DISSERTATION

ESSAYS ON REGIONAL ECONOMIC DEVELOPMENT: FINANCE, GROWTH, AND INCOME DISTRIBUTION

Submitted by

Luke Petach

Department of Economics

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Spring 2019

Doctoral Committee:

Advisor: Anita Alves Pena
Co-Advisor: Daniele Tavani

Stephan Weiler
Susan Opp

ABSTRACT

ESSAYS ON REGIONAL ECONOMIC DEVELOPMENT: FINANCE, GROWTH, AND INCOME DISTRIBUTION

This dissertation explores the relationship between financial markets and economic activity across regions within the United States. In particular, it examines how regional variation in financial markets may affect variables that play an important role in regional economic growth and income distribution. Despite the way the Great Recession of 2007-2009 made manifest the importance of financial markets in the determination of real economic outcomes, there is still much that is not understood about the relationship between financial markets and the rest of the economy. Financial markets may affect household decision making either indirectly via their influence over existing institutional arrangements and social norms, or directly via their influence over who can obtain access to the credit needed to start a business, purchase a home, or invest in human capital. These changes will alter the distribution of income—both directly via the re-distribution of funds through credit markets, and indirectly via the outcomes that access to credit (or lack there-of) engender.

ACKNOWLEDGEMENTS

I would first and foremost like to thank my wife, Josie, without whom I could not have completed this dissertation. I cannot fully express how grateful I am for all her love and support. Thank you Josie, I love you beyond words. I would also like to thank my parents, who have done everything they can to help further my education. Thank you for loving me unconditionally.

A big thanks goes out to my entire committee, who—no matter the circumstances—made time for my countless office visits, emails, questions, and concerns. Thank you Anita, Daniele, and Stephan. I could not have asked for a better set of mentors.

TABLE OF CONTENTS

ABSTRACT ii
ACKNOWLEDGEMENTS iii
LIST OF TABLES vi
LIST OF FIGURES vii

Chapter 1 Introduction 1
1.1 Financialization and the Geography of Banking 3
1.2 A Regional Approach to Wage- vs. Profit-Led Growth 4

Chapter 2 Local Financialization, Household Debt, and the Great Recession 6
2.1 Introduction 6
2.2 Finance, Debt, and Regional Economic Performance 8
2.3 Data and Estimation Strategy 13
2.4 Results 24
2.4.1 State-level Results 24
2.4.2 County-level Results 38
2.4.3 Sub-Sector Analysis 39
2.4.4 Distributional Implications 42
2.5 Conclusion 46

Chapter 3 Geography Matters: The Impact of Geographic Expansion on Bank Performance During the Great Recession 48
3.1 Introduction 48
3.2 Literature Review 51
3.3 A Model of Bank Branching and Systemic Risk 52
3.4 Data and Methods 55
3.5 Results 61
3.5.1 Robustness Checks 67
3.5.2 Discussion 67
3.6 Conclusion 72

Chapter 4 Distribution and Capacity Utilization in the United States: Evidence from State-level Data 74
4.1 Introduction 74
4.2 Demand and Distribution, Revisited 75
4.3 Data 82
4.4 Estimation 87
4.4.1 Demand 87
4.4.2 Demand: Sensitivity Analysis 92
4.4.3 Distribution 95
4.4.4 Discussion 98

4.5 Conclusion 103

Bibliography 106

Appendix A Appendix to Chapter 2 117
A.1 County-level Sample Means 117
A.2 State-level Continuous Treatment 117
A.3 County-level Discrete Treatment 119

Appendix B Appendix to Chapter 4 121
B.1 Stability Analysis 121
B.2 Hamilton (2017) Filter 122
B.3 Non-Parametric Techniques 123
B.4 Personal Income Inequality and the Demand Regime 124

LIST OF TABLES

2.1 Sample Means, State-level Data: 1999-2013 . . . 14
2.2 High-Finance States . . . 17
2.3 Estimation Results - Specification 1 . . . 25
2.4 Estimation Results - Specification 2 . . . 26
2.5 Estimation Results - Specification 3 . . . 27
2.6 Robustness Checks . . . 30
2.7 Dose Responsiveness Test . . . 32
2.8 Placebo Test . . . 33
2.9 Serial Correlation Adjustments . . . 35
2.10 Difference-in-Discontinuities Estimates . . . 36
2.11 County-level Results . . . 38
2.12 Sub-Sector Results . . . 40
2.13 Correlations Between Overall and Sub-Sector Treatment Variables . . . 42

3.1 Summary Statistics, *FDIC* Data, 1999-2011 . . . 57
3.2 Estimation Results 1 . . . 62
3.3 Estimation Results 2 . . . 64
3.4 Robustness Checks . . . 68

4.1 Sample Means . . . 83
4.2 Estimation Results - *NF* Specification . . . 88
4.3 Estimation Results - Minimum Wage Instruments . . . 90
4.4 Estimation Results - Demand Nullcline Sensitivity Analysis . . . 93
4.5 Estimation Results - Employment Rate Specification . . . 94

A.1 Sample Means, County-level Data . . . 117
A.2 Estimation Results: Continuous Treatment, State-level . . . 118
A.3 Continuous Treatment, State-level Employment Share . . . 119
A.4 Estimation Results, County-level Data, Discrete Treatment . . . 120

B.1 Estimation Results - Personal Income Inequality and the Demand Regime . . . 125

LIST OF FIGURES

2.1 Changes in State Financial Markets: 1963-2014 9
2.2 The Growth of Consumption and Consumer Credit in the United States: 1975-Present . 10
2.3 Composition of Per-Capita Borrowing in the United States: 1999-Present 11
2.4 Per-Capita Debt v. FIRE/GDP ratio 15
2.5 Trend Plots: Debt Per-Capita by Relative Size of the Local Financial Sector, 1999-2013 19
2.6 County Debt-to-Income Ratios, by Level of Financialization: 2001-2007 22
2.7 Debt-to-Income Ratios, by Own-County and Spatially-Lagged Financialization 24
2.8 Randomization Test - Placebo Distribution and True Effect 34
2.9 Estimated Share of Households with Positive Debt, by Income Percentile 43
2.10 Estimated Average Household Debt-to-Income Ratio, by Income Percentile 44
2.11 Estimated Treatment Effect - by Income Percentile 46

3.1 Geographic Variation in Bank Performance, 1966-2017 49
3.2 Geographic Dispersion and Banking Consolidation 57
3.3 Bank Performance Measures, by Extent of Geographic Expansion 59
3.4 Marginal Effects, from Estimating Equation (3.15) 66

4.1 Profit-Led, Profit-Squeeze Dynamics in the Goodwin Model 77
4.2 Comparison of Hamilton (2017) and HP Filter. Capacity Utilization and the Labor Share: 1974-2014. 83
4.3 Mapping the Labor Share Across U.S. States: 1974 85
4.4 Mapping the Labor Share Across U.S. States: 2014 86
4.5 Non-Parametric Estimates of the Distributive Curve: Pooled Sample 96
4.6 Non-Parametric Estimates of the Distributive Curve — State-by-State (1) 97
4.7 Non-Parametric Estimates of the Distributive Curve — State-by-State (2) 97
4.8 Non-Parametric Estimates of the Distributive Curve — State-by-State (3) 98
4.9 Demand and Distribution System with Wage-Led Demand and Non-Linear Distributive Curve 100
4.10 State Tax Rates 102

Chapter 1

Introduction

The traditional view of finance sees markets for credit and insurance as allowing individuals to smooth consumption optimally over time (Deaton, 1991, 1992; Friedman, 1957), to make costly human capital investments (Bardhan & Udry, 1999), or for firms to obtain the liquidity necessary to make new purchases of physical capital equipment (Fazzari, Hubbard, and Peterson, 1988). I refer to these three features as the “consumption smoothing view" of credit, as each is concerned with substituting consumption across periods. Additionally, asset markets—as long as monopolistic access to information is restricted—are thought to satisfy the “efficient market hypothesis,"—that is, asset markets are thought of as markets in which prices fully reflect all relevant information (Fama, 1969).

The consumption smoothing view argues in the absence of perfect risk pooling, sufficient entrepreneurial capital, or education funding, financial markets can be Pareto-improving by providing insurance and credit to households, linking savers to borrowers, and pushing households toward their optimal consumption path. This view is complicated when models are allowed to account for asymmetric information. The introduction of moral hazard or adverse selection may result in credit rationing, increased interest rates, and increased default risk (Stiglitz & Weiss, 1981). In this case, the competitive equilibrium allocation is not socially optimal, and government intervention in the market for credit may be Pareto-improving[1]. Additionally, the models considered in the consumption smoothing view rarely offer explicit consideration of household debt—a factor important both for its macroeconomic implications and the role it plays in (potentially) constraining household consumption as the cost of debt servicing increases. The consumption smoothing view is thus best thought of as an accurate—albeit partial—description of the role played by financial markets in the real economy.

[1]Although an imperfect credit market—in the consumption smoothing view—may be Pareto-superior to an equilibrium with no credit at all.

In contrast to the consumption smoothing view, Post Keynesian economics posits a macroeconomic theory of financial markets in which credit markets may be necessary, but–at times—inimical to economic growth and an equal distribution of income. Keynes (1936) argued that the separation of ownership and control facilitated by modern stock markets creates incentives which undermine the ability of private investment decisions to obtain a socially optimal outcome. He writes: "[T]here is no clear evidence from experience that the investment policy which is socially advanageous coincides with that which is most profitable" (157). More recently, work by Post Keynesian authors is concerned with the phenomena of financialization, or “the increasing role of financial motives, financial markets, financial actors and financial institutions in the operation of domestic and international economies" (Epstein, 2005, p.3). The characteristics of a financialized economy include: a declining corporate retention ratio, a decline in the rate of new equity issue, an increase in the pressure to “keep up with the Joneses," captured by a rise in personal consumption expenditure relative to personal disposable income, and an increase in household borrowing (Palley, 2007; Skott and Ryoo, 2008; Wisman, 2013).

The last characteristic—an increase in household borrowing—is of fundamental importance. It is an increase in household leverage through which some authors have linked financial markets to regional variation in unemployment during the Great Recession. Mian and Sufi (2010) make the case that county-level differences in household leverage ratios can explain the difference in severity of the Great Recession across regions within the United States. However, the authors offer little explanation—other than noting the existence of a “credit-supply shock"—for why borrowing increased in some counties more than others. Indeed, the consumption smoothing view of credit in general offers little in the way of explanation for why households residing in economically similar counties should have vastly different borrowing behavior. Chapter 2 of this dissertation takes a Post Keynesian perspective on this issue. Chapter 2 argues the “credit-supply shock" identified by Mian and Sufi (2010, 2011, 2012) can be explained by a number of deliberate shifts in federal policy that preceded the housing bubble in the United States. Chapter 2 then applies a difference-in-differences analysis to a consistent panel series for household debt to make the case that household borrowing increased more during the U.S. housing market boom in states and counties that were highly financialized, relative to less financialized areas.

1.1 Financialization and the Geography of Banking

The incentives faced by agents on the demand-side of credit market transactions (i.e. households and firms) are not the only important variable affected by financialization. The structure of financial markets—including the prevailing type and number of banks—is itself subject to change. As the institutions governing financial markets shift in response to the growing influence of financial actors, the conditions for successful operation of a financial enterprise are altered. Two primary examples of institutional change in financial markets in response to financialization are the recent, rapid de-regulation of banking and other forms of financial enterprise and the proliferation of financial engineering and new financial instruments (Brown, 2007; Wisman, 2013).

One effect of these changes has been an increase in banking consolidation. The passage of the Riegle-Neal Act in 1994 allowed for the first time the possibility of interstate branch banking. Immediately following this legislation the number of interstate mergers of large banks increased dramatically (Kwan, 2004). As of 2016, the "big four" U.S. banks (JP Morgan Chase, Bank of America, Wells Fargo, and Citigroup) held approximately 40% of all U.S. consumer deposits (*FDIC Bank Branch Data*, Author's Calculations). At the same time, consolidation brings about important changes in the spatial structure of banking. In particular, the past three decades have been characterized by a dramatic increase in the geographic concentration of banking activity. As bank headquarters become more geographically concentrated, the network of large bank activities becomes spread out over an increasingly large geographic footprint.

The aforementioned changes are not welfare neutral. If large, national banks and small, local banks adopt different organizational strategies for solving the typical asymmetric information problem inherent in credit markets, this will alter who gets access to credit. If small, local banks have access to soft-information via a "relationship lending" channel then the presence of these banks may extend the availability of credit to groups otherwise under-served by large banks. If large banks can exploit returns to scale in screening and monitoring then the resulting cost reductions may induce more lending than would otherwise occur at small banks.

Additionally, if the information attainable through relationship lending helps small banks avoid otherwise hard-to-detect risks, then an increase in the geographic concentration of bank headquarters may increase the amount of information asymmetry in the banking system. As a result, even if the simultaneous concentration of bank headquarters and expansion of bank branch networks reduces idiosyncratic risk in the short-run by allowing large banks to diversify lending portfolios geographically over a larger pool of assets, in the long-run it may increase systemic risk and make large banks more vulnerable to negative macroeconomic shocks. Chapter 3 of this dissertation tests the impact of geographic branch network expansion on the vulnerability of banks to systemic risk. Using the Great Recession as a source of variation in bank performance, I find that banks with larger geographic networks had larger performance declines during the Great Recession, even after controlling for other size-related measures of banking activity.

1.2 A Regional Approach to Wage- vs. Profit-Led Growth

Since Bhaduri and Marglin (1990) demonstrated the use of more general investment functions undermines the “wage-led" characteristic of Kaleckian models, a topic of perennial interest among economists working on growth and distribution has been the type of “growth regime" that characterizes an economy. An economy is said to be wage-led, “stagnationist," or “underconsmptionist" if a redistribution toward wages increases economic activity, and profit-led or “exhilarationist" if the opposite is true. The theoretically ambiguous results of Bhaduri and Marglin (1990) spawned a cottage industry of practitioners attempting to empirically assess whether certain economies are wage- or profit-led, leading to numerous and at times conflicting results (Barbosa-Filho and Taylor, 2006; Nikiforos and Foley, 2012; Rada and Kiefer, 2015; Stockhammer, 2017). One problem with existing studies is a lack of satisfactory econometric identification (Nikiforos and Foley (2012) is perhaps, the exception—although their approach is not without flaws). Chapter 4 addresses this problem by adopting a regional approach to the question of wage- vs. profit-led growth for the

United States. In particular, Chapter 4 exploits variation in minimum wage policy across U.S. states as a instrumental variable for the labor share as a means of identifying the type of growth regime present at a regional level. I find that states appear wage-led on average, but the simultaneous non-linearity of the distributive curve suggests difficulties for wage-led growth akin to a coordination problem.

Chapter 2

Local Financialization, Household Debt, and the Great Recession

2.1 Introduction

The common explanation for heterogeneous geographic responses to the 2007-2009 recession is rooted in differing levels of debt-financed economic activity across regions in the pre-recession period (Fort et al., 2013; Mian and Sufi, 2012). One strand of this literature focuses on a decline in debt-financed consumption, particularly in the non-tradeables sector, for U.S. counties with high pre-recession leverage ratios (Mian and Sufi, 2010, 2011, 2012). Another strand of literature focuses on the decline in debt-financed investment—particularly small business investment and new entrepreneurship facilitated by home-equity loans (Fort et al., 2013; Pressman and Scott, 2017). Whether one thinks it is debt-financed consumption or investment that explains the differential geographic response to the recession, neither the debt-financed consumption story nor the debt-financed investment story offer an adequate explanation for the rapid increase in household leverage in the pre-recession period. Mian and Sufi (2010) make note of a “credit-supply shock" linked to securitization in the sub-prime mortgage market, but are admittedly agnostic as to its source, chalking it up to “government programs, moral hazard on behalf of originators...and the enormous capital inflows into the United States" (p. 79).

If the increase in the supply of credit that occurred prior to the recession was the result of macroeconomic shocks (i.e. securitization, capital inflows) there is no *a priori* obvious reason why changes in household borrowing in response to these shocks should differ systematically across states and counties within the United States. The source of differences in household borrowing across otherwise similar regions matters if economists wish to prevent similar crises in the future. Without a clear explanation for why borrowing during the pre-crisis period differed across regions in such a fashion, the narrative regarding the Great Recession is incomplete. In this chapter, I provide one such explanation: I link changes in regional per-capita indebtedness in the pre-recession period to variation in the relative size of the local financial sector. Evidence on the growth of intra-financial sector lending, the negative relationship between credit stocks and economic growth, the increased incidence of financial crises in response to “financial deepening," and the relationship between financial innovation and consumer borrowing all point to the influence of the local financial sector as a potential source of geographic heterogeneity in household borrowing behavior (Bezemer et al., 2016; Livshits et al. , 2016; Montecino et al., 2016; Rousseau and Wachtel, 2011). Changes in the dynamics of the relationship between the financial sector and the rest of the economy—such as the establishment of a norm of financialized corporate governance and the interaction between credit constraints and consumption habits—may serve to exacerbate this effect (Admanti, 2017; Pintus and Wen, 2013).

Further evidence suggesting a link between local financial markets and household borrowing arises from the literature on financialization—the growing influence of financial actors and institutions on outcomes in the real economy. Financialization has been conceptually linked to growing household debt by several authors (Epstein, 2005; Mason et al., 2018; Palley, 2007; Skott and Ryoo, 2008), but little empirical work has been done to test this conjecture. Mason et al. (2018) argue that by recasting social ties as financial claims, the growth of financial markets re-orients household behavior toward market income over non-market activities, increasing the pressure to borrow. Palley (2007) claims that the defining feature of financialization in the United States has been an increase in the volume of debt. In either case, what is missing is an empirical link between the growth of financial markets and household borrowing.

To test the relationship between the size of the local financial sector and household borrowing, I exploit the credit-supply shock that occurred prior to the Great Recession as a plausibly exogenous source of variation in borrowing behavior. Using a difference-in-differences set-up, I find that per-capita indebtedness increased by several thousand dollars—approximately 10%—more in states with a large financial sector (as a share of state GDP) as a result of the credit-supply shock. I then use data on county-level leverage ratios from Mian, Sufi, and Rao (2013), and county-level FIRE-sector employment from the County Business Patterns (CBP) database, to show that a similar, positive relationship between financialization and indebtedness holds at the county level. Finally, I use data on county-to-county commuting flows to construct a spatially-lagged measure of county-level financialization, which I use to demonstrate the existence of between-county spillover effects of financialization on debt. Taken together, these results suggest the growing influence of financial markets and institutions on local economic activity as a fundamental cause of the Great Recession.

The rest of this chapter is organized as follows. Section (2.2) discusses the relationship between finance, consumer borrowing, and the rest of the economy. Section (2.3) describes the data and outlines the empirical model. Section (2.4) discusses the results. Section (2.5) concludes.

2.2 Finance, Debt, and Regional Economic Performance

How has the average importance of the financial sector at the state level varied over time within the United States? There are two intuitive ways this might be measured. First, the share of the finance, insurance, and real-state (FIRE) sector in state GDP gives a sense of how important the financial sector is, relative to the rest of the economy, for a particular region. Second, the importance of the financial sector might be measured by the type of income received by households. In particular, the relative size of incomes from financial sources—rent, interest, and dividend payments—relative to the non-financial compensation paid to households working in non-financial businesses. Figure (2.1) presents a time-series plot of the average value for both series for all U.S. states from 1963 to 2014.

In either case it is clear that the importance of financial markets has increased over time. On average the FIRE sector now claims close to a 20% share in state GDP, compared with 13% at the beginning of the series. The relative importance of rentier incomes seems to have undergone an institutional shift in the late 1970's and early 1980s, with relative stability both before and after. This shift accords with the timing of the “shareholder value" revolution. Following Friedman (1970)'s dictum that “[t]he social responsibility of a business is to increase its profits," a regime

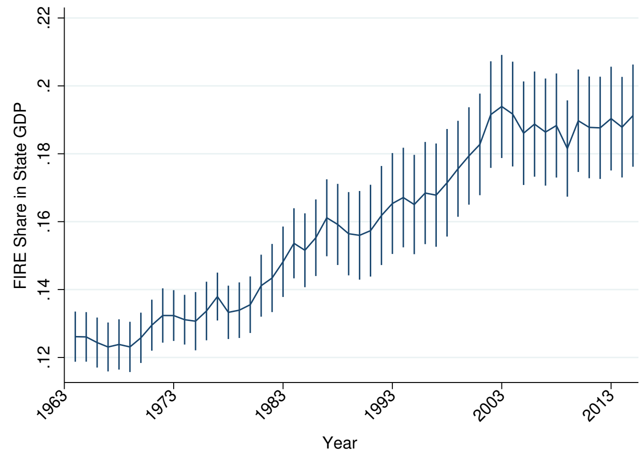

(a) Share of FIRE Sector in State GDP

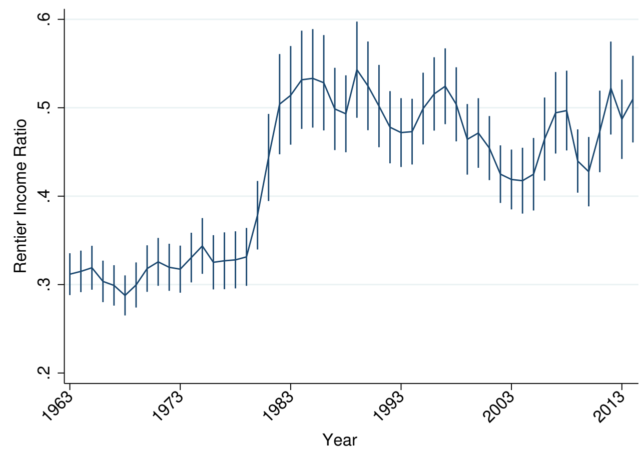

(b) Ratio of Rents, Interest, Dividends to the Non-Financial Compensation Paid to the Non-Financial Sector

Figure 2.1: Changes in State Financial Markets: 1963-2014

Source: BEA Regional Economic Accounts. Plot shows the average value for all fifty U.S. states. Vertical lines are 95% confidence intervals.

of financialized corporate governance came to prominence in the late 1970's and early 1980's (Admanti, 2017). The shareholder value revolution tied managerial compensation to measures of shareholder value such as stock prices and return on equity. At S&P 500 companies the share of corporate executive pay tied to performance measures grew from 16 to 47 percent between the 1970's and 1990's (Reda et al., 2017). These changes precipitated a fall in corporate retention ratios and rates of equity issue as management strategy shifted from 'retain and invest' to 'downsize and distribute' in an effort to improve pay-linked performance measures (Lazonick and O'Sullivan, 2000; Skott and Ryoo, 2008). The result of these changes was an increase in the size of financial payouts to shareholders, as captured by Figure (2.1b).

The increased importance of financial markets in the United States occurred simultaneously with increased access to consumer credit on behalf of households and a rising ratio of consumption to disposable income for the entire U.S. economy. Figure (2.2) displays the evolution of both the total amount of outstanding consumer credit and the ratio of real personal consumption expenditures to disposable income from 1975 to the present.

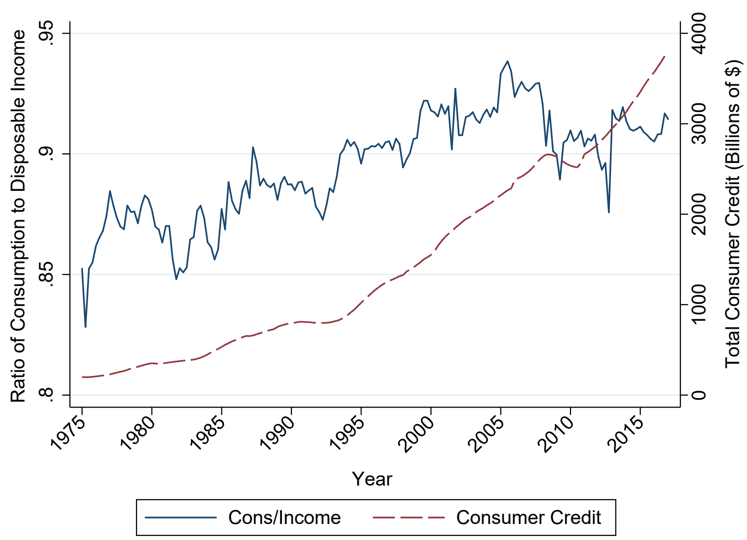

Figure 2.2: The Growth of Consumption and Consumer Credit in the United States: 1975-Present

Source: Federal Reserve Economic Data. Figure displays the ratio of real personal consumption expenditures to real disposable income and the total value of all outstanding consumer credit, owned and securitized.

The growth of personal consumption expenditures and total consumer credit prompts further questions. Does the increase in total consumer credit reflect an increase in per-capita borrowing? Has the increase in consumer credit been to finance short-term consumption, or does it reflect an increase in investment (or the cost of investment) in human capital or durable goods such as housing? Figure (2.3) sheds light on these questions by plotting trends in nominal per-capita borrowing in the years preceding the financial crisis and recession.

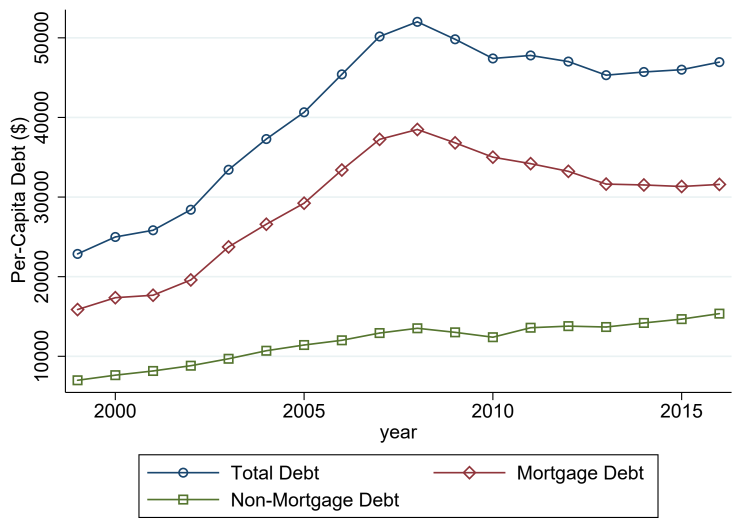

Figure 2.3: Composition of Per-Capita Borrowing in the United States: 1999-Present

Source: Federal Reserve Bank of New York, Economic Data. Figure displays the trends for total debt per-capita, housing debt per-capita, and non-housing debt per-capita.

The figure makes clear that per-capita borrowing in the United States has increased since 1999. The driving force behind this increase up until 2009 was the expansion of mortgage debt. After the financial crisis and recession total borrowing per-capita initially tapered off, but has since begun to increase again—this time due to a rise in non-housing debt. Recent changes in per-capita indebtedness have been largely due to increases in both student and auto debt. Student debt per-capita increased from $530 in 1999 to $4,920 in 2016. Similarly, auto debt per-capita increased from $1,830 in 1999 to $4,340 in 2016. In contrast, credit card debt per-capita has remained fairly constant, taking values of $2,370 in 1999 and $2,930 in 2016, respectively. Housing debt per-capita has also been fairly constant since 2013, leveling off around $31,500.[2]

What is not clear from Figure (2.3) is the source of the increase in consumer borrowing. The literature on the relationship between financial markets and the real economy suggests that the increasing influence of the financial sector itself may be the cause of the increase in the level of per-capita indebtedness in the United States. When financial markets are small, and the number of investment opportunities limited, the ability of households to access credit is limited by the liquidity preference of wealth owners. Because mortgage contracts and consumer receivables were historically illiquid—in the sense that loans were held by the originator—the supply of credit was limited by lender tolerance for illiquidity (Brown, 2007). The creation of markets for asset-backed securities makes liquid a formerly illiquid investment and increases the supply of credit available to households. On the one hand, this financial deepening may be beneficial to both consumers and investors—providing a way around liquidity constraints, increased access to mortgage financing, and an increased menu of profitable investment options. On the other hand, recent evidence suggests that excessive bank lending, over extension of mortgage credit, increases in the size of the credit stock, and growth in the total indebtedness of households may be related to increased economic instability, slower economic growth, and lower consumer welfare (Arcand et al., 2015; Bezemer et al., 2016; Dunn and Mirzaie, 2016; Rousseau and Wachtel, 2011).

Assuming that the share of asset-backed securities in financial portfolios was constant or increasing in the pre-recession period, an increase in the size of the local financial sector relative to the rest of the economy should—*ceteris paribus*—increase local financial sector demand for debt linked to such instruments. If financial markets cause households to prioritize market incomes over non-market activities by re-casting social ties as financial claims (Mason et al., 2018), then local financialization may result in increased pressure to borrow. Mason et al. (2018) point to student debt as the most prominent contemporary instance of a social tie—education—being recast as a financial claim, captured in the shift from direct public provisioning of education to provisioning

[2] *Source*: Federal Reserve Bank of New York, Economic Data.

based on debt-financing. However, the authors note that even activities innocuous as pet adoption may now involve the creation of a tradeable financial asset. The process of financialization erects barriers to social inclusion that can only be overcome by participating in financial markets, often via acquisition of financial liabilities (e.g. student debt). To the extent that households want or need to participate in these activities, they will experience increased pressure to borrow.

The combined effect of increased financial sector demand for assets created from household debt, and increased pressure to borrow created by the re-casting of social activities as financial claims, suggests that the macroeconomic credit-supply shock identified by Mian and Sufi (2010) should have differential effects on borrowing across regions, depending on the size of the local financial sector. This effect may be amplified by increased intra-financial sector lending and the formation of “consumption habits" related to household efforts to “keep up with the Joneses" (Montecino et al., 2016; Pintus and Wen, 2013; Wisman, 2013). The relationship between the size of local financial markets and demand for assets based on consumer debt on the one hand, and the recasting of social ties as financial assets on the other, provide explicit links between a more general process of financialization—the growing influence of financial markets, actors, and institutions on the real economy (Epstein, 2005; Lechevalier et al., 2017; Mason et al., 2018; Palley, 2007; Sawyer, 2013; Skott and Ryoo, 2008)—and household debt.

2.3 Data and Estimation Strategy

The description of the data used in this chapter is divided into two sections, arranged by level of geographic aggregation. I begin by looking at the state-level data, then move to the data that is organized at the county level. The state-level data in this chapter come from two main sources. Data on per-capita indebtedness at the state level comes from the Federal Reserve Board of New York/Equifax Consumer Credit Panel. Data on other state-level characteristics, including the size of the local financial sector, come from the Bureau of Economic Analysis (*BEA*) Regional Economic Accounts. The sample covers all 50 U.S. states and Washington D.C.. Time is measured in years and spans the period 1999-2013. The size of the local financial sector is measured as the share of the finance, insurance, and real-estate sector in state GDP. Sample means for this and all other relevant variables are included in Table (2.1).

Table 2.1: Sample Means, State-level Data: 1999-2013

Variable	Mean	Std. Dev.
$\frac{Finance}{GDP}_{it}$ (%)	18.8	5.47
Total Debt$_{it}$	\$30,962	10,197
Non-Mortgage Debt$_{it}$	\$9,100	1,463
$\frac{Debt}{Income}_{it}$	1.14	0.28
GDP Per-Capita$_{it}$	\$26,857	4,322
Consumption$_{it}$	\$24,946	3,712
$\%\ \Delta$ GDP$_{it}$	1.9	2.8
Labor Share$_{it}$ (%)	58.6	4.6
Gini$_{it}$	59.6	3.6
Population$_{it}$	5,848,258	6,529,520
Manufacturing Share$_{it}$ (%)	8.2	3.5
N	765	

Notes: Sample means for all 50 U.S. states and Washington D.C.. Data on debt, income, and consumption are in constant per-capita 1999 dollars. Manufacturing Share$_{it}$ reports the number of individuals employed in manufacturing as a share of total employment.

Both total debt and non-mortgage debt are measured in constant 1999 dollars. Total debt per-capita includes mortgage debt, auto debt, credit-card debt, and student loan debt. “GDP Per-Capita$_{it}$" and “Consumption$_{it}$" measure real per-capita disposable income and consumption in state i, at time t, in constant 1999 dollars, respectively, and are taken from the *BEA* data. “$\frac{Debt}{Income}_{it}$" gives the ratio of per-capita debt to per-capita income. State-level real GDP growth rates, “$\%\ \Delta$ GDP$_{it}$", are taken from the *BEA* regional economic accounts, with GDP calculated in chained 2009 dollars. The state labor share, “Labor Share$_{it}$", also comes from the *BEA* data, and is calculated by taking total wage and salary compensation over the sum of wage and salary compensation and the gross operating surplus of the business sector. “Gini$_{it}$" gives the state-level Gini index of inequality. Finally, I include measures of the total population and the share of individuals employed in manufacturing.

To isolate the causal effect of local financial market size on household borrowing behavior I rely on variation in credit-supply caused by federal policy shocks contributing the 2003-2007 U.S. housing market boom. During this period there were several federal policy changes that altered the prevailing incentives affecting financial actors and institutions. These changes include—but are not limited to—the following. In June of 2003 the Federal Reserve lowered the Fed Funds Rate to 1%, at the time the lowest rate in nearly 45 years. In December of 2003 the federal government passed the American Dream Downpayment Assistance Act, expanding government provision of housing subsidies with assistance for downpayments and closing costs. Throughout 2003 Fannie Mae and Freddie Mac ramped up their purchases of subprime mortgage securities, buying a total of $81 billion in 2003 (Leonnig, 2008). Finally, in 2004 the Securities and Exchange Commission (*SEC*) suspended the “net capital rule"—a statute effectively limiting the maximum allowable leverage ratio—for large banks (those with tentative net capital of $5 billion or more) (Labaton (2008), *SEC* Release 34-49830). For some banks—such as Bear Stearns—this resulted in leverage ratios rising as high as 33 to 1 (Labaton, 2008).

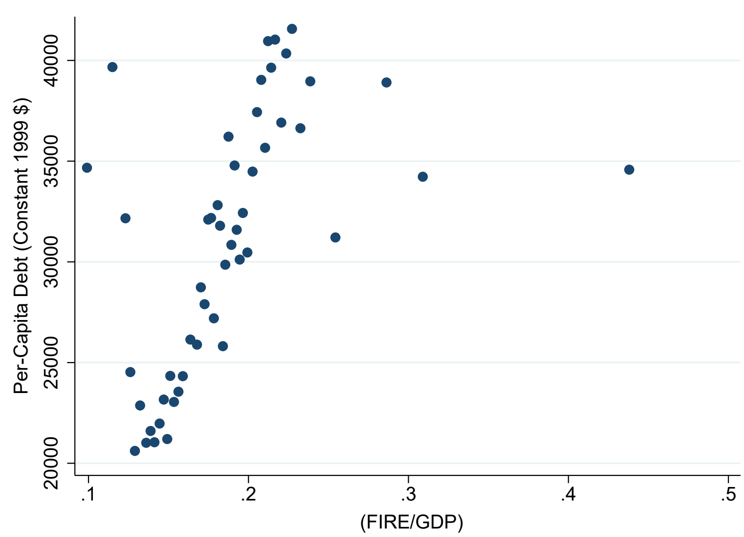

Figure 2.4: Per-Capita Debt v. FIRE/GDP ratio

Source: Federal Reserve Bank of New York and Bureau of Economic Analysis Regional Economic Accounts. Binned scatterplot of real state per-capita debt versus the state FIRE/GDP ratio.

The cumulative impact of these policy changes is the credit-supply shock addressed by Mian and Sufi (2010). The identifying assumption is—although the *de jure* impact of federal policy applies equally to all states—the *de facto* impact of the credit-supply shock on household borrowing should be more severe in states with a larger local financial sector, to the extent that a larger local financial sector (relative to the rest of the economy) resulted in more aggressive marketing of mortgages, a larger impetus to pecuniary emulation, a prioritization of market income over non-market activity, and greater financial sector demand for asset backed securities. Testing this hypothesis requires a delineation of states into a treated and untreated group based on the size of the state financial sector, but any such designation inevitably requires the choice of a threshold which—*a priori*—may appear somewhat arbitrary. To motivate the choice of threshold, Figure (2.4) presents a binned scatterplot of per-capita debt against the share of the FIRE sector in state GDP. The results confirm that per capita-indebtedness at the state level is strongly positively correlated with the size of the local financial sector during the sample period.

Given the positive relationship between the size of the local financial sector and per-capita indebtedness I sort the sample into two groups. A highly financialized, or "high-finance", group of states, for which the state value of the FIRE share in GDP in 1999 exceeded the sample mean in 1999 ($(\overline{\frac{FIRE}{GDP}})_{1999} = 0.179$), and a control group for which the value of the FIRE share of GDP in 1999 fell below the sample mean. This way of sorting ensures that the subset of states falling in the high-finance category is constant throughout the sample period. Because the mean in the first sample year is not the only possible cutoff, I employ several alternative definitions of the treated group in Section (2.4) to test for robustness. Table (2.2) presents the subset of high-finance states. I focus on the share of finance in state GDP, rather than the rentier income ratio, because the former potentially captures the influence of the local financial sector on both households and firms, while the latter may be increasing locally via financialization of the firm, even if the influence of the financial sector on households is otherwise small in a particular state.

Figure (2.5) presents trend-plots for per-capita indebtedness for both high-finance and control states. Because the strategy for identifying the causal effect of the local financial sector on

Table 2.2: High-Finance States

State	$\frac{Finance}{GDP}_{1999}$
Arizona	0.201
California	0.197
Colorado	0.207
Connecticut	0.274
Delaware	0.444
Florida	0.208
Hawaii	0.209
Illinois	0.221
Maine	0.184
Maryland	0.201
Massachusetts	0.229
Minnesota	0.196
Missouri	0.181
Nebraska	0.181
New Hampshire	0.188
New Jersey	0.210
New York	0.303
Oregon	0.184
Pennsylvania	0.183
Rhode Island	0.207
Virginia	0.184

Notes: A state is included in the high-finance group if the FIRE share in state GDP in 1999 exceeded the sample mean in 1999.

household borrowing adopted in this chapter takes a treatment effects approach, seeking to exploit variation in borrowing caused by a policy-induced credit-supply shock (the treatment), the trends for per-capita borrowing in both the treated (states with a large local financial sector) and control (states where the financial sector is small relative to the rest of the economy) groups must satisfy the parallel trends assumption. This assumption requires that—in the absence of the credit-supply shock—the average change in per-capita indebtedness would have been the same in states where finance is large relative to the rest of the economy and states where it is not. While the level of per-capita indebtedness may very well be greater in high-finance states—in the absence of an increased impetus to borrowing, the growth rate of per-capita indebtedness in both the high-finance group and the control group should be similar. The trends presented in Figure (2.5) allow an assessment of the validity of this assumption.

Figures (2.5a) and (2.5b) offer support for the parallel trends hypothesis[3]. The trend in debt per-capita appears parallel across treated and control states prior to 2002. Between 2002 and 2003, there is a shift in the trend for highly financialized states, with per-capita indebtedness increasing at a greater rate in states where local financial markets are large relative to the rest of the economy. The cumulative increase is indicated by Figure (2.5b). These figures suggest that insofar as local financial markets affected household borrowing in response to the credit-supply shock, they did so through the relative importance of finance to the economy as a whole.

To formally test this hypothesis I estimate the following difference-in-differences model:

$$Debt_{it} = \alpha_0 + \alpha_1 HighFinance_i + \alpha_2 After_t + \alpha_3 (HighFinance \times After)_{it} + \mathbf{X}_{it}^T \beta + \eta_i + \delta_t + \epsilon_{it} \tag{2.1}$$

Where α_0 is a constant, $Debt_{it}$ is per-capita indebtedness in state, i, in year, t, $HighFinance_i$ is an indicator for whether a state is in the highly financialized sub-group according to the FIRE share in state GDP, and $After_t$ is an indicator for the period after the credit-supply shock, taking a value of 0 prior to 2003 and a value of 1 thereafter. Given the definitions of $HighFinance_i$ and $After$, the

[3]A formal test of the validity of the parallel trends assumption is presented in (3.5).

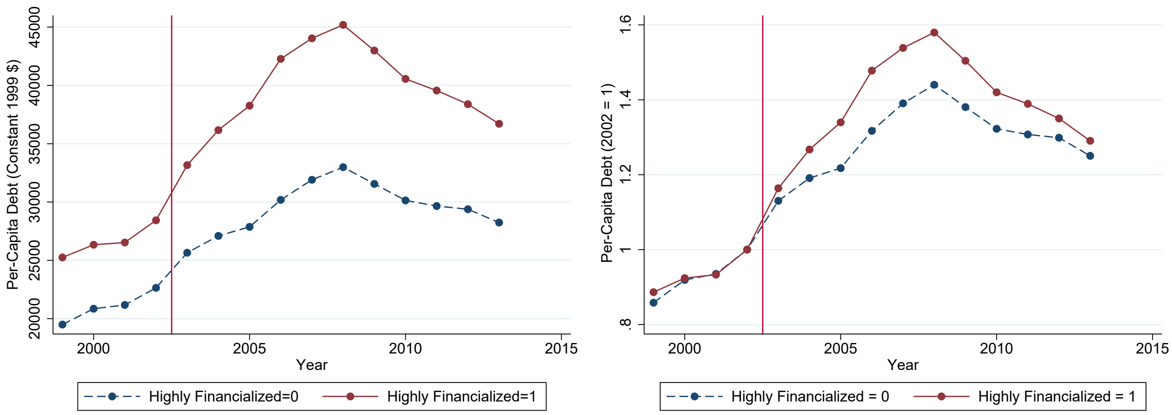

(a) Trend Plot: Levels

(b) Trend Plot: Cumulative

Figure 2.5: Trend Plots: Debt Per-Capita by Relative Size of the Local Financial Sector, 1999-2013

Notes: Plots present trends for debt per-capita in constant 1999 dollars for the time period 1999-2013, conditional on the relative size of the local financial sector.

regression coefficient on their interaction, α_3, gives the treatment effect of the credit-supply shock on high-finance states. $\mathbf{X}_{it}^T$ is a vector of state-level controls, such as real GDP growth, per-capita income, and the labor share of income, η_i and δ_t are state- and time-fixed effects, and ϵ_{it} is an idiosyncratic error term.

It is clear from Figure (2.5) that both treated and control states are affected by a second shock to per-capita indebtedness in the post-2008 period as a result of the financial crisis itself. To separate the effects of the initial credit-supply shock and the financial crisis, I estimate the following intermediate specification which acknowledges the initial increase in per-capita indebtedness during the pre-crisis period and the subsequent decline in debt post-2008. This specification allows the estimation of two separate effects: the effect of the credit-supply shock on high-finance states, and the effect of the financial crisis:

$$\begin{aligned} Debt_{it} = \alpha_0 + \alpha_1 HighFinance_i + \alpha_2 T_{1,t} + \alpha_3 T_{2,t} + \alpha_4 (HighFinance \times T_1)_{it} \\ + \alpha_5 (HighFinance \times T_2)_{it} + \mathbf{X}_{it}^T \beta \\ + \eta_i + \delta_t + \epsilon_{it} \end{aligned} \tag{2.2}$$

Where $T_{1,t}$ takes a value of 1 if $2003 \leq t < 2008$ and 0 otherwise, and $T_{2,t}$ takes a value of 1 if $t \geq 2008$ and 0 otherwise. This specification makes it possible to differentiate between the effect of the credit-supply shock on household debt and the effect of the financial crisis. α_4 gives the effect of the former, α_5 the latter. As before $HighFinance_i$ is an indicator for whether a state is in the high-finance group, α_0 is a constant, $\mathbf{X}_{it}^T$ is a vector of state-level controls, η_i and δ_t are state- and time-fixed effects, and ϵ_{it} is an idiosyncratic error.

Finally, I estimate the following alternative specification, which allows for flexible dynamics in the impact of the credit-supply shock on household debt by varying the treatment effect by year in the post-treatment period. Further, it allows for fully-flexible pre-treatment trend differentials between states in the high-finance group and those that are not, which is less restrictive than the usual assumptions required for difference-in-differences (Mora and Reggio, 2012):

$$Debt_{it} = \alpha_0 + \delta_t + \eta_i + \sum_{t \geq 2003} \Gamma_t (HighFinance \times \delta)_{it} + \mathbf{X}_{it}^T \beta + \epsilon_{it} \tag{2.3}$$

Where α_0 is a constant, δ_t and η_i are time- and state-fixed effects, Γ_t captures the treatment effect of the credit-supply shock on high-finance states in year t, $\mathbf{X}_{it}^T$ and ϵ_{it} are again state-level controls and an idiosyncratic error term, respectively.

The county-level data used in this chapter come from three main sources. First, I obtain data on county-level leverage ratios in the pre-recession period from Mian, Sufi, and Rao (2013). These publicly available data offer a sample of county debt-to-income ratios, constructed by Mian, Sufi, and Rao (2013) using a combination of *Federal Reserve* and *IRS* data, for the years 2001, 2002, and 2004-2007. I then match these data to the Census Bureau's County Business Patterns (CBP) data, from which I construct the FIRE-sector employment share for each county in the year 1999, in order to define the group of treated counties analogously to the state-level case. The CBP employment data provide an alternative measure for the relative size of the local financial sector, in the absence of data on the share of the FIRE sector in county-level output[4]. Finally, I include

[4]Although the *BEA* reports county-level estimates of both employment and compensation by sector, many observations are suppressed to avoid revealing the identity of particular firms or individuals. In contrast, the CBP reports

county-level controls for income, employment, and population from the *BEA*. The final sample consists of 2,175 counties, for the years 2001, 2002, and 2004-2007. Figure (2.6) plots trends in the county-level leverage ratios for counties with high (exceeding the 75th percentile in 1999), moderate (between the 50th and 75th percentile in 1999), and low (less than the 50th percentile) levels of financialization prior to the housing market boom. The figure suggests a relationship between the relative size of the local financial sector and household borrowing during the pre-recession period also holds at the county level.

To test the relationship between the share of the FIRE sector in county employment prior to the housing market boom and within county household borrowing during the boom, I estimate the following model:

$$\frac{Debt}{Income}_{it} = \alpha_0 + \alpha_1 FIRE\_Emp_i^{1999} + \alpha_2 After_t + \alpha_3 (FIRE\_Emp^{1999} \times After)_{it} + \delta_t + \Lambda_i + \epsilon_{it} \tag{2.4}$$

Where $FIRE\_Emp_i^{1999}$ gives the share of the finance, insurance, and real-estate sector in total county employment in county i in 1999. $After_t$ takes a value of 1 after 2003, and 0 otherwise. δ_t is a year-fixed effect, Λ_i is a county-fixed effect, and ϵ_{it} is an idiosyncratic error term. α_3 thus gives an estimate of the effect of increasing financialization in the period prior to the housing market boom on county debt-to-income ratios during the boom[5].

Finally, I use data on county-to-county commuter flows to construct a spatially-lagged measure of county-level financialization. This measure is used to test whether financialization in neighboring geographies significantly impacts county borrowing patterns—even if the county in question is not itself highly financialized. To do this, I obtain data on resident-workplace commuting from the 2000 Census. I then construct a spatial-weights matrix, **W**, where each entry, ϕ_{ij}, is given by:

interval estimates for observations where anonymity is a concern, a procedure which allows for an estimate of the true value to be imputed. In this case, I simply impute the median of the assigned interval.

[5]In Appendix (A.3), I estimate a model at the county level where the treatment is assigned discretely, as in the state-level case. Similarly, Appendix (A.2) presents results from a continuous treatment version of the state-level model.

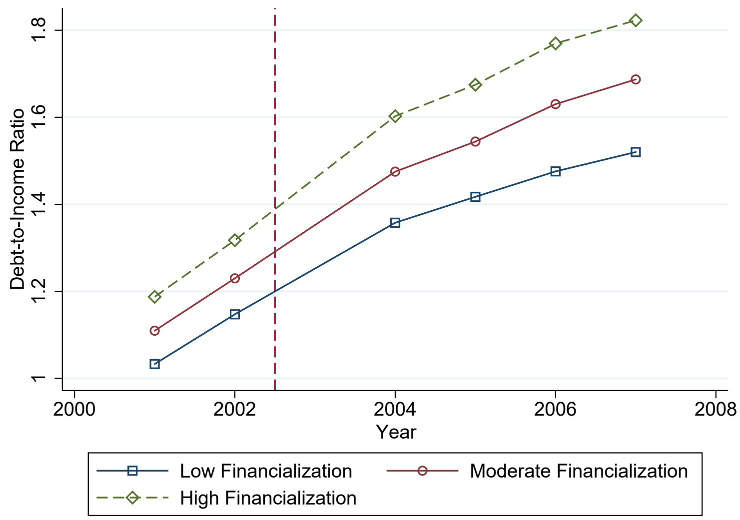

Figure 2.6: County Debt-to-Income Ratios, by Level of Financialization: 2001-2007

Notes: The "Low Financialization" category corresponds to counties which had a FIRE employment-share less than the median in 1999 ($\approx 4\%$), the "Moderate Financialization" category corresponds to counties with a FIRE employment-share between the median and seventy-fifth percentile ($\approx 5.5\%$) in 1999, and the "High Financialization" category to counties with a FIRE employment-share exceeding the seventy-fifth percentile in 1999.

$$\phi_{ij} = \frac{Commuting_{ij}}{\sum_{j=1}^{J} Commuting_{ij}} \tag{2.5}$$

where $Commuting_{ij}$ gives the flow of individuals commuting from county i to county j. ϕ_{ij} thus gives the share of county j in total commuting from county i. Pre-multiplying the vector containing county-level FIRE sector employment shares by **W** thus gives an estimate of the spatial lag of financialization. For an individual county, i, this is given by:

$$S\_FIRE\_Emp_{it} = \sum_{j=1}^{J} \phi_{ij} FIRE\_Emp_{jt} \tag{2.6}$$

The use of commuting data as the basis for the spatial weighting scheme, rather than conventional spatial weighting metrics like distance or contiguity, is preferable from an economic standpoint. Commuting data captures the network of economic activity that exists between regions, allowing for more complex interactions (such as long-distance commuting) that may affect an individual's social and economic network. Commuting data also avoids making the assumption that geographic proximity implies economic proximity, which is problematic in cases where there are topographical features inhibiting economic interaction. In other words, commuting flows make it possible to trace the contours of economic interactions across space in a way that simple distance metrics do not. After constructing spatially-lagged values for county financialization I estimate the following specification:

$$\begin{aligned}\frac{Debt}{Income}_{it} = \alpha_0 + \alpha_1 FIRE\_Emp_i^{1999} + \alpha_2 S\_FIRE\_Emp_i^{1999} + \alpha_3 After_t \\ + \alpha_4 (FIRE\_Emp^{1999} \times After)_{it} \\ + \alpha_5 (S\_FIRE\_Emp^{1999} \times After)_{it} + \delta_t + \Lambda_i + \epsilon_{it}\end{aligned} \tag{2.7}$$

Where $S\_FIRE\_Emp_i^{1999}$ gives the value of spatially-lagged financialization for county i in 1999, and the co-efficient on the interaction term between $S\_FIRE\_Emp_i^{1999}$ and $After_t$—α_5—gives the effect of increasing spatially-lagged financialization during the pre-treatment period on household borrowing during the housing market boom. All other variables are defined as in Equation (2.4). Equation (2.7) has the advantage of teasing out the effect of regional financialization on borrowing in counties with low-levels of own-county FIRE sector employment by exploiting the fact that greater exposure to the credit-supply shock in neighboring counties (via an increased value for spatially-lagged financialization) may impact county borrowing patterns, even if the county itself has low levels of FIRE sector employment.

Financialization in neighboring counties may serve to exacerbate the effects of own-county financialization on household debt in counties with high levels of FIRE sector employment. Figure (2.7) plots trends for the county debt-to-income ratio over counties with low values (below mean) for both own-county and spatially lagged FIRE sector employment, high values for own-county FIRE sector employment and low values for spatially-lagged FIRE sector employment, and high values for both own-county and spatially lagged FIRE sector employment. The figure suggests that counties with large values for both own-county and spatially-lagged FIRE sector employment

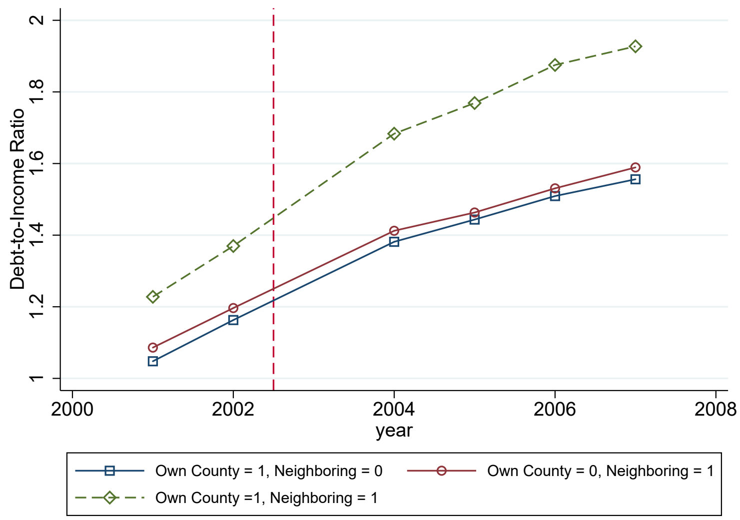

Figure 2.7: Debt-to-Income Ratios, by Own-County and Spatially-Lagged Financialization

Notes: Figure (2.7) presents trend plots for county debt-to-income ratios by both own-county and spatially-lagged financialization. “Own County" is an indicator for whether own-county FIRE sector employment share in 1999 is above or below the sample mean ($\approx 4.8\%$), and “Neighboring" is an indicator for whether the spatially-lagged FIRE sector employment share in 1999 is above or below the sample mean ($\approx 5.7\%$).

experienced a larger increase in borrowing during the housing market other boom than other counties.

2.4 Results

2.4.1 State-level Results

Table (2.3) presents the results from estimating equation (2.1). Column (1) displays the regression coefficients for the initial specification without any controls. Column (2) adds a variety of state-level controls. Column (3) adds time-fixed effects. Column (4) adds both time- and state-fixed effects. The standard errors are clustered at the state level to address the possibility of serial correlation over time within each state, the presence of which would bias the errors downward in the absence of clustering[6].

[6]The regression co-efficient on $HighFinance_i$ is not reported because state-fixed effects wash out all state-specific time-invariant variables. In specifications without state-fixed effects, this coefficient is positive and statistically significant. Random effects versions of the model—which allow for the inclusion of the $HighFinance_i$ term—were also tested. No meaningful difference was found in the size or significance of the treatment effect.

Table 2.3: Estimation Results - Specification 1

	(1) Debt$_{it}$	(2) Debt$_{it}$	(3) Debt$_{it}$	(4) Debt$_{it}$
$After_t$	8474.8***	4957.9***	6479.9***	8664.1**
	(681.1)	(1083.6)	(1991.8)	(3326.9)
$(HighFinance \times After)_{it}$	4642.8***	3661.7***	3156.2***	2910.2***
	(1016.7)	(1031.8)	(973.7)	(941.4)
N	765	765	765	765
R^2	0.40	0.66	0.87	0.87
Controls	N	Y	Y	Y
Time FE	N	N	Y	Y
State FE	N	N	N	Y

Notes: Standard errors in parenthesis, clustered at the state-level. * 0.10 ** 0.05 *** 0.01. Dependent variable is real per-capita indebtedness at the state level, measured in constant 1999 dollars. Controls include the growth rate of real GDP, the labor share of state income, log income and consumption per-capita, state population, the Gini coefficient, and the share of manufacturing in total state employment.

The results indicate that the credit-supply shock had an economically and statistically significant effect on borrowing in high-finance states. Per-capita indebtedness increased by approximately \$3,000 more—an 11% increase relative to the pre-treatment mean—in highly financialized states. This result is robust to the inclusion of both time- and state-fixed effects, as well as a wide variety of relevant controls. This suggests that the growing influence of finance at the state level is in part responsible for the increase in the level of indebtedness that helped engender the financial crisis and ensuing recession.

Although the results from estimating (2.1) support a positive relationship between the degree of local financialization and the size of the local credit-supply shock, it is not clear that the impact of the credit-supply shock was as uniform as the results suggest. Figures (2.3) and (2.5) show that borrowing peaked between 2008-2009, and the increases in borrowing experienced in both highly financialized states and control states were of different magnitude in different years, a fact which is not captured in the usual difference-in-differences framework. Tables (2.4) and (2.5) present results from the estimation of equations (2.2) and (2.3) which address this problem by acknowledging the distinct impacts of the credit-supply shock and financial crisis, and allowing the effect of these shocks to vary across periods.

Table 2.4: Estimation Results - Specification 2

	(1)
	Debt$_{it}$
$T_{1,t}$	14,434.7***
	(2,631.5)
$T_{2,t}$	8,321.4**
	(3,449.3)
$(HighFinance \times T_1)_{it}$	2,059.4**
	(853.2)
$(HighFinance \times T_2)_{it}$	3,462.2***
	(1,045.1)
N	765
R^2	0.875
Controls	Y
Time FE	Y
State FE	Y

Standard errors in parenthesis, clustered at the state level. * 0.10 ** 0.05 *** 0.01. Dependent variable is real per-capita indebtedness at the state level, measured in constant 1999 dollars. Controls include the growth rate of real GDP, the labor share of state income, log income and consumption per-capita, state population, the Gini coefficient, and the share of manufacturing in total state employment.

The results from estimating equation (2.2) suggest that the credit-supply shock and financial crisis had distinct impacts on per-capita indebtedness in highly financialized states. The coefficient on $(HighFinance \times T_1)_{it}$ indicates that the magnitude of the effect of the credit-supply shock in highly financialized states is slightly smaller than the effect estimated in Table (2.3) (however the credit-supply shock did coincide with a large increase—relative to previous periods—in per-capita indebtedness that was common to all states, as suggested by the coefficient on $T_{1,t}$). The coefficient on $(HighFinance \times T_2)_{it}$ indicates that the difference in indebtedness between treated and control group was increasing in the post-recession period, despite the decline in levels of indebtedness in both groups. This suggests that the de-leveraging processes proceeded slower in highly financialized states than in states where the financial sector was small relative to the total size of the economy.

Table (2.5) presents the results for the fully flexible specification given by equation (2.3). The parameter estimates on the interaction between the high-finance indicator and time dummies give

Table 2.5: Estimation Results - Specification 3

	(1)
	Debt$_{it}$
$(HighFinance \times \delta)_{i,2003}$	-226.8
	(767.8)
$(HighFinance \times \delta)_{i,2004}$	741.9
	(910.8)
$(HighFinance \times \delta)_{i,2005}$	1677.7
	(1123.4)
$(HighFinance \times \delta)_{i,2006}$	3337.1**
	(1375.0)
$(HighFinance \times \delta)_{i,2007}$	3311.2**
	(1604.8)
$(HighFinance \times \delta)_{i,2008}$	3384.0**
	(1642.3)
$(HighFinance \times \delta)_{i,2009}$	3763.0***
	(1313.5)
$(HighFinance \times \delta)_{i,2010}$	3655.1***
	(1254.8)
$(HighFinance \times \delta)_{i,2011}$	3211.8***
	(1168.2)
$(HighFinance \times \delta)_{i,2012}$	2502.6**
	(1060.5)
$(HighFinance \times \delta)_{i,2013}$	2202.6**
	(1001.8)
N	765
R^2	0.881
Time FE	Y
State FE	Y
Controls	Y

Standard errors in parenthesis, clustered at the state level. * 0.10 ** 0.05 *** 0.01. Dependent variable is real per-capita indebtedness at the state level, measured in constant 1999 dollars. Controls include the growth rate of real GDP, the labor share of state income, log income and consumption per-capita, state population, the Gini coefficient, and the share of manufacturing in total state employment.

the magnitude of the credit-supply shock (and subsequent crisis) on treated states in year t. The estimates tell how much greater borrowing was in high-finance states, relative to those states in the control group, in a given year.

The regression coefficients on the interactions of interest are only statistically significant from 2006 onward—despite the upward trend in per-capita indebtedness in both state groups in the years prior—indicating there may have been a lag between policy implementation and effectiveness. Further, borrowing in highly financialized states increased the most relative to other states in 2009, confirming the results from table (2.4) that the expansion of personal debt continued into the initial years of the financial crisis—possibly exacerbated by the crisis itself.

Given the results, it is clear that the credit-supply shock had a large positive effect on household borrowing in highly financialized states. The average increase in per-capita borrowing in high-finance states between 2003 and 2013 was approximately $2,000 to $3,000 greater than the increase in all other states.

Robustness Checks

Table (2.6) begins a series of robustness checks. Column (1) tests an alternative threshold for the definition of the treated group: namely, whether a state had a value for the share of the FIRE sector in state GDP in 1999 that exceeded the 75th percentile in 1999. The resulting group of high-finance states is a sub-sample of those originally listed in Column (1) of Table (2.2). The estimated coefficient is similar in magnitude and significance to previous estimates, suggesting that the mean is an appropriate threshold for delineating between a large and small local financial sector. If the relationship between financialization and debt observed via the credit-supply shock were a spurious one, increasing the threshold used to define the treated group may be expected to reduce or eliminate the treatment effect. This does not appear to be the case in this instance. However, one would expect excluding states between the 50th and 75th percentiles from the treated group to raise the mean value of per-capita debt in the untreated group. The fact that the estimated treatment effect is not diminished when these states are excluded suggests the size of the increase in per-capita debt is increasing the "dose" of the treatment. This is explored further in the next sub-section.

Column (2) tests a second alternative threshold by altering the time period in which the treated group is defined. Namely, it uses the mean for the entire pre-treatment period (all years prior to 2003) as the threshold. According to this threshold, a state is considered to have a large local financial sector (relative to the rest of its economy) if the mean value for the FIRE share in state GDP in that state for the years prior to 2003 exceeds the sample mean prior to 2003. This specification allows for a test of the sensitivity of the result to the time horizon that is used to define the highly financialized sub-sample. If the mean value is an appropriate cutoff, then extending the time-frame to which that cutoff is applied in the pre-treatment period—even if the treated group is slightly altered, to the extent that the average share of finance in state GDP is increasing over time in most states as Figure (2.1) suggests— should not drastically alter the significance or magnitude of the result. The estimated coefficient is close in both magnitude and significance to the original treatment effect, offering support for the original definition of the treated group.

Columns (3) and (4) of Table (2.6) drop states which may be seen as outliers that potentially drive the main result. Column (3) drops the five states which had the highest housing cost during the month of the bubble in which average home values in the United States peaked[7]—based on the Seasonally Adjusted Case-Shiller U.S. National Home Price Index and the Zillow Home Value Index for single family homes, the latter measured at the state level[8]—and New York, which is not featured among the top five states in terms of housing cost according to the Zillow measure, but is in the top five most expensive states by median sales price and average listing price on the website Trulia, a competitor of Zillow. As a result, I drop New York, New Jersey, California, Massachusetts, Hawaii, and Washington D.C. in Column (3). The purpose of this specification is to test whether the increase in borrowing in states with large local financial markets is purely explained by differences in housing cost across states, rather than the influence of the financial

[7]July 2006.

[8]See: https://www.zillow.com/research/data/

sector more broadly. The parameter estimate in Column (3) does not support this hypothesis, as the difference in borrowing as a result of the credit-supply shock between the highly financialized sub-sample and all other states is still statistically significant and of a similar magnitude as previous estimates.

Table 2.6: Robustness Checks

	(1) Debt$_{it}$	(2) Debt$_{it}$	(3) Debt$_{it}$	(4) Debt$_{it}$	(5) Non-Mortgage$_{it}$	(6) $\frac{Debt}{Income}_{it}$	(7) $\frac{Debt}{Income}_{it}$
$(HighFinance \times After)_{it}$	2838.9***	2560.1***	2854.5***	2988.5***	652.0***	0.0613**	—
	(892.6)	(890.4)	(885.7)	(919.9)	(241.5)	(0.0259)	
$(FIRE\_Emp^{1999} \times After)_{it}$	—	—	—	—	—	—	0.0199**
							(0.00890)
N	765	765	675	690	765	765	765
R^2	0.870	0.871	0.892	0.893	0.809	0.850	0.849
Controls	Y	Y	Y	Y	Y	Y	Y
Time FE	Y	Y	Y	Y	Y	Y	Y
State FE	Y	Y	Y	Y	Y	Y	Y

Notes: Standard errors in parenthesis, clustered at the state level. * 0.10 ** 0.05 *** 0.01. Variables follow the same conventions as Table (2.3). Columns (1) and (2) present results for alternative treatment thresholds. Column (3) excludes the highest housing cost states. Column (4) excludes the five states with the highest average share of homes sold for a loss during the recession. Column (5) uses only non-mortgage debt as the dependent variable. Column (6) uses the state-level debt-to-income ratio as the dependent variable. Column (7) uses the debt-to-income ratio as the dependent variable again, but this time looks at the impact of increasing the FIRE sector share of *employment* in 1999, to arrive at an estimate immediately comparable to the county-level estimates presented below.

Column (4) of Table (2.6) drops the five states with the highest average share of homes sold for a loss throughout the recession, based on Zillow home sales data. This test controls for the possibility that differences in observed borrowing are driven solely by variation in the severity of the housing market crash across states, rather than the credit-supply shock. The states excluded thus are California, Nevada, Florida, Arizona, and Michigan. The result of this test does not support the hypothesis that borrowing necessitated by the fall in housing prices in states where the crash was most severe is driving the original result. The estimated regression coefficient is of similar magnitude and statistical significance as the original estimates.

Column (5) of Table (2.6) estimates the treatment effect of the credit-supply shock on non-mortgage (student, auto, and credit card) debt. This test is to address concerns about whether the observed difference in borrowing across states with large and small local financial sectors is driven by increased mortgage borrowing in states with a large real-estate sector—which may simply reflect a high level of mortgage debt—rather than by the pressures to borrow associated with financialization. The results indicate that the credit-supply shock had a statistically significant effect on non-mortgage borrowing in highly financialized states (although of a smaller magnitude than the effect on total debt inclusive of mortgages). This result supports the hypothesis that the increase in borrowing in highly financialized states was driven in part by the recasting of social ties as financial claims and the subsequent efforts by households to "keep up with the Joneses." To the extent that the *de jure* effect of the credit-supply shock lay largely on borrowing for housing consumption, any increase in borrowing for other purposes must be the result of a subsequent behavioral response, rather than a direct result of a policy change.

Finally, Columns (6) and (7) of Table (2.6) look at the impact of financialization on the debt-to-income ratio, rather than per-capita debt. The estimates thus arrived at are directly comparable to the county-level results obtained below. Column (6) looks at the impact of the high financialization treatment on the state debt-to-income ratio, Column (7) uses the share of the FIRE sector in state-level *employment*—rather than output—in 1999, in a continuous fashion, such that the coefficient obtained in column (7) is the state-level equivalent of the estimated county-level coefficients. In both cases, financialization appears to have a statistically and economically significant impact on the state-level debt-to-income ratio. For states in the highly financialized group the debt-to-income ratio increased by approximately 6 percentage points more after the credit-supply shock than other states. Further, a 1 percentage-point increase in the share of the FIRE sector in state-level employment in 1999 is associated with a 2 percentage point increase in the debt-to-income ratio after the credit-supply shock.

Dose Responsiveness

A second category of robustness check, which expands on Column (1) of Table (2.6), looks at whether the treatment is dose responsive. When a treatment is not discretely assigned a valid concern regards whether the cut point of the continuous variable used to define the treated group influences the estimation results. This hypothesis can be tested by examining the magnitude of

Table 2.7: Dose Responsiveness Test

	(1)
	Debt$_{it}$
$(Low\_Treatment \times After)_{it}$	2774.2**
	(1120.3)
$(High\_Treatment \times After)_{it}$	3997.7***
	(985.3)
N	765
R^2	0.88
Controls	Y
Time FE	Y
State FE	Y

Notes: Standard errors in parenthesis, clustered at the state level. * 0.10 ** 0.05 *** 0.01. The variables $Low\_Treatment_{it}$ and $High\_Treatment_{it}$ denote states in the 50th-75th and 75th+ percentiles for the share of the FIRE sector in state GDP in 1999, respectively.

the effect under different “doses" of the treatment—in this case, different sizes of the share of the financial sector in value added at the state level. If debt is dose responsive to the size of the local financial sector, the treatment effect should be larger in states with a greater degree of financialization. Table (2.7) presents the results from a regression which tests this hypothesis. Column (1) splits the sample into three groups: an untreated group, a lightly treated group (those between the 50th and 75th percentile for the size of the local financial sector in 1999), and a highly treated group (those above the 75th percentile for the size of the local financial sector in 1999). The results indicate that per-capita debt is indeed dose responsive to the treatment, with the increase in per-capita borrowing increasing relative to the control group in the degree of the treatment received.

Placebo Test

A third category of robustness check is possible in the difference-in-differences framework. The so-called “placebo" test assigns an alternative treatment date in the pre-treatment period, using the same treated group, and estimates the resulting treatment effect. Under the parallel trends assumption, the expectation is that the regression coefficient should not be statistically significantly different from zero. If the placebo effect differs from zero this indicates that the pre-treatment

Table 2.8: Placebo Test

	(1)
	Debt_{it}
$After_t$	2610.4**
	(1162.5)
$(HighFinance \times After)_{it}$	-237.8
	(375.8)
N	204
R^2	0.788
Time FE	Y
State FE	Y
Controls	Y

Notes: Standard errors in parenthesis, clustered at the state level. * 0.10 ** 0.05 *** 0.01. Variables follow the same conventions as Table (2.3).

trends are not parallel, and the difference-in-difference estimates are likely biased. Table (2.8) presents the results from this test. The estimated placebo effect is negative in magnitude and statistically insignificant, a result that lends support to the validity of the parallel trends assumption.

Randomization Test

A second placebo- or falsification-type test concerns what would happen if the treatment were randomly assigned to non-treated units. Randomly assigning the treatment to non-treated units allows an assessment of whether the estimated difference-in-differences coefficient for treated states is large relative to states chosen at random. In the context of this chapter, one would like to ascertain whether or not the effect of the credit-supply shock on highly financialized states is large relative to the distribution of effects for regions not as exposed to financialization. Examples of this type of falsification test are common in treatment effects studies (Abadie et al., 2010; Angrist and Krueger, 1999; DiNardo and Pischke, 1997).

To implement the randomization test I adopt the following procedure. First, I assign treatment status to states in both treated and untreated groups by replacing the value of the FIRE sector share in state GDP with a value randomly drawn from a normal distribution matched to sample moments[9]. I then estimate the treatment effect for the randomly assigned treatment for n^{10} such random draws and compare the distribution of estimated placebo effects to the size of the "true" effect on highly financialized states. Figure (2.8) presents the results of this test. The results indicate that the size of the effect on treated units far exceeds even the upper bound of the implied placebo distribution, lending support to the validity of the size and significance of the impact of the credit-supply shock in highly financialized states.

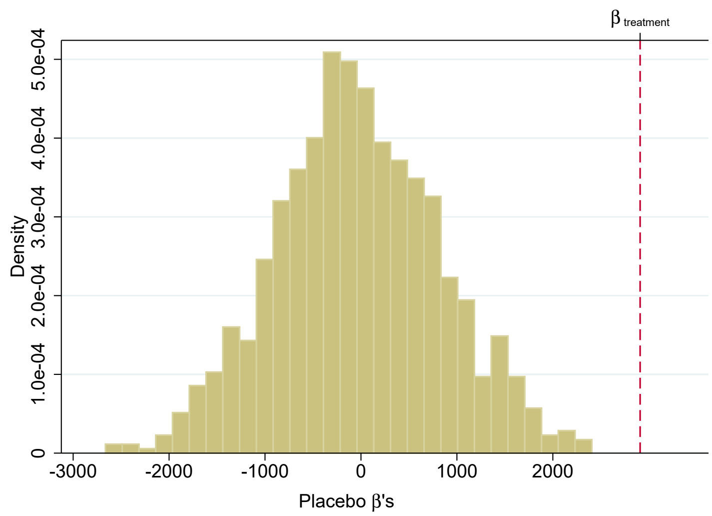

Figure 2.8: Randomization Test - Placebo Distribution and True Effect

Notes: The histogram displays the distribution of estimated placebo effects from the randomization test. The red dotted line indicates the size of the treatment effect on actually treated units.

Serial Correlation

Although I make use of clustered standard errors in an attempt to address the issue of serial correlation, Bertrand, Duflo, and Mullainathan (2002) show that this alone may not be sufficient to fully solve the problem in a difference-in-differences context. Two alternative means of addressing serial correlation are subsequently suggested by Bertrand, Duflo, and Mullainathan (2002). The

[9]The results of this test do not appear sensitive to the choice of distribution. For example, similar results are found when an uniform distribution is used.

[10]I set $n = 1000$.

first is estimating the standard errors via block bootstrapping. This technique involves bootstrapping the standard errors in such a way that the within-state auto-correlation structure is maintained by keeping all the residuals of a particular state together. Essentially, this involves re-sampling the residual distribution of ϵ_{it} to estimate a sequence of parameters from which the test statistic of interest is drawn (Bertrand, Duflo, and Mullainathan, 2002, 16)

The second alternative means of addressing serial correlation is averaging the data into two periods, before and after the treatment. The treatment effect is then estimated using this two-period panel. A slightly modified version of this technique—also suggested by Bertrand, Duflo, and Mullainathan (2002)—can be used to estimate the treatment effect in the presence of different pre-treatment period trends (or period lengths) while still addressing the possibility of serial correlation. This two-step procedure is as follows. First, using the full sample, regress the dependent variable of interest (debt per-capita) on state- and time-fixed effects and a vector of controls. Second, split the residuals from the above regression for only the treated group into residuals from before and after the initial treatment period, and obtain an estimate of the treatment effect by regression of these residuals on an “after" dummy in a two-period panel regression.

I re-estimate the effect of the credit-supply shock on high-finance states, addressing the possibility of serial correlation, using both block bootstrap standard errors and the two versions of the two-period panel technique outlined above. Table (2.9) presents the results.

Table 2.9: Serial Correlation Adjustments

	(1) Debt_{it}	(2) Debt_{it}	(3) DebtResid_{it}
Treatment	2,910.2***	3,284.2***	1,452***
	(909.5)	(983)	(417.3)
N	765	102	42
R^2	0.87	0.94	0.21

Notes: Standard errors in parenthesis. * 0.10 ** 0.05 *** 0.01. Column (1) reports the block bootstrap standard error. Columns (2) and (3) report standard errors from the simple and residual aggregation methods. Controls are included in the first stage of the residual aggregation method in Column (3) (i.e. the stage in which the residuals are calculated). A full set of controls is included in Columns (1) and (2).

Column (1) shows the effect of the credit-supply shock on highly financialized states where the standard error is estimated by block bootstrap, for 200 replications, using the same model specification as Column (4) of Table (2.3). The standard error actually decreases as a result. Column (2) estimates the treatment effect in the simple two-period panel. Again, the coefficient is statistically significant and of a similar magnitude as previous regressions. Finally, Column (3) estimates the treatment effect via the alternative, residual aggregate method outlined above. The resulting parameter estimate is smaller—however, this may be a result of the small sample size. Nonetheless, the result statistically and economically significant.

Difference-in-Discontinuities

Table 2.10: Difference-in-Discontinuities Estimates

	$Debt_{it}$ $h = 0.052$	$Debt_{it}$ $h = 0.025$	$Debt_{it}$ $h = 0.01$
Treatment	4259.6***	3791.7***	3154.9*
	(1130.7)	(1193.7)	(1625.86)
N	750	510	420
R^2	0.85	0.86	0.87
Time FE	Y	Y	Y
State FE	Y	Y	Y

Notes: Standard errors in parenthesis, clustered at the state level. * 0.10 ** 0.05 *** 0.01.

Grembi et al. (2016) propose an alternative treatment effects estimator for situations where a confounding discontinuity arises at the treatment threshold that may affect observations in both the treated and untreated groups. As an example, suppose there exists some factor—say a change in the relative supply of skilled workers—that affects per-capita household borrowing in states near the highly financialized threshold. In this case, the traditional difference-in-difference estimator fails to properly identify the treatment effect of the credit-supply shock on household debt, and estimates of the treatment effect will be biased by the confounding factor that occurs near the discontinuity. The solution to this problem is a "difference-in-discontinuities" approach, which generalizes the insights of difference-in-differences and regression discontinuity to control for the possibility of unobserved heterogeneity at the treatment threshold. The follow sketch will help clarify the approach. Following Grembi et al. (2016), Define $Y_{it}(0)$ as the value of the outcome variable (per-capita indebtedness at the state level) in the untreated group. Define $Y_{it}(1)$ similarly for the treated group (high-finance states). Assume the treatment occurs at time t_0 and at some threshold for the share of the FIRE sector in state GDP, F_c. Then, it is possible to define $y_k^- \equiv lim_{\epsilon \to 0} E[Y_{it}(k)|F_{it} = F_c - \epsilon, t \geq t_0]$ for an observation with treatment status k, and y_k^+ similarly (adding ϵ instead of subtracting). The presence of a confounding discontinuity at the treatment threshold implies $y_0^+ - y_0^- \neq 0 \equiv \Gamma_0$, where Γ_0 is the size of the confounding discontinuity. Under these conditions the standard difference-in-differences estimator will be biased.

Grembi et al. (2016) show that if the confounding discontinuity is roughly constant over time and affects both treated and un-treated groups, then the unbiased “difference-in-discontinuity" estimator for the treatment effect (in this case, the impact of the credit-supply shock on household debt) is given by:

$$\begin{aligned} Debt_{it} = \alpha_0 + \alpha_i F_{it}^* + HighFinance_i(\gamma_0 + \gamma_1 F_{it}^*) \\ + After_t[\gamma_0 + \gamma_1 F_{it}^* + HighFinance_i(\beta_0 + \beta_1 F_{it}^*)] + \eta_i + \delta_t + \epsilon_{it} \end{aligned} \tag{2.8}$$

Where $Debt_{it}$, $After_t$, $HighFinance_i$, η_i, and δ_t, are per-capita indebtedness at the state level, an indicator for the post-treatment period, an in indicator for treated states, time-fixed, and state-fixed effects, as before, and $F_{it}^* = \frac{Finance}{GDP}_{it} - \overline{F}_{1999}$ is the share of finance in state GDP normalized by the treatment threshold. The coefficient β_0 identifies the treatment effect of of the credit-supply shock on highly financialized states. The above model is estimated on the sample of states in the interval $F_{i,1999} \in [\overline{F}_{1999} - h, F_{1999,\max}]$, where $F_{i,1999}$ is the value of the share of finance in state GDP in 1999. For robustness, I test multiple values of h, beginning with a value equal to the standard deviation of $\overline{F}_{1999}$. Table (2.10) presents the results.

The results of the difference-in-discontinuities estimation support the original model. The credit-supply shock has a large positive effect on per-capita indebtedness in highly financialized states. These results are robust to multiple values of h.

2.4.2 County-level Results

Table (2.11) presents results from the county-level estimating equations. Column (1) presents the results from equation (2.4), Column (2) presents the results from equation (2.7), and Column (3) presents the results from estimating equation (2.7) on a reduced sample that excludes counties with high (above the sample mean) own-county FIRE sector employment shares. The findings are consistent with those at the state level. The results support hypothesis that local financialization had a significant effect on household borrowing during the housing market boom. In particular, a one-percentage-point increase in the own-county FIRE sector employment share in 1999 increases the county debt-to-income ratio following the credit-supply shock by approximately one percentage point.

Table 2.11: County-level Results

	(1) $\frac{Debt}{Income}_{it}$	(2) $\frac{Debt}{Income}_{it}$	(3) $\frac{Debt}{Income}_{it}$
$(FIRE\_Emp^{1999} \times After)_{it}$	0.0117**	0.00817**	0.0142
	(0.00480)	(0.00376)	(0.0161)
$(S\_FIRE\_Emp^{1999} \times After)_{it}$		0.0222***	0.0163***
		(0.00667)	(0.00565)
N	13,050	13,050	9,408
R^2	0.702	0.707	0.719
Time FE	Y	Y	Y
County FE	Y	Y	Y
Controls	Y	Y	Y

Notes: Standard errors in parenthesis, clustered at the state level. * 0.10 ** 0.05 *** 0.01. Column (1) presents results from estimating equation (2.4), Column (2) from estimating equation (2.7), and Column (3) from estimating equation (2.7) on a reduced sample excluding counties with high (above mean) levels of own-county FIRE sector employment shares. Control variables include (log) levels and growth rates of county income, population, and employment, and an index of county industry employment specialization.

The effects of local financialization are also characterized by significant spatial spillovers. The estimated co-efficient on $(S\_FIRE\_Emp^{1999} \times After)_{it}$ in Column (2) suggests that a one-percentage-point increase in spatially-lagged financialization in 1999 increases the county debt-to-income ratio during the housing market boom by approximately two percentage points. Further, these effects are not limited to counties with high-levels of own-county FIRE sector employment. Column (3) excludes all counties with FIRE sector employment shares greater than or equal to the sample mean. The results show that even for regions which are not characterized by high relative levels of financial sector employment, the proximity to high FIRE sector employment share counties has a positive effect on within county household debt.

2.4.3 Sub-Sector Analysis

A simple way of assessing the mechanism by which an increase in the size of the FIRE sector impacts household borrowing is to investigate whether there are differences in the observed effect across various sub-sectors of the FIRE industry. The *BEA* data on the FIRE sector can be subdivided into six industry sub-categories corresponding to grouped three-digit NAICS industries. These sub-categories include: *monetary authorities—central banking, credit intermediation, and related services*, which groups establishments that perform central banking functions as well as those that (1) lend funds raised from depositors, (2) lend funds raised from credit market borrowing, or (3) facilitate the lending of funds or issuance of credit by engaging in such activities as mortgage and loan brokerage, clearing house and reserve services, and check cashing services; *securities, commodity contracts, and other financial investments and related activities*, which groups establishments that are primarily engaged in (1) underwriting securities and/or making markets for securities and commodities, (2) acting as agents (i.e., brokers) between buyers and sellers of securities and commodities, (3) providing securities and commodity exchange services, or (4) providing other services, such as managing portfolios of assets, providing investment advice, or trust, fiduciary, and custody services; *insurance carriers and related activities,* which groups establishments engaged in (1) underwriting annuities and insurance policies or (2) facilitating such underwriting by selling insurance policies and by providing other insurance and employee benefit related services; *funds, trusts, and other related vehicles,* which groups legal entities organized to pool securities or other assets on behalf of shareholders; *real estate*, which groups establishments engaging in renting or leasing real estate to others, managing real estate for others, selling, buying, or renting real estate for others, or appraising real estate; and finally, *rental and leasing services and lessors of nonfinancial tangible assets,* which groups establishments that rent consumer goods and equipment, those that rent business machinery and equipment, and those that assign rights to intangible assets, such as patents, trademarks, and brand names.

For each of the above categories I repeat the main exercise of the chapter, examining the impact of financialization on the severity of the credit-supply shock by taking as treated states above the mean in each sub-sector in 1999. This exercise acts as an additional falsification test, as one would not expect insurance carriers and related activities or rental and leasing services and lessors of non-financial assets to have a large positive impact on household debt. In contrast, if financialization is truly responsible for geographic differences in household borrowing, one would expect securities, commodity contracts, and other financial investments and related activities to be strongly related to household debt, as these establishments are exactly those which demand assets linked to household debt. Table (2.12) presents the results from this exercise. Columns (1)-(6) examine the impact of each sub-sector separately, Column (7) includes each sub-sector in a single regression.

Table 2.12: Sub-Sector Results

	(1) Debt$_{it}$	(2) Debt$_{it}$	(3) Debt$_{it}$	(4) Debt$_{it}$	(5) Debt$_{it}$	(6) Debt$_{it}$	(7) Debt$_{it}$
$(CreditIntermediation \times After)_{it}$	1101.7						709.8
	(930.6)						(723.9)
$(Securities \times After)_{it}$		4393.9***					3645.6***
		(1012.8)					(1184.3)
$(RealEstate \times After)_{it}$			3633.4***				2557.7***
			(818.3)				(814.4)
$(FundsTrusts \times After)_{it}$				1995.3*			665.2
				(1010.0)			(708.6)
$(Insurance \times After)_{it}$					-215.5		-1686.4**
					(816.1)		(704.7)
$(RentalLeasing \times After)_{it}$						624.1	243.0
						(1030.4)	(757.2)
N	765	765	765	765	765	765	765
R^2	0.864	0.881	0.879	0.867	0.863	0.863	0.892
Time FE	Y	Y	Y	Y	Y	Y	Y
State FE	Y	Y	Y	Y	Y	Y	Y
Controls FE	Y	Y	Y	Y	Y	Y	Y

Notes: Standard errors in parenthesis, clustered at the state level. * 0.10 ** 0.05 *** 0.01.

The results indicate that the two sub-sectors of FIRE most strongly related to household borrowing are real estate and securities, commodity contracts, and other financial investments and related activities. This finding supports the notion that the increase in household borrowing in the pre-recession period is at least partially explained by increased financial sector demand for instruments linked to mortgage debt. Surprisingly, the results offer little evidence of a link between traditional banking institutions and household borrowing. The coefficient on the interaction between the $CreditIntermediation$ variable—a dummy indicating whether or not a state had an above the mean value for the monetary authorities—central banking, credit intermediation, and related services sub-sector in 1999—and the $After$ variable is insignificant in both Column (1) and Column (7). Given the significance of the impact of the securities sub-sector, one interpretation of this result is that the intensity with which traditional banking institutions marketed mortgage loans may have been contingent on the presence financial sector demand for assets derived from those loans. Further, it is worth nothing that the securities treated group is almost a perfect sub-set of the real estate group, with only Pennsylvania and Washington D.C. included in the former but not the latter. It is also unlikely that the sub-sector results are driven by paper institutions—those present in a state in only a legal sense—as the two states where this form of institution is most common—Delaware and South Dakota—are excluded from both the securities sub-sector treated group and the real estate sub-sector treated group.

The results also indicate that neither the funds, trusts, and other related vehicles, insurance carriers and related activities, nor rental and leasing services and lessors of nonfinancial tangible assets sub-sectors have significant positive impacts on household debt. The first is statistically significant only in the regression where all other sub-sectors are excluded, and becomes statistically insignificant in Column (7). When all sub-sectors are included the insurance sub-sector has a statistically significant *negative* impact on household debt, which suggests that inclusion of the insurance sector in the FIRE sector in the main analysis may bias the results downward.

Additional light can be shed on the sub-sector results by looking at correlations between the $HighFinance_i$ treatment indicator variable and each of the sub-sector treatment indicator vari-

Table 2.13: Correlations Between Overall and Sub-Sector Treatment Variables

Sub-Sector	Correlation
Credit Intermediation	0.29
Securities	0.57
Real Estate	0.57
Insurance	0.43
Funds and Trusts	0.20
Rental Leasing	0.15

ables. Table (2.13) presents these simple correlations. The correlations suggest that the securities and real estate sub-sectors drive inclusion in the aggregate highly financialized category.

2.4.4 Distributional Implications

The aggregate nature of both the state- and county-level data precludes a direct analysis of the distributional implications of the increase in household borrowing in highly financialized regions before and after the recession. Although the estimated effect of an increase in the relative size of the local financial sector is economically and statistically significant, it is certain that the actual increase in borrowing was not born equally by all households. For example, not all households have debt, and the proportion of households that do varies by the position of those households in the income distribution. To tease out some of the distributive implications of the increase in borrowing in regions with large local financial sectors, I make use of data from the Survey of Consumer Finances—a tri-annual cross sectional survey of U.S. families containing comprehensive information on income, assets, net worth, and other financial variables of interest—for the years 2001 to 2007. Figure (2.9) begins by plotting the estimated share of households which have debt balances greater than zero for each decile of the income distribution. The figure is arranged so that each point along the x-axis represents a bin corresponding to an income decile (e.g. the point at the 20th percentile represents the share of households with a positive debt balance with household incomes between the 20th and 30th percentile of the income distribution).

Figure (2.9) confirms what was thought to be the case: borrowing patterns are not uniform across households, and vary substantially with income. The pattern observed confirms some of the

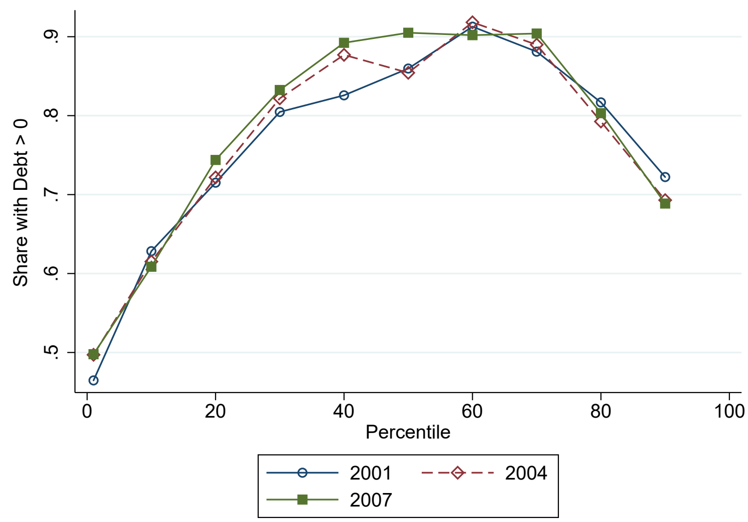

Figure 2.9: Estimated Share of Households with Positive Debt, by Income Percentile

Notes: Plot presents estimates of the share of households with debt balances greater than zero for each decile of the income distribution, by year, from the Survey of Consumer Finances.

claims of Mason (2017), namely that it is not households at the bottom end of the income distribution who make up the lion's share of borrowers. The 40th to the 70th percentile of the income distribution have the largest share of households with positive debt balances—close to 90% of all households in that range. Further, the share of households with positive debt balances remained fairly constant for both low- and high-income households during the pre-crisis credit expansion. The only segment of the income distribution that saw significant entry of new borrowers during the housing market boom were households with incomes between the 30th and 60th percentile, as evidenced by the way the debt-share curve gradually bows outward between 2001 and 2007 along this segment.

Another question regards the estimated debt-burden to individual households. While it is clear that middle class households are the single largest group of borrowers—as estimated by the share of these households with positive debt balances—does the size of the average debt-burden born by households in each decile of the income distribution vary substantially? Figure (2.10) answers this question by plotting the estimated average household debt-to-income ratio for each decile of the income distribution, from 2001 to 2007.

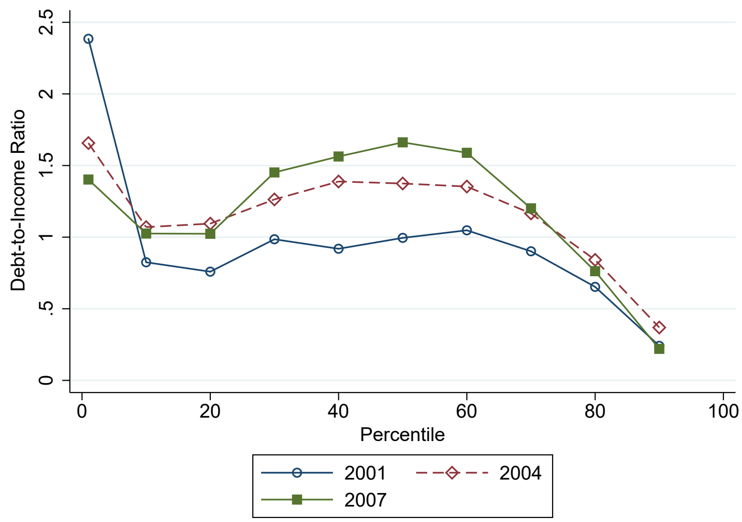

Figure 2.10: Estimated Average Household Debt-to-Income Ratio, by Income Percentile

Notes: Plot presents estimates of the average household debt-to-income ratio for each decile of the income distribution, by year, from the Survey of Consumer Finances.

First, it is worth noting that because the Survey of Consumer Finances oversamples households from the top end of the wealth distribution, the large debt-to-income ratios for households in the bottom half of the income distribution in Figure (2.10) may not reflect the borrowing behavior of truly poor households because households which report little to no income but have high wealth may show up in this segment. Second, the figure displays a clear increase in the size of household borrowing relative to income between 2001 and 2007 for households residing in the middle of the income distribution, and little to no change for households near the top. This is consistent with the narrative advanced by Mason (2017) (in that low-income households are not the primary source of new borrowing), and suggests that the debt effects of the credit-supply shock and recession fell hardest on middle-class households.

Given the distributional characteristics suggested by Figures (2.9) and (2.10), a third question relates to the size of the treatment effect on households with debt. That is, although it is clear that the aggregate increase in in-debtedness was largest among households in the middle of the income distribution, the fact that so few households near the bottom of the income distribution have debt implies that the increase in borrowing for infra-marginal households may have been larger for poor households. To get an estimate of the size of the treatment effect on households with positive debt-balances by income decile, note that the state-level per-capita effect estimated above can be decomposed as follows:

$$\hat{\beta}_{treatment} = \Delta\left(\frac{Debt}{Population}\right) \approx \left(\frac{\Delta Debt}{Population}\right) \tag{2.9}$$

Where the first equality comes from the definition of the treatment effect (the increase in per-capita indebtedness in states with a large local financial sector) and the second approximation will hold as long as state population is roughly constant during the sample period (an assumption that is not too heroic, given that average annual population growth for the United States was less than 1% for much of the sample).

Then, the treatment effect on households with a positive debt-balance is given by:

$$\hat{\beta}_{Debt} \approx \left(\frac{\Delta Debt}{Population}\right)\left(\frac{Population}{Population_{Debt>0}}\right) = \hat{\beta}_{treatment} \times \left(\frac{Population}{Population_{Debt>0}}\right) \tag{2.10}$$

Similarly, the treatment effect on households with debt, by income percentile, is given as:

$$\hat{\beta}_{Debt,p_i} \approx \frac{\Delta Debt}{Population_{Debt>0,p_i}} = \left(\frac{\Delta Debt}{Population}\right)\left(\frac{Population}{Population_{Debt>0}}\right)\left(\frac{Population_{Debt>0}}{Population_{Debt>0,p_i}}\right) \tag{2.11}$$

$$\hat{\beta}_{Debt,p_i} \approx \hat{\beta}_{Debt} \times \left(\frac{Population_{Debt>0}}{Population_{Debt>0,p_i}}\right) \tag{2.12}$$

Figure (2.11) presents estimates of the size of the treatment effect across infra-marginal households by income percentile. Implicitly, this figure assumes that the per-capita increase in debt was the same across the income distribution (that is, I am not estimating a separate treatment effect for each percentile), and that the variation that arises on the infra-marginal effect is due to variation in the share of households within each income percentile that take on debt. While unlikely to hold in a strict sense, this decomposition is nonetheless informative. Unsurprisingly, the plot of the estimated size of the treatment effect is inversely proportional to the share of households with positive debt balances in each decile of the income distribution. Intuitively, if the population of households with positive debt balances in each income decile was roughly constant throughout the credit expansion period, the estimated per-capita increase in borrowing must be larger for households along portions of the income distribution where a lower share of overall households have debt. Although the aggregate increase in debt was larger for middle-class households, because comparatively few households in the top and bottom portions of the income distribution have debt, the estimated per-capita increase for those households that do is larger.

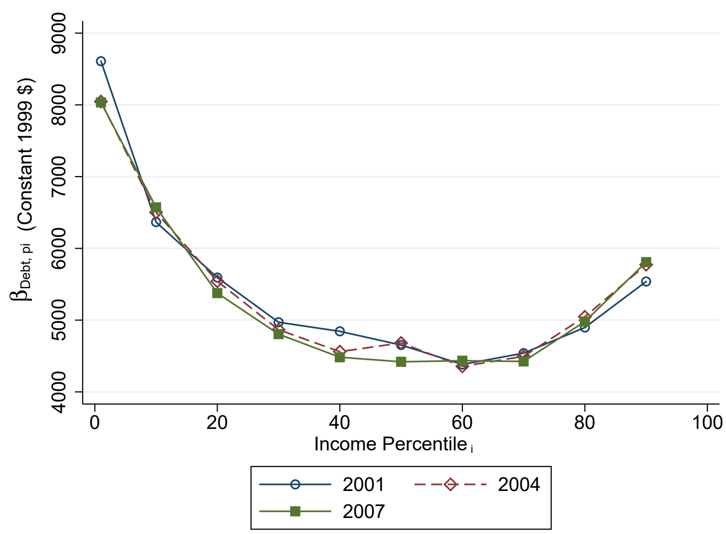

Figure 2.11: Estimated Treatment Effect - by Income Percentile

Notes: Plot presents estimates of the treatment effect of the credit-supply shock on household borrowing for each decile of the income distribution, by year, from the Survey of Consumer Finances.

2.5 Conclusion

The influence of local financial markets matters for household borrowing decisions. In this chapter, I test the relationship between the relative size of the local financial sector and per-capita indebtedness at the state level for the United States. Specifically, I adopt a difference-in-differences approach that exploits a number of policy changes that occurred at the start of the 2003-2007 housing bubble, the sum of which constitute a “credit-supply shock." I also make use of variation in indebtedness provided by the financial crisis itself. The central contention of this approach is that—although the *de jure* impact of the credit-supply shock is the same across states, insofar as the policy changes that facilitated the credit-supply shock occurred at the federal level—the *de facto* impact of the credit-supply shock, and the impact of the financial crisis, should be larger in highly financialized states. I find that this is indeed the case. On average, per-capita indebtedness—including mortgage, auto, credit-card, and student-loan debt—increased by $2,000-$3,000 more in states with large financial sectors.

I find that a similar relationship between FIRE sector *employment* shares and indebtedness holds at the county level. On average, a one-percentage point increase in the county FIRE sector employment share in the period preceding the housing market boom results in a one percentage point increase in the county debt-to-income ratio during the boom. The relationship between county FIRE sector employment and household debt also features significant spatial spillovers. A one percentage point increase in spatially-lagged financialization increases the county debt-to-income ratio by approximately two percentage points.

The results in this study have important implications for the relationship between financial markets, debt, and economic growth. First, the results offer empirical evidence that finance plays an important role as a *social* actor via its influence over household behavior. Models that treat financial markets as an entity that benignly intermediates between savers and borrowers will therefore tend to understate the impact that financial markets have on the economy. Second, the link between financialization and indebtedness has important implications for growth and employment volatility. In their work on the Great Recession, Mian and Sufi (2010) demonstrate that household leverage growth up to 2006 strongly predicts the severity of the recession across counties, despite giving little explanation for what may have caused this variation in borrowing. The results presented here suggest the influence of *local* financial markets—via their impact on household debt—as a fundamental cause of the Great Recession.

Chapter 3

Geography Matters: The Impact of Geographic Expansion on Bank Performance During the Great Recession

3.1 Introduction

A central question in banking concerns the impact of geographic expansion on bank performance. Modern portfolio theory suggests that increasing the geographic breadth of a bank's activity should improve bank outcomes. Intuitively, if geographic expansion allows a bank to diversify its portfolio of assets, exposure to idiosyncratic risk should fall as a result. Other literature argues that geographically diverse banks also attain cost-efficiencies which—in addition to lower idiosyncratic risk—improve bank stability (Boyd and Prescott, 1986; Diamond, 1984).

The notion that geographic expansion improves bank performance has had a powerful influence on banking policy. In the 1980's various states entered into agreements allowing interstate branch banking, culminating in the Riegle-Neal Interstate Banking and Branching Efficiency Act of 1994, which legalized interstate branching at the federal level (Krozner and Strahan, 1999; Mian and Sufi, 2017). Kane (1996) argues that the poor performance of geographically confined banks in the pre-interstate branching era partially explains support for deregulation of interstate branching, and Berger and DeYoung (2006) provide evidence that technological advances that reduce the agency costs of distance helped facilitate the growth of bank branch networks.

A secondary implication of the modern portfolio theory approach to geographic expansion is the gradual elimination of geographic differences in rates of return (Dymski, 2009). As restrictions on bank branching are eliminated, geographic differences in rates of return should be arbitraged away as banks move funds to their most productive use. Contrary to this implication, geographic variance in bank returns increased in the period following the removal of interstate branching restrictions. Figure (3.1) plots the mean absolute deviation in state-level return on assets and (log) net-income per-institution from 1966 to 2017. The figure shows an upward trend in geographic variance, with clear spikes during the banking crises of the late 1980's and late 2000's.

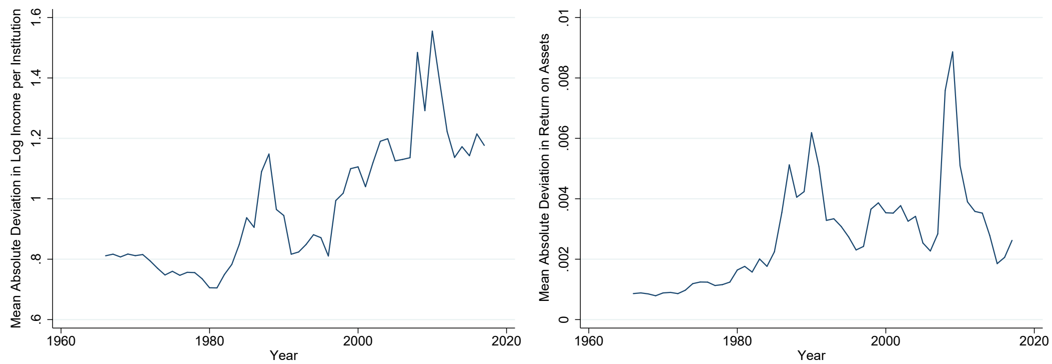

(a) Variance in Net Income per-Institution **(b)** Variance in Return on Assets

Figure 3.1: Geographic Variation in Bank Performance, 1966-2017

Source: FDIC Historical Statistics on Banking, authors' calculations.

Increasing geographic variance in returns casts doubt on the conclusions of modern portfolio theory. Recent literature suggests a number of reasons why geographic expansion may produce less than beneficial results. First, geographically concentrated banks are known to rely on "soft information" about potential borrowers obtained through relationship lending (Berger and Udell, 2002). Evidence suggests that relationship accounts exhibit lower rates of default and higher rates of utilization, and that the soft information necessary to maintain relationship accounts is difficult to transmit across either geographic or hierarchical distance (Agarwal and Hauswald, 2010; Agarwal et al., 2018; Liberti and Mian, 2009). Second, the existence of geographic information asymmetries may expose banks to negative performance shocks resulting from adverse selection in mergers and acquisitions. If a bank expands into a new geography by acquiring an existing branch in the new location, the outgoing institution may be privy to local economic conditions about which the acquiring bank is unaware. The literature provides mixed evidence of "lemons" effects in mergers and acquisitions, suggesting that under certain conditions mergers may negatively impact future bank performance (Beccalli and Frantz, 2009; Carlson and Mitchener, 2009; DeLong and DeYoung, 2007). Finally, Petkov (2016) provides evidence that banks exposed to regionally-specific negative shocks export changes in liquidity by reducing lending at distant branches, suggesting that branch networks act as a mechanism for transmitting negative shocks from one region of the country to another. In this case, even if geographic expansion reduces idiosyncratic risk, geographic bank expansion may contribute to greater systemic risk by amplifying national exposure to otherwise localized downturns.

In this chapter, we argue that geographic expansion negatively impacts bank performance during crises. Using panel data from the Federal Deposit Insurance Corporation (*FDIC*), we show that increases in the deposit-weighted average distance between a bank headquarters and its branches is associated with a small performance decline during the Great Recession. Although the marginal effect of a change in distance is small, the *total effect* of distance is large for banks with large geographic networks. During the recession, our estimated coefficients imply that geographic expansion resulted in a 1.1 percentage-point decline in return on assets, a 13 percentage-point decline in return on equity, and a 2.4 percentage-point increase in the share of non-current assets for the largest banks, relative to pre-recession levels. We also show that the *marginal effect* of distance on performance is declining in branch network size, suggesting that the greatest losses from geographic expansion come from initial bank expansion—consistent with the hypothesis that small banks rely soft information not easily transferable across distance. Finally, we document a separate, direct effect of branching on performance.

The rest of the chapter is organized as follows. Section (3.2) briefly reviews the empirical literature on the relationship between geographic expansion and bank performance. Section (3.3) outlines a model of bank branching and systemic risk. Section (3.4) describes the data and estimation strategy. Section (3.5) presents the results. Section (3.6) concludes.

3.2 Literature Review

The evidence on the impact of geographic expansion on bank performance is mixed. A number of papers find that increases in “functional" distance—the economic distance of a branch location from the operational center of a bank—decreases bank performance (Alessandrini, Croci, and Zazzaro, 2005; Torluccio, Cotungo, and Strizzi, 2011). Feng (2018) finds that bank competition is associated with greater risk taking, such that if expansion increases competition it may lead to an increase in risky behavior. Chong (1991), Demsetz and Strahan (1997), and Archarya, Hasan, and Saunders (2006) all find that geographic expansion is associated with an increase in risky lending activities. Akhigbe and Whyte (2003) and Deng and Elyasiani (2008) provide evidence to the contrary—namely, that geographic expansion reduces risk.

The study closest in spirit to ours is Goetz, Laeven, and Levine (2016). Using panel data on bank holding companies for the period 1986-1997, Goetz, Laeven, and Levine (2016) find that geographic expansion reduces bank risk. There are at least two reasons why the findings of Goetz, Laeven, and Levine (2016) might differ from ours. First, the authors adopt an instrumental variables strategy based on a “gravity-deregulation" model, which attempts to exploit the timing of the removal of branching restrictions by individual states in the 1980’s. In every case, the first-stage F-statistic reported by the authors for the instrumental variables regression falls below 10, suggesting the estimates may be biased due to weak instruments (Staiger and Stock, 1997).

Second, our study looks at the impact of distance on bank performance during a recession, while Goetz, Laeven, and Levine (2016) focus their analysis on the entire 1986-1997 period. Our study therefore focuses on a channel absent in Goetz, Laeven, and Levine (2016)’s analysis: the impact of geographic expansion on exposure to macroeconomic shocks. Even if it is the case that geographic expansion improves bank performance in normal times, it may be undesirable if it leads to increased exposure to macroeconomic downturns. It is possible that geographic expansion is characterized by a trade-off between reduced idiosyncratic risk and increased systemic risk. The latter mechanism is the focus of this chapter.

3.3 A Model of Bank Branching and Systemic Risk

In this section we consider what type of mechanisms might be at play in a world where branch banking negatively impacts bank performance. Given that our empirical estimates indicate large branch networks do negatively impact bank performance, and that the average size of bank branch networks appears to be increasing over time, we are interested in exploring what sort of spillovers might be consistent with both these facts. That is, we attempt to answer the question: what sort of incentive structure is consistent with expansion of branch network size, despite the negative impact it has on performance? We show that the existence of unacknowledged geographic information asymmetries, which result in spillovers from the *aggregate* branch network to the probability of borrower repayment (i.e. systemic risk), are consistent with larger than optimal branch network size.

We consider the following banking environment. To isolate the channel of expanding branch networks on bank risk, we abstract from the possibility of banks engaging in trading or other non-interest income generating activities and assume that banks earn all their income via lending at some volume, L. Banks are assumed to be price takers, and receive a rate of return, r, on their loans with probability q (capturing the probability of borrower repayment), the latter of which may vary with the riskiness of the bank's portfolio of lending activities, but—assuming each bank has chosen its desired risk portfolio *ex ante*—is, for the time being, exogenous from the perspective of an individual bank. Finally, banks face costs associated with origination, $C(L)$, where $C(\cdot)$ is increasing in L. The bank's profit function is:

$$\Pi = qrL - C(L) \tag{3.1}$$

By choosing L to maximize (3.1), the profit-maximizing solution is

$$L^* = C'^{-1}(qr) \tag{3.2}$$

Which is increasing in the rate of return and probability of repayment. The above formulation of the problem assumes that banks may vary their volume of lending continuously along the internal margin. If instead banks are lending constrained by the breadth of their branch network, the above solution will be insufficient to characterize bank behavior. Thus, we consider an alternative characterization of the problem, where a bank's lending is constrained by the number of branches it operates. We assume the following relationship:

$$L = L(n_j) = n_j^{\alpha} \tag{3.3}$$

Where n_j is the number of branches operated by bank j, and $\alpha > 0$ is the branch elasticity of loan volume. The bank's cost function will also now depend on n_j. We assume the relationship takes the form:

$$C(n_j) = \gamma n_j \tag{3.4}$$

Where $\gamma > 0$ is a parameter. Given this set-up, the bank's problem is now to choose n_j to maximize:

$$\Pi = qrn_j^{\alpha} - \gamma n_j \tag{3.5}$$

Written this way, the problem is essentially one of "optimal firm size," in the tradition of Coase (1937) and Williamson (1967). Setting $\frac{\partial \Pi}{\partial n_j} = 0$ gives the solution to the bank's problem:

$$\alpha qrn_j^{\alpha-1} = \gamma \tag{3.6}$$

Assuming no further complications, solving (3.6) for n_j, and plugging the resulting value back into (3.3) will give the equilibrium values for the number of branches and lending volume as a function of parameters only:

$$n_j = (\frac{\alpha qr}{\gamma})^{\frac{1}{1-\alpha}} \tag{3.7}$$

However, this solution is unlikely to paint a complete picture. In most cases expansion of the branch network, n_j, will require the bank to begin operating branches in new geographies—either through mergers and acquisitions, or by opening a new branch location. In either situation, the expansion of the bank's branch network will have an impact on the amount of information required for the bank to successfully achieve its *ex ante* desired risk portfolio, and thus the subsequent likelihood of actually matching the portfolio it desires *ex post*. If geographically expanding banks face either (1) lemons effects via mergers and acquisitions, (2) a loss of a relationship lending channel, or (3) exposure to a greater variety of previously isolated regional downturns via the branch network, then branch network expansion will be characterized by a negative externality. Further, to the extent that these effects are likely to spillover across banks, the probability of repayment will depend not only on the amount of branches operated by bank j, but on the total number of branches in the system. Assuming there are N total banks, this number is given by: $n = \sum_{j=1}^{N} n_j$

In terms of the simple model above, exposure to negative shocks via lemons effects and soft information loss, and transmission of the resulting negative performance shocks both within and between branch networks, imply that the probability of borrower repayment, q, is negatively impacted by the exposure to systemic risk resulting from the overall expansion of branch networks, n, regardless of the bank's *ex ante* desired risk portfolio. That is—because of the information asymmetries inherent in expansion of the branch network—banks will only ever be able to achieve their desired risk portfolio imperfectly. q can thus be re-written as:

$$q = \overline{q}n^{\beta} \tag{3.8}$$

Where $\overline{q}$ is the *ex ante* desired risk portfolio (probability of repayment) chosen by the bank, n is the *aggregate* number of branches, and $\beta < 0$ represents the magnitude of the systemic risk externality. $\beta = 0$ corresponds to the no externality case from above. If all banks are identical, such that the optimal value for n_j is the same for each bank, then $n = n_j N$. This implies that in equilibrium:

$$q = \overline{q}(Nn_j)^{\beta} \tag{3.9}$$

Assuming the bank internalizes no portion of the externality, using equations (3.6) and (3.9) to solve for the number of branches chosen by the bank in equilibrium gives:

$$\hat{n_j} = [\frac{\alpha \overline{q} r N^{\beta}}{\gamma}]^{\frac{1}{1-\alpha-\beta}} \tag{3.10}$$

In contrast, if the bank's profit maximization problem were solved by a social planner who fully internalized the extent of systemic risk, the number of branches chosen would be given by:

$$n_j^* = [\frac{(\beta+\alpha) \overline{q} r N^{\beta}}{\gamma}]^{\frac{1}{1-\alpha-\beta}} \tag{3.11}$$

Which, given $\beta < 0$ immediately implies the following:

$$\hat{n_j} > n_j^* \tag{3.12}$$

If expansion of the branch network results in negative spillovers from exposure to and transmission of negative performance shocks, the number of branches chosen by any particular bank in competitive equilibrium will always exceed the socially optimal amount.

3.4 Data and Methods

Two main data sources are used in this chapter. First, we use the *FDIC* Summary of Deposits (SOD). From 1994 on, the SOD provide annual information at the bank branch level on total deposits at each *FDIC*-insured branch as of June 30th of each survey year. Second, we use the *FDIC* Statistics on Depository Institutions (SDI). The SDI contain information on bank assets and liabilities, income and expense, performance and condition ratios, and other characteristics for all *FDIC*-insured institutions. The information contained in the SDI are collected quarterly. For stock variables, such as assets and liabilities, the reported values are the current value for the variable of interest as of the time of the survey. For flow variables, such as income and expense, values are reported each quarter on either a quarterly or a year-to-date basis. To construct a consistent annual series for the income and expense variables we take the sum of the quarterly values for the past four quarters as of June 30th of each year.

To construct our measure of geographic expansion we match the bank-level SDI data to the branch-level SOD data using the unique certification number assigned to each bank by the *FDIC*. We then use the address information in the SOD and SDI to calculate the distance in miles between each branch and the corresponding bank headquarters[11]. Finally, we aggregate the data to the bank level, creating a weighted average of the distance between a bank headquarters and its branches[12]. We weight each branch observation by the share of total bank deposits at the branch, to adjust our measure for the dispersion of bank activity across branches. Figure (3.2a) plots the movement of our distance measure over time. The figure shows a clear upward trend. This trend is not surprising, given the amount of recent consolidation in the banking industry. This consolidation is shown in Figure (3.2b) which illustrates a sharp decline in the total number of *FDIC*-insured institutions since 1980. It is also worth pointing out the deviation from trend that occurs in our distance measure during the Great Recession—a result of branch closures occurring throughout the 2007-2011 period.

To form our estimation sample, we restrict the analysis to the 1999-2011 period, focusing on the effect of the Great Recession from 2007 to 2011. Our final sample consists of an unbalanced panel of 105,915 bank-year observations, covering 10,873 separate banking institutions. Table (3.1) presents summary statistics for the main variables of interest. For ease of interpretation, we present our analysis in terms of a normalized distance variable, which we call "$ZDist_{it}$", that transforms the weighted average distance variable to a z-score using the sample mean and standard deviation.

"$Distance_{it}$" measures the weighted average distance in miles between a bank's headquarters and its branches, for bank i in year t. "$Branches_{it}$" gives the natural log of the total number

[11]For branches where latitude-longitude coordinates are missing, we use the batch geocoder from Texas A&M Geoservices to assign a latitude-longitude coordinate to the address of the branch.

[12]Because the *FDIC* data is aggregated at the bank-level, rather than the bank holding company-level, we control for inclusion in a bank holding company as a robustness check in Section (3.5).

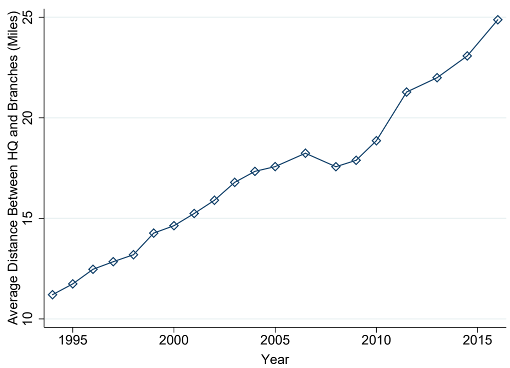

(a) Geographic Dispersion, 1994-2016

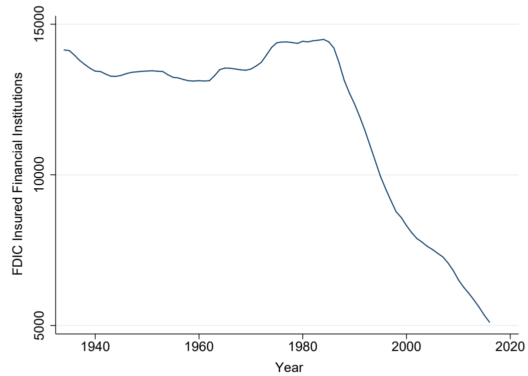

(b) Banking Consolidation, 1934-2016

Figure 3.2: Geographic Dispersion and Banking Consolidation

Notes: Source: FDIC Data, authors' calculations. Figure (3.2) presents the weighted average distance between a bank headquarters and its branches, in miles. Weights constructed using the share of bank deposits at each branch.

Table 3.1: Summary Statistics, *FDIC* Data, 1999-2011

	Mean	Std. Dev.	Min	Max
$Distance_{it}$	17.16	60.56	0	2463.5
$ZDist_{it}$	0	1.000	-0.283	40.40
$Branches_{it}$	1.103	1.060	0	8.745
$Assets_{it}(000's)$	\$991,010.0	\$17,279,235.4	760.0	\$13,25,384,417.08
$Deposits_{it}(000's)$	\$662,887.2	\$11,120,624.1	\$0	\$846,676,930.9
$Loans_{it}(000's)$	\$569,767.1	\$8,548,991.6	\$0	\$560,653,721.3
$\frac{Capital}{Asset}_{it}$	0.110	0.0707	0.00880	1.005
$NonInterestIncomeShare_{it}$	0.111	0.100	0	8.333
$\frac{NonCurrentAssets}{Assets}_{it}$ (%)	1.037	1.640	0	49.67
ROA_{it} (%)	0.876	1.353	-39.48	34.96
ROE_{it} (%)	9.245	8.976	-49.95	49.89
N	105,915			

of branches in a bank's branch network. “$Assets_{it}$," “$Deposits_{it}$," and “$Loans_{it}$," give the current dollar value of bank i's total assets—which includes cash, loans, securities, bank premises, and other assets, but does not include off-balance-sheet accounts— deposits, and net loans and leases, in constant 1999 dollars, as of year t. “$\frac{Capital}{Asset}_{it}$" gives the ratio of a bank's tier one capital—defined as common equity plus noncumulative perpetual preferred stock plus minority interests in consolidated subsidiaries less goodwill and other intangible assets—to the total value of its assets. “$NonInterestIncomeShare_{it}$" gives the value of non-interest income—income from fiduciary activities, plus service charges on deposit accounts in domestic offices, plus trading gains (losses) and fees from foreign transactions, plus other foreign transaction gains, plus other gains and fees from trading assets and liabilities—to the sum of non-interest and interest income. “$\frac{NonCurrentAssets}{Assets}_{it}$" gives the value of assets that are 90 days or more past due plus assets placed in nonaccrual status plus other real estate owned, over the value of total assets. Finally “ROA_{it}" and “ROE_{it}" give annualized net income as a percentage of average total assets and average total equity, respectively. Because there are large differences between the average institution—with relatively small geographic coverage in its branch network—and large institutions (e.g. Wells Fargo, JP Morgan, Citi, etc.) we analyze differences across bank-size groups in Section (3.5).

To get at the central question of the chapter—the impact of geographic expansion on bank performance during the Great Recession—Figure (3.3) plots trends in three main performance measures: return on assets, return on equity, and the share of non-current assets in total assets, for banks with operations of varying geographic scale. The figure previews the key empirical finding: the total effect of the Great Recession is largest for banks with wide-reaching branch networks. In particular, banks with a standardized distance measure greater than or equal to three standard deviations experienced the largest decline in performance during the recession.

To assess the impact of the geographic expansion on bank performance during the recession more formally, we adopt a difference-in-differences approach. In particular, we begin by regressing our measures of bank performance on a continuous version of our standardized distance measure interacted with an indicator variable for the recession:

(a) Non-Current Assets (% of Total Assets)

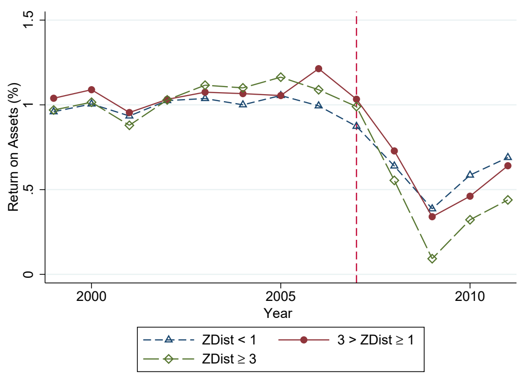

(b) Return on Assets (%)

(c) Return on Equity (%)

Figure 3.3: Bank Performance Measures, by Extent of Geographic Expansion

Notes: "$ZDist$" gives the standardized measure of weighted average distance between a bank headquarters and its branches. The vertical line is plotted at 2007.

$$Y_{it} = \beta_0 + \beta_1 ZDist_{it} + \beta_2 Recession_t + \beta_3 (ZDist \times Recession)_{it} + \mathbf{X}_{it}^T \beta + \epsilon_{it} \qquad (3.13)$$

Where "Y_{it}" is the bank performance measure of interest—return on assets, return on equity, or the share of non-current assets in total assets—"$ZDist_{it}$" is the standardized distance measure, "$Recession_t$" is a dummy variable taking a value of 1 after 2007 and 0 otherwise, and $\mathbf{X}_{it}$ is a vector of controls. β_3 then gives the treatment effect of the recession on performance, and variation in "$ZDist$" captures the "dose" of the treatment.

We also estimate a restricted version of (3.13), where only those banks for which there are observations before and after the onset of the recession are included. In this restricted specification we require banks to have at least two years of observations on either side of the treatment. The distance variable is then modified so that it captures average standardized distance in the pre-recession period:

$$Y_{it} = \beta_0 + \beta_1 ZDistPre_i + \beta_2 Recession_t + \beta_3 (ZDistPre \times Recession)_{it} + \mathbf{X}_{it}^T \beta + \epsilon_{it} \quad (3.14)$$

Finally, we estimate a more flexible specification, which allows for the effect of the recession to vary by year, while also allowing non-linearity in the effect of distance:

$$Y_{it} = \beta_0 + \beta_1 ZDist_{it} + \beta_2 ZDist_{it}^2 + \delta_t + \sum_{y=1999}^{2011} (\delta \times ZDist)_{it} + \sum_{y=1999}^{2011} (\delta \times ZDist^2)_{it} + \mathbf{X}_{it}^T + \epsilon_{it} \quad (3.15)$$

Where δ_t are year-fixed effects. The interactions of "$ZDist_{it}$" and its squared term with the year-fixed effects allows an assessment of the impact of distance on performance both during and prior to the recession. Further, inclusion of the squared term makes it possible for the marginal effect of distance on performance to vary according to the current size of the bank's branch network. If soft-information is both important for small bank lending decisions and difficult to transmit across distance, the marginal effect of geographic expansion on bank performance should be largest for these institutions. For ease of interpretation, the results from this specification are presented margins plots.

3.5 Results

Table (3.2) presents results from estimating equation (3.13). Standard errors are clustered at the bank level. Control variables include (log) assets, (log) branches, the share of non-interest income in total income, and the capital-asset ratio. Because assets, deposits, and loans are highly collinear, we choose to include only one, although altering which is included makes no meaningful difference. Column (1) presents results from our baseline specification with just controls. Column (2) adds bank- and year-fixed effects. Column (3) adds fixed effects for the state of the bank headquarters and the bank's asset specialization. The latter are categories assigned to banks by the *FDIC* that capture the institution's primary specialization in terms of asset concentration. For example, a bank is classified as a mortgage lending specialist if residential mortgage loans plus mortgage backed securities total more than 50 percent of the bank's total assets. The *FDIC* assigns indicators for nine such categories, and we include fixed-effects for each category. Column (4) replaces state- and year-fixed effects with state-by-year-fixed effects. Finally, Column (5) clusters the standard errors at the bank-post-2007 level, to address the fact that the performance measures' variance could be affected by the recession. Following Correia (2015) we exclude singleton groups in regressions with many fixed-effects, the inclusion of which overstates statistical significance.

For each of the chosen performance measures, geographic expansion is associated with worse performance during the recession. A one-standard deviation increase in distance results in a decline in return on assets of 0.06 percentage points, a decline in return on equity of 0.72 percentage points, and an increase in non-current assets by 0.13 percentage points. This decline in performance is relatively small. Pre-recession means for each of these variables are 1.0 for return on assets, 10.6 for return on equity, and 0.69 for non-current assets as a share of total assets. The estimated regression coefficients therefore represent a 6% decline in return on assets, a 6.8% decline in return

Table 3.2: Estimation Results 1

	(1)	(2)	(3)	(4)	(5)
Return on Assets					
$(ZDist \times Recession)_{it}$	-0.0529***	-0.0588***	-0.0591***	-0.0560***	-0.0560***
	(0.0144)	(0.0153)	(0.0158)	(0.0153)	(0.0135)
$Branches_{it}$	-0.166***	-0.482***	-0.495***	-0.499***	-0.499***
	(0.0162)	(0.0438)	(0.0433)	(0.0425)	(0.0395)
Return on Equity					
$(ZDist \times Recession)_{it}$	-0.558***	-0.717***	-0.718***	-0.649***	-0.649***
	(0.148)	(0.175)	(0.182)	(0.166)	(0.137)
$Branches_{it}$	-0.548***	-3.180***	-3.237***	-3.252***	-3.252***
	(0.1000)	(0.221)	(0.213)	(0.210)	(0.186)
Non-Current Assets					
$(ZDist \times Recession)_{it}$	0.148***	0.124***	0.128***	0.107***	0.107***
	(0.0299)	(0.0270)	(0.0286)	(0.0251)	(0.0205)
$Branches_{it}$	0.0423***	0.324***	0.323***	0.309***	0.309***
	(0.0164)	(0.0371)	(0.0361)	(0.0357)	(0.0301)
N	105,915	105,407	105,407	105,407	105,407
Controls	Y	Y	Y	Y	Y
Bank FE	N	Y	Y	Y	Y
Year FE	N	Y	Y	N	N
State FE	N	N	Y	N	N
Specialization FE	N	N	Y	Y	Y
State × Year FE	N	N	N	Y	Y

Notes: Standard errors in parenthesis. * 0.10 ** 0.05 *** 0.01. Columns (1)-(4) cluster standard errors at the bank level. Column (5) clusters standard errors at the bank-post-2007 level. Control variables include (log) assets, (log) branches, the share of non-interest income in total income, and the capital-asset ratio.

on equity, and a 19% increase in non-current assets, relative to the pre-recession mean, respectively. However, the total effect of geographic expansion on bank performance implied by the estimated coefficients is large for the banks with the largest geographic networks. GMAC (now Ally Bank), Citigroup, Countrywide, Bank of America, and Wells Fargo had values of "$ZDist_{it}$" between 15 and 22 during the recession. For these banks, the average decline in performance—associated with an average value of "$ZDist_{it}$" of 18.5—amounts to a 1.1 percentage-point decline in return on assets, a 13.32 percentage-point decline in return on equity, and a 2.4 percentage-point increase in the share of non-current assets in total assets. When the value of the total assets owned by these banks is considered, the macroeconomic implications of the decline in performance due to geographic expansion are large—despite the small number of banks affected in this fashion.

Our results also indicate a direct, negative impact of branching on bank performance—after other size measures are controlled for. One possible interpretation of this result is the existence of information costs associated with opening a branch in a new location. For a bank operating on a given geographic scale, a unit increase in the natural log of the number of branches is associated with a larger performance decline than an increase in geographic scale for a bank with a fixed number of branches. Thus, even if a bank is not—on average—expanding its geographic coverage, branching may nonetheless be associated with increased risk.

To further assess the impact of geographic expansion on bank performance, Table (3.3) presents results from estimating equation (3.14) on a reduced panel which covers banks with at least two years of observations before and after the onset of the recession. The results presented in Table (3.3) support our previous findings. In particular, banks with large geographic networks in the period prior to the recession experienced larger performance declines during the recession. The estimated regression coefficients from this specification are larger than those obtained from estimating equation (3.13), although the confidence intervals implied by the standard errors overlap in most cases, suggesting the coefficeints are unlikely to be statistically significantly different. The estimated negative impact of branching also appears similar in magnitude to the results presented in Table (3.2).

Table 3.3: Estimation Results 2

	(1)	(2)	(3)	(4)	(5)
Return on Assets					
$(ZDistPre \times Recession)_{it}$	-0.0523***	-0.0724***	-0.0704***	-0.0666***	-0.0666***
	(0.0143)	(0.0166)	(0.0172)	(0.0166)	(0.0133)
$Branches_{it}$	-0.174***	-0.485***	-0.490***	-0.499***	-0.499***
	(0.0181)	(0.0437)	(0.0435)	(0.0428)	(0.0378)
Return on Equity					
$(ZDistPre \times Recession)_{it}$	-0.712***	-0.922***	-0.944***	-0.861***	-0.861***
	(0.174)	(0.205)	(0.220)	(0.202)	(0.159)
$Branches_{it}$	-0.718***	-3.396***	-3.405***	-3.446***	-3.446***
	(0.116)	(0.226)	(0.228)	(0.226)	(0.192)
Non-Current Assets					
$(ZDistPre \times Recession)_{it}$	0.136***	0.129***	0.139***	0.115***	0.115***
	(0.0322)	(0.0320)	(0.0348)	(0.0296)	(0.0234)
$Branches_{it}$	0.0492***	0.368***	0.362***	0.352***	0.352***
	(0.0189)	(0.0412)	(0.0413)	(0.0413)	(0.0341)
N	85,281	85,281	85,281	85,281	85,281
Controls	Y	Y	Y	Y	Y
Bank FE	N	Y	Y	Y	Y
Year FE	N	Y	Y	N	N
State FE	N	N	Y	N	N
Specialization FE	N	N	Y	Y	Y
State × Year FE	N	N	N	Y	Y

Notes: Standard errors in parenthesis. * 0.10 ** 0.05 *** 0.01. Columns (1)-(4) cluster standard errors at the bank level. Column (5) clusters standard errors at the bank-post-2007 level. Control variables include (log) assets, (log) branches, the share of non-interest income in total income, and the capital-asset ratio.

Finally, we present results from estimating equation (3.15), which allows for a non-linear effect of geographic expansion on bank performance. This specification allows us to test whether the marginal effect of expansion on performance varies depending on the current geographic scale of a bank's operations. This feature allows for a direct assessment of the relationship lending hypothesis. If geographically concentrated banks rely on soft information obtained through relationship lending, and this information is difficult to transmit across distance, the marginal performance decline resulting from geographic expansion should be larger for these banks—even if the *total* performance decline is largest for banks with expansive branch networks. Figures (3.4a)-(3.4c) present the marginal effect of expansion on performance for each of our performance measures, for banks with operations of various geographic scale, before and after the recession.

Three key findings emerge from Figure (3.4). First, there is a clear distinction between the effect of expansion on performance before and after the recession. During "normal" times, expansion ranges from having no effect on performance to having a small, positive impact on performance. The latter effect is concentrated among the smallest banks, as made evident by the impact of expansion on non-current assets in Figure (3.4c). In contrast, expansion has a large negative impact on performance during the recession. However—the second insight from Figure (3.4)—this effect is declining in the geographic spread of a bank's branch network. The marginal effect of an increase in the average distance between a bank's headquarters and its branches falls as the average distance is increased. Thus, despite the fact that the estimation results in Tables (3.2) and (3.3) suggest the *total* effect is largest for large banks, the marginal effect of additional geographic expansion for these banks is small. Finally, the size of the marginal effect of geographic expansion on performance for banks with small geographic networks offers support for the relationship lending hypothesis. Expansion is costly for small banks if it results in a decline in the ability of these banks to accurately assess risk due to a loss of soft-information obtained from relationship lending. Figure (3.4) is consistent with this interpretation, as the marginal performance decline during the recession is largest for the banks most likely to make use of soft information obtained from relationship lending.

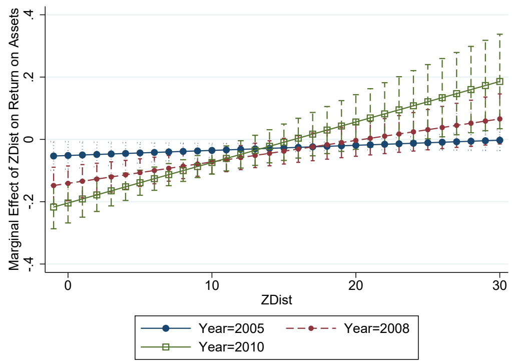

(a) Return on Assets

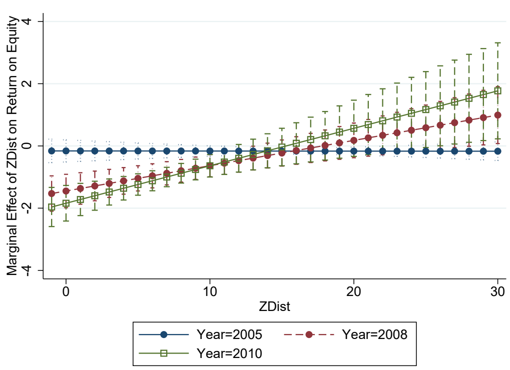

(b) Return on Equity

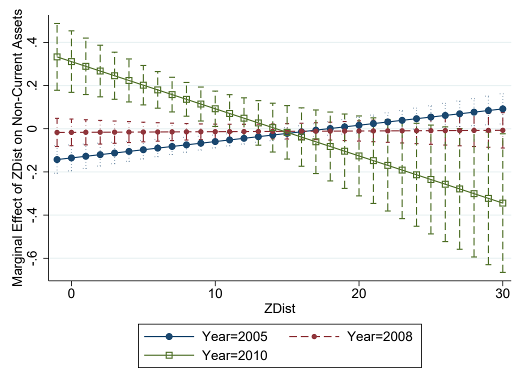

(c) Non-Current Assets

Figure 3.4: Marginal Effects, from Estimating Equation (3.15)

Notes: Figure presents marginal effects from estimating equation (3.15). Regressions include controls for the (log) number of branches, the capital-asset ratio, the share of non-interest income in total income, (log) total assets, bank-fixed effects, and year-fixed effects.

3.5.1 Robustness Checks

Table (3.4) begins a series of robustness checks. Row (1) excludes banks classified under either the *FDIC*'s mortgage lending specialization category or the *FDIC*'s consumer lending specialization category. This excludes banks with residential mortgage loans plus mortgage backed securities totaling more than 50 percent of the bank's total assets as well as banks with mortgage loans, credit card loans, and other loans to individuals in excess of 50 percent of the bank's total assets. This test is intended to make sure that mortgage lenders—the institutions hardest hit by the recession—are not driving our results. Row (2) includes controls for the deposit-weighted share of a bank's branches in the states hardest hit by the recession: California, Nevada, Arizona, Florida, and Massachusetts. This test is designed to assess whether or not our results are purely driven by a bank's direct exposure to disaster areas, rather than the potentially indirect effects of geographic information asymmetries more generally. Finally, Row (3) addresses the question of bank holding companies. Row (3) analyzes only institutions that are part of a larger bank holding company, and includes separate bank holding company-specific fixed-effects, with standard errors clustered at the bank holding company level. We include this test to address concerns about whether or not expansion of separate institutions belonging to the same bank holding company might have different effects than expansion at a single institution separate from any holding company.

The results from all of our robustness checks support our initial findings. In every case banks with a larger average distance between the bank headquarters and the bank branches experienced a larger decline in performance during the recession. This effect persists despite exclusion of mortgage lending specialists, controls from exposure to the regions most impacted by the crisis, and controls for bank organizational structure, in the form of bank holding companies.

3.5.2 Discussion

Does the negative impact of geographic expansion on bank performance during the Great Recession have broader economic implications or connections to a wider literature beyond banking? We believe that it does. In particular, this chapter's findings emphasize the (re-)emerging consen-

Table 3.4: Robustness Checks

	Return on Assets	Return on Equity	Non-Current Assets
(1) Exclude Mortgage Specialists			
$(ZDist \times After)_{it}$	-0.0493***	-0.563***	0.109***
	(0.0158)	(0.168)	(0.0281)
N	91,466	91,466	91,466
R^2	0.636	0.609	0.531
Controls	Y	Y	Y
Bank FE	Y	Y	Y
Specialization FE	Y	Y	Y
State $\times$ Year FE	Y	Y	Y
(2) Exposure to Crisis Areas			
$(ZDist \times After)_{it}$	-0.0562***	-0.657***	0.105***
	(0.0149)	(0.158)	(0.0244)
N	105,407	105,407	105,407
R^2	0.602	0.608	0.537
Controls	Y	Y	Y
Crisis Area Branch Share Control	Y	Y	Y
Bank FE	Y	Y	Y
Specialization FE	Y	Y	Y
State $\times$ Year FE	Y	Y	Y
(3) Bank Holding Companies			
$(ZDist \times After)_{it}$	-0.0574***	-0.682***	0.104***
	(0.0120)	(0.129)	(0.0302)
N	80,888	80,888	80,888
R^2	0.611	0.614	0.532
Controls	Y	Y	Y
Bank FE	Y	Y	Y
Bank Holding FE	Y	Y	Y
Specialization FE	Y	Y	Y
State $\times$ Year FE	Y	Y	Y

Notes: Standard errors in parenthesis. * 0.10 ** 0.05 *** 0.01. Rows (1) and (2) cluster standard errors at the bank level. Row (3) clusters standard errors at the bank holding company level. Control variables include (log) assets, (log) branches, the share of non-interest income in total income, and the capital-asset ratio.

sus that geography *does* matter for economic development, and the joint-picture painted by the empirical results and theoretical model offers some potential policy prescriptions in light of this.

In contrast to early predictions of “flattening" economies, regional and national, in the era of information technology (e.g. Friedman, 2005; Fukuyama, 1992), geography seems to be determining economic prospects to an ever-greater degree (e.g. Ganong and Shoag, 2018; Giannone, 2018). The twin elections of 2016 in the UK for Brexit and in the US for president served to underscore the effect of regional disparities—not just in influencing local economies—but in sentiments towards national and international politics. These disparities manifest themselves in institutional structures that make the asymmetries between regions all the more endemic, as in the case of banking explored in this chapter. The aggregate productivity slowdown in the US has been tied to the dramatic decline in small young firms, which are also the source the majority of job creation (Alon et al., 2018). Entrepreneurs often find themselves dependent on bank loans and credit lines for financing at critical moments in their maturation process, and the access and pricing of such loans in turn depend critically on loan officers’ assessment of the relative riskiness of the project in question.

The finding that distance between bank headquarters and branches has a systematically negative impact on bank performance indicates that local information on project prospects is valuable for banks themselves — as well as the regions in which they are located. The soft information on local contexts, trends, and networks has long been seen as crucial in determining the relative success of bank loan portfolios (e.g. Agarwal et al., 2018; Pham, Talavara, and Tsapin, 2018; Uchida and Udell, 2012), yet the consolidation of the banking industry is effectively generating a natural experiment on the impact of the loss of those local banking relationships. Our findings suggest that soft information is indeed a valuable intangible asset, and that recent consolidation needs to better assess these costs against the widely-touted scale economy benefits of a more concentrated banking system.

The finance findings also reinforce previous work on the impact of a geographical form of Akerlof (1970)’s classic information asymmetry conception. Recent work (e.g. Bunten et al., 2015; Moore, Petach, and Weiler, 2018) has demonstrated that the information generated by local market business activity, the highly-correlated opening and closings of establishments, has significant implications for local employment growth. More dynamic markets have greater gross streams of information through both the successes and failures of establishments, providing more information to bankers and other potential investors (Weiler, 2000). This information crucially shapes risk perceptions in such settings as determined by the perceived variance of project returns, evaluated on past experiences in similar settings. A greater number of data points leads to lower variance and thus a better chance for financing approval and favorable loan conditions. Insurance policy decisions follow similar criteria.

The ability to tap sources of soft information depends on having creditors that understand and can evaluate such local prospects fully (e.g. Weiler, Hoag, and Fan, 2006). Bank headquarters using standardized loan scoresheets for their branch networks are much less likely—in large part because their generalized structures are far less able—to incorporate the intangible contextual information about a particular project in a particular geographic market. As Anenberg et al. (2018) note in their recent Fed Notes policy brief, “the local branch is becoming less important for acquiring soft information, possibly caused by changes in underwriting technology (e.g. credit scoring)."

Therefore, the consolidation of the banking system is reinforcing systemic *geographic information asymmetries*, which are not only yielding sub-optimal regional growth trajectories but also sub-optimal performance in broadly-branched banks during macroeconomic downturns. The geographic consolidation of headquarters is distancing banks from critical local information in precisely those often-rural areas where there are fewer scoresheet-ready data points but also where bank loans are most effective in sparking successful entrepreneurial projects (Conroy, Low, and Weiler, 2017).

Given the empirical relevance of the negative spillovers generated by bank expansion for bank performance, the model presented in Section (3.3) can be used to evaluate the policy implications of bank expansion. In particular, the model suggests that any policy which reduces the size of the branch network may be welfare improving. Consider the model again under the assumption of a linear cost function. It is easy to show that the optimal allocation can be decentralized by setting a unit-tax on the number of branches. The tax that implements the optimal allocation, under the assumption of linear costs (taking $\gamma = 1$ to simplify), is given by:

$$1 + \tau = \frac{\alpha}{\alpha + \beta} \tag{3.16}$$

The size of the tax is proportional to the two key elasticities in the model. The size of the tax is decreasing in the branch elasticity of loan volume (as long as $|\beta| < |\alpha|$, which must hold if the optimal allocation is interior), and increasing in the absolute magnitude of the branch elasticity of the probability of repayment (note that since the latter is assumed to be < 0 an increase in β corresponds to a decrease in the extent of the externality, so that a decrease in β, corresponding to an increase in the negative spillover from the number of branches to the probability of repayment, results in a higher optimal tax rate). Further, note that as β increases in absolute value we have:

$$\lim_{\beta \to -\alpha} 1 + \tau = \infty \tag{3.17}$$

Which provides intuition for why $|\beta| < |\alpha|$ must hold at an interior solution. As β increases in absolute value the additional benefit from a new branch—captured by the increased lending capacity, α—is more than offset by the decrease in the probability of repayment that results from negative spillovers which may occur as a result of the lemons effects, within branch network effects, between branch network effects, and loss of soft information, that accompany branch network expansion. If β becomes large enough in absolute value, such that $|\beta| \geq |\alpha|$, the optimal solution is to set $n = 0$, which is accomplished by setting the tax as in (3.17). While it is unlikely that the size of the externality is this large in practice, the scenario is nonetheless instructive.

Several caveats to the above discussion is required. First, the empirical results in this chapter—despite demonstrating that a negative relationship between the number of branches a bank operates and bank performance does exist—indicate that it is not just the number of branches that matter for bank performance, but the size of the geography over which those branches are spread. The existence of geographic information asymmetries, and the ability of the branch network to act as a vehicle for the transmission of economic shocks from one location to another, are the source of the negative externality, suggesting that rather than a proportional tax on the number of branches, a tax which scales in value with the degree of spatial expansion of the branch network is preferable.

Second, alternative methods of regulation are possible. The obvious sort would be something akin to the branch banking restrictions which existed prior to the 1980s—although the political feasibility of such regulations is questionable, and direct regulation is generally less efficient than the Pigouvian-type taxation discussed above. Nonetheless, the results of the empirical section of this chapter make clear that banks with large geographic networks were susceptible to increased negative performance shocks during the crisis, suggesting that—although regulators may have taken into account reductions in idiosyncratic risk that occur with a geographically diversified lending portfolio—insufficient attention was paid to the potential for *systemic* risk when the bulk of branch banking restrictions were removed, and restoration of similar provisions may be welfare improving.

3.6 Conclusion

Geographic expansion negatively impacts bank performance during crises. Using *FDIC* data on bank performance and branch office locations, this chapter shows that an increase in the deposit-weighted average distance between a bank headquarters and its branches results in worse performance during the Great Recession—as measured by the share of non-current assets in total assets, the return on assets, and the return on equity. To the extent that banks with large branch networks fail to internalize the risk associated with branch network expansion, some form of preventative regulation may be socially desirable. While modern portfolio theory provides a rationale for removing restrictions on branch banking by highlighting reductions in idiosyncratic risk via diversified lending portfolios, the results in this chapter suggest that insufficient attention has been paid to the information frictions involved with branching and the associated costs of systemic risk created by branch banking deregulation.

Bank performance is fundamentally determined by the accurate assessment of risk, which itself rests on the quality and quantity of information a bank possesses about its portfolio. Geographic information asymmetries generated by branch networks are a novel and empirically significant basis on which to more fully understand the costs and benefits of banking consolidation. A natural extension of the research trajectory evaluating the contribution of branching to performance would be to explore the parallel impact of mergers. Assessing the effect of pre-recession mergers on post-recession performance by the distance between the acquiring and outgoing institution would provide an additional test of the geographic information asymmetries hypothesis suggested here.

Chapter 4

Distribution and Capacity Utilization in the United States: Evidence from State-level Data

4.1 Introduction

This chapter is yet another entry in the seemingly interminable debate on wage- versus profit-led growth. It is difficult to forge new paths on well-tread ground, but doing so is easier if one has new instruments with which to set out. In this chapter I bring along some such instruments—both figuratively and literally. Using a panel of U.S. states for the years 1974-2014, I estimate a demand and distribution system. To identify the effect of changes in the labor share on capacity utilization, I make use of variation in state-level redistributive policy—in particular, I exploit variation in the statutory minimum wage across states as an instrumental variable. In the absence of such an instrument for capacity utilization, I estimate the distributive curve non-parametrically. I find the U.S. is strongly wage-led at the state-level, and there are significant non-linearities in the distributive curve, with the U.S. exhibiting profit-squeeze dynamics at low levels of capacity utilization and wage-squeeze dynamics at high levels (although the shape of the distributive curve may vary across states). I discuss the implications of these results as offering two policy suggestions. First, even if the United States is profit-led at the national level, the wage-led results from the state-level panel suggest that regional governments should pursue redistribution in favor of wages, regardless of national macroeconomic policy. Second, given the first policy suggestion, state-level wage-led policy resembles a coordination problem, where regional policymakers—potentially in states operating at different equilibria—must take into account not only the direct impact of redistributive policy on capacity utilization, but the behavioral responses of firms and households that move capacity utilization in the opposite direction. That is—because the labor share is not a policy variable—it may be difficult to bring about increases in demand through redistribution, even when demand is wage-led.

The rest of the chapter is organized as follows. Section 2 reviews the wage- versus profit-led growth debate. Section 3 describes the data. Section 4 outlines the empirical model and presents the results. Section 5 concludes.

4.2 Demand and Distribution, Revisited

With the publication of their seminal paper, Bhaduri and Marglin (1990) opened the floodgates to an army of empirical researchers seeking to answer the question: is demand wage- or profit-led? The central insight of Bhaduri and Marglin (1990) was that the investment decision of a firm may depend independently on each of the individual components of the profit rate—capacity utilization and the profit-share—rather than just the profit-rate itself. The rationale for this change is that firms care about both the potential money earned from new capacity (the profit share) and the likelihood of selling additional goods (capacity utilization) (Marglin, 2017). Allowing the investment function to vary in this way, the usual results of the Kaleckian model are turned on their head. Rather than being unambiguously wage-led, the model generates theoretically ambiguous results regarding the relationship between income distribution, aggregate demand, and growth. The impact of a redistribution toward wages on the level economic activity becomes an empirical question.

Stockhammer (2017) identifies two different approaches to the estimation of demand and distribution systems: the “Neo-Kaleckian" approach, and the “Neo-Goodwinian" approach. Neo-Kaleckians estimate the relationship between demand and distribution via a series of behavioral equations corresponding to the components of total value added—i.e. a consumption function, an investment function, a government spending function, and a net export function—and identify the impact of a redistribution toward wages as the sum total effect from each of the different behavioral equations. The Neo-Kaleckian approach may be critiqued on a number of grounds related to the validity of the econometric “identification" strategy adopted by its authors. First, the parameter estimates in such models are likely biased due to endogeneity arising from the simultaneity of the right-hand side variables. Second, common shocks to the components of value-added may mean cross-equation correlation in the error terms is unaccounted for. As Blecker (2016) notes, some of these problems may be addressed by a simultaneous equations approach (such as seemingly unrelated regression), but in most cases those adopting the Neo-Kaleckian methodology have opted to simply estimate the system of behavioral equations via equation-by-equation ordinary least squares (OLS), thus not addressing the bias. Further, even in the case where cross-equation correlation in the error terms is controlled for, it is not clear there exist instrumental variables at the level of national accounts with which endogeneity can be satisfactorily be addressed.

In contrast, the Neo-Goodwinian approach—which I adopt in this chapter—starts from a reduced form system of demand and distribution based on Goodwin (1967)'s theoretical model of the business cycle. The Goodwin business-cycle model can be summarized in the following two-dimensional system of differential equations in capacity utilization (u) and the labor share (ψ) (using hats to denote growth rates):

$$\hat{u} = f(u, \psi) \tag{4.1}$$

$$\hat{\psi} = g(u, \psi) \tag{4.2}$$

The solution to this system can be depicted by a plot of the nullclines in $u - \psi$ space. The slopes of the $\hat{u} = 0$ and $\hat{\psi} = 0$ curves are characterized by several different possible regimes. A downward sloping demand nullcline—such that utilization is decreasing in the wage-share—is said to be "profit-led" or "exhilarationist", while an upward sloping demand nullcline is "wage-led", "underconsumptionist", or "stagnationist." On the other hand, an upward sloping distribution curve—such that the wage share is increasing in the rate of utilization—is described as a "profit-squeeze" regime, while a downward sloping distribution curve is of the "wage-squeeze" variety. Figure (4.1) presents an example of a profit-led/profit-squeeze regime, with the dotted line illustrating the transitional dynamics (Rada and Kiefer, 2015).

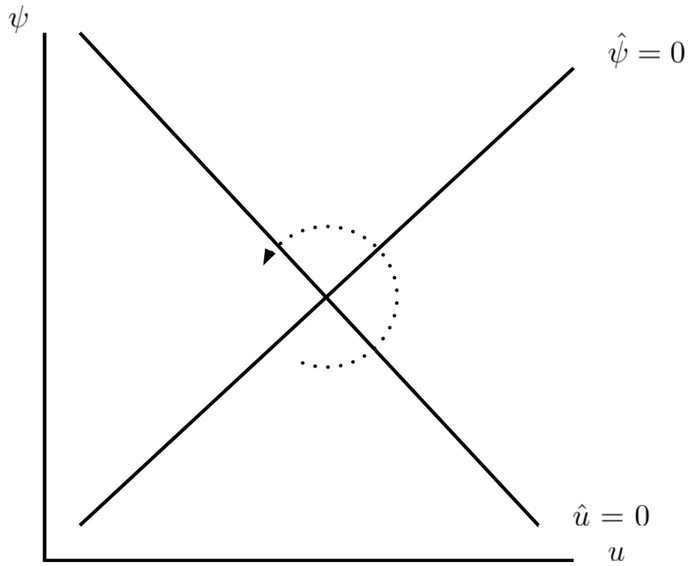

Figure 4.1: Profit-Led, Profit-Squeeze Dynamics in the Goodwin Model

The simultaneous nature of the demand- and distribution- system means that the Neo--Goodwinian approach is not without its own identification challenge. Various different identification strategies have been adopted by authors taking the Neo-Goodwinian approach. Examples include Barbosa-Filho and Taylor (2006), Rada and Kiefer (2015), and Nikiforos and Foley (2012). The first two of the aforementioned papers adopt a vector auto-regression approach to address endogeneity, the latter paper estimates the demand and distribution system separately, using lagged values of the dependent variables as instrumental variables for their current values. Barbosa-Filho and Taylor (2006) and Nikiforos and Foley (2012) estimate the demand and distribution system for U.S. data, Rada and Kiefer (2015) use a panel of OECD countries. They all reach the conclusion of a profit-led growth regime for their given sample. Further, although the sample demand and distribution system plotted in Figure (4.1) is linear, Nikiforos and Foley (2012) and Tavani, Flaschel, and Taylor (2011) provide evidence that the distributive curve may have substantial non-linearities about it, suggesting the possibility of multiple equilibria in the demand and distribution system.

Other recent papers have extended the analysis of demand and distribution regimes to allow for an independent effect of changes in the *personal distribution of income* on the relationship between the functional distribution of income and capacity utilization. Carvalho and Rezai (2016) and Palley (2016) allow wage inequality to impact both the level of demand and the demand-regime via differential marginal propensities to save across varying classes of wage-earners—a channel absent in the standard Kaleckian model. In Carvalho and Rezai (2016)'s model, higher wage inequality moves the demand regime in the direction of profit-led growth. The same result obtains in Palley (2016)'s model, however Palley (2016) distinguishes between inequality in wages and inequality in profits. Greater wage inequality pushes the economy in the direction of a profit-led regime, but greater profit inequality does the opposite: increasing capitalists' ownership share makes the economy more likely to be wage-led as a result of lower capitalist propensity to consume. In either case, the key insight is that changes in the personal distribution of income may have significant bearing on the relationship between the functional distribution of income and demand. This is especially important for the empirical exercise in this chapter, because changes in the minimum wage may impact the personal, as well as functional distribution of income. In Appendix (B.4), I extend the empirical model to allow for interactions between the functional and personal distribution of income. I find results consistent with those in the body of the chapter, but—in line with Carvalho and Rezai (2016)—I find that the state-level demand, and the demand regime, may be sensitive to the level personal income inequality.

Finally, Barbosa-Filho (2016) extends the empirical analysis of demand and distribution regimes by adopting the *employment rate*—the ratio of the employed population to the labor force—as a proxy for aggregate demand, rather than the commonly used measure of capacity utilization (obtained as the ratio between actual and potential output). This avoids the critique of the standard Neo-Goodwinian approach that the use of capacity utilization obtained via a statistical filter implies mean-reversion in demand following a distributive shock, due to the stationarity of the capacity utilization variable. Because one is presumably interested in the impact of permanent or quasi-permanent shifts in the labor share (induced by policies which alter the institutional arrangements governing capitalist-worker bargaining, for example) the use of this type of measure may be unsatisfactory. In contrast, the employment rate implies no-such mean reversion, which allows for the simultaneous appearance of a cyclical profit-led/profit-squeeze pattern, accompanied by a long-run reduction in both demand and the labor share. In Section (4.4.2) I address this critique by estimating a version of the demand curve with the employment rate, rather than capacity utilization, as the dependent variable. Unlike Barbosa-Filho (2016), who finds evidence of a profit-led employment curve for the United States, state-level data suggest wage-led employment, consistent with the results found for capacity utilization.

Despite the dominance of profit-led results in the empirical literature, the long-run negative trend in the wage-share found by Rada and Kiefer (2015) coincides with a persistent downward trend in the estimated equilibrium level of capacity utilization (measured by the output-gap) for OECD countries over time. The authors suggest the presence of a 'race to the bottom' among OECD countries in unit-labor costs in the pursuit of increased exports. The set of policies adopted in pursuit of these gains include reductions in corporate tax rates, deregulation of financial markets, an erosion of anti-trust laws, and the continued application of inflation targeting monetary policy. The net impact of these policies seems to have been a persistent under-utilization of capacity, the results of which may have a hysteresis about them which bears implications for the long-run. The 'race to the bottom' is characterized by a fallacy of composition, whereby the quest to increase utilization through export demand stimulated via lower unit-labor costs results in a lower equilibrium level of utilization when the strategy is pursued by all OECD countries simultaneously.

Arnim, Carvalho, and Tavani (2014) rationalize this fallacy of composition via an analysis of the distributional make-up of global demand. In particular, the authors make the the case that global demand is likely to be wage-led. The profit-led demand channel appears in open economies because a redistribution toward wages increases demand for imports (and hence lowers net-exports), decreasing capacity utilization. However, because the world balance of trade must sum to zero, this channel is absent for global aggregate demand. Even though individual countries may be profit-led, the world as a whole faces a wage-led effective demand problem—a result which holds for a wide range of behavioral assumptions. Razmi (2018) challenges this conclusion, arguing that compositional effects of trade may skew the world to be profit-led, even though global trade is balanced.

Despite this disagreement, the argument presented by Razmi (2018) offers support for the possibility suggested by Arnim, Carvalho, and Tavani (2014), namely: that if some (or possibly a majority) of the constituent parts of an economy are characterized by one type of demand and distribution regime, the same regime need not apply to the aggregate economy. This dichotomy arises because economic aggregates (the world, a country) are not isomorphic to their constituent parts (a country, a region), and may have emergent properties which do not appear at lower levels of aggregation.

In this chapter, I suggest that within countries the demand and distribution system may exhibit dynamics which invert the fallacy of composition outlined above. The empirical findings in Section (4.4) show that a country that appears profit-led at a national level—e.g. the United States—can have regions within it which are characterized by wage-led demand. This will happen, for example, if the country is profit-led through the trade-balance, but at a regional (state, province, metropolitan area, etc.) level demand for local non-tradeable goods is highly dependent on wages. However, if capital is highly mobile within countries, and tax policy varies across regions, difficulties for regional wage-led policy arise. Any regional government which raises taxes to finance a redistribution toward wages will face the prospect that firms shed jobs, cut wages, and move production activities to another region within the country with lower tax rates. In this case the problem of wage-led policy is characterized by a dual coordination failure. At the country level, the pressure to lower unit-labor costs may produce the 'race to the bottom' dynamics discussed by Rada and Kiefer (2015)—despite the fact that all countries would be better off if they could achieve coordinated redistribution toward wages (assuming global demand is, in fact wage-led). At a regional level, no government has an incentive to unilaterally raise tax rates to finance redistributive policy (for example, greater expenditure on worker benefits) if they fear this may lead to a loss of local jobs (or cuts in wages). Even if regions within a country can successfully coordinate tax policy to re-distribute toward wages, if the country is profit-led at a national level this may lower national growth unless there is some form of pre-existing cross-country coordination.

These concerns are motivated by empirical findings which rely on an instrumental variables estimation strategy—specifically, the use of state-level variation in minimum wage policy. I argue that this strategy constitutes an improvement over past attempts to estimate demand and distribution systems for the United States. The estimation strategies used in the literature generally lack a source of truly exogenous variation in the labor share. This is a problem, because the demand-driven Goodwin model used in the literature (Nikiforos and Foley, 2012; Rada and Kiefer, 2015) suggests demand and distribution are simultaneously determined. As an example, Nikiforos and Foley (2012) use lagged-values of the endogenous variables as instruments in a 2SLS regression, which—while predetermined—are not truly exogenous. Further, commonly used, seemingly innocuous data adjustments may actually have a dramatic impact on attempts to estimate the demand and distribution regime for an economy. Consider the following two techniques used by Nikiforos and Foley (2012). First, they obtain estimates of capacity utilization and the cyclical component of the labor share using the Hodrik-Prescott (HP) filter. The HP filter is known to introduce spurious dynamic relations that have no basis in the underlying data-generating process, create large differences in filtered values depending on their position in the sample, and to make use of smoothing parameters that are at odds with those suggested by a statistical formulation of the problem (e.g. the commonly used smoothing parameter of 1600 for quarterly data) (Hamilton, 2017).

Second, Nikiforos and Foley (2012) look only at the effect of the cyclical component of the labor share on capacity utilization. This is at odds with the intuition behind both the Kaleckian and Goodwinian presentation of the problem, which are presumably focused on how a policy-induced permanent or quasi-permanent shift in the labor share affects demand. Estimating the demand and distribution system using only the cyclical component of the labor share strips the results of any policy implication, because such cyclical movements are presumably *not* induced by the kind of institutional change implicit in the original theoretical analysis. While these techniques are not unique to Nikiforos and Foley (2012), they are representative of potential sources of bias in past estimates of the demand and distribution system for the United States (or other countries).

I overcome these problems by applying an alternative filtering method suggested by Hamilton (2017), and exploiting policy variation across U.S. states as a source of plausibly exogenous variation in the actual labor share (rather than just the cyclical component). I also estimate the distributive curve non-parametrically, with an eye toward potential non-linearity. Finally, I extend the empirical analysis to alternative measures of demand (the employment rate) and long-run growth (growth of state-level per-capita income), as well as testing the interaction between the personal distribution of income and the demand regime, â la Carvalho and Rezai (2016) and Palley (2016) (See Appendix (B.4)).

4.3 Data

The data in this chapter come from two main sources. First, data on state-level output, labor share, and tax revenue come from the Bureau of Economic Analysis (*BEA*) Regional Economic Accounts. The labor share is calculated as the ratio of wage and salary compensation to the sum of wage and salary compensation and the gross operating surplus of the business sector at the state-level[13]. Capacity utilization is derived from the *BEA* series for state-level output. Second, I make use of the annual maximum state-level minimum wage value in the series reported by Vaghul and Zipperer (2016). The final sample is comprised of annual data for the fifty U.S. states and Washington D.C. for the years 1974-2014. Table (4.1) presents the sample means for the main variables. All dollar values are converted to real terms using the U.S. CPI for all urban consumers from the Bureau of Labor Statistics[14].

To construct a series for capacity utilization that avoids the problems associated with the HP filter I make use of an alternative method for obtaining the transient component of a time series suggested by Hamilton (2017). Technical details are included in Appendix (B.2). Figure (4.2)

[13]Since the wage-share is calculated in terms of GDP (in particular, GDP less net taxes on production and imports) a change in business current transfer payments or depreciation can impact it. To avoid this, one would like to work with GNP at producer prices. Unfortunately the author knows of no such series at the state-level.

[14]With the exception of the minimum wage, which is kept at its statutory value in order to capture the effect policy changes, rather than changes in real values which may occur due to inflation. Using nominal values for the minimum wage in this manner is consistent with common practice in the minimum wage literature, e.g. Dube et. al. (2010).

Table 4.1: Sample Means

	Mean	Std. Dev.
u_{it}	0.993	0.113
ψ_{it}	0.605	0.0568
MW_{it}	\$4.67	1.7
Taxes_{it}	\$326.4	209.6
N	2091	

Notes: Observations indexed by state, i, and year, t. u_{it} gives the value of capacity utilization—defined as the ratio of actual real output to potential real output. ψ_{it} gives the value of the state labor share. MW_{it}, gives the statutory value of the maximum annual state minimum wage. Taxes_{it} gives the per-capita real value of personal current taxes paid to the state government.

presents plots of capacity utilization against the cyclical component of the labor share for both the Hamilton (2017) filtering method and the HP filter. I set the smoothing parameter for the HP filter equal to 100, a value commonly used in the case of annual data.

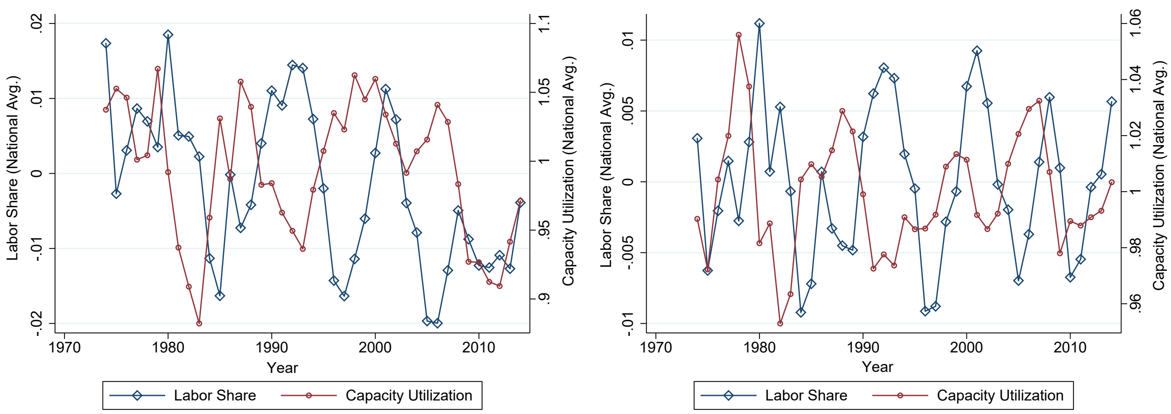

(a) Hamilton (2017) Filter

(b) HP Filter

Figure 4.2: Comparison of Hamilton (2017) and HP Filter. Capacity Utilization and the Labor Share: 1974-2014.

Comparison of the two figures generates three insights worth discussing. First, the HP filter tends to understate the extent of cyclical fluctuations relative to the Hamilton (2017) filter, a result which means that any attempt to capture business cycle dynamics with an HP filter will understate the severity of recessions and upswings. As an example, during the Great Recession, the trough value for capacity utilization generated by the HP filtered series is only 0.98, while the low point in the Hamilton (2017) series is closer to 0.9. Given the known severity of the Great Recession, the estimates generated by the HP filtered series seem like an understatement. Second, while the relationship between capacity utilization and the labor share produced by each filtering method is relatively similar at the beginning of the series, the HP filter generates markedly different dynamics at the end of the series (approximately post-1995) when compared with the Hamilton (2017) filter (a result that is in line with Hamilton (2017)'s critique). In particular, it appears that the HP filtered series overstates the movement of the labor share relative to capacity utilization in the later years of the sample, a dynamic which is bound to influence estimation results[15].

The third noteworthy feature of these figures is that in each case the nature of the relationship between the labor share and capacity utilization appears to change over time. Figure (4.2a) suggests that in the period beginning in 1974 up until the late 1980's, capacity utilization and the labor share moved together. Then, during the late 1980's and early 1990's, a shift takes place, with capacity utilization and the labor share begin to move inversely with one another. Finally, following the financial crisis and recession, it appears the two have again begun to move in tandem. This suggests the relationship between demand and distribution is likely not a fixed characteristic of a particular economy, but may be influenced by various historical and institutional factors.

Along with the improved capacity utilization measure, the second advantage of the data presented in this chapter is the ability of regional data to exploit variation in the labor share across states within the U.S., as well as variation across time. Figures (4.3) and (4.4) present state-level maps of the labor share by state for the years 1974 and 2016. The maps demonstrate what regional scholars have long known—the U.S. is not an economic monolith, and income distribution is no exception to this rule. Different regions demonstrate markedly different characteristics with respect to the distribution of wages and profits, and the policies and institutions that govern this distribution—rules governing collective bargaining, right to work laws, minimum wages, job

[15]A conjecture I confirm through empirical testing. Although not reported here, I find that—in regressions of the demand equation using the cyclical component of the labor share—use of the cyclical component generated by the HP filter biases the results in the direction of profit-led growth, with regression coefficients that are statistically significant, negative, and larger in absolute magnitude than the coefficients estimated using the cyclical component of the Hamilton (2017) series. Results available upon request.

training programs and educational subsidies, etc.—may vary substantially across states (one does not expect to find the same worker protections in Alabama or Mississippi—both right to work states[16]—that one finds in California).

In both 1974 and 2014 coastal states appear to have a slightly higher labor share on average, compared to states in the interior of the country. However, states on the East coast (and some of the Great Lakes states) outrank both their Midwestern and Western counterparts. As an example, states with larger-than-average values for the labor share in both years include Hawaii (not pictured), Maine, Massachusetts, New Hampshire, New Jersey, Rhode Island, Virginia, and Wisconsin. The region with the highest labor share in both 1974 and 2014 is Washington D.C., with values of approximately .79 and .71 for 1974 and 2014, respectively.

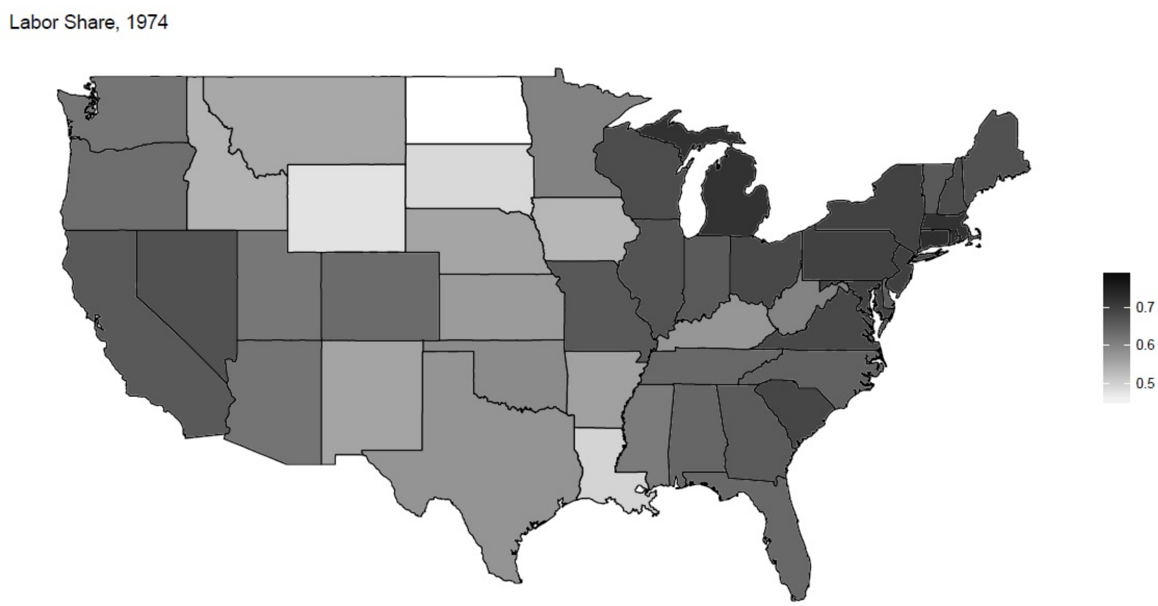

Figure 4.3: Mapping the Labor Share Across U.S. States: 1974

Forty-one of the fifty states and D.C. experienced a net decline in the labor share between 1974 and 2014. The ten states that did not experience a decline in the labor share were Arkansas, Idaho, Iowa, Kansas, Kentucky, Louisiana, Minnesota, North Dakota, South Dakota, and Wyoming. How-

[16]And not just in a statutory sense. Both the aforementioned states have adopted constitutional amendments in support of the “right to work".

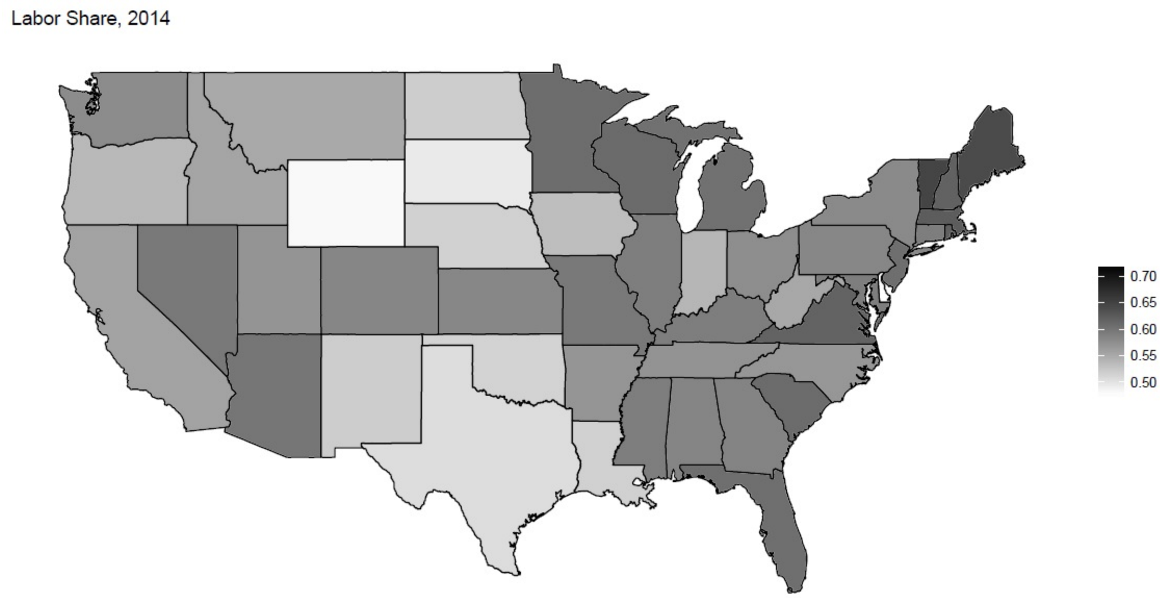

Figure 4.4: Mapping the Labor Share Across U.S. States: 2014

ever, the level of the labor share in most of these states is below average, despite experiencing an increase between the years 1974 and 2014. Further, in most cases the increase experienced in each state was small—less than one percentage point. The exception to this rule is North Dakota, which experienced close to a seven percentage point increase in the labor share between 1974 and 2014, likely due to the oil and natural gas boom experienced in the state in the latter part of that period. Finally, it is worth noting that the size of the total population in 2014 residing in states that experienced an increase in the labor share is approximately 27 million, or less than 10% of the entire U.S. population (the latter of which totaled 318.6 million in 2014).

Of the forty-one states that experienced a decline in the labor share, eleven states experienced a decline greater than ten percentage points. These states include Alaska (not pictured), California, Connecticut, Delaware, Indiana, Maryland, Michigan, New York, Ohio, Pennsylvania, and Rhode Island. Not surprisingly, several of the states on this list are “Rust Belt" states that were former centers of United States manufacturing employment (i.e. Michigan, Ohio, Pennsylvania, Indiana). The size of the total population residing in states that experienced at least a ten percentage point decline in the labor share is approximately 111 million, or 35% of the United States’ population in 2014. Looking at regional data on the labor share thus illuminates patterns not visible when looking at the aggregate data. Estimates from aggregate data have the share of labor compensation in GDP for the U.S. declining from approximately 0.65 in 1974 to 0.60 in 2014[17]. This decline does not capture the full extent of the magnitude felt by individuals residing in heterogeneous geographies. To the extent that local labor markets are a better measure of an individual's economic opportunities, the past forty years have been worse for many wage earners than the aggregate data suggests.

4.4 Estimation

4.4.1 Demand

I begin by estimating the following demand equation characterizing the relationship between capacity utilization and the labor share:

$$u_{it} = \beta_0 + \beta_1 u_{i,t-1} + \beta_2 u_{i,t-2} + \beta_3 \psi_{it} + \beta_4 ln(Taxes_{it}) + \eta_i + \delta_t + \epsilon_{it} \tag{4.3}$$

Where observations are indexed by state, i, and year t, u_{it} gives the rate of capacity utilization, ψ_{it} gives the state labor share, $Taxes_{it}$ gives the real per-capita personal taxes collected by the state, η_i and δ_t give state- and year-fixed effects, respectively. Tax revenues collected by the state are expected to move pro-cyclically as the state budget surplus (deficit) increases during booms (recessions). Following previous authors, I include lagged values for capacity utilization as independent variables in the estimation equation.

First, I estimate (4.3) using the method adopted by Nikiforos and Foley (2012) (from here on, *NF*). That is, I de-trend the labor share and state tax revenue data (using the Hamilton (2017) filter), and make use of lagged values of the labor share and de-trended labor productivity as instruments for the contemporaneous labor share value. I omit any explicit control for time in the form of fixed-effects or linear trend, to keep the specification consistent with *NF* whose only control for

[17] *Source:* Penn World Table 9.0.

time is in the use of the cyclical components of the variables of interest. Table (4.2) presents the results for this baseline estimation.

Table 4.2: Estimation Results - *NF* Specification

	(1) OLS	(2) *NF* IV	(3) *NF* IV + FE
$u_{i,t-1}$	0.833***	0.928***	0.874***
	(0.0209)	(0.0231)	(0.0606)
$u_{i,t-2}$	-0.0993***	-0.126***	-0.138***
	(0.0203)	(0.0219)	(0.0477)
$\psi_{i,t}$	-0.980***	-0.101	-0.259***
	(0.0489)	(0.0693)	(0.0863)
$ln(Taxes_{it})$	-0.000206	-0.000634	0.0316***
	(0.00130)	(0.00140)	(0.0117)
N	1989	1989	1989
R^2	0.755	0.716	0.697
State FE	N	N	Y
Time FE	N	N	N
First-Stage F-Stat	—	1353.28	1289.74
Sargan-Hansen Test (p-value)	—	0.0024	0.0413

Notes: Standard errors in parenthesis. Column (3) uses cluster-robust errors, clustered at the state-level. * 0.10 ** 0.05 *** 0.01. ψ_{it} gives the cyclical component of the labor share, based on the Hamilton (2017) filter.

Column (1) presents results from OLS estimation, Column (2) presents instrumental variables estimates, Column (3) produces IV estimates with the inclusion of state-fixed effects and cluster robust standard errors. I report the results of two tests of IV robustness. First, I report the first-stage F-Statistic, which Staiger and Stock (1997) argue needs to exceed 10 to avoid problems associated with weak instruments. Second, I report the p-value from the Sargan-Hansen test of overidentification. For instruments to be valid they must satisfy the orthogonality assumption, for a vector of instruments Z_{it}: $E[Z_{it}^T \epsilon_{it}] = 0$—that is, the instruments must be uncorrelated with the regression error term. The Sargan-Hansen test assess the null hypothesis that the instruments are in-fact uncorrelated with the errors. A rejection of the null hypothesis indicates that the orthogonality assumption is violated.

The results from the *NF* specification suggest that demand is profit-led, in line with the original findings of Nikiforos and Foley (2012). An increase in the labor share is associated with a decline the rate of capacity utilization in all three columns, with the estimates in Column (1) and Column (2) taking on statistical significance. However, in both IV specifications the null hypothesis of the Sargan-Hansen test is rejected at the 5% significance level, indicating that the lagged values for the cyclical component of the state-level labor share may be correlated with the error term, potentially influencing the estimates.

The control variables behave mostly as expected, and in a manner similar to previous estimates of demand and distribution systems in earlier papers. Once state-fixed effects are controlled for in Column (3), the tax revenue variable moves in the same direction as the rate of capacity utilization, capturing the counter-cyclical nature of government budget deficits. Current capacity utilization is positively associated with lagged values, on net.

In addition to the somewhat unsatisfactory performance of the *NF* instrumental variables in the state-level panel, estimation of (4.3) using only the cyclical component of the labor share faces a larger methodological critique: it is not only cyclical fluctuations in the labor share that matter for capacity utilization and growth. In fact, both the Neo-Kaleckian and Neo-Goodwinian formulation of problem are motivated in-part by a larger interest in how permanent or quasi-permanent variation in the institutional environment through changes in *policy* may impact the distribution of income and subsequently growth. To the extent that this source of variation is excluded by looking only at the cyclical component of the labor share, the model is stripped of its most interesting component.

To address the flaws in previous attempts to estimate the demand and distribution system for the United States, I adopt a different strategy for identifying the causal effect of the labor share on capacity utilization. I use state-level variation in wage-policy—specifically, the statutory minimum wage—as an instrumental variable for the labor share. The identifying assumption is that changes in the minimum wage only affect demand through their impact on distribution. This effect may be positive or negative, and may operate directly and indirectly (e.g. increasing workers' reservation wage)—all that matters is that the policy change does in some way alter the distribution of income. This strategy improves over past estimates of demand and distribution systems in at least two respects. First, neither attempts to estimate the demand and distribution system via vector autoregression (Barbosa-Filho and Taylor, 2006; Rada and Kiefer, 2015) nor IV estimates via lagged endogenous variables (Nikiforos and Foley, 2012) identify variation in the labor share as an explicit result of policy variation. Second, the use of the full value of the labor share—rather than just its cyclical component—avoids the potential issues related to filtering the data and allows a policy-based interpretation of the results. In the place of using the cyclical component, time is controlled for via the use of time-fixed effects. Table (4.3) presents the results from estimating equation (4.3) with minimum wage based instruments. I use the current value for the maximum annual minimum wage in state i, in year t, and an indicator variable for whether or not the state experienced an increase in the statutory minimum wage $\geq \$0.25$.

Table 4.3: Estimation Results - Minimum Wage Instruments

	(1) Cyclical Spec. IV	(2) IV1	(3) IV2	(4) MW IV + *NF* IV
$u_{i,t-1}$	0.935***	0.904***	0.936***	0.928***
	(0.0782)	(0.0615)	(0.0653)	(0.0526)
$u_{i,t-2}$	-0.153**	-0.125**	-0.0557	-0.0996**
	(0.0649)	(0.0600)	(0.0457)	(0.0467)
$\psi_{i,t}$	0.283	0.470***	0.903*	0.230***
	(0.276)	(0.109)	(0.533)	(0.0637)
$ln(Taxes_{it})$	0.0308***	0.0323***	0.00687	0.0120*
	(0.0113)	(0.00830)	(0.00583)	(0.00642)
N	1989	1989	1989	1989
R^2	0.631	0.640	0.709	0.762
Time FE	N	N	Y	Y
State FE	Y	Y	Y	Y
First Stage F-Stat	29.44	22.50	4.33	792
Sargan-Hansen Test (p-value)	0.20	0.29	0.72	0.22

Notes: Standard errors in parenthesis, clustered at the state-level. * 0.10 ** 0.05 *** 0.01. Column (1) presents results for minimum wage instruments on cyclical data as a baseline. All other columns use the full value of the labor share as the variable of interest. Column (4) includes both minimum wage instruments and a single lag of the labor share as a means of addressing the low first-stage F-statistic in Column (3).

In every specification the results suggest the conclusion of wage-led growth at the state-level for the United States. Column (1) uses the cyclical component of the labor share for comparison with the *NF* specification tested above. It is clear that the minimum wage instruments perform better in the cyclical specification as compared to the *NF* instrument vector of lagged values and labor productivity. The first stage F-stat and the Sargan-Hansen test make it possible to rule out the possibility of weak instruments and correlation of the instruments with the error term. The resulting parameter estimate is positive (0.28) but statistically insignificant. However, because use of only the cyclical component of the labor share makes it difficult to apply a policy interpretation to the results in Column (1), Columns (2) and (3) present results from the preferred specification, which uses the full value of the labor share as the variable of interest.

The results in Column (2) and (3) indicate the labor share is statistically and economically significantly positively related to capacity utilization at the state-level for the United States. The results in Column (2) suggest that a 10 percentage point increase in the labor share would correspond to a 5 percentage point increase in the rate of capacity utilization at the state-level, and the effect is even larger when time-fixed effects are included in Column (3). However, the stationarity of the capacity utilization series requires that these estimates be interpreted with caution: as long as capacity utilization is mean-reverting, these effects should be interpreted as temporary shocks. The specifications in Column (2) and (3) fail to reject the null hypothesis in the Sargan-Hansen test, supporting the conclusion that the instrumental variables are uncorrelated with the error term. However, the first-stage F-stat falls when time fixed-effects are included[18]. Thus, in Column (4) I add a lagged value of the labor share as an instrumental variable in addition to the minimum wage instruments. Column (4) thus presents a “mixed" specification, including both the minimum wage instruments and an *NF*-style instrument. The first-stage F-statistic is subsequently increased, and the conclusion of the Sargan-Hansen test that the instruments are uncorrelated with the error term is unaffected. However, the magnitude of the regression coefficient drops. Column (4) suggests

[18] A low first-stage F-stat is in some ways of less concern than a rejection of the null-hypothesis in the Sargan-Hansen Test. Bound et al. (1995) show that weak instruments tend to bias IV estimates in the direction of OLS, which means that Column (3) can be interpreted as close to the OLS estimate of the demand equation.

that a 10 percentage point increase in the labor share corresponds to roughly a 2 percentage point increase in capacity utilization at the state-level. Across all columns the general result is clear: demand is wage-led at the state-level.

4.4.2 Demand: Sensitivity Analysis

One obvious response to the above analysis is that the use of capacity utilization—a cyclical variable—still gives little insight into the long-run effects of changes in the labor share on the economy, despite the fact that I make use of the full value of the labor share, rather than its cyclical component. To address this concern, I re-estimate equation (4.3), taking the growth rate of state-level real GDP between time t and $t+1$ as the dependent variable, using the same set of instruments for the labor share as before. Columns (1)-(3) present the results of this specification.

The results from Columns (1), (2), and (3) indicate that changes in the state-level labor share in time t are positively related to the growth rate of real GDP in the state between t and $t + 1$. In each case the overidentification test fails to reject the null hypothesis that the instruments are uncorrelated with the errors at the 5% level, offering support for the exogeneity of the instruments. The first-stage F-stat exceeds 10 in Columns (1) and (3), and is slightly below 10 in Column (2), indicating that—in at least two out of the three specifications—the instruments are unlikely to suffer from weak-instrument bias. Given the wage-led results on capacity utilization, the results from the growth regression suggest that both demand and growth are wage-led at the state-level in the United States.

A second concern regarding the above regressions is the absence of a control for the political climate in various states, a variable which obviously influences the institutional environment governing the dynamics of the wage-share (and the likelihood of increases in the minimum wage in particular). While state-fixed effects will pick up any time-invariant characteristics of this sort, the political environment in any particular state changes over time. To control for this I make use of the state-level partisan balance series from Klarner (2013). In particular, I include a control for the share of Democratic state legislators (state senators) in a particular year. This series extends only

Table 4.4: Estimation Results - Demand Nullcline Sensitivity Analysis

	(1)	(2)	(3)	(4)	(5)
	g_{it}	g_{it}	g_{it}	u_{it}	g_{it}
ω_{it}	0.185***	0.547**	0.206***	0.872**	0.531***
	(0.0673)	(0.241)	(0.0406)	(0.409)	(0.198)
$g_{i,t-1}$	0.304***	0.352***	0.299***		0.354***
	(0.0464)	(0.0572)	(0.0447)		(0.0540)
$g_{i,t-2}$	0.0336*	0.0895**	0.0808***		0.0845**
	(0.0195)	(0.0360)	(0.0264)		(0.0355)
$u_{i,t-1}$				0.939***	
				(0.0687)	
$u_{i,t-2}$				-0.0658	
				(0.0503)	
$ln(Taxes_{it})$	0.00821*	0.0100***	0.0114**	0.00654	0.0103***
	(0.00446)	(0.00329)	(0.00508)	(0.00677)	(0.00321)
Democrat Share$_{it}$				0.0140	-0.000945
				(0.0203)	(0.0110)
N	1938	1938	1938	1764	1764
R^2	0.131	0.374	0.407	0.710	0.387
Time FE	N	Y	Y	Y	Y
State FE	Y	Y	Y	Y	Y
First Stage F-Stat	32.83	7.41	1409	6.10	8.15
Sargan-Hansen Test (p-value)	0.07	0.29	0.06	0.86	0.46

Notes: Standard errors in parenthesis, clustered at the state-level. * 0.10 ** 0.05 *** 0.01. Columns (1)-(3) present results for regressions where the state-level annual growth rate of RGDP between t and $t+1$ is the dependent variable. Columns (1) and (2) use a vector of minimum wage instruments. Column (3) adds the lagged value of the labor share as an additional instrument. Columns (4) and (5) add an additional control for the political composition of state-legislatures to the second-stage regression, in specifications using both utilization and growth, respectively.

to 2011, so 2012 to 2014 are dropped in the Columns including this variable. Columns (4) and (5) present results of these regressions. Including the Democrat share appears to have little impact on the regression results, taking on values that are both statistically and economically insignificant, and leaving the relationship between the labor-share, utilization, and growth unchanged.

Finally, Barbosa-Filho (2016) estimates a demand and distribution system using the employment rate as the demand variable, rather than capacity utilization. Unlike capacity utilization, which—as a statistical artifact obtained from either the Hodrick-Prescott or Hamilton (2017) filter—is stationary by definition, the long-run value of the employment rate can move over time. The employment rate therefore allows one to capture more than the just cyclical fluctuations expressed by movements in capacity utilization. Using the employment rate as a measure of demand may provide further insight—on top of that obtained via the use of the real GDP growth rate above—on the effect of changes in the labor share on demand in the long-run. Table (4.5) presents the results from estimating the demand equation with the employment rate as the dependent variable, where the employment rate, e_{it}, is given as one minus the state-level unemployment rate, obtained from the Bureau of Labor Statistics (*BLS*). The employment rate data extend back to only 1976, so 1974 and 1975 are dropped from the sample.

Table 4.5: Estimation Results - Employment Rate Specification

	(1) IV1	(2) IV2	(3) IV3	(4) IV4
$e_{i,t-1}$	1.163***	1.109***	1.126***	1.112***
	(0.0287)	(0.0347)	(0.0309)	(0.0319)
$e_{i,t-2}$	-0.463***	-0.309***	-0.340***	-0.323***
	(0.0199)	(0.0338)	(0.0275)	(0.0277)
ψ_{it}	0.0743**	0.124*	-0.0000912	0.0369***
	(0.0312)	(0.0733)	(0.00643)	(0.00952)
$ln(Taxes_{it})$	0.00911***	0.00245*	0.00112	0.00233**
	(0.00283)	(0.00136)	(0.00101)	(0.00110)
N	1887	1887	1887	1836
R^2	0.736	0.882	0.912	0.910
Time FE	N	Y	Y	Y
State FE	Y	Y	Y	Y
First Stage F-Stat	37.05	5.65	432.9	41.39
Sargan-Hansen Test (p-value)	0.27	0.40	0.02	0.12

Notes: Standard errors in parenthesis, clustered at the state-level. * 0.10 ** 0.05 *** 0.01. Column (1) omits time-fixed effects, and uses only the minimum wage instruments. Column (2) adds time-fixed effects, and uses only the minimum wage instruments. Column (3) adds a lagged value of the labor share as an instrument. Column (4) adds a deep lag (three periods) of the labor share as a instrument.

The results lend support to the wage-led findings obtained when capacity utilization is used as the demand variable. Column (1) omits time-fixed effects, and uses only the minimum wage instruments. Column (2) adds time-fixed effects, and uses only the minimum wage instruments. Column (3) adds a lagged value of the labor share as an instrument. Column (4) adds a deep lag (three periods) of the labor share as a instrument. With the exception of Column (3), which adds a one-period lagged value of the labor share as an additional instrument, the results indicate a statistically significant, positive impact of the labor share on the employment rate. The vector of instruments used in Column (3) result in a rejection of the null-hypothesis in the overidentification test, suggesting that the instruments are correlated with the regression error term. Thus, in Column (4) I test a deeper lag of the labor share as an additional instrument. In this specification I cannot reject the null-hypothesis that the instruments are orthogonal to the error term, and the resulting regression coefficient is positive and statistically significant.

The relative magnitude of the estimated effect of the labor share on the employment rate is comparable to the size of the effect found when capacity utilization is used as the demand variable. Comparing the estimates in Column (4) of Tables (4.3) and (4.5), which—for both the employment rate and capacity utilization—provide a lower bound to the positive range of estimated effects, the regression coefficient on the labor share in Table (4.3) is approximately twice the standard deviation of capacity utilization (0.11), and the regression coefficient on the labor share in Table (4.5) is approximately twice the standard deviation of the employment rate (0.02), such that the impact of a distributive shock—relative to the variance of either demand variable—is similar. However, in the case of the employment rate one need not be tied to the interpretation that the distribution shock is temporary, because—unlike in the case of capacity utilization, which is stationary by construction—the long-run value of the employment rate can move over time.

4.4.3 Distribution

In the absence of a viable identification strategy for obtaining parametric estimates of the distributive curve, I opt to estimate the distributive nullcline non-parameterically. I adopt two non-parametric techniques: Fan (1992)'s local regression smoother and Cleveland (1979)'s locally weighted regression. Technical details on the non-parametric estimators are included in the appendix. Figures (4.5), (4.6), (4.7), and (4.8) present the results of the estimation. Figure (4.5) presents results from the pooled sample of all fifty states and D.C., Figures (4.6)-(4.8) present the results obtained from estimating the distributive curve non-parametrically state-by-state. The confidence intervals presented are those calculated for the Fan (1992) estimator.

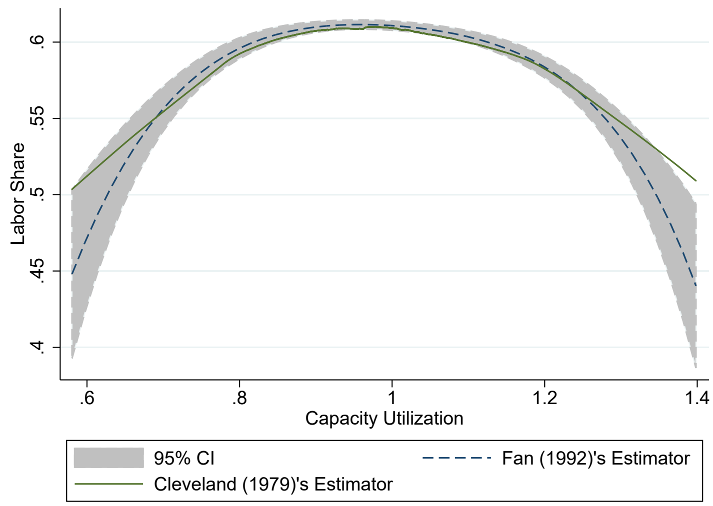

Figure 4.5: Non-Parametric Estimates of the Distributive Curve: Pooled Sample

The results suggest the presence of substantial non-linearities in the distributive curve—although the extent and shape of this non-linearity varies across states. States such as Connecticut, Delaware, and Oregon demonstrate clear non-linearities, whereas a state such as Wyoming has a distributive curve that is essentially flat. Further, the results suggest that whether or not the distributive curve is characterized by profit-squeeze or wage-squeeze dynamics varies across states. Washington D.C. appears to be characterized by wage-squeeze dynamics, whereas a state like Nevada demonstrates profit-squeeze dynamics at low levels of capacity utilization and wage-squeeze dynamics at high levels of utilization. This latter pattern is the same that appears in the pooled sample of all fifty states and Washington D.C.. The possibility of a distributive curve that demonstrates profit-squeeze dynamics at low-levels of capacity utilization and wage-squeeze dynamics at high-levels of capacity utilization has important implications for the feasibility of increasing demand through redistribution toward wages, especially if one adopts the alternative definition of wage- or profit-led that Nikiforos and Foley (2012) argue is implied by the presence of a non-linear distributive curve, namely that an economy is profit-led if an exogenous technological or distributive change against

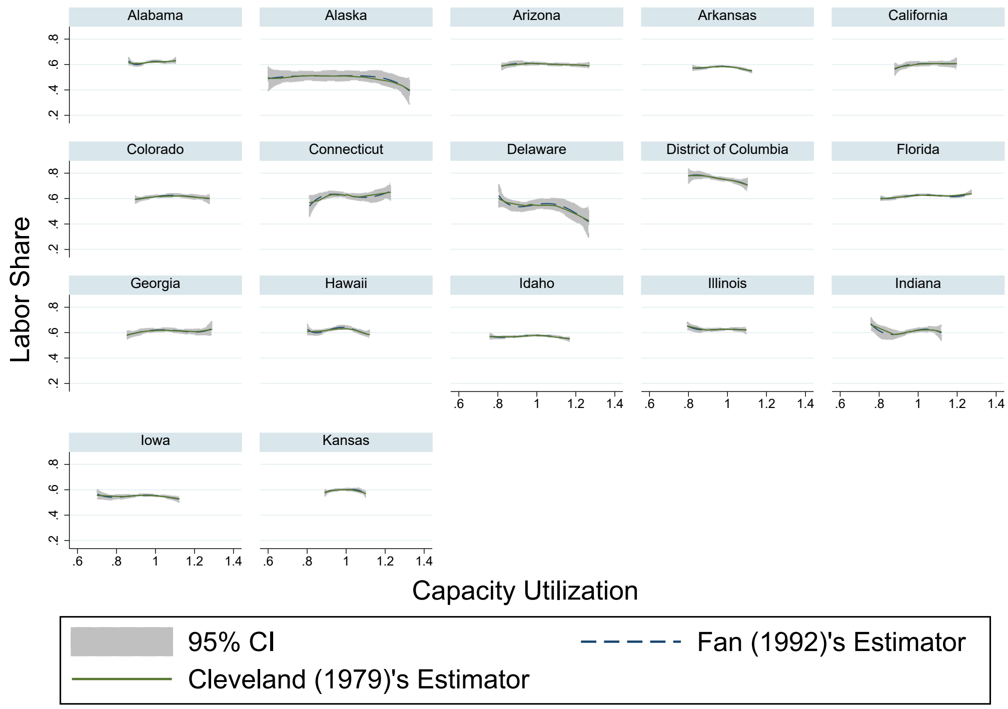

Figure 4.6: Non-Parametric Estimates of the Distributive Curve — State-by-State (1)

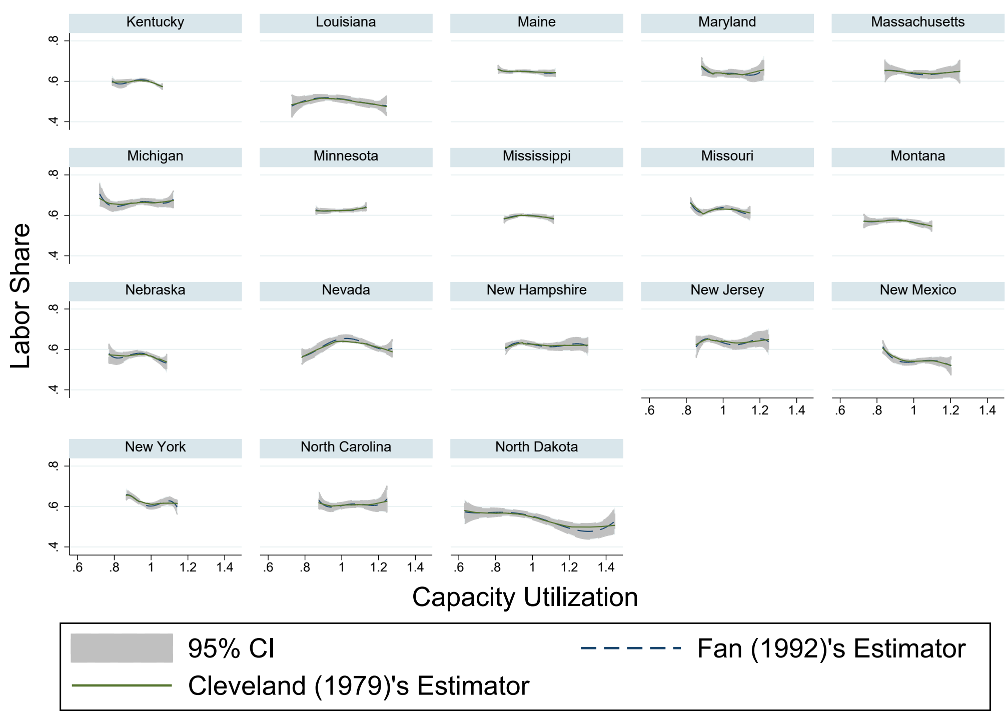

Figure 4.7: Non-Parametric Estimates of the Distributive Curve — State-by-State (2)

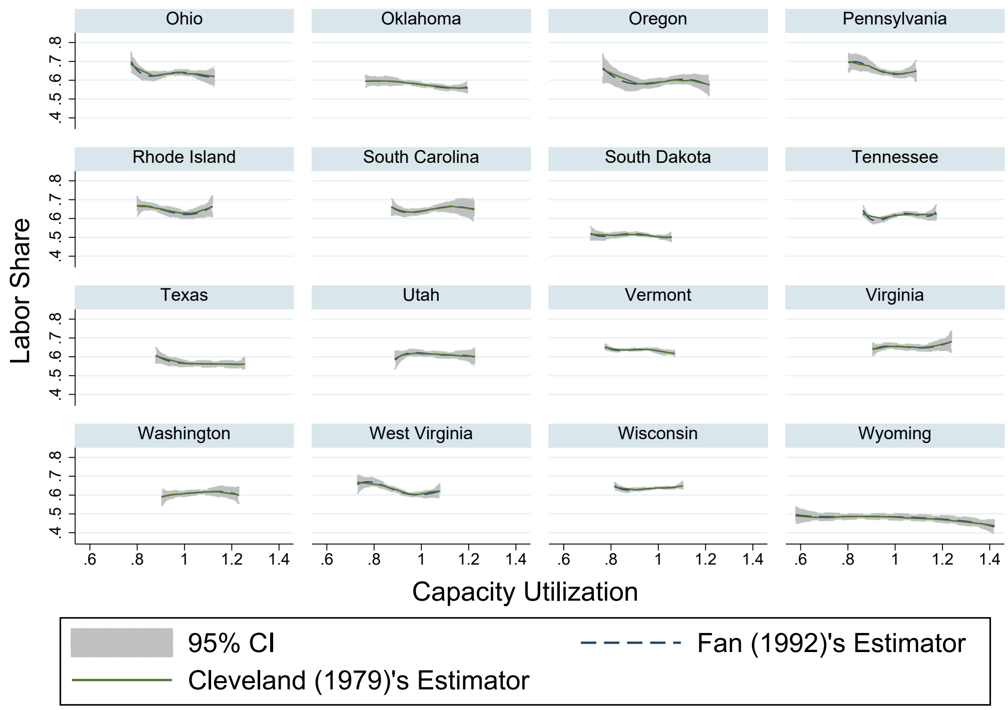

Figure 4.8: Non-Parametric Estimates of the Distributive Curve — State-by-State (3)

the labor share increases capacity utilization (e.g. if a shift of the distributive curve corresponding to a reduction in the labor share increases capacity utilization), and wage-led if the opposite is true. In this case, wage-led demand may not be sufficient to ensure that a redistribution in favor of wages increases capacity utilization. I move now to a discussion of some of these issues.

4.4.4 Discussion

The first important implication of the exercise undertaken above is the possibility of a divergence between the national and regional demand regimes within a country. The results suggest that at the state-level demand in the United States is primarily wage-led, despite profit-led results for demand in studies using national data such as Barbosa-Filho and Taylor (2006) and Nikiforos and Foley (2012). This result can be explained in terms of the paradox of costs for firms (or in this case, regional governments).[19] Raising wages may increase demand and profits when there is idle

[19]Credit goes to a helpful, anonymous referee for this insight.

capacity, provided regional governments act to do this in a coordinated fashion. However, no state would do this alone, for fear of losing economic activity to other, lower-tax states. A profit-led system thus arises in the aggregate, due to the coordination failure. Even though individual states are wage-led on average, the inability to coordinate redistribution across states creates a fallacy of composition such that—when the behavior is aggregated to the national level—leads to the appearance of profit-led features in the data. This result is consistent with the insight suggested by both Arnim, Carvalho, and Tavani (2014) and Razmi (2018): economic aggregates (the world, a country) are not isomorphic to their component parts (a particular country, a region). Economic aggregates may have emergent properties that do not appear at lower levels of aggregation, and the process of aggregation itself may produce trends in aggregate data which disappear when the data are dis-aggregated.

The second result in need of greater attention is the non-linearity of the distributive curve, and the divergence of the shape of this non-linearity from previous findings. In a system with linear demand and distributive curves, the wage-led property of state-level economic activity implies that regional policy makers should seek to implement redistribution toward wages, regardless of national macroeconomic policy. However, estimates of the distributive curve indicate that the demand and distribution system at the state-level is *not* linear in distribution, and thus—even with wage-led demand, it will be seen that increasing demand through redistribution may be difficult. Figure (4.9) plots the demand and distribution system implied by the estimation results for the demand curve and the pooled estimate of the distributive curve. Those results suggest a wage-led demand system and a distributive curve that is profit-squeeze at low levels of capacity utilization and wage-squeeze at high levels. The system has two possible equilibria —a high- wage-share high-capacity equilibrium (**A**), and a low- wage-share low-capacity equilibrium (**B**). For a state residing at **B**, a redistribution or technological change in favor of the wage-share—represented by an outward shift of the distributive nullcline from $\dot{\psi} = 0$ to $(\dot{\psi} = 0)'$— results in a decline in both equilibrium capacity utilization and the equilibrium wage-share (a movement to **B′**). Analysis of the stability of these equilibria is presented in Appendix (B.1).

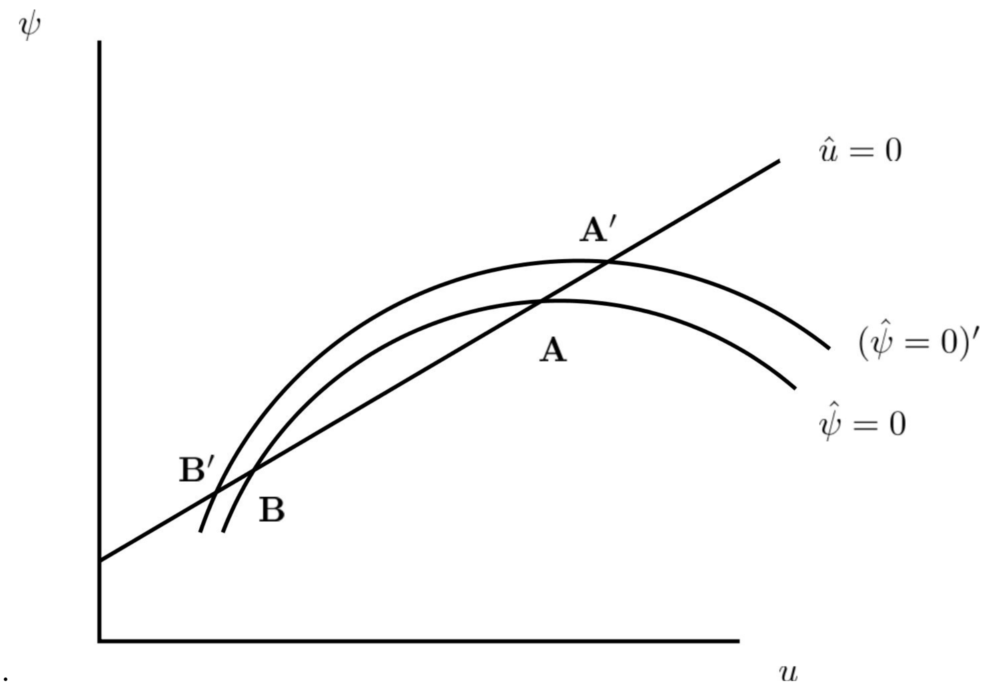

Figure 4.9: Demand and Distribution System with Wage-Led Demand and Non-Linear Distributive Curve

Notes: The stability conditions for each equilibria are checked in Appendix (B.1).

There are several possible reasons for the particular shape of the distributive curve and the two equilibria that result. First, the wage-squeeze at high capacity may be partially explained by the operation of inflation-targeting monetary policy in the United States. The monetary authority raises interest rates when capacity utilization is high, putting downward pressure on wages. A second explanation can be found by looking at the definition of the profit rate. Suppose that profits are taxed proportionately at rate τ. The after-tax rate of profit is given by:

$$r = (1 - \tau)(1 - \psi)u \tag{4.4}$$

In the long-run, it is not unreasonable to assume that within-country capital mobility ensures the profit-rate is fixed at the state-level at a rate equal to the national average, $\overline{r}$. In this case, re-arranging the above expression for the labor share, we arrive at:

$$\psi = \frac{u(1 - \tau) - \overline{r}}{u(1 - \tau)} \tag{4.5}$$

which, for a given tax rate, is a concave function of capacity utilization, in the shape of the distributive curve in Figure (4.9).

This explanation for the shape of the distributive curve is appealing, because the behavior that gives rise to equalization of the after-tax rate of profit across states can also explain why an attempted redistribution toward wages may not increase the level of economic activity, even when the demand regime is wage-led (that is, why a state might end up in equilibrium **B**). A large literature suggests that “tax competition" between regions resulting from the desire to attract new firms and residents may push taxation on capital toward inefficiently low levels, while homogenizing tax rates across regions (Leiser, 2017; Oates and Schwab, 1988; Wilson, 1999; Wilson and Wildasin, 2004). At the same time, tax competition has been found to reduce public revenue streams below the level necessary to finance the socially optimal level of public goods provision, and to have little impact on the rate of capital formation at the state-level (Break, 1967; Chirinko and Wilson, 2008). If a state were to raise taxes above the national average, capital would eventually migrate out of the state, seeking higher after-tax returns elsewhere, leaving the state in question worse off. Olsen (2010) gives reason to think that this behavior is to be expected from capitalist firms, because the objective of a capitalist firm's location choice is to secure a distribution of income that favors capital at the expense of labor.

The resulting configuration looks like a prisoners' dilemma. Wanting to avoid out-migration of capital, states have an incentive to reduce tax rates when other states do, despite the fact that all could be better off by redistributing towards wages. In addition to explaining the shape of the distributive curve at the state-level, the equalization of profit-rates across states is therefore consistent with the coordination failure that gives rise to the appearance of profit-led behavior in the aggregate. The shape of the distributive curve that results from the equalization of the after-tax rate of profit across states is also the shape one would expect if tax competition between states were present. This occurs because tax competition is itself a response to the behavior that leads to the equalization of the rate of profit (namely, the outflow of capital from states with higher-than-average tax rates). Figure (4.10) provides suggestive evidence of the operation of this type of tax competition in the United States. Since 1990, state-level tax policy has become more regressive, suggesting a “race to the bottom." Both the average effective state corporate income tax rate and the average effective state personal income tax rate have declined, while the average sales tax rate has increased.

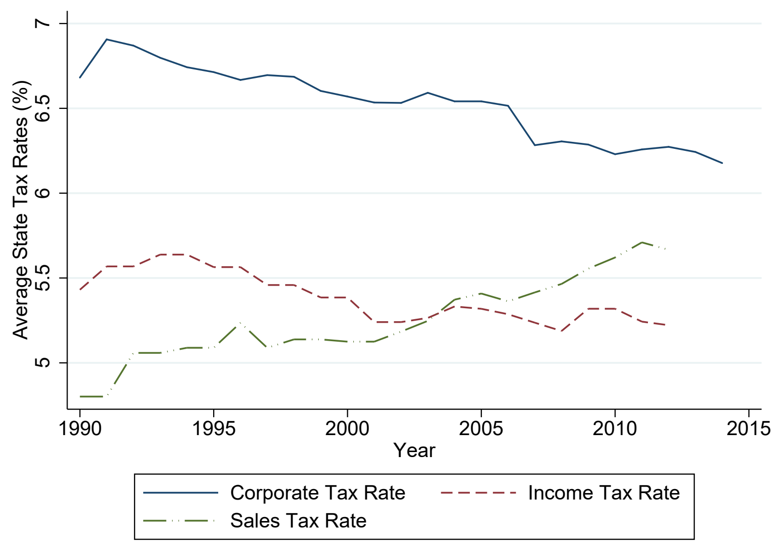

Figure 4.10: State Tax Rates

Sources: Chirinko and Wilson (2008), Tax Foundation, and The Book of the States.

Redistributive policy induces behavioral responses on behalf of both firms and households, such as the migration of capital between states. As a result, an attempted redistribution toward wages may not have the desired effect (indeed it may have actually reduce the equilibrium wage share!). Unless all states (or at least, a group of states in a particular region) decide to coordinate the tax schedules used to finance redistributive policy, firms will always have an incentive to shift capital and resources across state borders in response to a unilateral tax increase—behavior which may result in a reduction in utilized capacity and the revenues available for public expenditure. The evidence on tax evasion and transfer pricing suggests that the elasticity of reported income with respect to the corporate tax rate is quite large, and that the size of “missing" tax revenues globally is substantial (Bartelsman and Beetsma, 2003; Zucman, 2013). The ability of an individual state to boost demand through a unilateral increase in its expenditures on redistributive policy financed by a tax on either top income earners or corporate profits is constrained by the tax rates set by all other states. At a regional level wage-led policy therefore appears to be characterized by a coordination failure. If country-level policy is constrained by ‘race to the bottom’ dynamics, then wage-led policy appears as a set of nested coordination problems.

States may be stuck in an inferior equilibrium in **B**[20], where an attempted redistribution toward wages decreases both demand and the wage-share due to the combined effects of tax competition and the behavioral responses of firms and households, despite demand being wage-led. There exists a superior equilibrium in **A**, where a redistribution toward wages increases both the wage-share and capacity utilization, but this equilibrium requires coordinated redistribution among states to avoid the otherwise perverse effects that occur in equilibrium **B**. This result highlights an important complication for the feasibility of wage-led policy that is often ignored in discussions about wage- versus profit-led growth: the labor share is not a policy variable. Even when demand is wage-led, it is not clear that policies which re-distribute toward wages will successfully increase demand if those policies induce simultaneous supply-side behavioral responses that push demand in the opposite direction. Tax competition among states within countries, competition for lower unit-labor costs between countries, and the profit-led tendency of trade-deficit countries all create difficulties for successfully increasing demand through a redistribution toward wages.

4.5 Conclusion

In this chapter I have cataloged another entry in the wage- versus profit-led growth debate. Using regional data for a panel of U.S. states, I estimate a demand and distribution system using an IV 2SLS approach, exploiting variation in state-level minimum wage policy. I find that when variation in the full value of the labor share is properly identified, the United States appears to be characterized by wage-led demand at the state-level. This implies that even if the United States is profit-led at the country-level, state policy makers should strive to enact redistribution in favor

[20]Appendix (B.1) analyzes the stability of this equilibrium.

of wages, regardless of macroeconomic policy. However, the labor share is not a policy variable. Because demand and distribution are simultaneously determined, non-linearity in the distributive curve may make it difficult to increase demand through redistribution. If state governments respond to the firm's location choice problem by competing over tax rates in an effort to attract capital, an increase in regional taxation in order to finance redistribution may actually reduce both capacity utilization and the labor share. Wage-led policy at the state-level thus exhibits the characteristics of a coordination failure. These dynamics are further complicated by the fact that many open economies appear individually profit-led—despite the wage-led nature of global demand. Even if regional governments within countries can successfully cooperate to raise taxes and redistribute toward wages, the net effect may be a reduction of demand in the aggregate—unless a similar cooperative agreement exists between countries to ensure that 'race to the bottom' dynamics are avoided. Wage-led policy thus presents itself as a set of nested coordination problems.

This last point represents a significant obstacle to enacting a successful redistribution (that is, one which achieves the goal of both increasing labor's share of income and the level of economic activity in a region), and points to new directions for future research. In a world where the labor share is not free to vary as a benevolent social planner would dictate, the sign of the demand schedule matters less than the general equilibrium effect of a policy change. In a model with government infrastructure investment and research and development spending, Tavani and Zamparelli (2017) show there may be a range of policy options which increase both the rate of economic growth and the labor share, even when the economy is profit-led. In this chapter, I suggest that strategic interaction between governments may produce the opposite result—redistribution towards wages may decrease both demand and the labor share, even when the economy is wage-led. With respect to future theoretical work, this suggests that researchers interested in growth and distribution should attempt to incorporate elements of either behavioral responses to policy, strategic interaction among workers, capitalists, and governments, or both if they are to obtain realistic predictions about the demand and growth effects of redistribution. Future empirical research should direct itself toward investigating which policies actually bring about a successful increase in the labor share and level of economic activity, regardless of whether or not demand is wage- or profit-led.

Bibliography

Abadie, A., Diamond, A., and Hainmueller. 2010. “Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California’s Tobacco Control Program." *Journal of the American Statistical Association*, 105(490): 493-505.

Admanti, A. 2017. “A Skeptical View of Financialized Corporate Governance." *Journal of Economic Perspectives*, 31(3): 131-150.

Agarwal, S., and Hauswald, R. 2010. “Distance and Private Information in Lending." *Review of Financial Studies*, 23(7): 2757-2788.

Agarwal, S., Chomsisengphet, S., Liu, C., Song, C., and Souleles, N. 2018. “Benefits of Relationship Banking: Evidence from Consumer Credit Markets." *Journal of Monetary Economics*, 96: 16-32.

Akerlof, G. 1970. “The Market for Lemons: Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics*, 84(3): 488-500.

Akhigbe, A., and Whyte, A. 2003. “Changes in Market Assessment of Bank Risk Following the Riegle-Neal Act of 1994. *Journal of Banking and Finance*, 27, 87-102.

Alessandrini, P., Croci, M., and Zazzaro, A. 2005. “The Geography of Banking Power: The Role of Functional Distance." *BNL Quarterly Review*, 58(235): 129-167.

Alon, T., Berger, D., Dent, R., and Pugsley, B. 2018. “Older and Slower: The Startup Deficit’s Lasting Effects on Aggregate Productivity Growth." *Journal of Monetary Economics*, 9: 68-85.

Angrist, J., and Krueger, A. 1999. “Empirical Strategies in Labor Economics." in *Handbook of Labor Economics*, Vol 3., eds. A. Ashenfelter and D. Card, Amsterdam: Elsevier.

Anenberg, E., Chang, A., Grundl, S., Moore, K., and Windle, R. 2018. “The Branch Puzzle: Why are There Still Bank Branches?" *FED Notes No. 2018-08-20*.

Arcand, J., Berkes, E., and Panizza, U. 2015. “Too Much Finance?" *Journal of Economic Growth*, 20(2): 105-148.

Archarya, V., Hasan, I., and Saunders, A. 2006. “Should Banks be Diversified? Evidence from Individual Bank Loan Loan Portfolios." *Journal of Business*, 79: 1355-1411.

Arnim, R., Tavani, D., and Carvalho, L. 2014. Redistribution in a Neo-Kaleckian Two-Country Model, *Metroeconomica*, vol. 65, no. 3, 430-459.

Bardhan, P. & Udry, C. (1999). *Development Microeconomics*. Oxford: Oxford University Press.

Barbosa-Filho, N. 2016. “Elasticity of Substitution and Social Conflict: A Structuralist Note on Piketty’s *Capital in the Twenty-First Century*." *Cambridge Journal of Economics*, vol 40, no. 4, 1167-1183.

Barbosa-Filho, N., and Taylor, L. 2006. Distributive and Demand Cycles in the US Economy—A Structuralist Goodwin Model, *Metroeconomica*, vol. 57, no. 3, 389-411.

Bartelsman, E., and Beetsma, R. 2003. Why Pay More? Corporate Tax Avoidance Through Transfer Pricing in OECD Countries. *Journal of Public Economics*, vol. 87, no.9-10, 2225-2252.

Beccalli, E., and Frantz, P. 2009. “M&A Operations and Performance in Banking." *Journal of Financial Services Research*, 36: 203-226.

Berger, A., and DeYoung. 2006. “Technological Progress and the Geographic Expansion of the Banking Industry." *Journal of Money, Credit, and Banking*, 38(6): 1483-1513.

Berger, A., and Udell, G. 2002. “Small Business Credit Availability and Relationship Lending: The Importance of Bank Organizational Structure." *Economic Journal*, 112: F32-F53.

Bertrand, M., Duflo, E., and Mullainathan, S. 2002. “How Much Should We Trust Difference-in-Differences Estimates?" *NBER Working Paper* No. 8841.

Bertrand,M., Duflo, E., and Mullainathan, S. 2004. “How Much Should We Trust Difference-in-Differences Estimates?" *Quarterly Journal of Economics*, 119(1): 249-275.

Bezemer, D., Grydaki, M., and Zhang, L. 2016. “More Mortgages, Lower Growth?" *Economic Inquiry*, 54(1): 652-674.

Bhaduri, A. and Marglin, S. 1990. Unemployment and the Real Wage: The Economic Basis for Contesting Political Ideologies, *Cambridge Journal of Economics*, vol. 14, no. 4, 375-393.

Blecker, R. 2016. Wage-led versus Profit-Led Demand Regimes: the Long and the Short of it, *Review of Keynesian Economics*, vol. 4, no. 4, 373-390.

Bound, J., Jaeger, D., and Baker, R. 1995. Problems with Instrumental Variables Estimation When the Correlation Between the Instruments and the Endogenous Explanatory Variable is Weak, *Journal of the American Statistical Association*, vol. 90, no. 430, 443-450.

Break, G. 1967. *Intergovernmental Fiscal Relations in the United States*, Washington, D.C.: Brookings Institution.

Boyd, J., and Prescott, E. 1986. “Financial Intermediary-Coalitions." *Journal of Economic Theory*, 38: 211-232.

Brown, C. 2007. “Financial Engineering, Consumer Credit, and the Stability of Effective Demand." *Journal of Post Keynesian Economics*, 29(3): 427-450.

Bunten, D., Weiler, S., Thompson, E., and Zahran, S. 2015. “Entrepreneurship, Information, and Growth." *Journal of Regional Science*, 55(4): 560-584.

Carlson, M., and Mitchener, K. 2009. “Branch Banking as a Device for Discipline: Competition and Bank Survivorship during the Great Depression." *Journal of Political Economy*, 117(2): 165-210.

Carvalho, L, and Rezai, A. 2016. “Personal Income Inequality and Aggregate Demand." *Cambridge Journal of Economics*, vol. 40, no. 2, 491-505.

Chirinko, R., and Wilson, D. 2008. “State Investment Tax Incentives: A Zero-Sum Game? *Journal of Public Economics*, vol. 92, no. 12, 2362-2384.

Cleveland, W. 1979. Robust Locally Weighted Regression and Smoothing Scatterplots, *Journal of the American Statistical Association*, vol. 74, 829-836.

Chong, B. 1991. “The Effect of Interstate Banking on Commercial Banks’ Risk and Profitability." *Review of Economics and Statistics*, 73(1): 78-84.

Coase, R. 1937. “The Nature of the Firm." *Economica*, 4(16): 386-405.

Conroy, T., Low, S., and Weiler, S. 2017. “Fueling Job Engines: Impacts of Small Business Loans on Establishment Births in Metropolitan and Nonmetro Counties." *Contemporary Economic Policy*, 35(3): 578-595.

Correia, S. 2015. “Singletons, Cluster-Robust Standard Errors and Fixed Effects: A Bad Mix." *Working Paper*. http://scorreia.com/research/singletons.pdf.

Deaton, A. 1991. “Saving and Liquidity Constraints." *Econometrica*, 59(5): 1221-1248.

Deaton, A. 1992. *Understanding Consumption*. Oxford: Clarendon Press.

DeLong, G., and DeYoung, R. 2007. “Learning by Observing: Information Spillovers in the Execution and Valuation of Commercial Bank M&As." *Journal of Finance*, 62(1): 181-216.

Demsetz, R., and Strahan , P. 1997. “Diversification, Size, and Risk at Bank Holding Companies. *Journal of Money, Credit, and Banking*, 29: 300-313.

Deng, S., and Elyasiani, E. 2008. “Geographic Diversification, Bank Holding Company Value, and Risk. *Journal of Money, Credit, and Banking*, 40: 1217-1238.

Diamond, D.W. 1984. “Financial Intermediation and Delegated Monitoring." *Review of Economic Studies*, 51: 393-414.

DiNardo, J., and Pischke, J. 1997. “The Returns to Computer Use Revisited: Have Pencils Changed the Wage Structure To?" *Quarterly Journal of Economics*, 112, 291-303.

Dube, A., Lester, W., and Reich, M. 2010. Minimum Wage Effects Across State Borders: Estimates Using Contiguous Counties, *Review of Economics and Statistics*, vol. 92, no. 4, 945-964.

Dunn, L., and Mirzaie, I. 2016. “Consumer Debt Stress, Changes in Household Debt, and the Great Recession." *Economic Inquiry*, 54(1): 201-214.

Dymski, G. 2009. “The Global Financial Customer and the Spatiality of Exclusion after the ‘End of Geography’" *Cambridge Journal of Regions, Economy, and Society*, 2: 267-285.

Epstein, G. 2005. “Introduction." in, Epstein, G. ed. *Financialization and the World Economy*, Northampton: Edward Elgar.

Fama, E. 1969. “Efficient Capital Markets: A Review of Theory and Empirical Work." *Journal of Finance*, 25(2): 383-417.

Fan, J. 1992. Design-Adaptive Nonparametric Regression, *Journal of the American Statistical Association*, vol. 87, 998-1004.

Fazzari, S., Hubbard, G., and Peterson, B. 1988. “Financing Constraints and Corporate Investment." *Brookings Papers on Economic Activity*, 19(1): 141-206.

Feng, A. 2018. “Bank Competition, Risk Taking, and their Consequences: Evidence from the U.S. Mortgage and Labor Markets." *IMF Working Paper*, WP/18/157.

Fort, T., Haltiwanger, J., Jarmin, R. and Miranda, J. 2013. “How Firms Respond to Business Cycles: The Role of Firm Age and Firm Size." *IMF Economic Review*, 61(3): 520-559.

Frank, M. 2009. “Inequality and Growth in the United States: Evidence From a New State-level Panel of Income Inequality Measures." *Economic Inquiry*, vol. 47, no. 1, 55-68.

Friedman, M. 1957. *A Theory of the Consumption Function*. Princeton: Princeton University Press.

Friedman, M. 1970. “The Social Responsibility of Business is to Increase its Profits." *New York Times Magazine*, September 13, 1970.

Friedman, T. 2005. *The World is Flat.* New York: Farrar, Strauss, and Giroux.

Fukuyama, F. 1992. *The End of History and the Last Man*. New York: Simon and Schuster.

Ganong, P., and Shoag, D. 2018. “Why Has Regional Income Convergence in the US Declined?" *Journal of Urban Economics*, 102: 76-90.

Giannone, E. 2018. “Skill-Biased Technical Change and Regional Convergence." *International Economics Section Working Paper*. Princeton University.

Goetz, M., Laeven, L., and Levine, R. 2016. “Does the Geographic Expansion of Banks Reduce Risk?" *Journal of Financial Economics*, 120: 346-362.

Goodwin, R. 1967. A Growth Cycle, in Feinstein, C.H., ed. *Socialism, Capitalism, and Economic Growth*, Cambridge: Cambridge University Press, 54-58.

Grembi, V., Nannicini, T., and Troiano, U. 2016. “Do Fiscal Rules Matter." *American Economic Journal: Applied Economics*, 8(3): 1-30.

Hamilton, J. 2017. “Why You Should Never Use the Hodrick-Prescott Filter," *Review of Economics and Statistics*. Advanced Access, published online August 11, 2017, doi: 10.1162/REST_a_00706.

Kane, E. 1996. “De Jure Interstate Banking: Why Only Now?" *Journal of Money, Credit, and Banking*, 28(2): 141-161.

Keynes, J.M. 1936 [1964]. *The General Theory of Employment Interest, and Money*. New York: Harcourt.

Klarner, C. 2013. “State Partisan Balance Data, 1937-2011," http://hdl.handle.net/1902.1/20403. IQSS Dataverse Network [Distributor] V1 [Version].

Krozner, R., and Strahan, P. 1999. “What Drive Deregulation? Economics and Politics of the Relaxation of Bank Branching Restrictions." *Quarterly Journal of Economics*, 114(4): 1437-1467.

Kwan, S. 2004.“Banking Consolidation." *Federal Reserve Bank of San Francisco*, Economic Letter 2004-15.

Lechevalier, S., Debanes, P., and Shin, W. 2017. “Financialization and Industrial Policies in Japan and Korea: Evolving Institutional Complementarities and Loss of State Capabilities." *Structural Change and Economic Dynamics*, Forthcoming.

Labaton, S. 2008. “Agency’s ’04 Rule Lets Banks Pile Up New Debt." *The New York Times*, October 2.

Lazonick, W., and O’Sullivan, M. 2000. “Maximising Shareholder Value: A New Ideology for Corporate Governance." *Economy and Society*, 29(1): 13-35.

Leiser, S. 2017. “The Diffusion of State Tax Incentives for Business." *Public Finance Review*, vol 45., no. 3, 334-363.

Leonnig, C. 2008. “How HUD Mortgage Policy Fed the Crisis." *Washington Post*, June 10.

Liberti, J., and Mian, A. 2009. “Estimating the Effect of Hierarchies on Information Use." *Review of Financial Studies*, 22(10): 4057-4090.

Livshits, I., Mac Gee, J., and Tertilt, M. 2016. “The Democratization of Credit and the Rise in Consumer Bankruptcies." *Review of Economic Studies*, 4(1): 1673-1710.

Marglin, S. 2017. Wages, Prices, and Employment in a Keynesian Long-Run," *Review of Keynesian Economics*, vol. 5, no. 3, 360-425.

Mason, J. (2017). “Income Distribution, Household Debt, and Aggregate Demand: A Critical Assessment." *Working Paper*. Draft prepared for the Center for Household Financial Stability & Private Debt Project âĂIJTipping PointsâĂİ Conference, June 22, 2017.

Mason, J., Jayadev, A., and Schroder, E. 2018. “The Political Economy of Financialization in the United States, Europe, and India." *Development and Change*, Early View.

Mian, A. and Sufi, A. 2009. “The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis." *Quarterly Journal of Economics*, 124(4): 1449-1496.

Mian, A. and Sufi, A. 2010. “Household Leverage and the Recession of 2007-09." *IMF Economic Review*, 58(1): 74-114.

Mian, A. and Sufi, A. 2011. “House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis." *American Economic Review*, 101(5): 2132-2156.

Mian, A. and Sufi, A. 2012. “What Explains High Unemployment? The Aggregate Demand Channel." *NBER Working Paper*, No. 17830.

Mian, A., and Sufi, A. 2017. “How Does Credit Supply Expansion Affect the Real Economy? The Productive Capacity and Household Demand Channels." *Working Paper.* https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2971086.

Mian, A., Sufi, A., and Rao, K. 2013. “Household Balance Sheets, Consumption, and the Economic Slump." *Quarterly Journal of Economics*, 128(4): 1687-1726.

Montecino, J., Epstein, G., and Levina, I. 2016. “Long-Term Trends in Intra-Financial Sector Lending in the United States." *Eastern Economic Journal*, 42(4): 611-629.

Moore, J., Petach, L., and Weiler, S. 2018. “Great Expectations? Nonemployer Establishments as Sparks for Urban Start-Ups." *REDI@CSU Working Paper.* Fort Collins, CO: Colorado State University.

Mora, R., and Reggio, I. (2012). “Treatment Effect Identification Using Alternative Parallel Assumptions." *UC3M Working Paper, we1233*.

Nikiforos, M., and Foley, D. 2012. Distribution and Capacity Utilization: Conceptual Issues and Empirical Evidence, *Metroeconomica*, vol. 63, no. 1, 200-229.

Oates, W., and Schwab, R. 1988. Economic Competition Among Jurisdictions: Efficiency Enhancing or Distortion Inducing, *Journal of Public Economics*, vol. 35, 333-354.

Olsen, E. 2010. “Class Conflict and Industrial Location." *Review of Radical Political Economics*, vol. 42, no. 3, 344-352.

Palley, T. 2007. “Financialization: What it is and Why it Matters." *Levy Institute*, Working Paper No. 525.

Palley, T. 2016. “Wage- vs. Profit-led Growth: the Role of the Distribution of Wages in Determining Regime Character." *Cambridge Journal of Economics*, vol. 31, no. 1, 49-61.

Petkov, I. “Small Business Lending and the Bank Branch Network." *Working Paper*, Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2754138

Pham, T., Talavara, O., and Tsapin, A. 2018. “Branch Networking Structure and Lending Behavior." *NBU Working Paper*. Kiev: National Bank of Ukraine.

Pintus, P., and Wen, Y. 2013. “Leveraged Borrowing and Boom-Bust Cycles." *Review of Economic Dynamics*, 16(4): 617-633.

Pressman, S., and Scott, R. 2017. “House Arrest: The Effects of Underwater and Low-Equity Mortgages on Small Business Failure and Mobility." *Review of Social Economy*, 75(2): 231-249.

Rada, C., and Kiefer, D. 2015. Profit Maximizing Goes Global: The Race to the Bottom, *Cambridge Journal of Economics*, vol. 39, no. 5, 1333-1350.

Razmi, A. 2018. “Is Planet Earth as a Whole Likely to be Wage-led?" *Review of Keynesian Economics*, *forthcoming*.

Reda, J., Schmidt, D., Glass, K. 2017. Study of 2015 Short- and Long-Term Incentive Design Criteria Among Top 200 S&P 500 Companies." Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2916206.

Rousseau, P., and Wachtel, P. 2011. “What is Happening to the Impact of Financial Deepening on Economic Growth?" *Economic Inquiry*, 49(1): 276-288.

Sawyer, M. 2013. “What is Financialization?" *International Journal of Political Economy*, 42(4): 5-18.

Skott, P., and Ryoo, S. 2008. “Macroeconomic Implications of Financialisation." *Cambridge Journal of Economics*, 32(6): 827-862.

Staiger, D., and Stock, J. 1997. “Instrumental Variables Regression with Weak Instruments." *Econometrica*, 65(3): 557-586.

Stiglitz, J. & Weiss, A. 1981.“Credit Rationing in Markets with Imperfect Information." *American Economic Review*, 71(3): 393-410.

Stockhammer, E. 2017. Wage-led versus Profit-led Demand: What have We Learned? A Kaleckian-Minskyan view, *Review of Keynesian Economics*, vol. 5, no. 1, 25-42.

Tavani, D., Flaschel, P., and Taylor, L. 2011. Estimated Non-Linearities and Multiple Equilibria in a Model of Distributive-Demand Cycles, *International Review of Applied Economics*, vol. 25, no. 5, 519-538.

Tavani, D., and Zamparelli, L. 2017. Growth, Income Distribution, and the ‘Entrepreneurial State’, New School for Social Research Working Paper. http://www.economicpolicyresearch.org/econ/2017/NSSR_WP_112017.pdf.

Torluccio, G., Cotungo, M., and Strizzi, S. 2011. “How is Bank Performance Affected by Functional Distance?" *European Journal of Economics, Finance and Administrative Services*, 39: 79-93.

Uchida, H., and Udell, G. 2012. “Loan Officers and Relationship Lending to SMEs." *Journal of Financial Intermediation*, 21(1): 97-122.

Vaghul, K., and Zipperer, B. 2016. Historical State and Sub-State Minimum Wage Data, *Washington Center for Equitable Growth Working Paper.*

Weiler, S. 2000. "Pioneers and Settlers in LoDo Denver: Private Risk and Public Benefits in Urban Redevelopment." *Urban Studies*, 37(1): 167-179.

Weiler, S., Hoag, D., and Fan, C. 2006. "Prospecting for Economic Returns to Research: Adding Informational Value at the Market Fringe." *Journal of Regional Science*, 46(2): 289-312.

Williamson, O. 1967. "Hierarchical Control and the Optimal Firm Size." *Journal of Political Economy*, 75(2): 123-138.

Wilson, J. 1999. "Theories of Tax Competition." *National Tax Journal*, vol. 52., no. 2, 269-304.

Wilson, J., and Wildasin, D. 2004. "Capital Tax Competition: Bane or Boon." *Journal of Public Economics*, vol. 88, 1065-1091.

Wisman, J. 2013. "Wage Stagnation, Rising Inequality, and the Financial Crisis of 2008." *Cambridge Journal of Economics*, 37(4): 921-945.

Zucman, G. 2013. The Missing Wealth of Nations, Are Europe and the US net Debtors or Creditors?, *Quarterly Journal of Economics*, vol. 128, no. 3, 1321-1364.

Appendix A

Appendix to Chapter 2

A.1 County-level Sample Means

Table A.1: Sample Means, County-level Data

	Mean	Std. Dev.
$FIRE\_Emp_i^{1999}(\%)$	4.817	2.228
$S\_FIRE\_Emp_i^{1999}$ (%)	5.742	1.774
$\frac{Debt}{Income}_{it}$	1.398	0.534
ln(Population$_{it}$	10.90	1.109
ln(Income$_{it}$)	2.681	0.222
ln(Employment$_{it}$)	10.20	1.222
$\%$ Δ Population$_{it}$	0.00782	0.0207
$\%$ Δ Income$_{it}$	0.0118	0.0373
$\%$ Δ Employment$_{it}$	0.00981	0.0339
Industry Specialization$_{it}$	0.493	0.159
N	13,050	

Table (A.1) presents sample means for the key county-level variables. $FIRE\_Emp_i^{1999}$ gives the share of the finance, insurance, and real-estate sector in total county employment in 1999, $S\_FIRE\_Emp_i^{1999}$ gives the spatially-lagged value of the FIRE sector employment share, and $\frac{Debt}{Income}_{it}$ gives the county-level debt-to-income ratios from Mian, Sufi, and Rao (2013). Industry Specialization$_{it}$ gives the sum of the absolute deviations of county-level industry employment shares from national industry employment shares, where large values indicate that county employment is concentrated in relatively few industries, relative to the national average.

A.2 State-level Continuous Treatment

Table (A.2) presents results from estimating a continuous treatment version of the state level model, where the treatment variable is $\frac{FIRE^{1999}}{GDP}$: the share of the FIRE sector in state GDP in 1999. The treatment effect is economically and statistically significant in every specification except

Table A.2: Estimation Results: Continuous Treatment, State-level

	(1) Debt$_{it}$	(2) Debt$_{it}$	(3) Debt$_{it}$	(4) Debt$_{it}$	(5) Debt$_{it}$	(6) Debt$_{it}$	(7) Debt$_{it}$
$(\frac{FIRE^{1999}}{GDP} \times After)_{it}$	291.0**	291.0*	284.2*	146.7	280.0*	282.3**	268.7*
	(143.9)	(145.1)	(150.0)	(115.4)	(143.2)	(138.2)	(146.6)
Gini$_{it}$				-604.5***			
				(86.61)			
Theil$_{it}$					-25.30		
					(23.98)		
Top 1% Share$_{it}$						-394.5***	
						(104.4)	
Rel. Mean Deviation$_{it}$							-413.7***
							(130.8)
N	765	765	765	765	765	765	765
R^2	0.640	0.804	0.809	0.865	0.836	0.844	0.841
Controls	N	N	Partial	Y	Y	Y	Y
Time FE	N	Y	Y	Y	Y	Y	Y
State FE	Y	Y	Y	Y	Y	Y	Y

Notes: Standard errors in parenthesis, clustered at the state level. * 0.10 ** 0.05 *** 0.01.

Column (4), which uses the Gini co-efficient as a control for inequality. Columns (5)-(7) present results for alternative inequality measures, including the Theil Index, the Top 1% Income Share, and the Relative Mean Deviation. Each of these measures is highly correlated with the Gini Co-efficient (simple correlation co-efficients between 0.6 and 0.8), and the treatment effect is robust to each of these measures.

To make the continuous treatment specification at the state level fully analogous to the county-level estimates, Table (A.3) presents results from a regression where the treatment variable is defined based on the state-level FIRE sector employment share, rather than the output share. The results confirm that the relationship between FIRE sector employment and debt at the state level is similar to the one observed at the county level. A 1% increase in the share of the FIRE sector in total state employment in 1999 corresponds to an increase in per-capita indebtedness of approximately $970 during the housing market boom. The symmetry between on county- and state-level results on the one hand, and the results using output and employment shares on the other, is not surprising. FIRE sector employment and output shares are highly correlated, and both act as imperfect measure for the true institutional variable of interest: the level of local financialization.

Table A.3: Continuous Treatment, State-level Employment Share

	(1)
	Debt_{it}
$(FIRE\_Emp^{1999} \times After)_{it}$	792.2**
	(370.9)
Gini_{it}	-559.1***
	(90.27)
N	765
R^2	0.869
Controls	Y
Time FE	Y
State FE	Y

Notes: Standard errors in parenthesis, clustered at the state level. * 0.10 ** 0.05 *** 0.01.

A.3 County-level Discrete Treatment

Table (A.4) presents results from estimating the effect of financialization on county-level debt-to-income ratios, using a discrete treatment variable akin to that used with the state-level data. Column (1) presents the estimated treatment effect of the credit-supply shock on highly financialized counties, where a county is included in the treated group if the share of the FIRE sector in county employment in 1999 exceeded the sample mean in 1999. Column (2) estimates a dose responsiveness test, where $Low\_Treatment_i$ is an indicator variable taking a value of 1 if a county has a FIRE sector employment share between the median and seventy-fifth percentile in 1999, and 0 otherwise. $High\_Treatment_i$ is defined similarly for states above the seventy-fifth percentile. Column (3) includes indicators for whether county i is treated, as well as whether the spatial-lag of county i is treated. Finally, Column (4) estimates the impact of spatial spillovers on counties who have own-county FIRE sector employment shares below the treatment threshold.

In every specification the results of the discrete treatment mirror that of the continuous treatment. Highly financialized counties—according to both own-county FIRE sector employment and the spatial-lag of FIRE sector employment—borrowed more during the U.S. housing market boom, than counties which were not themselves highly financialized or exposed to highly financialized neighboring counties.

Table A.4: Estimation Results, County-level Data, Discrete Treatment

	(1) $\frac{Debt}{Income}_{it}$	(2) $\frac{Debt}{Income}_{it}$	(3) $\frac{Debt}{Income}_{it}$	(4) $\frac{Debt}{Income}_{it}$
$(HighFinance \times After)_{it}$	0.0606**		0.0479**	
	(0.0229)		(0.0191)	
$(Low\_Treatment \times After)_{it}$		0.0507*		
		(0.0295)		
$(High\_Treatment \times After)_{it}$		0.0911***		
		(0.0318)		
$(S\_HighFinance \times After)_{it}$			0.0848***	0.0655***
			(0.0260)	(0.0231)
N	13,050	13,050	13,050	9,408
R^2	0.702	0.704	0.708	0.719
Time FE	Y	Y	Y	Y
County FE	Y	Y	Y	Y
Controls	Y	Y	Y	Y

Notes: Standard errors in parenthesis, clustered at the state level. * 0.10 ** 0.05 *** 0.01. Column (1) presents the estimated treatment effect of the credit-supply shock on highly financialized counties, where a county is included in the treated group if the share of the FIRE sector in county employment in 1999 exceeded the sample mean in 1999. Column (2) estimates a dose responsiveness test, where $Low\_Treatment_i$ is an indicator variable taking a value of 1 if a county has a FIRE sector employment share between the median and seventy-fifth percentile in 1999, and 0 otherwise. $High\_Treatment_i$ is defined similarly for states above the seventy-fifth percentile. Column (3) includes indicators for whether county i is treated, as well as whether the spatial-lag of county i is treated. Finally, Column (3) estimates the impact of spatial spillovers on counties that have own-county FIRE sector employment shares below the treatment threshold.

Appendix B

Appendix to Chapter 4

B.1 Stability Analysis

Equations (4.1) and (4.2) define a dynamic system in capacity utilization and the labor share. By definition, the labor share can be written: $\psi = \frac{w}{x}$ where w is the real wage and x is labor productivity. Assuming that capacity utilization varies to equilibribrate saving and investment, as is commonplace in Keynesian models of growth and distribution, the dynamical system can be re-written:

$$\hat{u} = I(\psi, u) - S(\psi, u) \tag{B.1}$$

$$\hat{\psi} = \hat{w}(\psi, u) - \hat{x}(\psi, u) \tag{B.2}$$

The slopes of the demand and distributive nullclines are thus given by:

$$\frac{du}{d\psi}|_{\hat{u}=0} = -\frac{I_\psi - S_\psi}{I_u - S_u} \tag{B.3}$$

$$\frac{du}{d\psi}|_{\hat{\psi}=0} = -\frac{\hat{w}_\psi - \hat{x}_\psi}{\hat{w}_u - \hat{x}_u} \tag{B.4}$$

Assuming the usual Keynesian stability condition holds (namely, that the sensitivity of savings with respect to utilization is greater than the sensitivity of investment), such that the denominator of (B.3) is negative, a wage-led demand nullcline implies a positive value for the numerator of (B.3). The signs of the numerator and denominator are more difficulty to ascertain in the case of the distributive nullcline than the demand nullcline. At equilibrium **B** on figure (4.9), the distributive curve has a positive slope, which is to say that (B.4) is positive. Thus, either the numerator or the denominator of (B.4) must be negative, with the other positive, given the observed slope. If the distributive curve is stable—such that $\frac{d\hat{\psi}}{d\psi} < 0$—this implies a negative value for the numerator of

(B.4). A positively sloped distributive curve then requires $\hat{w}_u - \hat{x}_u > 0$, so that real wages respond more strongly to increases in capacity utilization than labor productivity.

Given the signs of the relevant arguments, the Jacobian matrix of first derivatives can be written:

$$\mathbf{J} = \begin{pmatrix} I_u - S_u & I_\psi - S_\psi \\ \hat{w}_u - \hat{x}_u & \hat{w}_\psi - \hat{x}_\psi \end{pmatrix} \tag{B.5}$$

Stability requires $Tr|\mathbf{J}| = (I_u - S_u) + (\hat{w}_\psi - \hat{x}_\psi) < 0$ and $Det|\mathbf{J}| = (I_u - S_u)(\hat{w}_\psi - \hat{x}_\psi) - (I_\psi - S_\psi)(\hat{w}_u - \hat{x}_u) > 0$. The trace is confirmed to be negative, given that both entries are negative from the signs of (B.3) and (B.4). The signs of both entries in the determinant are positive, which means the value of the determinant may be either positive or negative, depending on the relative magnitudes of the slopes of the demand and distributive curves. The equilibrium at **B** will thus either be stable or exhibit saddle-path instability, conditional on whether the determinant is positive or negative. The equilibrium at **A** is obviously stable. For the same signs on the entries in the first row of the Jacobian, a negatively sloped distributive curve implies $\hat{w}_u - \hat{x}_u < 0$. The sign of the trace is thus unchanged, but the sign of the determinant becomes unambiguously positive.

B.2 Hamilton (2017) Filter

For a given outcome variable y—in this case, real state-level GDP—indexed by time, t, Hamilton (2017) shows the cyclical component of the series can be obtained from the residuals of a regression of y at time $t + h$ on the four most recent values of y as of date t:

$$y_{t+h} = \beta_0 + \beta_1 y_t + \beta_2 y_{t-1} + \beta_3 y_{t-2} + \beta_4 y_{t-3} + v_{t+h} \tag{B.6}$$

Using the parameter estimates from the above regression, the cyclical component is given by:

$$\hat{v}_{t+h} = y_{t+h} - Y_t^T \hat{\beta} \tag{B.7}$$

where Y_t^T is a vector containing the four most recent values of y as of time t, and $\hat{\beta}$ is the vector of estimated parameters. An estimate of potential output is obtained by subtracting the cyclical component from actual real GDP: $\hat{Y}^p_{t+h} = Y_{t+h} - \hat{v}_{t+h}$. Finally, capacity utilization in time t is calculated as the ratio of actual to potential output: $u_t = \frac{Y_t}{\hat{Y}^p_t}$. Using this procedure, I construct estimates of capacity utilization for each state, i, for every year in the sample.

B.3 Non-Parametric Techniques

For two variables y, x, such that y can be expressed as a function of x, $y = m(x) + \epsilon$, Fan (1992)'s local regression smoother estimates $m(x)$ as an n^{th} order local polynomial function in $(x - x_o)$, where x_0 is a realization of the random variable X, and x is in the neighborhood of x_0. $m(x_o)$ is equal to $\hat{\beta}_0$, the intercept of a weighted linear regression of y on x at x_0. Each i^{th} observation is given a weight:

$$w_i = \frac{1}{h}\kappa(\frac{x_i - x_0}{h}) \tag{B.8}$$

Where $\kappa(z)$ is the kernel function, in this case the standard Epanechnikov kernel:

$$\kappa(z) = (\frac{3}{4}(1 - \frac{1}{5}z^2)/\sqrt{5})\alpha(z) \tag{B.9}$$

Where $\alpha(z)$ is an indicator function taking a value of 1 if $|z| < \sqrt{5}$ and 0 otherwise. The bandwidth, h, is chosen as that which minimizes the conditional mean weighted integrated squared error. Based on inspection of several different specifications, I report the results of the Fan (1992) estimator that fits the data as a fourth-order polynomial.

Cleveland (1979)'s locally weighted regression estimates a value, $\hat{y}_i$, for each y_i that gives the smoothed value of y_i, given the corresponding value of x. The locally weighted regression is similar to Fan (1992)'s smoothing technique, except that the regression is weighted so that the central point, (x_i, y_i), gets the highest weight and points that are farther away get less weight. The estimated regression is used to predict $\hat{y}_i$ for y_i only, and is then carried out on all points in the data.

B.4 Personal Income Inequality and the Demand Regime

Following Carvalho and Rezai (2016) and Palley (2016), it is possible that differential marginal propensities to save across classes of wage-earners create a linkage between personal income inequality and the likelihood that demand is wage- or profit-led with respect to the functional income distribution. This linkage presents an issue for using the minimum wage as an instrument for the labor share, because wage-led results may reflect the impact of the minimum wage on personal income inequality, rather than its impact on the functional income distribution. Excluding a measure of personal income inequality from the analysis will therefore cause the estimates to suffer from omitted variable bias. To address this source of potential bias, I estimate the following modified version of the demand equation:

$$u_{it} = \beta_0 + \beta_1 u_{i,t-1} + \beta_2 u_{i,t-2} + \beta_3 \psi_{it} + \beta_4 Gini_{it} + \beta_5 (\psi \times Gini)_{it} + \beta_6 ln(Taxes_{it}) + \eta_i + \delta_t + \epsilon_{it} \tag{B.10}$$

Where $Gini_{it}$ is the Gini coefficient for household income from Frank (2009), constructed using IRS data on adjusted gross income at the state-level. I begin by estimating the model without the interaction term, and then include the interaction between the labor share and the Gini coefficient subsequently. In order to satisfy the IV exclusion restrictions, I include lags of the Gini coefficient and the interaction term as additional instruments. The results are presented in Table (B.1).

The results indicate that state-level personal income inequality exerts a negative effect on aggregate demand separate from the impact of changes in the functional income distribution. In every specification, including OLS and IV estimates, an increase in the Gini coefficient produces a decline in the rate of capacity utilization. However, IV estimates in Columns (2) and (3) suggest that the labor share still has a positive impact on capacity utilization, such that the main wage-led results of the paper are unaltered. In Column (2) the results for the labor share are statistically and economically significant, although the regression coefficient on the labor share provides a lower

Table B.1: Estimation Results - Personal Income Inequality and the Demand Regime

	(1) OLS	(2) IV1	(3) IV2
$u_{i,t-1}$	0.908***	0.920***	0.920***
	(0.0542)	(0.0548)	(0.0546)
$u_{i,t-2}$	-0.136***	-0.105**	-0.104**
	(0.0370)	(0.0440)	(0.0445)
ψ_{it}	-0.247***	0.192***	0.282
	(0.0511)	(0.0569)	(0.203)
$Gini_{it}$	-0.368***	-0.197***	-0.0953
	(0.113)	(0.0735)	(0.197)
$(\psi \times Gini)_{it}$	—	—	-0.173
			(0.384)
$ln(Taxes_{it})$	0.0174	0.0132**	0.0134**
	(0.0125)	(0.00656)	(0.00662)
N	1989	1989	1989
R^2	0.776	0.766	0.766
Time FE	Y	Y	Y
State FE	Y	Y	Y
Sargan-Hansen Test (p-value)	—	0.19	0.20

Notes: Standard errors in parenthesis, clustered at the state-level. * 0.10 ** 0.05 *** 0.01. Column (1) presents the results from OLS estimation of the demand equation, including the state-level Gini coefficient as an additional control. Columns (2) and (3) present the results from IV estimation, using lagged values of the endogenous variables in addition to the minimum wage variables as instruments. In every case the F-statistic from the first stage regression, for each endogenous variable, exceeds 10.

bound to the estimates found in the paper. Column (3), which includes an interaction term between the labor share and the Gini coefficient, weakly implies a relationship between personal income inequality and the demand regime similar to that argued for by Carvalho and Rezai (2016) and Palley (2016). The regression coefficient on the interaction term suggests that as the level of personal income inequality increases the marginal effect of an increase in the labor share becomes increasingly negative—pushing the economy toward profit-ledness. However, the results from the estimation with the interaction term are statistically insignificant. This may suggest that the true effect of the interaction term is zero, but may also be an artifact of inflated variance resulting from collinearity introduced by the interaction term[21].

[21]The unconditional correlation between the Gini coefficient and the labor share is -0.37. The variance inflation factor (VIF) on the interaction term is 942, nearly three times the largest VIF in the model which excludes the interaction term.