer file], U.S. Census Bureau and Cornell University, Synthetic Data Server [distributor],

Washington,DC and Ithaca, NY, USA.

URL <http://www2.vrdc.cornell.edu/news/data/lbd-synthetic-data/>

Valletta, R. G., 1991. Job tenure and joblessness of displaced workers. *The Journal of Human Resources* 26 (4), 726–741.

Walsh, C. E., 2005. Labor market search, sticky prices, and interest rate policies. *Review of economic Dynamics* 8 (4), 829–849.

Wasmer, E., 2004. Labor supply dynamics, unemployment and experience in the labor market. *Recherches conomiques de Louvain* 70 (4), 461–482.