

THE UNIVERSITY OF CHICAGO

ESSAYS ON THE MACROECONOMIC AND FINANCIAL CAUSES OF THE GREAT
RECESSION

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE SOCIAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF ECONOMICS

BY
JUAN JOSE OSPINA TEJEIRO

CHICAGO, ILLINOIS

JUNE 2017

Copyright © 2017 by JUAN JOSE OSPINA TEJEIRO
All Rights Reserved

Dedication Text

This dissertation and my PhD are dedicated to all the people who made possible that I could finish my doctoral studies successfully, become a better economist, and a better person. First and foremost I want to dedicate this dissertation to my wife Ana Maria Serrano Angel. With her support, encouragement, and patience, she made sure that I could go through all the challenges, achievements, frustrations, and deceptions that characterize the experience of a graduate student. At the same time, she made sure my years in Chicago were the best of my life and that I got to every milestone. This PhD is hers as much as it is mine. All this work is also dedicated to Olivia Ospina Serrano, our daughter, who was born at the end of my studies. It is my hope that this work serves as example of dedication, effort, and perseverance, qualities that I will strive to leave her with.

I want to thank my parents, Clemencia Tejeiro and Jorge Alberto Ospina, for giving me all the foundations that I needed to get to this point and for inculcating me with discipline and love for learning. I also want to give thanks to my brothers and sister Alejandro, Ana Maria y Daniel, for being supportive over all these years of this adventure and for your unconditional, genuine interest in my success.

I am indebted to my advisors, Erik Hurst, Harald Uhlig and Michael Weber. They have been my mentors throughout and have made a difference in my career as an economist and a researcher. I want to thank them for their teachings, for giving me the opportunity to learn how to be a rigorous scholar and how to perform quality research, and also for their support, patience, and encouragement. I will always be grateful to Arnold Zellner, Pietro Veronesi, Christian Broda, and Alberto Carrasquilla for having given me the experience, advice and recommendations I needed to enter the PhD.

What I have achieved and learned during this time is to a very large degree due to my classmates at The University of Chicago: Pietro Bonaldi, Sumiko Hayasaka, Ignacia

Mercadal, Chenfei Lu, Jorn Boehnke, Philip Barrett, Seth Blumberg, Martin Beraja, Sara Moreira, Fabrice Tourre, Paulo Mateus, Andres Hojman, Artur Bezerra, Santiago Caicedo, Chen Yeh, Aaron Pancost, Felipe Labbe, Alexander McKay, Sergiy Verstyuk, Serginio Sylvain, Chattrin Laksanbunsong, David Coble, Chanont Banterngansa, Nicolas Castro Cienfuegos, Nuno Paixao, and Elisa Giannone.

My work over these years is also dedicated to the following people who made a difference in my life and my intellectual development during my time in Chicago: Gary Becker, Victor Lima, Virginia Bova, Christine Gramhofer, Amy Lee Boonstra, Lars Hansen, Drew Creal, Federico Bandi, George Constantinides, John Cochrane, Nicholas Paulson, Mark Hendricks, Roger Lee, Maria Rios, Andres Naranjo, Juan Carlos Fonnegra, Natasha Bloch, Mariana Laverde, Luciana Godoy, Valeria Schewk, Esperanza Johnson, Claudia Angel, Robert Herbst and all the staff of the Executive MBA program at Chicago Booth in Chicago, Hong Kong, Singapore and London. I am also very grateful to the University of Chicago, the Department of Economics and the Booth School of Business for giving me the opportunity to embark on this intellectual challenge at one of the most outstanding academic institutions in the world. Finally, I am very thankful to the Banco de la República de Colombia for providing me with the financial resources required to complete my studies successfully; I look forward to giving back to Colombia and being instrumental to our economic well-being and development.

TABLE OF CONTENTS

LIST OF FIGURES	viii
LIST OF TABLES	xii
ACKNOWLEDGMENTS	xiv
ABSTRACT	xv
1 REGIONAL BUSINESS CYCLE ACCOUNTING AND THE GREAT RECESSION	1
1.1 Abstract	1
1.2 Introduction	1
1.3 A Benchmark Model	8
1.3.1 Households	9
1.3.2 Firms	11
1.3.3 Market Clearing Conditions, Government, and Equilibrium	12
1.4 A friction and the corresponding wedge	13
1.5 Data	15
1.5.1 Regional Data	15
1.5.2 Aggregate Data	25
1.6 Model Estimation	29
1.6.1 A discussion on heterogeneity	34
1.7 Business Cycle Accounting Exercise	35
1.8 Sensitivity of Results	43
1.9 Discussion and Conclusions	46
1.10 Figures and Tables	49
1.11 Appendix	87
1.11.1 Aggregation	87
1.11.2 Stochastic Process Estimation	90
1.11.3 Regional Price Data	92
1.11.4 Computing the Value of the Housing Stock	93
1.11.5 List of Data Sets	97
1.11.6 County, MSA, and Other Local Data	98
1.11.7 Appendix Tables and Figures	105
1.12 On-Line Appendix	109
1.12.1 Additional Data Plots	110
1.12.2 Alternative Business Cycle Accounting Metrics	120
1.12.3 Solution to Regional Model	131
1.12.4 Solution to Aggregate Model	152
1.12.5 Aggregation Results	168
1.12.6 Aggregation Result 2 - Regional Economies	172

2	THE AGGREGATE IMPLICATIONS OF REGIONAL BUSINESS CYCLES . . .	178
2.1	Abstract	178
2.2	Introduction	179
2.3	Creating State-Level Price And Wage Indices	187
2.3.1	State-Level Wage Index	187
2.3.2	State-Level Price Index	189
2.4	Comparing Cross-State Patterns to Aggregate Time-series Patterns	197
2.5	A Monetary Union Model	200
2.5.1	Firms and Households	202
2.5.2	Sticky Wages	204
2.5.3	Equilibrium	205
2.5.4	Shocks	206
2.5.5	Aggregation	207
2.5.6	Aggregate vs. Local Shock Elasticities	209
2.6	A Procedure for Identifying Aggregate Shocks	214
2.7	Estimating the Wage Setting Equation Using Regional Data	218
2.7.1	Estimating Equation and Identification Assumptions	219
2.7.2	Estimates of λ and ϕ	224
2.8	The US Great Recession: From Regional to Aggregate	227
2.8.1	Findings in the Aggregate	227
2.8.2	Sensitivity to Alternative Parameters λ and ϕ	231
2.8.3	Sensitivity to Alternative Identifying Assumptions	233
2.8.4	Discussion: Aggregate v. Regional Shock Decomposition	236
2.9	Conclusion	237
2.10	Figures and Tables	240
2.11	Appendix	252
2.11.1	Proof of Lemma 1 and 2	252
2.11.2	Alternative wage setting specifications	254
2.12	Appendix Figures	258
2.13	Online Appendix	259
2.13.1	Descriptive Statistics For Retail Scanner Data	259
2.13.2	Creating the Scanner Data Price Index	260
2.13.3	Creating Composition Adjusted Wage Measures in the ACS and CPS	262
2.13.4	BLS Metro Area Price Indices	265
2.13.5	Cross-State Wage Patterns in the QEW	267
2.13.6	Accounting for Measurement Error in Wages and Prices	268
2.13.7	Variation in Government Policy Changes Across States During Great Recession	272
2.13.8	Controlling for Industry Controls in Estimates λ and ϕ	276
2.13.9	IV Procedure to Estimate λ and ϕ	277
2.13.10	Parameter Calibration	278
2.13.11	Estimating the Regional Shocks	280
2.13.12	Findings at the regional level	282

2.13.13 Online Appendix Figures	285
3 MORTGAGE-BACKED SECURITIES AND THE FINANCIAL CRISIS OF 2008: A POST MORTEM	294
3.1 Abstract	294
3.2 Introduction	294
3.3 The data	299
3.3.1 Database description	302
3.3.2 Classifying Securities into Market Segments	303
3.4 Analysis	304
3.4.1 Ratings	305
3.4.2 Losses	308
3.4.3 Credit Rating Rankings and Reversals	316
3.4.4 Returns	321
3.4.5 Prices	326
3.5 Discussion and Conclusion	328
3.6 Figures and Tables	330
3.7 Appendix Figures	360
REFERENCES	365

LIST OF FIGURES

1.1	Regional Variation in Employment During the Great Recession: 2007-2010	49
1.2	Regional Variation in Employment and Consumption over the Business Cycle: 2007-2014	50
1.3	Household Side of the Labor Wedge and the Business Cycle	51
1.4	Migration, Consumption Growth and Employment Growth: 2007-2010	52
1.5	State Level Consumption Growth vs Employment Growth: 2007-2010	53
1.6	State Level Wage Growth vs Employment Growth: 2007-2010	54
1.7	State Level Price Growth vs Employment Growth: 2007-2010	55
1.8	Net Asset Growth vs Employment Growth: 2007-2010	56
1.9	Great Recession in California	58
1.10	Relative Employment and Employment Components in California . . .	59
1.11	Relative Consumption and Consumption Components in California . .	60
1.12	Great Recession in Nevada	61
1.13	Relative Employment and Employment Components in Nevada	62
1.14	Relative Consumption and Consumption Components in Nevada . . .	63
1.15	Great Recession in New York	64
1.16	Relative Employment and Employment Components in New York . .	65
1.17	Relative Consumption and Consumption Components in New York . .	66
1.18	Decomposition of Employment One-Wedge Economy: 2007-2014 . . .	67
1.19	Decomposition of Employment Three-Wedge Economy: 2007-2014 . .	68
1.20	Relative Employment and Employment Components in Tennessee . . .	69
1.21	Decomposition of Consumption One-Wedge Economy: 2007-2014 . . .	72
1.22	Decomposition of Consumption Three-Wedge Economy: 2007-2014 . .	73
1.23	Sensitivity Analysis - Decomposition of Employment Three-Wedge Econ- omy: 2007-2014. Elasticity of Labor Supply $\phi = 1$	76
1.24	Sensitivity Analysis - Decomposition of Consumption One-Wedge Econ- omy: 2007-2014. Elasticity of Labor Supply $\phi = 1$	77
1.25	Sensitivity Analysis - Decomposition of Employment Three-Wedge Econ- omy: 2007-2014. Elasticity of Labor Supply $\phi = 4$	78
1.26	Sensitivity Analysis - Decomposition of Consumption One-Wedge Econ- omy: 2007-2014. Elasticity of Labor Supply $\phi = 4$	79
1.27	Sensitivity Analysis - Decomposition of Employment Three-Wedge Econ- omy: 2007-2014. Intertemporal Elasticity of Substitution $\sigma = 1$	80
1.28	Sensitivity Analysis - Decomposition of Consumption One-Wedge Econ- omy: 2007-2014. Intertemporal Elasticity of Substitution $\sigma = 1$	81
1.29	Sensitivity Analysis - Decomposition of Employment One-Wedge Econ- omy: 2007-2014. GHH Preferences	82
1.30	Sensitivity Analysis - Decomposition of Employment Three-Wedge Econ- omy: 2007-2014. GHH Preferences	83
1.31	Sensitivity Analysis - Decomposition of Consumption One-Wedge Econ- omy: 2007-2014. GHH Preferences	84

1.32	Sensitivity Analysis - Decomposition of Consumption Three-Wedge Economy: 2007-2014. GHH Preferences	85
1.33	MSA Level Wage Growth vs Employment Growth: 2007-2010	98
1.34	County Level Wage Growth vs Employment Growth: 2007-2010	99
1.35	MSA Level Price Growth vs Employment Growth: 2007-2010	100
1.36	County Level Price Growth vs Employment Growth: 2007-2010	101
1.37	State Level Real Wage Growth vs Employment Growth: 2007-2010	102
1.38	MSA Level Real Wage Growth vs Employment Growth: 2007-2010	103
1.39	County Level Real Wage Growth vs Employment Growth: 2007-2010	104
A1	Sensitivity Analysis - Decomposition of Employment Three-Wedge Economy: 2007-2014. Intertemporal Elasticity of Substitution $\sigma = 4$	105
A2	Sensitivity Analysis - Decomposition of Consumption One-Wedge Economy: 2007-2014. Intertemporal Elasticity of Substitution $\sigma = 4$	106
A3	Sensitivity Analysis - Decomposition of Employment Three-Wedge Economy: 2007-2014. Intertemporal Elasticity of Substitution $\sigma = 4$ and Elasticity of Labor Supply $\phi = 1$	107
A4	Sensitivity Analysis - Decomposition of Consumption One-Wedge Economy: 2007-2014. Intertemporal Elasticity of Substitution $\sigma = 4$ and Elasticity of Labor Supply $\phi = 1$	108
A5	State Level Consumption Growth vs Employment Rate Change: 2007-2010	110
A6	State Level Wage Growth vs Employment Rate Change: 2007-2010	111
A7	MSA Level Wage Growth vs Employment Rate Change: 2007-2010	112
A8	County Level Wage Growth vs Employment Rate Change: 2007-2010	113
A9	State Level Price Growth vs Employment Rate Change: 2007-2010	114
A10	MSA Level Price Growth vs Employment Rate Change: 2007-2010	115
A11	County Level Price Growth vs Employment Rate Change: 2007-2010	116
A12	State Level Real Wage Growth vs Employment Rate Change: 2007-2010	117
A13	MSA Level Real Wage Growth vs Employment Rate Change: 2007-2010	118
A14	County Level Real Wage Growth vs Employment Rate Change: 2007-2010	119
A15	Net Asset Growth vs Employment Rate Change: 2007-2010	120
A16	Metric 1- Average Distance from Zero for Employment. Baseline Parameters	121
A17	Metric 2 - Movement towards Zero for Employment. Baseline Parameters	122
A18	Metric 3 - Variance Decomposition for Employment. Baseline Parameters	123
A19	Metric 1- Average Distance from Zero for Consumption. Baseline Parameters	124

A20	Metric 2 - Movement towards Zero for Consumption. Baseline Parameters	125
A21	Metric 3 - Variance Decomposition for Consumption. Baseline Parameters	126
A22	Metric 3 - Variance Decomposition for Employment. Parameters ($\sigma = 1, \phi = 2$)	127
A23	Metric 3 - Variance Decomposition for Employment. Parameters ($\sigma = 2, \phi = 1$)	128
A24	Metric 3 - Variance Decomposition for Consumption. Parameters ($\sigma = 1, \phi = 2$)	129
A25	Metric 3 - Variance Decomposition for Consumption. Parameters ($\sigma = 2, \phi = 1$)	130
2.1	Nielsen Retail Price Index vs. CPI Food Price Index	240
2.2	The Evolution of Aggregate Real and Nominal Composition Adjusted Wages	241
2.3	State Employment Growth vs. State Nominal and Real Wage Growth, 2007-2010	242
2.4	Impulse Response to a Discount rate Shock	243
2.5	Impulse Response to a Productivity / Markup Shock	243
2.6	Impulse Response: Leisure Shock	244
2.7	Shock time-series	244
2.8	Counterfactual Employment Response	245
2.9	Counterfactual Price Response	245
2.10	Counterfactual Wage Response	246
B1	State Net Migration Rate 2009-2010 vs. State Employment Growth 2007-2010	258
B2	QEW Nominal Per Worker Earnings Growth vs. ACS Adjusted Nominal Wage Growth, 2007-2010	287
B3	QEW Nominal Per Worker Earnings Growth vs. QEW Employment Growth, 2007-2010	288
B4	Change in State Tax Rate vs. State Employment Growth: 2007-2010	289
B5	State SNAP Growth vs. State Employment Growth: 2007-2010	290
B6	State HAMP Take-Up vs. State Employment Growth 2007-2010	291
B7	Max Unemployment Benefit Receipt in 2010 vs. State Employment Growth 2007-2010	292
B8	Routine Share of Employment in 2007 vs. State Employment Growth 2007-2010	293
3.1	Distribution of Mean FICO Score, Loan Size, and LTV by Type of Loan	333
3.2	Losses and Unweighted Probability of Loss in RMBS Over Time	336
3.3	Losses Over Time by Type of Mortgage for AAA-Rated RMBS	337
3.4	Dollar Amount of Losses in Non-Agency RMBS	338
3.5	Distribution of Loss Size for All RMBS	339

3.6	Vintage Fixed Effects on Weighted Losses	342
3.7	Vintage Fixed Effects on Weighted Losses by Type of Mortgage Loan	343
3.8	State-Level House Price Boom and Bust From 2000-Q1 to 2009-Q2 . .	344
3.9	State-Level Dummies on Loss Rates: with and without controls	345
3.13	Subprime RMBS Price Indexes	346
3.10	Ex-Ante vs Ex-Post Ideal Ratings	347
3.11	”Right” and ”Wrong” Ratings Based on Moody’s Ideal Ratings	348
3.12	”Misratings” Ratings Based on Moody’s Ideal Ratings	349
3.14	Summary Statistics of Prices Collected by FINRA	350
3.15	Distribution of FICO Scores by Type of Mortgage	360
3.16	Vintage Fixed Effects on Probability of Loss	361
3.17	Vintage Fixed Effects on Weighted Probability of Loss by Type of Mortgage Loan	362
3.18	Non-weighted Probability of Loss Over Time by Type of Mortgage and by Credit Rating	363
3.19	Moody Idealized Cumulative Expected Loss Rates	364

LIST OF TABLES

1.1	Parameter Values	57
1.2	ψ -statistic for Employment: 2007-2014	70
1.3	Peak to Trough Changes in Employment and Components and Wedge-based Classification of States	71
1.4	ψ -statistics for Consumption: 2007-2014	74
1.5	Peak to Trough Changes in Consumption and Components and Wedge-based Classification of States	75
1.6	Cross-Region Employment and Wedges	86
1.7	Nielsen’s Retail Scanner Database Description	92
1.8	Data Sources to Value the Housing Stock	94
1.9	Data Sources	97
2.1	Comparison of Cross-State and Time-Series Estimates of Wage Elasticities During the Great Recession	247
2.2	Calibration	248
2.3	Aggregate v. Regional Employment Impact Elasticities	248
2.4	Aggregate Elasticities ($\lambda=0.7, \phi=2$)	249
2.5	Regional Elasticities ($\lambda=0.7, \phi=2$)	249
2.6	Estimates of λ and $\frac{\lambda}{\phi}$ using Cross- Region Data	250
2.7	Discount/Interest rate (γ) and Productivity/Markup (z) shocks’ contribution to aggregate employment change	251
B8	Descriptive Data for the Nielsen Scanner Price Data	285
B9	Discount rate (γ) and Productivity/Markup (z) shocks’ contribution to change in regional variables	286
B10	Regional counterfactual statistics	286
3.1	RMBS Database: Deals, Securities, Nominal Amounts by Year of Issuance	330
3.2	Database Variables	331
3.3	Database Coverage of the universe of Non-Agency RMBS	332
3.4	Credit Rating Activity of Non-Agency Residential Mortgage Backed Securities: 1987-2013	334
3.5	Non-Agency Residential Mortgage Backed Securities: Credit Rating Composition 1987-2013	334
3.6	RMBS Losses Six Years After The Beginning of the Financial Crisis	335
3.7	Losses and Credit Ratings by Vintage Group and Type of Mortgage Loan	340
3.8	Principal-Weighted Losses in RMBS and Credit Ratings	341
3.9	House Prices and Loss Rates	351
3.10	Credit Rating Reversals	352
3.11	Principal-Weighted Probability of Loss in RMBS and Credit Ratings	353

3.12	Internal Rate of Return Calculations From Issuance to 2013 by Credit Rating. For AAA: also Type of Mortgage.	354
3.13	Discount Margin (Return Over 3-month Tbill) Calculations From Issuance to 2013 by Credit Rating	355
3.14	Principal Value Weighted Discount Margin (Return Over 3-month Tbill) Calculations by Vintage	356
3.15	Discount Margin (Return Over 3-month Tbill) Calculations From Issuance to 2013 by Type of Mortgage Loan	357
3.16	Median Discount Margin (Return Over 3-month Tbill) for AAA-rated RMBS by Vintage, Loan Type and Bond Type	358
3.17	Subprime RMBS Returns for Deals Underlying ABX.HE Indexes . . .	359

ACKNOWLEDGMENTS

I am thankful for the advice and guidance of my advisers, Erik Hurst, Harald Uhlig, and Michael Weber. I thank Fernando Alvarez, Peter Klenow, Ellen McGrattan, and Robert Shimer for their insightful perspective. I thank participants of the Capital Theory Workshop, Applied Macro Theory Group, and International Trade Working Group at The University of Chicago. I also thank Daniel Tannenbaum and seminar participants at the Minnesota Workshop in Macroeconomic Theory, NBER's Summer Institute EF&G, NBER's Summer Institute Prices Program, NBER's Summer Institute Capital Markets and the Economy Program, Bank of England, Banco de la República de Colombia, University of Nebraska, the Federal Reserve Bank of Boston, UC Santa Barbara, Claremont McKeena College, Universidad de los Andes and Universidad del Rosario. I owe my classmates Martin Beraja, Pietro Bonaldi, Santiago Caicedo, David Coble, Sumiko Hayasaka, Felipe Labbe, Paulo Mateus, Sara Moreira, and Chen Yeh for their valuable suggestions. I especially thank Ana Serrano for her support, comments and help to write these papers.

ABSTRACT

This dissertation is composed of three essays that study the macroeconomic and financial causes of the Great Recession. The first chapter focuses on understanding some of the business cycle dynamics of different regions in the United States. In particular, I seek to understand what shocks and frictions are the drivers of consumption and employment differences across subnational economies, particularly states. I find that the shocks and frictions that drive the aggregate business cycle are not enough to understand regional business cycle dynamics. In this chapter I develop methodological contributions that can help researchers guide the construction of models whose goal is to understand regional business cycle dynamics and how it relates to aggregate business cycle dynamics. The second chapter focuses on understanding the link between regional and aggregate business cycles. We find that the shocks that we can identify using cross-sectional variation are insufficient to understand the joint dynamics of prices, wages and employment at business cycle frequencies. In particular, demand shocks identified using cross-region variation are insufficient to explain the persistent decline in aggregate employment. This chapter develops methodological contributions to identify shocks in macroeconomic models and to construct regional indexes for prices and wages. The third chapter is an empirical analysis of the non-agency mortgage backed securities market, which has been at the core of the explanations of the causes of the Great Recession. By carefully studying the cash flows, returns, and how they relate to the credit ratings, we find that contrary to the conventional narrative of the crisis, AAA-rated subprime mortgage backed securities performed remarkably well. This calls into question some key aspects of the explanations that have been given as triggers of the crisis of 2008, and points at the need to better understand the forces behind this event in order to have a more accurate understanding and be able to prescribe appropriate policies.

CHAPTER 1

REGIONAL BUSINESS CYCLE ACCOUNTING AND THE GREAT RECESSION

1.1 Abstract

I extend the business cycle accounting methodology to a setting of a monetary union. I construct a novel dataset on prices, wages, employment, net assets, and consumption that using both aggregate and regional data allows for the application of the methodology at three different levels of geographic aggregation: states, MSAs, and counties. Applied to the Great Recession at the state level, the business cycle accounting exercise provides two main findings. First, for 40 out of 48 states the labor wedge played a primary role in accounting for the differences between employment at the state level and employment at the aggregate level. Second, for 42 states the intertemporal wedge played a prominent role in accounting for the differences between consumption at the state level and consumption at the aggregate level. These results suggest that models using regional variation to study the business cycle of the Great Recession would need mechanisms generating fluctuations in more than one wedge to account for relative fluctuations in employment and consumption of a given region; however, in principle, such mechanisms need not be different for different regions.

1.2 Introduction

Economic outcomes across regions within a country may vary greatly over the business cycle. A recent, representative example is the United States during the Great Recession. Figures 1.1 and 1.2 illustrate this fact in a cross sectional and time series fashion respectively. Figure 1.1 shows the percentage point change in employment rate between 2007 and 2010 for states and counties. The cross region variation is large, a state like Nevada suffered nearly a 9

percentage point decrease in the fraction of people working while North Dakota only saw a negative change slightly higher than 1 percentage point. For counties the cross-region variation is orders of magnitude bigger even within a given state. Figure 1.2 plots the cyclical components of employment measured as total hours worked per person and real per capita consumption for three sample states. Not only the amplitude of the cycle was different across states but also the speed of recovery. While employment and consumption in New York were almost back with respect to trend by 2014, employment was still 4 percentage points below trend for Nevada and consumption was about 1.5 percentage points below trend for both California and Nevada.

What are the reasons behind this geographically heterogenous aggregate fluctuations? There are several candidates. One reason may be that the size of the shocks hitting local economies might be different.¹ A second possibility is that different regions have different economic structures, which may make them more or less sensitive to certain types of shocks.² Finally, another possible explanation is that regions may differ in certain institutions and/or frictions that affect economic outcomes. For example, labor income taxes and minimum wages vary by state, and certain market frictions such as the degree of wage rigidity could also be different across different regional labor markets.

There have been a large number of studies using cross-region variation to explore some of these possibilities.³ However, Beraja, Hurst, and Ospina (2016) have shown that cross-sectional, partial equilibrium studies are limited in their ability to inform us about the

1. Some good examples of this are provided by Mian et al. (2013) and Mian and Sufi (2014) who use regional variation in declining house prices and household leverage to identify the so called “housing net worth shock” and explore its effects on consumption and employment via demand channels.

2. An example would be Charles et al. (2016) who show that MSAs exhibited differences in the level and composition of manufacturing, which resulted in differential effects of the decline of manufacturing in employment, wages, and population. This paper is more about long run changes than changes at business cycle frequencies, but it does illustrate the importance of variation in the productive structure of the economy in understanding economic outcomes dynamics

3. Examples include Stroebel and Vavra (2016), Mian et al. (2013), Mian and Sufi (2014), Hagedorn et al. (2016), Mondragon (2015), Mehrotra and Sergeyev (2015), Autor et al. (2013), Charles et al. (2016),

aggregate economy and that the use of cross-region variation needs theory to be able to inform us about aggregate business cycles.

With this in mind, this paper seeks to contribute to our understanding of the sources of aggregate fluctuations by answering two questions that would help guide theory. First, what type of frictions/shocks are most promising in explaining these differential business cycle fluctuations across regions? How similar are these types of shocks across regions? To answer these questions I take a business cycle accounting approach and extend the methodology put forth by Chari, Kehoe, and McGrattan (2007a) to a setting of a monetary union in which different regions are connected through the actions of a monetary authority and trade of intermediate goods.

The business cycle accounting methodology is based on the so-called “wedges”. Wedges are discrepancies in the relationships between economic variables/quantities implied by either resource constraints or first order conditions of the decision problems of agents in a model economy when taken to the data. I take the view that wedges are related to meaningful shocks or frictions and are not just errors that result from the failure of theory to fit the data. This view is the one that makes the accounting exercise meaningful⁴. To support this view I find useful the fact that wedges exhibit systematic behavior over the business cycle⁵, so that understanding their behavior at business cycle frequencies could prove useful to learn about aggregate fluctuations and improve theory. An example is provided by the labor wedge, the difference between the marginal rate of substitution between consumption and leisure from the households’ theoretical labor decision and the real wage.⁶ To be specific,

4. See Brinca et al. (2016) for a detailed exposition

5. Clearly, one would not expect to see that theory holds exactly in data just from the fact that the variables we use are measured with error

6. In this paper the labor wedge will be just the wedge in the first order condition of the households with respect to their labor decision, while the productivity wedge will represent the firm’s side. The literature on the labor wedge started by looking at the labor wedge as the difference between the marginal product of labor from the firm’s labor decision and the marginal rate of substitution from the consumer’s labor supply decision

the labor wedge is countercyclical, it increases in recessions which makes look recessions as periods when people dislike to work, in the same way as a tax on labor income would theoretically affect labor. Figure 1.3 shows this fact in two ways. First by relating one-period (time series) changes in the cyclical component of aggregate labor to one-period changes in the aggregate labor wedge from 1976 through 2014. Second, by showing that this negative relationship is also a feature of the Great Recession period between 2007-2010 using the cross-state variation in the labor wedge and in state-level employment, per-capita income, and real per-capita consumption. The countercyclical nature of the aggregate labor wedge is a very well documented fact (see for example Shimer (2009)). The fact that the labor wedge was important during the Great Recession has also been documented in the literature (see for example Ohanian (2010)). Here I just show that this fact is also present when one looks at the cross section of states.

The accounting part of the methodology provides a general equilibrium aspect to the wedges. Clearly, to study the household side of labor wedge one only needs the first order condition of the households. However, different shocks/frictions and equations are involved in determining the variables in the equation with the labor wedge. Furthermore, wedges are correlated in the data. Therefore, recovering the other wedges is necessary to gauge the relative importance of the shocks/friction working through the labor wedge in accounting for fluctuations in the data.

I start by adapting the business cycle methodology of Chari, Kehoe, and McGrattan (2007a) to a context of regional economies within a monetary union. To do so, in section 1.3 I first write a simple model of a monetary union in which regions, represented by islands, are connected through trade of intermediate goods and a central monetary authority. The model is intended to play the role of a benchmark model in the same way as the Neoclassical Growth Model plays the role of benchmark for national economies in Chari et al. (2007a). I write the model to have little structure (no explicit shocks or frictions) as it is the idea

behind the wedge accounting procedure, while at the same time capturing some facts of the Great Recession, for example that wages and prices behaved differentially across regions (see Beraja et al. (2016)). The way the model is written and solved, allows one to apply the methodology for different levels of regional aggregation (States, MSAs, and Counties) since all the variables can be created from existing data for the three types of regions.

I then proceed to construct all the variables at the aggregate and regional level required to estimate the model and recover the wedges: wages, prices, employment, and net assets. For states I also employ data on consumption. The application of the business cycle methodology involves the estimation of a relatively large number of parameters. I use both aggregate data that goes back to 1976 and the cross-region data to maximize the available data for the estimation. Section 1.5 explains in detail the construction of the different variables. The construction of the dataset involves the use of about 25 different datasets, some of which are micro and some of which are already aggregated at some geographic level. For a complete list of the datasets I use in the paper see appendix table 1.9.

With the model and the data, in section 1.6 I proceed to estimate the parameters of the stochastic processes that govern the evolution of wedges. To this effect following the implementation proposed by Chari, Kehoe, and McGrattan (2007a), I log-linearize the model around a deterministic steady state, and then using the policy functions of the model I estimate the parameters by maximum likelihood via a Kalman Filter. In this section I lay out the conditions that make both the estimation and the application of the accounting methodology tractable. These conditions are to a large degree also imposed by the fact that data at the regional level is available only from 2005 and at annual frequencies. Under these conditions the aggregation properties of the model can be used to 1) facilitate the estimation by making the hidden states of the Kalman Filter of a given region independent from those of other regions 2) enable the use of cross-sectional data as well as aggregate data to estimate a relatively large number of parameters.

In section 1.7 I apply the business cycle accounting methodology by conducting the exercise with state-level data. The exercise is done by using the estimated parameters and the decision rules of the model to measure and recover the wedges. Then the wedges are fed back into the model one at a time or in groups to produce the wedge-components of employment and consumption and assess which wedge is most important in accounting for the fluctuations of these variables for each state. In section 1.8 I explore the sensitivity of results to changes in parameters, preferences and data.

The results from the exercise and from the sensitivity exercises indicate that the labor wedge played a prominent role in accounting for the fluctuations in employment for most states. In the baseline specification, that labor wedge accounts for 50% or more of the movements in employment for 31 states (60% of the regions) when fed into the model by itself. For only 7 states, the labor wedge accounts for less than 20% of the movements in output. The importance of the labor wedge is reinforced when one studies an economy with three wedges except for the labor wedge. For only 6 states, an economy without the labor wedge accounts for more than 50% of fluctuations in employment. Analyzing the decomposition state by state I find that the labor wedge was the main driver of differences in employment fluctuations across states, with 37 states having the labor wedge as the main peak-trough (2007-2010) wedge and 40 states having it as the main driver of employment fluctuations for the period 2007-2014.

For consumption I find that the intertemporal wedge by itself accounts for more than 40% of consumption for about 34 states. For the majority of states (31 out of 48), the three-wedge economy that does not contain the intertemporal wedge is unable to account for more than 50% of the movements in relative consumption. This pattern is stronger for a range of different preference parameters and for a change in preferences. A detailed analysis of the business cycle decomposition state by state shows that the intertemporal wedge was the main wedge driving fluctuations in consumption during the downturn (2007-2010) for 37

states and played a secondary but important role for another 5 states, while it play a key role for 40 states when the entire business cycle period (2007-2014) is analyzed.

Overall, the results indicate three main implications for models that attempt to capture or explain the cross-regional variation of employment and consumption. First, models with frictions and shocks that produce fluctuations in the labor wedge will have a better chance of matching the differential behavior of employment while models with frictions and shocks that produce fluctuations in the intertemporal wedge will be most promising in explaining the variation in consumption. Second, even though the Great Recession at the aggregate level looks like a labor wedge recession (see Ohanian (2010) and Brinca et al. (2016)), at the regional level variation in only one wedge will not be able to produce the employment and consumption variation that we see in the data. The implication is that model builders should aim to have features in the models working through two wedges. Third, the fact that one wedge accounts for a large fraction of the movements in employment for most states and that one wedge accounts for most of the movements in consumption for most states, indicates that models in principle do not need a zoo of frictions and shocks to explain the regional variation that we observe in the data.

Crucially, these results are in contrast with the results from a business cycle accounting exercise applied to aggregate data and aggregate models. From this exercise we would not rule out models that have undistorted Euler equations (no intertemporal wedge) and one would conclude that the Great Recession was a labor-wedge recession. However, when one opens the aggregate black box to study regional economies and takes into account that markets are incomplete, that it is difficult at business cycle frequencies to shift people and consumption around so that marginal utilities equalize, the Great Recession no longer looks like a labor wedge-only recession. From regional data and regional variation, shocks and frictions that distort the Euler Equation are crucial to understand the dynamics of consumption and we would rule out models that have undistorted Euler equations.

1.3 A Benchmark Model

The purpose of this section is to present a model that serves as a benchmark for the class of models that I will study. This model plays the same role for an economy composed of many regions as the Neoclassical Growth Model for a national economy in Chari et al. (2007a). The model is not a multi-region version of the Neoclassical Growth Model due to data limitations. Specifically, the model does not have investment and capital as these variables are unavailable at the local level. The model has nominal variables like prices and wages and the state variable represents net savings. The idea is to have a simple model for a monetary union with no explicit frictions or shocks. Instead, the model contains time varying wedges. These wedges are reduced form for shocks and frictions in otherwise explicit models that can be mapped into the benchmark. In section 1.4 I illustrate via one example how these models with explicit frictions and shocks can be mapped into the benchmark model of this section. The insight from business cycle accounting is that variation in the wedges represents variation related to the shocks and how they get amplified by frictions; when variation in a wedge accounts for variation in a variable of the model, then a explicit model that produces such wedge provides a promising mechanism to explain the fluctuations that we see in the data at business cycle frequencies.

The general characteristics of the regional benchmark model are that it represents an economy composed of a large number of ex-ante identical⁷ islands (local economies or regions). These islands belong to a monetary union with a monetary authority setting an interest rate common for all islands. Islands are also connected through trade of intermediate goods, which are used to produce a final non-tradable good. Examples of models with this structure can be found in the literature that studies the Great Recession, for instance in Beraja et al. (2016) and Midrigan and Philippon (2016). These models vary in the shocks

7. In the section 1.6 I discuss the implications of moving away from this assumption

and frictions they entertain. For example Beraja et al. (2016) have sticky wages with labor supply, discount rate, and productivity shocks, while Midrigan and Philippon (2016) have price rigidities, wage rigidities, and financial frictions.

The idea of the model is to impose little structure while having a model that can be fully estimated with existing data and that is truthful to empirical facts in the data like the wage and price dispersion across regions and which I document in section 1.5.⁸ The equilibrium conditions of the model are used in conjunction with the data to measure the wedges. Even though the model has little structure in the sense that heterogeneity, frictions, and shocks are not modeled explicitly, the measurement of the wedges will in general depend on the chosen functional forms for utility and production functions.⁹

1.3.1 Households

The islands (indexed by k) in the economy are inhabited by infinitely lived households who have standard preferences given by

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} e^{-\rho_{kt}} \left(\frac{C_{kt}^{1-\sigma}}{1-\sigma} - \frac{\phi}{1+\phi} N_{kt}^{\frac{1+\phi}{\phi}} \right) \right]$$

where C_{kt} is consumption of the final good, N_{kt} is labor, and ρ_{kt} is the discount rate. This preferences allow for wealth effects on labor supply. In section 1.8 I explore the consequences of moving towards GHH preferences and not allowing for wealth effects.

Households choose how much to consume of the final good, how much to work and how

8. Wage and price differences across regions are also required for the implementation of the methodology as I will write the model for the regional economies in deviations from the aggregate. If there is no differences in prices and wages across regions, these variables would vanish from the equilibrium equations.

9. This is always the case, even in the more standard context of Chari et al. (2007a), both parameter values and functional forms will matter.

much to save while facing the following period-by-period budget constraint

$$P_{kt}C_{kt} + B_{kt+1} \leq e^{\delta_{kt}^\omega} B_{kt}(1 + i_t) + e^{-\epsilon_{kt}^\omega} W_{kt}N_{kt} + \Pi_{kt} + T_t$$

with households spending their financial income $B_{kt}i_t$, labor income $W_{kt}N_{kt}$, profits from the firms Π_{kt} , and transfers from the government T_t on consumption and updating their nominal bond holdings at the beginning of period t , B_{kt} . There are two time-varying wedges in the budget constraint, one resembles is a tax on labor income (ϵ_{kt}^ω) and the other one resembles a tax on savings in bonds (δ_{kt}^ω). Given that this is a monetary union with bonds that trade freely, i_t , the nominal interest rate, is equalized across islands. It is worth noticing that this economy is isomorphic to one in which the wedge on savings is modeled as a wedge on the discount factor in preferences and the wedge on labor income is modeled as a wedge multiplying labor in preferences. I put them in the budget constrain as it is the standard approach in the business cycle accounting literature. For the model to be internally consistent, I assume that the government collects taxes across islands and rebates them back to consumers in the islands in equal amounts, T_t . It is possible to modify the model and add a more complex system of transfers from the government like in Beraja (2015). With existing data from the Bureau of Economic Analysis on fiscal transfers at the state level one can apply the methodology to a context of a fiscal union. At the MSA and county level it may not be possible depending on data availability.

An important, technical condition is that to induce stationarity the discount factor includes an endogenous component that satisfies $\rho_{kt+1} = \rho_{kt} + \Phi(\cdot)$ for some function $\Phi(\cdot)$ of the average per-capita variables of an island.¹⁰

10. This economy features market incompleteness since the residents of the islands only have access to a nominal bond, whose rate of return is determined by the monetary authority exogenously (from the island's perspective) as it depends on aggregate variables that an island cannot influence since the number of islands is large. This induces non-stationary distributions for variables in the islands, which makes it problematic to study business cycle dynamics in the context of small open economies log-linearized around a deterministic steady state. One way to solve this problem and induce stationarity is to include an endogenous component in the agent's discount factor. Following Beraja et al. (2016) I use the following function, $\Phi(\cdot) = \Phi_0(c_{kt} - c_t)$

1.3.2 Firms

In each island profit maximizing firms operate in two different, competitive sectors. There is an intermediate-good sector, which I denote with the superscript x that produces a tradable input to be used in the production of the final good, which I denote with the superscript y . Every period, producers in the intermediate sector choose the amount of local labor N_{kt}^x , pay nominal wages W_{kt} , and receive the price Q for their product. I assume that trade is free and therefore the law of one price holds. These producers face the following problem:

$$\max_{N_{kt}^x} e^{z_{kt}^{x\omega}} Q_t (N_{kt}^x)^\theta - W_{kt} N_{kt}^x$$

where $z_{kt}^{x\omega}$ is a time varying tax (wedge) on the sales of the intermediate good¹¹ and $\theta < 1$ is the labor share in the production of tradables. The final good is produced in each island using as inputs local labor N_{kt}^y and the tradable intermediate good. The final good is non-tradable and has price P_{kt} . The final good producer's problem is:

$$\max_{N_{kt}^y, X_{kt}} P_{kt} e^{z_{kt}^{y\omega}} (N_{kt}^y)^\alpha (X_{kt})^\beta - W_{kt} N_{kt}^y - Q_t X_{kt}$$

where $z_{kt}^{y\omega}$ is a time varying tax (wedge) on the sales of the final good and $(\alpha, \beta) : \alpha + \beta < 1$ are the labor and intermediates shares of output, and X_{kt} is the amount of intermediate goods used as inputs in island k . The assumption of decreasing returns to scale results in endogenous variation in markups and profits will not be zero in equilibrium.

The assumption that the final good is non-tradable results in price dispersion across islands. Additionally, I assume that there is labor mobility across sectors¹² but not across

using the notation of section 1.6. This solution follows Schmitt-Grohé and Uribe (2003) who show that the endogenous discount factor yields identical dynamics at business cycle frequencies to those produced by other alternative solutions.

11. Notice that this wedge looks like time varying productivity $z_{kt}^{x\omega}$. In the business cycle literature it has been called efficiency wedge. More structural interpretations would refer to it as productivity shocks

12. Notice that one does not need to assume that labor is mobile across sectors. However, each layer of

islands, which results in wages being equal across sectors within an island and in wage dispersion across islands¹³. These features of the model capture the fact that prices and wages varied differentially across regions over the business cycle as illustrated in Figures [Figures for prices, wages, and appendix]. Even though there is migration across regions, for the purposes of this paper the assumption of no labor mobility across islands seems reasonable, as such mobility does not seem to be the relevant margin of adjustment. Using migration data from the IRS¹⁴ I compute net migration rates. Figure 1.4 shows that both consumption growth and employment growth are uncorrelated with net migration rates at business cycle frequencies. In the Great Recession, the housing market was affected differentially across regions and house prices fell across the country; this should make it even harder for people to move, especially in the short run.

1.3.3 Market Clearing Conditions, Government, and Equilibrium

The remaining equations required are the market clearing conditions of the final goods market, the labor market, the intermediates good market, and the bond markets, which respectively are given by:

$$C_{kt} = e^{z_{kt}^{y\omega}} (N_{kt}^y)^\alpha X_{kt}^\beta$$

$$N_{kt} = N_{kt}^y + N_{kt}^x$$

complexity increases the data demands. In the case of labor mobility across sectors one would need to build local wage indexes by classifying industries as tradable and non-tradable, or use other measures of wages that require less data, especially for counties and cities

13. There are alternatives in the literature that allows one to produce geographic wage dispersion without assuming complete labor mobility. An example is provided by Suárez Serrato and Zidar (2016).

14. The data is available here <https://www.irs.gov/uac/soi-tax-stats-migration-data>

$$\sum_k X_{kt} = \sum_k e^{z_{kt}^{x\omega}} (N_{kt}^x)^\theta$$

$$\sum_k B_{kt} = 0$$

Also, the amount of taxes rebated to consumers is given by

$$T_t = \frac{1}{K} \sum \left[(e^{\delta_{kt}^\omega} - 1) B_{kt} (1 + i_t) + (1 - e^{-\epsilon_{kt}^\omega}) W_{kt} N_{kt} \right]$$

Finally, the monetary authority sets an interest rate according to a rule $i_t = i_t(\cdot) e^{\mu_t}$, which will depend only on aggregate variables where μ_t could be a wedge between the interest rate rule and the interest rate.¹⁵ In this economy, an equilibrium is a collection of prices $\{P_{kt}, W_{kt}, Q_t\}$ and quantities $\{C_{kt}, N_{kt}, B_{kt}, N_{kt}^x, N_{kt}^y, X_{kt}\}$ for each island k and time t that are consistent with household utility maximization and firm profit maximization, such that the market clearing conditions above hold, given an interest rate set by the rule and exogenous processes $\{\delta_{kt}^\omega, z_{kt}^{y\omega}, z_{kt}^{z\omega}, \epsilon_{kt}^\omega, \mu_t\}$.

1.4 A friction and the corresponding wedge

The model above has no explicit frictions. In this section I show how a friction in an explicit model creates a wedge in one of the first order conditions of the model and how shocks and frictions in an explicit model can be presented as a frictionless model with a time-varying wedge. For example, Beraja et al. (2016), write a model that has a shock to labor supply in the utility function and a shock to the discount rate to capture types of shocks that have been commonly cited in the literature as important forces during the Great Recession. With these shocks, the household problem would involve maximizing

15. The wedge in the interest rate will show up in the Euler equation alongside the investment wedge. Therefore, there will not be two separate wedges, I will recover just one that it is the sum of the two.

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} e^{-\rho_{kt} - \delta_{kt}} \left(\frac{C_{kt}^{1-\sigma}}{1-\sigma} - e^{\epsilon_{kt}} \frac{\phi}{1+\phi} N_{kt}^{\frac{1+\phi}{\phi}} \right) \right]$$

subject to

$$P_{kt}C_{kt} + B_{kt+1} \leq B_{kt}(1 + i_t) + W_{kt}N_{kt} + \Pi_{kt} + T_{kt}$$

where δ_{kt} and ϵ_{kt} are shocks to the household's discount factor and the disutility of labor, respectively, and all the other variables are the same as in the model of section 1.3. Now let us introduce a friction that will introduce a wedge in this model. A friction commonly used in the literature are sticky wages. A simple way to illustrate how sticky wages can become a wedge is by using a wage setting rule in which the wage today depends on a target wage and the past wage. If we assume that the target wage comes from the first order conditions of the households, the wedge immediately appears. Suppose that wages are set according to the following rule:

$$W_{kt} = (MRS(C_{kt}, N_{kt}))^\lambda (W_{kt-1})^{1-\lambda}$$

where $MRS(C_{kt}, N_{kt})$ is the marginal rate of substitution between consumption and leisure and λ represents the degree of wage stickiness. In particular, when $\lambda = 1$ wages are fully flexible and correspond to the efficient allocation, and when $\lambda = 0$ wages are completely rigid. In general workers will be off their labor supply curves for $\lambda < 1$. In the case of the specific model with shocks:¹⁶

$$W_{kt} = (P_{kt}e^{\epsilon_{kt}}(N_{kt})^{\frac{1}{\phi}}C_{kt}^\sigma)^\lambda (W_{kt-1})^{1-\lambda}$$

In log-linearized form, this could be written as

16. In the case of GHH preferences $W_{kt} = (P_{kt}e^{\epsilon_{kt}}(N_{kt})^{\frac{1}{\phi}})^\lambda (W_{kt-1})^{1-\lambda}$

$$w_{kt}^r = \left(\frac{1}{\phi} n_{kt} + \sigma c_{kt} \right) + (1 - \lambda) \left(w_{kt-1}^r - \pi_{kt} - \frac{1}{\phi} n_{kt} - \sigma c_{kt} \right) + \lambda \epsilon_{kt} \quad (1.1)$$

The first order condition of the model with wedges is $e^{-\epsilon_{kt}^\omega} W_{kt} = P_{kt} N_{kt}^{\frac{1}{\phi}} C_{kt}^\sigma$ and in log-linearized form:

$$-\epsilon_{kt}^\omega + w_{kt}^r = \frac{1}{\phi} n_{kt} + \sigma c_{kt} \quad (1.2)$$

The right hand side of (1.2) is equal to the right hand side of (1.1) when wages are perfectly flexible ($\lambda = 1$) and there are no shocks ($\epsilon_{kt} = 0$). Sticky wages thus create a wedge in the first order condition of the model and this wedge is equal to:

$$wedge^{consumer} \equiv -\epsilon_{kt}^\omega = (1 - \lambda) \left(w_{kt-1}^r - \pi_{kt} - \frac{1}{\phi} n_{kt} - \sigma c_{kt} \right) + \lambda \epsilon_{kt}$$

In this sense, wedges can be thought of as reduced form for frictions (sticky wages in this example) and shocks (labor supply shocks in the explicit model).

1.5 Data

1.5.1 Regional Data

One of the main challenges macroeconomic researchers face when studying sub-national economies is the lack of readily available data for all the aggregate series that a general equilibrium framework requires. For the United States, a country for which the availability of data is outstanding in comparison with most other countries, the availability of aggregate economic time series for different levels of regional aggregation is still precarious. For example, state-level time series on consumption were inexistent until August, 2014, when the Bureau of Economic Analysis (BEA) released prototype statistics on Personal Consumption Expenditures by State. These series, however, are not independent measures of consumption,

as some of the data is imputed from employment figures. Another example are price indexes; the Bureau of Labor Statistics (BLS) does not produce price indexes at the state or county levels and, even though it does produce price indexes for 27 MSAs, this number of MSAs clearly does not help increase power in studies using cross-sectional regional variation. Other time series that are used in macro models such as investment and capital are not available from official sources at the local level.¹⁷

The model above is designed taking into account these data limitations. The result is a general equilibrium model that we can estimate for different levels of regional aggregation. The cost is that the model may be simpler than what one could have for a national economy. To be able to estimate some of the parameters of the model, recover the stochastic processes and build the wedges I need data on prices, wages, employment, and net assets (see section 1.6) at the regional level. In this paper I construct all the required time series at the state, MSA and county levels. I am able to construct the endogenous state variables (wages and net assets) from 2005 to 2014 and the endogenous control variables (prices and employment) from 2006 through 2014 and thus we can study some of the differences in regional business cycles during the Great Recession and its aftermath.

1.5.1.1 Wages

In order to get wage measures at the local level I build wage indexes using a similar approach to Beraja, Hurst, and Ospina (2016). The goal of the wage indexes is to have measures of wages that 1) are as comparable as possible across regions, and 2) do not vary across regions and over time due to differences and changes in the composition of the labor force, which may occur as a result of both long term trends and business cycle fluctuations.

To accomplish this I employ the micro data of the *American Community Survey* (ACS)

17. Yamarik (2013) produced state-level measures of investment and capital from 1990 through 2007, using NAICS one-digit industry data to divide up the national capital stock based on the relative income generated within each state

available at the Minnesota Population Center (IPUMS-USA, Ruggles et al. (2015))¹⁸. Between 2005 and 2014, the ACS includes information of about 3 million people per year on average. For each year I calculate hourly nominal wages for working-age workers, both men and women, with a strong attachment of the labor force. More specifically, I restrict the sample to people between the ages of 16 and 64, who do not live in group quarters, who reported earning at least 5,000 dollars the prior year, who were employed at the time of the survey, who reported working usually at least 30 hours a week, and who worked for at least 48 weeks during the prior 12 months. Then for each person in the sample I calculate an hourly wage by dividing annual labor income by a measure of annual hours worked.¹⁹

In order to have a more uniform wage measure across geographies and better capture business cycle variation in nominal wages not coming from the composition of the workforce, which could vary across locations and over time, I adjust the wages by creating a measure that excludes the effect of observable characteristics. Specifically, I run the following cross-sectional regression:

$$\ln w_{itk} = \gamma_t + \Gamma_t X_{it} + \eta_{itk} \quad (1.3)$$

where $\ln w_{itk}$ is the log wage of person i in year t who resides in location k (State, MSA or county) and X_{it} is a vector of person/household-specific controls. To be precise, the vector of controls includes a dummy variable for sex (with female being the omitted group), a set of three dummy variables for hours usually worked (with “40-49 hours per week” being the omitted group), a series of nine age dummies (with “40-44” being the omitted group), a set of four dummies for educational attainment (with “some college” being the omitted group), three citizenship dummies (with “native born” being the omitted group), and a race dummy

18. The data is available here <https://usa.ipums.org/usa/>

19. Total labor income during the prior 12 months includes both wage and salary earnings and business earnings. Total hours worked during the previous 12 months is the product of the total number of weeks that the respondent worked during the prior 12 months and the respondent’s hours usually worked per week

(with "white" being the omitted group). I run these regressions for each year separately to allow for the possibility that the constant (γ_t) and the vector of controls (Γ_{it}) vary over time. Having controlled for the set of observables, for each individual I compute $w_{itk}^{adj} = e^{\gamma_t + \eta_{itk}}$ as the adjusted wage. Adding the constant to the regression allows the adjusted wage measures to reflect differences in average log-wages over time. Finally, to compute the adjusted wage index for a given location k I compute weighted averages of w_{itk}^{adj} across individuals in region k .

With this procedure, I am able to compute wage indexes for 48 states, 234 MSAs and 379 counties. For MSAs and counties these measures of adjusted wages will be noisier. Figures 1.6, 1.33, 1.34, and on-line appendix figures A6, A7, A8 present the wage indexes and show first that there was wage dispersion across regions over the business cycles, and that such dispersion was strongly correlated with local economic conditions. Moreover, appendix Figure 1.37 shows that the cross-sectional patterns documented by Beraja et al. (2016) for states, while Appendix Figures 1.38 and 1.39 show that the same strong relationship between real wage growth and local economic conditions is present for MSAs and counties²⁰. Even though this does not fully address the concern of measurement error, it does show that the wages I measure at a more granular level do capture some of the business cycle characteristics of the Great Recession.

Another possibility to address the issue of measurement error would be to use data from the County Business Patterns to measure local wages. In this case, however, the wages could not be adjusted to account for labor force compositional differences.

1.5.1.2 Employment

I construct measures of labor at the regional level that capture both the extensive and intensive margin of the employment decision. For a given region, the labor measure is total

20. The relation is even stronger when measured with respect to the unemployment rate

annual hours worked per person (N) computed as the product of the employment (E) to population (P) ratio and average annual hours worked per worker (h).²¹ The employment (E) measure comes from the Bureau of Labor Statistics (BLS) Local Area Unemployment Statistics (LAUS) Program, from which I take average annual employment figures (number of employees) for states, counties, and MSAs.

At the state level the the population variable (P) is the non-institutionalized civilian population²² aged 16-64, which can be obtained also from the BLS' LAUS program.²³ For MSAs and counties I rely on the population estimates program for the US Census Bureau and obtain measures of the resident population²⁴ at the local level.²⁵

Finally, for hours worked (h) I rely on different datasets depending on the level of aggregation. At the state level, the BLS produces Average Weekly Hours of All Employees in the Private Sector²⁶ from the Current Employment Statistics (CES) database *Employment, Hours, and Earnings - State and Metro Area*.

These series are available from 2007 through 2014. In order to obtain a consistent measure of hours for 2006, I use an auxiliary regression based on hours worked information from the American Community Survey (ACS). With the ACS information I calculate weighted

21. This is standard in the literature, see for example Shimer (2009) or Ohanian and Raffo (2012)

22. See definitions here http://www.bls.gov/cps/eetech_methods.pdf

23. The non-institutionalized civilian population aged 16 and over can be directly found at <http://www.bls.gov/lau/staadata.txt>. To obtain the population only up 64 years of age, one can subtract the 65 years and older population found at <http://www.bls.gov/lau/#tables>. It is possible also to access the files directly; for example for 2006 the table can be directly accessed at <http://www.bls.gov/lau/table14full106.xlsx> and for 2011 at <http://www.bls.gov/lau/table14full111.xlsx>.

24. See definitions here <https://www.census.gov/popest/about/terms.html>

25. The data for the 2000-2010 period is available here <https://www.census.gov/popest/data/intercensal/county/files/CO-EST00INT-AGESEX-5YR.csv> whereas the data for the 2010-2014 period can be found in the Census Fact Finder. I obtain the population for MSAs by adding up the population of the counties that make up each MSA since the data for MSA is not readily available for the 2000-2009 period. As a check, for the 2010-2013 period, adding up the county level population to the MSA level yields exactly the same answer as getting the MSA data directly.

26. These series can be obtained through the identifiers SMUXYYYYYY0500000002, where XX is the state fips code and YYYYYY is the metro area code

averages of weekly hours worked at the local level and I estimate the following regression pooling together all years from 2007 to 2014:

$$h_{BLS,t}^k = \alpha_k + \beta h_{ACS,t}^k + \gamma u_t^k + \varepsilon_t^k$$

where $h_{BLS,t}^k$ are the hours from the BLS in region k in year t , α_k is a region k fixed effect, $h_{ACS,t}^k$ is the weighted average hours worked in year t in region k and u_t^k is the unemployment rate²⁷ in year in region k in year t . For the hours of ACS, I restrict the calculations to people with ages 16 to 64 and who worked at least 1 week in the previous 12 months. The regression yields an R-square of 85%. I use the regression to predict the 2006 hours.

At the MSA and county levels I use hours usually worked from the ACS computed as a weighted average for people aged 16 to 64, who are in the labor force, and worked for at least 1 week in the previous 12 months.

1.5.1.3 Net Assets

Net asset positions at the local level are not commonly used in the literature. In one of their seminal papers on the Great Recession, Mian et al. (2013) build zipcode, county and MSA level net asset positions for 2006 as part of their measure of what they call the housing net worth shock that hit the economy between 2006 and 2009. Also, Beraja (2015) constructs a similar variable for the period 2006-2011 at the state level by aggregating MSA-level net worth data from Mian et al. (2013) and iterating forward the law of motion of assets using national accounting identities. Ideally one would like to use the data and methodology by Saez and Zucman (2016) to make a better assignment of wealth, but limited access to the internal IRS SOI data precludes this possibility. As a result, in this paper I follow closely Mian et al. (2013) to create a time series of net assets between 2005 and 2014.

27. This variable is also available for all regions in the Local Unemployment Statistics Database

For the households living in a given region i we can define net assets as $B_t^i = Stocks_t^i + Bonds_t^i + Housing_t^i - Debt_t^i$. The four terms refer to market values and I abstract from human capital in this definition. To compute the market value of stocks and bonds I use IRS Statistics of Income (SOI)²⁸ at the county level (which also contains state level information and can be aggregated to obtain MSA level figures). The IRS reports dividends and income received by households in a fiscal year by county. We can use these figures to assign the fraction of financial assets from the Federal Reserve Flow of Funds data to each region. Under the assumption that the representative household in a region holds the market index for stocks and bonds, the share of total dividends and total interest going to a region gives the fraction of financial assets held by the region.²⁹

To estimate the value of the housing stock I use data from the Census, the American Community Survey (ACS), and the Population Estimates to build a measure of the housing stock owned by households using estimates of the number of housing units, homeownership rates, vacancy rates, and the median house value. Depending on the geographic unit of analysis I use data on house market prices from the Federal Housing Finance Agency, Zillow and/or Corelogic, to track changes in the value of the housing stock. Appendix 1.11.4 provides a precise description of the data, sources and computations. Finally, I measure debt using data from Equifax Predictive Services at the zip code level and aggregating it up to each local level.³⁰ The Federal Reserve Bank of New York also keeps measures of debt levels at the state and county levels.³¹

Figure 1.8 and on-line appendix figure A15 show that there is dispersion also in changes

28. The data can be downloaded directly from <https://www.irs.gov/uac/SOI-Tax-Stats-County-Data>

29. The Flow of Funds Data can be downloaded from <http://www.federalreserve.gov/datadownload/>. The series for financial assets I use is FL154090005.A

30. The Equifax and Corelogic data were kindly provided by the Fama-Miller Center, <http://research.chicagobooth.edu/famamiller/data>

31. The data at the state level can be downloaded from <https://www.newyorkfed.org/microeconomics/data.html>. The data at the county level only goes through 2011 and can be requested from the Center for Microeconomic Data at the New York Fed.

in net assets at the local level at business cycle frequencies. The fact that those changes are correlated with employment growth provide some assurance that the measures are meaningful.

1.5.1.4 Prices

I construct regional price indexes using Nielsen's Retail Scanner Database (RMS). To this purpose I follow the methodology first outlined in Beraja et al. (2016), who built indexes for 48 states (Alaska and Hawaii are not part of Nielsen's data) for the period 2006-2012. I extend the data through 2014 and build retail price indexes at state, MSA and county levels. The RMS data is collected by the Nielsen Company and it is available at The University of Chicago Booth School of Business³². This scanner data contains primarily weekly prices and sales volume for products sold at about 36 thousand stores around the United States belonging to about 85 store chains. It is collected at the point of sale through barcode scanning systems, which reduces data collection errors, and then reported by the stores to Nielsen. Also, once a participating retailer has agreed to participate, all the stores in the retail chain are included, reducing potential selection issues.³³

This data is massive, as of 2014 it contains more than 119 billion single observations, but its structure is relatively simple. The unit of observation is a (product, store) combination. Products are uniquely identified by a 12-digit number called Universal Product Code (UPC) that is not specific to the database. Stores are uniquely identified by a database specific code but the identity of the store is not available as part of the agreement with the participating stores. An observation consists of the number of units of a UPC sold over a given week

32. The data is made available through the Marketing Data Center at the University of Chicago Booth School of Business. Information on availability and access to the data can be found at <http://research.chicagobooth.edu/nielsen/>

33. According to the data documentation only in rare instances, a retailer may consider a small number of their stores as confidential and exclude them from the dataset. Also, it is worth mentioning that not all retailers in the US are part of Nielsen's database and not all Nielsen's contributing retailers have agreed to share their data.

in a store, the corresponding quantity-weighted average price, and a variable that allows to compute unit prices in case the products were part of a promotion. The database only includes items with strictly positive sales in a store-week and excludes certain products such as random-weight meat, fruits, and vegetables since they do not have a UPC assigned.

In the database there is also information about store environment information, the characteristics of the products, and more importantly for the purpose of this paper about the location of the stores. The geographic coverage of the RMS database is unparalleled and stable over time, which allows to build indexes for many locations for the entire period. The database contains information for 48 states, 361 Metropolitan statistical areas, 895 Metro and Micropolitan Statical Areas, and more than 2,350 counties (about 75% of all counties in the US) for each year in the period 2006-2014. Table 1.7 in the appendix presents some figures that summarize data and give a sense of its size and breadth.

In order to build the price indexes proceed in three stages. In the first stage I aggregate the data to go from a weekly to monthly frequency. To do this I for each (UPC,Store,Week)-Month combination I add the quantities over weeks to obtain a total quantity sold in the month, and then compute quantity-weighted unit prices for the month. With the transformed data I then follow a a two-stage procedure to construct the price indexes. The procedure is similar to the one employed by the BLS in constructing the chained CPI (C-CPI) explained in Cage et al. (2003)³⁴. Nielsen groups the UPCs into approximately 1,100 “categories” (or modules) that are relatively narrowly defined (e.g. for canned fruit, the categories are apples, berries, figs, grapes, cherries, plums, oranges among others).

The first stage of the procedure consists of building an index at the category level for a given geographic unit. I define the good for the price index as a (UPC,Store) combination. By taking each pair (UPC,Store) as a different good entering the price index I allow the index to capture unmeasured quality differences across stores and allow the index to capture the

34. A brief description of the two stage procedure in the C-CPI can be found here <http://www.bls.gov/cpi/cpisupqa.htm>

store-switching behavior documented by Coibion et al. (2015), who show that an important feature over the business cycle is that households substitute from high cost to low cost stores and as a result prices may change. In order to aggregate the individual monthly prices into a category-level price index I use the following chained Laspeyres price index formula:

$$P_{j,t}^{L,r} = P_{j,t-1}^{L,r} \times \frac{\sum_{i \in (j,r)} p_{i,t} q_{i,y}^-}{\sum_{i \in (j,r)} p_{i,t-1} q_{i,y}^-} \quad (1.4)$$

with $\{t, t-1\} \in y+1$ for all months except for $t = \text{January}$ and with $t \in y+1$ and $t-1 \in y$ when $t = \text{January}$, where $P_{j,t}^{L,r}$ is the time t chained Laspeyres index for category j in region r , $p_{i,t}$ is the price at time t of good i , $q_{i,y}^-$ is the average monthly quantity sold of good i in the base year y . Notice that the basket at this level of aggregation is fixed for a year. This has the effect of shutting down the possibility of substitution within categories but solves other technical problems in the construction of price indexes with scanner data (see Dube et al (2016)). The basket is, however, updated every year.

In the second stage I aggregate the regional category-level price indexes into regional indexes. Like the BLS, in this stage I use the Tornqvist aggregation formula, an index formula that as part of the family of superlative index formulas better approximates the cost of living and better captures substitution. The inputs are the category-level prices and the total expenditures of each category. For each region I compute:

$$\frac{P_t^r}{P_{t-1}^r} = \prod_{j=1}^N \left(\frac{P_{j,t}^{L,r}}{P_{j,t-1}^{L,r}} \right)^{\frac{S_{j,r}^t + S_{j,r}^{t-1}}{2}} \quad (1.5)$$

where $S_{j,r}^t$ is the share of expenditure of category j in month t in region r . I calculate the shares using total expenditure in each category each month, even though for the category-level indexes (from the first stage) some goods were not included due to missing data. In other words, the weight of a category in the final index is not affected by the fact that the indexes in the first stage imposed some restrictions on what (UPC,Store) combinations

entered the basket. Notice this chained index updates the shares every month, accounting for substitution across categories.

With this methodology and this data I am able to construct indexes for 48 states, 895 Metro and Micropolitan Statistical Areas, and 2,266 counties. Figures 1.7, 1.35, 1.36, and on-line appendix figures A9, A10, A11 present the price index and show first that there was price dispersion across regions over the business cycles, and that such dispersion was slightly, but positively correlated with employment growth.

1.5.1.5 Consumption

Data on consumption is only available at the state level through the Bureau of Economic Analysis (BEA). This data was released in late 2014 and corresponds to the Personal Consumption Expenditures data. The definitions in the data at the state level are identical to those of the Aggregate US Economy, and the data is produced in a way that the state level data adds up to the national level data. Following Karabarbounis (2014b) consumption is defined as the sum of real expenditures on non-durables and real expenditures on services. These two categories are clearly defined in the BEA data. Figure 1.5 and on-line appendix figure A5 show the correlation between employment and consumption at the state level. The positive correlation is expected and is reassuring of the quality of the data.

1.5.2 Aggregate Data

Unlike regional data, aggregate data is relatively simple to obtain and goes back in time several years. I employ aggregate data from 1975 through 2014. In order to estimate the aggregate exogenous process of the model I only need three aggregate time series: wages (w), a measure of labor (n), and prices (p). In general, to the possible extent the aggregate series should be consistent with the regional series. As explained in section 1.6, I will write the regional model in deviations from the aggregate, a transformation that makes the model

for the regions independent from each other and simplifies the estimation procedure. The data needs to be transformed in the same way and thus the requirement on the consistency of the data. This will impose some restrictions on the aggregate data that I can use.

Let us begin by describing the data that I use for wages. Wages provide a good example of the consistency that one wants to have between the regional and the aggregate data. The American Community Survey used to create regional wage indexes only goes back to 2001. In order to increase the number of observations for estimation purposes, I use the March Supplement of the *Current Population Survey (CPS)* from IPUMS³⁵, which allows for wage index construction going back to 1975. I construct the indexes by performing the same procedure as for the regional wages, using exactly the same econometric specification of equation (1.3). The controls and restrictions on the data (including age, individuals not living in group quarters and strong attachment to the labor force) are also the same as for the regional wages in order to make the aggregate and regional wages as consistent as possible.

The construction of these indexes differs from the one that uses the ACS in three main aspects. (1) The timing of the information. While the ACS surveys people every month about their earnings over the previous 12 months, the CPS asks specifically for the earnings over the previous calendar year. Therefore, to adjust for this timing issue, I lag the wage indexes by one year with respect to the survey year; the CPS data available from 1976 through 2015 can be used to produce wage indexes from 1975 through 2014. (2) Citizenship controls. The regression for the CPS does not include citizenship dummies in the regression since citizenship is not consistently measured over the sample period.³⁶(3) The estimation is not made year by year but by pooling years together in three time periods, 1975-1995, 1996-2000, and 2000-2014. There are a couple of reasons for these choices. First, to increase power particularly in the early years of the sample. Second, to adjust for changes in the

35. The CPS data (Flood et al. (2015)) is available at <https://cps.ipums.org/cps/>

36. A specification using citizen dummies produces very similar indexes. However, I use the index that does not control for citizenship to avoid introducing measurement noise

survey: in 1995 the CPS earnings questions used to compute wages changed, and in 2000 the sample was expanded.³⁷ Beraja et al. (2016) show that the index constructed in this way using the CPS, have a correlation with the aggregate index constructed using the ACS and the 2000 Census of 0.99 for the period 2000-2012, which gives the required reassurance about the comparability between the regional and the aggregate information.

The aggregate measure for labor, in agreement with the regional measures, captures both the intensive and the extensive margin. The intensive margin is captured by a measure of hours worked per person (h). I use the *average weekly hours of all employees in the private sector*³⁸ from the database *Employment, Hours, and Earnings - National* of the BLS' Current Employment Statistics (CES) program. I use this series because it is the exact same counterpart at the regional level for states. Since this series only starts in 2006, in order to use as much data as possible for the estimation of the model I will extend it going back to 1975 using as index the variable *average weekly hours of production and non-supervisory employees in the private sector*.³⁹ I use this variable as index and not other alternatives (for example measures from the Major Sector Productivity Program⁴⁰) because its changes after 2006 track almost exactly the changes in the chosen series.

This data is similar but not exactly the same as what is used in other studies. The alternatives are presented in Karabarounis (2014b) who uses four different measures of aggregate hours worked in his study of the labor wedge. The first measure, the closest to the one used in this paper, is an index of aggregate weekly hours of production and non-supervisory employees⁴¹ available also in the CES database of the BLS. I do not use this measure that

37. Not pooling results in an index that has a correlation with the one I use of 89% in log changes and 99% in levels

38. Specifically, the variable I use has series ID CEU0500000002

39. The series ID of this variable is CEU0500000007 and goes back to 1964. This is the per employee per week counterpart of the series CES0500000034 used by Karabarounis (2014b).

40. Series PRS84006023 and PRS85006023 provide measures of average weekly hours worked for the business and non-farm business sector

41. The reference of this BLS Series is CES0500000034

already incorporates the extensive margin because it is not available at the local level. A second measure comes from Cociuba et al. (2012) who use the Current Population Survey to estimate total hours worked in the United States and hours worked per person aged 16 to 64 and add the hours worked by military personnel. For the purposes of this study the inclusion of military personnel at the aggregate level would call for the same adjustment at the local level. Even if this adjustment is possible, the local differences in military population will create differences across regions that are not coming from business cycle fluctuations, which we want to abstract from. The other two measures use the BLS' Major Sector Productivity Program⁴², which produces productivity measures for the U.S. business, non-farm business, and manufacturing sectors. However these sector-based variables are not available at the local level, and therefore I do not use them at the aggregate level.⁴³

In order to capture the extensive margin, I complement the hours per worker (h) with employment figures (E) produced by the BLS based on the Current Population Survey, and with data on the civilian non-institutionalized population (P) aged 16 and over upon which the employment statistics are based.⁴⁴ The data for non-institutionalized civilian population aged 16 to 64 can be obtained by subtracting the 65 and over population also from the BLS.⁴⁵⁴⁶

42. The exact name of the database is called Major Sector Productivity and Costs

43. One of these measures is an index of total hours worked in the business sector (BLS Series ID is PRS84006033), and the other one is a per-worker measure used by Ohanian and Raffo (2012).

44. The data is available here <http://www.bls.gov/cps/cpsaat01.htm>,

45. To get access to the two series go to <http://data.bls.gov/cgi-bin/srgate> and obtain the civilian non-institutionalized population 65 and over through the Series ID LNU00000097 and the 16 and over population through the Series ID LNU00000000. For resident population the data can be obtained here <https://www.census.gov/popest/data/historical/index.html> by adding up state-level population from the US Census Population Estimates Program. National data on resident population by age are not available for the 1980s, and for years prior to 1980 it includes military personnel, thus the need to get it from State-level data

46. Other specifications of the variables are possible. For example Beraja et al. (2016) build wage indexes only for prime age (25-56) males, and consequently they use the employment population ratio of this group. For different sex-age specifications the employment population ratios are available from the BLS and the OECD. From the BLS, series LNS12300000 and LNS12300001 directly provide the employment rate for all people and men aged 16 and over respectively. The St. Louis Federal Reserve Bank provides the OECD

Finally, for prices I use the Consumer Price Index (CPI) from the *All Urban Consumers Database* of the BLS.⁴⁷ Ideally one would like to use the chained version of the CPI but it is only available from 1999, so we use the fixed base index. Since the price indexes at the local level have a large food component, which is the result of having built them using scanner data mostly from supermarkets, the food component of the CPI is the one that is better tracked by the aggregate index that results from the aggregation of the regional price indexes. The CPI and the CPI food have tracked each other almost identically since 1975 exhibiting almost the same long run inflation. However during the Great Recession food prices had larger increases than other prices (for example housing-related prices). As a robustness check, in order to minimize differences between the data I use for the aggregate and the data I use for the regions, I also estimate the model using the CPI up to 2006 and complementing it with the food component⁴⁸ from 2006 onwards. The results are almost identical.

1.6 Model Estimation

The goal of this section is to explain how I estimate the parameters required to recover the wedges and perform the business cycle accounting exercise. For the the reader familiar with the business cycle accounting literature, I follow as guideline the procedure outlined in the seminal paper by Chari, Kehoe, and McGrattan (2007a) and adapt it to the class of models to which the model of section 1.3 belongs. As in Chari et al. (2007a) I assume that the wedges follow an AR(1) process in which the are allowed to be correlated contemporaneously an

series for other population groups (for example 15-64 and 25-56 by sex). For example, for prime age workers the series has ID LREM25TTUSM156N and is available here <https://research.stlouisfed.org/fred2/series/LREM25TTUSM156N>

47. Specifically I use the variable with BLS ID CUUR0000SA0.

48. This variable has BLS ID SUUR0000SAF.

over time.⁴⁹ Further, I allow for the wedges to have an aggregate component and a local (island-specific) component, and assume that the island specific component sums to zero in all periods:

$$[\gamma_{kt}^\omega, z_{kt}^{y\omega}, z_{kt}^{x\omega}, \epsilon_{kt}^\omega]' = P_0 + P [\gamma_{t-1}^\omega, z_{t-1}^{y\omega}, z_{t-1}^{x\omega}, \epsilon_{t-1}^\omega]' + B [u_t^\gamma, u_t^y, u_t^x, u_t^\epsilon]' + \tilde{B} [v_{kt}^\gamma, v_{kt}^y, v_{kt}^x, v_{kt}^\epsilon]' \quad (1.6)$$

with $\sum_k \frac{1}{K} v_{kt}^y = \sum_k \frac{1}{K} v_{kt}^x = \sum_k \frac{1}{K} v_{kt}^\gamma = \sum_k \frac{1}{K} v_{kt}^\epsilon = 0$.

The first step is to log-linearize the economy. I do so around a constant price and wage inflation steady state. In terms of notation, variables in lower case represent log-deviations from the steady state. So, for example, $n_{kt} \equiv \log(N_{kt}/N_k)$ is the log-deviation of employment at time t in island k from its steady state value. I denote aggregate variables by removing the subscript k , and they are the result of taking averages across islands. For example, aggregate employment is given by $n_t \equiv \sum_k \frac{1}{K} n_{kt}$.

Under some assumptions this log-linearized economy has two properties that will prove useful to facilitate estimation. Below, in section 1.6.1, I discuss the consequences of moving away from these assumptions and how one can go about doing the estimation. The first assumption is that the islands are identical in steady state. The second assumption is that the model parameters in the utility and production functions as well as the parameters of the process driving the wedges are the same across islands. Finally the third assumption is that the joint distribution of island-specific shocks that drive the stochastic processes is such that its cross-sectional sum is zero. From these assumptions the log-linearized economy has the following two properties:

1. The economy aggregates: by taking averages across the log-linearized equations for the

49. Like in Chari et al. (2007a) the assumption here is that the state-vector, call it s_t , follows an AR(1) process and that the wedges in period t can be used to uncover the event s_t uniquely, in the sense that the mapping from the event s_t to the wedges is one to one and onto. Thus, given these assumptions, here I write directly the process for the wedges

islands we get an aggregate economy that we can write in terms of wage inflation (π_t^w), the real wage (w^r), and employment (n_t), has only a final goods sector, and only three exogenous processes for the wedges ($\gamma_t^\omega, z_t^\omega, \epsilon_t^\omega$) with $z_t^\omega \equiv z_t^{y^\omega} + \beta z_t^{x^\omega}$. The evolution of the aggregate economy is characterized by the following equations

$$0 = \mathbb{E}_t [mu_{t+1} - mu_t - \pi_{t+1} - \gamma_{t+1}^\omega] + \varphi_\pi \mathbb{E}_t [\pi_{t+1}] + \varphi_y (w_t^r + n_t - y^*) \quad (1.7)$$

$$mu_t = -\sigma C(w_t^r + n_t) \quad (1.8)$$

$$\pi_t^w = \frac{\lambda}{1-\lambda} \left(\epsilon_t^\omega + \frac{1}{\phi} n_t + \sigma (w_t^r + n_t) - w_t^r \right) \quad (1.9)$$

$$w_t^r = -(1 - (\alpha + \theta\beta)) n_t + z_t^\omega \quad (1.10)$$

$$\pi_{t+1} \equiv \pi_{t+1}^w - (w_{t+1}^r - w_t^r) \quad (1.11)$$

$$y^* = \iota z_t^\omega + \eta \epsilon_t^\omega \quad (1.12)$$

$$\lambda \rightarrow 1$$

2. In deviations from aggregate variables, the economy of the islands behaves like small open economies where the interest rate and the price of intermediates are equal to their steady state values. Denoting the log deviation of island's k variable x from the aggregate at time t by $\tilde{x}_t = x_{kt} - x_t$, for a given island the evolution of the economy is characterized by the following equations

$$0 = \mathbb{E}_t [\tilde{m}u_{kt+1} - \tilde{m}u_{kt} - (\tilde{p}_{t+1} - \tilde{p}_t) - \Phi_0 (\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y) - \tilde{\gamma}_{t+1}^\omega] \quad (1.13)$$

$$\tilde{m}u_t = -\sigma C (\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y) \quad (1.14)$$

$$\tilde{w}_t = \tilde{z}_t^{x\omega} - (1 - \theta) \tilde{n}_t^x \quad (1.15)$$

$$\tilde{w}_t = \lambda \left(\tilde{p}_t + \tilde{\epsilon}_t^\omega + \frac{1}{\phi} \left(\frac{N^y}{N} \tilde{n}_t^y + \frac{N^x}{N} \tilde{n}_t^x \right) + \sigma (\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y) \right) + (1 - \lambda) \tilde{w}_{t-1} \quad (1.16)$$

$$\tilde{b}_t = (1 + r) \tilde{b}_{t-1} + \frac{X}{B} (\tilde{n}_t^x - \tilde{n}_t^y) \quad (1.17)$$

$$\tilde{w}_t = \tilde{p}_t - (1 - (\alpha + \theta\beta)) \tilde{n}_t^y - \beta (1 - \theta) (\tilde{n}_t^x - \tilde{n}_t^y) + \tilde{z}_t^{y\omega} + \beta \tilde{z}_t^{x\omega} \quad (1.18)$$

$$\lambda \rightarrow 1$$

Here I have assumed that in linearized form the interest rate rule is $i_{t+1} = \varphi_\pi \mathbb{E}_t[\pi_{t+1}] + \varphi_y(y_t - y_t^*) + \mu_{t+1}$ where π_t is the aggregate inflation rate and $y_t - y_t^*$ is the output gap, defined as the difference between actual output and the flexible wage equilibrium output for the same realization of shocks. These properties are not new⁵⁰, and they are stated formally and proved in Appendix 1.11.1 only for completeness. They are useful for two reasons. First, they allow to use aggregate data that goes back to the 1970s to help estimate the parameters of process 1.6. Second, they allow to easily separate the regional component of the wedges as the regional economies behave like independent small open economies.

With the log-linearized economies, I obtain the models' decision rules using the method of undetermined coefficients as in Uhlig (1995), so that the endogenous variables are expressed as a function of the exogenous state variables (the wedges). Online Appendices 1.12.3 and 1.12.4 provide the solution for the regional and aggregate models respectively. The next step is to estimate the parameters of process 1.6. As in Chari et al. (2007a), I set the parameters of preferences, production functions, and interest rate rule to values commonly used in the

50. These properties were first proved in Beraja et al. (2016)

literature or that I can compute from data⁵¹, and I estimate the parameters of the stochastic process for the states by MLE via Kalman Filter where the hidden states are the wedges. Table 1.1 contains the parameters I use and the moment in the data they target. To this end, it is convenient to express both the aggregate and the regional models in their state space form. For the aggregate model the state-space representation is:

$$[\gamma_t^\omega, z_t^\omega, \epsilon_t^\omega]' = A [\gamma_{t-1}^\omega, z_{t-1}^\omega, \epsilon_{t-1}^\omega]' + B [u_t^\gamma, u_t^z, u_t^\epsilon]' \quad (\text{Aggregate State Equation})$$

$$[\pi_t^w, w_t^r, n_t]' = Q[\gamma_t^\omega, z_t^\omega, \epsilon_t^\omega]' \quad (\text{Aggregate Observation Equation})$$

where Q is a matrix that depends on the parameters of the model and matrix A and it is the result of applying the method of undetermined coefficients. For the regional model the state-space form is given by

$$[\tilde{\gamma}_t^\omega, \tilde{z}_t^{y\omega}, \tilde{z}_t^{x\omega}, \tilde{\epsilon}_t^\omega]' = \tilde{A} [\tilde{\gamma}_{t-1}^\omega, \tilde{z}_{t-1}^{y\omega}, \tilde{z}_{t-1}^{x\omega}, \tilde{\epsilon}_{t-1}^\omega]' + \tilde{B} [v_{kt}^\gamma, v_{kt}^y, v_{kt}^x, v_{kt}^\epsilon]' \quad (\text{Regional State Equation})$$

$$[\tilde{w}_t, \tilde{b}_t, \tilde{p}_t, \tilde{n}_t^y]' = \tilde{P}[\tilde{w}_{t-1}, \tilde{b}_{t-1}]' + \tilde{Q}[\tilde{\gamma}_t^\omega, \tilde{z}_t^{y\omega}, \tilde{z}_t^{x\omega}, \tilde{\epsilon}_t^\omega]' \quad (\text{Regional Observation Equation})$$

where \tilde{P} and \tilde{Q} are the result of applying the method of undetermined coefficients to the regional model.

The objective is to estimate parameters \tilde{A} , \tilde{B} , A and B using aggregate data from 1978 through 2014 and regional data from 2005 through 2014. Given the aggregation properties of this economy, appendix 1.11.2 shows the relation between A , \tilde{A} , and P . For given parameters

51. Since the data is annual and for the regions there are only 9 years of data, I do not attempt to estimate all parameters.

\tilde{A}, \tilde{B}, A and B , I apply the Kalman filter and compute the likelihood for the aggregate economy and for each of the islands, and finally compute a total likelihood weighting each island and the aggregate economy by their population in 2006⁵². The weighting aims at reducing the impact of measurement error, under the assumption that variables are better measured for larger regions than for smaller regions. The MLE procedure looks for the values of $s, \tilde{A}, \tilde{B}, A$ and B , that maximize the likelihood of observing the data. In the estimation I allow for measurement error in the observation equation. The results are unchanged if I do not allow for measurement error.⁵³

1.6.1 *A discussion on heterogeneity*

So far I have made the assumption that islands are ex-ante identical and that they may potentially differ in the shocks that hit them and drive their wedges. This assumption may be too strong and limits the ability to make conclusions about regional heterogeneity and its consequences for business cycles. What would entail to move away from this assumption?

The main consequence of assuming that the islands are identical is that their model-implied response to shocks is the same. This is, given a shock, all islands will have the same response, which precludes the possibility of studying the differential effect of an aggregate shock on the regions. The easiest way to see this is by looking at the state-space representations above. The elements of matrices \tilde{P} , \tilde{Q} , and Q are functions of the parameters in the utility and productions functions, and in the case of Q and \tilde{Q} are also functions of the parameters of process (1.6). These matrices will be identical across islands. One possibility to move away from this is to allow the parameters of the stochastic process (1.6) driving the wedges to vary across islands. However, this alternative is problematic as the length of the time series data for local economies is too short for the number of parameters to be

52. The log-likelihood is given by $L_{total} = pop_{agg}^{2006} \times L_{agg} + \sum_k pop_k^{2006} \times L_k$

53. In Chari et al. (2006), which is the appendix to Chari et al. (2007a), the authors set the measurement error values to be zero for all their experiments.

estimated. In fact, the assumption that these parameters are the same across islands is what allows me to use the cross-region data alongside the aggregate data to estimate them.

Another possibility is to allow for some of the parameters either in the utility function or in the production functions to be different across regions. In the preferences, the parameters are the intertemporal elasticity of substitution (σ) and the Frisch elasticity of labor supply (ϕ). The measures of these two parameters vary greatly in the literature and have been subject of detailed studies. For example, in the case of σ while Hall (1988) does not find strong evidence of it being positive, Eichenbaum et al. (1988) obtain estimates as high as 10 depending on the data used. In the case of ϕ the measurement discussion in the literature is extensive. For instance, Rogerson and Wallenius (2009) compare micro and macro elasticities and find that the macro estimates are orders of magnitudes larger than the micro estimates, and the ranges of the estimates are also large. Providing a regional measure of these parameters is a hard task beyond the scope of this paper; differential measures of these parameters across regions may be as controversial as the assumption that they are the same. A more promising way to introduce heterogeneity would be through the parameters of the production function. In particular, using data on state-level GDP and compensation of employees, one could compute state-level measures of the labor share to parameterize the production functions for the final good in each state. I explore this option in the sections below.

1.7 Business Cycle Accounting Exercise

In this section I conduct the business cycle accounting exercise and analysis at the state level for the Great Recession. The purpose is threefold. First, for each state I aim to identify the wedge, if any, that plays a more prominent role in accounting for the relative behavior of employment and consumption. The result of this will point at the types of frictions and shocks that promising models should include to explain the relative business cycle fluctuations of an individual state. This is the use that the literature has made of

business cycle accounting for national economies. Second, I seek to gauge the extent to which wedges driving the differences across regions for a given variable (consumption or employment) tend to be the same. From a model building perspective this will give a guideline of how complex a model needs to be to capture the variation in real variables across regions. This will accomplish one of the main goals of this paper: help build models that can use cross-region variation and regional data to learn about the shocks that drive *aggregate* business cycles.

Third, I present different ways of gauging systematically the relevance of the different wedges in accounting for the fluctuations that we see in the data. The business cycle literature has developed a useful summary statistic to measure the importance of wedges. However, there are a number of reasons that prevent me from relying solely on the developed metric to reach the conclusions for the context of this paper. One reason is that so far there is not one single statistic that has been shown to measure precisely the importance of wedges. As suggested in Brinca et al. (2016), the assessment should be done in different ways. Another reason, is that the proposed statistic may not be as informative when taken in isolation for local economies as they are for national economies due to two issues. The first of these issues is that for regions within the United States the data can be constructed only at annual frequencies, which will make the statistic more imprecise for local economies than for national economies: while I can compute the statistic using only 9 data points for each region, Brinca et al. (2016) can use about 36 data points for each country as the data for national economies can be generally constructed at quarterly frequencies. The second issue is that the statistic will be more informative if all the wedges move the variables in the same direction. This is generally not the case, and more so when the accounting exercise for a given region is conducted in deviations from the aggregate economy.

Here I focus on results at the state level of geographic aggregation. States provide more flexibility in terms of the data available to them. In particular, consumption data at the

local level is hard to obtain. As the model is written in section 1.6, its estimation does not require consumption data, and therefore the accounting exercise can be performed for MSAs and counties with the data from section 1.5. Researchers using models to study shocks from county or MSA variation can guide their modeling choices using the accounting methodology with prices, net assets, wages, and employment. For states, however, one can use consumption data. This provides the opportunity of comparing the results for the economy composed by regions to the results that have been obtained in the literature using aggregated models and aggregate data.

It is worth reminding the reader that the business cycle accounting exercise of this section is performed using only the regional model to focus on the differences across regions. This means that the wedges will account for movements in the variables of the model expressed in deviations from the aggregate value of the variable. In order to explain how the analysis is conducted I use four states as examples: California, Nevada, New York, and Tennessee.

Figure 1.9 shows the absolute and relative performance of California over the business cycle for employment, real per-capita consumption, and real per-capita output. Panel A shows the cyclical component of the three variables relative to 2007. From peak to trough these variables decreased by about 5 percentage points. Panel B shows the performance of California relative to the aggregate, normalizing all variables to zero in 2007. Relative to the aggregate economy California had bigger declines in all three variables, so its relative position worsened between 2007 and 2010. The accounting exercise is related to this relative position and its objective, as explained by Brinca et al. (2016), is to isolate the marginal effects of the wedges.

To accomplish this objective, one must proceed in two steps. The first step consists of using the decision rules from the model (see Regional Observation Equation) to recover the wedges by solving for the wedges on the right hand side, plugging in the data on the left hand side, and inverting the matrix \tilde{Q} . By construction, the wedges account for the data. In the

case of California, for example, the data will enter as in panel B of Figure 1.9. The second step is to use again the decision rules in which one feeds the wedges back into the model one at a time (or in groups) while setting the other wedges to constants. This will produce again a sequence of employment and consumption, which in the case of this paper will be a sequence of employment and consumption in deviations from the aggregate economy. If the wedge being fed back into the model is the labor wedge, for example, the sequence of employment and consumption is referred to as the labor wedge component of employment and the labor wedge component of consumption, respectively. Here it is important to notice the following. The labor wedge component of employment represents how much employment would fluctuate if the only wedge that fluctuated was the labor wedge and the probability distribution of the labor wedge was the same as in the benchmark economy. This is achieved by holding the other wedges at constant. Since the wedges are correlated, by holding a wedge constant the the direct effect of the wedge is eliminated by its forecasting effect on other wedges is maintained. This guarantees that the expectations of the wedge (or wedges) are identical to those of the benchmark economy.⁵⁴ For a formal, detailed explanation of this, see the seminal paper by Chari et al. (2007a) and the follow-up paper by Brinca et al. (2016).

Figure 1.10 shows the evolution of employment for California in deviations from the aggregate and the contribution of each wedge to it. Each line represents a one-wedge economy. From employment perspective, this exercise suggests that the degree to which the economy in California differed from the aggregate economy was driven mainly by the intermediate wedge and to a lesser degree by the productivity wedge. Figure 1.11 performs the decomposition for consumption, showing that the intertemporal wedge, which is the distortion in the Euler Equation, tracks very closely the actual evolution of consumption, while the other wedges do

54. In this discussion I have followed very closely Brinca et al. (2016). The reason is that the business cycle accounting methodology has been subject to a number of misconceptions and precision in the language is in order.

not. This results suggests that a model with frictions/shocks generating fluctuations in two wedges, would be needed to produce the relative performance of California in consumption and employment that we see in the data.

The main question that this paper aims to answer is how similar are the different regions in the terms of this decomposition. How similar is California to the rest of the states? Let's take Nevada, a state that was the worst performer during the Great Recession. Panel A of Figure 1.12 shows that Nevada's employment moved 11 percentage points downwards from the peak of the recession through the trough while consumption fell by more than 4.5 percentage points. Panel B shows that the relative performance of Nevada was exceptionally poor, and that even by 2014, it had not recovered in these relative terms. Figure 1.13 shows the employment decomposition by contribution of each individual wedge. In this case it is the labor wedge the wedge that contributes the most to the relative evolution of employment. While the wedge responsible for employment is different from California's, the wedge that better accounts for the fluctuations of consumption is the same, the intertemporal wedge, as shown in Figure 1.14. Does taking a relatively good performer change the conclusion? Taking New York, and avoiding Texas or North Dakota due to their close links to the oil industry, one can see a much milder Great Recession in this state with only a 3 to 4 percentage-point peak-trough decline in employment, consumption and output as shown in Panel A of Figure 1.15. Panel B of the same figure shows that relative to the average state, New York did well, the opposite of Nevada and California. Despite this difference, figures 1.16 and 1.17 show that, like Nevada, New York had a recession with relative employment being driven primarily by the labor wedge and relative consumption largely explained by the intertemporal wedge. Thus, to explain the differential behavior of employment in New York and Nevada vis-a-vis the national average, models with frictions driving the labor wedge would be most promising, while for consumption models with fluctuations in the intertemporal wedge would be better.

So far the results suggest some common traits in business cycles across regions but also

differences. In order to assess systematically the importance of different wedges some summary statistics have been suggested in the business cycle accounting literature. One of those statistics, referred to in the literature as the ϕ -statistic, aims to measure how closely a particular component of a variable, say, the intertemporal wedge component of employment, tracks the realized value of employment in deviations from the aggregate. For notation purposes I rename the statistic here as the ψ -statistic to avoid confusion with the parameter I use for labor supply elasticity. For each variable of interest in the model (employment and consumption) one can compute such statistic. For employment, for example, the ψ -statistic is given by

$$\psi_{\omega}^n = \frac{1/\sum_t (n_t - n_{\omega t})^2}{\sum_j \left(1/\sum_t (n_t - n_{jt})^2\right)}$$

where $n_{\omega t}$ is the component of employment due to wedge $\omega = \{\tilde{\gamma}_t^{\omega}, \tilde{z}_t^{y\omega}, \tilde{z}_t^{x\omega}, \tilde{\epsilon}_t^{\omega}\}$.

I compute this statistic for the period 2007-2011 and 2007-2014 for all wedges individually, and for some groups of wedges. I only report the statistic for 2007-2014 as they show essentially the same patterns, but the 2007-2014 is computed with more data points. Figure 1.18, shows that for 60% of the states the labor wedge by itself accounts for more than 50% of the movements in employment. The productivity wedge it is important only for California and Washington. For only 7 states, the labor wedge accounts for less than 20% of the movements in output. The different colors and signs in the scatter plot are meant to reflect this classification. A number of states are labeled as no leading wedge, since no wedge accounts for a notably large fraction of the fluctuations using the statistic. Since the statistic is not computed with many data points, I do not fully rely on just one measure to assess the importance of a given wedge. Figure 1.19 presents the ψ -statistic for an economy with more than one wedge. The x-axis shows the ψ -statistic for an economy with no productivity wedges, it is a two wedge economy, mainly driven by the labor wedge. The y-axis shows the contribution of three wedges without the labor wedge. For only 6 states, an economy

without the labor wedge accounts for more than 50% of fluctuations in employment. Table 1.2 reports the exact values of the statistic for each state; the last two columns correspond to the multiple-wedge economies.

A state like Tennessee (TN) provides a good example of why I do not rely fully on the ψ -statistic to assess the most important wedge for a given state. As shown in figure 1.20, the labor wedge component tracks the fluctuations of employment quite well but as shown in Figure 1.18, from the measure of the ψ -statistic the labor wedge did not seem as important. As this example shows, the time series plots for the different states provide a way to help gauge the importance of a given wedge. Figures 1.10, 1.13, 1.16, show also that contrasting the change in the data from from peak (2007) to trough (2010) to the change in each of the components also provides information to value the significance of a given wedge in accounting for the movement in the underlying variable during the downturn. Table 1.3 reports for each state the trough year and the change in employment (in deviations from the aggregate) and each of its components between 2007 and 2010 (which is the trough year for the aggregate). In general these measures provide the same message as the ψ -statistic. In the last two columns of this table I combine the information from the first columns of the table, the ψ -statistic, and the time series decomposition of employment for each state, and classify each state by the wedge that seemed to have the most influence during the downturn (column “Peak-trough Wedge”) and by the wedge that seemed to have the most influence during the entire period 2007-2014 (column “2007-2014 Wedge”). Whenever there are two wedges mentioned, the one that appears first is judged to be the primary wedge. The evidence suggests that the labor wedge was the main driver of differences in employment fluctuations across states, with 37 states having the labor wedge as the peak-trough wedge and 40 states having the labor wedge as the main driver of fluctuations between 2007-2014.

Turning to consumption, the pattern is stronger. Figure 1.21 shows that two wedges, the intertemporal wedge and the productivity wedge account for more than two thirds of the

movements in relative consumption for all but six states. This is highlighted in the figure through the the downward-sloping lines, which represent combinations for which the sum of the productivity and intertemporal wedge components is 67% and 90%. More remarkable is the fact that the intertemporal wedge by itself accounts for more than 40% of consumption for about 34 states. Figure 1.22 highlights the importance of the intertemporal wedge. This figure shows the ψ -statistic for two three-wedge economies, one economy without the intertemporal wedge (on the y-axis) that I refer to as the productivity wedge economy as it is mainly driven by the productivity wedge, and one economy without the productivity wedge (on the x-axis) that I refer to as the interteporal wedge economy as it is primarily driven by the intertemporal wedge. For the majority of states (31 out of 48), the three-wedge economy that does not contain the intertemporal wedge is unable to account for more than 50% of the movements in relative consumption. Table 1.4 reports the exact values of the statistic for each state; the last two columns correspond to the multiple-wedge economies.

In Table 1.5 I report the changes in consumption and its wedge components from peak to trough (2007-2010). As with employment, changes in the components tell a similar story as the ψ -statistic, suggesting a primary role for the intertemporal wedge during the downturn. Using this information, the ψ -statistic, and the time series decomposition of consumption for each state, I perform the classification of each state by the wedge driving its consumption with respect to the aggregate. The last two columns contain the classification. During the downturn (column “Peak-trough Wedge”) the intertemporal wedge played a prominent role for 37 states and a secondary, but important role for another 5 states. Similarly, for the entire time period (column “2007-2014 Wedge”) the intertemporal wedge was key for 35 states, and it was important for another 5 states.

The metrics that I have used up to here are not perfect. In fact, the literature has not developed a single statistic that allows one to make a definite assessment about the importance of different wedges in producing the fluctuations in the data and to summarize

the results. In on-line appendix section 1.12.2, I present three alternative metrics for both employment and consumption and also for different values of parameters. These different ways of gauging the role of each wedge in accounting for the business cycle fluctuations of employment and consumption lead all to the same conclusion.

In sum, this business cycle accounting exercise shows a great degree of commonality across regions in terms of the shocks and frictions that drove the relative behavior of employment and consumption during the Great Recession. The results indicate that models that produce fluctuations in the labor wedge will be able to account for a large fraction of the cross-sectional employment, while models that produce fluctuations in the intertemporal wedge will be able to account for a large fraction of the cross-regional consumption. Overall, even though there is evidence from previous studies (for example Ohanian (2010) and Brinca et al. (2016)) that the Great Recession for the United States was a labor wedge driven cycle, the evidence from this exercise suggests that one wedge is not enough to account for the differences in both employment and consumption across regions. The good news, however, is that the fact that the same wedge tends to be important for all regions indicates that in principle one does not need multiple types of frictions and shocks in a model to capture the cross-regional variation.

1.8 Sensitivity of Results

In this section I perform some changes to the baseline specification in order to explore if the main results change. As explained in section 1.7, determining the most important wedge driving fluctuations at the local level is not a task that can be done through just one statistic. The reason is that the frequency of the data is not high enough to produce statistics (like the ψ -statistic) that by themselves give a definite, reliable answer. In this section, however, I do not perform a thorough classification of states by the wedges that drive their fluctuations as I did in the previous section. Instead and despite its limitations, I rely on the ψ -statistic

to assess if the same broad patterns persist as I make changes. The reason is that for the question in this paper it is more important to gauge the degree to which regions look similar among each other in terms of the main wedges and if the main wedges are generally the same across specifications, than to actually determine whether a specific state experienced a “labor wedge recession” or to measure precisely the contribution of a given wedge.

I explore the effects of three types of changes 1) Change in the value of parameters of the utility function ϕ and σ ; 2) Change in the functional form of the preferences; 3) Change in some of the data that goes into the model to address the concerns of measurement error. For each of these changes, I re-estimate the parameters governing the evolution of the wedges using the same procedure described in section 1.6 to get maximum likelihood estimates for every case.⁵⁵ Then I recover the wedges and perform the decompositions as explained in section 1.7.

Figures 1.23, 1.24 show the labor and consumption decompositions for a value of the elasticity of labor supply parameter $\phi = 1$, while Figures 1.25, 1.26 do it for $\phi = 4$. The general result does not change from the baseline specification ($\phi = 2$), the labor wedge plays a primary role for employment while the intertemporal wedge does it for consumption. A more elastic labor supply elasticity gives relative more importance to the labor wedge for employment. Figures 1.27, 1.28 vary the elasticity of substitution, changing it from $\sigma = 2$ (the baseline number) to $\sigma = 1$. This will reduce the impact in the first order condition of the wealth effect represented by the fall in consumption, leading to a less important role of the labor wedge for employment; however the wedge still preserves its leading role. In this case the intertemporal wedge play an even bigger role for consumption. Appendix figures A1, A2 perform the exercise for $\sigma = 4$. In this case the intertemporal wedge does not play such an important role for consumption; a value of $\phi = 1$ restores the importance of the

55. It is worth noting that each re-estimation is a computationally demanding exercise as I start the estimation for each case in thousands of different points to increase the probability of achieving the optimum every time.

intertemporal wedge in this case as shown in appendix figures A3, A4. Overall, for these variation in the parameters the results seem not to change much: the labor wedge plays a leading role in accounting for fluctuations in employment for the majority of states while the intertemporal wedge does it for consumption.

The wedges that we recover also depend on the functional forms of the utility and production functions. I explore the effect of changing to GHH preferences, not allowing for wealth effects in the labor decision of the households. Specifically, the functional form changes to

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} e^{-\rho kt} \frac{(C_{kt} - \frac{\phi}{1+\phi} N_{kt}^{\frac{1+\phi}{\phi}})^{1-\sigma}}{1-\sigma} \right]$$

This change will produce a different set of decision rules from the model. The on-line appendix makes also the derivations for this case. With the new decision rules I perform the business cycle accounting exercise using the same parameters as for the baseline case⁵⁶ Figures 1.29 and 1.30 present the results of the decomposition of employment. The results are different in this case than when one allows for wealth effects. In this case both the productivity and the intermediate wedge become the main wedges at the expense of the labor wedge. The assumption about income effects is clearly not innocuous. In the recession as consumption falls, there is a negative wealth effect that pushes people to work more, if we do not see this in the data the labor wedge increases further. Figure 1.3 shows that income and consumption fell sharply during the recession, with their cyclical components falling up 8.5 and 4.5 percentage points respectively for some states during the period 2007-2010. There is ample evidence that wealth effects were operational during the recession, mainly through falling house prices. Given this evidence I prefer not to use GHH as baseline. Figures 1.31 and 1.32 complete the analysis for consumption. In the case of consumption conclusions do

56. I change the technical parameter Φ_0 from 0.01 to 100, otherwise the model produces labor and intertemporal wedge components that are too large.

not change, the intertemporal wedge plays an even more important role in accounting for movements in consumption.

A common conclusion across all specifications is that there is not one single wedge that can be identified as the primary driver of fluctuations of consumption and employment. From regional variation, explaining fluctuations in real variables require shocks and frictions that work through two wedges; the labor wedge looks unable to play the primary role that it plays when a business cycle accounting exercise is performed for the aggregate US economy.

1.9 Discussion and Conclusions

During the Great Recession there was large regional heterogeneity in the behavior of economic variables like employment and consumption. In this paper I have set out to identify the types of shocks and frictions that may have produced such cross-regional variation by pursuing a business cycle accounting approach. To accomplish this goal I have extended the business cycle accounting methodology to allow for the study of regional data through the lenses of a model of a monetary union.

I have made three contributions. First, a methodological contribution by laying out conditions by which the application of the methodology does not involve solving and producing the wedge-related counterfactuals of all the regions simultaneously. This makes the business cycle accounting exercise very tractable. Second, I have constructed a dataset that allows one to use both aggregate and regional data simultaneously to estimate a relatively large number of parameters. More importantly, the dataset allows for the application of the wedge accounting methodology at three levels of geographic aggregation (states, MSAs, and counties) for the United States. This opens the possibility of studying regional business cycles using some of the tools that we have available for the aggregate economies. Our understanding of local business cycles is limited, largely due to data constraints. However, local economies behave differently over the business cycle, and large economic episodes like

the Great Recession may be characterized by large regional variation. Given this, our understanding of the Great Recession and of how shocks propagate across the country would benefit from our ability to study local business cycles. Furthermore, Beraja et al. (2016) have shown that the use of *regional data* and *regional variation* to understand the behavior of *aggregate variables* over the business cycle requires the development of models since the use of only cross-sectional data variation has very limited ability to inform us about the shocks that determine the dynamics of aggregate variables. In this paper, I have developed tools and results that help guide the construction of the required models that would allow researchers to use regional data and exploit regional variation in order to study aggregate business cycles.

The third contribution is the application of the methodology at the state level. The analysis shows that two different wedges are more prominent in accounting for the differential behavior of employment and consumption across states during the Great Recession. Shocks and frictions that produce fluctuations in the labor wedge are more promising to explain the cross-region variation in employment. Shocks and frictions that produce fluctuations in the intertemporal wedge are more promising to explain the cross-region variation in consumption. The implication for theory is that for a model to be able to capture and produce the fluctuations in employment and consumption that we see in regional data it needs theoretical mechanisms that produce variation in the two wedges. For example, in the case of the housing shock of Mian and Sufi (2014), one would need a combination of something like a discount rate shock (which produces variation in the intertemporal wedge) and sticky wages (which generate a labor wedge) to qualitatively capture that such shock can be the primary source of cross-region variation in both consumption and employment.

The reason these results are important to our understanding of the Great Recession and of both aggregate and regional business cycles more generally is the following. When one performs a similar accounting exercise using aggregate data and aggregate models (such as

in Brinca et al. (2016)) the Great Recession in the US looks like a labor-wedge recession. From this exercise we would not rule out models that have undistorted Euler equations (no intertemporal wedge). However, when we open the black box to study regional economies and we take into account that markets are incomplete, that it is difficult at business cycle frequencies to shift people and consumption around so that marginal utilities equalize, the Great Recession no longer looks like a labor wedge-only recession. From regional data and regional variation, shocks and frictions that distort the Euler Equation are crucial to understand the dynamics of consumption and we would rule out models that have undistorted Euler equations.

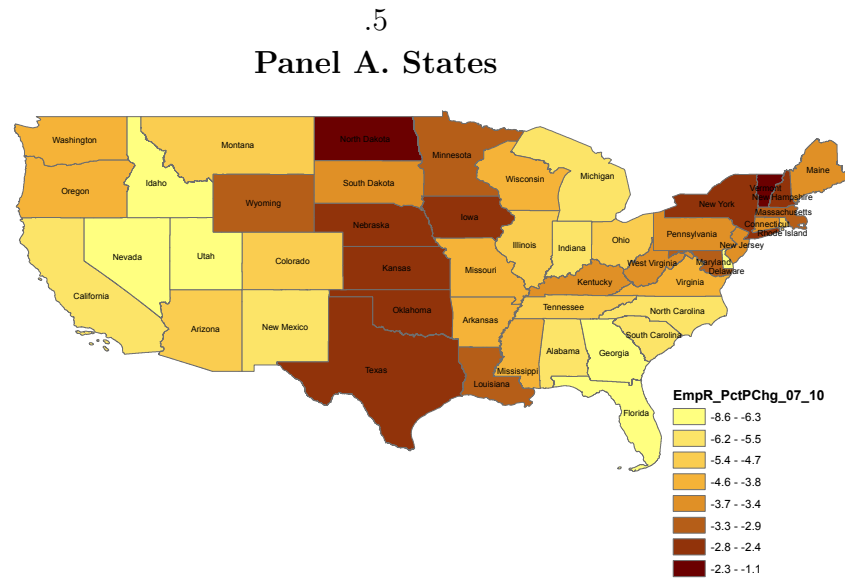
This is further evidence that the link between regional business cycles and aggregate business cycles is complex. In this paper I have shown that what we learn from aggregate business cycles (using business cycle accounting) does not give a full picture of regional business cycles. This complements the analysis by Beraja et al. (2016), who show that what we learn from regional business cycles does not give a full picture of aggregate business cycles.

Following Chari, Kehoe, and McGrattan (2007a), in order to be able to use the results of this paper to point at the mechanisms that are more likely to have created such differential business cycle fluctuations across regions, one needs to produce the corresponding equivalence results for the class of models that can be represented with the benchmark model of the paper. More research is required in this direction.


Finally, the methodology of the paper can be easily extended and applied to a context of a fiscal union using existing state-level data on taxes and subsidies. Regional data on investment and capital, which to my knowledge do not exist, would greatly improve the reach of the contributions in this paper to help us better understand sub-national business cycles.

1.10 Figures and Tables

Figure 1.1: Regional Variation in Employment During the Great Recession:
2007-2010

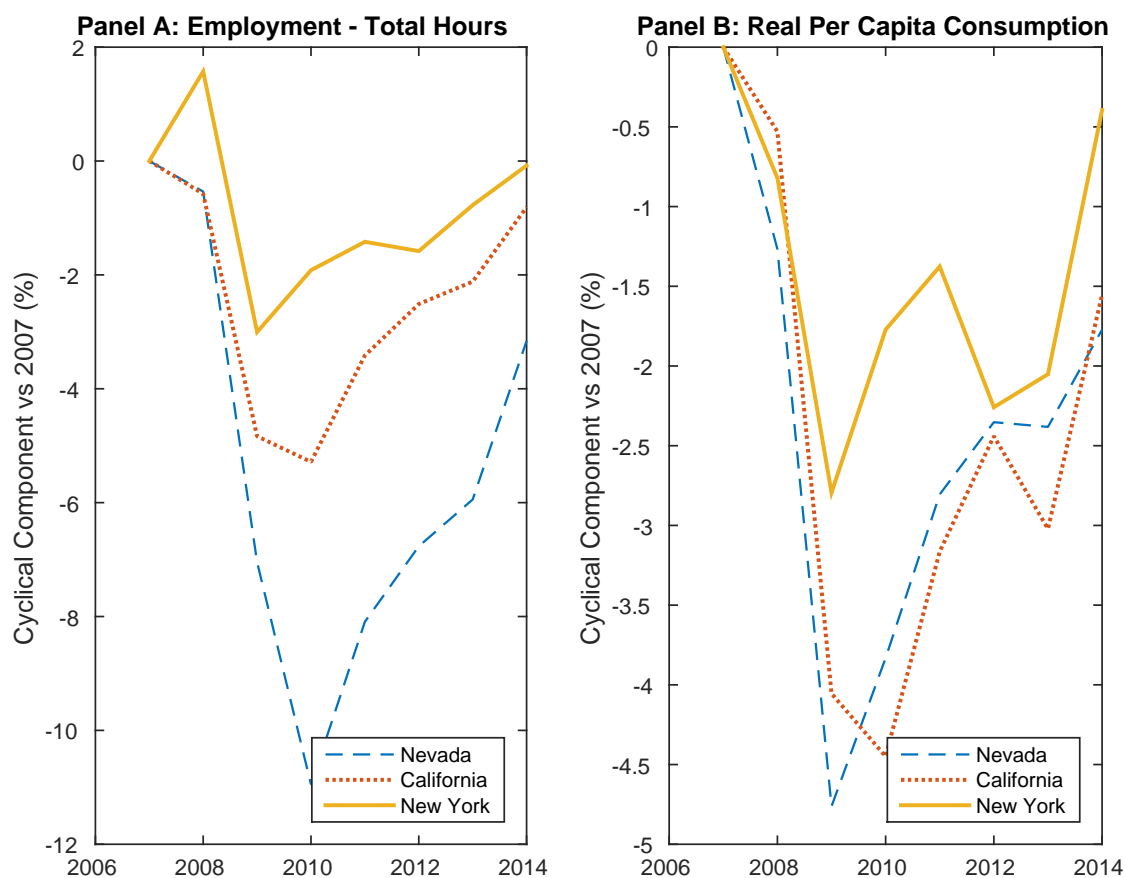


Panel B. Counties .5



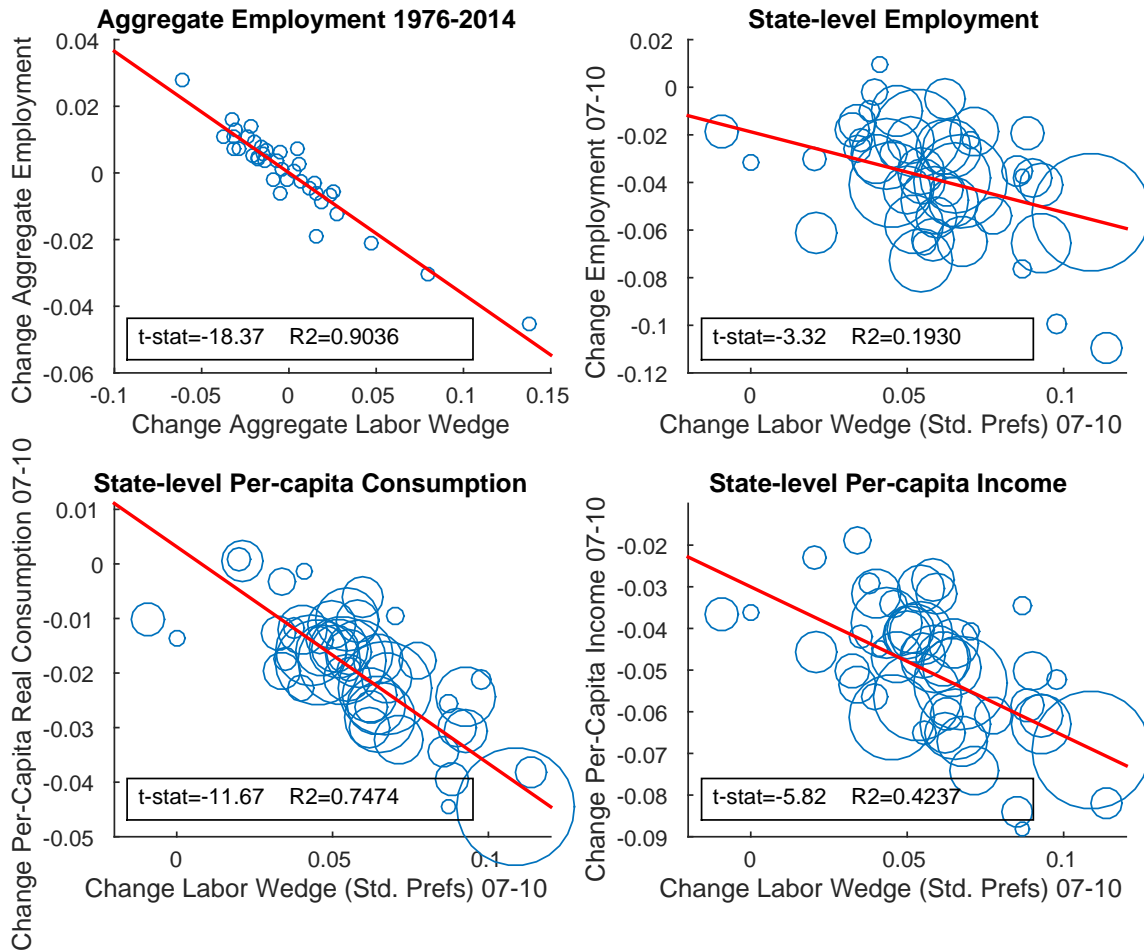
This figure shows the percentage point change in the employment rate for states (Panel A) and for counties (Panel B) in The United States between 2007 and 2010. The employment rate is computed as the number of employed people divided by the civilian population aged 16-64. The data comes from the LAUS program of the Bureau of Labor Statistics.

Figure 1.2: Regional Variation in Employment and Consumption over the Business Cycle: 2007-2014



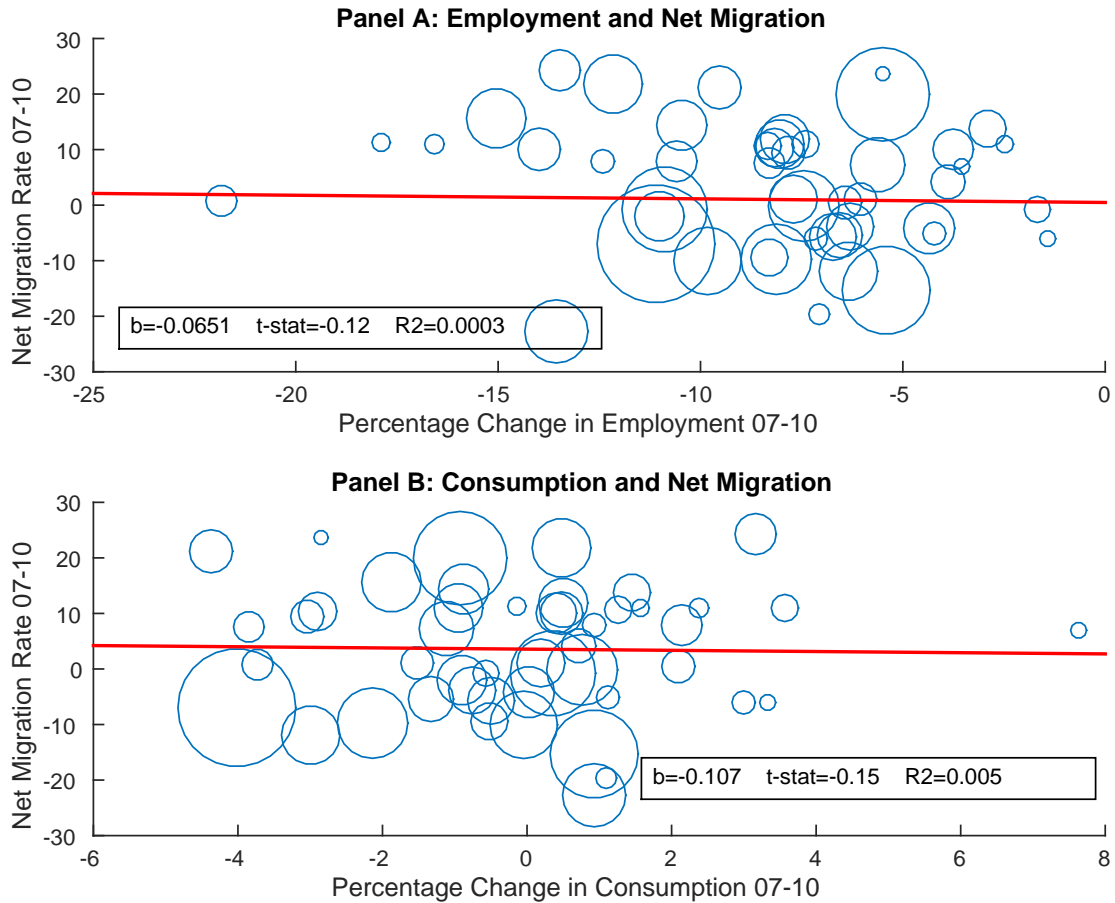
Panel A. presents the evolution of the cyclical component of total hours for three sample states, Nevada, California, and New York during the period 2007-2014. See section 1.5 for the exact construction of the variables. Panel B. shows the evolution of the cyclical component of consumption for the same states over the same time period. Consumption data comes from the Bureau of Economic Analysis.

Figure 1.3: Household Side of the Labor Wedge and the Business Cycle




The scatter plot at the top left corner of this figure shows one-period time series changes in the labor wedge computed from the aggregate model with one-period changes in the cyclical component of aggregate employment between 1976 and 2014. The other three scatter plots show the cross-sectional cyclical component of the labor wedge compute-at the state level with the cyclical component of state level employment, real per-capita consumption and per-capita personal income

Figure 1.4: Migration, Consumption Growth and Employment Growth: 2007-2010




Panel A shows a scatter plot of net migration rates against employment growth over the 2007-2010 period. Panel B. shows a scatter plot of net migration rates against consumption growth over the 2007-2010 period. Employment and consumption are constructed as explained in Section 1.5 while migration rates are computed using IRS data available at <https://www.irs.gov/uac/soi-tax-stats-migration-data>

Figure 1.5: State Level Consumption Growth vs Employment Growth:
2007-2010



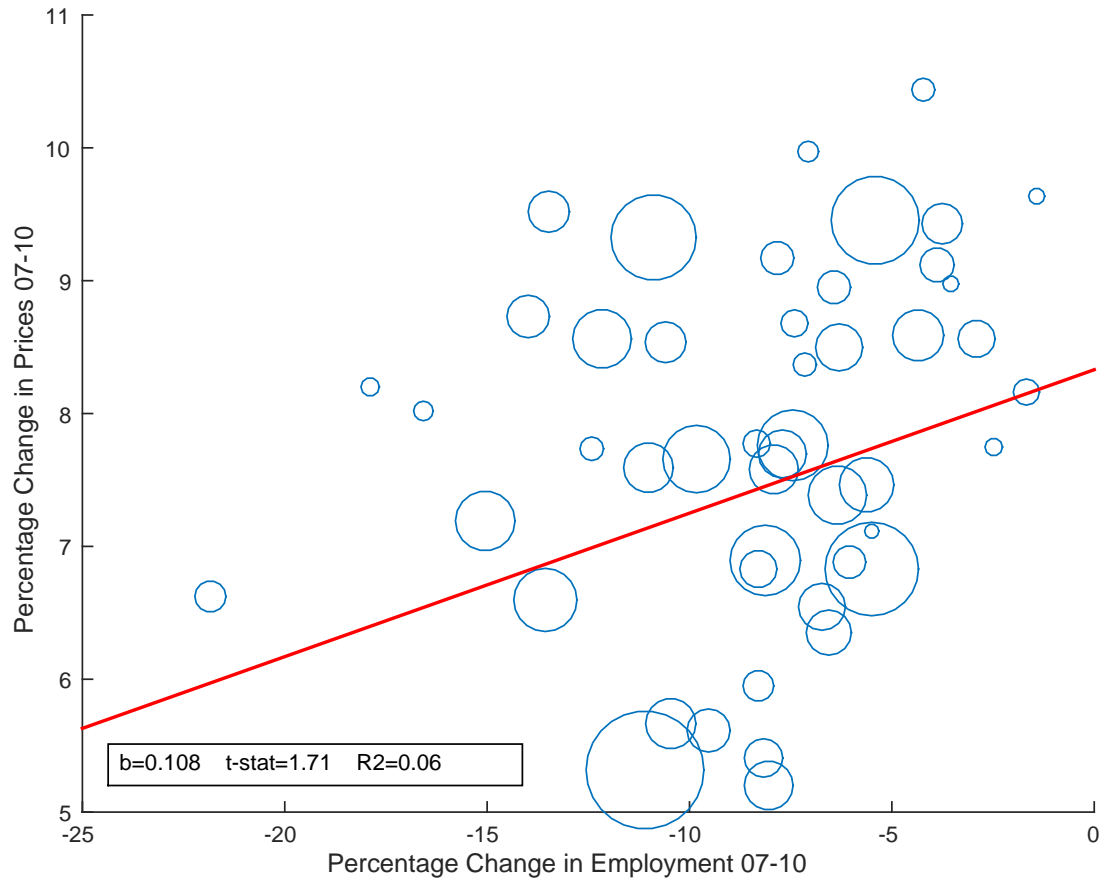
This figure shows a scatter plot of changes in consumption against changes in employment over the 2007-2010 time period. Employment is measured as in section 1.5 and captures both the intensive and the extensive margin of the labor decision.

Figure 1.6: State Level Wage Growth vs Employment Growth: 2007-2010




This figure shows a scatter plot of changes in wages at the state level against changes in employment over the 2007-2010 time period. Wages are computed as in section 1.5.

Figure 1.7: State Level Price Growth vs Employment Growth: 2007-2010



This figure shows a scatter plot of changes in local prices against changes in employment over the 2007-2010 time period. The geographic units of observation are states. Prices and employment are computed as in section 1.5.

Figure 1.8: Net Asset Growth vs Employment Growth: 2007-2010



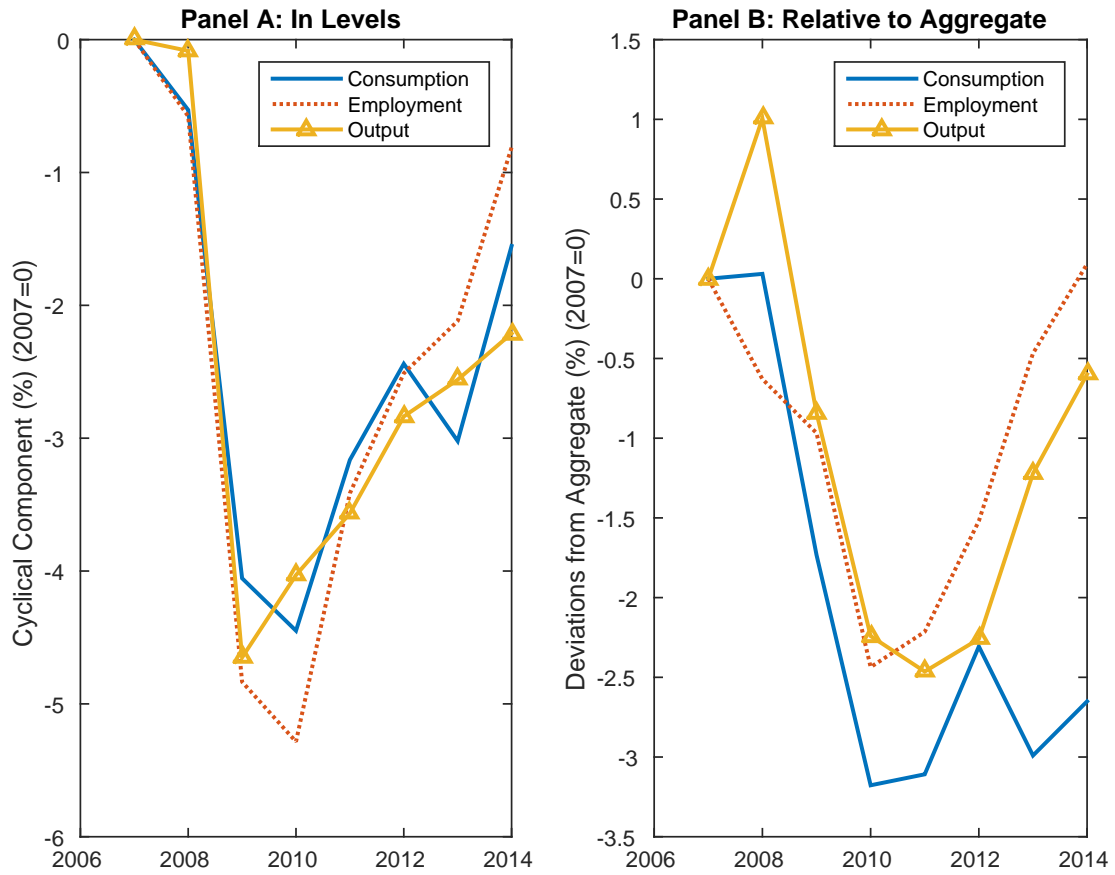
This figure shows a scatter plot of changes in local net assets against changes in employment over the 2007-2010 time period. The geographic units of observation are states. Net assets and employment are computed as in section 1.5.

Table 1.1: **Parameter Values**

This table shows the parameter values used in the baseline estimation of the model. Other parameters of the model are functions of these, consistent with the model solution in steady state

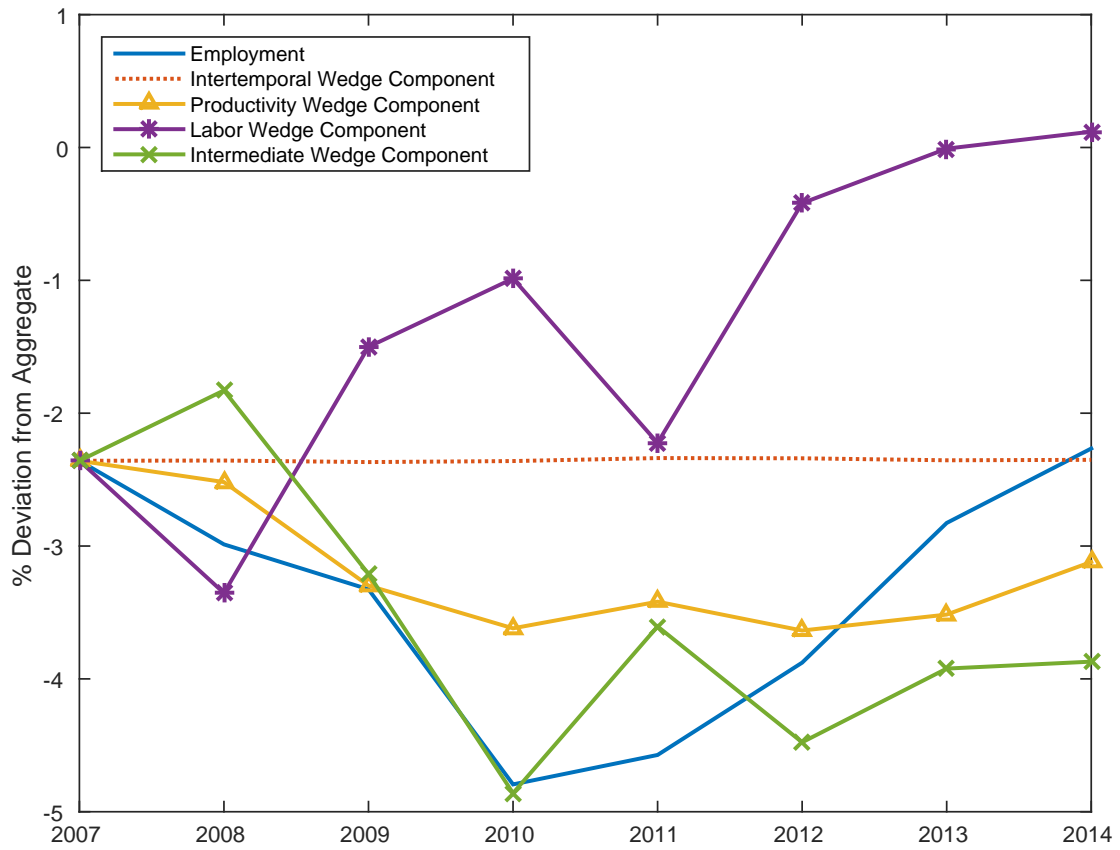
	Value	Target
α	0.62	Aggregate labor share in the non-tradable sector in 2006
θ	0.55	Aggregate labor share in the tradable sector in 2006
β	0.16	Aggregate labor share in 2006 (0.61). See on-line appendix
σ	2	Intertemporal elasticity of substitution. See Kehoe et al. (2016)
ϕ	2	Frisch elasticity. See Beraja et al. (2016) for estimation
λ	1	Wage stickiness.
φ_p	1.5	Taylor Rule from Galí (2011)
φ_y	0.5	Taylor Rule from Galí (2011)
Φ_0	0.01	Trade balance-output ratio volatility = 2 from Mendoza (1991)
R	0.03	Real interest rate
X	0.32	Intermediate inputs over output ratio US 2006
B	2.1	Median net worth to output ratio in the US 2006

Figure 1.9: Great Recession in California




Panel A shows the cyclical component of employment consumption and output for California during the period 2007 through 2014, with respect to 2007. Panel B shows the relative behavior of each of these variables for California in deviations from the aggregate economy.

Figure 1.10: Relative Employment and Employment Components in California



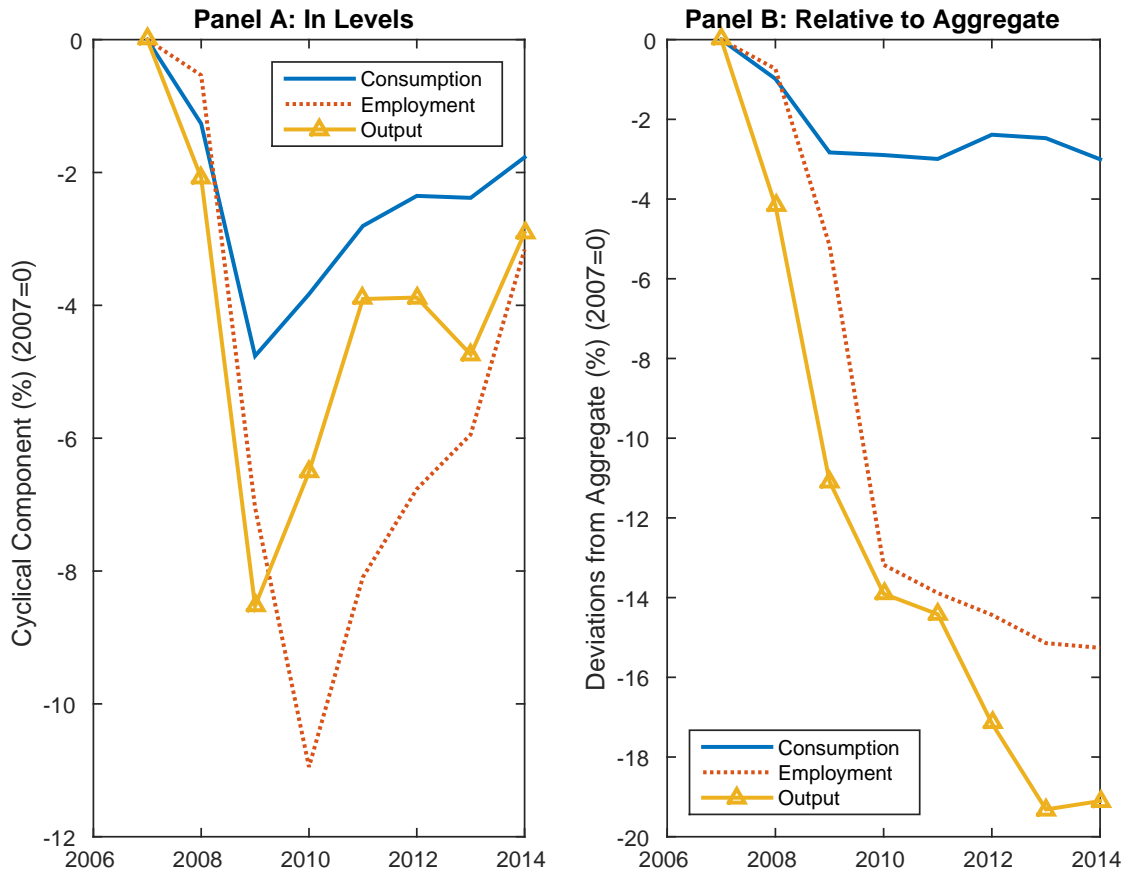
This figure shows the behavior of employment in California relative to the aggregate and its decomposition in wedge components for all wedges in the model.

Figure 1.11: Relative Consumption and Consumption Components in California




This figure shows the behavior of per-capita consumption in California relative to the aggregate and its decomposition in wedge components for all wedges in the model.

Figure 1.12: Great Recession in Nevada




Panel A shows the cyclical component of employment consumption and output for Nevada during the period 2007 through 2014, with respect to 2007. Panel B shows the relative behavior of each of these variables for Nevada in deviations from the aggregate economy.

Figure 1.13: Relative Employment and Employment Components in Nevada



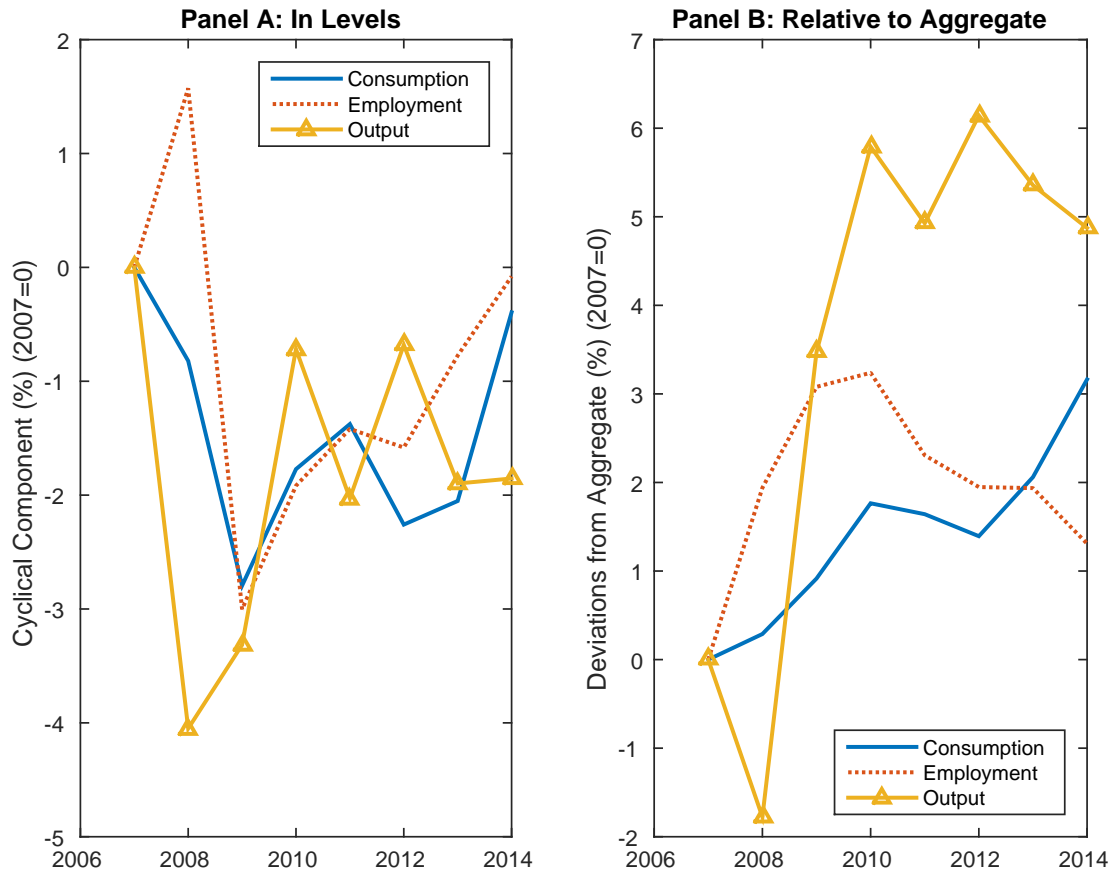
This figure shows the behavior of employment in Nevada relative to the aggregate and its decomposition in wedge components for all wedges in the model.

Figure 1.14: Relative Consumption and Consumption Components in Nevada




This figure shows the behavior of per-capita consumption in Nevada relative to the aggregate and its decomposition in wedge components for all wedges in the model.

Figure 1.15: Great Recession in New York



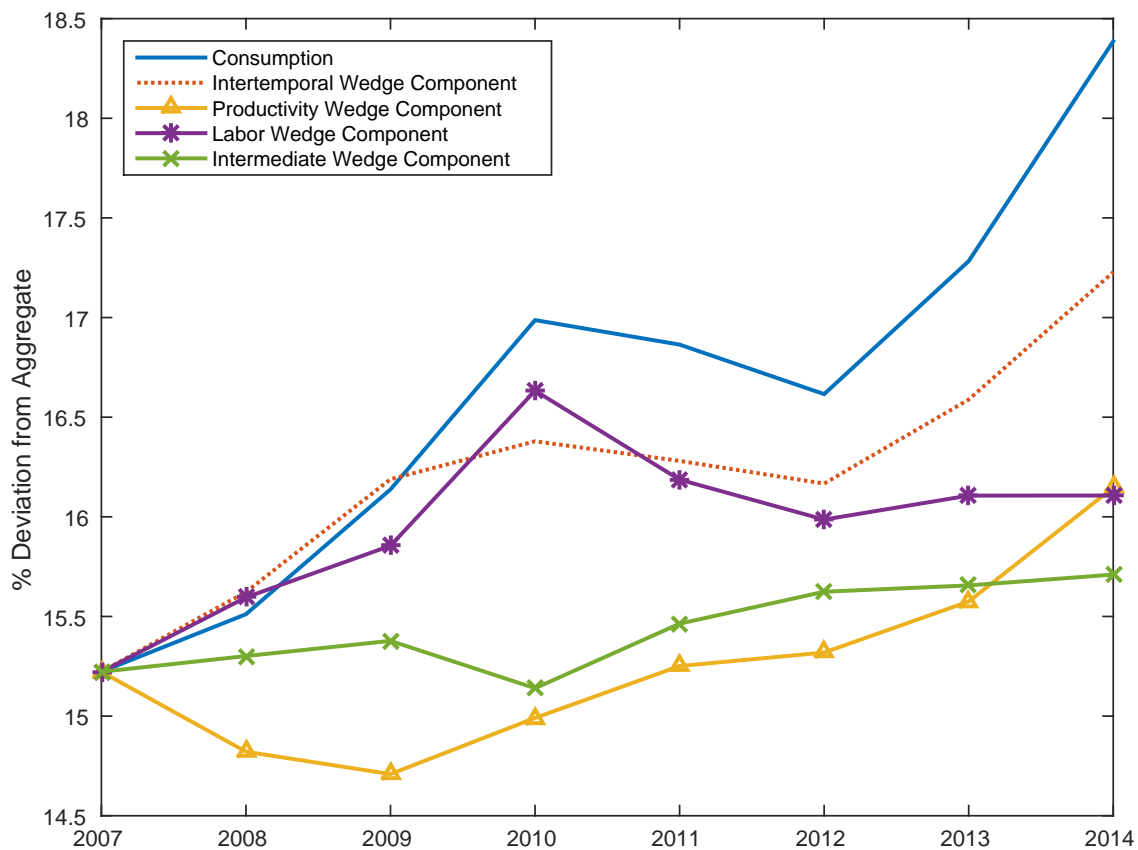
Panel A shows the cyclical component of employment consumption and output for New York during the period 2007 through 2014, with respect to 2007. Panel B shows the relative behavior of each of these variables for New York in deviations from the aggregate economy.

Figure 1.16: Relative Employment and Employment Components in New York




This figure shows the behavior of employment in New York relative to the aggregate and its decomposition in wedge components for all wedges in the model.

Figure 1.17: Relative Consumption and Consumption Components in New York



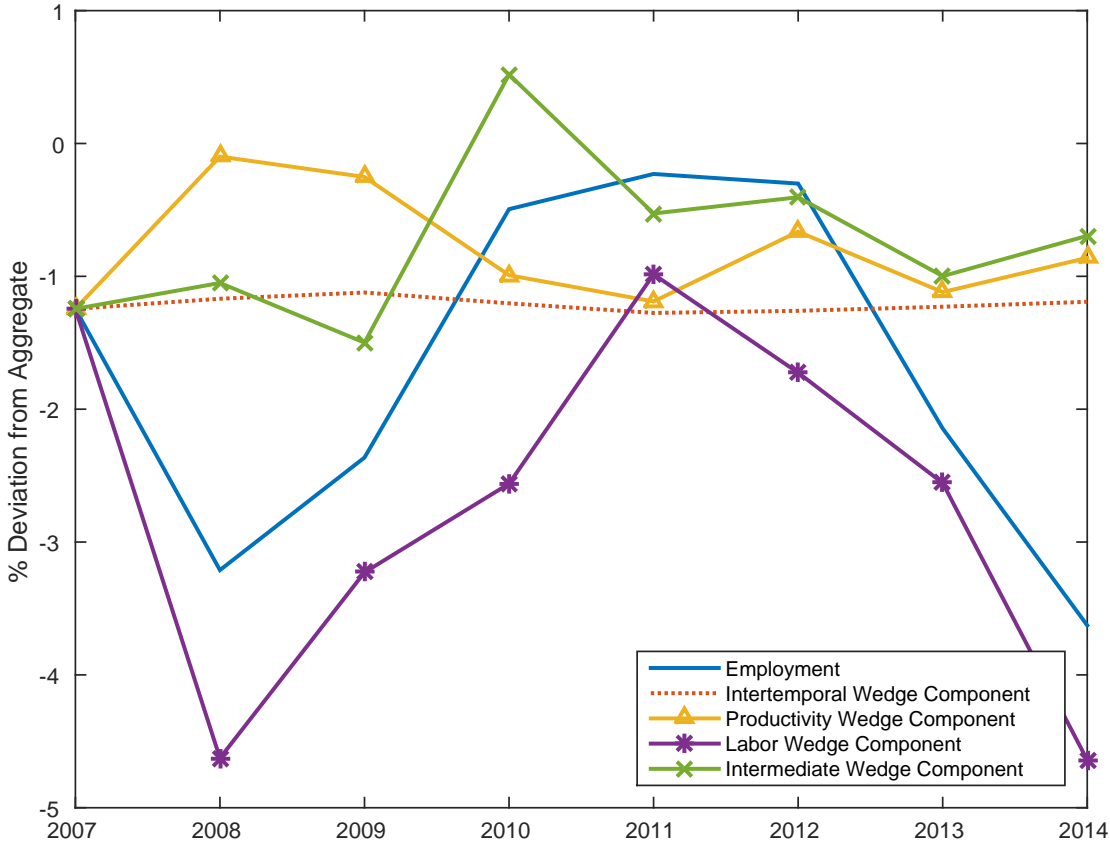
This figure shows the behavior of per-capita consumption in New York relative to the aggregate and its decomposition in wedge components for all wedges in the model.

Figure 1.18: Decomposition of Employment One-Wedge Economy: 2007-2014



This figure shows the contribution of each wedge to employment in the regional model using the ψ -statistic computed over the period 2007-2014. The x-axis corresponds to the statistic for the labor wedge-only economy while the y-axis corresponds to the statistic for the productivity-wedge only economy.

Figure 1.20: Relative Employment and Employment Components in Tennessee



This figure shows the behavior of employment in Tennessee relative to the aggregate and its decomposition in wedge components for all wedges in the model.

Table 1.2: ψ -statistic for Employment: 2007-2014

This table presents the exact values of the ψ -statistic for the decomposition of employment state by state. The first 4 columns correspond to the contribution of individual wedges while the last two columns correspond to the contribution of the indicated subset of wedges. The computations correspond to the period 2007-2014.


Fips Code	State	ϕ_{γ}^n	ϕ_y^n	ϕ_x^n	ϕ_{ϵ}^n	$\phi_{\gamma\epsilon}^n$	$\phi_{\gamma xy}^n$
1	Alabama	0.18	0.12	0.13	0.58	0.73	0.12
4	Arizona	0.19	0.36	0.23	0.21	0.26	0.54
5	Arkansas	0.19	0.16	0.12	0.53	0.67	0.15
6	California	0.14	0.49	0.32	0.04	0.05	0.73
8	Colorado	0.53	0.11	0.29	0.07	0.16	0.67
9	Connecticut	0.54	0.15	0.16	0.14	0.31	0.40
10	Delaware	0.11	0.08	0.09	0.73	0.84	0.08
12	Florida	0.23	0.06	0.42	0.29	0.43	0.20
13	Georgia	0.07	0.05	0.07	0.81	0.90	0.05
16	Idaho	0.21	0.10	0.19	0.50	0.66	0.15
17	Illinois	0.37	0.26	0.21	0.16	0.26	0.54
18	Indiana	0.07	0.07	0.06	0.80	0.71	0.05
19	Iowa	0.12	0.11	0.22	0.55	0.61	0.23
20	Kansas	0.17	0.07	0.25	0.51	0.65	0.14
21	Kentucky	0.16	0.17	0.10	0.57	0.70	0.14
22	Louisiana	0.13	0.16	0.20	0.51	0.57	0.28
23	Maine	0.23	0.16	0.22	0.39	0.52	0.25
24	Maryland	0.21	0.11	0.20	0.48	0.65	0.17
25	Massachussets	0.14	0.12	0.09	0.65	0.78	0.10
26	Michigan	0.16	0.13	0.10	0.61	0.75	0.12
27	Minnesota	0.14	0.08	0.16	0.62	0.76	0.11
28	Mississippi	0.27	0.25	0.33	0.15	0.22	0.51
29	Missouri	0.17	0.16	0.33	0.34	0.39	0.44
30	Montana	0.05	0.05	0.05	0.85	0.90	0.05
31	Nebraska	0.04	0.03	0.05	0.89	0.93	0.03
32	Nevada	0.03	0.02	0.04	0.90	0.95	0.03
33	New Hampshire	0.09	0.19	0.10	0.62	0.68	0.25
34	New Jersey	0.21	0.12	0.22	0.44	0.57	0.23
35	New Mexico	0.10	0.10	0.07	0.73	0.81	0.09
36	New York	0.05	0.07	0.07	0.81	0.80	0.07
37	North Carolina	0.31	0.14	0.19	0.36	0.55	0.20
38	North Dakota	0.19	0.09	0.21	0.50	0.66	0.15
39	Ohio	0.38	0.21	0.15	0.26	0.43	0.30
40	Oklahoma	0.05	0.04	0.07	0.84	0.89	0.05
41	Oregon	0.30	0.25	0.23	0.22	0.33	0.64
42	Pennsylvania	0.15	0.13	0.21	0.50	0.58	0.24
44	Rhode Island	0.42	0.13	0.27	0.18	0.34	0.32
45	South Carolina	0.33	0.11	0.25	0.30	0.48	0.22
46	South Dakota	0.17	0.09	0.19	0.56	0.71	0.13
47	Tennessee	0.26	0.16	0.23	0.35	0.49	0.26
48	Texas	0.20	0.10	0.19	0.51	0.66	0.15
49	Utah	0.50	0.08	0.31	0.11	0.24	0.40
50	Vermont	0.05	0.05	0.06	0.84	0.88	0.06
51	Virginia	0.09	0.07	0.08	0.76	0.85	0.08
53	Washington	0.12	0.46	0.07	0.34	0.41	0.48
54	West Virginia	0.29	0.28	0.07	0.35	0.51	0.27
55	Wisconsin	0.09	0.14	0.13	0.64	0.69	0.24
56	Wyoming	0.36	0.20	0.19	0.25	0.39	0.34

Table 1.3: **Peak to Trough Changes in Employment and Components and Wedge-based Classification of States**

This table shows the year in which the economy of each state wen through the trough of the recession. The next column shows the change in employment in deviations from aggregate employment in the data between 2007 and 2010, the trough-year for the US economy. The next 3 columns show the change in the wedge-components of employment. The last two columns show the main wedge driving fluctuations in employment for each state for the peak-trough period of 2007-2010 and for the entire period 2010-2014.

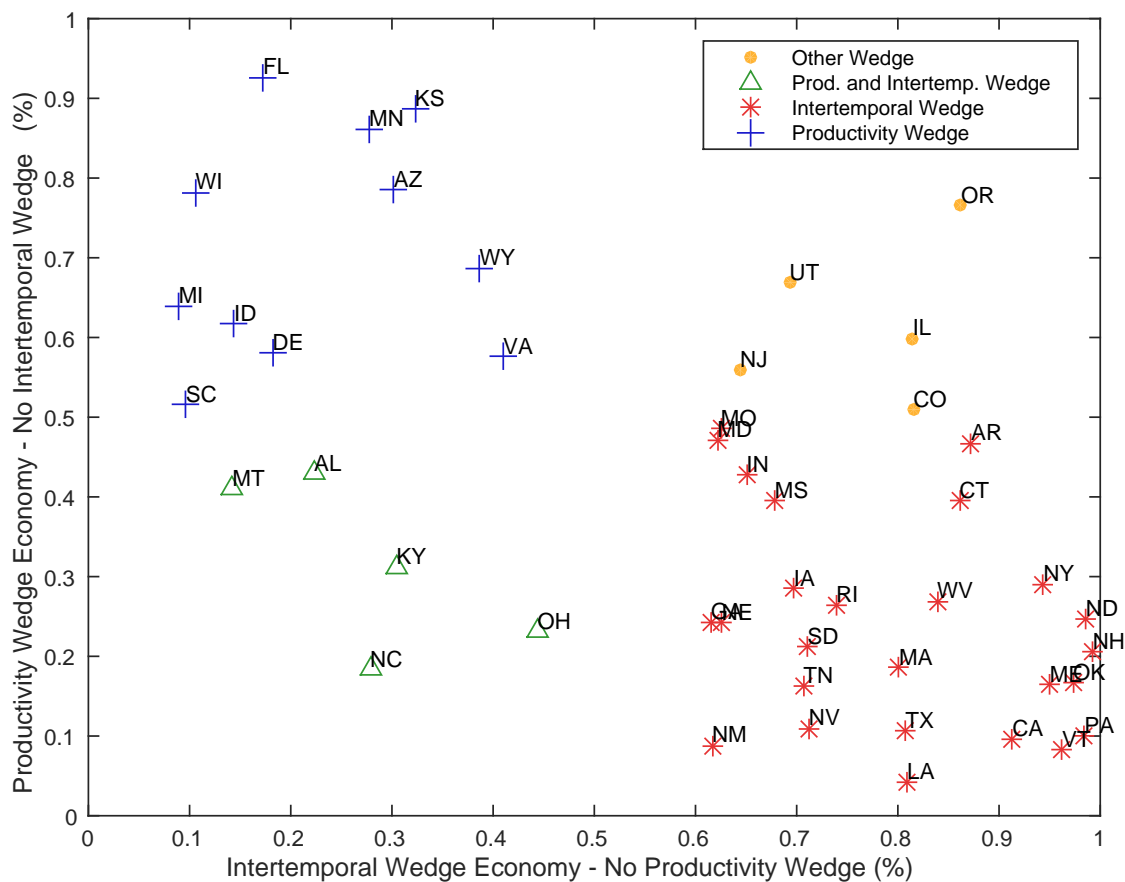
Fips Code	State	Trough Year	Δn	Δn_γ	Δn_{yx}	Δn_ϵ	Peak-Trough Wedge	Wedge 2007-2014
1	Alabama	2009	-5.3	0.2	4.0	-9.6	Labor	Labor
4	Arizona	2009	0.6	0.0	-3.5	4.2	Labor	Productivity/Labor
5	Arkansas	2010	0.8	0.0	0.9	-0.1	Labor/Productivity	Labor
6	California	2010	-2.4	0.0	-3.8	1.4	Intermediate	Intermediate/Productivity
8	Colorado	2010	-0.9	-0.1	-1.8	1.0	Productivity	Intermediate/Productivity
9	Connecticut	2010	0.4	-0.1	-4.5	5.0	Labor	Labor
10	Delaware	2010	-9.2	0.3	4.9	-14.4	Labor	Labor
12	Florida	2009	-2.2	0.3	1.8	-4.3	Labor	Labor/Intermediate
13	Georgia	2010	-6.4	0.2	2.1	-8.7	Labor	Labor
16	Idaho	2009	-3.8	0.2	4.3	-8.3	Labor	Labor
17	Illinois	2009	0.5	0.0	0.6	0.0	Intermediate	Intermediate/Productivity
18	Indiana	2009	-2.3	0.1	2.4	-2.8	Labor	Labor
19	Iowa	2010	4.8	-0.2	5.2	-0.3	Intermediate	Labor
20	Kansas	2010	2.6	-0.2	1.0	1.8	Labor	Labor
21	Kentucky	2010	-1.9	0.1	3.4	-5.4	Labor	Labor
22	Louisiana	2010	4.9	-0.2	1.8	3.3	Labor	Labor
23	Maine	2009	1.5	0.0	3.0	-1.5	Intermediate	Labor
24	Maryland	2010	1.9	-0.2	-0.7	2.8	Labor	Labor
25	Massachussets	2009	4.3	-0.1	-3.0	7.4	Labor	Labor
26	Michigan	2009	-4.9	0.3	4.2	-9.3	Labor	Labor
27	Minnesota	2009	2.1	-0.1	0.8	1.4	Intermediate/Labor	Labor
28	Mississippi	2010	2.2	0.1	5.1	-3.1	Intermediate/Productivity	Undetermined
29	Missouri	2010	1.0	0.0	3.3	-2.3	Intermediate	Intermediate
30	Montana	2009	-7.9	0.2	3.1	-11.3	Labor	Labor
31	Nebraska	2010	7.0	-0.3	1.3	5.9	Labor	Labor
32	Nevada	2010	-13.2	0.4	-0.9	-12.6	Labor	Labor
33	New Hampshire	2010	4.5	0.0	1.5	2.9	Labor/Productivity	Labor
34	New Jersey	2010	2.3	-0.1	0.9	3.1	Labor	Labor
35	New Mexico	2010	0.3	-0.2	2.1	-1.7	Labor	Labor
36	New York	2009	3.2	-0.1	-0.8	4.1	Labor	Labor
37	North Carolina	2009	-3.5	0.2	4.6	-8.3	Labor	Labor
38	North Dakota	2010	5.1	-0.3	2.5	2.8	Labor/Intermediate	Labor
39	Ohio	2009	-1.2	0.1	3.1	-4.4	Labor	Labor
40	Oklahoma	2010	5.7	-0.2	1.0	4.9	Labor	Labor
41	Oregon	2010	0.5	-0.2	0.2	0.5	Labor/Intermediate	Undetermined
42	Pennsylvania	2010	1.2	-0.1	1.0	0.2	Labor	Labor
44	Rhode Island	2010	1.6	0.1	4.0	-2.5	Intermediate/Productivity	Labor
45	South Carolina	2010	-4.8	0.3	6.0	-11.1	Labor	Labor
46	South Dakota	2009	6.2	-0.3	-0.1	6.6	Labor	Labor
47	Tennessee	2009	0.8	0.0	2.0	-1.3	Labor	Labor
48	Texas	2009	3.2	-0.1	-4.7	8.0	Labor	Labor
49	Utah	2010	0.4	-0.2	0.1	0.4	Labor/Intermediate	Intermediate
50	Vermont	2010	7.2	-0.3	1.8	5.7	Labor	Labor
51	Virginia	2009	3.1	-0.1	-0.1	3.2	Labor	Labor
53	Washington	2010	-1.8	-0.1	0.1	-1.8	Labor	Labor/Productivity
54	West Virginia	2011	1.3	-0.1	1.1	0.2	Labor	Labor/Productivity
55	Wisconsin	2009	2.3	0.0	2.3	0.1	Labor/Intermediate	Labor
56	Wyoming	2010	3.1	-0.2	-5.9	9.1	Labor/Intermediate	Labor/Intermediate

Figure 1.21: Decomposition of Consumption One-Wedge Economy: 2007-2014



This figure shows the contribution of each wedge to consumption in the regional model using the ψ -statistic computed over the period 2007-2014. The x-axis corresponds to the statistic for the intertemporal wedge-only economy while the y-axis corresponds to the statistic for the productivity-wedge only economy.

Figure 1.22: **Decomposition of Consumption Three-Wedge Economy: 2007-2014**



This figure shows the contribution of a subset of wedges to consumption in the regional model using the ψ -statistic computed over the period 2007-2014. The x-axis corresponds to the statistic for an economy without productivity wedge fluctuating (a three-wedge economy) while the y-axis corresponds to the statistic for an economy without the intertemporal wedge fluctuating (a three-wedge economy)

Table 1.4: ψ -statistics for Consumption: 2007-2014

This table presents the exact values of the ψ -statistic for the decomposition of consumption state by state. The first 4 columns correspond to the contribution of individual wedges while the last two columns correspond to the contribution of the indicated subset of wedges. The computations correspond to the period 2007-2014.


State Fips Code	State	ϕ_γ^c	ϕ_y^c	ϕ_x^c	ϕ_ϵ^c	$\phi_{\gamma x \epsilon}^c$	$\phi_{x y \epsilon}^c$
1	Alabama	0.61	0.20	0.15	0.03	0.22	0.43
4	Arizona	0.27	0.48	0.13	0.13	0.30	0.79
5	Arkansas	0.50	0.18	0.11	0.21	0.87	0.47
6	California	0.82	0.08	0.05	0.05	0.91	0.10
8	Colorado	0.46	0.25	0.12	0.17	0.82	0.51
9	Connecticut	0.52	0.19	0.10	0.20	0.86	0.39
10	Delaware	0.51	0.30	0.14	0.04	0.18	0.58
12	Florida	0.19	0.50	0.18	0.13	0.17	0.93
13	Georgia	0.74	0.03	0.14	0.08	0.62	0.24
16	Idaho	0.22	0.66	0.06	0.05	0.14	0.62
17	Illinois	0.41	0.26	0.13	0.20	0.81	0.60
18	Indiana	0.62	0.07	0.20	0.10	0.65	0.43
19	Iowa	0.44	0.10	0.23	0.23	0.70	0.28
20	Kansas	0.32	0.27	0.25	0.15	0.32	0.89
21	Kentucky	0.46	0.42	0.08	0.05	0.30	0.31
22	Louisiana	0.86	0.02	0.05	0.08	0.81	0.04
23	Maine	0.74	0.09	0.07	0.10	0.95	0.16
24	Maryland	0.53	0.17	0.15	0.15	0.62	0.47
25	Massachussets	0.67	0.02	0.07	0.24	0.80	0.19
26	Michigan	0.22	0.67	0.07	0.04	0.09	0.64
27	Minnesota	0.32	0.24	0.34	0.09	0.28	0.86
28	Mississippi	0.53	0.24	0.11	0.12	0.68	0.39
29	Missouri	0.38	0.28	0.21	0.13	0.63	0.49
30	Montana	0.32	0.58	0.06	0.04	0.14	0.41
31	Nebraska	0.40	0.03	0.32	0.24	0.63	0.24
32	Nevada	0.53	0.02	0.11	0.34	0.71	0.11
33	New Hampshire	0.66	0.07	0.07	0.19	0.99	0.21
34	New Jersey	0.41	0.33	0.12	0.14	0.64	0.56
35	New Mexico	0.84	0.08	0.05	0.03	0.62	0.09
36	New York	0.60	0.09	0.10	0.22	0.94	0.29
37	North Carolina	0.77	0.13	0.07	0.03	0.28	0.18
38	North Dakota	0.65	0.11	0.09	0.15	0.99	0.25
39	Ohio	0.57	0.30	0.07	0.05	0.44	0.23
40	Oklahoma	0.70	0.06	0.07	0.18	0.97	0.17
41	Oregon	0.32	0.23	0.15	0.30	0.86	0.77
42	Pennsylvania	0.81	0.05	0.06	0.08	0.98	0.10
44	Rhode Island	0.67	0.12	0.12	0.10	0.74	0.26
45	South Carolina	0.20	0.72	0.05	0.03	0.10	0.52
46	South Dakota	0.63	0.05	0.13	0.19	0.71	0.21
47	Tennessee	0.75	0.11	0.07	0.07	0.71	0.16
48	Texas	0.79	0.02	0.08	0.11	0.81	0.11
49	Utah	0.36	0.33	0.14	0.17	0.69	0.67
50	Vermont	0.81	0.03	0.04	0.11	0.96	0.08
51	Virginia	0.47	0.26	0.15	0.12	0.41	0.58
53	Washington	0.25	0.05	0.46	0.24	0.75	0.85
54	West Virginia	0.59	0.19	0.10	0.11	0.84	0.31
55	Wisconsin	0.15	0.73	0.08	0.04	0.11	0.78
56	Wyoming	0.31	0.46	0.13	0.11	0.39	0.69

Table 1.5: Peak to Trough Changes in Consumption and Components and Wedge-based Classification of States

This table shows the year in which the economy of each state wen through the trough of the recession. The next column shows the change in consumption in deviations from aggregate employment in the data between 2007 and 2010, the trough-year for the US economy. The next 3 columns show the change in the wedge-components of consumption. The last two columns show the main wedge driving fluctuations in consumption for each state for the peak-trough period of 2007-2010 and for the entire period 2010-2014.


Fips Code	State	Trough Year	Δc	Δc_γ	Δc_{yx}	Δc_ϵ	Peak-Trough Wedge	Wedge 2007-2014
1	Alabama	2009	1.32	0.95	2.32	-1.76	Intertemporal/Productivity	Intertemporal/Productivity
4	Arizona	2009	-0.10	-0.43	-0.75	0.43	Intertemporal	Productivity
5	Arkansas	2010	-2.20	-1.08	-0.85	0.00	Intertemporal/Productivity	Intertemporal/Productivity
6	California	2010	-3.18	-2.45	-1.00	-0.38	Intertemporal	Intertemporal
8	Colorado	2010	-3.53	-1.80	-0.98	-0.54	Intertemporal/Productivity	Intertemporal/Productivity
9	Connecticut	2010	0.32	-0.20	-0.50	1.21	Intertemporal	Intertemporal
10	Delaware	2010	0.69	0.25	3.16	-2.72	Productivity/Intertemporal	Productivity/Intertemporal
12	Florida	2009	1.20	0.08	0.69	-0.15	Productivity	Productivity
13	Georgia	2010	-1.04	-0.89	1.52	-1.80	Intertemporal	Intertemporal
16	Idaho	2009	1.77	1.09	1.84	-1.55	Intertemporal	Productivity/Intertemporal
17	Illinois	2009	-1.31	-0.59	-0.34	-0.08	Intertemporal	Intertemporal/Productivity
18	Indiana	2009	-0.08	0.02	1.03	-1.02	Intertemporal	Intertemporal
19	Iowa	2010	1.54	2.07	0.12	-0.01	Intertemporal	Intertemporal
20	Kansas	2010	-0.68	0.03	-0.53	0.16	Intertemporal	Productivity
21	Kentucky	2010	2.97	2.09	2.02	-0.90	Intertemporal	Intertemporal/Productivity
22	Louisiana	2010	1.24	1.56	-0.56	0.93	Intertemporal	Intertemporal
23	Maine	2009	3.83	2.74	1.24	-0.11	Intertemporal	Intertemporal
24	Maryland	2010	0.33	0.40	-0.39	0.26	Intertemporal	Intertemporal
25	Massachussets	2009	0.87	0.47	-1.31	1.65	Intertemporal/Labor	Intertemporal
26	Michigan	2009	1.77	0.91	2.71	-1.56	Productivity/Intertemporal	Productivity/Intertemporal
27	Minnesota	2009	-0.49	0.15	-0.40	-0.15	Productivity	Undetermined
28	Mississippi	2010	2.93	1.90	1.22	0.18	Intertemporal	Intertemporal
29	Missouri	2010	1.02	1.19	0.70	-0.40	Intertemporal	Intertemporal
30	Montana	2009	3.21	1.92	3.66	-2.26	Productivity/Intertemporal	Productivity/Intertemporal
31	Nebraska	2010	0.28	1.06	-1.41	1.06	Intertemporal	Intertemporal
32	Nevada	2010	-2.90	-3.28	1.78	-2.48	Intertemporal	Intertemporal
33	New Hampshire	2010	1.96	1.31	-0.59	1.31	Intertemporal	Intertemporal
34	New Jersey	2010	-2.15	-0.84	-1.10	-0.05	Intertemporal/Productivity	Intertemporal/Productivity
35	New Mexico	2010	2.10	1.85	0.98	-0.26	Intertemporal	Intertemporal
36	New York	2009	1.76	1.16	-0.48	1.41	Intertemporal	Intertemporal
37	North Carolina	2009	1.32	0.93	1.86	-1.33	Intertemporal	Intertemporal/Productivity
38	North Dakota	2010	8.47	5.59	2.24	1.45	Intertemporal	Intertemporal
39	Ohio	2009	0.78	0.72	1.14	-0.81	Productivity/Intertemporal	Intertemporal
40	Oklahoma	2010	2.28	1.89	-0.54	1.37	Intertemporal	Intertemporal
41	Oregon	2010	-2.07	-0.45	-0.55	-0.86	Productivity	Undetermined
42	Pennsylvania	2010	1.61	1.20	0.49	0.19	Intertemporal	Intertemporal
44	Rhode Island	2010	1.94	1.48	0.60	-0.03	Intertemporal	Intertemporal
45	South Carolina	2010	3.99	2.29	3.18	-1.41	Productivity/Intertemporal	Productivity/Intertemporal
46	South Dakota	2009	2.41	2.52	-0.58	0.87	Intertemporal	Intertemporal
47	Tennessee	2009	1.35	1.05	0.63	-0.17	Intertemporal	Intertemporal
48	Texas	2009	-0.09	-0.16	-1.28	1.66	Intertemporal	Intertemporal
49	Utah	2010	-3.02	-1.08	-0.96	-0.77	Intertemporal	Intertemporal
50	Vermont	2010	4.16	3.68	-0.40	1.14	Intertemporal/Productivity	Intertemporal/Productivity
51	Virginia	2009	-0.29	0.13	-0.88	0.33	Productivity	Intertemporal/Productivity
53	Washington	2010	-0.05	0.35	0.62	-0.70	Labor Wedge	Labor Wedge
54	West Virginia	2011	4.40	2.96	1.52	0.22	Intertemporal	Intertemporal
55	Wisconsin	2009	0.08	0.45	-0.27	0.08	Productivity	Productivity
56	Wyoming	2010	-2.01	-0.96	-2.31	1.84	Intertemporal	Productivity

Figure 1.23: Sensitivity Analysis - Decomposition of Employment Three-Wedge Economy: 2007-2014. Elasticity of Labor Supply $\phi = 1$



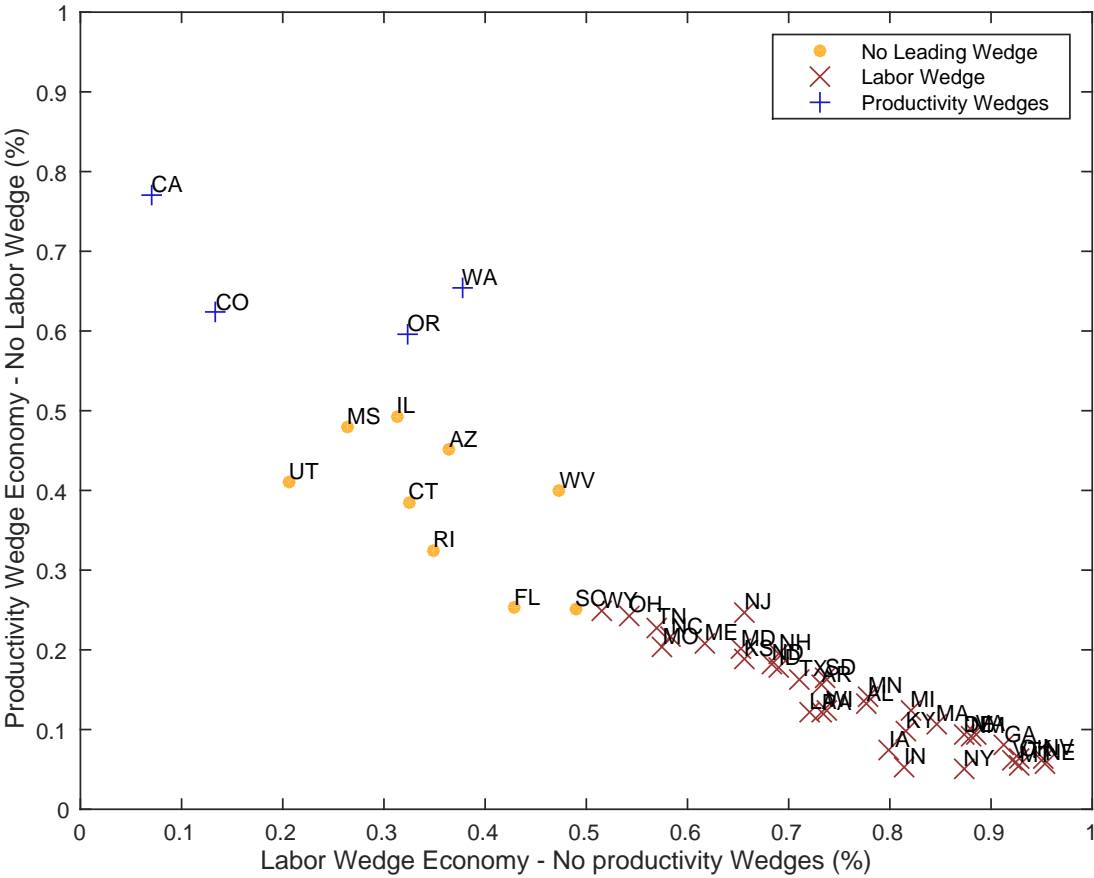
This figure shows the contribution of a subset of wedges to employment in the regional model using the ψ -statistic computed over the period 2007-2014 for a value of the elasticity of labor supply of $\phi = 1$. Recall that in the baseline specification the value was $\phi = 2$. The x-axis corresponds to the statistic for an economy without productivity wedges fluctuating (a two-wedge economy) while the y-axis corresponds to the statistic for an economy without the labor wedge fluctuating (a three-wedge economy)

Figure 1.24: Sensitivity Analysis - Decomposition of Consumption One-Wedge Economy: 2007-2014. Elasticity of Labor Supply $\phi = 1$




This figure shows the contribution of each wedge to consumption in the regional model using the ψ -statistic computed over the period 2007-2014 for a value of the elasticity of labor supply of $\phi = 1$. Recall that in the baseline specification the value was $\phi = 2$. The x-axis corresponds to the statistic for the intertemporal wedge-only economy while the y-axis corresponds to the statistic for the productivity-wedge only economy.

Figure 1.25: Sensitivity Analysis - Decomposition of Employment Three-Wedge Economy: 2007-2014. Elasticity of Labor Supply $\phi = 4$




This figure shows the contribution of a subset of wedges to employment in the regional model using the ψ -statistic computed over the period 2007-2014 for a value of the elasticity of labor supply of $\phi = 4$. Recall that in the baseline specification the value was $\phi = 2$. The x-axis corresponds to the statistic for an economy without productivity wedges fluctuating (a two-wedge economy) while the y-axis corresponds to the statistic for an economy without the labor wedge fluctuating (a three-wedge economy)

Figure 1.26: Sensitivity Analysis - Decomposition of Consumption One-Wedge Economy: 2007-2014. Elasticity of Labor Supply $\phi = 4$




This figure shows the contribution of each wedge to consumption in the regional model using the ψ -statistic computed over the period 2007-2014 for a value of the elasticity of labor supply of $\phi = 4$. Recall that in the baseline specification the value was $\phi = 2$. The x-axis corresponds to the statistic for the intertemporal wedge-only economy while the y-axis corresponds to the statistic for the productivity-wedge only economy.

Figure 1.27: Sensitivity Analysis - Decomposition of Employment Three-Wedge Economy: 2007-2014. Intertemporal Elasticity of Substitution $\sigma = 1$




This figure shows the contribution of a subset of wedges to employment in the regional model using the ψ -statistic computed over the period 2007-2014 for a value of the intertemporal elasticity of substitution of $\sigma = 1$. Recall that in the baseline specification the value was $\sigma = 2$. The x-axis corresponds to the statistic for an economy without productivity wedges fluctuating (a two-wedge economy) while the y-axis corresponds to the statistic for an economy without the labor wedge fluctuating (a three-wedge economy)

Figure 1.29: Sensitivity Analysis - Decomposition of Employment One-Wedge Economy: 2007-2014. GHH Preferences




This figure shows the contribution of each wedge to employment in the regional model using the ψ -statistic computed over the period 2007-2014. The x-axis corresponds to the statistic for the labor wedge-only economy while the y-axis corresponds to the statistic for the productivity-wedge only economy.

Figure 1.30: Sensitivity Analysis - Decomposition of Employment Three-Wedge Economy: 2007-2014. GHH Preferences




This figure shows the contribution of a subset of wedges to employment in the regional model using the ψ -statistic computed over the period 2007-2014. The x-axis corresponds to the statistic for an economy without productivity wedges fluctuating (a two-wedge economy) while the y-axis corresponds to the statistic for an economy without the labor wedge fluctuating (a three-wedge economy)

Figure 1.31: Sensitivity Analysis - Decomposition of Consumption One-Wedge Economy: 2007-2014. GHH Preferences



This figure shows the contribution of each wedge to consumption in the regional model using the ψ -statistic computed over the period 2007-2014. The x-axis corresponds to the statistic for the intertemporal wedge-only economy while the y-axis corresponds to the statistic for the productivity-wedge only economy.

Figure 1.32: Sensitivity Analysis - Decomposition of Consumption Three-Wedge Economy: 2007-2014. GHH Preferences



This figure shows the contribution of a subset of wedges to consumption in the regional model using the ψ -statistic computed over the period 2007-2014. The x-axis corresponds to the statistic for an economy without productivity wedge fluctuating (a three-wedge economy) while the y-axis corresponds to the statistic for an economy without the intertemporal wedge fluctuating (a three-wedge economy)

Table 1.6: **Cross-Region Employment and Wedges**

Panel A. in this table shows regressions of changes in employment at the state level on changes in the different wedges from the model with standard preferences for periods 2007-2010 and 2010-2014. Panel B. presents the results of the same regressions but using wedges that use GHH preferences.

Panel A: Standard Preferences								
	2007-2010				2010-2014			
	(1) b/se	(2) b/se	(3) b/se	(4) b/se	(5) b/se	(6) b/se	(7) b/se	(8) b/se
Δ Investment Wedge	0.001 (0.063)		2.830*** (0.674)		-0.051 (0.046)		0.834* (0.477)	
Δ Productivity Wedge		0.542*** (0.185)	-0.100 (0.222)			0.397*** (0.095)	0.135 (0.167)	
Δ Intermediate Prod. Wedge			0.795*** (0.189)				0.238* (0.129)	
Δ Labor Wedge				-0.295*** (0.024)				-0.216*** (0.020)
constant	0.003 (0.007)	0.003 (0.006)	-0.002 (0.006)	0.009*** (0.003)	-0.008* (0.004)	-0.011*** (0.004)	-0.013*** (0.004)	-0.011*** (0.002)
R^2	0.000	0.157	0.398	0.764	0.026	0.274	0.342	0.719
Observations	48	48	48	48	48	48	48	48

Panel B: GHH Preferences								
	2007-2010				2010-2014			
	(1) b/se	(2) b/se	(3) b/se	(4) b/se	(5) b/se	(6) b/se	(7) b/se	(8) b/se
Δ Investment Wedge	0.003 (0.020)		1.029*** (0.126)		-0.016 (0.014)		0.502*** (0.128)	
Δ Productivity Wedge		0.542*** (0.185)	0.091 (0.132)			0.397*** (0.095)	0.087 (0.112)	
Δ Intermediate Prod. Wedge			0.925*** (0.113)				0.446*** (0.111)	
Δ Labor Wedge				-0.169* (0.099)				-0.048 (0.069)
constant	0.003 (0.007)	0.003 (0.006)	-0.003 (0.004)	0.008 (0.007)	-0.008* (0.004)	-0.011*** (0.004)	-0.011*** (0.003)	-0.007 (0.004)
R^2	0.000	0.157	0.666	0.059	0.028	0.274	0.478	0.010
Observations	48	48	48	48	48	48	48	48

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses

1.11 Appendix

1.11.1 Aggregation

In this appendix I show two aggregation results that are useful to perform the business cycle accounting exercise and measure the investment wedge. The starting point to prove these results is the set of equations that describe the log-linearized equilibrium and that are derived in the online appendix section ??:

$$0 = \mathbb{E}_t [mu_{kt+1} - mu_{kt} - \pi_{kt+1} - \Phi_0 (c_{kt} - c_t) - \gamma_{kt+1}^\omega + \varphi_\pi \mathbb{E}_t [\pi_{t+1}] + \varphi_y (y_t - y^*)] \quad (1.19)$$

$$mu_{kt} = -\sigma C c_{kt} \quad (1.20)$$

$$w_{kt}^r = \lambda \left(\epsilon_{kt}^\omega + \frac{1}{\phi} n_{kt} + \sigma c_{kt} \right) + (1 - \lambda) (w_{kt-1}^r - \pi_{kt}) \quad (1.21)$$

$$0 = z_{kt}^{x\omega} - w_{kt} + q_t - (1 - \theta) n_{kt}^x \quad (1.22)$$

$$N n_{kt} = N^y n_{kt}^y + N^x n_{kt}^x \quad (1.23)$$

$$\sum_k x_{kt} = \sum_k (z_{kt}^x + \theta n_{kt}^x) \quad (1.24)$$

$$w_{kt}^r = w_{kt} - p_{kt} \quad (1.25)$$

$$\pi_{kt} = p_{kt} - p_{kt-1} \quad (1.26)$$

$$b_{kt+1} = (1+r)(b_{kt} + \tilde{i}_t) + \frac{X}{B} (z_{kt}^{x\omega} + \theta n_{kt}^x - x_{kt}) - r\tau_t \quad (1.27)$$

$$x_{kt} = n_{kt}^y + w_{kt} - q_t \quad (1.28)$$

$$c_{kt} = w_{kt}^r + n_{kt}^y \quad (1.29)$$

$$w_{kt}^r = -(1 - (\alpha + \theta\beta)) n_{kt}^y - \beta(1 - \theta)(n_{kt}^x - n_{kt}^y) + z_{kt}^{y\omega} + \beta z_{kt}^{x\omega} \quad (1.30)$$

In the case of GHH preferences, equations (1.20) and (1.21) are replaced by

$$mu_{kt} = -\frac{\sigma}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}\right)} \left(C c_{kt} - N^{\frac{1+\phi}{\phi}} n_{kt}\right) \quad (1.31)$$

$$w_{kt}^r = \lambda \left(\epsilon_{kt}^\omega + \frac{1}{\phi} n_{kt}\right) + (1 - \lambda)(w_{kt-1}^r - \pi_{kt}) \quad (1.32)$$

Aggregation Result 1 - Aggregate Economy Let's define the notation for the aggregate economy with lower-case letters with no k-subscript. For example $n_t \equiv \sum_k \frac{1}{K} n_{tk}$. The first aggregation result is that for the aggregate economy (the one we obtain after adding up across islands) the behavior of π_t^w, w_t^r, n_t in the log-linearized economy is identical to that of a representative economy with only a final goods sector with labor share in production $\alpha + \theta\beta$, no endogenous discount factor, and only 3 exogenous processes $\{z_t^\omega, \epsilon_t^\omega, \gamma_t^\omega\}$, where $z_t^\omega \equiv z_t^{y\omega} + \beta z_t^{x\omega}$. The detailed derivation is in online appendix 1.12.5.1. To get to this result add up equations (1.22), (1.24), and (1.28) to get that $n_t^y = n_t^x$. Using this, equation (1.23) implies $n_t = n_t^x = n_t^y$ and therefore equation (1.29) becomes:

$$c_t = w_t^r + n_t^y = w_t^r + n_t$$

For equations (1.19) (1.20), (1.21), and (1.30) just add up and apply the definition of the aggregate variable, and define $z_t^\omega \equiv z_t^{y\omega} + \beta z_t^{x\omega}$. The resulting equations that describe the behavior of the aggregate economy are shown in section 1.6, and are not shown here.

Aggregation Result 2 - Local Economies Let's define the notation for the local/regional economies. I write the regional economy in deviations from the aggregate and denote the regional variables with tilde, so that $\tilde{x}_t \equiv x_{kt} - x_t$ represents the deviation of variable x in island k from its aggregate counterpart as defined in the previous section. I drop the k subscripts to simplify notation. The second aggregation result is that For given $\{\tilde{z}_t^y, \tilde{z}_t^x, \tilde{\gamma}_t, \tilde{\epsilon}_t\}$, the behavior of $\{\tilde{p}_t, \tilde{w}_t, \tilde{n}_t^y, \tilde{n}_t^x\}$ in the log-linearized economy for each island in deviations from aggregates is identical to that of a small open economy where the price of intermediates and the nominal interest rate are at their steady state levels, i.e. $q_t = i_t = 0 \forall t$.

The resulting system of equations that describes the equilibrium of the economy are shown in section 1.6. This system is in essence identical to the original one setting $i_t = q_t = 0$ and dropping the market clearing condition in the intermediate goods market. To get to this result, take the equations above, add them across islands and then compute deviations from the aggregate. The detailed derivation is in online appendix 1.12.6.

1.11.2 Stochastic Process Estimation

The stochastic process driving the wedges is assumed to have the following form

$$\begin{bmatrix} \gamma_{kt} \\ z_{kt}^y \\ z_{kt}^x \\ \epsilon_{kt} \end{bmatrix} = \begin{bmatrix} \rho_\gamma & \rho_{\gamma y} & \rho_{\gamma x} & \rho_{\gamma \epsilon} \\ \rho_{y\gamma} & \rho_y & \rho_{yx} & \rho_{y\epsilon} \\ \rho_{x\gamma} & \rho_{xy} & \rho_x & \rho_{x\epsilon} \\ \rho_{\epsilon\gamma} & \rho_{\epsilon y} & \rho_{\epsilon x} & \rho_\epsilon \end{bmatrix} \begin{bmatrix} \gamma_{kt-1} \\ z_{kt-1}^y \\ z_{kt-1}^x \\ \epsilon_{kt-1} \end{bmatrix} + \begin{bmatrix} \sigma_\gamma & \sigma_{\gamma y} & \sigma_{\gamma x} & \sigma_{\gamma \epsilon} \\ \sigma_{y\gamma} & \sigma_y & \sigma_{yx} & \sigma_{y\epsilon} \\ \sigma_{x\gamma} & \sigma_{xy} & \sigma_x & \sigma_{x\epsilon} \\ \sigma_{\epsilon\gamma} & \sigma_{\epsilon y} & \sigma_{\epsilon x} & \sigma_\epsilon \end{bmatrix} \begin{bmatrix} u_t^\gamma \\ u_t^y \\ u_t^x \\ u_t^\epsilon \end{bmatrix} + \begin{bmatrix} \tilde{\sigma}_\gamma & \tilde{\sigma}_{\gamma y} & \tilde{\sigma}_{\gamma x} & \tilde{\sigma}_{\gamma \epsilon} \\ \tilde{\sigma}_{y\gamma} & \tilde{\sigma}_y & \tilde{\sigma}_{yx} & \tilde{\sigma}_{y\epsilon} \\ \tilde{\sigma}_{x\gamma} & \tilde{\sigma}_{xy} & \tilde{\sigma}_x & \tilde{\sigma}_{x\epsilon} \\ \tilde{\sigma}_{\epsilon\gamma} & \tilde{\sigma}_{\epsilon y} & \tilde{\sigma}_{\epsilon x} & \tilde{\sigma}_\epsilon \end{bmatrix} \begin{bmatrix} v_{kt}^\gamma \\ v_{kt}^y \\ v_{kt}^x \\ v_{kt}^\epsilon \end{bmatrix}$$

Then, in order to express this only as a function of $z_t = z_{t-1}^y + \beta z_{t-1}^x$ and $u_t^z = u_t^y + \beta u_t^x$ we need a restriction on the parameters. For the γ process, adding up across islands we get:

$$\gamma_t = \rho_\gamma \gamma_{t-1} + \rho_{\gamma y} z_{t-1}^y + \rho_{\gamma x} z_{t-1}^x + \rho_{\gamma \epsilon} \epsilon_{t-1} + \sigma_\gamma u_t^\gamma + \sigma_{\gamma y} u_t^y + \sigma_{\gamma x} u_t^x + \sigma_{\gamma \epsilon} u_t^\epsilon$$

In order to express this only as a function of $z_t = z_{t-1}^y + \beta z_{t-1}^x$ and $u_t^z = u_t^y + \beta u_t^x$ we need a restriction on the parameters. If $\rho_{\gamma x} = \beta \rho_{\gamma y}$ and $\sigma_{\gamma x} = \beta \sigma_{\gamma y}$, and letting $\rho_{\gamma y} \equiv \rho_{\gamma z}$ and $\sigma_{\gamma y} \equiv \sigma_{\gamma z}$ this can be written as:

$$\begin{aligned} \gamma_t &= \rho_\gamma \gamma_{t-1} + \rho_{\gamma y} z_{t-1}^y + \rho_{\gamma y} \beta z_{t-1}^x + \rho_{\gamma \epsilon} \epsilon_{t-1} + \sigma_\gamma u_t^\gamma + \sigma_{\gamma y} u_t^y + \sigma_{\gamma y} \beta u_t^x + \sigma_{\gamma \epsilon} u_t^\epsilon \\ &= \rho_\gamma \gamma_{t-1} + \rho_{\gamma z} z_{t-1} + \rho_{\gamma \epsilon} \epsilon_{t-1} + \sigma_\gamma u_t^\gamma + \sigma_{\gamma z} u_t^z + \sigma_{\gamma \epsilon} u_t^\epsilon \end{aligned}$$

Similarly, for the ϵ process if $\rho_{\epsilon x} = \beta \rho_{\epsilon y}$ and $\sigma_{\epsilon x} = \beta \sigma_{\epsilon y}$ and letting $\rho_{\epsilon y} \equiv \rho_{\epsilon z}$ and $\sigma_{\epsilon y} \equiv \sigma_{\epsilon z}$

$$\epsilon_t = \rho_{\epsilon \gamma} \gamma_{t-1} + \rho_{\epsilon z} z_{t-1} + \rho_{\epsilon \epsilon} \epsilon_{t-1} + \sigma_{\epsilon \gamma} u_t^\gamma + \sigma_{\epsilon z} u_t^z + \sigma_{\epsilon \epsilon} u_t^\epsilon$$

For the z process, after adding up across islands and that $z_t = z_t^y + \beta z_t^x$ we have that:

$$\begin{aligned} z_t &= (\rho_{y\gamma} + \beta\rho_{x\gamma}) \gamma_{t-1} + (\rho_y + \beta\rho_{xy}) z_{t-1}^y + (\rho_{yx} + \beta\rho_x) z_{t-1}^x + (\rho_{y\epsilon} + \beta\rho_{x\epsilon}) \epsilon_{t-1} \\ &+ (\sigma_{y\gamma} + \beta\sigma_{x\gamma}) u_t^\gamma + (\sigma_y + \beta\sigma_{xy}) u_t^y + (\sigma_{yx} + \beta\sigma_x) u_t^x + (\sigma_{y\epsilon} + \beta\sigma_{x\epsilon}) u_t^\epsilon \end{aligned}$$

If $\rho_{xy} = \rho_{yx} = 0$, $\rho_y = \rho_x \equiv \rho_z$, $\sigma_{xy} = \sigma_{yx} = 0$, $\sigma_y = \sigma_x \equiv \sigma_z$ then

$$z_t = (\rho_{y\gamma} + \beta\rho_{x\gamma}) \gamma_{t-1} + \rho_z z_{t-1}^y + \beta\rho_z z_{t-1}^x + (\rho_{y\epsilon} + \beta\rho_{x\epsilon}) \epsilon_{t-1} + (\sigma_{y\gamma} + \beta\sigma_{x\gamma}) u_t^\gamma + \sigma_y u_t^y + \beta\sigma_x u_t^x + (\sigma_{y\epsilon} + \beta\sigma_{x\epsilon}) u_t^\epsilon$$

$$z_t = (\rho_{y\gamma} + \beta\rho_{x\gamma}) \gamma_{t-1} + \rho_z z_t + (\rho_{y\epsilon} + \beta\rho_{x\epsilon}) \epsilon_{t-1} + (\sigma_{y\gamma} + \beta\sigma_{x\gamma}) u_t^\gamma + \sigma_z u_t^z + (\sigma_{y\epsilon} + \beta\sigma_{x\epsilon}) u_t^\epsilon$$

$$z_t = \rho_z \gamma_{t-1} + \rho_z z_t + \rho_z \epsilon_{t-1} + \sigma_z \gamma u_t^\gamma + \sigma_z u_t^z + \sigma_z \epsilon u_t^\epsilon$$

Therefore, under the proposed set-up the following processes are the ones I estimate:

$$\begin{bmatrix} \gamma_{kt} \\ z_{kt}^y \\ z_{kt}^x \\ \epsilon_{kt} \end{bmatrix} = \begin{bmatrix} \rho_\gamma & \rho_{\gamma z} & \beta\rho_{\gamma z} & \rho_{\gamma\epsilon} \\ \rho_{y\gamma} & \rho_z & 0 & \rho_{y\epsilon} \\ \rho_{x\gamma} & 0 & \rho_z & \rho_{x\epsilon} \\ \rho_{\epsilon\gamma} & \rho_{\epsilon z} & \beta\rho_{\epsilon z} & \rho_\epsilon \end{bmatrix} \begin{bmatrix} \gamma_{kt-1} \\ z_{kt-1}^y \\ z_{kt-1}^x \\ \epsilon_{kt-1} \end{bmatrix} + \begin{bmatrix} \sigma_\gamma & \sigma_{\gamma z} & \beta\sigma_{\gamma z} & \sigma_{\gamma\epsilon} \\ \sigma_{y\gamma} & \sigma_z & 0 & \sigma_{y\epsilon} \\ \sigma_{x\gamma} & 0 & \sigma_z & \sigma_{x\epsilon} \\ \sigma_{\epsilon\gamma} & \sigma_{\epsilon z} & \beta\sigma_{\epsilon z} & \sigma_\epsilon \end{bmatrix} \begin{bmatrix} u_t^\gamma \\ u_t^y \\ u_t^x \\ u_t^\epsilon \end{bmatrix} + \tilde{\Sigma} \begin{bmatrix} \tilde{v}_{kt}^\gamma \\ \tilde{v}_{kt}^y \\ \tilde{v}_{kt}^x \\ \tilde{v}_{kt}^\epsilon \end{bmatrix}$$

which for the aggregate economy means that I estimate

$$\begin{bmatrix} \gamma_t \\ z_t \\ \epsilon_t \end{bmatrix} = \begin{bmatrix} \rho_\gamma & \rho_{\gamma z} & \rho_{\gamma\epsilon} \\ \rho_{y\gamma} + \beta\rho_{x\gamma} & \rho_z & \rho_{y\epsilon} + \beta\rho_{x\epsilon} \\ \rho_{\epsilon\gamma} & \rho_{\epsilon z} & \rho_\epsilon \end{bmatrix} \begin{bmatrix} \gamma_{t-1} \\ z_{t-1} \\ \epsilon_{t-1} \end{bmatrix} + \begin{bmatrix} \sigma_\gamma & \sigma_{\gamma z} & \sigma_{\gamma\epsilon} \\ \sigma_{y\gamma} + \beta\sigma_{x\gamma} & \sigma_z & \sigma_{y\epsilon} + \beta\sigma_{x\epsilon} \\ \sigma_{\epsilon\gamma} & \sigma_{\epsilon z} & \sigma_\epsilon \end{bmatrix} \begin{bmatrix} u_t^\gamma \\ u_t^z \\ u_t^\epsilon \end{bmatrix}$$

Notice the relationship between the persistence parameters in the original process, which correspond to the parameters of the regional model, and the parameters for the aggregate model.

1.11.3 Regional Price Data

Table 1.7: Nielsen's Retail Scanner Database Description

This table provides some basic descriptive statistics of the AC Nielsen database used to construct regional price indexes

	Individual Years								Combined	
	2006	2007	2008	2009	2010	2011	2012	2013	Total	Average
Number of Obs. (million)	12,013.1	12,812.2	13,037.5	12,968.3	13,153.4	13,646.7	13,618.8	13,801.3	105,051.0	13,131.4
Number of UPCs	725,224	762,469	759,989	753,984	739,768	742,074	753,318	769,136	1,487,003	750,745
Number of Categories	1,085	1,086	1,086	1,083	1,085	1,081	1,105	1,113	1,113	1,091
Number of Chains	86	85	87	86	86	86	82	79	88	85
Number of Stores	32,642	33,745	34,830	35,343	35,807	35,645	36,059	36,316	40,350	35,048
Number of Zip Codes	10,869	11,123	11,357	11,476	11,589	11,639	11,626	11,553	11,797	11,404
Number of Counties	2,385	2,468	2,500	2,508	2,519	2,526	2,547	2,561	2,593	2,502
Number of MSAs	361	361	361	361	361	361	361	361	361	361
Number of States	49	49	49	49	49	49	49	49	49	49
Transaction Value (US billion)	187.9	207.8	219.6	223.7	227.6	235.2	239.5	238.7	1,779.9	222.5
Pct. Value used in Price Index	54.3%	50.0%	66.4%	66.0%	68.3%	68.0%	67.7%	67.2%	63.9%	63.5%

1.11.4 *Computing the Value of the Housing Stock*

In this appendix section I explain how I compute the value of the housing stock and what data I use to this purpose. I follow the approach by Mian et al. (2013). The strategy is to take the 2000 median home value, adjust it for every year using local house price indexes for every geography and apply it to the housing stock owned by households computed as $Housing_Units \times Occupancy_rate \times Homeownership_rate$. Table 1.8 lists the variables I use and the sources. First, I need a measure of the housing stock owned by households. From the Census 2000 we can get directly a measure of the number of housing units at the state and county level. To obtain MSA-level figures I aggregate the county-level figures.⁵⁷ Then using the intercensal population estimates from 2001 through 2013 I compute annual growth rates in housing units and apply them to the number of housing units in the Census 2000. In this way I get an estimate of the housing units year by year. For each of the three geographies I check how close the number of housing units obtained in this way is to the number of housing units in the Census 2010. At the state level the average absolute percentage difference between the two numbers is 0.29%, at the county level is 0.325% and at the MSA level is 3.3%. The difference at the MSA level is driven by a few outliers. The explanation may lie in different definitions of the MSA in between the two censuses. In these cases I replace the value of the housing units for 2010 computed from the county data of 2000 by the units from the Census 2010. To get to the units in other years I use the changes in housing units from the ACS estimates between 2005 and 2013.

⁵⁷. The walkover files linking MSAs and counties are here <http://www.census.gov/population/metro/data/def.html>

Table 1.8: **Data Sources to Value the Housing Stock**

This table provides the sources and different variables I use at different levels of aggregation to compute the value of the housing stock

Geography	Variable	Years	Dataset	Source	Link or Table
State, County	Housing Units	2000	Census 2000	Census Bureau	Table H001 - Housing Units
State, County, MSA	Housing Units	2010	Census 2010	Census Bureau	Table H1 - Housing Units
State, County	Housing Units	2001-2009	Population Estimates	Census Bureau	Housing Units Intercensal Estimates (1)
State, County	Housing Units	2011-2013	Population Estimates	Census Bureau	PEPANNHU
MSA	Housing Units	2005-2013	ACS	Census Bureau	B25001 - Housing Units
State, County	Occupancy Rate	2000	Census 2000	Census Bureau	Table H003 - Occupancy Status
State, County, MSA	Occupancy Rate	2010	Census 2010	Census Bureau	Table H3 - Occupancy Status
State, County, MSA	Occupancy Rate	2005-2013	ACS	Census Bureau	Table B25002 - Occupancy Status
State, County	Homeownership Rate	2000	Census 2000	Census Bureau	Table H004 - Tenure (2)
State, County, MSA	Homeownership Rate	2010	Census 2010	Census Bureau	Table H4 - Tenure
State, County, MSA	Occupancy Rate	2005-2013	ACS	Census Bureau	Table B25003 - Tenure
State, County	Median Home Value	2000	Census 2000	Census Bureau	Table H085 - Median Value
MSA	Median Home Value	2010	ACS	Census Bureau	Table B25077 - Median Value
State and MSA	House Prices	2000-2013	House Price Indexes	FHFA	Quarterly All Transaction Indexes(3)
County	House Prices	2000-2013	Median Home Value	Zillow	ZHVI All Homes(4)

(1) Download at <http://www.census.gov/popest/data/intercensal/housing/hu2010.html>

(2) Alternatively available at <http://www.census.gov/housing/hvs/data/ann14ind.html>

(3) Download at <http://www.fhfa.gov/DataTools/Downloads/Pages/House-Price-Index.aspx>

(4) Download at <http://www.zillow.com/research/data/#median-home-value>

The second step is to get a measure of the household units owned by the household. Fortunately, the US Census Bureau collects data and provide estimates of the homeownership,

which it is defined as follows:⁵⁸

$$\text{Homeownership_rate} = \frac{\text{Owner_occupied_housing_units}}{\text{Total_occupied_housing_Units}} * 100$$

The US Census Bureau also provides information on the vacancy rates for owner occupied units. For each geographic level (state, county, or MSA) I compute vacancy rates and homeownership rates using data from the American Community Survey (ACS) from 2005 through 2013 and the US census of 2010. Using the measures of the Census 2010 for these two rates as the base, I compute the homeownership rates and vacancy rates by applying changes in the same rates computed from the ACS. In this way I guarantee that the better measure (the one from the Census) is built in the computations. With the number of units, the vacancy rates, and the homeownership rates I compute each year the housing stock (measured in units) as the product of these three quantities.

The Census 2000 reports a value for the median home for both states and counties. Unfortunately the 2010 census does not report values for any geography. For MSAs, I compute the median value in 2000 as the house-unit weighted average of the counties that make up the MSA. For those MSAs for which the base housing units are those of the Census 2010 and not those of the Census 2000, I use the ACS 2010 median value as the base value. Then using market house prices from the Federal Housing Finance Agency, Zillow and/or Corelogic, I compute median house values by year for the different levels of geographic aggregation.

58. Source: http://quickfacts.census.gov/qfd/meta/long_HSG445213.htm

Finally, to get a dollar measure of the housing stock I multiply the median dollar home values by the household-owned stock of housing units.

1.11.5 List of Data Sets


Given the large number of datasets and data sources that I employ in the paper, as a guide for the interested reader in this section I list them along with their use

Table 1.9: **Data Sources**

Dataset	Source	Use
American Community Survey Micro Data	IPUMS	Regional wage indexes, hours worked, employment shares
Current Population Survey Micro Data	IPUMS	Aggregate wage indexes, measures related to search frictions
Quarterly Census of Employment & Wages - QCEW	Bureau of Labor Statistics	Local wages
Equifax Consumer Trends	Equifax - Fama Miller Center	Value housing stock, Measures of household leverage
State-level debt	New York Fed	Value of housing stock
Flow of Funds	Federal Reserve	Value of net assets
IRS Statistics of Income (SOI)	Internal Revenue Service	Value of Assets
American Community Survey	US Census Bureau	Value of housing stock
Census 2000	US Census Bureau	Value of housing stock
Census 2010	US Census Bureau	Value of housing stock, Regional population
Population Estimates Program	US Census Bureau	Value of housing stock, Regional population
All-transaction House Price Indexes	Federal Housing Finance Agency	Value of housing stock
Nielsen Retail Scanner Data	Nielsen - University of Chicago	Regional price indexes
All Urban Consumer Price Index	Bureau of Labor Statistics	Aggregate price index
Local Area Unemployment Statistics (LAUS)	Bureau of Labor Statistics	Regional employment, labor market tightness
Current Employment Statistics	Bureau of Labor Statistics	Aggregate hours worked
Labor Force Statistics - CPS	Bureau of Labor Statistics	Aggregate employment
Personal Consumption Expenditures by State	Bureau of Economic Analysis	State-level consumption
National Income and Product Account (NIPA)	Bureau of Economic Analysis	Aggregate and State-level consumption
Consumer Expenditure Survey	Bureau of Labor Statistics	State-level consumption
State and MSA Union Membership	www.unionstats.com	Unionization rates
HUD USPS ZIP Code Crosswalk Files	US Dept. of Housing and Urban Develop.	Household Leverage
Metropolitan and Micropolitan Delineations	US Census Bureau	MSA computations based on county data
County Business Patterns	US Census Bureau	Alternative county-level wage measures
IRS Migration Data	Internal Revenue Service	State-level net migration rates


1.11.6 County, MSA, and Other Local Data

Figure 1.33: MSA Level Wage Growth vs Employment Growth: 2007-2010




This figure shows a scatter plot of changes in wages at the state level against changes in employment over the 2007-2010 time period. Wages are computed as in section 1.5. The geographic unit of observation are states.

Figure 1.34: County Level Wage Growth vs Employment Growth: 2007-2010




This figure shows a scatter plot of changes in wages against changes in employment over the 2007-2010 time period. Wages are computed as in section 1.5. The geographic units of observation are counties.

Figure 1.35: MSA Level Price Growth vs Employment Growth: 2007-2010




This figure shows a scatter plot of changes in local prices against changes in employment over the 2007-2010 time period. The geographic units of observation are MSAs. Prices and employment are computed as in section 1.5.

Figure 1.36: County Level Price Growth vs Employment Growth: 2007-2010




This figure shows a scatter plot of changes in local prices against changes in employment over the 2007-2010 time period. The geographic units of observation are counties. Prices and employment are computed as in section 1.5.

Figure 1.37: State Level Real Wage Growth vs Employment Growth: 2007-2010




This figure shows a scatter plot of changes in local real wages against changes in employment over the 2007-2010 time period. The geographic units of observation are states. Wages, prices and employment are computed as in section 1.5.

Figure 1.38: MSA Level Real Wage Growth vs Employment Growth: 2007-2010



This figure shows a scatter plot of changes in local real wages against changes in employment over the 2007-2010 time period. The geographic units of observation are MSAs. Wages, prices and employment are computed as in section 1.5.


Figure 1.39: County Level Real Wage Growth vs Employment Growth:
2007-2010



This figure shows a scatter plot of changes in local real wages against changes in employment over the 2007-2010 time period. The geographic units of observation are counties. Wages, prices and employment are computed as in section 1.5.


1.11.7 Appendix Tables and Figures

Figure A1: Sensitivity Analysis - Decomposition of Employment Three-Wedge Economy: 2007-2014. Intertemporal Elasticity of Substitution $\sigma = 4$




This figure shows the contribution of a subset of wedges to employment in the regional model using the ψ -statistic computed over the period 2007-2014 for a value of the intertemporal elasticity of substitution of $\sigma = 4$. Recall that in the baseline specification the value was $\sigma = 2$. The x-axis corresponds to the statistic for an economy without productivity wedges fluctuating (a two-wedge economy) while the y-axis corresponds to the statistic for an economy without the labor wedge fluctuating (a three-wedge economy)

Figure A2: Sensitivity Analysis - Decomposition of Consumption One-Wedge Economy: 2007-2014. Intertemporal Elasticity of Substitution $\sigma = 4$




This figure shows the contribution of each wedge to consumption in the regional model using the ψ -statistic computed over the period 2007-2014 for a value of the intertemporal elasticity of substitution of $\sigma = 4$. Recall that in the baseline specification the value was $\sigma = 2$. The x-axis corresponds to the statistic for the intertemporal wedge-only economy while the y-axis corresponds to the statistic for the productivity-wedge only economy.

Figure A3: Sensitivity Analysis - Decomposition of Employment Three-Wedge Economy: 2007-2014. Intertemporal Elasticity of Substitution $\sigma = 4$ and Elasticity of Labor Supply $\phi = 1$



This figure shows the contribution of a subset of wedges to employment in the regional model using the ψ -statistic computed over the period 2007-2014 for a value of the elasticity of labor supply of $\phi = 1$ and a value of the intertemporal elasticity of substitution of $\sigma = 4$. Recall that in the baseline specification the values were $\phi = 2$ $\sigma = 2$. Recall that in the baseline specification the value was $\phi = 2$. The x-axis corresponds to the statistic for an economy without productivity wedges fluctuating (a two-wedge economy) while the y-axis corresponds to the statistic for an economy without the labor wedge fluctuating (a three-wedge economy)

Figure A4: Sensitivity Analysis - Decomposition of Consumption One-Wedge Economy: 2007-2014. Intertemporal Elasticity of Substitution $\sigma = 4$ and Elasticity of Labor Supply $\phi = 1$



This figure shows the contribution of each wedge to consumption in the regional model using the ψ -statistic computed over the period 2007-2014 for a value of the elasticity of labor supply of $\phi = 1$ and a value of the intertemporal elasticity of substitution of $\sigma = 4$. Recall that in the baseline specification the values were $\phi = 2$ $\sigma = 2$. Recall that in the baseline specification the value was $\phi = 2$. The x-axis corresponds to the statistic for the intertemporal wedge-only economy while the y-axis corresponds to the statistic for the productivity-wedge only economy.

1.12 On-Line Appendix


This document is the on-line appendix to the paper *Regional Business Cycle Accounting and The Great Recession*. The purpose of this document is fivefold:

1. Explain in detail how to solve the model of the paper. The model is written in a way that one can solve two separate models, a model for the aggregate economy, and a model for the regional economies expressed in deviations from the aggregate. The solution to the model is used in two parts. First to estimate the persistence parameters of the exogenous process of the model. Second, to recover these processes from the data, construct the wedges, and perform the business cycle accounting exercise
2. Provide a guide to build the code for the empirical estimation of the model
3. Provide additional elements of the paper that are left out from the main text but that were key in the development of the project
4. Facilitate replicability.
5. Add some plots and tables that are not part of the main text

In this online appendix, the equations that are numbered adding an OA are numbered within the appendix, while equations numbered without the OA particle come from the main text and have the same number as in the main text.

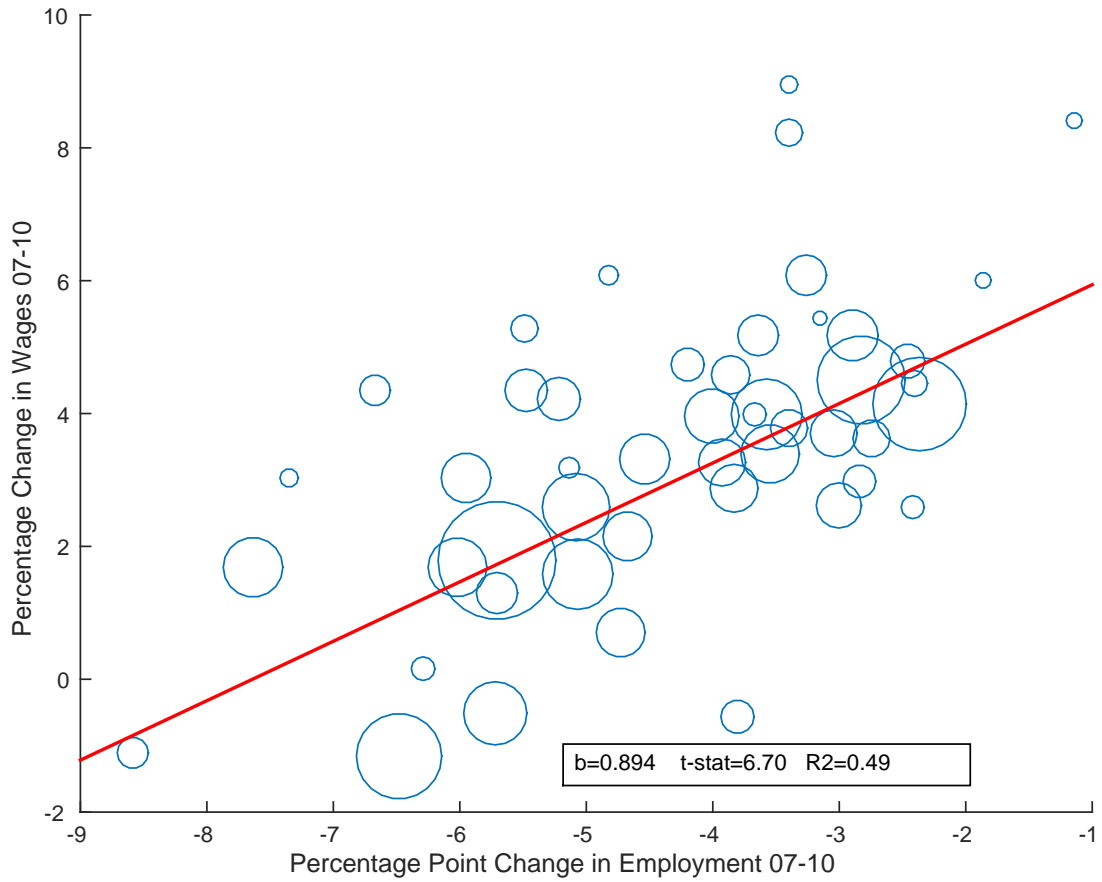
1.12.1 Additional Data Plots

Figure A5: State Level Consumption Growth vs Employment Rate Change: 2007-2010




This figure shows a scatter plot of changes in consumption against changes in employment rate over the 2007-2010 time period.

Figure A6: State Level Wage Growth vs Employment Rate Change: 2007-2010




This figure shows a scatter plot of changes in wages against changes in the employment rate over the 2007-2010 time period. Wages are computed as in section 1.5.

Figure A7: MSA Level Wage Growth vs Employment Rate Change: 2007-2010




This figure shows a scatter plot of changes in wages against changes in employment over the 2007-2010 time period for MSAs across the United States. Wages are computed as in section 1.5.

Figure A8: County Level Wage Growth vs Employment Rate Change: 2007-2010




This figure shows a scatter plot of changes in wages against changes in employment over the 2007-2010 time period. Wages are computed as in section 1.5. The geographic unit of observation are counties

Figure A9: State Level Price Growth vs Employment Rate Change: 2007-2010




This figure shows a scatter plot of changes in local prices against changes in the employment rate over the 2007-2010 time period. The geographic units of observation are states. Prices and the employment rate are computed as in section 1.5.

Figure A10: MSA Level Price Growth vs Employment Rate Change: 2007-2010



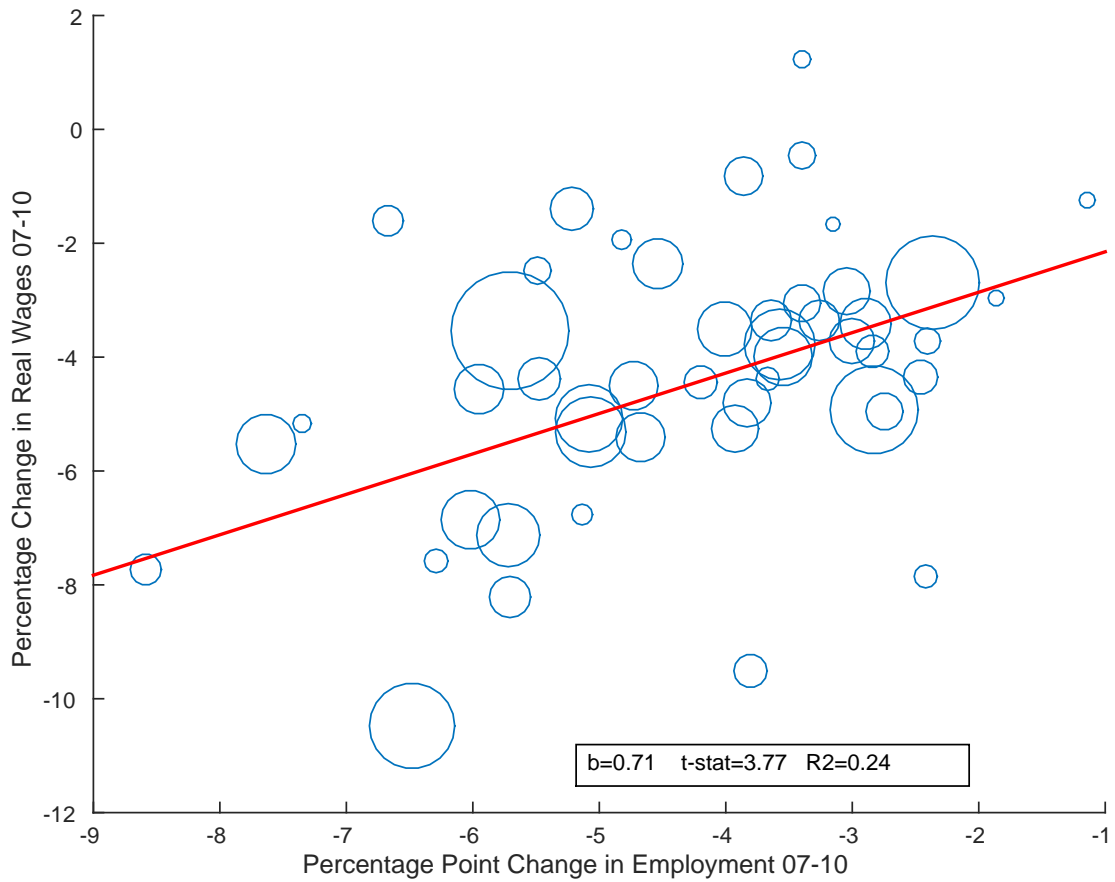
This figure shows a scatter plot of changes in local prices against changes in the employment rate over the 2007-2010 time period. The geographic units of observation are MSAs. Prices and the employment rate are computed as in section 1.5.

Figure A11: County Level Price Growth vs Employment Rate Change:
2007-2010




This figure shows a scatter plot of changes in local prices against changes in the employment rate over the 2007-2010 time period. The geographic units of observation are counties. Prices and the employment rate are computed as in section 1.5.

Figure A12: State Level Real Wage Growth vs Employment Rate Change: 2007-2010




This figure shows a scatter plot of changes in local real wages against changes in employment rate over the 2007-2010 time period. The geographic units of observation are states. Wages, prices and employment are computed as in section 1.5.

Figure A13: MSA Level Real Wage Growth vs Employment Rate Change: 2007-2010




This figure shows a scatter plot of changes in local real wages against changes in employment rate over the 2007-2010 time period. The geographic units of observation are MSAs. Wages, prices and employment are computed as in section 1.5.

Figure A14: County Level Real Wage Growth vs Employment Rate Change: 2007-2010



This figure shows a scatter plot of changes in local real wages against changes in employment rate over the 2007-2010 time period. The geographic units of observation are counties. Wages, prices and employment are computed as in section 1.5.

Figure A15: Net Asset Growth vs Employment Rate Change: 2007-2010




This figure shows a scatter plot of changes in local net assets against changes in the employment rate over the 2007-2010 time period. The geographic units of observation are states. Net assets and employment are computed as in section 1.5.

1.12.2 Alternative Business Cycle Accounting Metrics


In this section I present additional metrics to measure the prominence of different wedges in accounting for the fluctuations of employment and consumption. I present these metrics for different combination of parameter values.

Figure A16: Metric 1- Average Distance from Zero for Employment. Baseline Parameters




This figure shows a scatter plot of the average distance from zero of the labor-wedge component of employment against the average distance from zero of the productivity-wedge component of employment. Each dot represents a state. Values below the 45 degree line represent states for which according to the metric the labor wedge played a more prominent role than the productivity wedge.

Figure A17: Metric 2 - Movement towards Zero for Employment. Baseline Parameters




This figure shows a scatter plot of the movements toward zero of the counterfactual of employment when the labor-wedge is removed against the counterfactual of employment when the productivity wedge is removed. Each dot represents a state. Values below the 45 degree line represent states for which according to the metric the labor wedge played a more prominent role than the productivity wedge.

Figure A18: Metric 3 - Variance Decomposition for Employment. Baseline Parameters




This figure shows a scatter plot of the fraction of the data accounted for the labor-wedge component of employment against fraction of the data accounted for the productivity-wedge component of employment. Each dot represents a state. Values below the 45 degree line represent states for which according to the metric the labor wedge played a more prominent role than the productivity wedge.

Figure A19: Metric 1- Average Distance from Zero for Consumption. Baseline Parameters




This figure shows a scatter plot of the average distance from zero of the intertemporal-wedge component of consumption against the average distance from zero of the productivity-wedge component of consumption. Each dot represents a state. Values below the 45 degree line represent states for which according to the metric the intertemporal wedge played a more prominent role than the productivity wedge.

Figure A21: Metric 3 - Variance Decomposition for Consumption. Baseline Parameters




This figure shows a scatter plot of the fraction of the data accounted for the intertemporal-wedge component of consumption against fraction of the data accounted for the productivity-wedge component of consumption. Each dot represents a state. Values below the 45 degree line represent states for which according to the metric the intertemporal wedge played a more prominent role than the productivity wedge.

Figure A22: Metric 3 - Variance Decomposition for Employment. Parameters
 $(\sigma = 1, \phi = 2)$




This figure shows a scatter plot of the fraction of the data accounted for the labor-wedge component of employment against fraction of the data accounted for the productivity-wedge component of employment. Each dot represents a state. Values below the 45 degree line represent states for which according to the metric the labor wedge played a more prominent role than the productivity wedge. Here the parameter values are changed from the baseline values of $(\sigma = 2, \phi = 2)$ to the values of $(\sigma = 1, \phi = 2)$

Figure A23: Metric 3 - Variance Decomposition for Employment. Parameters
 $(\sigma = 2, \phi = 1)$




This figure shows a scatter plot of the fraction of the data accounted for the labor-wedge component of employment against fraction of the data accounted for the productivity-wedge component of employment. Each dot represents a state. Values below the 45 degree line represent states for which according to the metric the labor wedge played a more prominent role than the productivity wedge. Here the parameter values are changed from the baseline values of $(\sigma = 2, \phi = 2)$ to the values of $(\sigma = 2, \phi = 1)$

Figure A24: Metric 3 - Variance Decomposition for Consumption. Parameters
 $(\sigma = 1, \phi = 2)$



This figure shows a scatter plot of the fraction of the data accounted for the intertemporal-wedge component of consumption against fraction of the data accounted for the productivity-wedge component of consumption. Each dot represents a state. Values below the 45 degree line represent states for which according to the metric the intertemporal wedge played a more prominent role than the productivity wedge. Here the parameter values are changed from the baseline values of $(\sigma = 2, \phi = 2)$ to the values of $(\sigma = 1, \phi = 2)$

Figure A25: Metric 3 - Variance Decomposition for Consumption. Parameters
 $(\sigma = 2, \phi = 1)$



This figure shows a scatter plot of the fraction of the data accounted for the intertemporal-wedge component of consumption against fraction of the data accounted for the productivity-wedge component of consumption. Each dot represents a state. Values below the 45 degree line represent states for which according to the metric the intertemporal wedge played a more prominent role than the productivity wedge. Here the parameter values are changed from the baseline values of $(\sigma = 2, \phi = 2)$ to the values of $(\sigma = 2, \phi = 1)$

1.12.3 *Solution to Regional Model*

In the following two sections I explain in detail how to solve the baseline model of section 1.3 of the paper. This section presents the solution to the regional portion (in deviations from the aggregate), which I call the regional model, while the next section presents the solution to the aggregate portion, which I call the aggregate model. The model here allows for more general processes for the shocks, in particular, I allow for them to be correlated with each other.

The goal is to find a solution to the system of equations that describe each model; more specifically, for a given model we want to find expressions of the endogenous variables as a linear function of the exogenous processes. In the case of the regional model, the solution will be of the form

$$[\tilde{w}_t, \tilde{b}_t, \tilde{p}_t, \tilde{n}_t^y]' = \tilde{P} [\tilde{w}_{t-1}, \tilde{b}_{t-1}] + \tilde{Q} [\tilde{\gamma}_t^\omega, \tilde{z}_t^{y\omega}, \tilde{z}_t^{x\omega}, \tilde{\epsilon}_t^\omega]'$$

where \tilde{P} and \tilde{Q} are matrices whose elements are functions of the parameters of the model.

The goal of this section is to explain in detail how to find these two matrices.

1.12.3.1 Standard Preferences

Recall from section 1.6 that the log-linearized equilibrium conditions for the regional model are:

$$0 = \mathbb{E}_t [\tilde{m}u_{kt+1} - \tilde{m}u_{kt} - (\tilde{p}_{t+1} - \tilde{p}_t) - \Phi_0 (\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y) - \tilde{\gamma}_{t+1}^\omega] \quad (1.13)$$

$$\tilde{m}u_t = -\sigma C (\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y) \quad (1.14)$$

$$\tilde{w}_t = \tilde{z}_t^{x\omega} - (1 - \theta) \tilde{n}_t^x \quad (1.15)$$

$$\tilde{w}_t = \lambda \left(\tilde{p}_t + \tilde{\epsilon}_t^\omega + \frac{1}{\phi} \left(\frac{N^y}{N} \tilde{n}_t^y + \frac{N^x}{N} \tilde{n}_t^x \right) + \sigma (\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y) \right) + (1 - \lambda) \tilde{w}_{t-1} \quad (1.16)$$

$$\tilde{b}_t = (1 + r) \tilde{b}_{t-1} + \frac{X}{B} (\tilde{n}_t^x - \tilde{n}_t^y) \quad (1.17)$$

$$\tilde{w}_t = \tilde{p}_t - (1 - (\alpha + \theta\beta)) \tilde{n}_t^y - \beta (1 - \theta) (\tilde{n}_t^x - \tilde{n}_t^y) + \tilde{z}_t^{y\omega} + \beta \tilde{z}_t^{x\omega} \quad (1.18)$$

$$\lambda \rightarrow 1$$

where the wedges follow the following stochastic processes

$$\begin{bmatrix} \tilde{\gamma}_t^\omega \\ \tilde{z}_t^{y\omega} \\ \tilde{z}_t^{x\omega} \\ \tilde{\epsilon}_t^\omega \end{bmatrix} = \begin{bmatrix} \rho_\gamma & \rho_{\gamma y} & \rho_{\gamma x} & \rho_{\gamma \epsilon} \\ \rho_{y\gamma} & \rho_y & \rho_{yx} & \rho_{y\epsilon} \\ \rho_{x\gamma} & \rho_{xy} & \rho_x & \rho_{x\epsilon} \\ \rho_{\epsilon\gamma} & \rho_{\epsilon y} & \rho_{\epsilon x} & \rho_\epsilon \end{bmatrix} \begin{bmatrix} \tilde{\gamma}_{t-1}^\omega \\ \tilde{z}_{t-1}^{y\omega} \\ \tilde{z}_{t-1}^{x\omega} \\ \tilde{\epsilon}_{t-1}^\omega \end{bmatrix} + \begin{bmatrix} \tilde{\sigma}_\gamma & \tilde{\sigma}_{\gamma y} & \tilde{\sigma}_{\gamma x} & \tilde{\sigma}_{\gamma \epsilon} \\ \tilde{\sigma}_{y\gamma} & \tilde{\sigma}_y & \tilde{\sigma}_{yx} & \tilde{\sigma}_{y\epsilon} \\ \tilde{\sigma}_{x\gamma} & \tilde{\sigma}_{xy} & \tilde{\sigma}_x & \tilde{\sigma}_{x\epsilon} \\ \tilde{\sigma}_{\epsilon\gamma} & \tilde{\sigma}_{\epsilon y} & \tilde{\sigma}_{\epsilon x} & \tilde{\sigma}_\epsilon \end{bmatrix} \begin{bmatrix} \tilde{v}_t^\gamma \\ \tilde{v}_t^y \\ \tilde{v}_t^x \\ \tilde{v}_t^\epsilon \end{bmatrix}$$

1.12.3.2 Algebra Work

In this section I drop the tildes and the ω superscripts from the wedge processes to simplify notation. Let's rewrite the system only in terms of two control variables $[p_t^w, n_t^y]$ and two endogenous state variables $[w_t, b_t]$. Eliminate n_t^x by using equation equation (1.15):

$$\tilde{w}_t = \tilde{z}_t^x - (1 - \theta) \tilde{n}_t^x \implies n_t^x = -\frac{1}{1 - \theta} w_t + \frac{1}{1 - \theta} z_t^x \quad (1.15')$$

Euler Equation

The Euler Equation (equation (1.13)) is given by

$$0 = \mathbb{E}_t \left[\tilde{m} u_{kt+1} - \tilde{m} u_{kt} - (\tilde{p}_{t+1} - \tilde{p}_t) - \Phi_0 (\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y) - \tilde{\gamma}_{t+1} \right]$$

Replacing in the expression for marginal utility (1.14)

$$\begin{aligned} 0 &= \mathbb{E}_t \left[-\sigma C \left(w_{t+1} - p_{t+1} + n_{t+1}^y \right) + \sigma C \left(w_t - p_t + n_t^y \right) \right. \\ &\quad \left. - p_{t+1} + p_t - \Phi_0 \left(w_t - p_t + n_t^y \right) - \gamma_{t+1} \right] \end{aligned}$$

Multiplying out

$$\begin{aligned} 0 &= -\sigma C \mathbb{E}_t [w_{t+1}] + (\sigma C - 1) \mathbb{E}_t [p_{t+1}] - \sigma C \mathbb{E}_t \left[n_{t+1}^y \right] \\ &\quad + (\sigma C - \Phi_0) w_t + (-\sigma C + 1 + \Phi_0) p_t + (\sigma C - \Phi_0) n_t^y - \mathbb{E}_t [\gamma_{t+1}] \end{aligned}$$

Using the stochastic processes we get

$$\begin{aligned}
0 &= -\sigma C \mathbb{E}_t [w_{t+1}] + (\sigma C - 1) \mathbb{E}_t [p_{t+1}] - \sigma C \mathbb{E}_t [n_{t+1}^y] \\
&+ (\sigma C - \Phi_0) w_t + (-\sigma C + 1 + \Phi_0) p_t + (\sigma C - \Phi_0) n_t^y \\
&- \rho \gamma \gamma t - \rho \gamma y z_t^y - \rho \gamma x z_t^x - \rho \gamma \epsilon t
\end{aligned}$$

Wage Setting Equation

The wage setting equation (eq. (1.16)) is given by

$$w_t = \lambda \left(p_t + \epsilon_t + \frac{1}{\phi} \left(\frac{N^y}{N} n_t^y + \frac{N^x}{N} n_t^x \right) + \sigma (w_t - p_t + n_t^y) \right) + (1 - \lambda) w_{t-1}$$

Let's rearrange it. Multiply out

$$0 = -w_t + \lambda p_t + \lambda \epsilon_t + \frac{\lambda N^y}{\phi N} n_t^y + \frac{\lambda N^x}{\phi N} n_t^x + \lambda \sigma w_t - \lambda \sigma p_t + \lambda \sigma n_t^y + (1 - \lambda) w_{t-1}$$

Substitute in equation (1.15')

$$0 = -w_t + \lambda p_t + \frac{\lambda N^y}{\phi N} n_t^y + \frac{\lambda N^x}{\phi N} \left(-\frac{1}{1-\theta} w_t + \frac{1}{1-\theta} z_t^x \right) + \lambda \sigma w_t - \lambda \sigma p_t + \lambda \sigma n_t^y + \lambda \epsilon_t + (1 - \lambda) w_{t-1}$$

Group like terms

$$0 = - \left(1 + \frac{1}{1-\theta} \frac{\lambda N^x}{\phi N} - \lambda \sigma \right) w_t + \lambda (1 - \sigma) p_t + \left(\frac{\lambda N^y}{\phi N} + \lambda \sigma \right) n_t^y + \frac{1}{1-\theta} \frac{\lambda N^x}{\phi N} z_t^x + \lambda \epsilon_t + (1 - \lambda) w_{t-1}$$

Bond Equation

The bond equation (eq. (1.17)) is given by

$$b_t = (1 + r) b_{t-1} + \frac{X}{B} (n_t^x - n_t^y)$$

Rearranging it,

$$0 = -b_t + \frac{X}{B} n_t^x - \frac{X}{B} n_t^y + (1 + r) b_{t-1}$$

Substitute in equation (1.15')

$$0 = -b_t + \frac{X}{B} \left(-\frac{1}{1-\theta} w_t + \frac{1}{1-\theta} z_t^x \right) - \frac{X}{B} n_t^y + (1 + r) b_{t-1}$$

Group like terms

$$0 = -\frac{1}{1-\theta} \frac{X}{B} w_t - b_t - \frac{X}{B} n_t^y + \frac{1}{1-\theta} \frac{X}{B} z_t^x + (1 + r) b_{t-1}$$

Producer First Order Condition

The producers' choice is described by equation (1.18),

$$\tilde{w}_t = \tilde{p}_t - (1 - (\alpha + \theta\beta)) \tilde{n}_t^y - \beta(1 - \theta) (\tilde{n}_t^x - \tilde{n}_t^y) + \tilde{z}_t^y + \beta \tilde{z}_t^x$$

Rearranging, multiplying out

$$0 = -w_t + p_t - (1 - \alpha - \theta\beta - \beta + \theta\beta) n_t^y - \beta(1 - \theta) n_t^x + z_t^y + \beta z_t^x$$

Grouping like terms

$$0 = -w_t + p_t - (1 - \alpha - \beta) n_t^y - \beta(1 - \theta) n_t^x + z_t^y + \beta z_t^x$$

Substitute in equation (1.15')

$$0 = -w_t + p_t - (1 - \alpha - \beta) n_t^y - \beta(1 - \theta) \left(-\frac{1}{1 - \theta} w_t + \frac{1}{1 - \theta} z_t^x \right) + z_t^y + \beta z_t^x$$

Multiply out

$$0 = -w_t + p_t - (1 - \alpha - \beta) n_t^y + \beta w_t - \beta z_t^x + z_t^y + \beta z_t^x$$

Group like terms

$$0 = (\beta - 1) w_t + p_t - (1 - \alpha - \beta) n_t^y + z_t^y$$

1.12.3.3 Matrix Representation and Solution

The matrix representation of the regional model is the following

$$\begin{aligned}
0 = & \begin{bmatrix} -\sigma C & 0 & \sigma C - 1 & -\sigma C \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \mathbb{E}_t \begin{bmatrix} w_{t+1} \\ b_{t+1} \\ p_{t+1} \\ n_{t+1}^y \end{bmatrix} \\
+ & \begin{bmatrix} \sigma C - \Phi_0 & 0 & -\sigma C + 1 + \Phi_0 & \sigma C - \Phi_0 \\ -\left(1 + \frac{1}{1-\theta} \frac{\lambda N^x}{\phi} - \lambda \sigma\right) & 0 & \lambda(1-\sigma) & \frac{\lambda N^y}{\phi} + \lambda \sigma \\ \beta - 1 & 0 & 1 & -(1-\alpha-\beta) \\ -\frac{1}{1-\theta} \frac{X}{B} & -1 & 0 & -\frac{X}{B} \end{bmatrix} \begin{bmatrix} w_t \\ b_t \\ p_t \\ n_t^y \end{bmatrix} \\
+ & \begin{bmatrix} -\rho_\gamma & -\rho_{\gamma y} & -\rho_{\gamma x} & -\rho_{\gamma \epsilon} \\ 0 & 0 & \frac{1}{1-\theta} \frac{\lambda N^x}{\phi} & \lambda \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \frac{1}{1-\theta} \frac{X}{B} & 0 \end{bmatrix} \begin{bmatrix} \gamma_t \\ z_t^y \\ z_t^x \\ \epsilon_t \end{bmatrix} \\
+ & \begin{bmatrix} 0 & 0 & 0 & 0 \\ 1-\lambda & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 1+r & 0 & 0 \end{bmatrix} \begin{bmatrix} w_{t-1} \\ b_{t-1} \\ p_{t-1} \\ n_{t-1}^y \end{bmatrix}
\end{aligned}$$

Following Uhlig (1995), pages 9 and 10, we have a system of the form

$$0 = \mathbb{E}_t [F x_{t+1} + G x_t + H x_{t-1} + L z_{t+1} + M z_t]$$

$$z_{t+1} = N z_t + \epsilon_{t+1}$$

$$\mathbb{E}_t [\epsilon_{t+1}] = 0$$

where x_t is a vector with the endogenous variables (jump and state variables) and z_t is a vector of exogenous stochastic processes. In our case we have $L = 0$, and

$$F = \begin{bmatrix} -\sigma C & 0 & \sigma C - 1 & -\sigma C \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$G = \begin{bmatrix} \sigma C - \Phi_0 & 0 & -\sigma C + 1 + \Phi_0 & \sigma C - \Phi_0 \\ -\left(1 + \frac{1}{1-\theta} \frac{\lambda}{\phi} \frac{N^x}{N} - \lambda\sigma\right) & 0 & \lambda(1-\sigma) & \frac{\lambda}{\phi} \frac{N^y}{N} + \lambda\sigma \\ \beta - 1 & 0 & 1 & -(1-\alpha-\beta) \\ -\frac{1}{1-\theta} \frac{X}{B} & -1 & 0 & -\frac{X}{B} \end{bmatrix}$$

$$H = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 1-\lambda & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 1+r & 0 & 0 \end{bmatrix}$$

$$M = \begin{bmatrix} -\rho_\gamma & -\rho_{\gamma y} & -\rho_{\gamma x} & -\rho_{\gamma \epsilon} \\ 0 & 0 & \frac{1}{1-\theta} \frac{\lambda}{\phi} \frac{N^x}{N} & \lambda \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \frac{1}{1-\theta} \frac{X}{B} & 0 \end{bmatrix}$$

$$N = \begin{bmatrix} \rho_\gamma & \rho_{\gamma y} & \rho_{\gamma x} & \rho_{\gamma \epsilon} \\ \rho_{y\gamma} & \rho_y & \rho_{yx} & \rho_{y\epsilon} \\ \rho_{x\gamma} & \rho_{xy} & \rho_x & \rho_{x\epsilon} \\ \rho_{\epsilon\gamma} & \rho_{\epsilon y} & \rho_{\epsilon x} & \rho_\epsilon \end{bmatrix}$$

The solution we are looking for is of the form

$$x_t = \tilde{P}x_{t-1} + \tilde{Q}z_t$$

In the case of the regional model,

$$[\tilde{w}_t, \tilde{b}_t, \tilde{p}_t, \tilde{n}_t^y]' = \tilde{P} [\tilde{w}_{t-1}, \tilde{b}_{t-1}] + \tilde{Q} [\tilde{\gamma}_t, \tilde{z}_t^y, \tilde{z}_t^x, \tilde{\epsilon}_t]'$$

Theorem 1 from Uhlig (1995) shows that \tilde{P} satisfies the following matrix quadratic equation

$$0 = F\tilde{P}^2 + G\tilde{P} + H$$

Also, the same theorem shows that given \tilde{P} , matrix \tilde{Q} satisfies the following equation

$$V\tilde{Q} = -vec(LN + M)$$

where

$$V = N' \otimes F + I_k \otimes (F\tilde{P} + G)$$

In order to obtain \tilde{P} , I use Theorem 3 in Uhlig (1995), page 15. With matrix \tilde{P} , straight application of the above equations give us \tilde{Q} . The Matlab file PolicyRegionalStdPref.m computes these

expressions and solves for \tilde{P} and \tilde{Q} .

1.12.3.4 Matrix Representation and Solution With Consumption

To use consumption, I use that $\tilde{c}_t = \tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y$ and replace p_t by $\tilde{p}_t = \tilde{w}_t - \tilde{c}_t + \tilde{n}_t^y$. Matrix F , becomes

$$F = \begin{bmatrix} -1 & 0 & 1 - \sigma C & -1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

And matrix G becomes

$$G = \begin{bmatrix} 1 & 0 & \sigma C - 1 - \Phi_0 & 1 \\ -\left(1 + \frac{1}{1-\theta} \frac{\lambda}{\phi} \frac{N^x}{N} - \lambda\right) & 0 & -\lambda(1 - \sigma) & \frac{\lambda}{\phi} \frac{N^y}{N} + \lambda \\ \beta & 0 & -1 & (\alpha + \beta) \\ -\frac{1}{1-\theta} \frac{X}{B} & -1 & 0 & -\frac{X}{B} \end{bmatrix}$$

1.12.3.5 GHH Preferences

In section 1.12.6 of this appendix I derive the following log-linearized model equilibrium conditions for the aggregate economy with GHH preferences:

$$\begin{aligned}
0 &= \mathbb{E}_t [\tilde{m}u_{kt+1} - \tilde{m}u_{kt} - (\tilde{p}_{t+1} - \tilde{p}_t) - \Phi_0(\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y) - \tilde{\gamma}_{t+1}^\omega] \\
\tilde{m}u_t &= -\frac{\sigma}{\left(C - \frac{\phi}{1+\phi}N^{\frac{1+\phi}{\phi}}\right)} \left(C(\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y) - N^{\frac{1+\phi}{\phi}} \left(\frac{N^y}{N} \tilde{n}_{t+1}^y + \frac{N^x}{N} \tilde{n}_{t+1}^x \right) \right) \\
\tilde{w}_t &= \tilde{z}_t^{x\omega} - (1 - \theta) \tilde{n}_t^x \\
\tilde{w}_t &= \lambda \left(\tilde{p}_t + \tilde{\epsilon}_t^\omega + \frac{1}{\phi} \left(\frac{N^y}{N} \tilde{n}_t^y + \frac{N^x}{N} \tilde{n}_t^x \right) \right) + (1 - \lambda) \tilde{w}_{t-1} \\
\tilde{b}_t &= (1 + r) \tilde{b}_{t-1} + \frac{X}{B} (\tilde{n}_t^x - \tilde{n}_t^y) \\
\tilde{w}_t &= \tilde{p}_t - (1 - (\alpha + \theta\beta)) \tilde{n}_t^y - \beta(1 - \theta) (\tilde{n}_t^x - \tilde{n}_t^y) + \tilde{z}_t^{y\omega} + \beta \tilde{z}_t^{x\omega} \\
\lambda &\rightarrow 1
\end{aligned}$$

1.12.3.5.1 Algebra Work The Euler Equation is given by

$$\begin{aligned}
0 &= \mathbb{E}_t [mu_{kt+1} - mu_{kt} - (p_{t+1} - p_t) - \Phi_0(w_t - p_t + n_t^y) - \gamma_{t+1}] \\
\tilde{m}u_t &= -\frac{\sigma}{\left(C - \frac{\phi}{1+\phi}N^{\frac{1+\phi}{\phi}}\right)} \left(C(\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y) - N^{\frac{1+\phi}{\phi}} \left(\frac{N^y}{N} \tilde{n}_{t+1}^y + \frac{N^x}{N} \tilde{n}_{t+1}^x \right) \right)
\end{aligned}$$

Replacing in the expression for marginal utility

$$\begin{aligned}
0 &= \mathbb{E}_t \left[-\frac{\sigma}{\left(C - \frac{\phi}{1+\phi}N^{\frac{1+\phi}{\phi}}\right)} \left(C(w_{t+1} - p_{t+1} + n_{t+1}^y) - N^{\frac{1+\phi}{\phi}} \left(\frac{N^x}{N} n_{t+1}^x + \frac{N^y}{N} n_{t+1}^y \right) \right) \right. \\
&+ \frac{\sigma}{\left(C - \frac{\phi}{1+\phi}N^{\frac{1+\phi}{\phi}}\right)} \left(C(w_t - p_t + n_t^y) - N^{\frac{1+\phi}{\phi}} \left(\frac{N^x}{N} n_t^x + \frac{N^y}{N} n_t^y \right) \right) \\
&\left. - p_{t+1} + p_t - \Phi_0(w_t - p_t + n_t^y) - \gamma_{t+1} \right]
\end{aligned}$$

Multiplying out

$$\begin{aligned}
0 &= -\frac{\sigma C}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}\right)} \mathbb{E}_t [w_{t+1}] + \left(\frac{\sigma C}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}\right)} - 1 \right) \mathbb{E}_t [p_{t+1}] \\
&- \left(\frac{\sigma C}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}\right)} - \frac{\sigma N^{1/\phi} N^y}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}\right)} \right) \mathbb{E}_t [n_{t+1}^y] \\
&+ \frac{\sigma N^{1/\phi} N^x}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}\right)} \mathbb{E}_t [n_{t+1}^x] + \left(\frac{\sigma C}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}\right)} - \Phi_0 \right) w_t \\
&+ \left(-\frac{\sigma C}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}\right)} + 1 + \Phi_0 \right) p_t + \left(\frac{\sigma C}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}\right)} - \frac{\sigma N^{1/\phi} N^y}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}\right)} - \Phi_0 \right) n_t^y \\
&- \frac{\sigma N^{1/\phi} N^x}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}\right)} n_t^x - \mathbb{E}_t [\gamma_{t+1}]
\end{aligned}$$

To have a cleaner expression, let

$$m \equiv C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}$$

and substitute

$$\begin{aligned}
0 &= -\frac{\sigma C}{m} \mathbb{E}_t [w_{t+1}] + \left(\frac{\sigma C}{m} - 1 \right) \mathbb{E}_t [p_{t+1}] \\
&- \left(\frac{\sigma C}{m} - \frac{\sigma N^{1/\phi} N^y}{m} \right) \mathbb{E}_t [n_{t+1}^y] \\
&+ \frac{\sigma N^{1/\phi} N^x}{m} \mathbb{E}_t [n_{t+1}^x] + \left(\frac{\sigma C}{m} - \Phi_0 \right) w_t \\
&+ \left(-\frac{\sigma C}{m} + 1 + \Phi_0 \right) p_t + \left(\frac{\sigma C}{m} - \frac{\sigma N^{1/\phi} N^y}{m} - \Phi_0 \right) n_t^y \\
&- \frac{\sigma N^{1/\phi} N^x}{m} n_t^x - \mathbb{E}_t [\gamma_{t+1}]
\end{aligned}$$

Eliminate n_t^x by using equation (1.15')

$$\begin{aligned}
0 &= -\frac{\sigma C}{m} \mathbb{E}_t [w_{t+1}] + \left(\frac{\sigma C}{m} - 1 \right) \mathbb{E}_t [p_{t+1}] \\
&\quad - \left(\frac{\sigma C}{m} - \frac{\sigma N^{1/\phi} N^y}{m} \right) \mathbb{E}_t [n_{t+1}^y] \\
&\quad + \frac{\sigma N^{1/\phi} N^x}{m} \mathbb{E}_t \left[-\frac{1}{1-\theta} w_{t+1} + \frac{1}{1-\theta} z_{t+1}^x \right] + \left(\frac{\sigma C}{m} - \Phi_0 \right) w_t \\
&\quad + \left(-\frac{\sigma C}{m} + 1 + \Phi_0 \right) p_t + \left(\frac{\sigma C}{m} - \frac{\sigma N^{1/\phi} N^y}{m} - \Phi_0 \right) n_t^y \\
&\quad - \frac{\sigma N^{1/\phi} N^x}{m} \left(-\frac{1}{1-\theta} w_t + \frac{1}{1-\theta} z_t^x \right) - \mathbb{E}_t [\gamma_{t+1}]
\end{aligned}$$

Grouping like terms

$$\begin{aligned}
0 &= -\frac{\sigma C + \sigma \frac{1}{1-\theta} N^{1/\phi} N^x}{m} \mathbb{E}_t [w_{t+1}] + \left(\frac{\sigma C}{m} - 1 \right) \mathbb{E}_t [p_{t+1}] \\
&\quad - \left(\frac{\sigma C}{m} - \frac{\sigma N^{1/\phi} N^y}{m} \right) \mathbb{E}_t [n_{t+1}^y] + \left(\frac{\sigma C + \sigma \frac{1}{1-\theta} N^{1/\phi} N^x}{m} - \Phi_0 \right) w_t \\
&\quad + \left(-\frac{\sigma C}{m} + 1 + \Phi_0 \right) p_t + \left(\frac{\sigma C}{m} - \frac{\sigma N^{1/\phi} N^y}{m} - \Phi_0 \right) n_t^y \\
&\quad - \mathbb{E}_t [\gamma_{t+1}] + \frac{\sigma \frac{1}{1-\theta} N^{1/\phi} N^x}{m} (\mathbb{E}_t [z_{t+1}^x] - z_t^x)
\end{aligned}$$

Using the stochastic processes we get

$$\begin{aligned}
0 &= -\frac{\sigma C + \sigma \frac{1}{1-\theta} N^{1/\phi} N^x}{m} \mathbb{E}_t [w_{t+1}] + \left(\frac{\sigma C}{m} - 1 \right) \mathbb{E}_t [p_{t+1}] \\
&\quad - \left(\frac{\sigma C}{m} - \frac{\sigma N^{1/\phi} N^y}{m} \right) \mathbb{E}_t [n_{t+1}^y] + \left(\frac{\sigma C + \sigma \frac{1}{1-\theta} N^{1/\phi} N^x}{m} - \Phi_0 \right) w_t \\
&\quad + \left(-\frac{\sigma C}{m} + 1 + \Phi_0 \right) p_t + \left(\frac{\sigma C}{m} - \frac{\sigma N^{1/\phi} N^y}{m} - \Phi_0 \right) n_t^y \\
&\quad - \rho_\gamma \gamma_t - \rho_{\gamma y} z_t^y - \rho_{\gamma x} z_t^x - \rho_{\gamma \epsilon} \epsilon_t \\
&\quad + \frac{\sigma \frac{1}{1-\theta} N^{1/\phi} N^x}{m} (\rho_{x\gamma} \gamma_t + \rho_{xy} z_t^y + (\rho_x - 1) z_t^x + \rho_{x\epsilon} \epsilon_t)
\end{aligned}$$

Finally grouping like terms one more time

$$\begin{aligned}
0 &= -\frac{\sigma C + \sigma \frac{1}{1-\theta} N^{1/\phi} N^x}{m} \mathbb{E}_t [w_{t+1}] + \left(\frac{\sigma C}{m} - 1 \right) \mathbb{E}_t [p_{t+1}] \\
&\quad - \left(\frac{\sigma C - \sigma N^{1/\phi} N^y}{m} \right) \mathbb{E}_t [n_{t+1}^y] + \left(\frac{\sigma C + \sigma \frac{1}{1-\theta} N^{1/\phi} N^x}{m} - \Phi_0 \right) w_t \\
&\quad + \left(-\frac{\sigma C}{m} + 1 + \Phi_0 \right) p_t + \left(\frac{\sigma C - \sigma N^{1/\phi} N^y}{m} - \Phi_0 \right) n_t^y \\
&\quad + \left(-\rho_\gamma + \frac{\sigma \frac{1}{1-\theta} N^{1/\phi} N^x \rho_{x\gamma}}{m} \right) \gamma_t + \left(-\rho_{\gamma y} + \frac{\sigma \frac{1}{1-\theta} N^{1/\phi} N^x \rho_{xy}}{m} \right) z_t^y \\
&\quad + \left(-\rho_{\gamma x} + \frac{\sigma \frac{1}{1-\theta} N^{1/\phi} N^x (\rho_x - 1)}{m} \right) z_t^x + \left(-\rho_{\gamma \epsilon} + \frac{\sigma \frac{1}{1-\theta} N^{1/\phi} N^x \rho_{x\epsilon}}{m} \right) \epsilon_t
\end{aligned}$$

To write these expressions in a more concise way, let

$$f = \sigma \frac{1}{1-\theta} N^{1/\phi} N^x$$

$$g = \sigma N^{1/\phi} N^y$$

The equation becomes

$$\begin{aligned} 0 = & -\frac{\sigma C + f}{m} \mathbb{E}_t [w_{t+1}] + \left(\frac{\sigma C}{m} - 1 \right) \mathbb{E}_t [p_{t+1}] \\ & - \left(\frac{\sigma C - g}{m} \right) \mathbb{E}_t [n_{t+1}^y] + \left(\frac{\sigma C + f}{m} - \Phi_0 \right) w_t \\ & + \left(-\frac{\sigma C}{m} + 1 + \Phi_0 \right) p_t + \left(\frac{\sigma C - g}{m} - \Phi_0 \right) n_t^y \\ & + \left(-\rho_\gamma + \frac{f \rho_{x\gamma}}{m} \right) \gamma_t + \left(-\rho_{\gamma y} + \frac{f \rho_{xy}}{m} \right) z_t^y \\ & + \left(-\rho_{\gamma x} + \frac{f(\rho_x - 1)}{m} \right) z_t^x + \left(-\rho_{\gamma \epsilon} + \frac{f \rho_{x\epsilon}}{m} \right) \epsilon_t \end{aligned}$$

Wage Setting Equation

The wage setting equation is given by

$$w_t = \lambda \left(p_t + \epsilon_t + \frac{1}{\phi} \left(\frac{N^x}{N} n_t^x + \frac{N^y}{N} n_t^y \right) \right) + (1 - \lambda) w_{t-1}$$

Let's rearrange it. Multiply out

$$0 = -w_t + \lambda p_t + \frac{\lambda N^y}{\phi N} n_t^y + \frac{\lambda N^x}{\phi N} n_t^x + \lambda \epsilon_t + (1 - \lambda) w_{t-1}$$

Substitute in equation (1.15')

$$0 = -w_t + \lambda p_t + \frac{\lambda N^y}{\phi N} n_t^y + \frac{\lambda N^x}{\phi N} \left(-\frac{1}{1 - \theta} w_t + \frac{1}{1 - \theta} z_t^x \right) + \lambda \epsilon_t + (1 - \lambda) w_{t-1}$$

Group like terms

$$0 = - \left(1 + \frac{1}{1-\theta} \frac{\lambda N^x}{\phi N} \right) w_t + \lambda p_t + \frac{\lambda N^y}{\phi N} n_t^y + \frac{1}{1-\theta} \frac{\lambda N^x}{\phi N} z_t^x + \lambda \epsilon_t + (1-\lambda) w_{t-1}$$

Bond Equation

The bond equation (eq. (1.17)) is given by

$$b_t = (1+r) b_{t-1} + \frac{X}{B} (n_t^x - n_t^y)$$

Rearranging it,

$$0 = -b_t + \frac{X}{B} n_t^x - \frac{X}{B} n_t^y + (1+r) b_{t-1}$$

Substitute in equation (1.15')

$$0 = -b_t + \frac{X}{B} \left(-\frac{1}{1-\theta} w_t + \frac{1}{1-\theta} z_t^x \right) - \frac{X}{B} n_t^y + (1+r) b_{t-1}$$

Group like terms

$$0 = -\frac{1}{1-\theta} \frac{X}{B} w_t - b_t - \frac{X}{B} n_t^y + \frac{1}{1-\theta} \frac{X}{B} z_t^x + (1+r) b_{t-1}$$

Producer First Order Condition

The producers' choice is described by equation (1.18),

$$\tilde{w}_t = \tilde{p}_t - (1 - (\alpha + \theta\beta)) \tilde{n}_t^y - \beta(1 - \theta)(\tilde{n}_t^x - \tilde{n}_t^y) + \tilde{z}_t^y + \beta\tilde{z}_t^x$$

Rearranging, multiplying out

$$0 = -w_t + p_t - (1 - \alpha - \theta\beta - \beta + \theta\beta) n_t^y - \beta(1 - \theta) n_t^x + z_t^y + \beta z_t^x$$

Grouping like terms

$$0 = -w_t + p_t - (1 - \alpha - \beta) n_t^y - \beta(1 - \theta) n_t^x + z_t^y + \beta z_t^x$$

Substitute in equation (1.15')

$$0 = -w_t + p_t - (1 - \alpha - \beta) n_t^y - \beta(1 - \theta) \left(-\frac{1}{1 - \theta} w_t + \frac{1}{1 - \theta} z_t^x \right) + z_t^y + \beta z_t^x$$

Multiply out

$$0 = -w_t + p_t - (1 - \alpha - \beta) n_t^y + \beta w_t - \beta z_t^x + z_t^y + \beta z_t^x$$

Group like terms

$$0 = (\beta - 1) w_t + p_t - (1 - \alpha - \beta) n_t^y + z_t^y$$

1.12.3.6 Matrix Representation and Solution

The matrix representation of the regional model is the following

$$\begin{aligned}
0 &= \mathbb{E}_t \begin{bmatrix} -\frac{\sigma C+f}{m} & 0 & \frac{\sigma C}{m} - 1 & -\frac{\sigma C-g}{m} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} w_{t+1} \\ b_{t+1} \\ p_{t+1} \\ n_{t+1}^y \end{bmatrix} \\
&+ \begin{bmatrix} \frac{\sigma C+f}{m} - \Phi_0 & 0 & -\frac{\sigma C}{m} + 1 + \Phi_0 & \frac{\sigma C-g}{m} - \Phi_0 \\ -\left(1 + \frac{1}{1-\theta} \frac{\lambda N^x}{\phi N}\right) & 0 & \lambda & \frac{\lambda N^y}{\phi N} \\ \beta - 1 & 0 & 1 & -(1 - \alpha - \beta) \\ -\frac{1}{1-\theta} \frac{X}{B} & -1 & 0 & -\frac{X}{B} \end{bmatrix} \begin{bmatrix} w_t \\ b_t \\ p_t \\ n_t^y \end{bmatrix} \\
&+ \begin{bmatrix} -\rho_\gamma + \frac{f\rho_{x\gamma}}{m} & -\rho_{\gamma y} + \frac{f\rho_{xy}}{m} & -\rho_{\gamma x} + \frac{f(\rho_x-1)}{m} & -\rho_{\gamma\epsilon} + \frac{f\rho_{x\epsilon}}{m} \\ 0 & 0 & \frac{1}{1-\theta} \frac{\lambda N^x}{\phi N} & \lambda \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \frac{1}{1-\theta} \frac{X}{B} & 0 \end{bmatrix} \begin{bmatrix} \gamma_t \\ z_t^y \\ z_t^x \\ \epsilon_t \end{bmatrix} \\
&+ \begin{bmatrix} 0 & 0 & 0 & 0 \\ 1 - \lambda & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 1 + r & 0 & 0 \end{bmatrix} \begin{bmatrix} w_{t-1} \\ b_{t-1} \\ p_{t-1} \\ n_{t-1}^y \end{bmatrix}
\end{aligned}$$

where

$$m \equiv C - \frac{\phi}{1 + \phi} N^{\frac{1+\phi}{\phi}}$$

$$f = \sigma \frac{1}{1-\theta} N^{1/\phi} N^x$$

$$g = \sigma N^{1/\phi} N^y$$

Following Uhlig (1995), pages 9 and 10, we have a system of the form

$$0 = \mathbb{E}_t [F x_{t+1} + G x_t + H x_{t-1} + L z_{t+1} + M z_t]$$

$$z_{t+1} = N z_t + \epsilon_{t+1}$$

$$\mathbb{E}_t [\epsilon_{t+1}] = 0$$

where x_t is a vector with the endogenous variables (jump and state variables) and z_t is a vector of exogenous stochastic processes. In our case we have $L = 0$, and

$$F = \begin{bmatrix} -\frac{\sigma C + f}{m} & 0 & \frac{\sigma C}{m} - 1 & -\frac{\sigma C - g}{m} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$G = \begin{bmatrix} \frac{\sigma C + f}{m} - \Phi_0 & 0 & -\frac{\sigma C}{m} + 1 + \Phi_0 & \frac{\sigma C - g}{m} - \Phi_0 \\ -\left(1 + \frac{1}{1-\theta} \frac{\lambda N^x}{\phi N}\right) & 0 & \lambda & \frac{\lambda N^y}{\phi N} \\ \beta - 1 & 0 & 1 & -(1 - \alpha - \beta) \\ -\frac{1}{1-\theta} \frac{X}{B} & -1 & 0 & -\frac{X}{B} \end{bmatrix}$$

$$H = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 1 - \lambda & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 1 + r & 0 & 0 \end{bmatrix}$$

$$M = \begin{bmatrix} -\rho_\gamma + \frac{f\rho_{x\gamma}}{m} & -\rho_{\gamma y} + \frac{f\rho_{xy}}{m} & -\rho_{\gamma x} + \frac{f(\rho_x - 1)}{m} & -\rho_{\gamma\epsilon} + \frac{f\rho_{x\epsilon}}{m} \\ 0 & 0 & \frac{1}{1-\theta} \frac{\lambda}{\phi} \frac{N^x}{N} & \lambda \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \frac{1}{1-\theta} \frac{X}{B} & 0 \end{bmatrix}$$

$$N = \begin{bmatrix} \rho_\gamma & \rho_{\gamma y} & \rho_{\gamma x} & \rho_{\gamma\epsilon} \\ \rho_{y\gamma} & \rho_y & \rho_{yx} & \rho_{y\epsilon} \\ \rho_{x\gamma} & \rho_{xy} & \rho_x & \rho_{x\epsilon} \\ \rho_{\epsilon\gamma} & \rho_{\epsilon y} & \rho_{\epsilon x} & \rho_\epsilon \end{bmatrix}$$

The Matlab file PolicyRegionalGHHPref.m computes these expressions and solves for \tilde{P} and \tilde{Q} .

1.12.3.7 Matrix Representation and Solution With Consumption

To use consumption, I use that $\tilde{c}_t = \tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y$ and replace p_t by $\tilde{p}_t = \tilde{w}_t - \tilde{c}_t + \tilde{n}_t^y$. Matrix F , becomes

$$F = \begin{bmatrix} -1 - \frac{f}{m} & 0 & 1 - \frac{\sigma C}{m} & -1 + \frac{g}{m} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

And matrix G becomes

$$G = \begin{bmatrix} 1 + \frac{f}{m} - \Phi_0 & 0 & \frac{\sigma C}{m} - 1 - \Phi_0 & 1 - \frac{g}{m} \\ -\left(1 + \frac{1}{1-\theta} \frac{\lambda N^x}{\phi N} - \lambda\right) & 0 & -\lambda & \frac{\lambda N^y}{\phi N} + \lambda \\ \beta & 0 & -1 & \alpha + \beta \\ -\frac{1}{1-\theta} \frac{X}{B} & -1 & 0 & -\frac{X}{B} \end{bmatrix}$$

1.12.4 *Solution to Aggregate Model*

In this section I explain in detail how to solve the aggregate model. I will do so by allowing a more general process for the shocks, in which they are allowed to be correlated. The goal is to find a solution to the system of equations that describe the aggregate model in which the the endogenous variables π_t^w, w_t^r, n_t are expressed as a linear function of the exogenous processes:

$$[w_t^r, \pi_t^w, n_t]' = Q[\gamma_t^\omega, z_t^\omega, \epsilon_t^\omega]'$$

where Q is a matrix that is a function of the structural parameters of the model. The goal of this section is to find matrix Q .

1.12.4.1 Standard Preferences

Recall from the Estimation section in the paper (section 1.6) that the log-linearized model equilibrium conditions for the aggregate economy are given by:

$$0 = \mathbb{E}_t [mu_{t+1} - mu_t - \pi_{t+1} - \gamma_{t+1}^\omega] + \varphi_\pi \mathbb{E}_t [\pi_{t+1}] + \varphi_y (w_t^r + n_t - y^*) \quad (1.7)$$

$$mu_t = -\sigma C(w_t^r + n_t) \quad (1.8)$$

$$\pi_t^w = \frac{\lambda}{1-\lambda} \left(\epsilon_t^\omega + \frac{1}{\phi} n_t + \sigma(w_t^r + n_t) - w_t^r \right) \quad (1.9)$$

$$w_t^r = -(1 - (\alpha + \theta\beta)) n_t + z_t^\omega \quad (1.10)$$

$$\pi_{t+1} \equiv \pi_{t+1}^w - (w_{t+1}^r - w_t^r) \quad (1.11)$$

$$y^* = \iota z_t^\omega + \eta \epsilon_t^\omega \quad (1.12)$$

$$\lambda \rightarrow 1$$

Where the wedges follow the AR(1) process

$$\begin{bmatrix} \gamma_t^\omega \\ z_t^\omega \\ \epsilon_t^\omega \end{bmatrix} = \begin{bmatrix} \gamma_0 \\ z_0 \\ \epsilon_0 \end{bmatrix} + \begin{bmatrix} \rho_\gamma & \rho_{\gamma z} & \rho_{\gamma \epsilon} \\ \rho_{z\gamma} & \rho_z & \rho_{z\epsilon} \\ \rho_{\epsilon\gamma} & \rho_{\epsilon z} & \rho_\epsilon \end{bmatrix} \begin{bmatrix} \gamma_{t-1}^\omega \\ z_{t-1}^\omega \\ \epsilon_{t-1}^\omega \end{bmatrix} + \begin{bmatrix} \sigma_\gamma & \sigma_{\gamma z} & \sigma_{\gamma \epsilon} \\ \sigma_{z\gamma} & \sigma_z & \sigma_{z\epsilon} \\ \sigma_{\epsilon\gamma} & \sigma_{\epsilon z} & \sigma_\epsilon \end{bmatrix} \begin{bmatrix} u_t^\gamma \\ u_t^z \\ u_t^\epsilon \end{bmatrix}$$

1.12.4.2 Algebra Work

Let's rewrite the system only in terms of $[\pi_t^w, w_t^r, n_t]$.

Euler Equation

The Euler Equation (equation (1.7)) is given by

$$0 = \mathbb{E}_t (mu_{t+1} - mu_t - \pi_{t+1} - \gamma_{t+1}^\omega) + \varphi_p \mathbb{E}_t (\pi_{t+1}) + \varphi_y (w_t^r + n_t - y^*)$$

Replace the expression from equations (1.8), (1.11), and (1.12):

$$\begin{aligned}
0 &= \mathbb{E}_t [-\sigma C(w_{t+1}^r + n_{t+1})] \\
&+ \mathbb{E}_t [\sigma C(w_t^r + n_t) - \pi_{t+1}^w + w_{t+1}^r - w_t^r - \gamma_{t+1}^\omega] \\
&+ \varphi_p \mathbb{E}_t (\pi_{t+1}^w - w_{t+1}^r + w_t^r) + \varphi_y (w_t^r + n_t - \iota z_t^\omega - \eta \epsilon_t^\omega)
\end{aligned}$$

Replace equation (1.10) into the marginal utility expressions:

$$\begin{aligned}
0 &= \mathbb{E}_t [-\sigma C(-(1 - (\alpha + \theta\beta)) n_{t+1} + z_{t+1}^\omega + n_{t+1})] \\
&+ \mathbb{E}_t [\sigma C(-(1 - (\alpha + \theta\beta)) n_t + z_t^\omega + n_t) - \pi_{t+1}^w + w_{t+1}^r - w_t^r - \gamma_{t+1}^\omega] \\
&+ \varphi_p \mathbb{E}_t (\pi_{t+1}^w - w_{t+1}^r + w_t^r) + \varphi_y (w_t^r + n_t - \iota z_t^\omega - \eta \epsilon_t^\omega)
\end{aligned}$$

Group like terms in the expectation terms

$$\begin{aligned}
0 &= \mathbb{E}_t [-\sigma C((\alpha + \theta\beta) n_{t+1} + z_{t+1}^\omega)] \\
&+ \mathbb{E}_t [\sigma C((\alpha + \theta\beta) n_t + z_t^\omega) - \pi_{t+1}^w + w_{t+1}^r - w_t^r - \gamma_{t+1}^\omega] \\
&+ \varphi_p \mathbb{E}_t (\pi_{t+1}^w - w_{t+1}^r + w_t^r) + \varphi_y (w_t^r + n_t - \iota z_t^\omega - \eta \epsilon_t^\omega)
\end{aligned}$$

Group like terms again

$$\begin{aligned}
0 &= (1 - \varphi_p) \mathbb{E}_t (w_{t+1}^r) + (\varphi_p - 1) \mathbb{E}_t (\pi_{t+1}^w) - \sigma C (\alpha + \theta \beta) \mathbb{E}_t (n_{t+1}) \\
&+ (\varphi_y + \varphi_p - 1) w_t^r + (\sigma C (\alpha + \theta \beta) + \varphi_y) n_t \\
&- \mathbb{E}_t (\gamma_{t+1}^\omega) - \sigma C \mathbb{E}_t (z_{t+1}^\omega) + \sigma C z_t^\omega - \varphi_{y\ell} z_t^\omega - \varphi_y \eta \epsilon_t^\omega
\end{aligned}$$

Using the autoregressive expression for the exogenous processes we have everything in terms of these processes at time t

$$\begin{aligned}
0 &= (1 - \varphi_p) \mathbb{E}_t (w_{t+1}^r) + (\varphi_p - 1) \mathbb{E}_t (\pi_{t+1}^w) - \sigma C (\alpha + \theta \beta) \mathbb{E}_t (n_{t+1}) \\
&+ (\varphi_y + \varphi_p - 1) w_t^r + (\sigma C (\alpha + \theta \beta) + \varphi_y) n_t \\
&- \rho_\gamma \gamma_t^\omega - \rho_{\gamma z} z_t^\omega - \rho_{\gamma \epsilon} \epsilon_t^\omega - \gamma_0^\omega - \sigma C (\rho_{z\gamma} \gamma_t^\omega + \rho_z z_t^\omega + \rho_{z\epsilon} \epsilon_t^\omega + z_0^\omega) + \sigma C z_t^\omega - \varphi_{y\ell} z_t^\omega - \varphi_y \eta \epsilon_t^\omega
\end{aligned}$$

Finally grouping like terms

$$\begin{aligned}
0 &= (1 - \varphi_p) \mathbb{E}_t (w_{t+1}^r) + (\varphi_p - 1) \mathbb{E}_t (\pi_{t+1}^w) - \sigma C (\alpha + \theta \beta) \mathbb{E}_t (n_{t+1}) \\
&+ (\varphi_y + \varphi_p - 1) w_t^r + (\sigma C (\alpha + \theta \beta) + \varphi_y) n_t \\
&+ (-\rho_\gamma - \sigma C \rho_{z\gamma}) \gamma_t^\omega + (\sigma C (1 - \rho_z) - \rho_{\gamma z} - \varphi_{y\ell}) z_t^\omega \\
&+ (-\sigma C \rho_{z\epsilon} - \rho_{\gamma \epsilon} - \varphi_y \eta) \epsilon_t^\omega - \gamma_0^\omega - \sigma C z_0^\omega
\end{aligned}$$

Wage Setting Equation

The wage setting equation (eq. (1.9)) is given by

$$\pi_t^w = \frac{\lambda}{1-\lambda} \left(\epsilon_t^\omega + \frac{1}{\phi} n_t - w_t^r \right)$$

Multiply through by $(1-\lambda)$

$$(1-\lambda) \pi_t^w = \lambda \left(\epsilon_t^\omega + \frac{1}{\phi} n_t - w_t^r \right)$$

Rearranging we get

$$0 = \lambda(\sigma-1) w_t^r - (1-\lambda) \pi_t^w + \lambda \left(\frac{1}{\phi} + \sigma \right) n_t + \lambda \epsilon_t^\omega$$

$$0 = -\frac{\lambda}{1-\lambda} w_t^r - \pi_t^w + \frac{\lambda}{1-\lambda} \frac{1}{\phi} n_t + \frac{\lambda}{1-\lambda} \epsilon_t^\omega$$

Labor Demand Equation

The labor demand equation (eq. (1.10)) is given by

$$w_t^r = -(1 - (\alpha + \theta\beta)) n_t + z_t^\omega$$

Rearranging we get

$$0 = -w_t^r + ((\alpha + \theta\beta) - 1)n_t + z_t^\omega$$

1.12.4.3 Matrix Representation and Solution

Setting the mean of the stochastic processes equal to zero, the matrix representation of the system above is the following

$$\begin{aligned}
0 = & \begin{bmatrix} 1 - \varphi_p & \varphi_p - 1 & -\sigma C(\alpha + \theta\beta) \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \mathbb{E}_t \begin{bmatrix} w_{t+1}^r \\ \pi_{t+1}^w \\ n_{t+1} \end{bmatrix} \\
& + \begin{bmatrix} \varphi_p + \varphi_y - 1 & 0 & \sigma C(\alpha + \theta\beta) + \varphi_y \\ \lambda(\sigma - 1) & -(1 - \lambda) & \lambda\left(\frac{1}{\phi} + \sigma\right) \\ -1 & 0 & ((\alpha + \theta\beta) - 1) \end{bmatrix} \begin{bmatrix} w_t^r \\ \pi_t^w \\ n_t \end{bmatrix} \\
& + \begin{bmatrix} -\rho_\gamma - \sigma C\rho_{z\gamma} & \sigma C(1 - \rho_z) - \rho_{\gamma z} - \varphi_y l & -\sigma C\rho_{z\epsilon} - \rho_{\gamma\epsilon} - \varphi_y \eta & -\gamma_0^\omega - \sigma C z_0^\omega \\ 0 & 0 & \lambda & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} \gamma_t^\omega \\ z_t^\omega \\ \epsilon_t^\omega \\ 1 \end{bmatrix}
\end{aligned}$$

Following Uhlig (1995), pages 9 and 10, we have a system of the form

$$0 = \mathbb{E}_t [F x_{t+1} + G x_t + H x_{t-1} + L z_{t+1} + M z_t]$$

$$z_{t+1} = Nz_t + \epsilon_{t+1}$$

$$\mathbb{E}_t[\epsilon_{t+1}] = 0$$

where x_t is a vector with the endogenous variables (jump and state variables) and z_t is a vector of exogenous stochastic processes. In our case we have $H = 0$, $L = 0$, and

$$F = \begin{bmatrix} 1 - \varphi_p & \varphi_p - 1 & \sigma C(\alpha + \theta\beta) \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$G = \begin{bmatrix} \varphi_p + \varphi_y - 1 & 0 & C(\alpha + \theta\beta) + \varphi_y \\ \lambda(\sigma - 1) & -(1 - \lambda) & \lambda\left(\frac{1}{\phi} + \sigma\right) \\ -1 & 0 & ((\alpha + \theta\beta) - 1) \end{bmatrix}$$

$$M = \begin{bmatrix} -\rho_\gamma - \sigma C \rho_{z\gamma} & \sigma C(1 - \rho_z) - \rho_{\gamma z} - \varphi_y l & -\sigma C \rho_{z\epsilon} - \rho_{\gamma\epsilon} - \varphi_y \eta & -\gamma_0^\omega - \sigma C z_0^\omega \\ 0 & 0 & \lambda & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$N = \begin{bmatrix} \rho_\gamma & \rho_{\gamma z} & \rho_{\gamma\epsilon} & \gamma_0^\omega \\ \rho_{z\gamma} & \rho_z & \rho_{z\epsilon} & z_0^\omega \\ \rho_{\epsilon\gamma} & \rho_{\epsilon z} & \rho_\epsilon & \epsilon_0^\omega \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

The solution we are looking for is of the form

$$x_t = Px_{t-1} + Qz_t$$

In the case of the aggregate model, there are no endogenous state variables and therefore we have

$x_t = Qz_t$. In terms of the model notation this is

$$[w_t^r, \pi_t^w, n_t]^T = Q[\gamma_t^w, z_t^w, \epsilon_t^w, 1]^T$$

Using Theorem 1 from Uhlig (1995), matrix Q satisfies the following equation

$$VQ = -vec(LN + M)$$

where

$$V = N' \otimes F + I_k \otimes (FP + G)$$

Since for the aggregate model $P = 0$, $L = 0$, the solution is just

$$VQ = -vec(M)$$

with

$$V = N' \otimes F + I_k \otimes (G)$$

The Matlab file PolicyAggregateStdPref.m computes these expressions and solves for Q .

1.12.4.4 GHH Preferences

In section 1.12.5.1 of this appendix I derive the following log-linearized model equilibrium conditions for the aggregate economy with GHH preferences:

$$\begin{aligned}
0 &= \mathbb{E}_t [mu_{t+1} - mu_t - \pi_{t+1} - \gamma_{t+1}^\omega] + \varphi_\pi \mathbb{E}_t [\pi_{t+1}] + \varphi_y (w_t^r + n_t - y^*) \\
mu_t &= -\frac{\sigma}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}\right)} \left(C(w_t^r + n_t) - N^{\frac{1+\phi}{\phi}} n_t\right) \\
\pi_t^w &= \frac{\lambda}{1-\lambda} \left(\epsilon_t^\omega + \frac{1}{\phi} n_t - w_t^r\right) \\
w_t^r &= -(1 - (\alpha + \theta\beta)) n_t + z_t^\omega \\
\pi_{t+1} &\equiv \pi_{t+1}^w - (w_{t+1}^r - w_t^r) \\
y^* &= \iota z_t^\omega + \eta \epsilon_t^\omega \\
\lambda &\rightarrow 1
\end{aligned}$$

1.12.4.5 Algebra Work

Let's rewrite the system only in terms of $[\pi_t^w, w_t^r, n_t]$.

Euler Equation

The Euler Equation is given by

$$0 = \mathbb{E}_t (mu_{t+1} - mu_t - \pi_{t+1} - \gamma_{t+1}^\omega) + \varphi_\pi \mathbb{E}_t (\pi_{t+1}) + \varphi_y (w_t^r + n_t - y^*)$$

Replace the expression from equations of marginal utility for GHH preferences, (1.11), and (1.12):

$$\begin{aligned}
0 &= \mathbb{E}_t \left[-\frac{\sigma}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}\right)} \left(C(w_{t+1}^r + n_{t+1}) - N^{\frac{1+\phi}{\phi}} n_{t+1} \right) \right] \\
&+ \mathbb{E}_t \left[\frac{\sigma}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}\right)} \left(C(w_t^r + n_t) - N^{\frac{1+\phi}{\phi}} n_t \right) - \pi_{t+1}^w + w_{t+1}^r - w_t^r - \gamma_{t+1}^\omega \right] \\
&+ \varphi_p \mathbb{E}_t (\pi_{t+1}^w - w_{t+1}^r + w_t^r) + \varphi_y (w_t^r + n_t - \iota z_t^\omega - \eta \epsilon_t^\omega)
\end{aligned}$$

Replace equation (1.10):

$$\begin{aligned}
0 &= \mathbb{E}_t \left[-\frac{\sigma}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}\right)} \left(C(-(1 - (\alpha + \theta\beta)) n_{t+1} + z_{t+1}^\omega + n_{t+1}) - N^{\frac{1+\phi}{\phi}} n_{t+1} \right) \right] \\
&+ \mathbb{E}_t \left[\frac{\sigma}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}\right)} \left(C(-(1 - (\alpha + \theta\beta)) n_t + z_t^\omega + n_t) - N^{\frac{1+\phi}{\phi}} n_t \right) - \pi_{t+1}^w + w_{t+1}^r - w_t^r - \gamma_{t+1}^\omega \right] \\
&+ \varphi_p \mathbb{E}_t (\pi_{t+1}^w - w_{t+1}^r + w_t^r) + \varphi_y (w_t^r + n_t - \iota z_t^\omega - \eta \epsilon_t^\omega)
\end{aligned}$$

Group like terms in the expectation terms

$$\begin{aligned}
0 &= \mathbb{E}_t \left[-\frac{\sigma}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}\right)} \left(\left(C(\alpha + \theta\beta) - N^{\frac{1+\phi}{\phi}} \right) n_{t+1} + C z_{t+1}^\omega \right) \right] \\
&+ \mathbb{E}_t \left[\frac{\sigma}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}\right)} \left(\left(C(\alpha + \theta\beta) - N^{\frac{1+\phi}{\phi}} \right) n_t + C z_t^\omega \right) - \pi_{t+1}^w + w_{t+1}^r - w_t^r - \gamma_{t+1}^\omega \right] \\
&+ \varphi_p \mathbb{E}_t (\pi_{t+1}^w - w_{t+1}^r + w_t^r) + \varphi_y (w_t^r + n_t - \iota z_t^\omega - \eta \epsilon_t^\omega)
\end{aligned}$$

Group like terms again

$$\begin{aligned}
0 &= (1 - \varphi_p) \mathbb{E}_t (w_{t+1}^r) + (\varphi_p - 1) \mathbb{E}_t (\pi_{t+1}^w) - \frac{\sigma C (\alpha + \theta \beta) - \sigma N^{\frac{1+\phi}{\phi}}}{C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}} \mathbb{E}_t (n_{t+1}) \\
&+ (\varphi_y + \varphi_p - 1) w_t^r + \left(\frac{\sigma C (\alpha + \theta \beta) - \sigma N^{\frac{1+\phi}{\phi}}}{C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}} + \varphi_y \right) n_t \\
&- \mathbb{E}_t (\gamma_{t+1}^\omega) - \frac{\sigma C}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}} \right)} \mathbb{E}_t (z_{t+1}^\omega) + \frac{\sigma C}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}} \right)} z_t^\omega - \varphi_y \iota z_t^\omega - \varphi_y \eta \epsilon_t^\omega
\end{aligned}$$

Using the autoregressive expression for the exogenous processes we have everything in terms of these processes at time t

$$\begin{aligned}
0 &= (1 - \varphi_p) \mathbb{E}_t (w_{t+1}^r) + (\varphi_p - 1) \mathbb{E}_t (\pi_{t+1}^w) - \frac{\sigma C (\alpha + \theta \beta) - \sigma N^{\frac{1+\phi}{\phi}}}{C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}} \mathbb{E}_t (n_{t+1}) \\
&+ (\varphi_y + \varphi_p - 1) w_t^r + \left(\frac{\sigma C (\alpha + \theta \beta) - \sigma N^{\frac{1+\phi}{\phi}}}{C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}} + \varphi_y \right) n_t \\
&- \rho_\gamma \gamma_t^\omega - \rho_{\gamma z} z_t^\omega - \rho_{\gamma \epsilon} \epsilon_t^\omega - \gamma_0^\omega + \left(\frac{\sigma C (1 - \rho_z)}{C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}} - \iota \varphi_y \right) z_t^\omega \\
&- \frac{\sigma C}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}} \right)} (\rho_z \gamma \gamma_t^\omega + \rho_{z \epsilon} \epsilon_t^\omega + z_0^\omega) - \varphi_y \eta \epsilon_t^\omega
\end{aligned}$$

To have a cleaner expression, let

$$s \equiv \alpha + \beta \theta$$

$$m \equiv C - \frac{\phi}{1 + \phi} N^{\frac{1+\phi}{\phi}}$$

and substitute

$$\begin{aligned}
0 &= (1 - \varphi_p) \mathbb{E}_t (w_{t+1}^r) + (\varphi_p - 1) \mathbb{E}_t (\pi_{t+1}^w) - \frac{\sigma C s - \sigma N^{\frac{1+\phi}{\phi}}}{m} \mathbb{E}_t (n_{t+1}) \\
&+ (\varphi_y + \varphi_p - 1) w_t^r + \left(\frac{\sigma C s - \sigma N^{\frac{1+\phi}{\phi}}}{m} + \varphi_y \right) n_t \\
&- \rho \gamma \gamma_t^\omega - \rho \gamma z z_t^\omega - \rho \gamma \epsilon \epsilon_t^\omega - \gamma_0^\omega + \left(\frac{\sigma C (1 - \rho z)}{m} - \nu \varphi_y \right) z_t^\omega \\
&- \frac{\sigma C}{m} (\rho z \gamma \gamma_t^\omega + \rho z \epsilon \epsilon_t^\omega + z_0^\omega) - \varphi_y \eta \epsilon_t^\omega
\end{aligned}$$

Finally, grouping like terms

$$\begin{aligned}
0 &= (1 - \varphi_p) \mathbb{E}_t (w_{t+1}^r) + (\varphi_p - 1) \mathbb{E}_t (\pi_{t+1}^w) - \frac{\sigma C s - \sigma N^{\frac{1+\phi}{\phi}}}{m} \mathbb{E}_t (n_{t+1}) \\
&+ (\varphi_y + \varphi_p - 1) w_t^r + \left(\frac{\sigma C s - \sigma N^{\frac{1+\phi}{\phi}}}{m} + \varphi_y \right) n_t \\
&+ \left(-\rho \gamma - \frac{\sigma C}{m} \rho z \gamma \right) \gamma_t^\omega + \left(\frac{\sigma C (1 - \rho z)}{m} - \nu \varphi_y - \rho \gamma z \right) z_t^\omega \\
&+ \left(-\frac{\sigma C}{m} \rho z \epsilon - \varphi_y \eta - \rho \gamma \epsilon \right) \epsilon_t^\omega - \gamma_0^\omega - \frac{\sigma C}{m} z_0^\omega
\end{aligned}$$

Wage Setting Equation

The wage setting equation (eq. (1.9)) is given by

$$\pi_t^w = \frac{\lambda}{1 - \lambda} \left(\epsilon_t + \frac{1}{\phi} n_t - w_t^r \right)$$

Rearranging we get

$$0 = -\frac{\lambda}{1-\lambda}w_t^r - \pi_t^w + \frac{\lambda}{1-\lambda}\frac{1}{\phi}n_t + \frac{\lambda}{1-\lambda}\epsilon_t$$

Multiply through by $(1-\lambda)$

$$0 = -\lambda w_t^r - (1-\lambda)\pi_t^w + \frac{\lambda}{\phi}n_t + \lambda\epsilon_t^\omega$$

Labor Demand Equation

The labor demand equation (eq. (1.10)) is given by

$$w_t^r = -(1 - (\alpha + \theta\beta))n_t + z_t^\omega$$

Rearranging we get

$$0 = -w_t^r + ((\alpha + \theta\beta) - 1)n_t + z_t^\omega$$

1.12.4.6 Matrix Representation and Solution

Setting the mean of the shock processes equal to zero, the matrix representation of the system above is the following

$$\begin{aligned}
0 &= \begin{bmatrix} 1 - \varphi_p & \varphi_p - 1 & -\frac{\sigma C s - \sigma N^{\frac{1+\phi}{\phi}}}{m} \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \mathbb{E}_t \begin{bmatrix} w_{t+1}^r \\ \pi_{t+1}^w \\ n_{t+1} \end{bmatrix} \\
&+ \begin{bmatrix} \varphi_p + \varphi_y - 1 & 0 & \frac{\sigma C s - \sigma N^{\frac{1+\phi}{\phi}}}{m} + \varphi_y \\ -\lambda & -(1 - \lambda) & \frac{\lambda}{\phi} \\ -1 & 0 & s - 1 \end{bmatrix} \begin{bmatrix} w_t^r \\ \pi_t^w \\ n_t \end{bmatrix} \\
&+ \begin{bmatrix} -\rho_\gamma - \frac{\sigma C}{m} \rho_{z\gamma} & \frac{\sigma C(1-\rho_z)}{m} - \nu\varphi_y - \rho_{\gamma z} & -\frac{\sigma C}{m} \rho_{z\epsilon} - \varphi_y \eta - \rho_{\gamma\epsilon} & -\gamma_0^\omega - \frac{\sigma C}{m} z_0^\omega \\ 0 & 0 & \lambda & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} \gamma_t^\omega \\ z_t^\omega \\ \epsilon_t^\omega \\ 1 \end{bmatrix}
\end{aligned}$$

Following Uhlig (1995), pages 9 and 10, we have a system of the form

$$0 = \mathbb{E}_t [F x_{t+1} + G x_t + H x_{t-1} + L z_{t+1} + M z_t]$$

$$z_{t+1} = N z_t + \epsilon_{t+1}$$

$$\mathbb{E}_t [\epsilon_{t+1}] = 0$$

where x_t is a vector with the endogenous variables (jump and state variables) and z_t is a

vector of exogenous stochastic processes. In our case we have $H = 0$, $L = 0$, and

$$F = \begin{bmatrix} 1 - \varphi_p & \varphi_p - 1 & -\frac{\sigma C s - \sigma N \frac{1+\phi}{\phi}}{m} \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$G = \begin{bmatrix} \varphi_p + \varphi_y - 1 & 0 & \frac{\sigma C s - \sigma N \frac{1+\phi}{\phi}}{m} + \varphi_y \\ -\lambda & -(1 - \lambda) & \frac{\lambda}{\phi} \\ -1 & 0 & s - 1 \end{bmatrix}$$

$$M = \begin{bmatrix} -\rho_\gamma - \frac{\sigma C}{m} \rho_{z\gamma} & \frac{\sigma C(1-\rho_z)}{m} - \nu\varphi_y - \rho_{\gamma z} & -\frac{\sigma C}{m} \rho_{z\epsilon} - \varphi_y \eta - \rho_{\gamma\epsilon} & -\gamma_0^\omega - \frac{\sigma C}{m} z_0^\omega \\ 0 & 0 & \lambda & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$N = \begin{bmatrix} \rho_\gamma & \rho_{\gamma z} & \rho_{\gamma\epsilon} & \gamma_0^\omega \\ \rho_{z\gamma} & \rho_z & \rho_{z\epsilon} & z_0^\omega \\ \rho_{\epsilon\gamma} & \rho_{\epsilon z} & \rho_\epsilon & \epsilon_0^\omega \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

The solution we are looking for is of the form

$$x_t = P x_{t-1} + Q z_t$$

In the case of the aggregate model, there are no endogenous state variables and therefore we have $x_t = Qz_t$. In terms of the model notation this is

$$[w_t^r, \pi_t^w, n_t]^{\prime} = Q[\gamma_t, z_t, \epsilon_t]^{\prime}$$

Using Theorem 1 from Uhlig (1995), matrix Q satisfies the following equation

$$VQ = -vec(LN + M)$$

where

$$V = N' \otimes F + I_k \otimes (FP + G)$$

Since for the aggregate model $P = 0$, $L = 0$, the solution is just

$$VQ = -vec(M)$$

with

$$V = N' \otimes F + I_k \otimes (G)$$

The Matlab file PolicyAggregateGHHPref.m computes these expressions and solves for Q .

1.12.5 Aggregation Results

1.12.5.1 Aggregation Result 1 - Aggregate Economy

First, after adding up equations (1.22), (1.24), and (1.28) we have that $n_{kt}^y = n_{kt}^x$. Adding up equation (1.22):

$$\sum_k \frac{1}{K} 0 = \sum_k \frac{1}{K} z_{kt}^x - \sum_k \frac{1}{K} w_{kt} + \sum_k \frac{1}{K} q_t - (1 - \theta) \sum_k \frac{1}{K} n_{kt}^x \implies 0 = z_t^x - w_t + q_t - (1 - \theta) n_t^x$$

Adding up equation (1.28)

$$\sum_k \frac{1}{K} x_{kt} = \sum_k \frac{1}{K} n_{kt}^y + \sum_k \frac{1}{K} w_{kt} - \sum_k \frac{1}{K} q_t \implies x_t = n_t^y + w_t - q_t$$

Combining these,

$$w_t - q_t = z_t^x - (1 - \theta) n_t^x = x_t - n_t^y \implies x_t = z_t^x - (1 - \theta) n_t^x + n_t^y$$

Finally dividing through by K in equation (1.24)

$$\sum_k \frac{1}{K} x_{kt} = \sum_k \frac{1}{K} (z_{kt}^x + \theta n_{kt}^x) \implies x_t = z_t^x + \theta n_t^x$$

Combining these we get

$$x_t = z_t^x + \theta n_t^x = z_t^x - (1 - \theta) n_t^x + n_t^y \implies n_t^y = n_t^x$$

Now let's add up the other equations. Equation (1.19) becomes

$$0 = \mathbb{E}_t [mu_{t+1} - mu_t - \pi_{t+1} - \Phi_0 (c_t - c_t) - \gamma_{t+1} + \varphi_\pi \mathbb{E}_t [\pi_{t+1}] + \varphi_y (y_t - y^*)]$$

$$0 = \mathbb{E}_t [mu_{t+1} - mu_t - \pi_{t+1} - \gamma_{t+1}] + \varphi_\pi \mathbb{E}_t [\pi_{t+1}] + \varphi_y (y_t - y^*)$$

Equation (1.20) becomes:

$$mu_t = -\sigma C c_t$$

Equation (1.21) becomes:

$$w_t^r = \lambda \left(\epsilon_t^\omega + \frac{1}{\phi} n_t + \sigma c_t \right) + (1 - \lambda) (w_{t-1}^r - \pi_t)$$

Equation (1.23) just says $n_t = n_t^x = n_t^y$. Equation (1.29) becomes:

$$c_t = w_t^r + n_t^y = w_t^r + n_t$$

Equation (1.30) becomes:

$$w_t^r = -(1 - (\alpha + \theta\beta)) n_t^y + z_t^{y\omega} + \beta z_t^{x\omega} = -(1 - (\alpha + \theta\beta)) n_t + z_t^{y\omega} + \beta z_t^{x\omega} = -(1 - (\alpha + \theta\beta)) n_t + z_t^\omega$$

where we used that $z_t^\omega \equiv z_t^{y\omega} + \beta z_t^{x\omega}$. Replacing $y_t = c_t$ and $c_t = w_t^r + n_t$ and noting that

$\sum_k b_{kt} = 0$ we get the following 4 unknowns $(w_t^r, n_t, \pi_{t+1}, mu_{t+1})$ and 4 equations:

$$0 = \mathbb{E}_t [mu_{t+1} - mu_t - \pi_{t+1} - \gamma_{t+1}^\omega] + \varphi_\pi \mathbb{E}_t [\pi_{t+1}] + \varphi_y (w_t^r + n_t - y^*)$$

$$mu_t = -\sigma C c_t$$

$$w_t^r = \lambda \left(\epsilon_t^\omega + \frac{1}{\phi} n_t + \sigma c_t \right) + (1 - \lambda) (w_{t-1}^r - \pi_t)$$

$$w_t^r = -(1 - (\alpha + \theta\beta)) n_t + z_t^\omega$$

Notice that this system of recursive equations has three exogenous state variables and one endogenous state variable, w_{t-1} . In terms of solving the model it may be easier to get rid of the endogenous state variable. To this purpose, using that $\pi_{t+1} \equiv \pi_{t+1}^w - (w_{t+1}^r - w_t^r)$ rewrite the wage setting equation as follows

$$\begin{aligned} w_t^r &= \lambda \left(\epsilon_t^\omega + \frac{1}{\phi} n_t + \sigma c_t \right) + (1 - \lambda) (w_{t-1}^r - \pi_t) \\ w_t^r &= \lambda \left(\epsilon_t^\omega + \frac{1}{\phi} n_t + \sigma c_t \right) + (1 - \lambda) (w_t^r - \pi_t^w) \\ (1 - \lambda) \pi_t^w &= \lambda \left(\epsilon_t^\omega + \frac{1}{\phi} n_t + \sigma c_t \right) + (1 - \lambda) w_t^r - w_t^r \\ \pi_t^w &= \frac{\lambda}{1 - \lambda} \left(\epsilon_t^\omega + \frac{1}{\phi} n_t + \sigma c_t - w_t^r \right) \end{aligned}$$

GHH Preferences Equation (1.31) becomes:

$$mu_t = -\frac{\sigma}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}} \right)} \left(C c_t - N^{\frac{1+\phi}{\phi}} n_t \right)$$

Equation (1.32) becomes:

$$w_t^r = \lambda \left(\epsilon_t^\omega + \frac{1}{\phi} n_t \right) + (1 - \lambda) (w_{t-1}^r - \pi_t)$$

1.12.5.2 Summary

Standard Preferences In sum, the evolution of the aggregate economy is characterized

by:

$$0 = \mathbb{E}_t [mu_{t+1} - mu_t - \pi_{t+1} - \gamma_{t+1}^\omega] + \varphi_\pi \mathbb{E}_t [\pi_{t+1}] + \varphi_y (w_t^r + n_t - y^*) \quad (1.7)$$

$$mu_t = -\sigma C(w_t^r + n_t) \quad (1.8)$$

$$\pi_t^w = \frac{\lambda}{1 - \lambda} \left(\epsilon_t^\omega + \frac{1}{\phi} n_t + \sigma(w_t^r + n_t) - w_t^r \right) \quad (1.9)$$

$$w_t^r = -(1 - (\alpha + \theta\beta)) n_t + z_t^\omega \quad (1.10)$$

$$\pi_{t+1} \equiv \pi_{t+1}^w - (w_{t+1}^r - w_t^r) \quad (1.11)$$

$$y^* = \iota z_t^\omega + \eta \epsilon_t^\omega \quad (1.12)$$

$$\lambda \rightarrow 1 \quad (??)$$

GHH Preferences

$$\begin{aligned}
0 &= \mathbb{E}_t [mu_{t+1} - mu_t - \pi_{t+1} - \gamma_{t+1}^\omega] + \varphi_\pi \mathbb{E}_t [\pi_{t+1}] + \varphi_y (w_t^r + n_t - y^*) \\
mu_t &= -\frac{\sigma}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}\right)} \left(C(w_t^r + n_t) - N^{\frac{1+\phi}{\phi}} n_t\right) \\
\pi_t^w &= \frac{\lambda}{1-\lambda} \left(\epsilon_t^\omega + \frac{1}{\phi} n_t - w_t^r\right) \\
w_t^r &= -(1 - (\alpha + \theta\beta)) n_t + z_t^\omega \\
\pi_{t+1} &\equiv \pi_{t+1}^w - (w_{t+1}^r - w_t^r) \\
y^* &= \iota z_t^\omega + \eta \epsilon_t^\omega \\
\lambda &\rightarrow 1
\end{aligned}$$

1.12.6 Aggregation Result 2 - Regional Economies

Now we express the evolution of the regional economies as log-deviations from the aggregate.

Let $\tilde{x}_t \equiv x_{kt} - x_t$. Start with equation (1.23) as it will be used in some of the other equations:

$$Nn_{kt} = N^y n_{kt}^y + N^x n_{kt}^x$$

Summing over k and dividing through by K :

$$Nn_t = N^y n_t^y + N^x n_t^x$$

Subtracting we get:

$$N\tilde{n}_t = N^y\tilde{n}_t^y + N^x\tilde{n}_t^x \implies \tilde{n}_t = \frac{N^y}{N}\tilde{n}_t^y + \frac{N^x}{N}\tilde{n}_t^x \quad (1 \text{ OA})$$

Do the same for equation (1.29))

$$c_{kt} = w_{kt}^r + n_{kt}^y$$

and obtain

$$c_t = w_t^r + n_t^y$$

which can be written as

$$\tilde{c}_t = \tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y \quad (2 \text{ OA})$$

Finally, take equation (1.28) and write it also in deviations from the aggregate:

$$\tilde{x}_t = \tilde{n}_t^y + \tilde{w}_t \quad (3 \text{ OA})$$

Now re-write all the other regional equations in deviations from the aggregate. Subtracting the aggregate Euler Equation (1.7) from the regional Euler Equation (1.19), and obtain:

$$0 = \mathbb{E}_t [\tilde{m}u_{kt+1} - \tilde{m}u_{kt} - (\tilde{p}_{t+1} - \tilde{p}_t) - \Phi_0 (\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y) - \tilde{\gamma}_{t+1}^\omega]$$

Subtracting the aggregate marginal utility (1.8) from the regional (1.20) for standard preferences obtain:

$$\tilde{m}u_t = -\sigma C\tilde{c}_t$$

and for GHH preferences obtain

$$\tilde{m}u_t = -\frac{\sigma}{\left(C - \frac{\phi}{1+\phi}N^{\frac{1+\phi}{\phi}}\right)} \left(C\tilde{c}_t - N^{\frac{1+\phi}{\phi}}\tilde{n}_t\right)$$

Replacing (1 OA) and (2 OA), the expressions for marginal utility in both cases become

$$\tilde{m}u_t = -\sigma C (\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y)$$

$$\tilde{m}u_t = -\frac{\sigma}{\left(C - \frac{\phi}{1+\phi}N^{\frac{1+\phi}{\phi}}\right)} \left(C(\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y) - N^{\frac{1+\phi}{\phi}}\left(\frac{N^y}{N}\tilde{n}_{t+1}^y + \frac{N^x}{N}\tilde{n}_{t+1}^x\right)\right)$$

Equation (1.22) becomes

$$0 = \tilde{z}_t^{x\omega} - \tilde{w}_t - (1 - \theta)\tilde{n}_t^x \implies \tilde{w}_t = \tilde{z}_t^{x\omega} - (1 - \theta)\tilde{n}_t^x$$

Let's first write equations (1.21) and (1.32) in terms of wages and prices. For standard preferences

$$\begin{aligned} w_{kt}^r &= \lambda \left(\epsilon_{kt}^\omega + \frac{1}{\phi}n_{kt} + \sigma c_{kt} \right) + (1 - \lambda) (w_{kt-1}^r - \pi_{kt}) \\ w_{kt} - p_{kt} &= \lambda \left(\epsilon_{kt}^\omega + \frac{1}{\phi}n_{kt} + \sigma c_{kt} \right) + (1 - \lambda) (w_{kt-1} - p_{kt-1} - (p_{kt} - p_{kt-1})) \\ w_{kt} &= \lambda \left(p_{kt} + \epsilon_{kt}^\omega + \frac{1}{\phi}n_{kt} + \sigma c_{kt} \right) + (1 - \lambda) w_{kt-1} \end{aligned}$$

For GHH preferences

$$\begin{aligned}
w_{kt}^r &= \lambda \left(\epsilon_{kt}^\omega + \frac{1}{\phi} n_{kt} \right) + (1 - \lambda) (w_{kt-1}^r - \pi_{kt}) \\
w_{kt} - p_{kt} &= \lambda \left(\epsilon_{kt}^\omega + \frac{1}{\phi} n_{kt} \right) + (1 - \lambda) (w_{kt-1} - p_{kt-1} - (p_{kt} - p_{kt-1})) \\
w_{kt} &= \lambda \left(p_{kt} + \epsilon_{kt}^\omega + \frac{1}{\phi} n_{kt} \right) + (1 - \lambda) w_{kt-1}
\end{aligned}$$

Using (1 OA) and (2 OA), these two equations in deviations from the aggregate can be written as

$$\begin{aligned}
\tilde{w}_t &= \lambda \left(\tilde{p}_t + \tilde{\epsilon}_t^\omega + \frac{1}{\phi} \left(\frac{N^y}{N} \tilde{n}_t^y + \frac{N^x}{N} \tilde{n}_t^x \right) + \sigma(\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y) \right) + (1 - \lambda) \tilde{w}_{t-1} \\
\tilde{w}_t &= \lambda \left(\tilde{p}_t + \tilde{\epsilon}_t^\omega + \frac{1}{\phi} \left(\frac{N^y}{N} \tilde{n}_t^y + \frac{N^x}{N} \tilde{n}_t^x \right) \right) + (1 - \lambda) \tilde{w}_{t-1}
\end{aligned}$$

Equation (1.27) becomes

$$\tilde{b}_t = (1 + r) \tilde{b}_{t-1} + \frac{X}{B} (\tilde{z}_{kt}^{x\omega} + \theta \tilde{n}_t^x - \tilde{x}_t)$$

Using equation(3 OA)

$$\tilde{b}_t = (1 + r) \tilde{b}_{t-1} + \frac{X}{B} (\tilde{z}_{kt}^{x\omega} + \theta \tilde{n}_t^x - \tilde{n}_t^y - \tilde{w}_t)$$

From equation (1.22) in deviations from the aggregate (see above), we get

$$\tilde{b}_t = (1 + r) \tilde{b}_{t-1} + \frac{X}{B} (\tilde{n}_t^x - \tilde{n}_t^y)$$

Equation (1.30) becomes

$$\tilde{w}_t = \tilde{p}_t - (1 - (\alpha + \theta\beta)) \tilde{n}_t^y - \beta(1 - \theta) (\tilde{n}_t^x - \tilde{n}_t^y) + \tilde{z}_t^{y\omega} + \beta \tilde{z}_t^{x\omega}$$

1.12.6.1 Summary

Standard Preferences In sum, the evolution of the aggregate economy with standard preferences is characterized by:

$$0 = \mathbb{E}_t [\tilde{m}u_{kt+1} - \tilde{m}u_{kt} - (\tilde{p}_{t+1} - \tilde{p}_t) - \Phi_0 (\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y) - \tilde{\gamma}_{t+1}^\omega] \quad (1.13)$$

$$\tilde{m}u_t = -\sigma C (\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y) \quad (1.14)$$

$$\tilde{w}_t = \tilde{z}_t^{x\omega} - (1 - \theta) \tilde{n}_t^x \quad (1.15)$$

$$\tilde{w}_t = \lambda \left(\tilde{p}_t + \tilde{\epsilon}_t^\omega + \frac{1}{\phi} \left(\frac{N^y}{N} \tilde{n}_t^y + \frac{N^x}{N} \tilde{n}_t^x \right) + \sigma (\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y) \right) + (1 - \lambda) \tilde{w}_{t-1} \quad (1.16)$$

$$\tilde{b}_t = (1 + r) \tilde{b}_{t-1} + \frac{X}{B} (\tilde{n}_t^x - \tilde{n}_t^y) \quad (1.17)$$

$$\tilde{w}_t = \tilde{p}_t - (1 - (\alpha + \theta\beta)) \tilde{n}_t^y - \beta(1 - \theta) (\tilde{n}_t^x - \tilde{n}_t^y) + \tilde{z}_t^{y\omega} + \beta \tilde{z}_t^{x\omega} \quad (1.18)$$

$$\lambda \rightarrow 1$$

GHH Preferences In sum, the evolution of the aggregate economy with GHH preferences

is characterized by:

$$0 = \mathbb{E}_t \left[\tilde{m}u_{kt+1} - \tilde{m}u_{kt} - (\tilde{p}_{t+1} - \tilde{p}_t) - \Phi_0 (\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y) - \tilde{\gamma}_{t+1}^\omega \right]$$

$$\tilde{m}u_t = - \frac{\sigma}{\left(C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}} \right)} \left(C (\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y) - N^{\frac{1+\phi}{\phi}} \left(\frac{N^y}{N} \tilde{n}_{t+1}^y + \frac{N^x}{N} \tilde{n}_{t+1}^x \right) \right)$$

$$\tilde{w}_t = \tilde{z}_t^{x\omega} - (1 - \theta) \tilde{n}_t^x$$

$$\tilde{w}_t = \lambda \left(\tilde{p}_t + \tilde{\epsilon}_t^\omega + \frac{1}{\phi} \left(\frac{N^y}{N} \tilde{n}_t^y + \frac{N^x}{N} \tilde{n}_t^x \right) \right) + (1 - \lambda) \tilde{w}_{t-1}$$

$$\tilde{b}_t = (1 + r) \tilde{b}_{t-1} + \frac{X}{B} (\tilde{n}_t^x - \tilde{n}_t^y)$$

$$\tilde{w}_t = \tilde{p}_t - (1 - (\alpha + \theta\beta)) \tilde{n}_t^y - \beta (1 - \theta) (\tilde{n}_t^x - \tilde{n}_t^y) + \tilde{z}_t^{y\omega} + \beta \tilde{z}_t^{x\omega}$$

$$\lambda \rightarrow 1$$

CHAPTER 2

THE AGGREGATE IMPLICATIONS OF REGIONAL BUSINESS CYCLES

written jointly with Martin Beraja and Erik Hurst

2.1 Abstract

We argue that it is difficult to make inferences about the drivers of aggregate business cycles using regional variation alone because (i) the local and aggregate elasticities to the same type of shock are quantitatively different and (ii) purely aggregate shocks are differenced out when using cross-region variation. We highlight the importance of these issues in a monetary union model, and by contrasting the behavior of US aggregate time-series and cross-state patterns during the Great Recession. In particular, using household and retail scanner data for the US, we document a strong relationship across states between local employment growth and local nominal and real wage growth. These relationships are much weaker in US aggregates. In order to identify the shocks driving aggregate (and regional) business cycles we develop a methodology that combines regional and aggregate data. The methodology uses theoretical restrictions implied by a wage setting equation that holds in many monetary union models with nominal wage stickiness. We show how to estimate this equation using cross-state variation—thus linking particular regional patterns to particular aggregate shock decompositions. Applying the methodology to the US, we find that a combination of both "demand" and "supply" shocks are necessary to account for the joint dynamics of aggregate

prices, wages and employment during the 2007-2012 period while only "demand" shocks are necessary to explain most of the observed cross-state variation. We conclude that the wage stickiness necessary for demand shocks to be the primary cause of aggregate employment decline during the Great Recession is inconsistent with the flexibility of wages estimated from cross-state variation.

2.2 Introduction

A large and growing literature is exploiting regional variation to learn about the determinants of aggregate economic variables.¹ However, we argue that making inferences about the aggregate economy using only regional variation is complicated by two issues. First, we show that, in a monetary union model, local and aggregate elasticities to the same type of shock are quantitatively different both because of factor mobility and general equilibrium forces. This discrepancy makes it problematic to use local shock elasticities estimated from regional data to ascertain the importance of a given aggregate shock. Second, purely aggregate shocks get differenced out when using cross-region variation. As a result, it is not possible to learn anything about these aggregate shocks by exploiting variation across regions. Furthermore, we provide evidence of both these issues by contrasting the behavior of US aggregate time-series and cross-state patterns during the Great Recession. We document a strong relationship across states between local employment growth, and local nominal and real wage growth. These relationships are much weaker in US aggregates. In summary, we

1. For recent examples, see Autor et al (2013), Charles et al (2015), Hagedorn et al (2015), Mehrotra and Sergeyev (2015), Mian and Sufi (2014) and Mondragon (2015).

cannot expect to understand the joint evolution of aggregate variables by using cross-regional variation alone.

Therefore, we present a methodology that uses regional data along with aggregate data in order to identify aggregate shocks driving business cycles. The methodology exploits theoretical restrictions implied by a wage setting equation that hold in many monetary union models with wage stickiness. In turn, the extent to which aggregate wages are sticky is a key restriction in identifying the type of shocks driving aggregate fluctuations (e.g., "demand" vis a vis "supply" shocks)². Under certain conditions, we show how to use cross-region variation in wages, prices, and employment to estimate this wage setting equation—thus parameterizing the theoretical restrictions and linking regional business cycles to shock decompositions of aggregate business cycles.

Using household and retail scanner data for the US, we construct state-level wage and price indices as well as a measure of employment. Given the strong comovement of wages and employment across states, our estimates of the wage setting equation suggest that wages are relatively flexible—thus limiting the contribution of "demand" shocks to aggregate employment decline during the Great Recession. Instead, we find that a combination of "demand" and other shocks are necessary to account for the joint dynamics of aggregate prices, wages and employment during the 2007–2012 period. In particular, the relative stability of aggregate wages in the time-series compared to state-level wages is not caused by wage stickiness,

2. We refer to a "demand" shock as a shock that moves employment and real wages in opposite directions and moves employment and prices in the same direction. In the model of the monetary union we develop below, these shocks can be formalized as shocks to the household's discount rate or as shocks to the aggregate nominal interest rate rule. Our model also allows for a productivity/markup shock and a shock to household preference for leisure.

but because different aggregate shocks have relatively offsetting effects on aggregate wages. We conclude that the wage stickiness necessary for demand shocks to be the primary cause of aggregate employment decline during the Great Recession is inconsistent with the flexibility of wages estimated from cross-state variation.

The paper is organized as follows. In Sections 2.3 and 2.4, we begin by documenting a series of new facts about the variation in nominal and real wages across US states during the Great Recession. Using data from the *2000 US Census* and the 2000 - 2012 *American Community Surveys* (ACS), we construct state-level nominal wage indices during the 2000 to 2012 period. We restrict our sample to full time workers with a strong attachment to the labor force. We adjust our wage measures to cleanse them from observable changes in labor force composition over the business cycle. In order to construct a measure of real wages we deflate our nominal wage indices with state-level price indices created using data from Nielsen's *Retail Scanner Database*. The *Retail Scanner Database* includes weekly prices and quantities for given UPC codes at over 40,000 stores from 2006 through 2012. While the price indices we create from this data are based mostly on consumer packaged goods, we show how under certain assumptions the indices can be scaled to be representative of a composite local consumption good. Furthermore, we show that an aggregate price index created with the retail scanner data matches the BLS's Food CPI nearly identically.

Using our indices, we show that states that experienced larger employment declines between 2007 and 2010 had significantly lower nominal and real wage growth during the same time period. These cross-state patterns stand in sharp contrast with the well documented aggregate time-series trends for prices and wages during the same time period. As both ag-

aggregate output and employment contracted sharply in the US during the 2007-2012 period, aggregate nominal wage growth remained robust and real wage growth did not break trend.³ In sum, while aggregate wages appear to be sticky during the Great Recession, state-level wages do not.

In Section 2.5, we present a monetary union model that we use for two purposes. First, a calibrated version of the model allows us to sign the elasticities to a given shock and quantify the differences between aggregate and local elasticities. Second, the model makes explicit assumptions that are sufficient to estimate the parameters in an aggregate wage setting equation using cross-state variation in employment, wages and prices. As we highlight below, these parameters help us identify the underlying aggregate drivers of the joint dynamics of employment, wages and prices.

The model has many islands linked by trade in intermediate goods which are used in the production of a non-tradable final consumption good. The only asset is the economy is a one-period, non-state contingent nominal bond. The nominal interest rate on this asset follows a rule that endogenously responds to aggregate variables and is set at the union level. Labor is the only other input in production, which is not mobile across islands. We assume that nominal wages are only partially flexible. This is the only nominal rigidity in the model. Finally, the model includes multiple shocks: a shock to the household's discount rate, shocks to non-tradable and tradable productivity/markup, a shock to the household's preference for leisure, and a monetary policy shock. Aside from the monetary policy shock, all shocks

3. The robust growth in nominal wages during the recession is viewed as a puzzle for those that believe that the lack of aggregate demand was the primary cause of the Great Recession. For example, this point was made by Krugman in a recent New York Times article ("Wages, Yellen and Intellectual Honesty", NYTimes 8/25/14).

have both local and aggregate components. By definition the weighted average of the local shocks sum to zero. We show that, under relatively few assumptions, the log-linearized economy aggregates. This allows us to study the aggregate and local behavior separately, a property that we will exploit when estimating the aggregate and regional shocks through our methodology.

Using a calibrated version of the model, we show that local employment elasticities to a local discount rate shock are two to three times larger than the aggregate employment elasticity to a similarly sized aggregate discount rate shock. This implies that elasticities often estimated for "demand" shocks (i.e., our discount rate shock) using cross-region variation are likely to dramatically overstate the elasticities of aggregate variables to "demand" shocks in the aggregate. The key general equilibrium forces in the model that may dampen aggregate elasticities are the endogenous response of nominal interest rates to aggregate variables and trade in the intermediate input. We show that local and aggregate elasticities get much closer together when the interest rate does not endogenously respond to changes in aggregate prices or employment (as when the economy is close to the zero lower bound).⁴

In Section 2.6, we turn to estimation of aggregate shocks. We present a procedure that allow us estimate the shocks in a larger class of monetary union models than the benchmark model outlined above, thus imposing less a-priori structure and making the analysis more persuasive. In particular, we consider models where the aggregate and local equilibria can be represented as a structural vector autoregression (SVAR) in price inflation, nominal wage inflation, and employment with three shocks. We refer to the three shocks as the discount

4. A similar point is made in Nakamura and Steinsson(2014) with respect to local estimates of fiscal multipliers.

rate shock (which is a combination of the discount rate and monetary policy shock), the productivity/markup shock (which is a combination of the productivity/markup shocks in the tradable and non tradable sectors) and the leisure shock (which is the shock to leisure preference). In order to identify the aggregate shocks, we estimate a SVAR and impose certain properties of our benchmark monetary union model. Our results will be consistent with monetary union models that satisfy all of these. First, we use the aggregate wage setting equation to derive a series of linear restrictions linking the reduced form errors to the underlying structural shocks. Second, we use the sign of the joint-response of employment, wages and prices (on impact) to a discount rate and a productivity/markup shock.⁵ These two, together with the usual shock-orthogonality conditions, are sufficient to identify the structural shocks.

The methodology requires parameterizing the structural wage setting equation. We use state-level data on prices, wages and employment during the 2006-2012 period to estimate the two parameters in our base specification, i.e., the Frisch elasticity of labor supply and the degree of wage stickiness. Across a variety of specifications and identification procedures, including instrumenting for local labor demand shocks, we estimate only a modest degree of wage stickiness. These estimates are much smaller than estimates of wage stickiness obtained using only aggregate time-series data.

With the parameterized aggregate wage setting equation, we use the SVAR identification

5. We view this methodology as an additional contribution of our paper. Beraja (2015) presents an extension of this scheme to a more general class of models. These are part of a growing literature developing “hybrid” methods that, for instance, constructs optimal combinations of econometric and theoretical models (Carriero and Giacomini (2011), Del Negro and Schorfheide (2004)) or uses the theoretical model to inform the econometric model’s parameter (An and Schorfheide (2007), Schorfheide(2000)). Our procedure is closest in spirit to the procedure recently developed in Baumeister and Hamilton (2015).

procedure described above to estimate the shocks driving aggregate employment, prices, and wages during the Great Recession. Our results suggest that during the early part of the recession (2008-2009) roughly 30 percent of the aggregate employment decline can be attributed to the discount rate shock (i.e., the "demand" shock). The leisure shock explains roughly 30 percent of the decline in aggregate employment while the productivity/markup shock explains the remaining 40 percent. Over a longer period (2008-2012), however, the discount rate shock cannot explain any of the persistence in employment decline. Instead, it is the productivity/markup and labor supply shocks that explain why employment remained low from 2010-2012. In sum, while "demand" shocks may have been important in the early part of the recession, they cannot explain the persistently low levels of employment in the US after 2009.⁶ Furthermore, we find that the aggregate leisure shock - not sticky wages - explains why aggregate wages did not fall during the Great Recession.

Our paper contributes to many literatures. First, our work contributes to the recent surge in papers that have exploited regional variation to highlight mechanisms of importance to aggregate fluctuations. For example, Mian and Sufi (2011 and 2014), Mian, Rao, and Sufi (2013) and Midrigan and Philippon (2011) have exploited regional variation within the US to explore the extent to which household leverage has contributed to the Great Recession.⁷

6. Christiano et al (2015a) estimate a New Keynesian model using data from the recent recession. Although their model and identification are different from ours, they also conclude that something akin to a supply shock is needed to explain the joint aggregate dynamics of prices and employment during the Great Recession. Likewise, Vavra (2014) and Berger and Vavra (2015) document that prices were very flexible during the Great Recession. They also conclude that something more than a demand shock is needed to explain aggregate employment dynamics given the missing aggregate disinflation.

7. There has been an explosion of papers using regional data to better understand aggregate dynamics during the Great Recession. Some recent papers include: Giroud and Mueller (2015), Hagedorn et al. (2015), Mehrotra and Sergeyev (2015), and Mondragon (2015).

Nakamura and Steinsson (2014) use sub-national US variation to inform the size of local government spending multipliers. Blanchard and Katz (1991), Autor et al. (2013), and Charles et al. (2015) use regional variation to measure the responsiveness of labor markets to labor demand shocks. Our work contributes to this literature on two fronts. First, we show that local wages also respond to local changes in economic conditions at business cycle frequencies. Second, we provide a procedure where local variation can be combined with aggregate data to learn about the nature and importance of certain mechanisms for aggregate fluctuations. With respect to the latter innovation, our paper is similar in spirit to Nakamura and Steinsson (2014).

Second, our paper contributes to the recent literature trying to determine the causes of the Great Recession. In many respects, our model is more stylized than others in this literature in that we include a broad set of shocks without trying to uncover the underlying micro-foundations for these shocks. However, the shocks we chose to focus on were designed to proxy for many of the popular theories about the drivers of the Great Recession. For example, our discount rate shock can be thought of as reduced form representation of tightening of household borrowing limits. For example, such shocks have been proposed by Eggertsson and Krugman (2012), Guerrieri and Lorenzoni (2011) and Mian and Sufi (2014) as an explanation of the 2008 recession. Likewise, our productivity/markup shock can be interpreted as anything that changes firms' demand for labor. In a reduced form sense, credit supply shocks to firms, such as those proposed by Gilchrist et al (2014), would be similar to our productivity/markup shock. Finally, our leisure shock can be seen as a proxy for increased distortions in the labor market due to changes in government policy (e.g., Mulligan

(2012) or as a reduced form representation of a skill mismatch story within the labor market (e.g., Charles et al. (2013, 2015)).

2.3 Creating State-Level Price And Wage Indices

2.3.1 State-Level Wage Index

To construct nominal wage indices at the state level, we use data from the 2000 *Census* and the 2001-2012 *American Community Surveys* (ACS). The 2000 *Census* includes 5 percent of the US population while the 2001-2012 ACS's includes around 600,000 respondents per year between 2001 and 2004 and around 2 million respondents per year between 2005 and 2012. The large coverage allows us to compute detailed labor market statistics at the state level. For each year of the *Census/ACS* data, we calculate hourly nominal wages for prime-age males with a strong attachment to the labor force. In particular, we restrict our sample to only males between the ages of 21 and 55, who were employed at the time of the Census, who reported usually working at least 30 hours per week, and who worked at least 48 weeks during the prior 12 months. Then, for each individual in the resulting sample, we divide total labor income earned during the prior 12 months by a measure of annual hours worked during prior 12 months.⁸

Despite our restriction to prime-age males with a strong attachment to the labor force, the composition of workers on other dimensions may still differ across states and within

8. Total labor income during the prior 12 months is the sum of both wage and salary earnings and business earnings. Total hours worked during the previous 12 month is the product of total weeks worked during the prior 12 months and the respondents report of their usual hours worked per week.

a state over time. The changing composition of workers could be explaining some of the variation in nominal wages across states over time. To cleanse our wage indices from these compositional issues, we create a composition adjusted wage measure (at least based on observables) by running the following regression on the ACS data:

$$\ln(w_{itk}) = \gamma_t + \Gamma_t X_{it} + \eta_{itk}$$

where $\ln(w_{itk})$ is log nominal wages for household i in period t residing in state k and X_{it} is a vector of household specific controls. The vector of controls include a series of dummy variables for usual hours worked (with "40-49 hours per week" being the omitted group), a series of five year age dummies (with "40-44" being the omitted group), four educational attainment dummies (with "some college" being the omitted group), three citizenship dummies (with "native born" being the omitted group), and a series of race dummies (with "white" being the omitted group). We run these regressions separately for each year so that both the constant, γ_t , and the vector of coefficients on the controls, Γ_t , can differ for each year. Then, we take the residuals from these regressions, η_{itk} , and add back the constant, γ_t . Adding back the constant from the regression preserves differences over time in average log-wages. To compute average wages in a state holding composition fixed, we average $e^{\eta_{itk} + \gamma_t}$ across all individuals in state k . We refer to this measure as the "adjusted nominal wage index" in time t in state k . This is the series we use to exploit cross-state variation in wages during the Great Recession.

The benefit of the *Census/ACS* data is that it is large enough to compute detailed labor

market statistics at state levels. However, one drawback of the *Census/ACS* data is that it is not available at an annual frequency prior to 2000. To complement our analysis, we use data from the March Supplement of the *Current Population Survey (CPS)* to examine longer run aggregate trends in both nominal and real wages. These longer run trends are an input into our aggregate shock decomposition procedure discussed below. We compute the wage indices using the *CPS* data analogously to the way we computed the wage indices within the *Census/ACS* data.⁹ For the remainder of the paper, we use the *Census/ACS* data to explore regional wage variation and the *CPS* data to examine aggregate time series wage variation. However, for the 2000-2012 period, we can compare the time-series variation in aggregate wages using the *Census/ACS* data with the time series variation in aggregate wages using the *CPS* data. The two series have a correlation of 0.99 during this time period.

2.3.2 State-Level Price Index

2.3.2.1 Price Data

State-level price indices are necessary to measure state-level real wages. In order to construct state-level price indices we use the *Retail Scanner Database* collected by AC Nielsen and made available at The University of Chicago Booth School of Business.¹⁰ The Retail Scanner data

9. In particular, we compute hourly wages for men 21-55 with a strong attachment to the labor force (those currently working at least 30 hours a week and those who worked at least 48 weeks during the prior year). Again, like for the ACS data, we adjust the wages to account for a changing vector of observables over time. A full discussion of our methodology to compute composition adjusted wages in the *CPS* can be found in the Online Appendix that accompanies the paper.

10. The data is made available through the Marketing Data Center at the University of Chicago Booth School of Business. Information on availability and access to the data can be found at <http://research.chicagobooth.edu/nielsen/>. Contemporaneously, Coibion et al. (2015), Kaplan and Menzio (2015) and Stroebel and Vavra (2014) also use local scanner data/household price data to estimate that

consists of weekly pricing, volume, and store environment information generated by point-of-sale systems for about 90 participating retail chains across all US markets between January 2006 and December 2012. As a result, the database includes roughly 40,000 individual stores selling, for the most part, food, drugs and mass merchandise.

For each store, the database records the weekly quantities and the average transaction price for roughly 1.4 million distinct products. Each of these products is uniquely identified by a 12-digit number called Universal Product Code (UPC). To summarize, one entry in the database contains the number of units sold of a given UPC and the weighted average price of the corresponding transactions, at a given store during a given week. The database only includes items with strictly positive sales in a store-week and excludes certain products such as random-weight meat, fruits, and vegetables since they do not have a UPC assigned. Nielsen sorts the different UPCs into over one thousand narrowly defined "categories". For example, sugar can be of 5 categories: sugar granulated, sugar powdered, sugar remaining, sugar brown, and sugar substitutes. We use these categories when defining our price indices.

Finally, the geographic coverage of the database is outstanding and is one of its most attractive features. It includes stores from all states except for Alaska and Hawaii. Likewise, it covers stores from 361 Metropolitan Statistical Areas (MSA) and 2,500 counties. The data comes with both zip code and FIPS codes for the store's county, MSA, and state. Over the seven year period, the data set includes total sales across all retail establishments worth over \$1.5 trillion. In this paper, we aggregate data to the level of US states and compute

local prices vary with local economic conditions at business cycle frequencies. Our paper complements this literature by actually making price indices using the Nielsen scanner data for each state at the monthly frequency and using those price indices to estimate structural parameters of the local wage setting equation.

state-level retail scanner data price indices. Online Appendix Table R1 shows summary statistics for the retail scanner data for each year between 2006 and 2012 and for the sample as a whole.¹¹

2.3.2.2 A Retail Scanner Data Price Index

In order to construct state-level price indices we follow the BLS construction of the CPI as closely as possible.¹² While we briefly outline the price index construction in this sub-section, the full details of the procedure are discussed in the Online Appendix that accompanies our paper. Our retail scanner price indices are built in two stages. In the first stage, we aggregate the prices of goods within the roughly 1,000 categories described above. For our base indices, a good is a given store-UPC pair such that a UPC in store A is treated as a different good than the same UPC sold in store B. This allows for the possibility that prices may change as households substitute from a high cost store (that provides a different shopping experience) to a low cost store when local economic conditions deteriorate. Then, we compute, for each good, the average price and total quantity sold in a given month and state. Next, we construct the quantity weighted average price for all goods in each detailed category in a given month and state. We aggregate our index to the monthly level to reduce the number

11. The Online Appendix is available at http://faculty.chicagobooth.edu/erik.hurst/research/regional_online_appendix.pdf

12. There is a large literature discussing the construction of price indices. See, for example, Diewert (1976). Cage et al (2003) discuss the reasons behind the introduction of the BLS's Chained Consumer Price Index. Melsner (2011) discuss problems that arise with the construction of price indices with scanner data. In particular, if the quantity weights are updated too frequently the price index will exhibit "chain drift". This concern motivated us to follow the BLS procedure and keep the quantity weights fixed for a year when computing the first stage of our indices rather than updating the quantities every month. Such problems are further discussed in Dielwert et al. (2011).

of missing values.¹³

Specifically, for each category, we compute:

$$P_{j,t,y,k} = P_{j,t-1,y,k} \times \frac{\sum_{i \in j} p_{i,t,k} \bar{q}_{i,t-1,k}}{\sum_{i \in j} p_{i,t-1,k} \bar{q}_{i,t-1,k}} \quad (1)$$

where $P_{j,t,y,k}$ is category level price index for category j , in period t , with a base year y , in state k . $p_{i,t,k}$ is the price at time t of the specific good i (from category j) in state k and $\bar{q}_{i,t-1,k}$ is the average monthly quantity sold of good i in the prior year in state k . By fixing quantities at their prior year's level, we are holding fixed household's consumption patterns as prices change. We update the basket of goods each year and produce the chained index for each category in each state.

In the second stage of our construction we aggregate the category-level price indices into an aggregate index for each state k . The inputs are the category-level prices and the total expenditures of each category. Specifically, for each state we compute:

$$\frac{P_{t,k}}{P_{t-1,k}} = \prod_{j=1}^N \left(\frac{P_{j,t,y,k}}{P_{j,t-1,y,k}} \right)^{\frac{\bar{S}_{j,k}^t + \bar{S}_{j,k}^{t-1}}{2}} \quad (2)$$

where $\bar{S}_{j,k}^t$ is the share of expenditure of category j in month t in state k averaged over the year.

Finally, as a consistency check, we compare our retail scanner price index for the aggregate

13. One issue discussed in greater depth in the Online Appendix is how we deal with missing data when computing the price indices. Monthly prices may be missing, for instance, in the case of seasonal goods, the introduction of new goods, and the phasing out of existing goods. When computing our price indices, we restrict our sample to only include (1) goods that had positive sales in the prior year and (2) goods that had positive sales in every month of the current year. Online Appendix Table R1 shows the share of sales included in the price index for each sample year.

US to the BLS's CPI for food. We choose the BLS Food CPI as a benchmark given that approximately 60 percent of the goods in our database can be classified as food.¹⁴ Figure 2.1 shows that our retail scanner aggregate price index matches nearly exactly the BLS's Food CPI at the monthly level between 2006 and 2012.

2.3.2.3 A State-Level Price Index from the Retail Scanner Price Index

The previous subsection described the construction of a state-level price index for goods sold in retail grocery and mass merchandising stores. However, our goal is to construct state-level price indices that are representative of the composite basket of consumer goods and services. In this subsection, we describe conditions under which our retail scanner price index and a composite local price index differ only by a scaling factor. We then propose to estimate this scaling factor using available data from the BLS. Nonetheless, as we highlight throughout, using this scaling factor (as opposed to using our retail scanner price indices directly) has little effect on the quantitative results of the paper.

Most goods in our sample are produced outside a local market and are simultaneously sold to many local markets. These intermediate production costs represent the traded portion of local retail prices. If there were no additional local distribution and/or trade costs, one would expect little variation in retail prices across states; the law of one price would hold. However, these "non-tradable" costs do exist, including the wages of workers in the retail establishments, the rent of the retail facility, and expenses associated with local warehousing

14. The non-food goods in our sample include health and beauty products (13 percent), alcoholic beverages (6 percent), and paper products and household cleaning supplies (13 percent). The remaining items includes batteries, cutlery, pots and pans, candles, cameras, small consumer electronics, office supplies, and small household appliances.

and transportation.¹⁵

Assuming that the shares of these non-tradable costs are constant across states and identical for all firms in the retail industries, we can express local retail scanner prices, P^r , in region k during period t as:

$$P_{t,k}^r = (P_t^T)^{1-\kappa_r} (P_{t,k}^{NT})^{\kappa_r}$$

where P_t^T is the tradable component of local retail scanner prices in period t (which does not vary across states) and $P_{t,k}^{NT}$ is the non-tradable component of local retail prices in period t (which potentially does vary across states). κ_r represents the share of non-tradable costs in the total price for the retail scanner goods in our sample.

Analogously, we can express local prices in other sectors for which we do not have data as:

$$P_{t,k}^{nr} = (P_t^T)^{1-\kappa_{nr}} (P_{t,k}^{NT})^{\kappa_{nr}}$$

where $P_{t,k}^{nr}$ is local prices in these sectors outside of the grocery/mass-merchandising sector and κ_{nr} is the share of non-tradable costs in the total price for these other sectors.¹⁶

Next, assume that the price of household's composite basket of goods and services in a

15. Burstein et al (2003) document that distribution costs represent more than 40 percent of retail prices in the US.

16. The grocery/mass-merchandising sector is only one sector within a household's local consumption bundle. For example, there are other sectors where the non-tradable share may differ from those in our retail-scanner data. For example, many local services primarily use local labor and local land in their production (e.g., dry-cleaners, hair salons, schools, and restaurants). Conversely, in other retail sectors, the traded component of costs could be large relative to the local factors used to sell the good (e.g., auto dealerships).

state can be expressed as a composite of the prices in the retail scanner sectors ($P_{t,k}^r$) and prices in the other sectors (P_t^{nr}):

$$P_{t,k} = (P_t^{nr})^{1-s} (P_{t,k}^r)^s \equiv (P_t^T)^{1-\bar{\kappa}} (P_{t,k}^{NT})^{\bar{\kappa}}$$

where s is expenditure share of grocery/mass-merchandising goods in an individuals consumption bundle and $\bar{\kappa} \equiv (1-s)\kappa_{nr} + s\kappa_r$ is the non-tradable share in the aggregate consumption good, constant across all states.

Given these assumptions, we can transform the variation in retail scanner prices across states into variation in the broader consumption basket across states. Taking logs of the above equations and differencing across states we get that the variation in log-prices of the composite good between two states k and k' , $\Delta \ln P_{t,k,k'}$, is proportional to the variation in log-retail scanner prices across those same states, $\Delta \ln P_{t,k,k'}^r$. Formally,

$$\Delta \ln P_{t,k,k'} = \left(\frac{\bar{\kappa}}{\kappa_r} \right) \Delta \ln P_{t,k,k'}^r$$

If $\frac{\bar{\kappa}}{\kappa_r} > 1$, the local grocery/mass-merchandising sector will use a lower share of non-tradables in production than the composite local consumption good. In order to construct the scaling factor $\frac{\bar{\kappa}}{\kappa_r}$, it would be useful to have local indices for both grocery/mass-merchandising goods and for a composite local consumption good. While we do not have such indices for every US state, we can compare the relationship between local food inflation and local total inflation using BLS metro area price indices. These indices are only available

for 27 MSAs at varying degrees of frequency (monthly, bi-monthly, semi-annually).¹⁷ As a result, they are not overly useful in measuring prices for a broad set of local areas. However, for the MSAs covered, the BLS creates both a local food price index and a price index for the total local consumption basket. One approach to estimate $\frac{\bar{\kappa}}{\bar{\kappa}_r}$, therefore, would be to estimate a regression of local food inflation on local total inflation using data for these 27 MSAs. However, the BLS cautions against such a regression because they report that the local price indices contain a substantial amount of measurement error.¹⁸ Such measurement error will bias our estimate of $\frac{\bar{\kappa}}{\bar{\kappa}_r}$ towards zero.

To get around the measurement error problem, we follow the lead of Fitzgerald and Nicolini (2014) and regress food (total) inflation on some measure of local economic activity that is measured with relatively more precision. Taking a ratio of the coefficients from these two separate regressions can yield an estimate of $\frac{\bar{\kappa}}{\bar{\kappa}_r}$. Specifically, we regress the 3-year inflation rate (either for food or total CPI) at the MSA level on the 3 year change in the unemployment rate during the 2007-2010 period. Within the BLS data, we find that a 1 percentage point increase in the local unemployment rate is associated with a 0.34 percentage point decline in the local food inflation rate (standard error = 0.22) and 0.47 percentage point decline the local composite inflation rate (standard error = 0.15). These estimates are very similar to those reported by Fitzgerald and Nicolini (2014) who use data over a longer time period. The fact that the coefficient on the change in unemployment rate is smaller in the food inflation regression than the total inflation regression is consistent

17. In the online appendix that accompanies this paper, we discuss the BLS local price indices in greater depth.

18. See, for example, <http://www.bls.gov/opub/btn/volume-1/pdf/consumer-price-index-data-quality-how-accurate-is-the-us-cpi.pdf>

with our belief that the tradable share of food is higher than the tradable share of the local composite consumption good. Given these coefficients, the BLS data suggests a measure of $\frac{\bar{\kappa}}{\kappa_r}$ of 1.4 (-0.47/-0.34). We will use this as our base adjustment factor throughout the paper. However, our main decompositions later in the paper are robust to any scaling factor between 1.0 and 2.0.

2.4 Comparing Cross-State Patterns to Aggregate Time-series Patterns

The goal of this section is to contrast the strong co-movement of wages and economic activity at the local level to the relatively weaker co-movement at the aggregate level, during the Great Recession.

The left hand panel of Figure 2.3 shows the log-change in our demographic adjusted nominal wage indices between 2007 and 2010 across states against the log-change in the employment rate. As seen from the figure, nominal growth was strongly, positively correlated with employment growth in the 2007-2010 period. A simple linear regression through the data (weighted by the state's 2006 labor force) suggests that a 1 percent change in a state's employment rate is associated with a 0.62 percent change in nominal wages (standard error = 0.10). These findings are consistent with the extensive literature in labor economics and public finance showing that local labor demand shocks cause both employment and wages to vary together in the short to medium run. For example, Blanchard and Katz (1991), Autor, Dorn and Hanson (2013) and Charles, Hurst and Notowidigdo (2013) all find that

negative local labor demand shocks cause substantial declines in local wages over the three to five year horizon. Our results further suggest that wages are fairly flexible in response to labor demand shocks at the local level. However, we illustrate the patterns at business cycle frequencies.¹⁹

The right hand panel of Figure 2.3 shows similar patterns for real wage variation. We compute local real wages by deflating local nominal wage growth with the growth in the prices of a composite local consumption good ($P_{t,k}$).²⁰ A simple linear regression through the data (weighted by the state's 2006 labor force) suggests that a 1 percent change in a state's employment rate is associated with a 0.52 percent change in real wages (standard error = 0.15). Growth in local nominal and real wages were highly correlated with changes in many other measures of state economic activity during the 2007-2010 period as well. Although not shown, lower GDP growth, lower unemployment growth, lower hours growth and lower house price growth were all strongly correlated with lower nominal and real wage growth during the recent recession.

Figure 2.2 shows our composition adjusted aggregate wage indices for the 2000 to 2012 period calculated using CPS data. To construct aggregate composition adjusted real wages, we deflate the aggregate nominal adjusted wages from the CPS by the aggregate June CPI-U

19. The patterns we document in Figure 2.3 also show up in other wage series. While there are no government data sets that produce broad based composition adjusted wage series at the local level, the Bureau of Labor Statistics's Quarterly Census of Employment and Wages (QEW) collects firm level data on employment counts and total payroll at local levels. In Online Appendix Figures R1 and R2 we present results using local wage indices constructed from the QEW data instead. In these data, a one percent increase in a state's employment growth between 2007 and 2010 was associated with a roughly 0.5 increase in the state's nominal per capita earnings growth during the same time period.

20. As discussed in the previous section, we scale the growth in the retail scanner price index by a factor of 1.4 to account for the fact that grocery/mass merchandising goods have a higher tradable share than the composite local consumption good.

with 2000 as the base year. Between 2007 and 2010, average composition adjusted nominal wages in the US increased by roughly 4 percent despite aggregate employment falling substantially. The patterns in our data replicate the aggregate nominal wage growth patterns documented by many others in the literature.²¹ Given that consumer prices increased by 5 percent during the same period, aggregate real wages in the US fell by roughly 1 percent between 2007 and 2010. This was similar to the trend in real wages prior to the start of the recent recession. As seen from Figure 2.2 nominal wages increased slightly and real wage growth did not seem to break trend during the Great Recession. The "puzzle" is why aggregate wages did not decline relative to trend despite the very weak aggregate labor market. Wage stickiness is one potential explanation. However, as seen from Figure 2.3, local nominal and real wages moved quite a bit with changes in local employment during the same time period.

Table 2.1 compares these cross-state elasticities with the corresponding aggregate time-series elasticities during the Great Recession.²² The top panel displays the local wage elasticities from the simple scatter plots shown in Figure 2.3. The bottom panel provides an estimate of similar elasticities over the same time period at the aggregate level. In particular, the last row shows the aggregate nominal (and real) wage elasticity with respect to changes in employment between 2007 and 2010. To construct these elasticities we use our adjusted nominal wage measure from the *CPS* (in the case of real wages, we deflate them with June CPI-U) and the aggregate employment-to-population ratio from the BLS. We de-trend all

21. See, for example, Daly and Hobijn (2015).

22. We thank Bob Hall for giving us the idea for this table. We base it on the analysis he did as part of his discussion of our paper at the 2015 NBER summer EFG program meeting.

variables by estimating a linear trend between 2000 and 2007. The de-trended employment decline between 2007 and 2010 was 6.8 percent whereas the de-trended nominal wage decline was 1.7 percent. De-trended real wages actually increased by 1.2 percent during the 2007-2010 time period. Therefore, the implied aggregate wage elasticities with respect to employment during the Great Recession are 0.25 for nominal wages $(-1.7/-6.8)$ and -0.17 for real wages $(1.2/-6.8)$.

Our main empirical finding comes from comparing the cross-state wage elasticities with the aggregate wage elasticities. The response of wages to changes in employment were *much* stronger at the state level during the Great Recession than at the aggregate level. For example, the local nominal wage elasticity with respect to employment changes was over twice as big as the aggregate elasticity (0.62 vs. 0.25). It is these differences in the relationships between wages and employment at the local level and at the aggregate level that forms the basis of the remainder of this paper. Why did local wages adjust so much when local employment conditions deteriorated during the Great Recession while aggregate wages hardly responded at all despite a sharp deterioration in aggregate employment conditions? Can aggregate wages be sticky when local wages adjust so much? We turn to answering these questions next.

2.5 A Monetary Union Model

In this section we present a monetary union model with several goals in mind. First, the model allows us to discuss the patterns we documented in the previous section in a formal environment where local economies aggregate. Second, the model makes explicit our as-

assumptions on how wages are set. The nominal wage stickiness we specify will be essential to our identification strategy in later parts of the paper. Third, a calibrated version of the model allows us to quantify differences in aggregate vis a vis local elasticities to a variety of different shocks. While the theoretical possibility of these differences are known, much less is known about their magnitudes. The calibration exercise provides guidance to researchers who want to take an estimated local elasticity to a given shock and apply it to the aggregate economy. Fourth, the model provides an example of an economy that is encompassed by our SVAR procedure in Section 5 of the paper. The SVAR approach will allow us to estimate shocks for a larger set of models than the one we write down in this section. Finally, the model provides us with theoretical co-movements between variables that help us identify the shocks in the SVAR as well as give them an economic interpretation.

Formally, our model economy is composed of many islands inhabited by infinitely lived households and firms in two distinct sectors that produce a final consumption good and intermediates that go into its production. The only asset in the economy is a one-period nominal bond in zero net supply where the nominal interest rate is set by a monetary authority. We assume intermediate goods can be traded across islands but the consumption good is non-tradable.²³ Finally, we assume labor is mobile across sectors but not across islands.²⁴ Throughout we assume that parameters governing preferences and production are identical across islands and that islands only differ, potentially, in the shocks that hit them.

23. The final good can be thought of as being retail goods and services purchased in places such as: restaurants, barbershops and stores; and the intermediate sector providing physical goods such as: food ingredients, scissors and cellphones.

24. We explore the issue of labor mobility during the Great Recession when we take the model to the data.

2.5.1 Firms and Households

Producers of tradable intermediate goods x in island k use local labor N_k^x and face nominal wages W_k (equalized across sectors) and prices Q (equalized across islands k). The time subscripts are omitted for clarity. Their profits are

$$\max_{N_k^x} Q e^{z_k^x} (N_k^x)^\theta - W_k N_k^x$$

where z_k^x is a tradable productivity shock in island k and $\theta < 1$ is the labor share in the production of tradables. Final (retail) goods y producers face prices P_k and obtain profits

$$\max_{N_k^y, X_k} P_k e^{z_k^y} (N_k^y)^\alpha (X_k)^\beta - W_k N_k^y - Q X_k$$

where z_k^y is a final good (retail) productivity shock and $(\alpha, \beta) : \alpha + \beta < 1$ are the labor and intermediates shares. Unlike the tradable goods prices, final good prices (P_k) vary across islands.²⁵

Households preferences are given by

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} e^{-\rho_{kt} - \delta_{kt}} \frac{(C_{kt} - e^{\epsilon_{kt}} \frac{\phi}{1+\phi} N_{kt}^{\frac{1+\phi}{\phi}})^{1-\sigma}}{1-\sigma} \right]$$

25. It is worth noting that all model shocks will generate endogenous variation in markups given our assumption of decreasing returns to scale. Additionally, what we call a "productivity shock" is isomorphic to any shifter of unit labor costs and, hence, labor demand schedules. Later we will refer to it as the productivity/markup shock. We do not attempt to distinguish between the different interpretations of this shock in this paper.

where C_{kt} is consumption of the final good, N_{kt} is labor, and δ_{kt} and ϵ_{kt} are exogenous processes driving the household's discount factor and the disutility of labor, respectively. Our base preferences abstract from income effects on labor supply. However, we show in section 7.4 that relaxing this assumption does not quantitatively change the conclusions of the paper.

Households are able to spend their labor income $W_{kt}N_{kt}$, profits accruing from firms Π_{kt} , financial income $B_{kt}i_t$, and transfers from the government T_t , where B_{kt} are nominal bond holdings at the beginning of the period and i_t is the nominal interest (equalized across islands given our assumption of a monetary union where the bonds are freely traded). Thus, they face the period-by-period budget constraint

$$P_{kt}C_{kt} + B_{kt+1} \leq B_{kt}(1 + i_t) + W_{kt}N_{kt} + \Pi_{kt} + T_t$$

A well known issue in the international macroeconomics literature is that under market incompleteness of the type we just described there is no stationary distribution for bond holdings across islands in the log-linearized economy; and all other island variables in the model have unit roots. We follow Schmitt-Grohe and Uribe (2003) and let ρ_{kt} be the endogenous component of the discount factor that satisfies $\rho_{kt+1} = \rho_{kt} + \Phi(\cdot)$ for some function $\Phi(\cdot)$ of the average per capita variables in an island. As such, agents do not internalize this dependence when making their choices. This modification induces stationarity for an appropriately chosen function $\Phi(\cdot)$. Schmitt-Grohe and Uribe (2003) show that alternative stationary inducing modifications (a specification with internalization, a debt-elastic inter-

est rate or convex portfolio adjustment costs) all deliver similar quantitative results in the context of a small open economy real business cycle model.

2.5.2 *Sticky Wages*

We allow for the possibility that nominal wages are sticky and use a partial-adjustment model where a fraction λ of the gap between the actual and frictionless wage is closed every period. Formally:

$$W_{kt} = (P_{kt}e^{\epsilon_{kt}}(N_{kt})^{\frac{1}{\phi}})^{\lambda}(W_{kt-1})^{1-\lambda}$$

Given our assumption on household preferences, $P_{kt}e^{\epsilon_{kt}}(N_{kt})^{\frac{1}{\phi}}$ corresponds to the marginal rate of substitution between labor and consumption and the parameter λ measures the degree of nominal wage stickiness. In particular, when $\lambda = 1$ wages are fully flexible and when $\lambda = 0$ they are fixed. This implies that workers will be off their labor supply curves whenever $\lambda < 1$. A similar specification has been used by Shimer (2010) and, more recently, by Midrigan and Philippon (2011). Shimer (2010) argues that in labor market search models there is typically an interval of wages that both the workers are willing to accept and firms willing to pay. To resolve this wage indeterminacy he considers a wage setting rule that is a weighted average of a target wage and the past wage. The target wage in our case is the value of the marginal rate of substitution.

Popular alternatives in the literature include the wage bargaining model in the spirit of Hall and Milgrom (2008) as in Christiano, Eichenbaum and Trabandt (2015b); and the

monopsonistic competition model where unions representing workers set wages period by period as in Gali (2009). The key difference with the partial adjustment model is that both alternatives result in a forward looking component in the wage setting rule that is absent in our specification. In fact, this wage setting rule can be derived from the monopsonistic competition setup in the case where agents are myopic about the future; or from the labor market search setup in the special case where firms make take-it-or-leave-it offers and the probability of being employed in the future is independent of the current employment status.²⁶

2.5.3 *Equilibrium*

An equilibrium is a collection of prices $\{P_{kt}, W_{kt}, Q_t\}$ and quantities $\{C_{kt}, N_{kt}, B_{kt}, N_{kt}^x, N_{kt}^y, X_{kt}\}$ for each island k and time t such that, for an interest rate rule $i_t = i(\cdot)e^{\mu t}$ and given exogenous processes $\{z_{kt}^x, z_{kt}^y, \epsilon_{kt}, \delta_{kt}, \mu_t\}$, they are consistent with household utility maximization and firm profit maximization and such that the following market clearing conditions hold:

$$C_{kt} = e^{z_{kt}^y} (N_{kt}^y)^\alpha X_{kt}^\beta$$

$$N_{kt} = N_{kt}^y + N_{kt}^x$$

$$\sum_k X_{kt} = \sum_k e^{z_{kt}^x} (N_{kt}^x)^\theta$$

26. While there is no forward looking component in the reset wage in our base specification, we consider the implications of including forward looking behavior in Section 7.4.

2.5.4 Shocks

We assume exogenous processes are AR(1) processes, with an identical autoregressive coefficient across islands (and sectors in the case of productivity), and that the innovations (i.e., shocks) to these processes are iid, mean zero, random variables with an aggregate and island specific component. Let $\gamma_{kt} \equiv \delta_{kt} - \delta_{kt-1} - \mu_t$ be a combination of the discount rate exogenous growth and the monetary policy exogenous process that shows up as a wedge in the Euler equation. Then, we write the exogenous processes as:

$$\begin{aligned} z_{kt}^y &= \rho_z z_{kt-1}^y + \sigma_z u_t^y + \tilde{\sigma}_y v_{kt}^y \\ z_{kt}^x &= \rho_z z_{kt-1}^x + \sigma_z u_t^x + \tilde{\sigma}_x v_{kt}^x \\ \gamma_{kt} &= \rho_\gamma \gamma_{kt-1} + \sigma_\gamma u_t^\gamma + \tilde{\sigma}_\gamma v_{kt}^\gamma \\ \epsilon_{kt} &= \rho_\epsilon \epsilon_{kt-1} + \sigma_\epsilon u_t^\epsilon + \tilde{\sigma}_\epsilon v_{kt}^\epsilon \end{aligned}$$

with $\sum_k v_{kt}^y = \sum_k v_{kt}^x = \sum_k v_{kt}^\gamma = \sum_k v_{kt}^\epsilon = 0$. By assumption, we assume the weighted average of the island specific shocks sums to zero in all periods.

Let $u_t^z \equiv u_t^y + \beta u_t^x$ be a combination of productivity shocks in both sectors. We will call u_t^z , u_t^γ and u_t^ϵ the aggregate *Productivity/Markup*, *Discount rate* and *Leisure shocks* respectively. These are the shocks that the econometric procedure aims to identify. Analogously, v_{kt}^y , v_{kt}^x , v_{kt}^γ , v_{kt}^ϵ are the *Regional shocks*. The interpretation of the Leisure and Productivity/Markup shocks is relatively straightforward given our model environment. They are shifters of households and firms' labor supply (wage setting) and labor demand schedules,

respectively. On the other hand, what we identify as a "discount rate shock" (γ_{kt}) is the combination of two more fundamental shocks. First, a shock to the marginal rate of substitution between consumption in consecutive periods. Second, a shock to the nominal interest rate rule set by the monetary authority. Our procedure is unable to distinguish between the two given that they both show up in the household's Euler equation, thus we treat them as a single shock.

2.5.5 Aggregation

Our first key assumption for aggregation is that all islands are identical with respect to their underlying production parameters (α , β , and θ), their underlying utility parameters (σ and ϕ) and the degree of wage stickiness (λ).²⁷ Our second assumption is that islands are identical in the steady state and that price and wage inflation are zero. The last assumption is that the joint distribution of island-specific shocks is such that its cross-sectional sum is zero. If K , the number of islands, is large this holds in the limit because of the law of large numbers. We log-linearize the model around this steady state and show that it aggregates up to a representative economy where all aggregate variables are independent of any cross-sectional considerations to a first order approximation.²⁸ We denote with lowercase letters

27. When implementing our procedure using data on US states, we discuss the plausibility of this assumption. Given that the broad industrial composition at the state level does not differ much across states, the assumption that productivity parameters and wage stickiness are roughly similar across states is not dramatically at odds with the data. As a robustness exercise, we estimate our key equations with industry fixed effects and show that our key cross section estimates are unchanged.

28. The model we presented has many islands subject to idiosyncratic shocks that cannot be fully hedged because asset markets are incomplete. By log-linearizing the equilibrium, we gain in tractability but ignore these considerations and the aggregate consequences of heterogeneity. The approximation will be good as long as the underlying volatility of the idiosyncratic shocks is not too large. If our unit of study was an individual, as for example in the precautionary savings literature with incomplete markets, the use of linear approximations would likely not be appropriate. However, since our unit of study is an island the size

a variable's log-deviation from its steady state. Variables without a k subscript represent aggregates. For example, $n_{kt} \equiv \log\left(\frac{N_{kt}}{N}\right)$ and $n_t \equiv \sum_k \frac{1}{K} n_{kt}$. We assume that the monetary authority announces a nominal interest rate rule which is a function of aggregate variables. In log-linearized form, the rule is: $i_{t+1} = \varphi_\pi \mathbb{E}_t[\pi_{t+1}] + \varphi_y(y_t - y_t^*) + \mu_{t+1}$ where π_t is the aggregate inflation rate and $y_t - y_t^*$ is the output gap, defined as the difference between actual output and the flexible wage equilibrium output for the same realization of shocks. Finally, we assume that the endogenous component of the discount factor is $\Phi(\cdot) = \Phi_0(c_{kt} - c_t)$.²⁹

The following lemmas present a useful aggregation result and show that we can write the island-level equilibrium in log-deviation from the aggregate union equilibrium. Let w_t^r be real wage growth and π_t^w be nominal wage growth. Formally, $w_t^r \equiv \log\left(\frac{W_t/P_t}{W/P}\right)$ and $\pi_t^w = w_t^r - w_{t-1}^r + \pi_t$.

Lemma 1. *The behavior of π_t^w, w_t^r, n_t in the log-linearized economy is identical to that of a representative economy with only a final goods sector with labor share in production $\alpha + \theta\beta$, no endogenous discount factor, and only 3 exogenous processes $\{z_t, \epsilon_t, \gamma_t\}$.*

Denote any variable $\tilde{x}_t \equiv x_{kt} - x_t$ as corresponding to island's k log-deviation from aggregates at time t , where the subscript k is dropped for notational simplicity.

Lemma 2. *For given $\{\tilde{z}_t^y, \tilde{z}_t^x, \tilde{\gamma}_t, \tilde{\epsilon}_t\}$, the behavior of $\{\tilde{p}_t, \tilde{w}_t, \tilde{n}_t^y, \tilde{n}_t^x\}$ in the log-linearized economy for each island in deviations from aggregates is identical to that of a small open*

of a small country or a state, we believe this is not too egregious of an assumption. The volatilities of key economic variables of interest at the state or country level are orders of magnitude smaller than the corresponding variables at the individual level.

29. $\Phi_0 > 0$ is enough to induce stationary of island-level variables in log-deviations from the aggregate. Furthermore, since $\Phi(\cdot)$ depends only on these deviations, the aggregate equilibrium will feature a constant endogenous discount factor ρ .

economy where the price of intermediates and the nominal interest rate are at their steady state levels, i.e. $q_t = i_t = 0 \forall t$.

Proof. See Appendix 2.11.1 for a proof of Lemma 1 and 2. □

2.5.6 Aggregate vs. Local Shock Elasticities

Having described the model, we now explore the extent to which aggregate employment, price and wage elasticities to a given shock differ from local employment, price and wage elasticities to the same shock. Many researchers use clever identification strategies exploiting regional variation to estimate local elasticities to a given shock. For example, Mian and Sufi (2014) uses variation in debt across US metropolitan areas to isolate the extent to which a local "demand" shock (i.e., our discount rate shock) affects local employment. We show in this sub-section that the local employment (price, wage) elasticity to a given discount rate shock (productivity shock, leisure shock) is different, in general, from the aggregate elasticity to the same shock. Moreover, we calibrate the model in order to quantify the difference.

To gain some intuition as to the difference between local and aggregate elasticities in our model, we first consider the special case where there is an endowment of the tradable good and no labor is used in its production, i.e. $\theta = 0$. Focusing on a discount rate shock in this special case makes the comparison very transparent. We let $\xi_0^{agg} \equiv \frac{dn_0}{d\gamma_0}$ and $\xi_0^{reg} \equiv \frac{d\tilde{n}_0}{d\tilde{\gamma}_0}$ be the employment elasticities to the discount rate shock on impact. By solving for the

recursive laws of motion in equilibrium we obtain,

$$\xi_0^{agg} = \frac{(1 - \lambda)}{(1 - \alpha + \frac{\lambda}{\phi})(\varphi_p - 1) + (\varphi_y \alpha - (\varphi_p - 1)(1 - \alpha)) \frac{1 - \lambda}{\rho_\gamma}}$$

$$\xi_0^{reg} = \frac{d\tilde{n}_0}{d\tilde{\gamma}_0} = \frac{(1 - \lambda + \lambda\beta)}{\left(1 - \alpha + \frac{\lambda}{\phi} + \left(\frac{\sigma(1 - \lambda(\alpha - \frac{1}{\phi}))}{1 - \frac{\alpha\phi}{1 + \phi}} - (1 + \frac{\lambda}{\phi})\right) \beta\right) \left(\frac{1 + r}{\rho_\gamma} - 1\right)}$$

These expressions help understand the general equilibrium forces that make local and aggregate elasticities different. From the perspective of the closed economy, the endogenous response of the nominal interest rate rule $\{\varphi_p$ and $\varphi_y\}$ reduces the aggregate employment impact elasticity to an unanticipated discount rate shock. A negative discount rate shock puts downward pressure on employment and prices. The monetary authority can lower interest rates to offset such a shock. The parameters of the interest rate rule are entirely absent in the expression for the regional elasticity. Therefore, the aggregate employment elasticity to a discount rate shock is typically smaller than the local employment elasticity to a local discount rate shock.

From the local perspective, since island level economies in deviations from the aggregate are small open economies, there are two extra margins of adjustment that are absent in the aggregate closed economy. First, the possibility to substitute labor for intermediate goods in the production of final consumption goods ($\beta > 0$) decreases the regional employment elasticity to the shock (as long as the term $\left(\frac{\sigma(1 - \lambda(\alpha - \frac{1}{\phi}))}{1 - \frac{\alpha\phi}{1 + \phi}} - 1\right)$ is positive). Second, the possibility to transfer resources intertemporally through saving/borrowing at the interest rate r , as seen in the term $\left(\frac{1 + r}{\rho_\gamma} - 1\right)$, decreases the regional employment elasticity. Theoretically,

therefore, the aggregate employment elasticity to an aggregate discount rate shock can be either greater or smaller than the local employment elasticity to a local discount rate shock.

It is also interesting to compare how these discount rate elasticities change with the degree of nominal wage stickiness. Our identification procedure allows us to do this exercise when we estimate the impulse response to a discount rate shock. When $\varphi_p > 1$, both elasticities are decreasing in λ . In particular, employment does not respond to discount rate shocks at all in the limit when wages are perfectly flexible ($\lambda \rightarrow 1$).

While it is generally understood that local and aggregate elasticities can differ, there has been little quantitative work assessing the potential size of these differences. A parameterized version of our model allows us to directly compute the local and aggregate employment elasticities to different types of shocks. To this end, Table 2.3 quantifies the employment impact elasticities to each of the shocks in the full model. Table 2.2 presents and explains the parameterization of our model. Most of the parameters are standard from the literature or are chosen to match the labor share in the tradable and non-tradable sectors. The Online Appendix has an extended discussion of our baseline parameter choice. For our base specification, we use estimates of λ and ϕ of 2 and 0.7, respectively. These are the parameters that show up in the aggregate and local wage setting equations. The value of these parameters are the ones that we estimate using local variation in Section 6.

Column 1 of Table 2.3 shows our base estimates of the local and aggregate employment elasticities. In columns 2 - 8 of Table 2.3, we show how the elasticities change across alternate parameterization. Specifically, in column 2, we re-compute the elasticities reducing the Frisch elasticity of labor supply (ϕ) from 2 to 1. In column 3, we make wages more sticky by reducing

λ from 0.7 to 0.5 (returning the Frisch elasticity to our base parameterization). In column 4, we set $\beta = 0$, thus shutting down the possibility to substitute labor for intermediate goods in the production of final goods. In the next two columns, we shut down the endogenous feedback in the nominal interest rate to changes in the employment gap such that φ_y is set to zero. In the first of those two columns, we leave the response of the nominal interest rate to the inflation target (φ_p) at its base parameterization. In the second of those two columns, we lower φ_p such that the local and aggregate responses to a discount rate shock are the same on impact. Finally, in the last two columns, we explore how the elasticities change as the persistence of the demand shock changes.

In our base specification, we find that the regional employment elasticity to a discount rate shock is 2.3 times larger than the aggregate employment elasticity to a discount rate shock. This implies that using cross-region variation to estimate local employment elasticities to demand shocks dramatically overstates employment responses when those local elasticities are applied to the aggregate. The conclusion remains unchanged across the different parameterizations of the wage setting rule, as shown in columns 2 and 3. Local employment elasticities to discount rate shocks are always two to three times larger than the aggregate employment elasticities. In columns 4 to 6, we see the importance of general equilibrium forces. As we shut down the ability to substitute labor for intermediate goods ($\beta = 0$), the gap between the regional and aggregate elasticities gets larger. The ability to trade intermediates across regions dampens the local employment elasticity to discount rate (demand) shocks. In columns 5 and 6, we see that the endogenous monetary policy response also dramatically dampens the aggregate response to a discount rate shock. This

suggests that in periods where the economy is at the zero lower bound, aggregate and local employment elasticities to a demand shock are more similar, a point also made in Nakamura and Steinsson (2014). The last column explores the sensitivity to changes in the persistence of the discount rate shock. The less persistent is the local discount rate shock, the smaller the local employment elasticity because the regions can borrow from and lend to each other. Table 2.3 also shows the local and aggregate employment response to local and aggregate productivity/markup and leisure shocks. For these two shocks, the local employment elasticities are usually *smaller* than their aggregate counterparts. For the most part, this results from the particular specification of the nominal interest rate rule. To summarize, the quantitative difference between aggregate and local employment elasticities depends on the type underlying shock and can be quite large.

Tables 2.4 and 2.5 summarize the aggregate and regional impulse responses, respectively, for all variables and shocks in our benchmark calibration. We show results upon impact (the "short-run" elasticities) and after 5 years (the "long-run" elasticities). These tables allow us to assess the model's predictions. We use the same parameterization as in Table 2.2. The short run responses in Columns 1 of Table 2.4 and Table 2.5 just restate the employment elasticities in column 1 of Table 2.3. The remainder of the tables show the estimates for the price, nominal wage and real wage elasticities to all the underlying shocks in the model upon impact. As seen from Table 2.4, an aggregate negative discount rate shock (households become less patient) lowers aggregate employment, lowers aggregate prices, and lowers (slightly) aggregate real wages. Conversely, an aggregate negative productivity shock lowers aggregate employment, raises aggregate prices, and raises aggregate real wages. We

will use the sign of these impact elasticities to help identify the shocks in the SVAR in Section 2.6.

2.6 A Procedure for Identifying Aggregate Shocks

The above model was designed to (1) link aggregate and regional economies, (2) specify the local and aggregate wage setting equations, (3) provide a quantitative assessment of aggregate and local elasticities to the shocks embedded in the model, and (4) guide our interpretation of the shocks. In this section, we develop a procedure that allow us estimate the shocks in a larger class of monetary union models than the benchmark model outlined above, thus imposing less a-priori structure and making the analysis more persuasive.³⁰ Specifically, we consider models where aggregate equilibria can be represented as a structural vector autoregression (SVAR) in price inflation, nominal wage inflation, and employment with three shocks. In order to identify the shocks, we use three properties of our benchmark monetary union model: the wage setting equation, the sign of the impact elasticities to a discount rate and productivity/markup shocks, and the orthogonality of shocks. Our results will be consistent with monetary union models that satisfy all of these. Beraja (2015) discusses this identification procedure in detail, as well as its application to more general SVARs and

30. We also performed a business cycle accounting exercise by solving the model and using the data to recover the exogenous stochastic processes. By doing so, we learned that different "wedges" are required to explain the joint dynamics of employment prices and wages. We find that the labor wedge is quite important in explaining variation in employment in the early stages of the Great Recession. Like our SVAR results, the Euler equation wedge only explained less than half of the employment decline during the early part of the recession and explained essentially none of the persistence. However, unlike the SVAR, the business cycle accounting does not allow us to recover the fundamental shocks and, therefore, we do not report more specific results here. As has been shown in Buera and Moll (2015), slight changes in model specification can alter the mapping between underlying structural shocks and corresponding aggregate wedges.

theoretical models than the ones in this paper.

We begin by noting that the recursive solution to the equilibrium system of equations in Lemma 1 can be written as a SVAR(∞) in $\{\pi_t, \pi_t^w, n_t\}$.³¹

$$(I - \rho(L)) \begin{bmatrix} \pi_t \\ \pi_t^w \\ n_t \end{bmatrix} = \Lambda \begin{bmatrix} u_t^\epsilon \\ u_t^z \\ u_t^\gamma \end{bmatrix}$$

Knowledge of $\rho(L)$ and an invertible matrix Λ together with aggregate data on prices, nominal wages and employment allow recovering the structural shocks.

The first step in our procedure consists of estimating the reduced form VAR to obtain the autoregressive matrix $\rho(L)$ and the reduced form errors covariance matrix V . In practice we will truncate $\rho(L)$ to be of finite order as it is typically done in the literature. The second step involves deriving a set of theoretical restrictions to identify the structural shocks from the reduced form errors.

As a reminder, the wage setting equation³² in log-linearized form is:

$$\pi_t^w = \lambda(\pi_t + \epsilon_t - \epsilon_{t-1} + \frac{1}{\phi}(n_t - n_{t-1})) + (1 - \lambda)\pi_{t-1}^w$$

Applying the conditional expectation operator $\mathbb{E}_{t-1}(\cdot)$ on both sides and constructing ex-

31. The exogenous processes are AR(1) and the system of equations characterizing the equilibrium is of first order. When written in matrix form it is easy to show that there is a representation as a SVAR(∞).

32. At the end of Section 7, we show the sensitivity of our estimation procedure to alternative wage setting equations.

pectational errors, we obtain:

$$\begin{bmatrix} \lambda & -1 & \frac{\lambda}{\phi} \end{bmatrix} \Lambda \begin{bmatrix} u_t^\epsilon \\ u_t^z \\ u_t^\gamma \end{bmatrix} + \lambda \sigma_\epsilon u_t^\epsilon = 0 \quad (3)$$

Similarly, constructing $\mathbb{E}_{t-1}(\cdot) - \mathbb{E}_{t-2}(\cdot)$, we obtain:

$$\left(\begin{bmatrix} \lambda & -1 & \frac{\lambda}{\phi} \end{bmatrix} \rho_1 + \begin{bmatrix} 0 & 1 - \lambda & 0 \end{bmatrix} \right) \Lambda \begin{bmatrix} u_{t-1}^\epsilon \\ u_{t-1}^z \\ u_{t-1}^\gamma \end{bmatrix} + \lambda(\rho_\epsilon - 1)\sigma_\epsilon u_{t-1}^\epsilon = 0 \quad (4)$$

where ρ_1 is the matrix collecting the first order autoregressive coefficients in the reduced form VAR.

The above equations (3) and (4) have to hold for all realizations of the shocks. In particular, equation (3) gives us two linear restrictions in the elements of Λ for given parameters in the wage setting equation when there are either contemporaneous discount rate or productivity/markup shocks. These two restrictions, together with the six restrictions coming from the orthogonalization of the shocks, are sufficient to identify the column in the impulse response matrix Λ corresponding to the leisure shock (u_t^ϵ). In order to identify the discount rate and productivity/markup shocks (u_t^γ, u_t^z) we proceed as follows. From equation (4), we obtain two extra linear restrictions that hold when there is a lagged discount rate shock or a lagged productivity/markup shock. However, these restrictions alone cannot "separate" the discount rate from the productivity/markup shocks because they are identical for both.

Therefore, we use the sign of the impact elasticities from our model to a discount rate and productivity/markup shock (u_t^γ and u_t^y), respectively. Specifically, we search over all linear combinations $\psi \in [0, 1]$ of the independent restrictions coming from equation (4) such that a discount rate (productivity shock/markup) shock: (i) moves prices and employment in the same (opposite) direction on impact, and (ii) moves real wages and employment in opposite (same) direction on impact. If more than one linear combination of the restrictions satisfy these, we pick the one that is closer to giving equal weighting to both restrictions.

For completeness, the matrix Λ solves the system:

$$\begin{aligned} \begin{bmatrix} \lambda & -1 & \frac{\lambda}{\phi} \end{bmatrix} \Lambda \begin{bmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} &= \begin{bmatrix} 0 & 0 \end{bmatrix} \\ \left(\begin{bmatrix} \lambda & -1 & \frac{\lambda}{\phi} \end{bmatrix} \rho_1 + \begin{bmatrix} 0 & 1 - \lambda & 0 \end{bmatrix} \right) \Lambda \begin{bmatrix} 0 \\ \psi \\ 1 - \psi \end{bmatrix} &= 0 \\ \Lambda \Lambda' &= V \end{aligned}$$

It is worth noting that, when $\lambda = 1$, this procedure cannot identify all columns in the impulse response matrix Λ because the system above is underdetermined (i.e., (4) implies linear restrictions that are merely linear combinations of the restrictions implied by equation (3)). Therefore, some degree of wage stickiness is key for identification of the shocks through this procedure. The next section shows how to estimate λ using regional data—thus linking

particular regional patterns to particular aggregate shock decompositions when combined with the procedure in this section.

2.7 Estimating the Wage Setting Equation Using Regional Data

In this section, we discuss how we estimate λ and ϕ which are necessary inputs in our shock identification procedure. Given the above assumptions, the aggregate and local wage setting equations can be expressed as:

$$\pi_t^w = \lambda(\pi_t + \frac{1}{\phi}(n_t - n_{t-1})) + (1 - \lambda)\pi_{t-1}^w + \lambda(u_t^\epsilon - (1 - \rho_\epsilon)\epsilon_{t-1})$$

$$\pi_{kt}^w = \lambda(\pi_{kt} + \frac{1}{\phi}(n_{kt} - n_{kt-1})) + (1 - \lambda)\pi_{kt-1}^w + \lambda(u_t^\epsilon - (1 - \rho_\epsilon)\epsilon_{t-1}) + \lambda v_{kt}^\epsilon$$

The aggregate and local wage setting curves are functions of the Frisch elasticity of labor supply (ϕ) and the wage stickiness parameter (λ). There is a literature on estimating micro and macro labor supply elasticities. However, it is hard to estimate the degree of wage stickiness using aggregate data given the small degrees of freedom inherent to aggregate data and given that at the aggregate level it is hard to isolate movements in employment growth and price growth that are arguably uncorrelated with the aggregate leisure shock (u_t^ϵ). In some instances, regional data can be used to estimate these parameters.

In order for regional data to be used to estimate λ and ϕ , one of the following must hold: either (1) the leisure shock has no regional component ($v_{kt}^\epsilon = 0$) or (2) the regional

component of the leisure shock must be uncorrelated with changes in local economic activity (i.e., $cov(v_{kt}^\epsilon, (n_{kt} - n_{kt-1})) = 0$ and $cov(v_{kt}^\epsilon, \pi_{kt}) = 0$). The latter condition holds if a valid instrument can be found that isolates movement in $n_{kt} - n_{kt-1}$ and π_{kt} that is orthogonal to v_{kt}^ϵ . In this section, we estimate λ and ϕ using regional data on prices, wages and employment growth during the Great Recession. We argue that state-level leisure shocks were small during the Great Recession, thus allowing us to estimate λ and ϕ by OLS. Additionally, we use state-level house price variation during the early part of the Great Recession as an instrument to isolate movements in $n_{kt} - n_{kt-1}$ and π_{kt} that are orthogonal to local leisure shocks. Both procedures yield estimates of λ and ϕ that are fairly similar.

2.7.1 *Estimating Equation and Identification Assumptions*

Formally, we estimate the following specification using our state-level data:

$$\pi_{kt}^w = b_0 + b_1\pi_{kt} + b_2(n_{kt} - n_{kt-1}) + b_3\pi_{kt-1}^w + \Psi D_t + \Gamma X_k + e_{kt}$$

where $b_1 = \lambda$, $b_2 = \lambda/\phi$, $b_3 = (1 - \lambda)$, and $b_0 = \lambda(u_t^\epsilon - (1 - \rho_\epsilon)\epsilon_{t-1})$. Any aggregate leisure shocks are embedded in the constant term. The local error term includes λv_{kt}^ϵ as well as measurement error for the local economic variables. We estimate this equation pooling together all annual employment, price and wage data for years between 2007 and 2011. When estimating the above regression, we include year fixed effects (D_t). This ensures that we are only using the cross-state variation to estimate the parameters. We estimate this equation annually because we only have annual measures of wages at the state level. Our

annual nominal wage measures at the state level are the composition adjusted nominal log wages computed from the American Community Survey discussed above. π_{kt}^w , therefore, is the log-growth rate in adjusted nominal wages in the state between year t and $t - 1$. Our measure of employment growth at the state level is calculated using data from the US Bureau of Labor Statistics. The BLS reports annual employment counts and population numbers for each state in each year. We divide employment counts by population to make an annual employment rate measure for each state. $n_{kt} - n_{kt-1}$ is the log-change in the employment rate between year t and $t - 1$. π_{kt} is log-change in the average price index in each state i in year t . In our base specification, we use the retail scanner data local inflation rate scaled to account for the difference between the local non-tradable share in the retail sector and the composite consumption good. In alternative specifications, we use the raw inflation rate from the retail scanner data as our measure of local inflation. Finally, in some specifications, we include controls for the state's industry mix in 2007. This allows for the possibility that local leisure shocks, to the extent that they exist, may be correlated with the state's industry structure. Given that we have observations on 48 states for 4 years of growth rate data, our estimating equation includes 192 observations in our base specification. We also show results restricting our data to the period from 2007 to 2009, before the large changes in unemployment benefits extension starting in 2010.

Two additional comments are needed about our estimating equation. First, the theory developed above implies that $b_1 + b_3 = 1$. We impose this condition when estimating the cross-state regression. Second, we believe our local wage and price indices are measured with error. The measurement error, if classical, will attenuate our estimates of b_1 and b_3 .

Additionally, because we are regressing wage growth on lagged wage growth, any classical measurement error in wages in year t will cause a negative relationship between wage growth today and lagged wage growth. We proceed as follows to deal with this issue. Given the large sample sizes on which our wage indices (price indices) are based, we split the sample in each year and compute two measures of wage indices (price indices) for each state in each year. For example, if we have 1 million observations in the 2007 American Community Survey, we split the sample into two distinct samples with 500,000 observations each. Within each sub-sample, we compute a wage measure for each state. Both these wage measures are measured with error. Then, we use the growth rates in wages in one sub-sample as an instrument for growth rate in wages in the other sub-sample. We discuss these procedure in detail in the Online Appendix that accompanies the paper. As we show in that appendix, the procedure corrects the attenuation bias from measurement error in our estimates.³³

In order to recover unbiased estimates of λ and ϕ via OLS, we must assume $v_{kt}^\epsilon = 0$. The assumption that there are no local leisure shocks cannot generically be true. However, in the Online Appendix, we provide some evidence suggesting that this assumption may be roughly valid during the 2007-2011 period. We show that many potential leisure shocks highlighted in the literature to explain the Great Recession had large aggregate components but varied little in across US states. For example, the decline in routine jobs (Jaimovich and Siu (2014), Charles et al (2013, 2015)) was dramatic at the aggregate level during the 2007-2011 period, but occurred in all US states with roughly equal propensity. We show these results in Online Appendix Figure R7. Likewise, some have argued that the expansion of government policies

33. This procedure is similar to the split sample instrumental variable estimation in Angrist and Krueger (1995).

acted like a leisure shock that discouraged work (Mulligan (2012)). We show that many of these government policies—such as the expansion of the Supplemental Nutrition Assistance Program (SNAP) —was large at the aggregate level but had little cross state variation. In Online Appendix Figure R4, we show that SNAP benefits per recipient increased by roughly 30 percent between 2007 and 2011. Because the increase in per recipient benefit occurred at the federal level, there was statutorily no variation in per recipient benefits across US states during this time period.³⁴

One policy that has received considerable attention in its potential to act as a labor supply shock is the differential extension of the duration of unemployment benefits across states during the Great Recession. By law in 2010, weeks of unemployment benefits were tied to the state’s unemployment rate. However, as of 2010, most US states met triggers that resulted in the duration of unemployment benefits being close to the maximum of 99 weeks. These states comprised the bulk of the US population. However, some smaller states, mostly in the Plains region of the US, had smaller employment declines and, as a result, had a smaller extension of unemployment benefits.³⁵ Despite the fact that there was little population-weighted variation across states in unemployment benefit extensions during the Great Recession, we still perform two additional robustness exercises to account for the fact that the small policy differences across states that did occur may have discouraged

34. In the Online Appendix, we also show that there was no systematic variation in state labor income tax rates during the 2007-2010 period. Additionally, we show that there was little state variation in Federal programs to help underwater homeowners (like HAMP) that occurred during the 2007-2010 period. The reason is that take up rates of the program during this time period were very low (with take up rates being essentially zero prior to 2010).

35. States also had some discretion as to whether they opted into the program. This explains why some states did not have the maximum weeks of unemployment benefits even when their unemployment rate was higher. We discuss these policies and how they varied across states in detail in the Online Appendix.

labor supply. First, when using our full time period, we exclude any state that had less than 85 weeks of unemployment benefit extensions leaving us with a sample of states that had essentially no remaining variation in unemployment benefit extensions. Given that the exclude states were small in population terms, such exclusion had essentially no effect on our estimates. Additionally, we re-estimate our key parameters using only data prior to 2010. Prior to 2010, the duration of extended unemployment benefits were the same across all states. We discuss these results below.

While we defend that OLS estimation of the above equation yields unbiased estimates of λ and ϕ using cross state variation during the Great Recession, it is impossible to completely rule out that leisure shocks are causing some of the variation in state business cycles during this period. To further explore the robustness of our results, we also estimate IV specifications of the above equation. Following the work of many recent papers, including Mian and Sufi (2014), we use contemporaneous and lagged variation in local house prices as our instruments for local employment and price growth. The argument is that local house price variation during the 2007-2011 period (in our base specification) or during the 2007-2009 period (in our restricted specification) is orthogonal to movements in local leisure shocks. This seems like a plausible assumption for the 2007-2009 period as state policy changes did not occur prior to 2009. In the Online Appendix, we discuss the IV procedure in detail. We also show that contemporaneous housing price growth strongly predicts contemporaneous employment growth and lagged measures of housing growth predicts price growth. We find that IV and OLS estimates are very similar.

Before turning to the estimation, it is also worth discussing the no cross-state migration

assumption that we have imposed throughout. If individuals were more likely to migrate out of poor performing states and into better performing states, our estimated labor supply elasticities from the state regression may be larger than the aggregate labor supply elasticity. While theoretically interstate migration could be problematic for our results, empirically it is not the case. Using data from the 2010 American Community Survey, we compute migration flows to and from each state and, then, construct a net-migration rate for each state. As documented by others, we find that the net migration rate was very low during the Great Recession (Yagan 2014). This can be seen from Appendix Figure A1. Both the low level of inter-state migration and the fact that it is uncorrelated with employment growth during this period makes us confident that our estimated parameters of our local wage setting curve can be applied to the aggregate.

2.7.2 Estimates of λ and ϕ

Column 1 of Table 2.6 shows the estimates of our base OLS specification where we use all data from 2007-2011 and do not include any additional controls. Our base estimates are $b_1 = 0.69$ (standard error = 0.13) and $b_2 = 0.31$ (standard error = 0.08). As noted above, b_1 is λ and b_2 is λ/ϕ . Given our base estimates, the cross sectional variation in prices and wages implies a labor supply elasticity of 2.2. Standard macro models imply a labor supply elasticity of 2 to 4 based on time-series variation. The estimates from the cross-section of states are in-line with these macro time-series estimates. Our base estimate of $\lambda = 0.69$ suggests only a modest amount of wage stickiness. Perfectly flexible wages imply $\lambda = 1$ while perfectly sticky wages imply $\lambda = 0$. In other words, lagged wage growth predicts

current wage growth conditional on current employment and price growth, but the effect is much less than one for one (as would be implied by perfectly sticky wages). Below, we show that similar regressions run on aggregate data yield much smaller estimates of λ implying a greater amount of wage stickiness.

Columns 2 and 3 of Table 2.6 show a variety of robustness checks for our base estimates. In column 2 we include industry controls. Specifically, we include the share of workers in 2007 working in manufacturing occupations or in routine occupations. This allows us to proxy for different degrees of wage stickiness or different potential leisure shocks that are correlated with industrial mix. In column 3, we use the actual retail scanner data price index as opposed to the scaled price index representative of the composite local consumption good. Neither the inclusion of controls for local industry mix nor changing the scaling on local retail price variation affect our estimates of λ and ϕ in any meaningful way. In columns 4 and 5, we re-estimate our base specification with and without industry controls using only data from 2007-2009 prior to the changes in national policy extending unemployment benefit duration. Again, our estimates λ and ϕ remain 0.73 and 1.9, respectively.³⁶ Finally, in columns 6 and 7, we show our IV estimates for the 2007-2011 period and the 2007-2009 period where we instrument local employment growth and local price growth with contemporaneous and one lag of local house price growth. Our estimates of λ and λ/ϕ are 0.77 (standard error = 0.13) and 0.76 (standard error = 0.17) implying an estimated Frisch elasticity of 1.0.

Regardless of our specification we estimate labor supply elasticities of between roughly 1.0 and 2.0. More importantly, all of our estimates imply a fair degree of wage flexibility

36. Additionally, we estimated our base specification excluding CA, NV, AZ, and FL. In both cases, our estimates were nearly identical to our base specification in column 1 of Table 3.

with our estimates of λ ranging from about 0.7 to 0.8. These results are consistent with the patterns shown in Figure 2.3 where local wages co-moved strongly with local employment during the Great Recession. The estimation that wages are fairly flexible is a key insight that is important for our main results in the next section and has broader implications for the literature. In the context of the methodology we presented in the previous section as well as our monetary union model, it is hard to get aggregate "demand" shocks to be the primary shock driving economic conditions during the Great Recession if wages are fairly flexible. In other words, if wages were sticky enough in the aggregate to have "demand" shocks be the primary driver of aggregate employment decline during the recent recession, we would not have observed wages moving as much as they did in the cross-section of states during the same time period.

To show the stark difference between local and aggregate estimates of wage stickiness, we use aggregate data on prices, nominal wages, and employment between 1976 and 2012. This is the same data that we will use in our SVAR estimation in the next section. Given the short time-series sample, power is an issue. However, across all specifications we explored, estimates of λ using aggregate data ranged from about 0.4 to 0.6. These estimates are below the estimates of 0.7 to 0.8 using local variation. If aggregate leisure shocks occur along with shocks that shift labor demand, wages will appear sticky in the aggregate time-series. The assumption of no aggregate leisure shocks is a common one when estimating wage stickiness using aggregate data (see, for example, Christiano et al. (2015b)).

2.8 The US Great Recession: From Regional to Aggregate

The cross-regional facts presented above represent a puzzle. Aggregate nominal wages did not fall much (relative to trend) during the Great Recession. However, local nominal wages and employment were significantly, negatively correlated. Why did aggregate wages respond so little during the Great Recession while there was a strong relationship across states?

One potential explanation is that a series of shocks made aggregate employment fall. Some of these shocks put downward pressure on wages while others put upward pressure, thus making wages seem unresponsive. However, if the shocks putting upward pressure on wages were purely aggregate, they would be differenced out when considering variation across states—thus resulting in the observed negative correlation between employment and wages across states. Our methodology allows us to quantify the relative magnitudes of these shocks and to assess their contributions to the behavior of prices, wages and employment.

2.8.1 Findings in the Aggregate

We follow the procedure described in Section 2.6. We first estimate the VAR with two lags in aggregate employment growth, price growth and nominal wage growth via OLS equation by equation using annual data from 1976 to 2012. We obtain sample estimators of the covariance matrix $\hat{V} = \frac{UU'}{\text{Years} - \#\text{Variables} \# \text{Lags}}$ from reduced form errors U .

We construct aggregate variables that are comparable to our regional measures. Given that our cross-sectional equations are estimated using annual data, we analogously define our aggregate data at annual frequencies. We use data from the CPI-U to create our measure of

aggregate prices. Specifically, we take log-change in the CPI's between the second quarter of year t and $t - 1$ for our measure of p_t . For n , we use BLS data on the aggregate employment to population rate of all males 25-54. We choose this age range so as to abstract from the downward trend in employment rates due to the aging of the population over the last 30 years.³⁷ Finally, we use data from the Current Population Survey (CPS), discussed above, to construct our aggregate composition adjusted wage measure. As with the CPI, we take the log-change in this wage measure between t and $t - 1$ for our measure of w_t . For all data, we use years between 1976 and 2012.

Figures 2.4, 2.5, and 2.6 report the impulse response of aggregate employment, nominal wages and price growth to each of the shocks using our benchmark estimates for λ and ϕ reported in column 1 of Table 2.6 ($\lambda = 0.69$ and $\phi = 2.2$). Figure 2.4 shows their behavior after a one-standard deviation discount rate (γ). Qualitatively, after a discount rate shock both prices and employment increase sharply relative to trend while real wages decline slightly relative to trend. These results are identical to the theoretical predictions shown in Table 2.4. Figure 2.5 shows the impulse responses to a one-standard deviation productivity/markup (z) shock. Prices decrease on impact while employment increases sharply. Nominal wages, however, only decline slightly. Again, these predictions match the predictions of our simple theoretical model shown in Table 2.4. While both γ and z shocks increase employment, the γ shock puts upward pressure on prices while the z shock puts downward pressure on prices.

37. We detrend all data when estimating the VAR. Specifically, we allow for a linear trend in the employment to population ratio between 1978 and 2007. For the price inflation rate and the nominal wage inflation rate, we use an HP filter (with a smoothing parameter of 100). Given that we detrend the data, our results are essentially unchanged when we use the employment to population ratio for all individuals as opposed to using it just for prime age males.

Figure 2.6 shows the impulse response of employment, prices, and nominal wages to the leisure shock. Upon impact, the leisure shock reduces employment and prices while it increases wages. Again, these predictions match the predictions from the benchmark monetary union model.

We now turn to quantifying the contribution of each shock to explaining the behavior of the aggregate US economy during the Great Recession. To do so, we present the counterfactual cumulative response of each individual variable when we feed the VAR with the sequence of shock realizations between 2008 and 2012, one at a time.³⁸ Our analysis suggests that "demand" shocks cannot be solely responsible for the employment decline during the Great Recession. If "demand" shocks (i.e., discount rate shocks) were solely responsible, price and wage inflation would have been lower. Instead a combination of "supply" shocks (i.e., productivity/markup and leisure shocks) explain the "missing price and wage deflation".³⁹

Figure 2.8 presents the counterfactual employment response. Employment fell by more than 4 percent between 2008-2009 (relative to trend) and remained at this low level thereafter. The counterfactual exercise shows that the productivity/markup, discount rate, and leisure shocks contributed about the same amount to the initial decline during the 2008-2009 period (each explaining roughly one-third of the aggregate employment decline). However, the discount rate and leisure shocks do not explain any of the persistence in the employment decline post 2009. Instead, it is the productivity/markup shock that explains most of the

38. For the interested reader, the actual realizations of the shocks we estimate can be seen in Figure 2.7

39. The robust growth in consumer prices during the recession is also viewed as a puzzle for those that believe that the lack of aggregate demand was the primary cause of the Great Recession. For discussions of the "missing deflation", see Hall (2011), Ball and Mazumder (2011), Stock and Watson (2012), and Del Negro et al. (2015).

sluggish response of employment post 2009.

Figures 2.9 and 2.10 help understand the "missing price and wage deflation puzzle". Figure 2.9 shows the counterfactual price response to each of the shocks. Aggregate prices fell relative to trend between 2008 and 2009 and quickly stabilized thereafter despite the weak employment situation post-2009. This is the sense in which there was "missing price deflation". Both the discount rate and the leisure shock put downward pressure on aggregate prices. However, the productivity/markup shock put upward pressure on aggregate prices. The counterfactual analysis shows that if the economy had only been hit by the productivity/markup shock, prices would have risen (by upwards of 1 percent) relative to trend during the Great Recession. Instead, we find that it is this countervailing productivity/markup shock that arises as the explanation for the missing deflation puzzle—particularly post 2009. This finding is consistent with the results of Christiano et al. (2015a).

Figure 2.10 shows the cumulative nominal wage response to each of the shocks. Again, the figure shows the "missing wage deflation puzzle" during the Great Recession. Throughout the recession, nominal wage growth was close to zero (relative to trend). However, if the economy had only experienced the discount rate shock, nominal wages would have fallen by roughly 1.5 percent relative to trend by 2009 and would have remained below trend in 2011. It is the leisure and productivity/markup shocks that explain why nominal wages did not fall during the Great Recession.

2.8.2 Sensitivity to Alternative Parameters λ and ϕ

How do our estimated parameters affect our employment, price and wage counterfactuals? In Table 2.7, we report the contribution of each shock to aggregate employment declines implied by different combinations of $\{\phi, \lambda\}$. We do this for both the initial years of the recession (2008 to 2009) and over the longer period encompassing the recovery (2008 to 2012). Each cell in Table 2.7 shows how much of the employment change during the time period can be attributed to the discount rate shock (γ) and how much can be explained by the productivity/markup shock (z). The sum of all three shocks sums to 100 percent. So, the difference between the sum of the γ and z contributions and 100 percent is attributed to the leisure shock (ϵ). The qualitative conclusions of the previous section still hold for the range of $\{\phi, \lambda\}$ estimates in Table 2.6. These go from roughly 0.7 to 0.8 for λ and from roughly 1.0 to 2.5 for ϕ .

Table 2.7 offers several further results worth discussing. First, we observe that the relative importance of the leisure shock vis-a-vis the discount rate and productivity/markup shocks combined is governed by the Frisch labor supply elasticity (ϕ). We estimate a relatively large elasticity, in the range of that used to calibrate standard macro models.⁴⁰ However, suppose we used a much lower elasticity instead, e.g., $\phi = 0.5$, which is in line with some microeconomic estimates in the literature. In this case, the leisure shock would account for

40. This result may be of independent interest to the reader familiar with the macro v. micro labor supply elasticities (see Chetty, Guren, Manoli, and Weber (2011)). Using cross-sectional data (same as in most of the micro labor-supply elasticity literature) we arrive at an estimate similar to the macro elasticity (estimated from aggregate time-series data). We believe this is because the regional variation in employment rates that we use to estimate this elasticity only incorporates the extensive margin adjustment in the labor supply, which is the same margin that is most important in accounting for aggregate fluctuations in total hours over the business cycle.

a much larger fraction of the employment decline in the Great Recession. In other words, if labor supply is fairly elastic, large movements in employment are consistent with relatively small movements in real wages (as we observe in US data), without the need of large leisure shocks. While this sensitivity analysis results from re-estimating the shocks under different parameterizations for ϕ using our procedure, the intuition is in line with the benchmark monetary union model from Section 4.

The intuition for the decomposition between discount rate and productivity/markup shocks is more subtle but also in line with the our benchmark monetary union model. We find that the degree of wage flexibility (λ) affects their relative contribution to the remaining, unexplained part by the leisure shock alone. For example, if we increased the degree of wage flexibility,⁴¹ the productivity/markup shock would account for a much larger fraction of the employment decline in the Great Recession. Theoretically, it is clear that when λ is large, the discount rate shock should not matter much for the determination of employment. To see this, consider the extreme case where wages are perfectly flexible and the discount rate shock is only composed of the monetary shock. Then the equilibrium in the simple theoretical model satisfies monetary neutrality. We formalized this point in Section 2.5.6 when we derived the model's implied elasticity of aggregate employment to a discount rate shock. Conversely, when wages are very rigid ($\lambda = 0.1$), our procedure suggest that discount rate shocks explain essentially all of the decline in the early part of the recession and much of the persistence in employment decline during the 2008-2012 period.

41. It is worth mentioning that for large values of λ and small values of ϕ , the results in Table 2.7 become rather sensitive to small variations in parameters. This is because our shock identification procedure needs a certain degree of wage stickiness, as explained in the Section 2.6. For example, for values of λ around 0.95 it is not possible to identify the productivity/markup and discount rate shocks.

2.8.3 Sensitivity to Alternative Identifying Assumptions

The above results are based on the particular functional form of our wage setting equation:

$$W_{kt} = (P_{kt}e^{\epsilon_{kt}}(N_{kt})^{\frac{1}{\phi}})^{\lambda} (W_{kt-1})^{1-\lambda}$$

This wage setting equation reflects our assumption of GHH preferences as well as no forward looking behavior when wages are reset. Both of these assumptions were made for tractability. In this sub-section, we explore the sensitivity of our results to relaxing both of these assumptions.

In Appendix B1, we derive the aggregate and local wage setting equations under a broad set of utility functions where consumption and leisure are non-separable. This class of utility functions allows for arbitrarily large income and substitution effects. As we show in the appendix, the use of local consumption data allows us to estimate the extent of wage stickiness as well as to estimate the parameters that encompass both the income and substitution effects on labor supply. In particular, we can estimate the following equation using local data:

$$\pi_{kt}^w = \tilde{b}_t + \tilde{b}_1\pi_{kt} + \tilde{b}_2(n_{kt} - n_{kt-1} + \tilde{b}_3\pi_{kt-1}^w + \tilde{b}_4(c_{kt} - c_{kt-1})) + \Psi D_t + \Gamma X_k + e_{kt}.$$

This equation is identical to our estimating equation above aside from the addition of local consumption growth (i.e., $c_{kt} - c_{kt-1}$). As outlined in Appendix B1, the coefficients \tilde{b}_1

and \tilde{b}_3 sum to 1 even under the broader preference specification. We impose this restriction when estimating the modified equation. We measure local real consumption growth using the log-change in real retail expenditures at the state level computed within the Nielsen sample. We obtain real expenditures by deflating nominal expenditures with our local price indices.⁴² For our base specification, our estimates of \tilde{b}_1 , \tilde{b}_2 , and \tilde{b}_4 are 0.72 (standard error = 0.12), 0.25 (standard error = 0.08), and 0.16 (standard error = 0.06), respectively. A positive and significant coefficient on real consumption growth (\tilde{b}_4) reflects the presence of income effects on labor supply. Controlling for this income effect, our estimate of wage flexibility (\tilde{b}_1) is slightly higher than our base specification where income effects are not allowed.

In the aggregate wage setting equation, we can substitute out consumption growth using the model definition ($c_t = w_t + n_t - p_t$). Appendix B1 shows that the aggregate wage setting equation still takes the following form:

$$\pi_t^w = \lambda\pi_t + \frac{\lambda}{\phi}(n_t - n_{t-1}) + (1 - \lambda)\pi_{t-1}^w + \frac{\lambda}{1 - \omega}\epsilon_t$$

where ω is a parameter that represents the strength of the income effect on labor supply (and maps directly to \tilde{b}_4 from the above local labor supply regression, see equation (5) in Appendix B1). Aside from the coefficient scaling the aggregate leisure shock, this equation is identical to the identification restriction we imposed when estimating the aggregate SVAR. The only difference is that there is no longer a direct mapping between λ and λ/ϕ —in

42. This measure of real expenditures is (1) highly correlated with measures of local employment and (2) highly correlated with the BEA's recent state level personal expenditures measure. Our results are similar if we use the BEA's local consumption measure. However, we prefer our measure given that much of the BEA's local consumption measure is imputed (where the imputation uses local employment measures).

the above aggregate wage setting equation that we impose to identify the SVAR—and the reduced-form parameters \tilde{b}_1 and \tilde{b}_2 —in the local wage setting equation. However, as shown in Appendix B1, there is still a one-to-one mapping between the parameters we estimate from the local regression (\tilde{b}_1 , \tilde{b}_2 , and \tilde{b}_4) and the aggregate parameters we need to identify the SVAR (λ and ϕ). With the correctly specified λ and ϕ , we can just use the matrix in Table 2.7 to read off the decomposition of shocks during the Great Recession. While λ and ϕ are no longer structural parameters (instead being combinations of structural parameters), knowing them still helps identifying the aggregate SVAR. Using our estimates of \tilde{b}_1 , \tilde{b}_2 , and \tilde{b}_4 and the procedure developed in Appendix B1, we estimate λ and ϕ (allowing for income effects on labor supply) to be 0.68 and 2.0, respectively. These parameters are nearly identical to our base specification without income effects. The take-away from this sensitivity exercise is that abstracting from income effects on labor supply is not biasing our decomposition of the shocks driving aggregate employment declines during the Great Recession in any meaningful way.

In Appendix B2, we specify an alternative wage setting equation allowing for forward looking behavior when wages are reset. We show that ignoring forward looking wage setting behavior biases up our estimates of wage flexibility. That is, the amount of wage flexibility that we estimate using cross-state variation is too large relative to the true amount of wage flexibility in the aggregate. We show that the bias depends on two parameters: (1) the extent to which firms put weight on forward looking behavior when setting wages (which we call κ) and (2) the underlying persistence process of local wages (which we call $\bar{\rho}_w$). Assuming that at the aggregate level the monetary authority wants to stabilize expected nominal wage

growth, we quantify the extent to which our estimates of λ from the local regressions are biased upwards as a function of these two parameters. Under a range of plausible parameter estimates for κ and $\bar{\rho}_w$, we show that the bias is quite small. For our baseline parameter estimates, we show that $\kappa(1 - \bar{\rho}_w)$ must exceed 2.5 for the true λ to be lower than 0.4. This is an order of magnitude larger than any plausible parametrization for either κ or $\bar{\rho}_w$. Moreover, as seen in Table 2.7, our estimates of the role of demand shocks in explaining employment decline during the Great Recession are broadly similar for values of λ between 0.4 and 0.69. These results suggest that our abstraction from including expectations in our wage setting equation is not quantitatively altering the paper's conclusions.

2.8.4 Discussion: Aggregate v. Regional Shock Decomposition

The above results suggest that aggregate wages appeared sticky during the Great Recession because a combination of aggregate shocks resulted in relatively offsetting effects on aggregate wages. By extending our procedure to allow for the estimation of regional shocks in a regional SVAR, we find that the discount rate shock explains roughly 80 percent of the change in non-tradable employment between 2007 and 2010 across states. This is consistent with results found in Mian and Sufi (2014) suggesting that housing price declines explained much of the cross-state variation in non-tradable employment. We conclude that even though discount rate shocks only explained a portion of aggregate employment decline early in the Great Recession and very little of its persistence, the discount rate shock was primarily responsible for much of the cross-state variation during this time period.

We relegate the discussion of the estimation procedure and identification assumptions

to the Online Appendix—given that the estimation of regional shocks is not central to the paper. However, this exercise allow us to highlight the difficulty in extrapolating findings from cross-region variation to interpret aggregate time-series patterns. The fact that local discount rate shocks explain much of the cross-state variation in employment during the Great Recession does not imply that an aggregate discount rate shock explains much of the aggregate time-series variation in employment during the Great Recession. If the discount rate shock would have been the main driver of aggregate employment decline during the Great Recession then aggregate wages would have behaved similarly to state-level wages.

2.9 Conclusion

Regional business cycles during the Great Recession in the US were strikingly different than their aggregate counterpart. This is the cornerstone observation on which we built this paper. Yet, the aggregate US economy is just a collection of these regions connected by trade of goods and assets. We argued that their aggregation cannot be arbitrary and that regional business cycle patterns have interesting implications for aggregate business cycles.

Our paper offers four takeaways. The first is that the relationship between wages and employment in the aggregate time-series during the 2006-2011 period is very different than the cross-state relationship between these variables during the same time period. For example, while aggregate wages appeared to be sticky despite aggregate employment falling sharply, both local nominal and real wages co-varied strongly with local employment growth in the cross-section of US states. Both documenting the regional facts and the creation of the underlying local price and wage indices are the first innovations of the paper.

The second take-away is that wages seem to be modestly sticky when using cross-state variation to estimate our wage setting equation. The amount of wage stickiness is often a key parameter in many macro models. Despite its importance, there are not many estimates of the frequency with which wages adjust (particularly relative to estimates of price adjustments). We develop a procedure to estimate the amount of wage stickiness using cross-region variation. The wage stickiness parameter is key to our empirical methodology to estimate the underlying shocks and elasticities. The fact that we estimate that wages are only modestly sticky limits the importance of "demand" shocks at the aggregate level in explaining the Great Recession. If wages are only modestly sticky, aggregate demand shocks should have resulted in falling aggregate wages—which was not observed in the aggregate time-series. Regardless of the use of this parameter in our empirical work, our estimate of wage stickiness could be of independent interest to researchers.

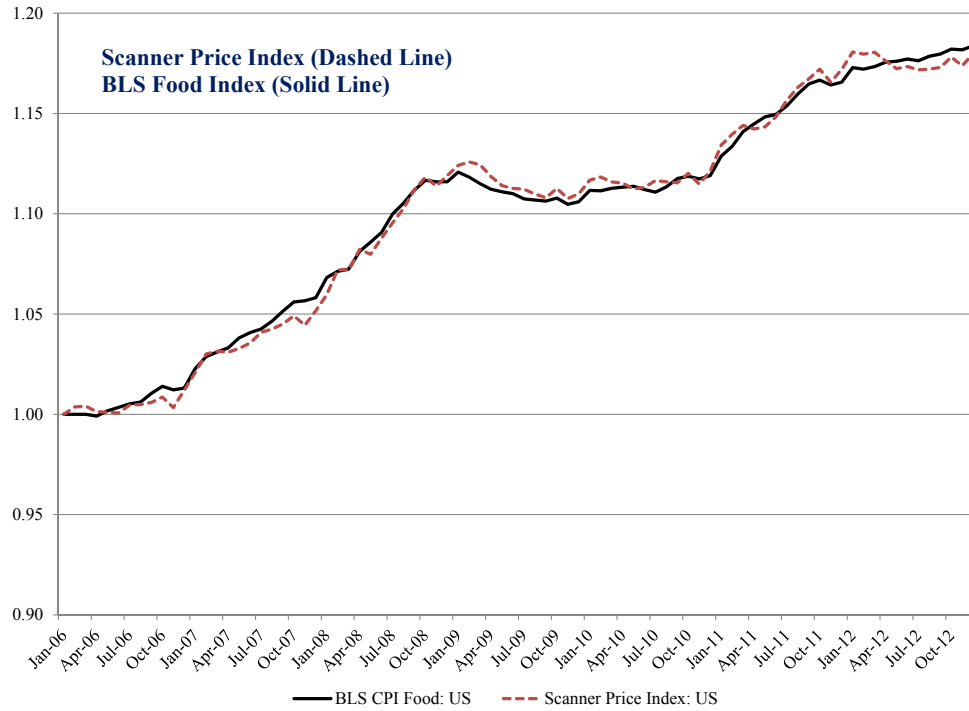
The third take-away from this paper is developing a methodology that allows us to estimate aggregate shocks by combining aggregate and regional data. This methodology is a hybrid method that merges restrictions imposed by a theoretical model with aggregate and cross-sectional data when estimating a SVAR and identifying the corresponding shocks. We view this as a contribution to the growing literature that uses model-based structure to estimate SVARs.

Finally, the fourth take-away is perhaps the most important for the goals of the paper. We show that a combination of both "demand" and "supply" shocks are necessary to account for the joint dynamics of aggregate prices, wages and employment during the 2007-2012 period in the US. In contrast with the aggregate results, we find that discount rate shocks explain most

of the observed employment, price and wage dynamics across states. These results suggest that solely using cross-region variation to explain aggregate fluctuations is insufficient when some shocks do not have a substantive regional component. The fact that aggregate wages did not fall cannot be explained by large degrees of wages stickiness. The reason aggregate wages did not fall is that the series of shocks experienced by the aggregate economy were such that some shocks put downward pressure on prices and wages (discount rate shocks) while other shocks put upward pressure on prices and wages (productivity/markup and leisure shocks). In the cross-section, however, the discount rate shocks caused prices, wages and employment to move in the same direction. Lastly, in our calibrated monetary union model, we show that the local employment elasticity to a local discount rate shock is substantially larger than the aggregate employment elasticity to an identically-sized aggregate discount rate shock. These results suggest that even when the aggregate and regional shocks are the same, it is hard to draw inferences about the aggregate economy using regional variation. Collectively, our results suggest that researchers should be cautious when extrapolating cross-sectional variation to make statements about aggregate business cycles.


2.10 Figures and Tables

Figure 2.1: Nielsen Retail Price Index vs. CPI Food Price Index




Note: In this figure, we compare our monthly retail scanner price index for the U.S. as a whole (dashed line) to the CPI's aggregate monthly food price index (solid line). We normalize both indices to 1 in January of 2006.

Figure 2.2: The Evolution of Aggregate Real and Nominal Composition Adjusted Wages




Note: Figure shows the evolution of aggregate real and nominal log wages within the U.S. between 2000 and 2012 using data from the Current Population Survey. The sample is restricted to only males between the ages of 21 and 55, who are currently employed, who report usually working 30 hours per week, and who worked at least 48 weeks during the prior 12 months. As discussed in the text, we adjust wages for the changing labor market condition over time by controlling for age, race, education, and usual hours worked. We compute real wages by deflating our nominal wage index by the CPI-U of the corresponding year.

Figure 2.3: State Employment Growth vs. State Nominal and Real Wage Growth, 2007-2010




Note: Figure shows a simple scatter plot of the percent growth in the state employment rate between 2007 and 2010 against nominal wage growth (left panel) and real wage growth (right panel) during the same period. The state employment rate comes dividing state employment from the BLS by total state population from the BLS. Nominal wages are computed from the ACS and are adjusted for the changing labor market composition of workers within each state over time. We restrict wage measures to a sample of men between the ages of 21 and 55 with a strong attachment to the labor market. Our composition adjustment controls for age, education, race, nativity and usual hours worked. See text for details. To compute real wages, we adjust our nominal wage measures by our local price indices created using the retail scanner data. The size of the underlying state is represented by the size of the circle in the figure. The line represents a weighted regression line from the bi-variate regression.

Figure 2.4: Impulse Response to a Discount rate Shock




Note: Figure shows the impulse response to a one standard deviation discount rate shock. The horizontal axis are years after the shock.

Figure 2.5: Impulse Response to a Productivity / Markup Shock



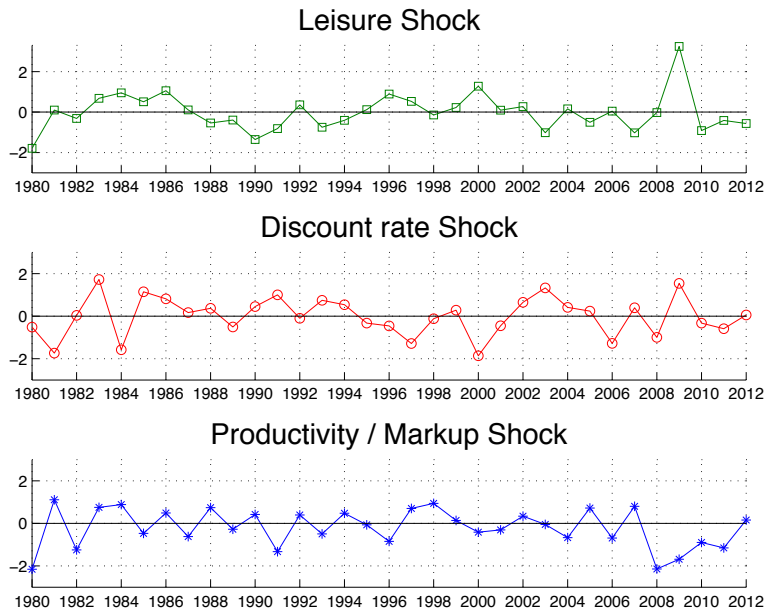
Note: Figure shows the impulse response to a one standard deviation productivity/markup shock. The horizontal axis are years after the shock.

Figure 2.6: Impulse Response: Leisure Shock




Note: Figure shows the impulse response to a one standard deviation leisure shock . The horizontal axis are years after the shock.

Figure 2.7: Shock time-series




Note: Figure shows the estimated aggregate shock realizations from 1980 to 2012.

Figure 2.8: Counterfactual Employment Response




Note: Figure shows the cumulative response of employment when we feed the VAR with the sequence of shocks between 2008 and 2012; one at a time.

Figure 2.9: Counterfactual Price Response



Note: Figure shows the cumulative response of Prices when we feed the VAR with the sequence of shocks between 2008 and 2012; one at a time.

Figure 2.10: Counterfactual Wage Response



Note: Figure shows the cumulative response of Wages when we feed the VAR with the sequence of shocks between 2008 and 2012; one at a time.

Table 2.1: Comparison of Cross-State and Time-Series Estimates of Wage Elasticities During the Great Recession

	Employment Rate	Nominal Wage	Real Wage
Cross-State			
Cross-State Wage Elasticity With Respect to Employment, 2007-2010		0.62 (0.10)	0.52 (0.15)
Aggregate			
Actual Aggregate Growth, 2007-2010	-7.7 percent	3.8 percent	-0.9 percent
Expected Aggregate Growth, 2007-2010 (Based on 2000-2007 Trend)	-0.9 percent	5.5 percent	-2.1 percent
Aggregate Deviation from Expected Growth, 2007-2010	-6.8 percent	-1.7 percent	1.2 percent
Aggregate Wage Elasticity With Respect to Employment, 2007-2010		0.25	-0.17

Note: Table compares the wage elasticity to a one percent change in the employment rate estimated off of cross-state data (top panel) to a similarly defined wage elasticity estimated off of aggregate time-series data during the 2007 to 2010 period (bottom panel). The cross-state elasticities come from the simple scatter plots shown in Figure 3. Standard errors from the regression line in the scatter plots are shown in parentheses. The aggregate time-series elasticity is computed using aggregate data. For the aggregate nominal wage data, we use the adjusted wage series we created using data from the CPS. See text for details. For aggregate real wages, we adjust the nominal wage data by the June CPI-U. To get predicted nominal and real wage growth between 2007 and 2010, we take a simple linear prediction of the corresponding nominal and real growth between the 2000 and 2007 period. Once we get the deviation between actual wage growth and predicted wage growth between 2007 and 2010, we divide that difference by -6.8 percent. -6.8 percent is the decline in the aggregate employment rate between 2007 and 2010 above and beyond what would have been predicted from changes in the employment rate between 2000 and 2007. We use aggregate data from the BLS to compute the employment rate in 2000, 2007 and 2010.

Table 2.2: **Calibration**

	Value	Target
α	0.62	Aggregate labor share in the non-tradable sector in 2006
θ	0.55	Aggregate labor share in the tradable sector in 2006
β	0.16	Aggregate labor share in 2006 (0.61). See on-line appendix
σ	2	Intertemporal elasticity of substitution
ϕ	2	Frisch elasticity. See Section 2.7.2 for estimation
λ	0.7	Wage stickiness. See Section 2.7.2 for estimation
φ_p	1.5	Taylor Rule from Galí (2011)
φ_y	0.5	Taylor Rule from Galí (2011)
Φ_0	0.01	Trade balance-output ratio volatility = 2 from Mendoza (1991).
R	0.03	Real interest rate
X	0.32	Intermediate inputs over output ratio US 2006
B	2.1	Median net worth to output ratio in the US 2006
ρ_γ	0.9	Persistence of discount rate shock
ρ_z	0.76	Persistence of productivity shock from the model. See on-line appendix
ρ_ϵ	0.66	Persistence of leisure shock from the model. See on-line appendix

Table 2.3: **Aggregate v. Regional Employment Impact Elasticities**

	ρ_γ	0.9	0.9	0.9	0.6	0.1			
	β	0.16	0	0.16	0.16	0.16			
	φ_p	1.5	1.5	1.5	1.31	1.5			
	φ_y	0.5	0.5	0	0.5	0.5			
	(λ, ϕ)	(0.7, 2)	(0.7, 1)	(0.5, 2)	(0.7, 2)	(0.7, 2)			
Aggregate	γ	0.77	0.53	1.29	0.74	1.11	1.76	0.71	0.31
	z	0.29	0.20	-0.35	0.28	-0.24	-1.07	0.29	0.29
	ϵ	-0.65	-0.51	-0.28	-0.69	-0.26	0.41	-0.65	-0.65
Regional	γ	1.76	1.26	2.95	2.85	1.76	1.76	0.36	0.03
	z^y	0.44	0.38	-0.18	0.10	0.44	0.44	0.44	0.44
	z^x	0.08	-0.06	0.23	0.05	0.08	0.08	0.08	0.08
	ϵ	-0.64	-0.52	-0.16	-0.43	-0.64	-0.64	-0.64	-0.64

Note: The table summarizes the response of employment on impact to each of the model shocks under our base calibration (column 1) and alternate calibrations (columns 2-8). The rows represent the employment response to different aggregate shocks (top three rows) and different local shocks (bottom three rows). Columns 2 and 3, explore the elasticities under different calibrations of wage stickiness and the Frisch elasticity of labor supply. Column 4 examines the robustness to changes in the tradable share of the intermediate good. Columns 5 and 6 examine the results under alternate Taylor Rule parameters. The final two columns change the persistence of the demand shock. The units are percentage deviations from the steady state, in the case of aggregate employment, and percentage deviations from the aggregate in the case of regional employment.

Table 2.4: **Aggregate Elasticities** ($\lambda=0.7, \phi=2$)

		n	π	π^w	w^r
Short run	γ	0.77	1.64	1.42	-0.22
	z	0.29	-2.73	-1.81	0.92
	ϵ	-0.65	0.94	1.13	0.19
Long run	γ	0.56	1.01	1.03	-0.16
	z	0.12	-0.66	-0.79	0.40
	ϵ	-0.18	0.35	0.32	0.05

Note: The table summarizes the response of each aggregate variable to the shocks in percentage deviations from the steady state. The "short" elasticity is the response at date $t = 0$. The "long" elasticity is the response after 5 years.

Table 2.5: **Regional Elasticities** ($\lambda=0.7, \phi=2$)

		n	π	π^w	w^r
Short run	γ	1.76	2.54	2.40	-0.14
	z^y	0.44	-2.00	-1.25	0.75
	z^x	0.08	0.29	0.31	0.02
	ϵ	-0.64	0.71	0.97	0.26
Long run	γ	0.12	-0.28	-0.30	0.19
	z^y	0.63	0.31	0.23	0.21
	z^x	-0.06	-0.03	-0.04	0.04
	ϵ	-0.44	-0.12	-0.20	0.14

Note: The table summarizes the response of each island level variable to the shocks in percentage deviations from the steady state. The "short" elasticity is the response at date $t = 0$. The "long" elasticity is the response after 5 years.

Table 2.6: Estimates of λ and $\frac{\lambda}{\phi}$ using Cross- Region Data

	Specification						
	OLS			IV			
	2007-2011			2007-2009		2007-2011	2007-2009
(1)	(2)	(3)	(5)	(6)	(7)	(8)	
λ	0.69 (0.13)	0.69 (0.13)	0.75 (0.13)	0.73 (0.17)	0.73 (0.17)	0.77 (0.13)	0.79 (0.18)
λ/ϕ	0.31 (0.08)	0.32 (0.08)	0.31 (0.07)	0.39 (0.09)	0.39 (0.10)	0.76 (0.17)	0.99 (0.25)
Implied ϕ	2.2	2.2	2.4	1.9	1.9	1.0	0.8
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry Controls	No	Yes	Yes	No	Yes	Yes	Yes
Scaling Factor of Prices	1.4	1.4	1.0	1.4	1.4	1.4	1.4

Note: Table shows the estimates of λ and $\frac{\lambda}{\phi}$ from our base wage setting specification using the regional data. Each observation in the regression is state-year pair. Each column shows the results from different regressions. The regressions differ in the years covered and additional control variables added. The first three columns show the OLS results using all local data between 2007 and 2011. Columns 4 and 5 show OLS results using only data from 2007 through 2009. The final two columns show IV results for the different time periods. In the IV specifications, we instrument contemporaneous employment and price growth with contemporaneous and lagged house price growth. We adjust for measurement error in wage growth, lagged wage growth, and price growth using the split sample methodology discussed in the Online Data Appendix. All regressions included year fixed effects. All standard errors are clustered at the state level.

Table 2.7: Discount/Interest rate (γ) and Productivity/Markup (z) shocks' contribution to aggregate employment change

λ	2008 to 2009						2008 to 2012				
	ϕ						ϕ				
		0.5	1	2	3	4	0.5	1	2	3	4
0.1	γ	103	83	108	107	108	40	-34	46	57	52
	z	-3	22	-3	-1	-2	47	126	48	38	43
0.3	γ	47	66	51	29	101	-13	-25	-33	-20	9
	z	2	16	45	71	1	98	123	134	121	92
0.5	γ	8	36	31	13	92	-1	-13	-19	2	47
	z	6	16	48	74	0	94	123	133	111	65
0.7	γ	0	0	30	21	11	47	48	-12	-12	0
	z	3	33	33	53	69	53	72	136	135	122
0.9	γ	0	0	0	2	3	45	47	41	15	9
	z	-1	24	54	64	69	58	79	91	114	118

Note: Table shows the percent contribution of the demand and supply shocks to the aggregate employment change implied by our procedure for different combinations of the parameters. For a given pair $\{\phi, \lambda\}$, the ' γ ' entry corresponds to the demand shock. The ' z ' entry to the supply shock. The percent contribution of the leisure shock can be calculated by subtracting the sum of both entries from 100. Entries with * are such that no decomposition of the shocks satisfy the identification restrictions for those parameter values.

2.11 Appendix

2.11.1 Proof of Lemma 1 and 2

The following equations characterize the log-linearized equilibrium

$$w_{kt}^r = \lambda(\epsilon_{kt} + \frac{1}{\phi}n_{kt}) + (1 - \lambda)(w_{kt-1}^r - \pi_{kt})$$

$$w_{kt}^r = -(1 - (\alpha + \theta\beta))n_{kt}^y - \beta(1 - \theta)(n_{kt}^x - n_{kt}^y) + z_{kt}^y + \beta z_{kt}^x$$

$$0 = \mathbb{E}_t (mu_{kt+1} - mu_{kt} - \pi_{kt+1} - \gamma_{kt+1} - \Phi_0(c_{kt} - c_t) + \varphi_p\pi_t + \varphi_y(c_t - c_t^*))$$

$$mu_{kt+1} = -\frac{\sigma}{C - \frac{\phi}{1+\phi}N^{\frac{1+\phi}{\phi}}} \left(Cc_{kt+1} - N^{\frac{1+\phi}{\phi}} \left(\frac{1+\phi}{\phi}\epsilon_{kt+1} + n_{kt+1} \right) \right)$$

$$Nn_{kt} = N^x n_{kt}^x + N^y n_{kt}^y$$

$$c_{kt} = w_{kt}^r + n_{kt}^y$$

$$b_{kt} = (1 + r)(b_{kt-1} + i_t) + \frac{X}{B}(z_{kt}^x + \theta n_{kt}^x - x_{kt}) - r\tau_t$$

$$0 = z_{kt}^x - (w_{kt} - q_t) - (1 - \theta)n_{kt}^x$$

$$x_{kt} = n_{kt}^y + (w_{kt} - q_t)$$

$$\sum_k x_{kt} = \sum_k (z_{kt}^x + \theta n_{kt}^x)$$

From the last 3 equations, after adding up, it holds that $n_t^x = n_t^y$. Then the aggregate log-linearized equilibrium evolution of $\{\pi_t^w, w_t^r, n_t\}$ is characterized by

$$\begin{aligned}
0 &= \mathbb{E}_t(mu_{t+1} - mu_t - \pi_{t+1} - \gamma_{t+1}) + \varphi_p \mathbb{E}_t[\pi_{t+1}] + \varphi_y(w_t^r + n_t) \\
\pi_t^w &= \frac{\lambda}{1-\lambda}(\epsilon_t + \frac{1}{\phi}n_t - w_t^r) \\
w_t^r &= -(1 - (\alpha + \theta\beta))n_t + z_t \\
mu_{t+1} &\equiv -\frac{\sigma}{C - \frac{\phi}{1+\phi}N^{\frac{1+\phi}{\phi}}}\left(C(w_{t+1}^r + n_{t+1}) - N^{\frac{1+\phi}{\phi}}\left(\frac{1+\phi}{\phi}\epsilon_{t+1} + n_{t+1}\right)\right) \\
\pi_{t+1} &\equiv \pi_{t+1}^w - (w_{t+1}^r - w_t^r)
\end{aligned}$$

which is equivalent to the system of equations characterizing the log-linearized equilibrium in a representative agent economy with a production technology that utilizes labor alone with an elasticity of $\alpha + \theta\beta$, no endogenous discounting and only 3 exogenous processes $\{z_t, \epsilon_t, \gamma_t\}$. The top equation is the aggregate Euler equation. The second equation is the aggregate wage setting equation. The third equation is effectively the aggregate labor demand curve.

To prove Lemma 2, just take log-deviations from the aggregate in the original system. This results in the system characterizing the evolution of $\{\tilde{p}_t, \tilde{w}_t, \tilde{n}_t^y, \tilde{n}_t^x\}$ for given

$$\{\tilde{z}_t^y, \tilde{z}_t^x, \tilde{\gamma}_t, \tilde{\epsilon}_t\},$$

$$\tilde{w}_t = \lambda \left(\tilde{p}_t + \tilde{\epsilon}_t + \frac{1}{\phi} \left(\frac{N^x}{N} \tilde{n}_t^x + \frac{N^y}{N} \tilde{n}_t^y \right) \right) + (1 - \lambda) \tilde{w}_{t-1}$$

$$\tilde{w}_t = \tilde{p}_t - (1 - (\alpha + \theta\beta)) \tilde{n}_t^y - \beta(1 - \theta)(\tilde{n}_t^x - \tilde{n}_t^y) + \tilde{z}_t^y + \beta \tilde{z}_t^x$$

$$\tilde{w}_t = \tilde{z}_t^x - (1 - \theta) \tilde{n}_t^x$$

$$0 = \mathbb{E}_t (\tilde{m}u_{t+1} - \tilde{m}u_t - (\tilde{p}_{t+1} - \tilde{p}_t) - \Phi_0(\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y) - \tilde{\gamma}_{t+1})$$

$$\tilde{m}u_{t+1} \equiv -\frac{\sigma}{C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}} \left(C(\tilde{w}_{t+1} - \tilde{p}_{t+1} + \tilde{n}_{t+1}^y) - N^{\frac{1+\phi}{\phi}} \left(\frac{1+\phi}{\phi} \tilde{\epsilon}_{t+1} + \left(\frac{N^x}{N} \tilde{n}_{t+1}^x + \frac{N^y}{N} \tilde{n}_{t+1}^y \right) \right) \right)$$

$$\tilde{b}_t = (1 + r) \tilde{b}_{t-1} + \frac{X}{B} (\tilde{n}_t^x - \tilde{n}_t^y)$$

This system is identical to the original where we have set $i_t = q_t = 0$ and dropped the market clearing condition in the intermediate goods market.

2.11.2 Alternative wage setting specifications

2.11.2.1 Preferences with wealth effects in labor supply

In our benchmark specification for the wage setting equation we assumed that the marginal rate of substitution between consumption and hours worked is independent of consumption (as is the case with GHH preferences). In this section we explore the consequences of moving away from this assumption for our econometric procedure in Section 2.6. For a general set of preferences represented by $u(c, n)$, we can write the marginal rate of substitution in log-

deviations from steady state as,

$$\begin{aligned} mrs_{kt} &= \left(\frac{u_{cn}c}{u_n} - \frac{u_{cc}c}{u_c} \right) c_{kt} + \left(\frac{u_{nn}n}{u_n} - \frac{u_{nc}c}{u_n} \right) n_{kt} \\ &\equiv \omega c_{kt} + \left(\omega + \frac{1}{\hat{\phi}} \right) n_{kt} \end{aligned}$$

which nest the special case with no wealth effects ($\omega = 0$) and so we obtain the marginal rate of substitution from our benchmark specification. The aggregate and state level wage setting equations become,

$$\begin{aligned} \pi_t^w &= \hat{\lambda}(\pi_t + \epsilon_t - \epsilon_{t-1} + \left(\frac{1}{\hat{\phi}} + \omega\right)(n_t - n_{t-1}) + \omega(c_t - c_{t-1})) + (1 - \hat{\lambda})\pi_{t-1}^w \\ \tilde{w}_{kt} &= \hat{\lambda}(\tilde{p}_{kt} + \epsilon_{kt} + \left(\frac{1}{\hat{\phi}} + \omega\right)\tilde{n}_{kt} + \omega\tilde{c}_{kt}) + (1 - \hat{\lambda})\tilde{w}_{kt-1} \end{aligned}$$

Replacing aggregate consumption with the model implied $w_t + n_t - p_t$ we obtain

$$\pi_t^w = \lambda\pi_t + \frac{\lambda}{\hat{\phi}}(n_t - n_{t-1}) + (1 - \lambda)\pi_{t-1}^w + \frac{\lambda}{1 - \omega}(\epsilon_t - \epsilon_{t-1})$$

where $\lambda \equiv \frac{\hat{\lambda}(1-\omega)}{1-\lambda\omega}$ and $\frac{1}{\hat{\phi}} \equiv \frac{\frac{1}{\hat{\phi}}+2\omega}{1-\omega}$. Also, we can re-write the state level equation as

$$\tilde{w}_{kt} = \lambda\tilde{p}_{kt} + \frac{\lambda}{\hat{\phi}}\tilde{n}_{kt} + (1 - \lambda)\tilde{w}_{kt-1} + \frac{\lambda\omega}{1 - \omega}(\tilde{p}_{kt} + \tilde{c}_{kt} - (\tilde{w}_{kt} + \tilde{n}_{kt})) + \frac{\lambda}{1 - \omega}\tilde{\epsilon}_{kt} \quad (5)$$

Since state level economies are open economies, in general, the term $\tilde{p}_{kt} + \tilde{c}_{kt} - (\tilde{w}_{kt} + \tilde{n}_{kt})$ will be different from zero. By omitting it in our cross-sectional regressions we could be obtaining biased estimates of $\lambda, \hat{\phi}$.

2.11.2.2 Forward looking wages

In our benchmark specification for the wage setting equation we assumed that there was no forward looking term in the target wage. In this section we explore the consequences of having a forward looking component in the wage setting equation for our econometric procedure in Section 2.6. In particular, consider the aggregate and state level wage setting equations

$$\begin{aligned}\pi_t^w &= \lambda(\pi_t + \epsilon_t - \epsilon_{t-1} + \frac{1}{\phi}(n_t - n_{t-1})) + \lambda\kappa\mathbb{E}_t[\pi_{t+1}^w] + (1 - \lambda)\pi_{t-1}^w \\ \tilde{w}_{kt} &= \lambda(\tilde{p}_{kt} + \tilde{\epsilon}_{kt} + \frac{1}{\phi}\tilde{n}_{kt}) + \lambda\kappa\mathbb{E}_t[\tilde{w}_{kt+1} - \tilde{w}_{kt}] + (1 - \lambda)\tilde{w}_{kt-1}\end{aligned}$$

where κ parameterizes the importance of the forward looking term. Also, let's consider the case where local wages follow an AR(1) process in equilibrium with coefficient $\tilde{\rho}_w$ and aggregate expected wage inflation is zero. Our model from Section 2.5, would imply this, for instance, when $\theta \rightarrow 1$ so that $\tilde{w}_{kt} = \tilde{z}_{kt}^x$ in equilibrium and $\tilde{\rho}_w = \rho_x$; and the monetary authority fully stabilizes expected aggregate nominal wage growth. We obtain,

$$\begin{aligned}\pi_t^w &= \lambda(\pi_t + \epsilon_t - \epsilon_{t-1} + \frac{1}{\phi}(n_t - n_{t-1})) + (1 - \lambda)\pi_{t-1}^w \\ \tilde{w}_{kt} &= \frac{\lambda}{1 + \lambda\kappa(1 - \tilde{\rho}_w)}(\tilde{p}_{kt} + \tilde{\epsilon}_{kt} + \frac{1}{\phi}\tilde{n}_{kt}) + \frac{1 - \lambda}{1 + \lambda\kappa(1 - \tilde{\rho}_w)}\tilde{w}_{kt-1}\end{aligned}$$

Then, we can write,

$$\lambda = \frac{1 - \beta_w}{1 + \beta_w\kappa(1 - \tilde{\rho}_w)}$$


where $\beta_w \equiv \frac{1-\lambda}{1+\lambda\kappa(1-\tilde{\rho}_w)}$. From this expression we see that our estimates for λ using cross-state variation are upward biased. However, we can get a notion on the magnitude of the bias by asking what would $\kappa(1 - \tilde{\rho}_w)$ have to be in order for λ to be less than some λ_0 . We obtain,

$$\kappa(1 - \tilde{\rho}_w) > \frac{1 - \beta_w - \lambda_0}{\beta_w \lambda_0}$$

For example, given our lower estimate for $\beta_w = 0.5$, in order for λ to be below 0.1 we would need a $\kappa(1 - \tilde{\rho}_w)$ larger than 8.

2.12 Appendix Figures

Figure B1: State Net Migration Rate 2009-2010 vs. State Employment Growth 2007-2010



Note: Figure shows state net migration rate between 2009 and 2010 against employment growth in the state during 2007-2010. Employment growth comes from the BLS and is defined in the text. State net migration rates come from American Community Survey.

2.13 Online Appendix

2.13.1 *Descriptive Statistics For Retail Scanner Data*

Online Robustness Appendix Table B8 shows descriptive statistics for the Nielsen Retail Scanner Database for each year between 2006 and 2012 and for the sample as a whole. A few things are of particular note. The sample sizes - in terms of stores covered - increased from 32,642 stores (in 2006) to 36,059 stores (in 2012). Second, notice that the number of observations (store*week*UPC code) is massive. The database includes over *90 billion* unique observations. Third, during the entire sample, there is about 1.4 million unique UPC codes within the database. On average, each year contains roughly 750,000 UPC codes. Fourth, the geographic coverage of the database is substantial in that it includes stores for about 80 percent of all counties within the United States. Moreover, the number of geographical units (zip codes, counties, MSAs, states) is very similar from year to year highlighting that the geographical coverage is consistent through time. Finally, the dataset includes between \$188 billion and \$240 billion of transactions within each year. For the time periods we study, this represents roughly 30 percent of total U.S. expenditures on food and beverages (purchased for off-premise consumption) and roughly 2 percent of total household consumption.⁴³

43. To make these calculations, we compare the total transaction value in the scanner data to BEA reports of total spending on food and beverages (purchased for off-premise consumption) and total household consumption.

2.13.2 Creating the Scanner Data Price Index

In this sub-section, we discuss our procedure for computing the Retail Scanner Price Indices. Formally, the first step is to produce a category-level price index which can be expressed as follows:

$$P_{j,t,y,k}^L = P_{j,t-1,y,k}^L \times \frac{\sum_{i \in j} p_{i,t,k} \bar{q}_{i,t-1,k}}{\sum_{i \in j} p_{i,t-1,k} \bar{q}_{i,t-1,k}}$$

where $P_{j,t,y,k}^L$ is price index for category j , in year t , with base year y , in geography k . For our analysis, geographies will either be U.S. states or the country as a whole. $p_{i,t,k}$ is the price at time t of the specific good i in geography k and $\bar{q}_{i,t-1,k}$ is the average monthly quantity sold of good i in the prior year in location k . By fixing quantities at their prior year's level, we are holding fixed household's consumption patterns as prices change. We update the basket of goods each year, and chain the resulting indices to produce one chained index for each category in each geography, denoted by $P_{j,t,k}^L$. In this way, the index for months in 2007 uses the quantity weights defined using 2006 quantities and the index for months in 2008 uses the quantity weights defined using 2007 quantities. This implies that the price changes we document below with changing local economic conditions is not the result of changing household consumption patterns. Fixing the basket also minimizes the well documented chain drift problems of using scanner data to compute price indices (Dielwert et al. (2011)). Notice, this procedure is very similar to the way the BLS builds category-level first stage for their price indices.

When computing our monthly price indices, one issue we confront is how to deal with

missing values from period to period. For example, a product that shows up in month m may not have a transacted price in month $m + 1$ making it impossible to compute the price change for that good between the two months. Missing values may be due to new products entering the market, old products withdrawing from the market, and seasonality in sales. Our results in the paper are robust to the various ways we dealt with missing values but clearly the price indices will generally differ depending on how one treats such data points. Although we could have used some ad hoc imputation methods like interpolation between observed prices or keeping a price fixed until a new observation appears, we chose to follow a more conservative approach. Looking at the above equation, we see that we can handle the missing values without imputation by restricting the goods that enter the basket to those that have positive sales over at least one month in the previous year and over the 12 months of the current year. This is what we do when creating our indices. For example, when computing the category prices in 2008 we use the reference basket for 2007. In doing so, we only take the goods that have $q_{i,2007,k} > 0$ and $q_{i,t,k} > 0$ for all $t \in 2008$.⁴⁴ This ensures that for a given product in the price index during year t , we will have a weight for this product based on $t - 1$ data and we will have a non-missing transaction price in all months in which the price index is computed during that year.⁴⁵ The bottom row of Appendix Table B8 includes the share of all expenditures (value weighted) that were included in our price index for a given year. In the five later years of the sample, our price index includes

44. The database starts in 2006. As a result, our baseline specification of the 2006 price indices only includes products that have positive sales in all months of 2006.

45. This procedure implies that we will miss products that are introduced within a given year. These products, however, will be incorporated in next year's basket as long as they have continuous sales during the subsequent calendar year.

roughly two-thirds of all prices (value weighted).

The second stage of our price indices also follows the BLS procedure in that we aggregate the category-level price indices into an aggregate index for each location k . The inputs are the category-level prices and the total expenditures of each category. Specifically, for each state we compute:

$$\frac{P_{t,k}}{P_{t-1,k}} = \prod_{j=1}^N \left(\frac{P_{j,t,y,k}^L}{P_{j,t-1,y,k}^L} \right)^{\frac{\bar{S}_{j,k}^t + \bar{S}_{j,k}^{t-1}}{2}}$$

where $\bar{S}_{j,k}^t$ is the share of expenditure of category j in month t in location k averaged over the year. We calculate the shares using total expenditure on all goods in each category, even though for the category-level indices some goods were not included due to missing data. For the purposes of this paper, we make our baseline specification one that fixes the weights of each category for a year in the same fashion as we did for the category-level indices. However, as a robustness specification, we allowed the weights in the second step to be updated monthly. The results using the two methods were nearly identical.

2.13.3 *Creating Composition Adjusted Wage Measures in the ACS and*

CPS

To make the composition adjusted wage measures in the 2000 U.S. Census and the 2001-2012 American Community Survey (ACS), we start with the raw annual data files that we downloaded directly from the IPUMS website.⁴⁶ For each year, we restrict the sample to

46. The ACS is just the annual survey which replaces the Census long form in off Census years. The national representative survey started in 2001. As a result, the Census and ACS questions are identical.

include male individuals between the ages of 25 and 55 (inclusive) who do not live in group quarters. We further restrict the sample to individuals who (1) are currently employed, (2) report working usually more than 30 hours per week (inclusive), (3) report working at least 48 weeks during the prior year, and (4) report earning at least 5000 dollars during the prior year. These latter restrictions select workers with a strong attachment to the labor force. For each individual, we create a measure of hourly wages. We do this by dividing annual labor income earned during the prior twelve month period by reported hours worked during that same time period. Our labor income also includes business income.⁴⁷ Hours worked are computed by multiplying weeks worked during the prior twelve month period by usual weekly hours worked. With the data, we compute wage measures for each year between 2000 and 2012. We wish to stress that within the ACS, the prior year refers to the prior 12 months before the survey takes place (not the prior calendar year). Individuals interviewed in January of year t report earnings and weeks worked between January and December of year $t - 1$. Individuals in June of year t report earnings between June of year $t - 1$ and May of year t . Given that the ACS samples individuals in every month, the wage measures we create for year t can be thought of as representing average wages between the middle of year $t - 1$ through middle of year t . This differs slightly from the timing in the Current Population Survey (CPS) which we discuss next.

To create measures of composition adjusted wages, we regress $\ln(wage_{it})$ on three dummies for usual weekly hours worked (hours worked 30-39, hours worked 50-59, and hours worked 60+), five age dummies (age 25-29, age 30-34, age 35-39, age 45-49, and age 50-

47. Our results were robust to excluding business income from our measure of labor earnings.

55), four education dummies (less than high school, exactly high school, exactly a bachelors degree, and more than a bachelors degree), a dummy if the individual's race is black, and dummies for the individual's citizenship. This regression is run separately for each year. We weight each regression using the individual survey weights provided by the Census and ACS. After running the regression, we compute the residuals for each individual. To rescale the residuals, we add back in the regression constant for all individuals. Finally, we convert the composition adjusted log wage residuals back to levels. To make the annual state level composition adjusted wage index, we just take the weighted average of wage residuals across individuals in each state separately for each year.

To examine longer aggregate trends in composition adjusted wages, we use data from the March Current Population Survey. We download the data directly from the IPUMS website. As with the ACS data, we restrict the sample to men between the ages of 25 and 55 who do not live in group quarters and who have a strong attachment to the labor force (currently employed, worked 48 weeks during the prior year, and currently report working at least 30 hours per week during a usual week). We also restrict the data to those individuals with positive sample weights. Our procedure for making composition adjusted wages are identical to above aside from the following three changes: (1) we do not control for citizenship given that citizenship is not consistently measured over our sample period, (2) our sample period is 1977 through 2012, and (3) we pool years together when running regressions to increase power. In particular, we run our regression on three separate time periods: 1977-1995, 1996-2000, and 2000-2012. We break the sample into three periods because the CPS earnings questions changed in 1995 and the CPS expanded its sample in

2000.

2.13.4 BLS Metro Area Price Indices

The U.S. Bureau of Labor Statistics produces 27 metro areas price indices at various degrees of time aggregation.⁴⁸ For each metro area, the BLS publishes multiple price indices (food, nondurables, etc.). Although there are only 27 MSAs, we can still explore the relationship between unemployment growth between 2007 and 2010 and the cumulative inflation rate between 2007 and 2010 using these data. We then can compare the relationship between local unemployment growth and local scanner price inflation that we document in Section 3 with the relationship between local unemployment growth and local inflation computed using the BLS data. Given the price indices are only provided semi-annually for many MSAs, we compare the change in the unemployment rate and prices from the first half of 2007 to the latter half of 2010. When data is provided monthly or bi-monthly, we simply take the geometric average over the first and last half of the year to make the data semi-annual. As with the results in the main paper, we use data from the BLS's local area unemployment statistics to measure the percentage point change in the local unemployment rate.

Using this data, we regress the 3-year inflation rate at the MSA level on the 3 year change in the unemployment rate. We run this regression for various inflation measures (food, services, all goods less housing, etc.). The results of these simple regressions using the BLS

48. The 27 MSAs are (in order of reporting frequency): Chicago, Los Angeles, New York, Atlanta, Boston, Cleveland, Dallas-Fort Worth, Detroit, Houston, Miami, Philadelphia, San Francisco, Seattle, Washington, Anchorage, Cincinnati, Denver, Honolulu, Kansas City, Milwaukee, Minneapolis, Phoenix, Pittsburgh, Portland, St. Louis, San Diego, and Tampa. The first three MSAs have price indices that are reported at monthly frequencies. The next 11 MSAs have price indices that are reported at bi-monthly frequencies. The last 13 price indices are reported semi-annually.

data are strikingly similar to our results using the scanner data. For example, the results using the scanner index find that a 1 percentage point increase in the local unemployment rate is associated with a 0.38 decline in the local inflation rate (for the specification where goods are defined as UPC-store pairs). Within the BLS data, we find that a 1 percentage point increase in the local unemployment rate is associated with a 0.34 percentage point decline in the local food inflation rate (standard error = 0.22).

Also, as predicted by our simple model in the paper, the relationship between the inflation rate for all goods and the local unemployment rate change should be higher than the relationship between food inflation and the local unemployment rate change if food is relatively more tradable than the local consumption good. We cannot test this within our scanner data. However, the BLS data allows us to test this prediction directly. Within the BLS data, the relationship between the inflation rate for all goods with the change in the unemployment rate is in fact higher at -0.47 (standard error = 0.15). The fact that the coefficient is larger in magnitude is consistent with our belief that the variation in packaged food data across regions should be a lower bound for the variation in the average consumption good given that the packaged goods in our dataset are relatively more tradable.

The patterns that we document for the 2007-2010 period with the BLS data are very consistent with the results in Fitzgerald and Nicolini (2014) that use the BLS MSA level price indices to show these relationships over a much longer time period. Fitzgerald and Nicolini (2014) find that over the period of 1976-2010, a 1 percentage point increase in the local unemployment rate is associated with a 0.3 percentage point decline in the local annual inflation rate. In summary, a limitation of our price data is that it only covers goods

sold in grocery, pharmacy and mass-merchandising stores. The fact that the patterns we uncover are nearly identical for similar goods using the BLS metro area price indices is very reassuring. Moreover, as predicted, using broader measures of goods in the BLS data seems to only strengthen the cross-region variation. Our data is an advance over the BLS metro level data in that can be calculated for every state and, if one desires, a much larger set of metro areas.

2.13.5 Cross-State Wage Patterns in the QEW

While there are no government data sets that produce broad based composition adjusted wage series at the local level, the Bureau of Labor Statistics's Quarterly Census of Employment and Wages (QEW) collects firm level data on employment counts and total payroll at local levels. These measures are broad based in that the underlying data are collected as part of the state and federal unemployment insurance programs and covers roughly 98 percent of workers in the U.S.. Using this data, yearly earnings-per-worker can be computed at the state level. This measure is an imperfect measure of wages in that it is not adjusted for cyclical movements in hours worked. Additionally, the measure does not adjust for changes in the composition of workers over the business cycle. Finally, the earnings measures reported include wages and salary as well as bonuses, stock options, and in some states, contributions to deferred compensation plans. These latter measures are not included in the ACS wage indices.

Despite these differences, the cross state correlation between growth in our composition adjusted wage index from the ACS and the growth in earnings-per-worker from QEW is

quite high. Online Appendix Figure B2 shows the simple scatter plot of the growth in the ACS and QEW wage measures between 2007 and 2010. If we fit a line through the scatter plot, the slope coefficient is 0.72 (standard error = 0.20). The correlation between the two measures is about 0.5. Online Appendix Figure B3 shows that even within the QEW data, there is a strong relationship between employment growth and earnings-per-worker growth during the Great Recession. The x -axis of Online Appendix Figure B3 is QEW employment growth between 2007 and 2010. QEW employment growth is essentially perfectly correlated with the employment growth measure we use from the BLS in Tables 1 and 2. The y -axis of Online Appendix Figure B3 is QEW nominal earnings per worker growth between 2007 and 2010. As seen from the figure, places with lower employment growth had lower nominal wage growth. The slope coefficient from the line in the scatter plot is 0.45 (standard error = 0.07). This is very similar to the estimated relationship between employment growth and ACS nominal growth during the same time period as show in Table 1 of the main paper.

2.13.6 Accounting for Measurement Error in Wages and Prices

Formally, in Section 6 of the main paper, we estimate the following specification using our regional data to obtain estimates of λ and ϕ :

$$\pi_{kt}^w = b_0 + b_1\pi_{kt} + b_2(n_{kt} - n_{k,t-1}) + b_3\pi_{k,t-1}^w + \Psi D_t + \Gamma X_k + e_{kt}$$

where $b_1 = \lambda$, $b_2 = \lambda/\phi$, $b_3 = (1 - \lambda)$, and $b_t = \lambda(u_t^\epsilon - (1 - \rho_\epsilon)\epsilon_{t-1})$. However, given our data construction procedures, the regional inflation rate (π_{kt}), regional nominal wage growth

(π_{kt}^w) , and regional lagged nominal wage growth (π_{kt-1}^w) are all measured with error. If the measurement error were classical, our estimates of b_1 and b_3 would be attenuated. Moreover, given that we measure nominal wages in each period and then compute the changes over time to get the growth rates, any measurement error in our nominal wage measure in a given year will induce a negative correlation between π_{kt}^w and π_{kt-1}^w . This type of measurement error would cause b_3 to be biased downward.

Given that we want to recover structural parameters, we take these measurement errors seriously. To account for the measurement error in π_{kt} , π_{kt}^w , and π_{kt-1}^w , we exploit the large sample sizes underlying our the construction of local price and wage measures. We begin by addressing the measurement error in prices. Using the underlying data in from Nielsen, we split the data in half by product categories. As discussed, the Nielsen data includes roughly 1,000 product categories. We split the underlying data into two groups of categories: sample 1 includes all the odd category numbers (1, 3, 5, etc.) while sample 2 includes all the even category numbers (2, 4, 6, ect.). We then construct state level prices indices for each month using the data from sample one (P_{kt}^{samp1}) and then separately using the data from sample two (P_{kt}^{samp2}). When running our key equations, we instrument the inflation rate using the price indices computed with data from sample 1 (π_{kt}^{samp1}) with the inflation rate using the price indices computed with the data from sample 2 (π_{kt}^{samp2}). Not surprisingly, the annual inflation rates obtained from each of the two separate data samples are highly correlated. Specifically, running the first stage equation on our annual data from 2007-2011:

$$\pi_{kt}^{samp1} = \psi_0 + \psi_1 \pi_{kt}^{samp2} + \psi_2 D_t + \psi_{error}$$

yields an estimate of $\psi_1 = 0.268$ (standard error = 0.06) with an adjusted R-squared of 0.84. The F-stat from including π_{kt}^{samp2} is 19.6. Using the above relationship, we can make a predicted inflation measure:

$$\hat{\pi}_{kt} = \hat{\psi}_0 + \hat{\psi}_1 \pi_{kt}^{samp2} + \hat{\psi}_2 D_t$$

It is this predicted inflation measure that we include in our estimating equation. We adjust standard errors accordingly to account for the predicted regressor.

We perform a similar methodology to adjust for the measurement error in wages. Specifically, we use a random number generator to split the underlying micro data from the ACS into two equal sized samples. Individuals within a year with a random number less than or equal to 0.5 go into sample 1 and individuals within a year with a random number greater than 0.5 go into sample 2. As a result, for each year, we will have three samples of data within each year: sample 1 individuals, sample 2 individuals, and all individuals. Within each sample, we can make measures of adjusted nominal wages (where the adjustment for observables as is discussed in the paper). With the adjusted nominal wages for each year and each sample, we can compute nominal wage growth for each year-sample pair. To account for the measurement error in nominal wage growth, we run the following two "first stage" equations using our annual data from 2007-2010:

$$\pi_{kt}^w = \psi_0^1 + \psi_1^1 \pi_{kt}^{w,samp1} + \psi_2^1 D_t + \psi_{error}^1$$

$$\pi_{kt-1}^w = \psi_0^2 + \psi_1^2 \pi_{kt-1}^{w,samp2} + \psi_2^2 D_t + \psi_{error}^2$$

Specifically, we use sample 1 data to predict contemporaneous wage growth and use sample 2 data to predict lagged wage growth. This ensures that any measurement error in our current wage growth measure will not be correlated with the measurement error in our lagged wage growth measure. Like with the price data, the wage data are highly correlated across the samples. For the contemporaneous wage growth equation, our estimate of ψ_1^1 is 0.80 (with a standard error of 0.03) and an adjusted R-squared of 0.89. The F-stat from including $\pi_{kt}^{w,samp1}$ is 78.3. For the lagged wage growth equation, our estimate of ψ_1^2 is 0.30 (with a standard error of 0.04) and an adjusted R-squared of 0.66. The F-stat from including $\pi_{kt-1}^{w,samp2}$ is 75.6. Using these regressions, we can make measurement error adjusted current and lagged wage growth measures:

$$\hat{\pi}_{kt}^w = \hat{\psi}_0^1 + \hat{\psi}_1^1 \pi_{kt}^{w,samp1} + \hat{\psi}_2^1 D_t$$

$$\hat{\pi}_{kt-1}^w = \hat{\psi}_0^2 + \hat{\psi}_1^2 \pi_{kt-1}^{w,samp2} + \hat{\psi}_2^2 D_t$$

It is these predicted measures that we use in our estimation equations. Specifically, our key estimating equation can be expressed as:

$$\hat{\pi}_{kt}^w = b_0 + b_1 \hat{\pi}_{kt} + b_2(n_{kt} - n_{k,t-1}) + b_3 \hat{\pi}_{k,t-1}^w + \Psi D_t + \Gamma X_k + e_{kt}$$

In the main text of the table, we suppress the hats on the variables. We adjust the standard errors by bootstrapping to account for the fact that the measures of price and wage growth are estimated.

2.13.7 Variation in Government Policy Changes Across States During Great Recession

In the main text, we argued that we could identify unbiased estimates of λ and ϕ using regional variation via OLS if $v_{kt}^\epsilon = 0$. This assumption would be violated if government policy changed differentially across states in a way that discouraged labor supply. In this section of the appendix we show that the difference in many such policy changes across states were small (particularly relative to the aggregate changes) and that these policy changes - to the extent that they did occur - occurred after 2009. Specifically, we focus our attention on four such government policies: state income tax rates, federal food assistance programs, federal programs to help underwater homeowners renegotiate their mortgage contract, and the extension of unemployment benefits. Our analysis shows that for these major policies, our assumption that these policies varied little across regions and the extent to which they did vary was uncorrelated with local measures of economic activity is not at odds with the data.

Using data on statutory tax rates by state from the tax foundation and micro data on

individual incomes by state from the American Community Survey, we compute the average marginal tax rate for each state in 2007 and 2010. State income tax changes were not that common during this time period. Roughly 90 percent of the states, population weighted, had essentially no change in their average marginal tax rate during this time period. As seen from Online Appendix Figure B4, the extent to which the average marginal tax rate changed between 2007 and 2010 was uncorrelated with state employment or price growth between 2007 and 2010.

Likewise, Online Appendix Figure B5 shows no correlation in the growth in benefits from the federal Supplemental Nutrition Assistance Program (SNAP) and employment growth across states during the 2007-2010 period. SNAP is the successor to the federal Food Stamps program. This program was expanded dramatically during the Great Recession. Given that the SNAP program is means tested, an expansion of the program can discourage work effort. This point is made by Mulligan (2012). Using summary data from the US Department of Agriculture, we measure the dollar per SNAP recipient for each state between 2007 and 2010. The average recipient received an increase about 33 percent during the 2007 to 2010 period. Yet, there was very little regional variation in the increase (a standard deviation of only 4 percent across the states, population weighted). As seen from Online Appendix Figure B5, the variation that did occur across states was uncorrelated with state unemployment growth. So, while the increase in SNAP benefits may have discouraged labor supply at the aggregate level, there is very little variation across U.S. states.⁴⁹ This

49. We also explored whether the eligibility for SNAP differed across states in a way that is correlated with the change in state economic conditions between 2007 and 2010. We found no evidence suggesting such a relationship.

is not to say that SNAP total dollars paid did not differ across states. It is saying that the dollar per recipient did not vary across states. If a local shock hit that resulted in people not working, the number of people eligible for SNAP would increase. We are showing that all the variation in SNAP dollars paid across states was due to the changing number of recipients across states NOT the change in the dollar per recipient.

Another new federal program that was means tested and was argued to possibly discourage work effort during the latter part of the Great Recession was the Home Affordable Modification Program (HAMP). HAMP was designed to help homeowners who were underwater renegotiate their mortgage. The program was authorized in early 2009 but there were no significant modifications taking place until mid 2010. By the end of 2010, only about 0.5% of households had participated in the program. As seen from Online Robustness Appendix Figure B6, the extent to which households participated in the program varied slightly with the state's employment growth between 2007 and 2010. A one percent decline in employment growth was only associated with a 0.07 percentage point increase in HAMP take up (i.e., from 0.50% to 0.57%). This effect is very small. However, given that there was some relationship between HAMP take up and underlying economic conditions within the state, we perform a robustness specification in our cross sectional estimation that excludes 2009 and 2010 data and therefore only focuses on the periods before HAMP went into effect. As discussed in the main text, these estimates were very similar to our baseline estimates. Additionally, we excluded states with the highest amount of loan modifications (CA, FL, NV and AZ) as a robustness specification. Our results were unchanged when these states were excluded.

Finally, we explored the extent to which unemployment benefits were differentially extended at the state level. During the 2007-2011 period, unemployment benefits increased from about 26 weeks per recipient to upwards of 99 weeks per recipient in some states. Some researchers have argued that this large increase in the duration of unemployment benefits can explain only a small portion of the aggregate decline in employment (Rothstein (2012)) while others have argued that it can explain a more substantive portion of the aggregate decline in employment (Hagedorn et al. (2013)). By law in 2010, weeks of unemployment benefits were tied to the state's unemployment rate. This implies that there will be a correlation between the total amount of unemployment benefit extension within the state and the change in the state employment rate between 2007 and 2010. Online Appendix Figure R6 shows this correlation. As of 2010, 70 percent of U.S. states had a duration of unemployment benefits that exceeded 86 weeks. These states represent roughly 90 percent of the U.S. population. However many smaller states, mostly in the Plains region of the U.S., had small employment declines and only an extension of unemployment benefits from 60-85 weeks.⁵⁰

A simple regression line through Online Robustness Figure B7 shows that a 1 percent decline in employment growth was associated with an additional 1.5 weeks of unemployment benefits. Again, this is a tiny change. However, as discussed in the main text, we can reestimate our model focusing only on data before the unemployment benefit extension occurred (i.e., prior to 2010). As we show in the main text, the results are nearly identical

⁵⁰. States also had some discretion as to whether they opted into the program. This explains why some states did not have the maximum weeks of unemployment benefits even when their unemployment rate was higher.

to our base specification. Additionally, we can include only those states that had an increase in unemployment benefits to at least 86 weeks (using all years of our data).⁵¹ If pooling together states that had large and small unemployment benefits were biasing our estimates of λ and ϕ , our estimates would change once we excluded the low unemployment benefit during states. It is comforting that our estimates did not change at all when these states were excluded. This says that the fact that unemployment benefit durations changed differentially across states is not biasing our results in any substantive way. This is not surprising when one realizes that essentially all states (population weighted) had increases in unemployment benefit duration to at least 86 weeks.

2.13.8 Controlling for Industry Controls in Estimates λ and ϕ

Charles et al. (2013) document that the secular decline in manufacturing depressed employment rates during the 2000s. Autor and Dorn (2013) show that declines in routine employment also depressed employment rates during the 2000s. Both of these prior papers exploit variation across either MSAs or commuting zones. As we show in Online Robustness Appendix Figure B8, there is very little variation in routine occupation shares across U.S. states as of 2007 and the extent to which there is variation, it is uncorrelated with employment growth during the 2007-2010 period. We define routine occupations to include all manufacturing and administrative occupations. Although not shown, a similar pattern exists for just manufacturing employment as of 2007. This is consistent with the results in

51. The excluded states are: Arkansas, Iowa, Louisiana, Maryland, Mississippi, Montana, Nebraska, New Hampshire, North Dakota, Oklahoma, South Dakota, Utah, and Wyoming.

Charles et al. that most of the decline in manufacturing employment within the U.S. during the last 15 years took place in the early 2000s. As a robustness exercise, we control for the 2007 share of routine jobs within each state and the 2007 share of manufacturing jobs in each state. As seen from the main text, adding these controls do not alter our estimates in any meaningful way. We also explored adding more detailed industry controls (at the one digit level). Again, these controls did not alter our estimates of λ and ϕ in any meaningful way.

2.13.9 IV Procedure to Estimate λ and ϕ

As a specification check on our estimation procedure, we attempted to isolate variation in labor demand across states that were orthogonal to shocks to labor supply. We build on the work of many others (including Mian and Sufi 2014) by looking at changes in housing prices. Instead of instrumenting housing prices with local housing supply elasticities to isolate a causal effect of housing price changes, we use housing prices directly as a instrument. We are not interested in getting a causal effect of housing prices on local labor markets. Instead, we are interested in isolating movements in wage growth, price growth, and employment growth that is orthogonal to changes in the taste for leisure at the local level. The key condition we need is that (1) the house price changes were not caused by changes in the taste for leisure and (2) the house price changes did not cause a change in the taste for leisure, our identification strategy will be valid. Like with our OLS estimates, we estimate the IV procedure over two time periods - the full 2007-2010 period and the shorter 2007-2009 period (prior to the potentially reactionary change in state policy variables).

Specifically, we use two instruments: contemporaneous and lagged housing price growth. We instrument for contemporaneous employment growth and price growth. Given we restrict the coefficient on lagged wage growth to be one minus the coefficient on price growth, we only need two instruments to estimate λ and ϕ . Not surprisingly, the housing variables predict both employment growth and lagged wage growth. The first stage F-stat for the instruments in predicting employment growth is 17.3.

2.13.10 *Parameter Calibration*

In this section we explain how we determined the values of the following parameters used in the paper (see Table 2 in the paper): ρ_z (persistence of productivity shock), ρ_ϵ (persistence of the labor supply shock), and β (intermediates share in the non-tradable sector).

2.13.10.1 Persistence Parameters in Shock Processes

To measure the persistence parameters of the productivity/mark-up shock and labor supply/leisure shock we use two equations from the model from which we can back them out. For the productivity/mark-up process, z_t , we use the aggregate labor demand equation:

$$w_t^r = -(1 - (\alpha + \theta\beta))n_t + z_t$$

For the labor supply/leisure process, ϵ_t , we use the aggregate wage-setting equation:

$$\pi_t^w = \frac{\lambda}{1 - \lambda}(\epsilon_t + \frac{1}{\phi}n_t - w_t^r)$$

In these equations, we plug in the values for the other parameters and the aggregate data that we used in our VAR and solve for z_t and ϵ_t . With the time series of these processes we measured their AR(1) persistence coefficient obtaining $\rho_z = 0.76$ and $\rho_\epsilon = 0.66$

2.13.10.2 Input Shares

Here we just explain how we computed the values for β , the share of intermediate inputs.

The labor shares in the non-tradable and tradable sector are, respectively:

$$\alpha = \frac{wN^y}{PY^y}$$

$$\theta\beta = \frac{wN^x}{PY^y}$$

The aggregate employment share in the economy is

$$\vartheta = \frac{wN^x + wN^y}{QY^x + PY^y}$$

Divide through by PY^y

$$\vartheta = \frac{\frac{wN^x}{PY^y} + \frac{wN^y}{PY^y}}{\frac{QY^x + PY^y}{PY^y}} = \frac{\theta\beta + \alpha}{1 + \frac{QY^x}{PY^y}}$$

Rearranging we get β .

$$\theta\beta + \alpha = \vartheta \left(1 + \frac{QY^x}{PY^y} \right)$$

$$\theta\beta + \alpha = \vartheta(1 + \beta)$$

$$\beta = \frac{\vartheta - \alpha}{\theta - \vartheta}$$

We used values for employment shares in the economy, the tradable and non-tradable sector in 2006 to get a value for β of 0.16.

2.13.11 *Estimating the Regional Shocks*

The procedure for estimating regional shocks and elasticities follows a similar logic than the one for aggregate shocks. The recursive solution to the equilibrium system of equations in Lemma 2 can be written as a SVAR(∞) in $\{\tilde{p}_t, \tilde{w}_t, \tilde{n}_t^y\}$ when $\tilde{v}_t^\epsilon = 0$. Given that there are four shocks at the local level, we need one further identification restriction. We set $\tilde{v}_t^\epsilon = 0$ and provide some evidence for this choice in the next section. The regional SVAR, therefore, can be expressed as follows:

$$(I - \tilde{\rho}(L)) \begin{bmatrix} \tilde{p}_t \\ \tilde{w}_t \\ \tilde{n}_t^y \end{bmatrix} = \tilde{\Lambda} \begin{bmatrix} \tilde{v}_t^y \\ \tilde{v}_t^x \\ \tilde{v}_t^\gamma \end{bmatrix}$$

Again, this vector autoregression representation is consistent with a much more general class of models than the one characterized in Lemma 2. From here on we will be working with a subset of these general class of models such that the regional wage setting equation

in log-linearized form holds,

$$\tilde{w}_t = \lambda \left(\tilde{p}_t + \frac{N^y}{N\phi} \tilde{n}_t^y + \frac{N - N^y}{N\phi(1 - \theta)} (\tilde{z}_t^x - \tilde{w}_t) \right) + (1 - \lambda) \tilde{w}_{t-1}$$

This equation is obtained from replacing the tradable goods labor demand and labor market clearing condition into the wage setting equation.

The first step in the procedure consists in estimating the reduced form VAR via OLS to obtain the autoregressive matrix $\tilde{\rho}(L)$ and the reduced form errors covariance matrix \tilde{V} . In practice we will truncate $\tilde{\rho}(L)$ to be of finite order.

We now derive the identification restrictions that allow us to estimate $\tilde{\Lambda}$ and the shocks. Applying the conditional expectation operator $\mathbb{E}_{t-1}(\cdot)$ on both sides of the wage setting equation and constructing the reduced form expectational errors we obtain,

$$\left[\begin{array}{ccc} 1 & - \left(\frac{1}{\lambda} + \frac{N - N^y}{N\phi(1 - \theta)} \right) & \frac{N^y}{N\phi} \end{array} \right] \Lambda \begin{bmatrix} \tilde{v}_t^y \\ \tilde{v}_t^x \\ \tilde{v}_t^\gamma \end{bmatrix} + \frac{N - N^y}{N\phi(1 - \theta)} \sigma^x \tilde{v}_t^x = 0 \quad (6)$$

Similarly, constructing $\mathbb{E}_{t-1}(\cdot) - \mathbb{E}_{t-2}(\cdot)$, we obtain

$$\left(\left[\begin{array}{ccc} 0 & \frac{1 - \lambda}{\lambda} & 0 \end{array} \right] + \left[\begin{array}{ccc} 1 & - \left(\frac{1}{\lambda} + \frac{N - N^y}{N\phi(1 - \theta)} \right) & \frac{N^y}{N\phi} \end{array} \right] \tilde{\rho}_1 \right) \Lambda \begin{bmatrix} \tilde{v}_{t-1}^y \\ \tilde{v}_{t-1}^x \\ \tilde{v}_{t-1}^\gamma \end{bmatrix} - \frac{1}{\phi} \frac{N - N^y}{N(1 - \theta)} \rho^x \sigma^x \tilde{v}_{t-1}^x = 0$$

where $\tilde{\rho}_1$ is the matrix collecting the first order autoregressive coefficients in the SVAR. As

with the procedure for identifying aggregate shocks, we can identify the impulse matrix $\tilde{\Lambda}$ with these extra linear restrictions, the orthogonality of regional shocks, and the model-implied co-movement on impact of the variables in the SVAR.

2.13.12 Findings at the regional level

We follow the procedure in Section 2.13.11. We first estimate the VAR with two lags in state-level non-tradable employment, price and nominal wage growth via OLS equation by equation.⁵² All variables are expressed in log-deviations from their average weighted by population in 2006. We pool all data between 2006 and 2011, and estimate common autoregressive coefficients and reduced form errors covariance matrix for all states. In our benchmark specification, we set $\{\phi, \lambda\}$ equal to our estimates from our base specification in the paper. We set $\theta = 0.55$ to match the labor share in the manufacturing sector in the US and $\frac{N^y}{N} = 0.85$ to match the share of total employment in the service sector plus self-employed/family workers as reported in the BLS.

Table B9 summarizes the contribution of the discount rate shock and the combined productivity/markup shocks to non-tradable employment, wages and prices. We define shock j 's contribution to the change in variable y between 2007 and 2010 as

$$\xi_y^j = \frac{\frac{1}{\sum_k \omega_k (\Delta^j \tilde{y}_k - \Delta \tilde{y}_k)^2}}{\sum_j \frac{1}{\sum_k \omega_k (\Delta^j \tilde{y}_k - \Delta \tilde{y}_k)^2}}$$

where ω_k are population weights in 2006; $\Delta \tilde{y}_k$ is the change in variable y in state k between

52. We define non-tradable employment in a state as the employment rate in the service and retail sectors combined.

2007 and 2010; and $\Delta^j \tilde{y}_k$ is the counterfactual change if only shock j would have occurred. Note that ξ_y^j is always in $[0, 1]$ and increases when the counterfactual and actual changes in variable y are close to each other.

For our benchmark $(\lambda, \phi) = (0.7, 2)$ we find that the discount rate shock contributed 79 percent to the change in non-tradable employment between 2007 and 2010 across all states; 41 percent to local price changes and 24 percent to local wage changes. We obtain similar numbers for $(\lambda, \phi) = (0.5, 1)$. We conclude that the discount rate shock was the main driver of regional variation in non-tradable employment.

Table B10 summarizes a second set of results from the regional counterfactuals. We characterize the joint distribution of cumulative growth rates between 2007 and 2010 for each variable across states with two statistics: the variance and the correlation with each other. We compare the actual statistics from our data in Section 3 with the counterfactuals obtained when simulating one shock at a time. Again, we find that the discount rate shock alone can generate 98 percent of the cross-state variance of non-tradable employment growth; 69 percent of the price growth variance; and 47 percent of the nominal wage growth variance. Moreover, it reproduces the right sign for the cross-state correlations of price growth and non-tradable employment growth; nominal wage growth and non-tradable employment growth; and nominal wage growth and price growth. Although, quantitatively, it generates a larger correlation between prices and non-tradable employment than in the data. Both productivity/markup shocks combined can explain only 35 percent of the non-tradable employment growth variance across states. They do as good a job as the discount rate in explaining 85 percent of the variation in price growth and a worse job in explaining 26 per-

cent of nominal wage growth variance. However, they imply negative correlations between price/wage growth and non-tradable employment growth. The opposite is observed in the data. We conclude that the discount rate shock alone does a fairly good job in reproducing the regional patterns that we documented in the paper. To improve the estimated fit, non-tradable and tradable productivity/markup shocks are needed. However, their contribution to aggregate employment changes are small.

2.13.13 Online Appendix Figures

Table B8: Descriptive Data for the Nielsen Scanner Price Data

	Individual Years							Combined Years	
	2006	2007	2008	2009	2010	2011	2012	Total	Average
Number of Obs. (million)	12,013.1	12,812.2	13,037.5	12,968.3	13,153.4	13,646.7	13,618.8	90,250.1	13,035.7
Number of UPCs	725,224	762,469	759,989	753,984	739,768	742,074	753,318	1,425,484	748,118
Number of Chains	86	85	87	86	86	86	82	88	85
Number of Stores	32,642	33,745	34,830	35,343	35,807	35,645	36,059	39,368	34,867.3
Number of Zip Codes	10,869	11,123	11,357	11,476	11,589	11,639	11,626	11,797	11,382.7
Number of Counties	2,385	2,468	2,500	2,508	2,519	2,526	2,547	2,568	2,493.3
Number of MSAs	361	361	361	361	361	361	361	361	361
Number of States	49	49	49	49	49	49	49	49	49
Transaction Value (US billion)	187.9	207.8	219.6	223.7	227.6	235.2	239.5	1541.2	220.2
Pct. Value used in Price Index	54.3%	50.0%	66.4%	66.0%	68.3%	68.0%	67.7%	63.4%	63.0%

Note: Table shows descriptive statistics for the underlying data that we used to create our Scanner Price Index using the Nielsen Retail Scanner Data.

Table B9: **Discount rate (γ) and Productivity/Markup (z) shocks' contribution to change in regional variables**


(λ, ϕ)		\tilde{w}	\tilde{p}	\tilde{n}^y
(0.7,2)	$\tilde{\zeta}^{\gamma}$	24	41	79
	$\tilde{\zeta}^{z^y}$	17	39	11
	$\tilde{\zeta}^{z^x}$	60	20	11
(0.5,1)	$\tilde{\zeta}^{\gamma}$	19	40	79
	$\tilde{\zeta}^{z^y}$	14	41	11
	$\tilde{\zeta}^{z^x}$	67	19	10

Table B10: **Regional counterfactual statistics**

	Data	Discount rate	Productivity/Markup
$\sigma_n^2 / (\sigma_n^{data})^2$	1	0.98	0.35
$\sigma_p^2 / (\sigma_p^{data})^2$	1	0.69	0.62
$\sigma_w^2 / (\sigma_w^{data})^2$	1	0.47	0.26
β_{p,n^y}	0.48	0.73	-1.82
β_{w,n^y}	0.51	0.47	-0.66
$\beta_{p,w}$	0.32	0.65	0.37

Note: The first three lines in the table show the counterfactual variance across states relative to the actual variance of the total percent change in each variable between 2007-2010. The last three lines show the population weighted OLS coefficient corresponding to each variable pair. For example, β_{p,n^y} is the coefficient in the regression of price growth between 2007-2010 onto employment growth in the non-tradable sector where each state is weighted by its population in 2006. The second column corresponds to the counterfactual with the γ shock alone. The third column corresponds to the counterfactual with both z^x, z^y shocks and no γ shock.

Figure B2: QEW Nominal Per Worker Earnings Growth vs. ACS Adjusted Nominal Wage Growth, 2007-2010




Note: Figure shows scatterplot of nominal wages from QEW against ACS nominal wage

Figure B3: QEW Nominal Per Worker Earnings Growth vs. QEW Employment Growth, 2007-2010



Note: Figure shows scatterplot of nominal wages from QEW against QEW employment growth.

Figure B4: Change in State Tax Rate vs. State Employment Growth: 2007-2010




Note: Figure shows the change in the average marginal tax rate in a state between 2007 and 2010 against employment growth in the state during the same period. Employment growth comes from the BLS and is defined in the text. To compute the average marginal tax rate in the state we use data from the American Community Survey and state tax rate formulas from taxfoundatoin.org. Using the American Community Survey, we compute the fraction of state residents in 15 labor income bins as well as the mean income within each bin. We then compute the marginal tax rate in that bin. Averaging over the bins, we get the state average marginal tax rate. Our procedure does not account for any state level deductions or exemptions. Additionally, it assumes no one files jointly. It is meant to give a summary statistic for the state average marginal tax rate.

Figure B5: State SNAP Growth vs. State Employment Growth: 2007-2010



Note: Figure shows the change in SNAP payment growth per recipient at the state level between 2007 and 2010 against employment growth in the state during the same period. Employment growth comes from the BLS and is defined in the text. SNAP growth per recipient was collected from <http://www.fns.usda.gov>

Figure B6: State HAMP Take-Up vs. State Employment Growth 2007-2010




Note: Figure shows number of households participating in HAMP programs in 2010 against employment growth in the state during 2007-2010. Employment growth comes from the BLS and is defined in the text. HAMP participation comes from <http://www.treasury.gov>.

Figure B7: Max Unemployment Benefit Receipt in 2010 vs. State Employment Growth 2007-2010



Note: Figure shows the maximum number of unemployment benefits allowed in state in 2010 against employment growth in the state during 2007-2010. Employment growth comes from the BLS and is defined in the text.

Figure B8: Routine Share of Employment in 2007 vs. State Employment Growth 2007-2010



Note: Figure shows the routine share of employment in the state in year 2007 against employment growth in the state during 2007-2010. Routine employment is defined as anyone working in a manufacturing or administrative job. Routine employment shares are computed from the 2007 American Community Survey.

CHAPTER 3

MORTGAGE-BACKED SECURITIES AND THE FINANCIAL CRISIS OF 2008: A POST MORTEM

written jointly with Harald Uhlig

3.1 Abstract

We examine the payoff performance, up to the end of 2013, of non-agency residential mortgage-backed securities (RMBS), issued up to 2008. For our analysis, we have created a new and detailed data set on the universe of non-agency residential mortgage backed securities, per carefully assembling source data from Bloomberg and other sources. We compare these payoffs to their ex-ante ratings as well as other characteristics. We establish five facts. First, the bulk of these securities was rated AAA. Second, AAA securities did ok: on average, their total cumulated losses up to 2013 are under six percent. Third, the subprime AAA-rated RMBS did particularly well. Forth, the bulk of the losses were concentrated on a small share of all securities. Fifth, later vintages did worse than earlier vintages. Together, these facts call into question the conventional narrative, that improper ratings of RMBS were a major factor in the financial crisis of 2008.

3.2 Introduction

post mortem: an examination of a dead body to determine the cause of death.

Gradually, the deep financial crisis of 2008 is in the rearview mirror. With that, standard narratives have emerged, which will inform and influence policy choices and public perception in the future for a long time to come. For that reason, it is all the more important to examine these narratives with the distance of time and available data, as many of these narratives were created in the heat of the moment.

One such standard narrative has it that the financial meltdown of 2008 was caused by an overextension of mortgages to weak borrowers, repackaged and then sold to willing lenders drawn in by faulty risk ratings for these mortgage backed securities. To many, mortgage backed securities and rating agencies became the key villains of that financial crisis. In particular, rating agencies were blamed for assigning the coveted AAA rating to many securities, which did not deserve it, particularly in the subprime segment of the market, and that these ratings then lead to substantial losses at institutional investors, who needed to invest in safe assets and who mistakenly put their trust in these misguided ratings.

In this paper, we re-examine this narrative. We seek to address two questions in particular. First, were these mortgage backed securities bad investments? Second, were the ratings wrong? We answer these questions, using a new and detailed data set on the universe of non-agency residential mortgage backed securities (RMBS), obtained by devoting considerable work to carefully assembling source data from Bloomberg and other sources. This data set allows us to examine the actual repayment stream and losses on principal on these securities up to 2014, and thus with a considerable distance since the crisis events. We find that the standard narrative needs substantial rewriting: the ratings and the losses were not nearly as bad as one is often led to believe.

Specifically, we calculate the ex-post realized losses as well as ex-post realized return on investing on par in these mortgage backed securities, under various assumptions of the losses for the remaining life time of the securities. We will also calculate these returns, when purchasing these securities at market prices in 2008 and 2009. We compare these realized returns to their ratings in 2008 and their promised loss distributions, according to tables available from the rating agencies. We shall investigate, whether ratings were a sufficient statistic (to the degree that a discretized rating can be) or whether they were, essentially, just “noise”, given information already available to market participants at the time of investing such as ratings of borrowers. We also compare the realized returns to the returns on government bonds, in order to judge whether investors shunning these securities in favor of the latter would have fared better, and how much, if so.

It is obvious, that such an ex post evaluation is different from the perhaps more appropriate evaluation ex ante. Ideally, one would take investments in these securities in 2008, evaluate their payoffs in all possible future states of the world, appropriately discounted and compare that to the original investment or their later market prices. Likewise, ideally we would compare those future random streams of payoffs to that of other securities with a similar rating. Crafting these counterfactuals and finding those comparison securities appears to require considerable structure and assumptions. Whether convincingly feasible or not, any such analysis would still need to be benchmarked to the actual and observed realization. That is why we view our contribution as the key step even for such an ex-ante analysis.

We establish five facts. First, the bulk of these securities was rated AAA. Second, AAA securities did ok: on average, their total cumulated losses up to 2013 are under six percent.

Table 3.12 presents more detailed results for these returns, depending on the market segment and assumptions regarding terminal value: these results are presented in greater detail in subsection 3.4.4. The most important result here may be that AAA securities provided an internal rate of return of about 2.44% to 3.31%, depending on the scenario. The yield on 10-year treasuries in 2008 was between 3 and 4 percent: the difference is surely smaller than what the standard crisis narrative seems to suggest. It mattered quite a bit, whether the mortgages were fixed rate or floating rate. Overall, though, these returns on AAA RMBS strike us as rather reasonable, and unlikely to have thrown the financial system into the abyss.

While we judge the losses on AAA securities to be modest, losses were considerably larger on securities with ratings other than AAA. We examine them in considerable detail. We find that ratings by and large did a good job in sorting securities into categories of risk, as measured by their losses ex-post, starting at a substantial increase in losses, when moving from AAA to AA securities. We find some rating reversals. For example, the accumulated value-weighted loss as a fraction of principal was higher for AA-rated securities than for securities rated C or below even as late as 2012.

Overall, however, these non-AAA-rated securities were of minor importance. The bulk received the AAA rating, so that the substantial losses on most of the non-AAA securities end up to be of comparable total magnitude to the modest losses on AAA-rated securities. Figure 3.4 shows the cumulative total losses across the universe of non-agency residential mortgage backed securities, up to 2013. The total losses amounted to no more than 350 billion dollars, with about 100 billion dollar of these losses attributed to the AAA-rated

securities. To put this in perspective, the total losses are less than 2.5 percent of US GDP or less than half a percent annualized over that period, and quite a bit less than the amount devoted to the 2009 American Recovery and Reinvestment Act or “stimulus” package. From a macro-economic perspective, it is hard to get very excited about losses of that magnitude.

It may be good to emphasize that we only examine non-agency residential mortgage backed securities. Agency-backed securities were backed implicitly by the tax payer and explicitly by programs of the Federal Reserve Bank, and therefore their role in the crisis was largely a matter of policy. We also do not investigate higher layers of leveraging and repackaging, such as, say, AAA-rated collateralized debt securities, backed by a basket of lower-rated mortgage-backed securities. Note that losses here are just redistributing the losses of the original rMBS. There are a variety of other securities that got their share of blame. None have received quite the attention of non-agency residential mortgage backed securities, though, which are the focus here.

The paper proceeds as follows. Section 3.3 discusses our unique and novel data set, and how we assembled it. Subsection 3.4 contains our analysis. Subsection 3.4.1 examines the ratings, in particular their relationship and their predictive value for future losses. Subsection 3.4.2 examines the depth, probability and distribution of losses. In subsection 3.4.3 we explore errors in rating from an ex post perspective, and the degree of rating reversals, where securities with higher ratings experienced larger losses than those with lower ones. Subsection 3.4.4 examines the resulting annualized returns on an investment at par value, under a range of assumptions on the terminal value. We revisit these return calculations, when examining investment at 2008 market value, in section 3.4.5, shedding light on the

question of underpricing and fire sales during that time. Section 3.5 concludes.

3.3 The data

We seek to investigate the market for residential non-agency mortgage-backed securities. These securities are excluded from guarantees or insurance by the government agencies “Fannie Mae” (FNMA), “Freddie Mac” (FHLMC) or “Ginnie Mae” (GNMA) due to certain characteristics, such as “jumbo loans” exceeding the limit of, say, 333700 \$ in 2004, loans on second properties such as vacation homes, insufficient documentation or borrowers with credit history problems. At the end of 2003, non-agency MBS/ABS had an outstanding amount of 842 billion \$, constituting 20% of the entire market for MBS, with agency-backed securities constituting the other 80%.

For our investigation, a major challenge was to obtain a suitable data set for these securities. The market is characterized by considerable decentralization. While the appointed trustees of a deal are responsible for providing investors with detailed information about the performance of the loans underlying the securities every month, there is no centralized repository that collects and organizes the data¹. In terms of prices, many of these securities do not trade very often, and when they do so the transactions are over-the-counter. This makes it impossible to obtain a suitable time series of transaction prices for individual deals².

As there was no readily available, organized data source, we constructed the main data

1. Some companies including Corelogic and Blackbox Logic collect and sell information and analytic tools to market participants

2. Now the Financial Regulatory Authority (FINRA) provides some summary statistics on prices and volume of daily transactions.

ourselves. We start from the Mortgage Market Statistical Annual 2013 Edition by Inside Mortgage Finance³. This publication in Volume II, Table A, Non-Agency MBS Activity, contains a complete list of the RMBS deals, completed over the years 2006-2012. For each deal, the name, the original issuer, the original amount and a few other characteristics are listed. There are a total of 2824 such deals. However, information such as cash flow or losses is not provided here. For our further data base construction, we obtain data from Bloomberg.

For each deal listed by the Mortgage Market Statistical Annual 2013 edition, we search for that deal on Bloomberg. The matching sometimes required a bit of a search, and we managed to find nearly 96 percent of the original list, by principal amount. Once we found the appropriate deal entry, we look for all deals that have similar names going forward and going back in time. Bloomberg lists the deal manager for the original deal. We then also search for all mortgage backed securities from this deal manager from 1987 onwards. Proceeding in that manner, we find a total of 8615 deals, going back to 1987 rather than just 2006, as in the Statistical Annual. In this way we hope to have minimized the number of deals that we may be leaving out. The technical appendix contains a detailed step by step description of how we built our data. Each deal generates approximately 17 separate securities or bonds on average, usually with different ratings, for a total of 143232 securities, each of which we seek to track. Their total principal amount is 5842 billion \$. Further details are in table 3.1. Table 3.2 provides an overview of the data we obtained for each security..

3. Information about this source can be found here <http://www.insidemortgagefinance.com/books/>

In this manner, we obtain as complete a universe of RMBS securities emerging from these deals, as seems possible, as well as information about their ratings and monthly cash flow and losses. We downloaded the various pieces of information, security by security, and assembled it into a spread sheet, readable by MATLAB for further analysis. The process took several month to complete, largely due to the download restrictions of Bloomberg. In order to understand our data base construction further, appendix ?? provides a sample of the information available from the Statistical Annual as well as from Bloomberg, how to read the available information and details on how we constructed our data base. A replication kit, including a more detailed description describing the construction, is available from the authors for those that seek to replicate our analysis.

Table 3.3 compares the deals in our final database with those in the Statistical Annual. Panel A provides evidence that our database contains about 94% of the deals and about 96% of the issued amount across different types of securities over the 2006-2012 period, which is the available period in the 2013 edition of the Statistical Annual. The fraction covered by our data is about the same across different market segments. Panel B shows the coverage by market segment over time to show that not only the coverage is high overall, but also that it is high consistently over time. The high matching rate for this time period, and the procedure that we followed to search for securities, give us confidence that our conclusions will not follow from having a selected sample.

We complemented this main data set with data on RMBS price indices as well as house prices. For RMBS prices we obtain the ABX.HE indexes from Markit⁴, which are built to

4. Information about these indexes and how to purchase the data is available here <https://www.markit.com/Product/ABX>

represent CDS transactions on Subprime RMBS issued in 2006 and 2007 for different credit rating levels. Finally, we use publicly available house price data at the state level from Zillow to build some of our control variables.⁵

3.3.1 Database description

Our constructed database contains information for more than 143 thousand RMBS, which were issued between 1987 and 2013 and are part of about 8,500 securitization deals. Table 3.1 shows the issuance activity over time. The table shows the boom in activity in terms of deals, bonds, market participants (issuers), and deal size from the early 2000s through 2007, and the corresponding collapse after 2008. Most of the deals after 2008 correspond to resecuritizations.

About 99% of the securities in our data, which represent 97% of the dollar principal amount, are private-label (non-agency), non-government backed,⁶ non-CDO securities. We will limit our analysis to these securities throughout the paper.

The collected information can be grouped into groups. The first group is the cash flow time series information. This constitutes the bulk of our data. Given downloading limits imposed by Bloomberg, we had to spend several months downloading this information chunk by chunk. For each security we observe the interest payments, principal payments, outstanding balance, the coupon rate and the losses each month after issuance. The second group of variables allows to identify the security and describe some of its characteristics. These

5. This data can be downloaded at <http://www.zillow.com/research/data/>

6. The government backed securities include agency securities and also non-agency securities whose underlying mortgages are backed by the Federal Housing Administration (FHA) and the U.S department of Veterans Affairs (VA)

include the Cusip ID, deal names, deal managers names, dates of issuance, coupon type and frequency, maturity date, type of tranche, notional amounts, as well as the credit rating assigned by up to 5 different credit rating agencies upon issuance. A third group of variables is related to the collateral of the securities, i.e. the underlying mortgages. These include information on the composition of the mortgages by type of rates (adjustable rates vs fixed rate mortgages), by type of occupancy (vacation home, family home, etc), by purpose of the mortgage (equity take out, refinance, purchase), or by geography (at the state level). This group of variables also includes information commonly used to assess the risk of pools of mortgages. We observe moments of the distribution of the credit scores, loan size, and loan to value ratios across the mortgage loans underlying a deal. A final group of variables include variables that can help us classify securities (for example agency vs non-agency, residential vs non-residential MBS) and commonly used metrics in mortgages backed security analysis such as weighted average maturity (WAM), weighted average coupon (WAC), and weighted average life (WAL). In the on-line appendix we list and describe all the variables in the raw data.

3.3.2 Classifying Securities into Market Segments

The most common classification used in the market and one that has determined the narratives of the crisis yields three main categories of MBS: Sub-prime, Alt-A, and Prime (or jumbo)⁷. This classification is available from the Mortgage Market Statistical Annual and

7. There are other classifications that we largely ignore. As one example, there is the Scratch & Dent category. These are loans of borrowers with the lowest FICO scores, which sometimes could have been originated outside the underwriting guidelines. These will generally fall under the sub-prime category

it is based on the classification of the bulk of the underlying loans into the same three categories.

Prime non-agency mortgages are jumbo loans that are not qualified for agency guarantees because of their size. Alt-A or Alternative A are loans in the middle of the credit spectrum, missing documentation and other characteristics. Subprime loans are loans further down the credit spectrum. The main characteristic that determines the classification of a loan is the FICO score, but a full description involves the loan-to-value ratio(LTV), loan size, and the documentation supporting the loan. An RMBS predominantly backed by subprime loans will be a Subprime RMBS.

Figure 3.1 compares that classification for deals issued after 2005 to the mean FICO scores, loan sizes and LTVs available from Bloomberg. Clearly the FICO score is key distinguishing characteristic, but it is not the only one. Figure 3.15 in the appendix further supports this claim.⁸

3.4 Analysis

We group our analysis and results into several segments. Subsection 3.4.1 examines the ratings, in particular their relationship and their predictive value for future losses. Subsection 3.4.2 examines the depth, probability and distribution of losses. In subsection 3.4.3 we explore errors in rating from an ex post perspective, and the degree of rating reversals, where securities with higher ratings experienced larger losses than those with lower ones. Subsec-

⁸. Similar pictures for loan size and LTV can be found in the technical appendix. They show that LTV and size do not provide a clear cut classification

tion ?? examines whether later RMBS have become more risky. Subsection 3.4.4 examines the resulting annualized returns on an investment at par value, under a range of assumptions on the terminal value. We revisit these return calculations, when examining investment at 2008 market value, in section 3.4.5, shedding light on the question of underpricing and fire sales during that time.

3.4.1 Ratings

Credit ratings are meant to provide guidance about the credit risk of a security. One of our goals is to assess the degree to which credit ratings were appropriate. It is useful then to take a look at what credit ratings agencies say about their ratings of RMBS. Even though different credit rating agencies present their criteria and definitions in different ways, the following four elements that come directly from agencies' documents are useful to guide and understand our analysis in section 3.4.2. First, the ratings in structured finance vehicles are long-term ratings: "Long-term ratings are assigned to issuers or obligations with an original maturity of one year or more and reflect both on the likelihood of a default on contractually promised payments and the expected financial loss suffered in the event of default" (Moody's, *Rating Symbols and Definitions*⁹). Second, comparability of ratings across asset classes or issuers is not straightforward. Credit rating agencies (at least after the crisis) strive for but

9. The document can be downloaded at <https://www.moodys.com/sites/products/AboutMoodyRatingsAttachments/MoodysRatingSymbolsandDefinitions.pdf>

do not guarantee comparability.¹⁰¹¹ Credit rating agencies have introduced additional tools such as stress scenarios to promote comparability of ratings across sectors, geographies, and over time. However, some of these tools were not in place before the 2008 crisis.

Third, the ratings are given to reflect expected losses. According to Moody's their structured finance ratings "reflect both the likelihood of a default and the expected loss suffered in the event of default. Ratings are assigned based on a rating committee's assessment of a security's expected loss rate (default probability multiplied by expected loss severity)."¹². Fourth, the ratings are meant to provide a relative measure of the credit risk and not an absolute one: "Standard & Poor's credit ratings are designed primarily to provide relative rankings among issuers and obligations of overall creditworthiness; the ratings are not measures of absolute default probability".¹³

Table 3.4 describes the credit rating activity in our database based on the assignment of a rating upon issuance. Besides the fact that a large fraction of our securities are rated, two other elements are worth noting. First, Moody's and Standard & Poor's have bigger market shares than Moody's. Apparently, investors in CDOs had a preference for CDOs

10. The following quote from Moody's in *Rating Symbols and Definitions* illustrates the point "Moody's differentiates structured finance ratings from fundamental ratings (i.e., ratings on nonfinancial corporate, financial institution, and public sector entities) on the global long-term scale by adding (sf) to all structured finance ratings. The addition of (sf) to structured finance ratings should eliminate any presumption that such ratings and fundamental ratings at the same letter grade level will behave the same. The (sf) indicator for structured finance security ratings indicates that otherwise similarly rated structured finance and fundamental securities may have different risk characteristics. Through its current methodologies, however, Moody's aspires to achieve broad expected equivalence in structured finance and fundamental rating performance when measured over a long period of time".

11. From Standard & Poor's article *General Criteria: Understanding Standard & Poor's Rating Definitions*, "Our rating symbols are intended to connote the same general level of creditworthiness for issuers and bonds in different sectors and at different times". This is available at https://www.standardandpoors.com/en_US/web/guest/article/-/view/sourceId/5435305

12. From Moody's document *Rating Symbols and Definitions*

13. From Standard & Poor's article *General Criteria: Understanding Standard & Poor's Rating Definitions*

rated by the first two firms. At the time of rating a CDO transaction, these two firms would take a conservative approach (assign a lower rating to those bonds rated by a third-party) if the CDO had purchased or used as reference RMBS bonds not rated by them. This opens the possibility that there are differences in the performance of bonds rated by different agencies via some investor-driven selection process. In our econometric analysis we will include controls for which credit rating agency rated the security whenever is required. Second, more than 62% of the securities (which represent 85% by value of principal) had at least 2 ratings.

In our analysis below we will present results by one summary credit rating, as follows. We abstract from the rating qualifiers “-” and “+”. So for example a BBB+ for us is a BBB and an A- is an A. This should not be too problematic since an A- should be closer to an A than to a BBB. Whenever a security has 2 or more ratings from different agencies we average the rating. For instance, if agency 1 rates it as AAA, and rating agency 2 as AA, and rating agency 3 as AAA, the bond will be AAA.¹⁴ For the case of two agencies, one rating a bond as AAA and one as AA, we solved the tie upwards, so the bond would be AAA. These discrepancies are not common in the data.

We can now document our *first fact*: **The great majority of non-agency RMBS securities were assigned a AAA rating upon issuance.** Almost 87% of the principal amounts and 57% of the bonds had the highest rating. Most of the other rated securities were investment grade securities (BBB or higher); only about 9% of the bonds, which represent a

14. This clearly requires a mapping of the different ratings across agencies. We used the mapping provided by the Bank of International Settlements, which is available here <http://www.bis.org/bcbs/qis/qisrating.htm>

little bit more than 1% by principal value had non-investment grade ratings upon issuance. Table 3.5 presents the numbers in detail. Moreover, the breakdown by credit rating agency is essentially the same (see the last three columns of Table 3.5).

3.4.2 Losses

In this section we analyze the losses incurred by non-agency RMBS. A loss occurs, when a scheduled payment is not made or when there is a complete default on the remaining principal and stream of payments. We observe the time series of the losses suffered month by month by each of the securities in our data. This allows us to calculate the cumulative losses at different points in time and study the differences across ratings, vintages, and market segments. In the main text, we present our results, where we weight our computations by the original principal amount of the RMBS. The technical appendix complements this with unweighted results.

Panel A of Figure 3.2 plots the cumulative losses as a fraction of original principal (weighted by this value) as of December 2013 for all RMBS in our database. Figure 3.3 provides the same plot for AAA securities alone, broken down by their three categories, prime, Alt-A and subprime. These figures together with table 3.6 document our ***second fact: AAA securities did ok: on average, their total cumulated losses up to 2013 are under six percent.*** They furthermore document our ***third fact: the subprime AAA-rated RMBS did particularly well.***

This fact requires some discussion. The fact that AAA securities had lower losses than the other securities does not make them a AAA security. It is possible that the loss rate

does not correspond to what a AAA security should exhibit. Here we find difficulties with some of what rating agencies say about their ratings. First, they are meant to provide a relative ranking, they are not to be associated with specific probability of default levels, and the ratings for structured finance securities may capture different credit risk than ratings for other securities. Moreover, almost mechanically, an average AAA security will exhibit lower losses than a lower-rated security since the AAA tranches are the last to face losses in the capital structure. Therefore, a pure relative comparison through a ranking will not be enough to say that a AAA security actually behaved as such. We therefore, think that the second fact must involve a comparison with other AAA securities, and a relative comparison across RMBS with different ratings that should be not only in terms of ranking, but also in terms of level differences in their loss rates. In terms of the relative comparison, the figure shows that the long-run loss rates of AAA-rated RMBS as of December 2013 were several orders of magnitude lower than those of other investment-grade RMBS.

These cumulative losses of 2.3% of principal may be considered large for securities with a AAA rating. Indeed, judging by a table¹⁵ provided by Moody in 2001 of ideal expected losses by credit rating over time, losses of this size should be expected for BBB securities and not for securities rated AAA. Put differently, losses of this size should be unlikely for AAA rated securities. Then again, the financial crisis of 2008 should be considered an unlikely disaster. Compared to the magnitude of this crisis, we consider the cumulated losses of 2.3% over the six years 2008 to 2013 (both included) as quite modest indeed. Figure 3.4 shows the

15. The table can be found in a document available at http://siteresources.worldbank.org/EXTECAREGTOPPRVSECDEV/Resources/570954-1211578683837/Bielecki_Moodys_Rating_SME_transactions.pdf

distribution of the corresponding total cumulated losses on all non-agency RMBS securities. The total for AAA-rated securities is around 100 billion US \$ over the course of six years, due to the large underlying principal. While not negligible, it would be a challenge to blame the magnitude of the financial crisis on these losses, as often seems to be done in the standard narratives. Even taking into account the total losses of 323 billion US \$ over the course of six years seem to us a magnitude that would need to take considerable amplification and addition by other factors to bring down the entire financial system.

We can break down the analysis by market segment defined by loan type (Prime, Alt-A, and Subprime). Panel A of Figure 3.2, figure 3.3, table 3.6 or the additional vintage-level detail provided by table 3.7 document the ***third fact: the subprime AAA-rated RMBS did particularly well.*** AAA-rated Subprime Mortgage Backed Securities were the safest securities among the non-agency RMBS market. As of December 2013 the principal-weighted loss rates AAA-rated subprime securities were on average 0.42% We consider this to be a surprising fact given the usual narrative for the causes of the financial crisis and its assignment of the considerable blame to the subprime market and its mortgage-backed securities. An example of this narrative is provided by Gelman and Loken (2014)¹⁶: “We have in mind an analogy with the notorious AAA-class bonds created during the mid-2000s that led to the subprime mortgage crisis. Lower-quality mortgages —that is, mortgages with high probability of default and, thus, high uncertainty—were packaged and transformed into financial instruments that were (in retrospect, falsely) characterized as low risk”.

16. We have chosen this quote because it is quite representative of the common narrative during the crisis and useful for our purposes. We have not chosen it as a critique of the article by Gelman and Loken (2014), whose subject of interest is not the RMBS market per se

Table 3.7 shows that the bad performance of AAA securities is primarily due to the Alt-A segment of the market, which in principle was backed by mortgage loans of borrowers with better credit risk prospects. The fact that Prime and Alt-A RMBS exhibit loss rates comparable to those of Subprime securities for ratings below AAA, and worse for the highest rated securities is evidence of the lower screening effort exerted by financial institutions that found easier to securitized and sell loans of higher-quality borrowers, as documented by Keys et al. (2009, 2010, 2012). For completeness, appendix figure 3.18 presents the incidence of losses by market segment, again with a relatively strong performance of AAA-rated subprime bonds in comparison with similar rated Prime and Alt-A bonds, and with a comparable (poor) performance for the other ratings.

Figure 3.2 and 3.3 reveal that losses in AAA securities picked up in 2011. This may seem late, but it is the result of the credit support protection enjoyed by AAA securities. These losses, however, are to a large degree predictable once delinquencies in the underlying loans pick up. If delinquencies are large enough to predict that lower tranches will be wiped out, then one can predict AAA securities facing losses. Since many lower-rated tranches have faced big losses, we should expect AAA securities to keep accumulating some losses beyond the final period in our plot. However, as of December 2013, AAA securities taken together still had \$341 billion of cushion coming from lower-rated bonds. Given that all the losses over the last 6 years do not reach this amount and given the recovery of the US economy, we do not believe our conclusions are going to change dramatically.

Panel B of figure 3.2 plots the fraction of securities facing losses. By this measure, we do see a relatively large incidence of losses, even in AAA securities, even though the size of the

losses is small relative to the principal amount. We can then ask, what does the distribution of losses look like across securities? Figure 3.5 shows that losses tend to be concentrated in a few securities; the outcome is almost binary, either the losses as a fraction of principal are low, or the security gets wiped out. With this, we have documented our ***fourth fact: the bulk of the losses were concentrated on a small share of all securities.***

Table 3.8 shows the losses as a fraction of principal as of 2013 for the different credit ratings and for different groups of vintages. Table 3.7 contains additional detail, broken down by type. While the average losses on AAA rated securities were fairly modest, regardless of the type and the vintage, later vintages did worse. Table 3.8 shows that principal-weighted losses on AAA rated securities issued before 2003 were less than 0.03% and even were below 0.4% for those issued from 2003 to 2005, but that this went up to a considerable 4.8% on securities issued from 2006 to 2008. Considerably larger rises in loss rates by vintages occur for lower-rated securities. For example, while A-rated securities issued before 2003 lost less than 0.6%, dramatic losses of nearly 66% occurred for the 2006-2008 vintage. With that we obtain our ***fifth fact: later vintages did worse than earlier vintages.***

Can we therefore conclude that ratings deteriorated over time and that rating agencies became more generous? This certainly has been a theme in much of the narrative of the crisis. The deterioration in performance could also have been due to bad luck, though. Consider a security issued long before the peak of the house price boom, and compare it to an otherwise identical security issued just at the peak. The former security is less likely to be subject to losses, since the 2013 value of the underlying home relative to the original

purchase price is higher for the former compared to the latter. If one views the arrival of the house price decline as a random event, unrelated to current level of house prices, one could argue that the resulting higher losses for the later vintages were just a stroke of bad luck, and not the result of a more liberal rating. While obviously a very benign interpretation, it is useful to check, how far the house price movements at the time of issuance can explain the subsequent loss performance. Similar remarks may apply to other covariates. We therefore run the cross-sectional regressions

$$l_i = \alpha + y + \beta X_i + \epsilon_i \quad (1)$$

$$dl_i = \alpha + y + \beta X_i + \epsilon_i \quad (2)$$

where l_i is the loss as a fraction of principal as of December 2013 for RMBS i , dl_i is a dummy that takes the value 1 if the the cumulative losses of RMBS i are strictly greater than zero as of December 2013, y is a set of vintage or issue-year fixed effects, and X_i are our covariates. We also run these cross-sectional regressions with dl_i as the dependent variable, where dl_i is a dummy that takes the value 1 if the the cumulative losses of RMBS i are strictly greater than zero as of December 2013. The covariates are the amount of credit support, credit ratings, moments of the distribution of the characteristics of the underlying mortgages (LTV, size, and FICO scores), the purpose of the mortgages (purchase, refinancing, equity takeout), the type of mortgage as it relates to the interest rate (Fixed-rate vs Adjustable Mortgage Rates), geography fixed effects¹⁷, and house price appreciation.

17. We use information on the concentration of mortgages backing the pool by state to include geography controls

Figure 3.6 presents the results for losses across all securities and by credit rating, where the regressions are weighted by the original principal amount. The coefficient corresponding to the 2001 vintage fixed effect is normalized to zero. Figure 3.7 plots the vintage fixed effects for Prime, Alt-A, and Subprime RMBS overall and by the credit rating assigned upon issuance. Comparing these results to table 3.8 shows that the inclusion of the covariates does little to explain the performance deterioration of the later vintages. Judging by these results, we therefore cautiously endorse the view that rating standards have deteriorated in the run-up to the crisis. Moreover, these results are consistent with the findings of Adelino et al. (2015), who argue that middle income borrowers had an increasing relative role in mortgage delinquencies and defaults in the run-up to the crisis. These results are also consistent with the idea that securitization contributed to bad lending by reducing incentives of lenders to carefully screen borrowers, and that lower screening standards happened for relatively high FICO scores as those loans were easier to securitize as argued by Keys et al. (2010).

It is instructive to investigate the relationship of the house price boom and busts, shown state-by-state in figure 3.8, and its relationship to the losses of the RMBS. For each security in our data set, we know the top five states in terms of the locations of the underlying mortgages, and the fraction of the total principal invested there. We thus run a regression of losses on these fractions multiplied with a state dummy, as well as possibly a set of covariates. For the covariates, we chose dummy variables for the credit rates as well as mortgage type (Prime, Alt-A and Subprime). More precisely, we use that classification of the Statistical Annual and match it to cut-offs in terms of FICO scores, as indicated in figure 3.15. We then use the FICO scores available from Bloomberg, to classify the securities into Prime, Alt-A

and Subprime for all other deals, since Bloomberg does not list that categorization.

The regression coefficient on each fraction times state dummy then provides an estimate of the losses that would have occurred, if an RMBS had invested in that state only. This exercise results in the maps shown in figure 3.9, sorting the loss state dummies again into quintiles, with and without controls. These figures may best be compared to the bottom map of figure 3.8. There are similarities, such as a fairly benign region from North Dakota to Texas, or the darker regions around Nevada.

Table 3.9 estimates that relationship more formally. There, we have calculated a linear regression of the cumulative loss as a fraction of initial principal on the change in house prices, both during the run-up phase from 2000-2006 as well as the crash-phase from 2006-2009. To find the house price change relevant for each RMBS, we have averaged the house price changes over the five top states in which that security was invested, using the relative investment shares to calculate these averages. Our preferred specifications are in columns (3) and (5). There, we find that the increase in house prices decreased losses, but that the subsequent decrease in house prices increased losses for the security. According to column (5), say, an additional increase of house prices from 2000 to 2006 decreased losses by 0.18 percent of principal, while an additional decline of house prices from 2006 to 2009 by one percent increased losses by 0.53 percent. Column (3) provides a rather similar answer. If only the price increase is included or if state dummies are included, with the weights given by the investment shares, these effects (rather naturally) disappear.

While we argue that the losses at least on the AAA securities have been modest, they certainly are larger than what should be expected for a typical AAA-rated security. Accord-

ing to a 2001 table by Moody's concerning the idealized expected losses over a period of 10 years by credit rating, AAA securities should be expected to have losses of 0.0055%, AA in a range from 0.055% to 0.22%, A is a range from 0.385% to 1%, and BBB in a range from 1.43% to 3.36%. Comparing these numbers to the implied levels from Table 3.8, we see that only securities issued before 2003 and perhaps also AAA-rated securities issued in 2003-2005 had losses roughly commensurate with these table values.

To summarize, even though AAA-rated RMBS bonds suffered some losses, the losses were relatively small. To a large degree, RMBS securities rated AAA and issued until 2005 behaved like within expectations for AAA securities, even in an environment of extreme stress. In particular, the much-maligned subprime segment did particularly well. The securitization and rating process did fail perhaps for non-AAA investment grade bonds: these ratings were, in retrospect, too high, and in particular so for the 2006-2008 vintage. However, it is the AAA securities and their role as safe assets for pension funds and other institutions that have received the bulk of the attention: they also constituted the bulk of the market.

3.4.3 Credit Rating Rankings and Reversals

In this section we explore the credit rating and their rankings, using their ex-post performance in relationship to ex-ante ratings, covariates available at the time of rating and corresponding expectations. With this, we seek to address two questions. How close to appropriate were the ratings, in hindsight? And: even if losses turned out to be higher than expected, did the rating agencies do a good job in ranking securities appropriately?

We present two exercises. In the **first exercise**, we compare the average realized losses of

securities to Moody's expected losses by rating. Moody's has published a table of "Idealized Cumulative Expected Loss Rate" which we present as reference in the appendix in Figure 3.19.¹⁸ For example, in 10 years a BBB- security would be expected to have a loss rate of approximately 3.35%. Each of our securities had an initial (ex-ante) credit rating and therefore an expected loss rate based on Moody's Table. For each security, we assign an ex-post rating based on its actual realized loss rate. So, if a given security had a realized loss rate between the AAA and the AA expected loss rate on Moody's table (between 0.0055% and 0.22% in 10 years), the security receives an ex-post rating of AA. Then we compare the ex-ante rating with the ex-post rating. Figure 3.10 presents the results. The solid line is the fraction of securities by original rating (ex-ante). The dotted line is the fraction of securities that would have received a given rating based on their realized loss rates. According to the figure, on average, AAA ratings were not given in excess, but too many securities were given investment grade ratings. This figure, however, hides the fact that having on average the right number of AAA securities does not mean that AAA securities were given the right rating as many ex-ante AAA could be ex-post AA while many ex-ante AA can be ex-post AAA. Figure 3.11 addresses the issue by showing the fraction of securities for which ex-ante and ex-post ratings coincide (labeled as Correct Rating), those for which the ex-post rating is higher than the ex-ante rating (labeled as Deflated Rating), and those for which the ex-post rating is lower than the ex-ante rating (labeled as Inflated Rating). While about 75% of AAA securities received a rating that was in line with the expected losses, AA securities were

18. The table is originally available here https://www.moody.com/sites/products/productattachments/marvel_user_guide1.pdf or here http://siteresources.worldbank.org/EXTECAREGTOPPRVSECDEV/Resources/570954-1211578683837/Bielecki_Moodys_Rating_SME_transactions.pdf

almost entirely wrongly rated, and securities with ratings A and below had a large fraction of inflated ratings. Figure 3.12 provides a three-dimensional overview of the given ex-ante rating and the ex-post rating, based on our exercise. Ideally, all the mass would accumulate along the diagonal. The off-diagonal mass gives a visual impression of “rating gone wrong”. The story of this picture is fairly clear. Many of the AAA-rated securities deserved their ratings, though a good portion should have been rated lower, in hindsight. For all the other securities, we find a bimodal result. The smaller share had no losses at all, justifying a AAA rating ex post. The larger share had substantial losses, that would have required a rating below CCC. There hardly is any mass on the diagonal, except for AAA ratings.

In the **second exercise**, we wish to understand, whether the ratings could have been improved upon at the time, aside from the overall extent of the losses examined above. We seek to calculate the extent to which the inclusion of additional covariates X , available at the time of rating, for a higher-ranked security predicts larger loss probabilities than observed on average for lower-ranked securities. We call this a ratings reversal.

More precisely, for $\alpha \in [0, 1]$ as well as for each rating, say AAA, we first seek to estimate $P(Loss > \alpha \mid AAA)$ and $P(Loss > \alpha \mid AAA, X)$, given the crisis of 2008. For the former, we estimate this probability with the fraction of AAA-securities, whose losses exceeded α at the end of 2013. For the latter and for each security i rated AAA and with covariates X_i , construct the observation

$$Y_i = 1_{Loss_i > \alpha}$$

indicating, whether the losses for security i exceeded α or not. We then estimate a linear

probability model¹⁹, per linear regressing these observations Y_i on the covariates X_i .

Note that we estimate this probability model separately for each α . One could, instead, formulate a model for losses and infer the probability of losses exceeding any given treshold α from that model and the observed covariates.

We then seek to estimate the gain from including covariates as well as the probability of ratings reversals. For each α , we define the gain from including covariates X compared to the raw difference between securities rated AAA and AA as

$$Gain_{AAA,AA}(\alpha) = \frac{E [|P(Loss > \alpha | AAA, X) - P(Loss > \alpha | AAA)|]}{P(Loss > \alpha | AAA) - P(Loss > \alpha | AA)} \quad (3)$$

where the outer expectation is taking an expectation over the random covariates X . We estimate the numerator by the sample average of $\hat{P}(Loss > \alpha | AAA, X_i) - \hat{P}(Loss > \alpha | AAA)$ for all AAA-rated securities i and the probability estimators explained above.

We likewise define $Gain_{AAA,A}(\alpha)$, $Gain_{AA,A}(\alpha)$, etc.. We define the probability of rating reversals for AAA-rated securities to AA securities as

$$Reversal_{AAA,AA}(\alpha) = P(P(Loss > \alpha | AAA, X) > P(Loss > \alpha | AA))$$

where the outer probability is likewise taken as an expectation over the random covariates X .

We estimate $Reversal(\alpha)$ by calculating the fraction of all AAA-rated securities i , for which $\hat{P}(Loss > \alpha | AAA, X_i)$ exceeds $\hat{P}(Loss > \alpha | AA)$, with $\hat{P}(\cdot)$ denoting the estimator of $P(\cdot)$

19. A linear probability model was computationally easier than a probit model, while yielding essentially the same results, when we cross-checked.

explained above. We likewise construct estimators for $Reversal_{AAA,A}(\alpha)$, $Reversal_{AA,A}(\alpha)$, etc..

As covariates, we made use of (essentially) all the additional information on these securities provided by Bloomberg. They are the dollar value of the principal of the security, the average FICO score, the mean Loan-To-Value ratio, the mean loan size, the original credit support, the weighted average coupon, the weighted average life, the fraction of loans with adjustable mortgage rates, the fraction of loans for single family homes, the fraction of loans for condos, the fraction of loans for first purchases, the fraction of loans for refinancing, a dummy for floating coupon securities, a dummy for floating rate securities, a set of dummy variables for the top 5 states represented in the underlying mortgage loans, a set of dummy variables for the credit rating agencies involved in rating each security.

Panel A of Table 3.10 reports estimates of the gains given by equation 3, for all the pairwise comparisons between a given rating and ratings below it for investment grade RMBS. We see that covariates did carry information that would have been useful to predict losses, and to assign ratings, particularly for the AA, A, and BBB ratings. For AAA ratings, only for low values of alpha we see some gains from the covariates. The estimates of rating reversals is reported in panel B of Table 3.10. It turns out that the value of α matters considerably. If $\alpha = 0$, then we find a 40 percent probability of rating reversal. To understand, consider the probability of the occurrence of any loss, as shown in table 3.11. It turns out that AAA securities were actually somewhat more likely to incur losses than AA securities: the overall fractions are 28 percent versus 16 percent. We know already, however, that losses on AAA securities are typically small, if they occur at all. Figure 3.5 shows that the distribution

for AAA securities puts more weight on small losses compared to the distribution for other investment grade securities. Thus, as α is increased to, say, 10%, we find a rating reversal probability of only 3%.

Given the financial crisis, a threshold of $\alpha = 0$ may then be too stringent to judge the appropriateness of the rankings. Overall we judge that the rating agencies got the rankings about right. This conclusion comes with a number of caveats, of course. First, the construction of the securities often implies mechanically, that lower-ranked securities will be hit with losses before that happens to higher-ranked securities. The ranking of securities for any given deal is therefore very unlikely to be incorrect (assuming that rating agencies did indeed check the loss sequencing): the comparison here is more interesting regarding the consistency of rankings for securities across deals. Second, all our inference is conditional on the crisis of 2008: this is the only set of observations we got. We obviously cannot infer anything here about the appropriateness of the ratings or their rankings across all potential futures from 2007 on forward. Finally, we have used the realized losses to estimate the weight on information available a priori, in order to check for rating reversals. Obviously, the rating agencies did not have that information at hand at the time when they had to give their assessments.

3.4.4 Returns

One of our goals in this paper is to assess the performance of RMBS as investments. In this section, we investigate the performance for a buy-and-hold investor, who purchased the RMBS at par at issuance, and plans to hold them until all payments have matured. Only

investment-grade securities were sold at par: so we focus on these.

We have calculated these returns in two ways. One is to calculate the internal rate of return. This is the rate r that solves net present value equation

$$P_0 = \sum_{t=1}^T \frac{i_t + p_t}{(1+r)^t} + \frac{TV_T}{(1+r)^T} \quad (4)$$

where P_0 is the initial value of the security, i_t is the monthly cash flow corresponding to interest payments, p_t is the monthly cash flow corresponding to principal paydown, and TV_T is the terminal value at some date T . The other approach and our preferred approach is to calculate the discount margin θ over a benchmark interest rate r_t . The discount margin can be interpreted as the average return above the benchmark. For the benchmark, we are using the 3-month Treasury Bill. The discount margin is the value for θ that solves

$$P_0 = \sum_{t=1}^T \frac{i_t + p_t}{\left(1 + r_t^{t_bill} + \theta\right)^t} + \frac{TV_T}{\left(1 + r_t^{t_bill} + \theta\right)^T} \quad (5)$$

We will present results mainly for the latter: results for the internal rate of return provide rather similar insights. Note that we do not take into account risk prices or term premia in either calculation.

We set T to be December 2013, given our data set. We observe the payments i_t as well as p_t , but we do need to make assumptions regarding the terminal value TV_T . The natural candidate for the terminal value is the outstanding principal balance at time T , which is part of the monthly information that we have for each security. To that end, it is important

to understand how past losses affect the outstanding principal value in the data. In a typical prospectus for an RMBS one can find the explanation: realized losses are applied to reduce the principal amount and “if a loss has been allocated to reduce the principal amount of your class of certificates, you will receive no payment in respect of that reduction.” From this we conclude that the principal balance recorded in the data at date T already incorporates losses on principal that have occurred previously rather than leaving them on the book. However it is possible that there needs to be some additional discounting of the outstanding principal value, because additional losses may be expected in the future. We therefore examine six different scenarios regarding the terminal value. The first three assume that all securities are valued at 80%, 90% and 100% of the principal outstanding as of December 2013.

For the fourth scenario, we assume that each security trades at a loss equal to the loss rate it has suffered up to that point. For the fifth, we assume that each security trades at a loss equal to the mean loss rate of the securities with the same original credit rating and same vintage. The sixth is similar to the fifth, except for using the median loss rate rather than the mean.

To provide a perspective for the first three scenarios, we consulted information provided by FINRA for the month of December 2013, see figure 3.14. In 2009, the Financial Industry Regulatory Authority (FINRA) made a proposal to collect data for ABS, CDO, and MBS securities.²⁰ Now daily reports going back to May 2011 with the number of transactions, trade volume, and statistics on transaction prices are publicly available.²¹ From these

20. <https://www.finra.org/newsroom/2009/finra-proposes-expanding-trace-reporting-asset-backed-securiti>

21. Reports are available and can be downloaded at <http://tps.finra.org/idc-index.html>

reports one can see that, as of December 2013, investment-grade securities were mostly trading with prices above 90, and non-investment grade with prices above 75 and generally above 80. We therefore consider our range from 80 to 100 percent to be reasonable.

Table 3.12 presents results for the realized internal rate of return calculations, for the first three scenarios regarding the terminal value. The most important result here may be that AAA securities provided an internal rate of return of about 2.44% to 3.31%, depending on the scenario. It mattered quite a bit, whether the mortgages were fixed rate or floating rate. For fixed rate mortgages, AAA securities returned between 3.6 and 4.8 percent, depending on the market segment and assumptions regarding the terminal value. For floating rate mortgages, AAA securities returned between 0.4 and 3.8 percent. Overall, though, these returns on AAA RMBS strike us as rather reasonable, and unlikely to have thrown the financial system into the abyss.

Tables 3.13, 3.14, 3.15, and 3.16 present the results of our discount margin computations. In general we see that under the different assumptions the computed rates of return are very stable, and change in the expected direction when changing the terminal value assumption from 80%, to 90% and to 100%, and change in a way that is consistent with the results and analysis of the losses that we have already presented in section 3.4.2 for the other assumptions. For AAA securities which on average presented small loss rates the returns are quite stable across assumptions and also for securities older vintages as they have more time to produce cash flows and pay down the principal.

The distribution of returns is left skewed, with a very long left tail, as evidenced by the mean quite far to the left of the median and large standard deviation. This is another version

of our fourth fact, that losses were concentrated on a small share of all securities.

We find that 65% of the AAA securities, under the assumption of a terminal value equal to 80% of the outstanding balance would yield a return equal or higher than the average Libor rates over this time period.

In Table 3.15 we break down the return by market segment as defined the mortgage loan type. These results show about a 2 percentage point realized premium of Prime over Subprime securities. This may be surprising at first given that we showed that losses in subprime securities were not particularly worse than in other segments and for AAA were actually lower. One reason behind this is the fact that the fraction of floating rate bonds (almost 90%) in the subprime segment was higher than the fraction of floating rate bonds in the Alt-A (about 62%) and Prime (about 46%) segments. In a period of low interest rates like the one we consider, floating rate bonds did worse than fixed rate bonds. To show this and obtain more comparable measures, we break down the computations for AAA bonds by vintage, by mortgage loan type and also by the type of bond according to the type of rate it would pay (fix or floating). We present weighted median rates in Table 3.16 along with the US 3-month Libor rate as benchmark. We see that across vintages and loan types, fixed rate AAA securities provided returns between 2 and 5 percentage points higher than the Libor. Floating rate bonds yielded lower returns, always higher than the Libor for Prime securities, and generally larger than the Libor but smaller than Prime returns for Alt-A and Subprime (except for vintages 2006 and 2007).

3.4.5 *Prices*

While we examined the returns to purchasing RMBS at issuance and holding them to maturity in the previous section, one may also wish to examine the returns when purchasing them at market prices, in particular at the height of the financial crisis of 2008. While individual price series for the RMBS do not seem to exist, time series for indices are available. On January 16th 2006, Markit launched a series of asset-backed credit default swap indexes on US home equity Asset Backed Securities. The indexes are tradable synthetic derivatives, which reference 20 subprime RMBS deals/bonds. There are four series of indexes, each of which corresponds to a different vintage of securities: series 06-01 references deals issued between June 2005 and January 2006, series 06-02 references deals issued between January 2006 and June 2006, series 07-01 references deals issued between June 2006 and January 2007, and series 07-02 references deals issued between January 2007 and June 2007. The 20 deals used in each series are determined at the inception of the index and they never change. These deals are selected among a list of fifty deals of the 25 largest issuers (2 deals per issuer) and they must meet the following main criteria: the deal size must be at least \$500 million, it must have tranches with all of 5 ratings (AAA, AA, A, BBB, and BBB-), it must have been rated by both S&P and Moody's, each tranche must have a weighted average life of at least 4 years with the AAA having a minimum of 5 years, the weighted average FICO credit score of the obligors on the assets backing the securities issued in the RMBS transaction must not exceed 660 as of its issuance date, and each Required Tranche must bear interest at a floating rate, with the base rate being one month LIBOR.

Each series corresponds to a set of 6 indexes by credit rating, from AAA to BBB-. In any transaction involving the index there is a protection buyer and a protection seller. The protection buyer makes two types of payments to the protection seller. One is a one-time payment upfront computed as the difference between par value and the index value multiplied by the notional amount²² The second type of payment is a coupon or spread, payable monthly. This coupon is fixed for a given index. For example, for the AAA.06-1 index, the coupon is fixed at 18 basis points per annum, while for the BBB.07-01 is 224 basis points. The protection buyer receives payments from the protection seller in the event of interest shortfalls, principal shortfalls, and writedowns. Based on the value (or “price”) of the index, one can compute excess returns between t and $t + 1$ as

$$r^e = \frac{price_{t+1} - price_t}{price_t} + spread \times \frac{day_count}{360}$$

One can think of this as an excess return over some risk-free benchmark (for example Libor) by thinking about the funded transaction: the protection seller faces the risk up to 100% of the notional, if after entering the transaction he sets aside this amount at risk, he will earn some benchmark rate like the libor. The running coupon and the index appreciation will make up the rest of the return.

Figure 3.13 plots the prices of indexes of the different series (vintages) by credit rating. Only the AAA tranches and the AA tranche of the 2006-01 vintage recover strongly after the crisis. The corresponding monthly returns computed with the formula above are shown in

22. the notional amount is adjusted by the so-called factor, which refers to the outstanding principal amount of the underlying bonds.

Table 3.17. What these numbers as well as figure 3.13 reveal is that substantial returns were earned by investors purchasing these securities in May 2009 and holding them to December 2013. On the other hand, excess returns were near zero or negative, when the purchase was made in June 2007. The excess returns on AAA rated securities was substantially negative for anyone purchasing them in, say, June 2007 and selling them in May 2009, implying losses of up to 6 percent per month, while for any long term investor, who held out until December 13, we obtain the still modest losses of 0.41 percent on a monthly basis as the worst of the AAA securities, i.e. the vintage 2007-2. Securities further down in the rating scale performed considerably worse. These results are generally in line with our findings of the payoff streams in the other sections.

3.5 Discussion and Conclusion

We have examined the payoff performance, up to the end of 2013, of non-agency residential mortgage-backed securities (RMBS), issued up to 2008. For our analysis, we have created a new and detailed data set on the universe of non-agency residential mortgage backed securities, per carefully assembling source data from Bloomberg and other sources. We have compared these payoffs to their ex-ante ratings as well as other characteristics. We have established five facts. First, the bulk of these securities was rated AAA. Second, AAA securities did ok: on average, their total cumulated losses up to 2013 are under six percent. Table 3.12 presents more detailed results regarding their returns, depending on the market segment and assumptions regarding terminal value. The most important result here may be that AAA securities provided an internal rate of return of about 2.44% to 3.31%, depending

on the scenario. The yield on 10-year treasuries in 2008 was between 3 and 4 percent: the difference is surely smaller than what the standard crisis narrative seems to suggest. It mattered quite a bit, whether the mortgages were fixed rate or floating rate. Overall, though, these returns on AAA RMBS strike us as rather reasonable, and unlikely to have thrown the financial system into the abyss. Third, the subprime AAA-rated RMBS did particularly well. Forth, the bulk of the losses were concentrated on a small share of all securities. Fifth, later vintages did worse than earlier vintages. Together, these facts call into question the conventional narrative, that improper ratings of RMBS were a major factor in the financial crisis of 2008.

3.6 Figures and Tables

Table 3.1: RMBS Database: Deals, Securities, Nominal Amounts by Year of Issuance

This table reports some figures that describe the size of our database of Residential Mortgage Backed Securities by year of issuance. All the information comes from Bloomberg.

Year	No. Deal Managers	No. Deals	No. MBS	Notional (\$ Billion)	Average Deal Size (\$ Million)
1987 - 1999	35	858	9,462	244.2	284.6
2000	20	227	2,724	93.8	413.2
2001	23	397	5,815	179.9	453.1
2002	30	574	8,255	314.0	547.1
2003	30	788	12,420	475.1	603.0
2004	30	1,106	15,787	723.4	654.1
2005	29	1,361	22,017	1,005.2	738.5
2006	39	1,563	27,184	1,237.4	791.7
2007	35	1,027	19,143	936.1	911.5
2008	20	108	1,541	103.3	956.4
2009	17	151	5,660	170.6	1,129.9
2010	17	135	6,089	155.9	1,154.5
2011	13	101	3,182	68.3	676.5
2012	11	92	1,789	36.5	396.9
2013	13	127	2,164	98.7	776.9
All Years	83	8,615	143,232	5,842.3	678.2

Table 3.2: Database Variables

<i>Security Identification</i>	<i>Credit Rating</i>
Cusip ID	Current and Original Ratings (5 agencies)
Deal Name	<i>Other Security Characteristics</i>
Deal Manager	Credit Support at Issuance
Issuer Company	Original Principal Amount
<i>Security Classification</i>	<i>Collateral Description</i>
Deal Type (eg. CMBS, RMBS)	Mortgage Purpose (% Equity Takeout, Refinance)
Collateral Type (eg. Home, Auto, Student)	LTV Distribution (min, max, mean, 25th, 50th,75th)
Collateral Type (eg. ARM vs FRM)	Credit Score Distribution (min, max, mean, 25th, 50th,75th)
Agency Backed (yes, no)	Mortgage Size Distribution (min, max, mean, 25th, 50th,75th)
Agency (eg. Fannie Mae, Freddie Mac)	MBS metrics 1: Weighted Average Coupon
<i>Dates</i>	MBS metrics 2: Weighted Average Life
Issue Date	MBS metrics 3: Weighted Average Maturity
Pricing Date	Fraction of ARM and FRM
Maturity Date	Occupancy (% of Owner, Investment, Vacation)
<i>Security Description</i>	<i>Geographic Information</i>
Bond type (e.g. Floater, Pass-through, Interest Only)	Fraction of mortgages in top 5 states
Tranche Subordination Description	<i>Cash Flow and Losses</i>
Coupon Type (e.g. Fixed, Floating)	Monthly Interest and Principal Payment
Coupon Frequency (e.g. Monthly)	Monthly Outstanding balance
Coupon Index Rate (e.g. 3M-libor)	Monthly Losses

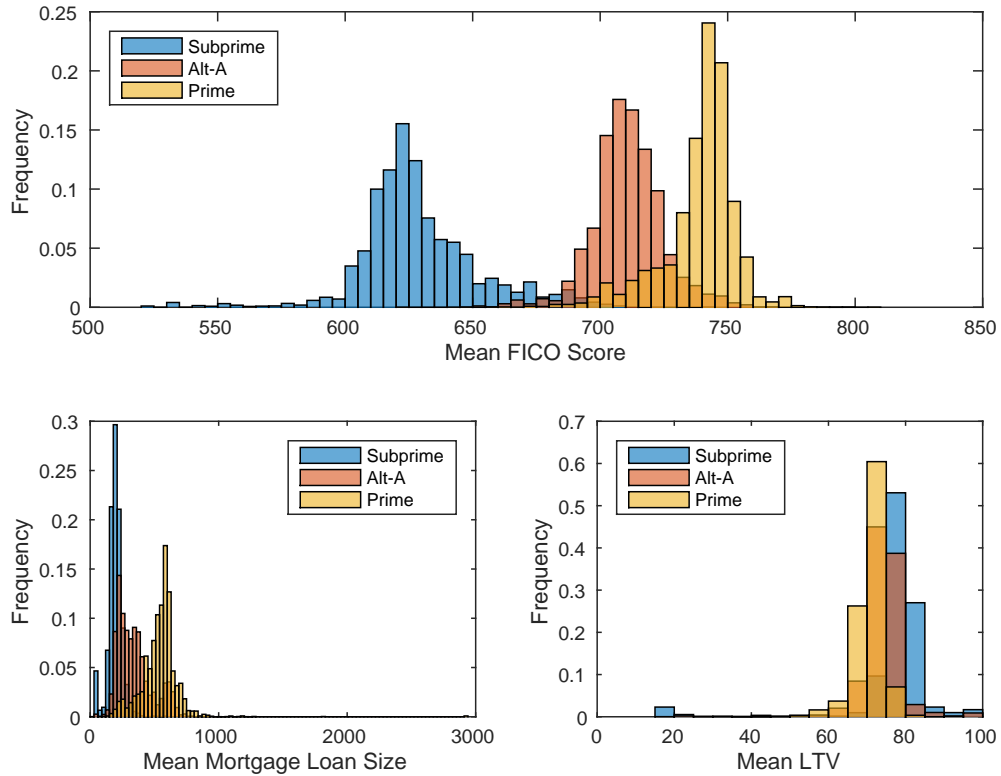
Table 3.3: Database Coverage of the universe of Non-Agency RMBS

This table compares our database to the universe of mortgage-backed securities for 2006 to 2012, as listed in the Mortgage Market Statistical Annual 2013 Edition.

Panel A: Principal Amount and Deals Coverage by Type of Mortgage-backed Security										
	Prime		Alt-A		Subprime		Other		Total	
	Amount	Deals	Amount	Deals	Amount	Deals	Amount	Deals	Amount	Deals
Unmatched	9.8	15	16.2	30	33.6	52	26.3	90	85.9	187
Matched	406.5	484	596.5	739	656.4	771	304.4	643	1,963.8	2,637
Total	416.3	499	612.7	769	690.0	823	330.7	733	2,049.8	2,824
Pct. Matched	97.7	97.0	97.4	96.1	95.1	93.7	92.0	87.7	95.8	93.4

Panel B: Principal Amount Coverage by Year and Type of Mortgage-backed Security									
Type of MBS	Year								
	2006	2007	2008	2009	2010	2011	2012	All Years	
All									
Principal Amount	1,129.3	697.0	58.6	60.4	63.8	27.6	13.2	2,049.8	
Pct. Matched	97.3	93.4	95.2	95.0	97.7	97.3	91.6	95.8	
Prime									
Principal Amount	218.79	180.4	7.0	5.5	0.5	0.7	3.5	416.3	
Pct. Matched	96.9	98.6	94.4	100.0	100.0	100.0	100.0	97.7	
Alt-A									
Principal Amount	362.79	247.4	1.9	-	0.7	-	-	612.7	
Pct. Matched	99.1	94.9	72.5	100.0	100.0	100.0	100.0	97.4	
Subprime									
Value	470.11	213.4	2.4	0.9	-	0.5	2.8	690.0	
Principal Amount	96.6	91.8	97.9	100.0	100.0	100.0	100.0	95.1	
Other									
Value	77.57	55.8	47.4	54.0	62.6	26.4	6.9	330.7	
Principal Amount	93.8	76.0	100.0	100	100	100	100	92.0	

Figure 3.1: Distribution of Mean FICO Score, Loan Size, and LTV by Type of Loan



The types of mortgages in the figure are Prime, Alt-A and Subprime. As there is no flag in our data that allows us to classify securities by type of mortgages, we rely on the classification provided in the Mortgage Market Statistical Annual 2013 Edition. Since we only have information in the Statistical Annual for the period 2006-2012, only securities issued in those years are included.

Table 3.4: **Credit Rating Activity of Non-Agency Residential Mortgage Backed Securities: 1987-2013**

This table presents some figures about the credit rating activity in the RMBS market between 1987 and 2013.

Rating Activity	MBS Bonds		Principal Amount	
	No.	Pct.	(\$ Billion)	Pct.
Rated by at least one agency	115,282	81.4	5,214.1	92.0
Rated by 2 or more agencies	87,937	62.1	4,812.3	84.9
Rated by all 3 big agencies	16,324	11.5	1,085.0	19.1
Rated by all agencies	0	0.0	0	0.0
Rated by Standard & Poors	90,006	63.6	4,518.0	79.7
Rated by Moody's	67,036	47.3	3,931.5	69.4
Rated by Fitch	58,692	41.4	2,530.6	44.6
Rated by Kroll (KBRA)	207	0.1	19.7	0.3
Rated by DBRS	7,179	5.1	366.6	6.5
Not Rated	26,348	18.6	453.8	8.0
All Bonds	141,630	100.0	5,667.9	100.0

Table 3.5: **Non-Agency Residential Mortgage Backed Securities: Credit Rating Composition 1987-2013**

This table shows the number of bonds and their corresponding principal amounts by credit rating. The credit rating corresponds to the rating assigned to a bond upon issuance. If several ratings were given, we have taken an average.

Rating	MBS Bonds		Principal Amount		Principal Amount By Agency		
	No.	Pct.	(\$ Billion)	Pct.	S&P	Moody's	Fitch
AAA	65,590.0	56.8	4,535.1	86.9	88.9	89.8	89.8
AA	13,298.0	11.5	297.0	5.7	4.8	5.4	4.1
A	13,355.0	11.6	212.3	4.1	2.9	2.5	3.7
BBB	13,062.0	11.3	118.4	2.3	1.9	1.8	1.8
BB	6,096.0	5.3	40.1	0.8	0.6	0.4	0.4
B	3,865.0	3.3	13.6	0.3	0.4	0.0	0.1
CCC	66.0	0.1	0.3	0.0	0.2	0.0	0.0
CC	22.0	0.0	0.6	0.0	0.2	0.0	0.0
C	51.0	0.0	3.3	0.1	0.1	0.0	0.0
Rated	115,405.0	81.2	5,220.5	91.7	4,523.4	3,931.5	2,531.4
Not Rated	26,774.0	18.8	472.1	8.3			


Table 3.6: **RMBS Losses Six Years After The Beginning of the Financial Crisis**

This table shows the number of securities with losses and the dollar size of the losses in December 2013, about six years after the beginning of the Subprime crisis in mid-2007.

Panel A: Losses by Credit Rating as of December 2013						
	Number of Securities			Dollar Amount		
	Total	W/ Losses	Pct. Losses	Total	W/ Losses	Pct. Losses
All RMBS	93,902	43,264	46.1	4,965.6	323.3	6.5
AAA	49,188	14,511	29.5	4,402.4	101.1	2.3
AA	12,087	6,425	53.2	263.6	89.6	34.0
A	11,144	6,752	60.6	144.9	56.3	38.8
BBB	12,015	8,309	69.2	101.6	48.0	47.3
NIG	9,468	7,267	76.8	53.1	28.3	53.2


Panel B: Losses by Mortgage Type and Credit Rating as of December 2013						
	Number of Securities			Dollar Amount		
	Total	W/ Losses	Pct. Losses	Total	W/ Losses	Pct. Losses
<i>All Securities</i>						
Prime	25,476.	11,987	47.1	1,238.7	37.5	3.0
Alt-A	27,135	17,763	65.5	1,327.3	145.2	10.9
Subprime	18,705	9,172	49.0	1,196.0	119.1	10.0
<i>AAA Rated Securities</i>						
Prime	15,610	5,418	34.7	1,172.7	14.8	1.3
Alt-A	14,851	7,646	51.5	1,210.0	78.9	6.5
Subprime	6,509	560	8.6	979.5	4.3	0.4
<i>Investment Grade Ex-AAA Securities</i>						
Prime	6,436	3,827	115.8	54.0	18.4	34.0
Alt-A	9,610	7,769	80.8	96.5	55.3	57.3
Subprime	10,893	7,456	68.4	203.9	104.8	51.4
<i>Non-Investment Grade Securities</i>						
Prime	3,430	2,742	79.9	12.0	4.3	36.2
Alt-A	2,674	2,348	87.8	20.8	10.9	52.7
Subprime	1,303	1,156	88.7	12.7	10.0	78.7

Figure 3.2: Losses and Unweighted Probability of Loss in RMBS Over Time




This figure shows the losses as a fraction of principal and the probability of losses incurred by the Residential Mortgage Backed securities in our database during the period 2000-2013.

Figure 3.3: Losses Over Time by Type of Mortgage for AAA-Rated RMBS



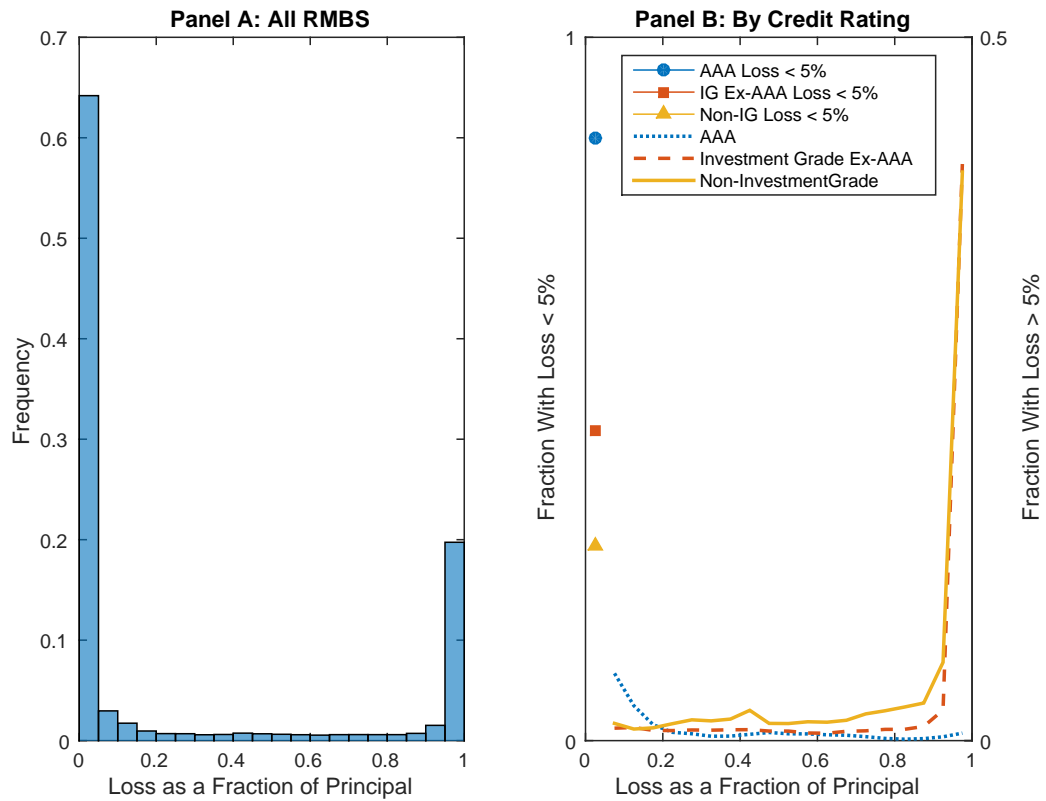
This figure plots the losses as a fraction of principal weighted by principal amount for the AAA-rated Residential Mortgage Backed Securities in our database during the period 2000-2013 by type of mortgage loan.

Figure 3.4: Dollar Amount of Losses in Non-Agency RMBS



This figure shows the cumulative Dollar amount of losses in RMBS up to December 2013 in billions of dollar. The category Investment Grade Ex-AAA includes AA, A, and BBB rated securities. The Non-Investment Grade Category includes all bonds rated BB and below

Figure 3.5: Distribution of Loss Size for All RMBS



Panel A presents the distribution of cumulative losses as of December 2013 as a fraction of the original principal amount for all the RMBS in our database issued from 1987 through 2008. Panel B shows the distribution of cumulative losses as of December 2013 as a fraction of the original principal amount for different groups of RMBS based on the type of the underlying mortgage loans.

Table 3.7: Losses and Credit Ratings by Vintage Group and Type of Mortgage Loan

This table presents principal-weighted regressions of the cumulative loss as fraction of initial principal as of December 2013 on credit rating dummy variables for all the RMBS in our database issued through 2008 classified by the type of mortgage loan underlying the securities.

Rating	Prime			Alt-A			Subprime		
	Before 2003	2003 - 2005	2006 - 2008	Before 2003	2003 - 2005	2006 - 2008	Before 2003	2003 - 2005	2006 - 2008
AAA	0.0000 (0.0001)	0.0032*** (0.0008)	0.0259*** (0.0015)	0.0002 (0.0003)	0.0076*** (0.0015)	0.0953*** (0.0015)	0.0004 (0.0008)	0.0013 (0.0018)	0.0068*** (0.0020)
AA	0.0001 (0.0007)	0.2242*** (0.0052)	0.5841*** (0.0093)	0.0045** (0.0019)	0.2097*** (0.0060)	0.7824*** (0.0080)	0.0007 (0.0040)	0.0348*** (0.0061)	0.6260*** (0.0061)
A	0.0006 (0.0011)	0.2837*** (0.0075)	0.3071*** (0.0117)	0.0080*** (0.0027)	0.3620*** (0.0094)	0.7260*** (0.0119)	0.0151*** (0.0050)	0.1566*** (0.0077)	0.7600*** (0.0078)
BBB	0.0039*** (0.0013)	0.3065*** (0.0088)	0.2957*** (0.0146)	0.0267*** (0.0033)	0.4728*** (0.0115)	0.5161*** (0.0117)	0.0718*** (0.0058)	0.3609*** (0.0091)	0.8654*** (0.0094)
BB	0.0163*** (0.0021)	0.2850*** (0.0095)	0.2303*** (0.0143)	0.0499*** (0.0054)	0.6415*** (0.0173)	0.3546*** (0.0144)	0.1113*** (0.0199)	0.5755*** (0.0191)	0.8861*** (0.0166)
B	0.0336*** (0.0026)	0.7159*** (0.0165)	0.8828*** (0.0344)	0.0863*** (0.0075)	0.7765*** (0.0254)	0.4816*** (0.0193)	0.2448*** (0.0419)	0.5133*** (0.0580)	0.7463*** (0.0495)
CCC	- (0.0754)	0.2474** (0.1091)	0.9484*** (0.2731)	0.9710*** (0.0866)	0.6840*** (0.1954)	0.8850** (0.3523)	- (0.0977)	0.3836** (0.1949)	0.9931*** (0.1512)
CC	- (0.1109)	0.0109 (0.0951)	- (0.2628)	- (0.1027)	- (0.0977)	0.6322*** (0.1078)	- (0.1027)	0.1189 (0.3823)	- (0.5566)
C or Below	- (0.1215)	0.7679*** (0.1359)	0.9687 (0.5928)	0.4963*** (0.0449)	- (0.0219)	0.3112*** (0.0241)	- (0.1369)	0.3775 (0.3712)	0.9932** (0.5020)
Observations	4,095	13,366	8,015	2,908	8,226	16,001	1,363	6,028	11,314
R-squared	0.0554	0.3468	0.4182	0.1571	0.4329	0.4975	0.1432	0.3217	0.7052
Weighted	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Standard errors in parentheses

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

Table 3.8: **Principal-Weighted Losses in RMBS and Credit Ratings**


This table shows regressions of the cumulative loss as fraction of initial principal as of December 2013 on credit rating dummy variables. The regressions are weighted by the principal dollar amount upon issuance of each RMBS.

Credit Rating	Full Sample	Before 2003	2003 - 2005	2006-2008
AAA	0.0218*** (0.0006)	0.0002 (0.0001)	0.0034*** (0.0007)	0.0483*** (0.0011)
AA	0.3096*** (0.0025)	0.001 (0.0008)	0.1180*** (0.0028)	0.5091*** (0.0043)
A	0.3620*** (0.0033)	0.0055*** (0.0008)	0.2000*** (0.0036)	0.6572*** (0.0062)
BBB	0.4480*** (0.0040)	0.0334*** (0.0013)	0.3152*** (0.0041)	0.6655*** (0.0072)
BB	0.4923*** (0.0064)	0.0653*** (0.0029)	0.4886*** (0.0075)	0.5136*** (0.0102)
B	0.5812*** (0.0117)	0.0938*** (0.0042)	0.6989*** (0.0147)	0.5619*** (0.0182)
CCC	0.7360*** (0.0867)	0.4125*** (0.0558)	0.4102*** (0.0987)	0.9465*** (0.1361)
CC	0.2036*** (0.0562)	0.1364 (0.0964)	0.0251 (0.1228)	0.2005*** (0.0719)
C or Below	0.3863*** (0.0225)	0.0661*** (0.0227)	0.6607*** (0.1665)	0.3604*** (0.0274)
Observations	93,902	19,230	38,381	36,291
R-squared	0.3217	0.0852	0.2972	0.485

Standard errors in parentheses

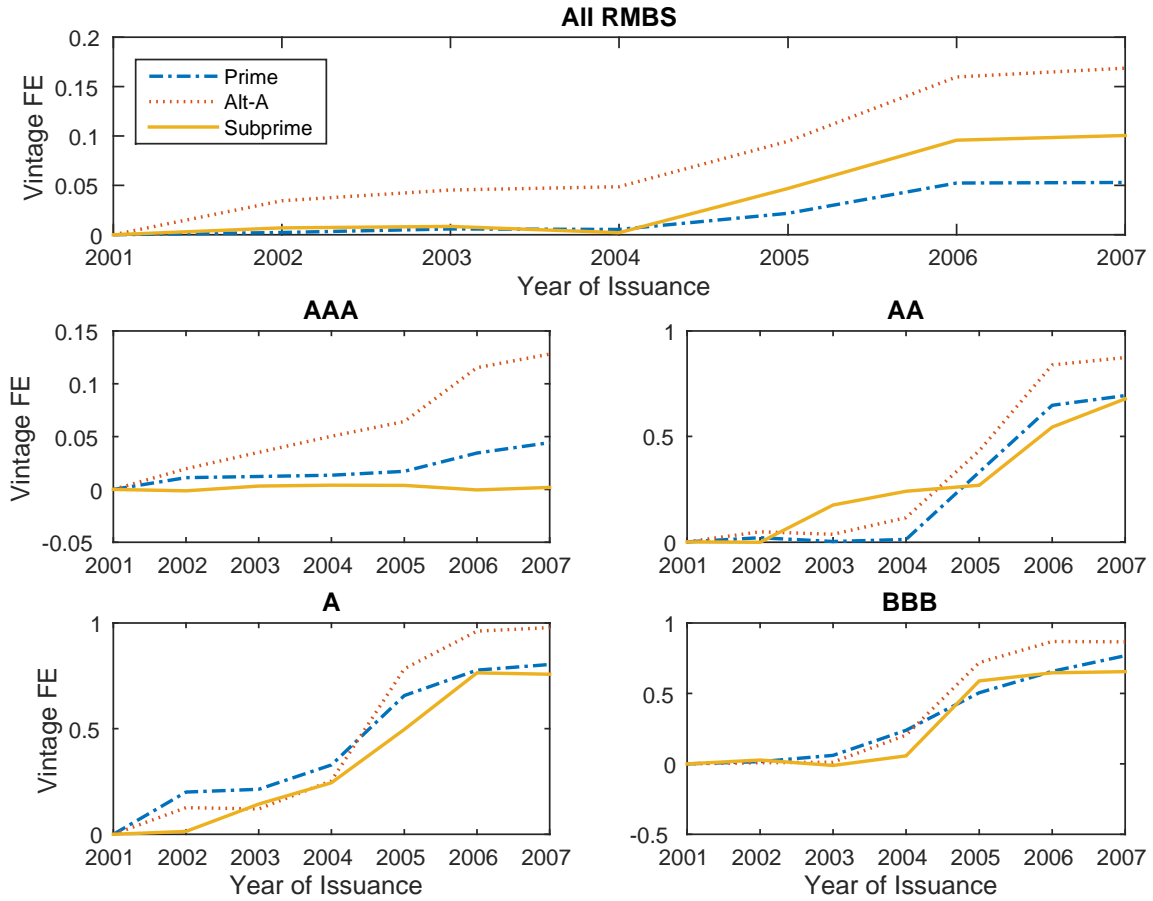
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 3.6: Vintage Fixed Effects on Weighted Losses



This figure plots the coefficient estimates corresponding to issue year (vintage) dummy variables in linear regressions that have as left hand side variable the cumulative losses as of December 2013 as a fraction of principal and on the right hand side have all the covariates available in our database as controls. The lines are the mean plus/minus one standard error.

Figure 3.7: Vintage Fixed Effects on Weighted Losses by Type of Mortgage Loan

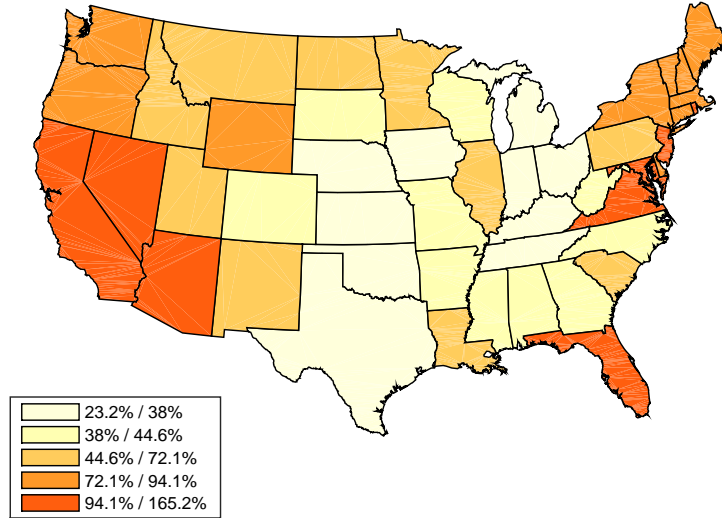


This figure plots the coefficient estimates corresponding to issue year (vintage) dummy variables in linear regressions that have as left hand side variable the cumulative losses as of December 2013 as a fraction of principal and on the right hand side have all the covariates available in our database as controls.

Figure 3.8: State-Level House Price Boom and Bust From 2000-Q1 to 2009-Q2

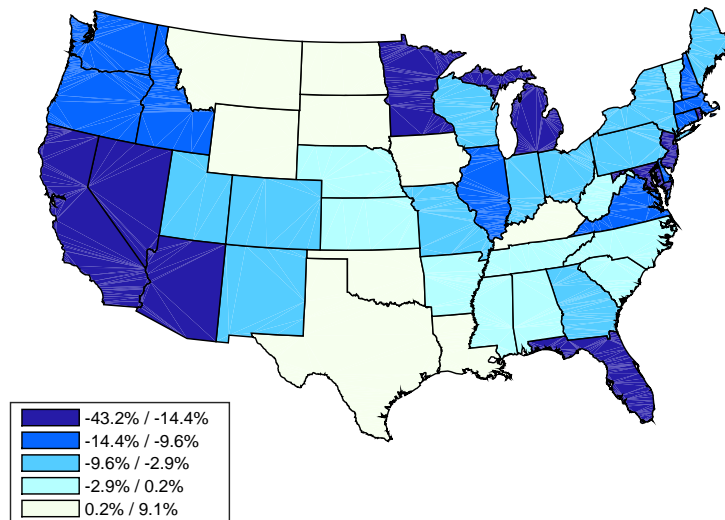
.8

Panel A. Boom: 2000-Q1 to 2006-Q4



.8

Panel B. Bust: 2006-Q4 to 2009-Q4

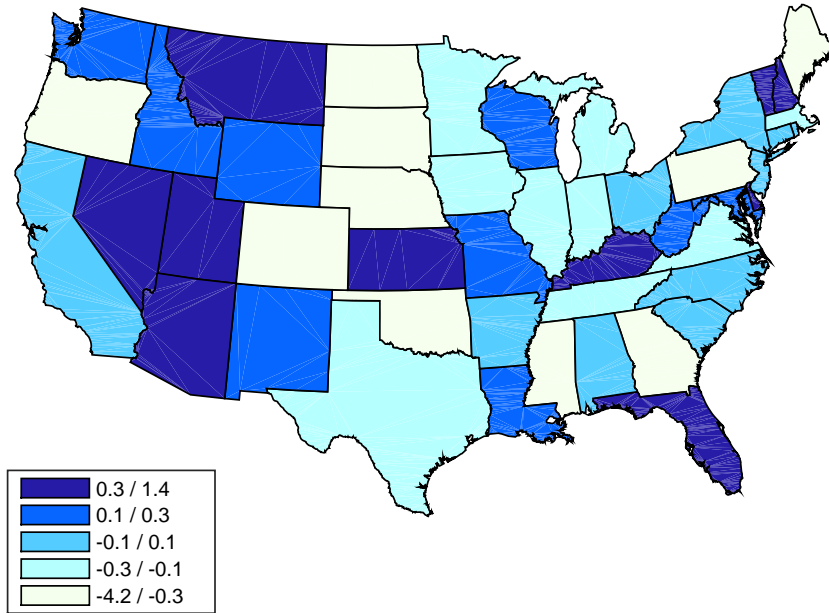


This map highlights the differences in house price increases between the first quarter of 2000 and the fourth quarter of 2006 and house price decreases between the fourth quarter of 2006 and the fourth quarter of 2009 across states. The house price data comes from the Federal Housing Agency and corresponds to the State-Level All-transaction indexes available here <http://www.fhfa.gov/DataTools/Downloads/Pages/House-Price-Index-Datasets.aspx#qexx>. States are grouped in quintiles according to the house price change the experienced in each period. Each color on the map represents a quintile, with darkest colors representing bigger absolute changes (red for the period of price increases and blue for the period of price decreases).

Figure 3.9: State-Level Dummies on Loss Rates: with and without controls

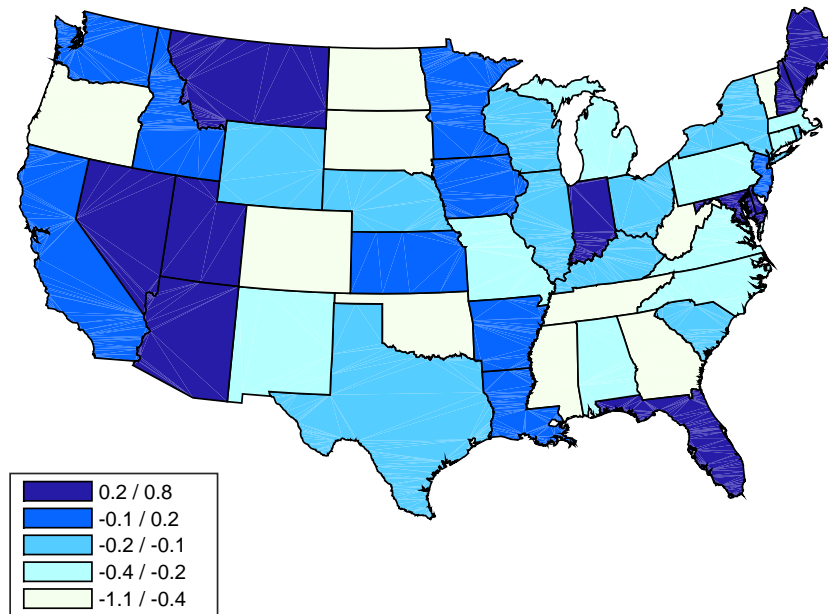
.8

Panel B. Dummy Coefficients Regression with Controls



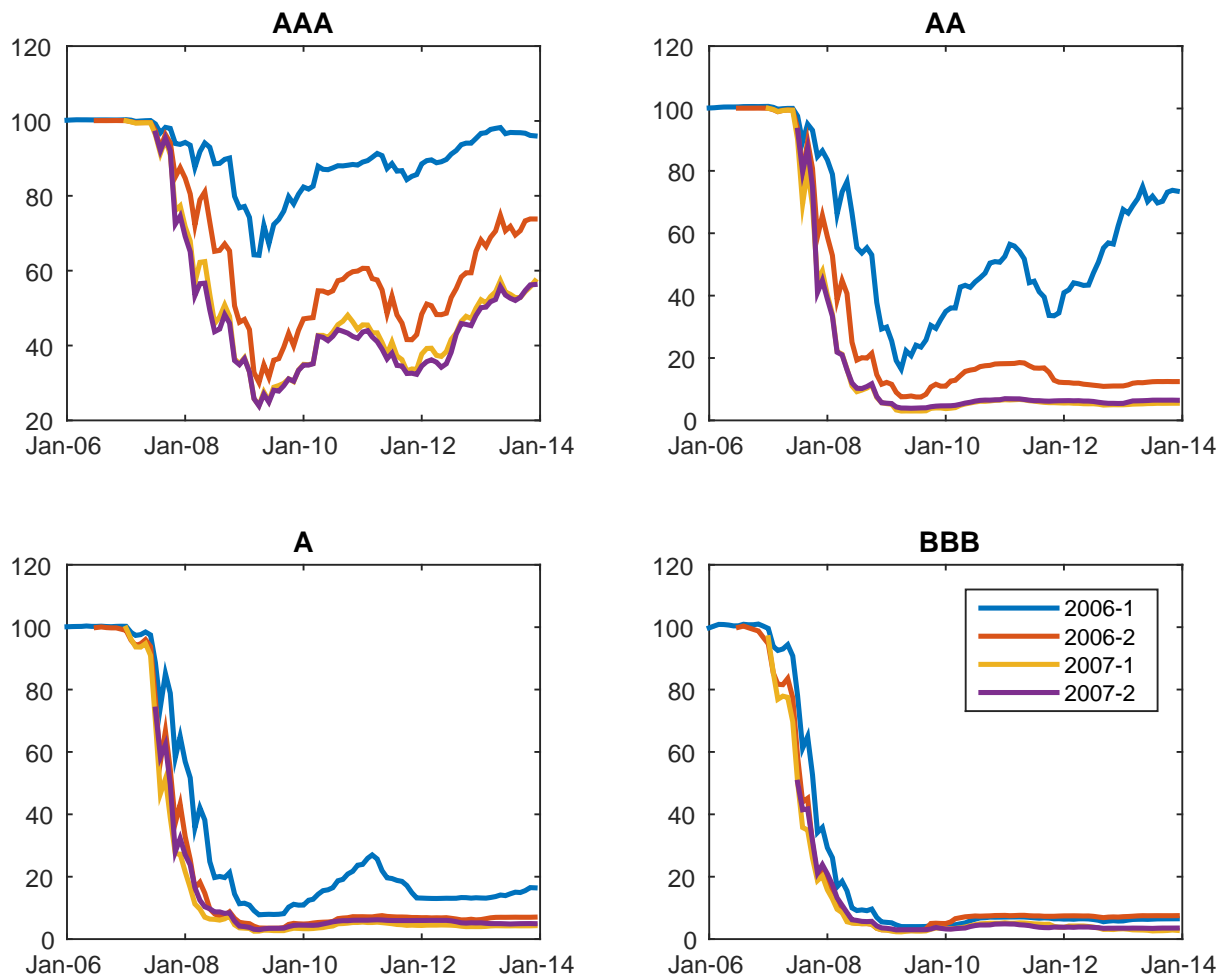
.8

Panel B. Dummy Coefficients Regression without Controls




The map is colored according to the regression of loss rates on the state dummy variables, with or without the inclusion of additional controls. The regression coefficients are sorted into quintiles to deliver the coloring scheme, with darker colors representing larger losses.

Figure 3.13: Subprime RMBS Price Indexes




This figure plots the prices of ABX.HE indexes by Markit. Each line represents a vintage of subprime RMBS and the Index. Each panel shows the evolution of prices over time by credit rating. These indexes are constructed based on 20 deals.

Figure 3.10: Ex-Ante vs Ex-Post Ideal Ratings




In this figure, for each security, we compare the original credit rating (which we call here Ex-Ante Rating) to the rating that ex-post we would have assigned given the security's realized loss using Moody's idealized Expected Loss Table by Rating. This table is available here http://siteresources.worldbank.org/EXTECAREGTOPPRVSECDEV/Resources/570954-1211578683837/Bielecki_Moodys_Rating_SME_transactions.pdf. The solid line shows the fraction of securities that was assigned each rating level. The dotted line shows the fraction of securities that should have gotten each rating level based on their loss as a fraction of original principal. We do this for all securities issued up to 2008.

Figure 3.11: "Right" and "Wrong" Ratings Based on Moody's Ideal Ratings




This figure compares the original rating of each security to the rating we would have assigned ex-post based on Moody's idealized Expected Loss Table by Rating.

Figure 3.12: “Misratings” Ratings Based on Moody’s Ideal Ratings



This figure classifies each security in a bin defined in two dimensions. One dimension is the ex-ante credit rating as determined by the original credit rating. The second dimension is the ex-post rating determined by Moody’s table for idealized expected losses. If all securities had behaved as expected, all the mass would be represented in bars on the diagonal running southwest-northeast in the plot. The height of the bar represents the number of securities.

Figure 3.14: Summary Statistics of Prices Collected by FINRA



This figure shows summary statistics of daily transaction prices collected by the Financial Industry Regulatory Authority from May 2011 through May 2016 on Non-Agency MBS. The plots at the top break up the statistics by Investment Grade and Non-Investment Grade, while the plots at the bottom break up the statistics by groups of vintages only for Investment Grade securities. FINRA produces this information daily since 2011. The lines in the different figures correspond to 22-day moving averages (daily monthly averages) of the daily values reported by FINRA. Here we report the principal weighted average and the 25th and 75th percentiles of the average transaction price. The daily reports are available here <http://tps.finra.org/idc-index.html>

Table 3.9: House Prices and Loss Rates

This table presents linear regressions to study the relation between the the cumulative loss as fraction of initial principal as of December 2013 and changes in house prices.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔHP 2000-2006	0.073*** 0.003		-0.218*** 0.010		-0.178*** 0.012	-0.021 0.027		
ΔHP 2006-2009		-0.203*** 0.006	-0.63*** 0.021		-0.532*** 0.020	0.342*** 0.061		
Price Reversal				-0.238*** 0.007				
AA					0.426*** 0.003	0.423*** 0.003	0.423*** 0.003	0.427*** 0.003
A					0.493*** 0.004	0.488*** 0.004	0.489*** 0.004	0.492*** 0.004
BBB					0.555*** 0.005	0.55*** 0.005	0.55*** 0.005	0.555*** 0.005
BB					0.5*** 0.007	0.492*** 0.007	0.494*** 0.007	0.502*** 0.007
B					0.599*** 0.013	0.594*** 0.012	0.594*** 0.013	0.606*** 0.013
CCC					0.749*** 0.087	0.74*** 0.086	0.739*** 0.086	0.747*** 0.088
CC					0.496*** 0.089	0.493*** 0.087	0.482*** 0.088	0.512*** 0.089
C or Below					0.324*** 0.023	0.305*** 0.022	0.307*** 0.022	0.348*** 0.023
Subprime					0.009*** 0.002	-0.003* 0.002	-0.004** 0.002	0.008*** 0.002
Alt-A					0.049*** 0.002	0.032*** 0.002	0.033*** 0.002	0.06*** 0.002
Contstant	0.011*** 0.002	-0.001 0.002	0.019*** 0.002	-0.121*** 0.006	-0.038*** 0.005	0.059*** 0.007	-0.026*** 0.004	0.004*** 0.001
State Dummies	No	No	No	No	No	Yes	Yes	No
Weighted Dummies	No	No	No	No	No	Yes	Yes	No
Observations	93,902	93,902	93,902	93,902	71,316	71,316	71,316	71,316
R-squared	0.0059	0.0107	0.0156	0.0128	0.4345	0.4513	0.4492	0.4272
Weighted Regression	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Standard errors in parentheses

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

Table 3.10: Credit Rating Reversals

Panel A presents the calculation of equation (3).

Panel A: Gains from Including Other Covariates							
	$\alpha = 0$	$\alpha = 0.05$	$\alpha = 0.1$	$\alpha = 0.25$	$\alpha = 0.5$	$\alpha = 0.75$	$\alpha = 0.9$
<i>Principal Value Weighted</i>							
AAA vs AA	1.65	0.47	0.24	0.12	0.08	0.04	0.03
AAA vs A	2.26	0.58	0.28	0.14	0.08	0.04	0.03
AAA vs BBB	0.82	0.32	0.17	0.09	0.06	0.03	0.02
AA vs A	-8.52	-6.76	-6.69	-11.29	-20.50	-59.50	-699.95
AA vs BBB	2.27	2.54	2.77	2.59	2.79	2.62	2.54
A vs BBB	1.90	1.94	2.07	2.32	2.82	3.02	3.11
Panel B: Credit Rating Reversals							
	$\alpha = 0$	$\alpha = 0.05$	$\alpha = 0.1$	$\alpha = 0.25$	$\alpha = 0.5$	$\alpha = 0.75$	$\alpha = 0.9$
<i>Principal Value Weighted</i>							
AAA switched to AA	49.9	13.9	1.0	0.3	0.3	0.0	0.0
AAA switched to A	54.5	21.4	2.3	0.3	0.3	0.0	0.0
AAA switched to BBB	31.2	2.8	0.4	0.3	0.2	0.0	0.0
AA switched to A	75.4	73.2	72.8	71.7	69.0	67.4	66.1
AA switched to BBB	63.8	63.1	63.1	62.8	60.9	59.2	57.8
A switched to BBB	67.4	67.7	67.9	68.1	68.5	67.3	65.4

Table 3.11: **Principal-Weighted Probability of Loss in RMBS and Credit Ratings**

This table presents linear regressions to study the relation between the probability of incurring losses and credit ratings.

Credit Rating	Full Sample	Before 2003	2003 - 2005	2006-2008
AAA	0.2747*** (0.0016)	0.0639*** (0.0019)	0.1667*** (0.0021)	0.4636*** (0.0028)
AA	0.1603*** (0.0066)	0.0207* (0.0107)	0.1010*** (0.0088)	0.1667*** (0.0106)
A	0.2073*** (0.0087)	0.0178 (0.0111)	0.1472*** (0.0114)	0.3351*** (0.0155)
BBB	0.3177*** (0.0104)	0.1690*** (0.0170)	0.3168*** (0.0131)	0.3136*** (0.0179)
BB	0.5943*** (0.0169)	0.2203*** (0.0389)	0.6514*** (0.0238)	0.4819*** (0.0253)
B	0.5885*** (0.0308)	0.2264*** (0.0560)	0.7656*** (0.0466)	0.4303*** (0.0452)
CCC	0.5822** (0.2272)	0.9361 (0.7475)	0.4627 (0.3134)	0.5364 (0.3381)
CC	0.6340*** (0.1472)	0.4459 (1.2919)	-0.0239 (0.3900)	0.5364*** (0.1785)
C or Below	0.7176*** (0.0589)	0.0696 (0.3045)	0.8333 (0.5286)	0.5364*** (0.0681)
Observations	93,902	19,230	38,381	36,291
R-squared	0.0370	0.0079	0.0451	0.0376

Standard errors in parentheses

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

Table 3.12: Internal Rate of Return Calculations From Issuance to 2013 by Credit Rating. For AAA: also Type of Mortgage.

Return Statistic	80% TV	90% TV	100% TV
<i>By Credit Rating</i>			
AAA	2.44	2.89	3.31
AA	-7.90	-7.01	-6.21
A	-10.92	-10.10	-9.35
BBB	-13.56	-12.80	-12.11
Inv. Grade Ex AAA	-9.01	-8.15	-7.38
<i>By Type of Mortgage</i>			
AAA Prime	3.61	3.98	4.33
AAA SubPrime	1.61	2.14	2.62
AAA AltA	1.37	2.01	2.61
<i>Fixed Rate MBS</i>			
AAA Prime Fixed	4.25	4.56	4.84
AAA SubPrime Fixed	4.86	4.96	5.04
AAA AltA Fixed	3.64	4.13	4.58
<i>Floating Rate MBS</i>			
AAA Prime Floating	3.03	3.45	3.83
AAA SubPrime Floating	1.45	1.97	2.44
AAA AltA Floating	0.42	1.12	1.76

Table 3.13: Discount Margin (Return Over 3-month Tbill) Calculations From Issuance to 2013 by Credit Rating

This table presents discount margin calculations for the RMBS in our database by credit rating using the 3-month Tbill rate as benchmark. The discount margin IRR solves equation 5.

Return Statistic	Terminal Value Assumption					
	80% TV	90% TV	100% TV	Same Loss	Mean Loss	Median Loss
Panel A: Principal Value - Weighted						
AAA-Rated						
Mean	0.0	0.3	0.7	0.6	0.7	0.6
Median	1.3	1.6	2.1	2.1	2.1	1.9
Std. Dev.	20.1	20.2	20.2	20.2	20.2	20.2
AA-Rated						
Mean	-1.6	-1.6	-1.6	-1.7	-1.8	-1.7
Median	-0.4	0.7	1.8	1.7	-1.1	-1.0
Std. Dev.	52.5	52.4	52.6	52.6	51.9	51.8
A-Rated						
Mean	-1.2	-1.2	-1.2	-1.2	-1.3	-1.2
Median	-0.1	0.4	1.0	1.0	-4.4	-3.6
Std. Dev.	52.5	52.4	52.6	52.6	51.9	51.8
BBB-Rated						
Mean	-1.0	-1.0	-1.0	-1.0	-1.1	-1.1
Median	-5.3	-7.5	-7.4	-8.2	-18.1	-11.3
Std. Dev.	67.6	66.9	67.1	67.0	65.5	66.1
Panel B: Unweighted						
AAA-Rated						
Mean	-1.3	-1.2	-1.0	-1.0	-1.0	-1.0
Median	2.2	2.6	2.8	2.7	2.8	2.7
Std. Dev.	34.3	34.8	34.8	34.8	34.8	34.8
AA-Rated						
Mean	-5.3	-5.5	-5.5	-5.5	-5.8	-5.6
Median	-1.2	-0.8	-0.4	-1.1	-7.8	-8.1
Std. Dev.	58.7	58.2	58.5	58.3	57.6	57.7
A-Rated						
Mean	-6.9	-7.2	-7.2	-7.2	-7.4	-7.3
Median	-18.0	-25.3	-25.3	-26.2	-32.8	-26.2
Std. Dev.	58.7	58.2	58.5	58.3	57.6	57.7
BBB-Rated						
Mean	-8.3	-8.6	-8.6	-8.6	-8.9	-8.7
Median	-34.3	-41.0	-41.0	-41.2	-44.1	-41.1
Std. Dev.	70.8	69.7	69.8	69.7	69.1	69.2

Table 3.14: **Principal Value Weighted Discount Margin (Return Over 3-month Tbill) Calculations by Vintage**

This table presents discount margin calculations for the RMBS in our database by vintage (year of issuance) using the 3-month Tbill rate as benchmark. The discount margin IRR solves equation 5.

Return Statistic	Terminal Value Assumption					
	80% TV	90% TV	100% TV	Same Loss	Mean Loss	Median Loss
<i>Vintage 2000</i>						
Mean	0.0	0.0	0.0	0.0	0.0	0.0
Median	5.1	5.1	5.1	5.1	5.1	5.1
Std. Dev.	22.1	22.1	22.1	22.1	22.1	22.1
<i>Vintage 2001</i>						
Mean	0.1	0.1	0.1	0.1	0.1	0.1
Median	5.4	5.4	5.4	5.4	5.4	5.4
Std. Dev.	19.8	19.8	19.7	19.7	19.7	19.7
<i>Vintage 2002</i>						
Mean	0.1	0.1	0.1	0.1	0.1	0.1
Median	3.9	3.9	3.9	3.9	3.9	3.9
Std. Dev.	21.5	21.5	21.5	21.5	21.5	21.5
<i>Vintage 2003</i>						
Mean	0.2	0.2	0.2	0.2	0.2	0.2
Median	3.4	3.6	3.7	3.7	3.7	3.7
Std. Dev.	10.6	10.7	10.7	10.7	10.7	10.7
<i>Vintage 2004</i>						
Mean	0.3	0.3	0.3	0.3	0.3	0.3
Median	3.2	3.3	3.5	3.5	3.5	3.5
Std. Dev.	17.1	17.1	17.2	17.2	17.2	17.2
<i>Vintage 2005</i>						
Mean	0.1	0.2	0.2	0.2	0.2	0.2
Median	3.4	3.8	4.0	4.0	4.0	4.0
Std. Dev.	16.7	16.6	16.5	16.6	16.5	16.5
<i>Vintage 2006</i>						
Mean	-1.9	-1.8	-1.7	-1.8	-1.9	-1.8
Median	1.6	2.2	2.8	2.7	2.8	2.5
Std. Dev.	33.8	33.9	34.0	34.0	34.6	34.0
<i>Vintage 2007</i>						
Mean	-1.3	-1.2	-1.1	-1.1	-1.2	-1.2
Median	1.3	1.9	2.5	2.4	2.3	2.0
Std. Dev.	39.9	40.0	40.2	40.2	40.2	40.1

Table 3.15: Discount Margin (Return Over 3-month Tbill) Calculations From Issuance to 2013 by Type of Mortgage Loan

This table presents discount margin calculations for the RMBS in our database by type of mortgage loan using the 3-month Tbill rate as benchmark. The discount margin IRR solves equation 5.

Return Statistic	Terminal Value Assumption					
	80% TV	90% TV	100% TV	Same Loss	Mean Loss	Median Loss
Panel A: Principal Value - Weighted						
<i>Subprime</i>						
Mean	-2.2	-2.2	-2.1	-2.1	-2.4	-2.3
Median	0.0	0.6	1.5	1.5	1.4	0.9
Std. Dev.	35.0	35.2	35.3	35.4	35.8	35.3
<i>Alt-A</i>						
Mean	-1.5	-1.4	-1.3	-1.4	-1.3	-1.3
Median	0.8	1.3	2.0	1.9	1.9	1.7
Std. Dev.	35.4	35.6	35.7	35.7	35.7	35.6
<i>Prime</i>						
Mean	0.4	0.4	0.5	0.5	0.5	0.5
Median	3.2	3.5	3.8	3.8	3.8	3.7
Std. Dev.	15.2	15.4	15.2	15.3	15.3	15.3
Panel B: Unweighted						
<i>Subprime</i>						
Mean	-8.5	-8.7	-8.6	-8.7	-9.2	-8.9
Median	-1.2	-0.3	0.0	-0.1	-1.8	-2.4
Std. Dev.	64.4	64.2	64.4	64.3	63.8	63.7
<i>Alt-A</i>						
Mean	-9.7	-9.9	-9.8	-9.9	-9.9	-9.9
Median	-0.4	0.2	0.8	0.4	0.6	0.2
Std. Dev.	64.1	63.7	63.9	63.8	63.8	63.8
<i>Prime</i>						
Mean	-2.0	-2.2	-2.1	-2.1	-2.1	-2.2
Median	2.9	3.3	3.6	3.6	3.6	3.4
Std. Dev.	37.2	38.0	38.1	38.1	38.1	38.1

Table 3.16: Median Discount Margin (Return Over 3-month Tbill) for AAA-rated RMBS by Vintage, Loan Type and Bond Type

This table presents discount margin calculations for the AAA RMBS in our database by type of mortgage loan, vintage, and by type of bond (floating rate or fixed rate) using the 3-month Tbill rate as benchmark. The discount margin IRR solves equation 5.

Loan Type	Fixed Rate			Floating Rate			3-Month Libor
	80% TV	90% TV	100% TV	80% TV	90% TV	100% TV	
<i>Vintage 2000</i>							
Subprime	4.5	4.5	4.5	1.6	1.6	1.6	
Alt-A	4.0	4.0	4.0	2.6	2.6	2.6	2.4
Prime	3.7	3.7	3.7	3.2	3.3	3.3	
<i>Vintage 2001</i>							
Subprime	2.9	2.9	2.9	0.3	0.4	0.4	
Alt-A	4.9	4.9	4.9	3.8	3.8	3.9	2.1
Prime	4.7	4.7	4.7	4.4	4.4	4.4	
<i>Vintage 2002</i>							
Subprime	2.8	2.8	2.8	-0.1	-0.1	-0.1	
Alt-A	4.6	4.7	4.7	1.1	1.2	1.2	2.0
Prime	4.4	4.4	4.4	3.0	3.0	3.0	
<i>Vintage 2003</i>							
Subprime	4.9	5.1	5.4	-0.7	-0.6	-0.6	
Alt-A	3.6	3.7	3.7	1.2	1.3	1.5	2.4
Prime	3.5	3.5	3.6	1.6	1.6	1.7	
<i>Vintage 2004</i>							
Subprime	4.4	4.6	4.7	-0.9	-0.8	-0.8	
Alt-A	4.0	4.1	4.2	0.6	0.9	1.2	2.0
Prime	3.9	4.0	4.1	2.1	2.4	2.7	
<i>Vintage 2005</i>							
Subprime	5.0	5.3	5.8	0.6	0.8	0.8	
Alt-A	3.9	4.4	4.8	0.9	1.5	2.0	2.1
Prime	4.2	4.5	4.8	2.5	3.0	3.5	
<i>Vintage 2006</i>							
Subprime	4.3	5.0	5.7	0.7	1.3	2.0	
Alt-A	3.1	3.9	4.6	-0.1	0.9	1.8	2.1
Prime	4.5	4.9	5.4	2.8	3.4	4.1	
<i>Vintage 2007</i>							
Subprime	4.3	5.0	5.8	-0.6	0.6	1.6	
Alt-A	2.9	3.8	4.6	-1.0	0.2	1.3	1.9
Prime	4.5	5.1	5.6	2.4	3.1	3.8	

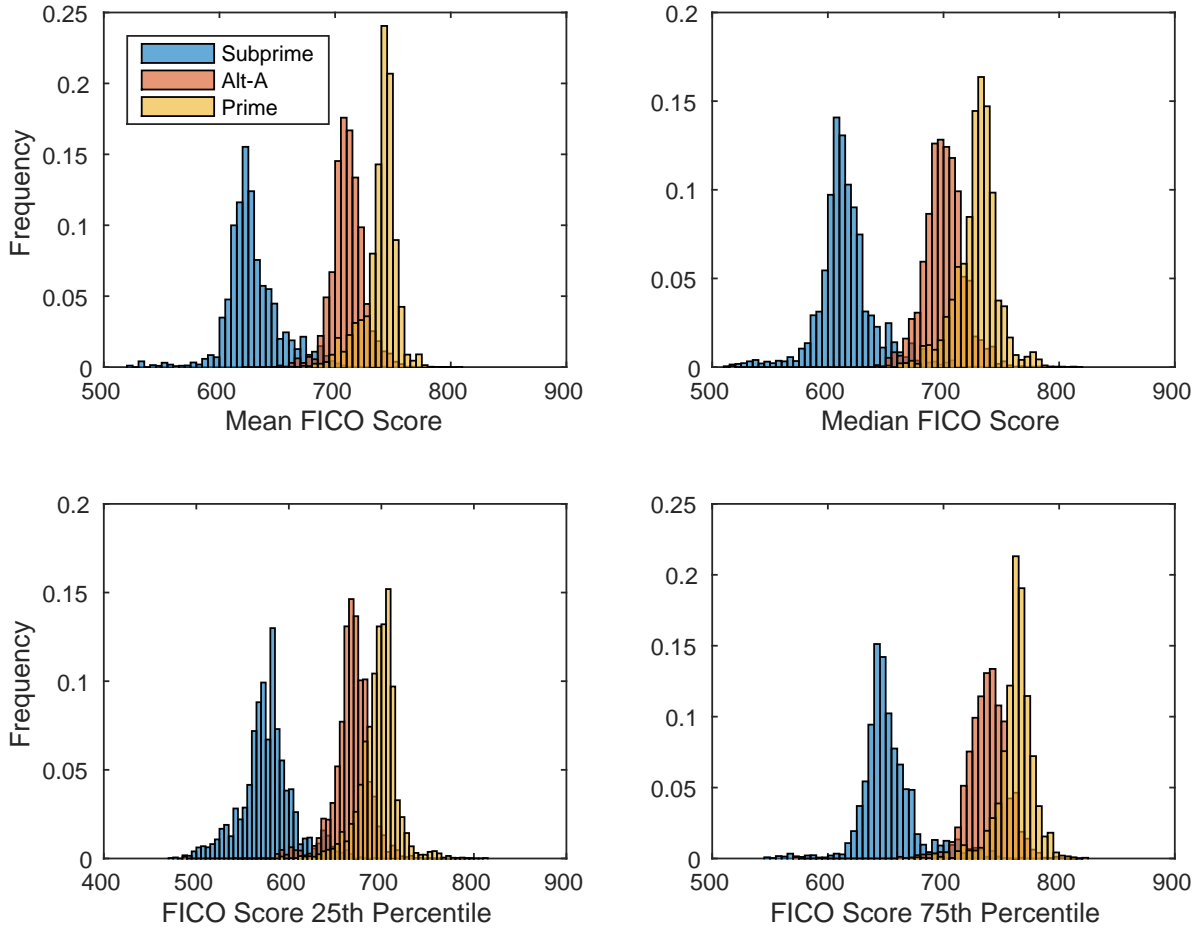
Table 3.17: Subprime RMBS Returns for Deals Underlying ABX.HE Indexes

This table reports monthly returns computed from the evolution of prices of the ABX.HE indexes. The calculations shows three time periods. The first period is the entire period of analysis, 2006-2013. The second period runs from January 2006 through May 2009, when the prices of AAA bonds bottomed out. The third period goes from May 2009 through December 2013.

	AAA	AAAp	AA	A	BBB	BBB-
<i>Vintage 2006-1</i>						
From Jan 06 to Dec 13	0.00	-0.37	0.09	-1.35	-2.36	-2.62
From Jan 06 to May 09	-1.08	-1.37	-4.11	-5.52	-7.07	-7.28
From May 09 To Dec 13	0.76	-0.17	3.01	1.54	0.92	0.63
<i>Vintage 2006-2</i>						
From June 06 to Dec 13	-0.09	-0.11	-1.77	-2.31	-2.17	-2.40
From June 06 to May 09	-3.30	-4.56	-6.55	-8.46	-9.70	-9.84
From May 09 To Dec 13	1.79	0.77	1.05	1.32	2.27	1.98
<i>Vintage 2007-1</i>						
From Jan 07 to Dec 13	-0.38	0.05	-2.73	-3.05	-3.62	-4.49
From Jan 07 to May 09	-4.76	-8.01	-10.86	-11.50	-12.13	-12.15
From May 09 To Dec 13	1.73	1.64	1.19	1.02	0.49	-0.80
<i>Vintage 2007-2</i>						
From June 07 to Dec 13	-0.41	0.24	-2.76	-2.86	-2.94	-2.83
From June 07 to May 09	-6.00	-7.94	-12.56	-12.98	-11.83	-11.31
From May 09 To Dec 13	1.68	1.85	0.91	0.94	0.39	0.35

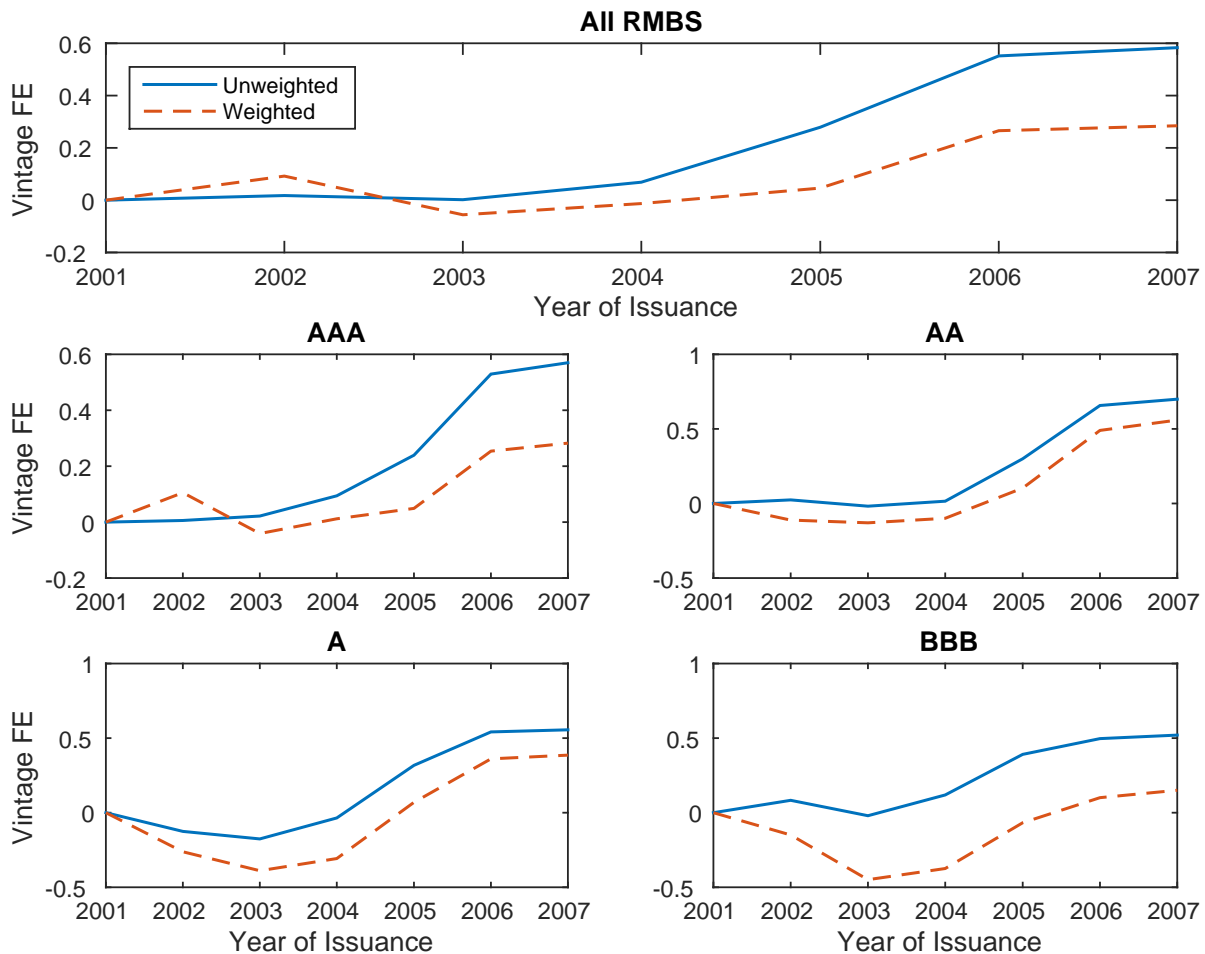
3.7 Appendix Figures

Figure 3.15: Distribution of FICO Scores by Type of Mortgage



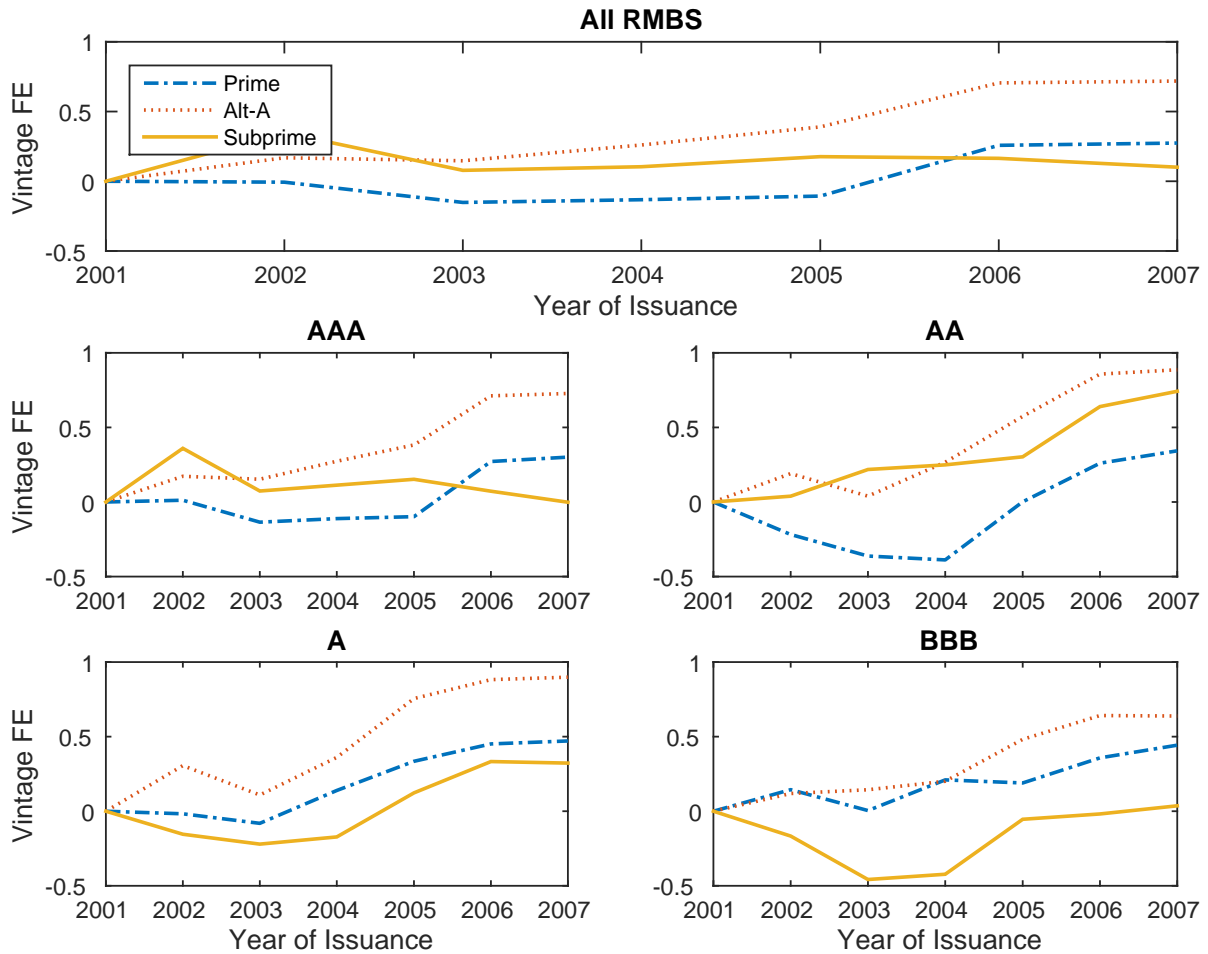
This figure plots histograms for different moments of the distribution of FICO Scores of the mortgage loans underlying the Residential Mortgage Backed Securities in our database. These moments correspond to the value-weighted average, the median, 25th and 75th percentiles of FICO scores upon issuance of the MBS. The corresponding Bloomberg fields are MTG_WAOCS, MTG_QRT_SCORE_MED, MTG_QRT_SCORE_25, and MTG_QRT_SCORE_75. The histograms are shown by type of mortgage loan (Prime, Alt-A, and Subprime) according to the classification provided by the Mortgage Market Statistical Annual 2013 Edition. Only securities issued in the period 2006-2012 are included.

Figure 3.16: Vintage Fixed Effects on Probability of Loss



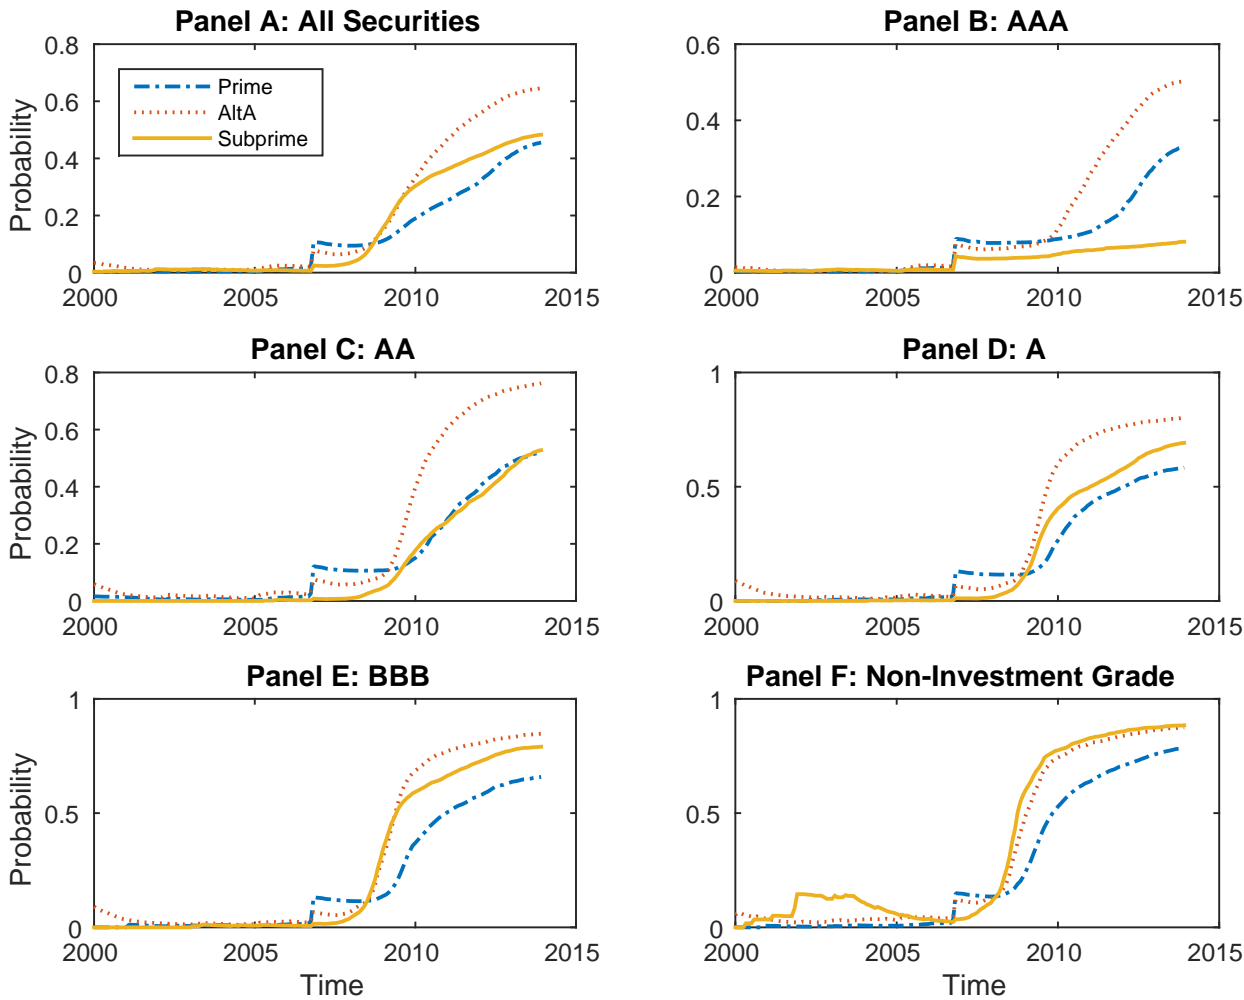
This figure plots the coefficient estimates corresponding to issue year (vintage) dummy variables in linear regressions that have as left hand side variable a dummy variable that takes the value one if the cumulative losses as of December 2013 are strictly greater than zero, and takes the value zero otherwise. The right hand side have all the covariates available in our database as controls, including issue year dummies. The solid line corresponds to the point estimate in non-weighted regressions and the dashed line to the point estimates in principal-weighted regressions. The coefficients are measured with respect to year 2001, whose coefficient is normalized to zero. The top panel shows the results including all credit ratings. The bottom four panels present the results for regressions for each credit rating. The LHS of the regression is computed by accumulating the time series of losses for each MBS (Bloomberg field HIST_LOSSES) up to December 2013 and assigning the value one to those securities with strictly positive losses. The credit rating is based on the credit ratings assigned upon issuance of the security (for example, for the rating by Moody's we use the Bloomberg field RTG_MDY_INITIAL). The calculations include all the bonds issued since 1987. We exclude all the MBS bonds for which the original principal amount is only a reference or that can distort our computations. The excluded bonds include bonds with zero original balance, excess tranches, interest-only bonds, and Net Interest Margin deals (NIM).

Figure 3.17: Vintage Fixed Effects on Weighted Probability of Loss by Type of Mortgage Loan



This figure plots the coefficient estimates corresponding to issue year (vintage) dummy variables in linear regressions that have as left hand side variable a dummy variable that takes the value one if the cumulative losses as of December 2013 are strictly greater than zero, and takes the value zero otherwise. The right hand side have all the covariates available in our database as controls, including issue year dummies. Each panel shows the coefficients for regressions run separately for different categories of mortgage loan. Each MBS is assigned to one of three categories (Prime, Alt-A, and Subprime) based on different moments of the distribution of FICO scores of the underlying mortgage loans using the classification of bonds issued after 2005 by the Mortgage Market Statistical Annual Edition 2013. See figures 3.1 and 3.15 to understand the basis of the classification. The regressions are weighted by the original principal amount and the coefficients are measured with respect to year 2001, whose coefficient is normalized to zero. Year 2001 was the earliest year for which we could estimate a coefficient for the three types. The top panel shows the results including all credit ratings. The bottom four panels present the results for regressions for each credit rating.

Figure 3.18: Non-weighted Probability of Loss Over Time by Type of Mortgage and by Credit Rating



This figure plots the fraction of securities that incurred losses for all Residential Mortgage Backed Securities in our database during the period 2000-2013 by credit rating and by type of mortgage loan. Each bond is assigned to one of three categories (Prime, Alt-A, and Subprime) based on different moments of the distribution of FICO scores of the underlying mortgage loans using the classification of bonds issued after 2005 by the Mortgage Market Statistical Annual Edition 2013. See figures 3.1 and 3.15 to understand the basis of the classification.. Each panel shows the calculations by different credit ratings or groups of ratings for the three categories. Both the principal amount and the ratings refer to their values upon issuance of the security. At every point in time, all securities that have been issued at that point in time or before enter the calculations. The calculations include all the bonds issued from 1987 through 2008 as long as they had a credit rating available, but we only plot the results starting in year 2000. We exclude all the MBS bonds for which the original principal amount is only a reference or that can distort our computations.

Figure 3.19: Moody Idealized Cumulative Expected Loss Rates

Rating	Year									
	1	2	3	4	5	6	7	8	9	10
Aaa	0.000028%	0.000110%	0.000390%	0.000990%	0.001600%	0.002200%	0.002860%	0.003630%	0.004510%	0.005500%
Aa1	0.000314%	0.001650%	0.005500%	0.011550%	0.017050%	0.023100%	0.029700%	0.036850%	0.045100%	0.055000%
Aa2	0.000748%	0.004400%	0.014300%	0.025850%	0.037400%	0.048950%	0.061050%	0.074250%	0.090200%	0.110000%
Aa3	0.001661%	0.010450%	0.032450%	0.055500%	0.078100%	0.100650%	0.124850%	0.149600%	0.179850%	0.220000%
A1	0.003196%	0.020350%	0.064350%	0.103950%	0.143550%	0.181500%	0.223300%	0.264000%	0.315150%	0.385000%
A2	0.005979%	0.038500%	0.122100%	0.189750%	0.256850%	0.320650%	0.390500%	0.455950%	0.540100%	0.660000%
A3	0.021368%	0.082500%	0.198000%	0.297000%	0.401500%	0.500500%	0.610500%	0.715000%	0.836000%	0.990000%
Baa1	0.049500%	0.154000%	0.308000%	0.456500%	0.605000%	0.753500%	0.918500%	1.083500%	1.248500%	1.430000%
Baa2	0.093500%	0.258500%	0.456500%	0.660000%	0.869000%	1.083500%	1.325500%	1.567500%	1.782000%	1.980000%
Baa3	0.231000%	0.577500%	0.940500%	1.309000%	1.677500%	2.035000%	2.381500%	2.733500%	3.063500%	3.355000%
Ba1	0.478500%	1.111000%	1.721500%	2.310000%	2.904000%	3.437500%	3.883000%	4.339500%	4.779500%	5.170000%
Ba2	0.858000%	1.908500%	2.849000%	3.740000%	4.625500%	5.373500%	5.885000%	6.413000%	6.957500%	7.425000%
Ba3	1.545500%	3.030500%	4.328500%	5.384500%	6.523000%	7.419500%	8.041000%	8.640500%	9.190500%	9.713000%
B1	2.574000%	4.609000%	6.369000%	7.617500%	8.866000%	9.839500%	10.521500%	11.126500%	11.682000%	12.210000%
B2	3.938000%	6.418500%	8.552500%	9.971500%	11.390500%	12.457500%	13.205500%	13.832500%	14.421000%	14.960000%
B3	6.391000%	9.135500%	11.566500%	13.222000%	14.877500%	16.060000%	17.050000%	17.919000%	18.579000%	19.195000%
Caa	14.300000%	17.875000%	21.450000%	24.134000%	26.812500%	28.600000%	30.387500%	32.175000%	33.962500%	35.750000%

This figure presents a table that relates credit ratings with the loss rates (loss as fraction of principal) that asset backed securities would be expected to have. The table was used up to the crisis as a reference and it was produced by Moody's. Importantly, Moody's would use this table as part of the risk and valuation analysis, but not as summary statistic that would completely determine its rating. The table is available here https://www.moodys.com/sites/products/productattachments/marvel_user_guide1.pdf

REFERENCES

- ADELINO, M., A. SCHOAR, AND F. SEVERINO (2015): “Loan Originations and Defaults in the Mortgage Crisis: The Role of the Middle Class,” Working Paper.
- AUCLERT, A. (2016): “Monetary Policy and the Redistribution Channel,” Working Paper, Princeton University.
- AUTOR, D. H. AND D. DORN (2013): “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market,” *American Economic Review*, 103, 1553–1597.
- AUTOR, D. H., D. DORN, AND G. H. HANSON (2013): “The China Syndrome: Local Labor Market Effects of Import Competition in the United States,” *American Economic Review*, 103, 2121–2168.
- BARNICHON, R. AND A. FIGURA (2015): “Labor Market Heterogeneity and the Aggregate Matching Function,” *American Economic Journal: Macroeconomics*, 7, 222–249.
- BERAJA, M. (2015): “A Semi-Structural Methodology for Policy Counterfactuals and an Application to Fiscal Unions,” Working Paper, University of Chicago.
- BERAJA, M., E. G. HURST, AND J. OSPINA (2016): “The Aggregate Implications of Regional Business Cycles,” Working Paper 21956, National Bureau of Economic Research.
- BESSEMBINDER, H., W. F. MAXWELL, AND K. VENKATARAMAN (2013): “Trading Activity and Transaction Costs in Structured Credit Products,” *Financial Analysts Journal*, 69, 55–67.
- BILS, M., P. J. KLENOW, AND B. A. MALIN (2014): “Resurrecting the role of the product market wedge in recessions,” Working Paper 20555, National Bureau of Economic Research.
- BRINCA, P. (2013): “Monetary business cycle accounting for Sweden,” *The B.E. Journal of Macroeconomics*, 13.
- (2014): “Distortions in the neoclassical growth model: A cross-country analysis,” *Journal of Economic Dynamics and Control*, 47, 1–19.
- BRINCA, P., V. V. CHARI, P. J. KEHOE, AND E. MCGRATTAN (2016): “Accounting for Business Cycles,” *Handbook of Macroeconomics*.
- BUERA, F. J. AND B. MOLL (2015): “Aggregate Implications of a Credit Crunch: The Importance of Heterogeneity,” *American Economic Journal: Macroeconomics*, 7, 1–42.
- CAGE, R., J. GREENLEES, AND P. JACKMAN (2003): “Introducing the chained consumer price index,” in *International Working Group on Price Indices (Ottawa Group): Proceedings of the Seventh Meeting*, Paris: INSEE, 213–246.

- CARLINO, G. A. AND R. H. DEFINA (1998): “The Differential Regional Effects of Monetary Policy,” *The Review of Economics and Statistics*, 80, 572–587.
- (1999): “The Differential Regional Effects of Monetary Policy: Evidence from the U.S. States,” *Journal of Regional Science*, 39, 339–358.
- (2004): “How strong is co-movement in employment over the business cycle? Evidence from state/sector data,” *Journal of Urban Economics*, 55, 298–315.
- CAVALCANTI, V, T. (2007): “Business Cycle and Level Accounting: The Case of Portugal,” *Portuguese Economic Journal*, 6, 47–64.
- CHAHROUR, R., S. CHUGH, A. FINKELSTEIN SHAPIRO, AND A. LARIAU (2016): “The Labor Wedge: A Search and Matching Perspective,” Working Paper.
- CHAKRABORTY, S. AND K. OTSU (2013): “Business cycle accounting of the BRIC economies,” *The B.E. Journal of Macroeconomics*, 13.
- CHARI, V., P. KEHOE, AND E. R. MCGRATTAN (2004): “Business Cycle Accounting,” Working Paper 10351, National Bureau of Economic Research.
- (2006): “Appendices: Business Cycle Accounting,” Research Department Staff Report 362, Federal Reserve Bank of Minneapolis.
- (2007a): “Business Cycle Accounting,” *Econometrica*, 75, 781–836.
- (2007b): “Comparing Alternative Representations, Methodologies, and Decompositions in Business Cycle Accounting,” Research Department Staff Report 647, Federal Reserve Bank of Minneapolis.
- CHARLES, K. K., E. HURST, AND M. J. NOTOWIDIGDO (2016): “Housing Booms, Manufacturing Decline, and Labor Market Outcomes,” Working Paper.
- COBLE, D. AND L. BUENAVENTURA (2016): “The Labor Wedge: New Facts based on US Microdata,” Working Paper.
- COCIUBA, S., E. PRESCOTT, AND A. UEBERFELDT (2012): “US hours and productivity behavior using CPS hours worked data: 1947-III to 2011-IV,” *Unpublished paper, University of Western Ontario*.
- COCIUBA, S. E. AND A. UEBERFELDT (2015): “Heterogeneity and long-run changes in aggregate hours and the labor wedge,” *Journal of Economic Dynamics and Control*, 52, 75–95.
- COIBION, O., Y. GORODNICHENKO, AND G. H. HONG (2015): “The Cyclicalities of Sales, Regular and Effective Prices: Business Cycle and Policy Implications [†],” *American Economic Review*, 105, 993–1029.

- EICHENBAUM, M. S., L. P. HANSEN, AND K. J. SINGLETON (1988): “A Time Series Analysis of Representative Agent Models of Consumption and Leisure Choice under Uncertainty,” *The Quarterly Journal of Economics*, 103, 51.
- FEENBERG, D. AND E. COUTTS (1993): “An Introduction to the TAXSIM Model,” *Journal of Policy Analysis and Management*, 12, 189.
- FLOOD, S., M. KING, S. RUGGLES, AND J. R. WARREN (2015): “Integrated Public Use Microdata Series, Current Population Survey: Version 4.0,” [Machine-readable database], Minneapolis: University of Minnesota.
- FRANCIS, N., M. T. OWYANG, AND T. SEKHOSYAN (2012): “The Local Effects of Monetary Policy,” *The B.E. Journal of Macroeconomics*, 12.
- FRATTO, C. AND H. UHLIG (2014): “Accounting for Post-Crisis Inflation and Employment: A Retro Analysis,” Working Paper 20707, National Bureau of Economic Research.
- GALI, J., M. GERTLER, AND J. D. LOPEZ-SALIDO (2007): “Markups, Gaps, and the Welfare Costs of Business Fluctuations,” *The Review of Economics and Statistics*, 89, 44–59.
- GELMAN, A. AND E. LOKEN (2014): “Ethics and Statistics: The AAA Tranche of Subprime Science,” *Chance*, 27, 51–56.
- HAGEDORN, M., F. KARAHAN, I. MANOVSKII, AND K. MITMAN (2016): “Unemployment Benefits and Unemployment in the Great Recession: The Role of Macro Effects,” Working Paper.
- HALL, R. E. (1988): “Intertemporal Substitution in Consumption,” *Journal of Political Economy*, 96, 339–357.
- (1997): “Macroeconomic Fluctuations and the Allocation of Time,” *Journal of Labor Economics*, 15, S223–S250.
- (2009): “Reconciling cyclical movements in the marginal value of time and the marginal product of labor,” *Journal of political Economy*, 117, 281–323.
- HIRSCH, B. T. AND D. A. MACPHERSON (2003): “Union membership and coverage database from the current population survey: Note,” *Industrial & labor relations review*, 56, 349–354.
- KARABARBOUNIS, L. (2014a): “Home Production, Labor Wedges, and International Business Cycles,” *Journal of Monetary Economics*, 64, 68–84.
- (2014b): “The Labor Wedge: MRS vs. MPN,” *Review of Economic Dynamics*, 17, 206–223.

- KEHOE, P., V. MIDRIGAN, AND E. PASTORINO (2016): “Debt Constraints and the Labor Wedge,” *The American Economic Review*, 106, 548–553.
- KERSTING, E. K. (2008): “The 1980s recession in the UK: A business cycle accounting perspective,” *Review of Economic Dynamics*, 11, 179–191.
- KEYS, B. J., T. MUKHERJEE, A. SERU, AND V. VIG (2009): “Financial regulation and securitization: Evidence from subprime loans,” *Journal of Monetary Economics*, 56, 700–720.
- (2010): “Did Securitization Lead to Lax Screening? Evidence from Subprime Loans,” *Quarterly Journal of Economics*, 125, 307–362.
- KEYS, B. J., A. SERU, AND V. VIG (2012): “Lender Screening and the Role of Securitization: Evidence from Prime and Subprime Mortgage Markets,” *Review of Financial Studies*, 25, 2071–2108.
- MACALUSO, C. (2016): “Right person, wrong place? Cities’ skill structure and post-layoff employment outcomes,” Working Paper, University of Chicago.
- MARINESCU, I. AND R. RATHELOT (2014): “Mismatch Unemployment and the Geography of Job Search,” Working Paper, University of Chicago and University of Warwick.
- MEHROTRA, N. AND D. SERGEYEV (2015): “Financial Shocks and Job Flows,” Working Paper, Brown University Working Paper.
- MIAN, A., K. RAO, AND A. SUFI (2013): “Household Balance Sheets, Consumption, and the Economic Slump,” *The Quarterly Journal of Economics*, 128, 1687–1726.
- MIAN, A. AND A. SUFI (2014): “What Explains the 2007-2009 Drop in Employment?: The 2007-2009 Drop in Employment,” *Econometrica*, 82, 2197–2223.
- MIDRIGAN, V. AND T. PHILIPPON (2016): “Household Leverage and the Recession,” Working Paper.
- MONDRAGON, J. (2015): “Household Credit and Employment in the Great Recession,” Working Paper, University of California, Berkeley.
- MULLIGAN, C. (2002): “A Century of Labor-Leisure Distortions,” Working Paper 8774, National Bureau of Economic Research.
- (2012): *The Redistribution Recession: How Labor Market Distortions Contracted the Economy*, Oxford University Press.
- OHANIAN, L. E. (2010): “The Economic Crisis from a Neoclassical Perspective,” *Journal of Economic Perspectives*, 24, 45–66.

- OHANIAN, L. E. AND A. RAFFO (2012): “Aggregate hours worked in OECD countries: New measurement and implications for business cycles,” *Journal of Monetary Economics*, 59, 40–56.
- OTSU, K. (2010): “A Neoclassical Analysis of the Asian Crisis: Business Cycle Accounting for a Small Open Economy,” *The B.E. Journal of Macroeconomics*, 10, 1–39.
- PIKETTY, T. AND G. ZUCMAN (2015): “Wealth and Inheritance in the Long Run,” in *Handbook of Income Distribution*, Elsevier, vol. 2, 1303–1368.
- PISKORSKI, T., A. SERU, AND J. WITKIN (2015): “Asset Quality Misrepresentation by Financial Intermediaries: Evidence from the RMBS Market: Asset Quality Misrepresentation by Financial Intermediaries,” *The Journal of Finance*, 70, 2635–2678.
- RAVN, M. O. AND H. UHLIG (2002): “Notes: On Adjusting The Hodrick-Pescott Filter For The Frequency Of Observations,” *Review of Economics & Statistics*, 84, 371–380.
- RODRIGUEZ-LOPEZ, J. AND M. SOLIS GARCIA (2016): “Accounting for Spanish Business Cycles,” *Macroeconomic Dynamics*, 20, 685–714.
- ROGERSON, R. AND J. WALLENIUS (2009): “Micro and macro elasticities in a life cycle model with taxes,” *Journal of Economic Theory*, 144, 2277–2292.
- RUGGLES, S., K. GENADEK, J. GROVER, AND M. SOBEK (2015): “Integrated Public Use Microdata Series: Version 6.0,” [Machine-readable database], Minneapolis: University of Minnesota.
- SAEZ, E. AND G. ZUCMAN (2016): “Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data,” *The Quarterly Journal of Economics*, qjw004.
- SAHIN, A., J. SONG, G. TOPA, AND G. L. VIOLANTE (2014): “Mismatch Unemployment,” *American Economic Review*, 104, 3529–3564.
- SCHMITT-GROHÉ, S. AND M. URIBE (2003): “Closing Small Open Economy Models,” *Journal of International Economics*, 61, 163–185.
- SHIMER, R. (2009): “Convergence in macroeconomics: The labor wedge,” *American Economic Journal: Macroeconomics*, 280–297.
- SIMONOVSKA, I. AND L. SÖDERLING (2015): “Business Cycle Accounting for Chile,” *Macroeconomic Dynamics*, 19, 990–1022.
- STROEBEL, J. AND J. VAVRA (2016): “House Prices, Local Demand, and Retail Prices,” Working Paper, NYU, University of Chicago.
- SUÁREZ SERRATO, J. C. AND O. ZIDAR (2016): “Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms,” *American Economic Review*, 106, 2582–2624.

- UHLIG, H. (1995): “A Toolkit for Analyzing Nonlinear Dynamic Stochastic Models Easily,” Discussion Paper Vol. 1995-97., CentER.
- YAMARIK, S. (2013): “State-Level Capital and Investment: Updates and Implications,” *Contemporary Economic Policy*, 31, 62–72.
- ŠUSTEK, R. (2011): “Monetary business cycle accounting,” *Review of Economic Dynamics*, 14, 592–612.