State and Local Economic Policy in Equilibrium

by

Matthew G. Wilson

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in The University of Michigan
2021

Doctoral Committee:

Professor Linda Tesar, Chair
Professor James R. Hines, Jr.
Professor Christopher House
Assistant Professor Pablo Ottonello

Matthew G. Wilson

mgwilson@umich.edu

ORCID iD: 0000-0002-2378-8828

For Jameson. “Many daughters have done nobly, but you excel them all.”

ACKNOWLEDGEMENTS

This dissertation would not have been possible without a multitude of people behind me and beside me throughout the process. I feel blessed and grateful to have had such a network, and I can merely attempt to offer sufficient thanks.

I am lucky for such a wonderful dissertation committee, each of whom brought a unique and valuable perspective to my work. Linda Tesar as chair has been a continual source of encouragement, ideas, and advice, refusing to let me settle and pushing me to organize my thoughts into concrete work. Pablo Ottonello is a trusted mentor from whom I learned first-hand how to be a good economist; his generosity with his time and ideas, patience in allowing me to fail, and passion for helping graduate students have had a lasting impact on my career. I thank Jim Hines for many pleasant hours discussing fiscal federalism, and for pushing me to publish my first paper. Chris House has offered a wealth of knowledge and ideas, pushing my work to better places I wouldn't have gone on my own; additionally, his first-year teaching helped nudge me into macroeconomics.

Many other faculty members and economists deserve acknowledgement for their role in this work. John Leahy and Andres Blanco provided pivotal conversations, and contributed greatly (along with Pablo Ottonello) to my graduate macro experience through teaching inspiring second year courses and participating in reading groups. The papers in this dissertation benefited from numerous individual conversations, along with a good deal of seminar and conference presentations, each of which has improved and refined my research.

Aaron Yelowitz has guided me since my first year as an undergraduate, encouraging me to do independent research, study graduate material, and publish my work. Chris Bollinger and Adib Bagh, among many other professors at Kentucky, were also invaluable resources for career advice and assistance as a budding economics student. Finally, I credit the generosity of Paul Coomes in meeting with me as a high school student with confirming my desire to pursue a career in economics.

The University of Michigan is a vibrant graduate department, and I credit numerous conversations with incredibly smart people for refining my thinking and keeping me grounded. I was blessed to share an office with Wenting Song, Guangye Cao, Ellen Stuart, and Mattan Alalouf for many years, each of whom added to a fun and productive environment. I thank Andrew Simon for crossing the hall to have a conversation about federalism and minimum wages, which turned into a fantastic collaboration I have enjoyed greatly. Xing Guo and Luis Baldomero Quintana especially were always generous with advice and suggestions as they went ahead of me in the program.

I am indebted to the support of family and friends throughout this process. I will never forget conversations with my dad as a child that sparked an interest in economics, and my mom has always encouraged and participated in my interests, including my current research. They, along with my brother, created an environment in which curiosity, hard work, integrity, humility, and faith were valued highly, and I carry these values with me always. I have also married into a fantastic family whose love and support has been incredible. Our friends here at New Life Church provided a much needed space for rest and community during even the most difficult parts of graduate school. Many friends around the country have provided much-needed encouragement from afar, and the "Not Summer Crew" group deserves special recognition for keeping me grounded.

Finally, I owe everything to the love of my wife Jameson and daughter Parker. Jameson gave up dreams to come on this journey with me, and for that I will be forever grateful. I don't deserve a woman of her caliber, and I fall more in love with her every day. It is impossible to be with Parker and not feel the light and joy every parent wants to feel. Her arrival pushed me across the finish line, and reminded me what's truly important. Jameson and Parker, this work is as much yours as it is mine.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

LIST OF APPENDICES

Appendix

ABSTRACT

This dissertation studies a subset of state and local economic policy issues and their interactions with federal government policy in the United States, in macroeconomic and general equilibrium models. The first chapter examines the transmission of monetary policy through state and local government public finance. Municipal bond yields increase by 22bp after a 100bp positive monetary shock, though the effect varies across states, and by risk and liquidity. To study the effects of these borrowing cost elasticities on local fiscal policy, I model U.S. localities as small open economies in a monetary union. Here, local governments conduct fiscal policy in response to borrowing costs and economic conditions. In a model calibrated to the U.S., median passthrough of monetary policy shocks to municipal borrowing costs implies a dampening of transmission to output of over half relative to a case which ignores the muni market. Realistic cross-sectional differences in borrowing cost responses result in up to a 50% difference in monetary transmission across localities, and account for 10-20% of observed monetary transmission differences in U.S. data.

The second chapter, jointly written with Andrew Simon, studies the relative merits of centralized and decentralized policy setting in the context of the minimum wage in the U.S. A binding policy is optimal if the benefits from redistribution outweigh the costs from migration, which are relatively steeper for local governments. Centralized

policy, though uniform in practice, reduces horizontal migration externalities, which improves decentralized minimum wage setting. Our results therefore indicate that decentralized and centralized policy setting exhibit strategic complementarity; the extent of which depends on mobility and regional heterogeneity. We then calibrate a model of the continental U.S. and find that joint policy setting leads to a small welfare gain over centralization, and closely resembles the social planner's optimal policies.

The final chapter explores the savings behavior of U.S. states over the business cycle. In the U.S., transfers from federal to state governments respond more strongly to aggregate cycles than state-level cycles. As a result, more U.S. states with more unique business cycles engage in more precautionary savings. A model in which two governments enact fiscal policy, the "regional" government has credit constraints, and the "central" government's information is imperfect, matches the data and implies a sizable information friction. The central government's inability to respond to idiosyncratic shocks implies a role for states in countercyclical policy, despite limitations, and provides caution to centralized fiscal unions.

CHAPTER I

The Municipal Government Channel of Monetary Policy

1.1 Introduction

State and local government debt is a significant sector in the U.S. economy, with important implications for monetary policy. In 2019, the market for state and local government (municipal) bonds was valued at $3.9 trillion, over 1/3 the size of the corporate bond market and greater than 3% of the valuation of the global bond market. Additionally, while most state and local governments in the U.S. have measures in place to prevent an excessive use of debt financing for expenditures, debt financing is nevertheless a key component of municipal public finance and has been increasing over time, especially as interest rates have fallen (Figure 1.1). State and local government debt outstanding is about 100% of state and local government to-

Figure 1.1: Municipal Debt and Treasury Rates, 1960-2020

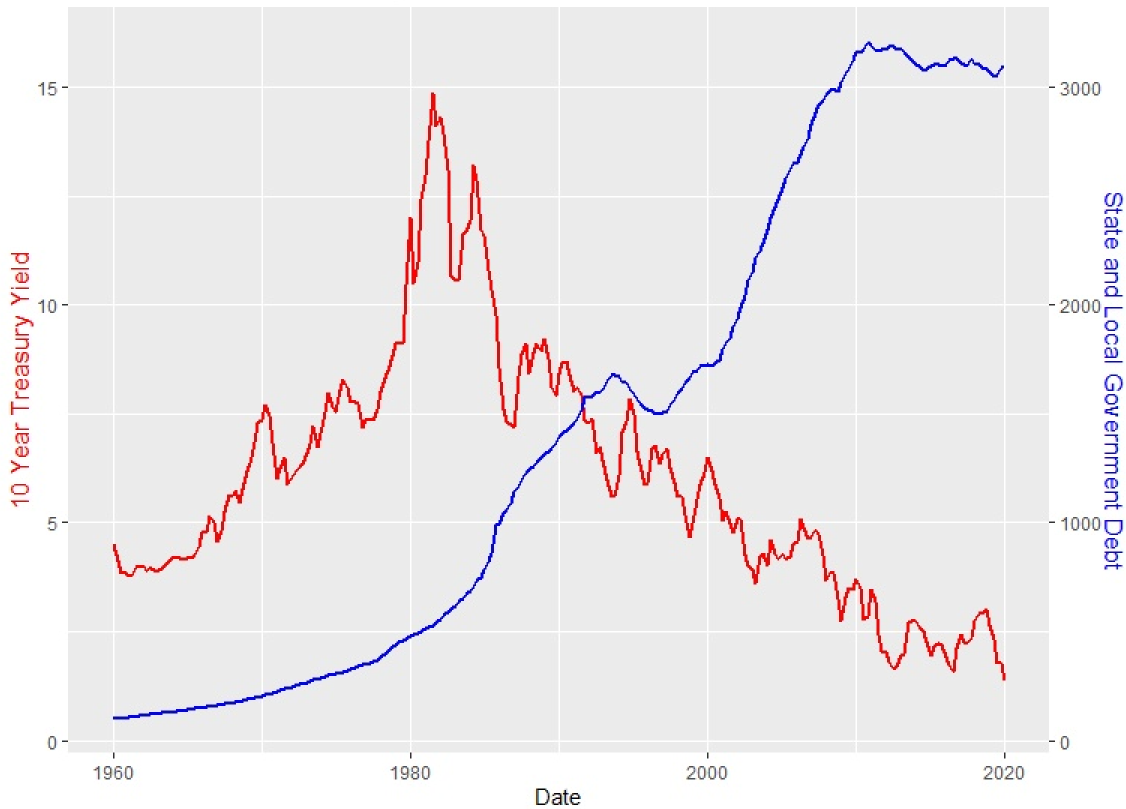

tal expenditures, and interest payments on municipal debt take up around 5% of annual general fund expenditures. Furthermore, rather than "leaning against" monetary expansions, state and local government spending tends to expand when interest rates fall[1], suggesting the presence of a municipal public finance channel of monetary policy transmission, by which national *monetary* policy affects local *fiscal* policy.

[1]Appendix A.3.1 present some time series evidence of the responses of some public finance variables to monetary shocks. A 1 s.d. expansionary monetary shock increases municipal government spending in the medium term by up to 0.25%, after an initial decrease which is possibly due to decreases in automatic stabilizers.

This paper seeks to be the first to describe and illuminate this "Municipal Government Channel" of monetary policy. To do this, I first outline a framework for modeling U.S. localities as small open New Keynesian economies in a monetary union, each with a representative household and fiscal authority. These local fiscal authorities choose debt and public goods spending on locally-produced goods, subject to borrowing costs which are *imperfectly* and *heterogeneously* linked to the risk-free interest rate. The elasticities of these borrowing costs to changes in the risk-fee rate determine the effect of monetary policy on the local government's budget; the more the government's budget relaxes after a monetary expansion, the more it can engage in stimulative spending. In the empirical section of the paper, I use time series and panel data to study the size and source of the response of municipal borrowing costs to monetary shocks, finding that municipal yields decrease (increase) by 22bp in response to a 100bp decrease (increase) in treasury yields, though there is sizeable heterogeneity across U.S. states, driven in part by liquidity and default risk factors. Finally, I calibrate the baseline model to reflect U.S. localities, showing that realistic municipal borrowing cost elasticities result in a significant departure from a "risk-free rate" assumption, imply monetary transmission heterogeneity of up to 50% across states, and can account for 10-20% of transmission heterogeneity in the data.

The paper begins by outlining a model of state and local governments as small open economies in a monetary union, with local governments facing potentially heterogeneous borrowing costs. Each locality is populated by two types of agents, intertemporal and hand-to-mouth, which work and consume to optimize the representative household's utility. Consumption is comprised of tradable and non-tradable goods; the non-tradable goods are produced locally by a New Keynesian production market, with Calvo-style price setting. The locality's representative government receives a stream of tax revenue, and chooses debt issuance and public goods spending to maximize household utility, which also depends on public goods consumption on non-tradable goods. The municipal government does not borrow at the risk-free rate set by the monetary authority; rather, its idiosyncratic borrowing cost is determined on an external financial market and is imperfectly linked to the risk-free rate.

When the national monetary authority lowers the risk-free rate, borrowing costs for the local fiscal authority decrease, relaxing the government's budget constraint and incentivizing public spending. Because public spending occurs with non-tradable goods produced by a New Keynesian market, public spending is stimulative for both output and employment, amplifying the existing expansionary effects of a drop in interest rates. Consequently, the extent to which monetary policy passes through to municipal borrowing costs is of crucial importance for determining the size of this channel. Heterogeneity in municipal debt pricing could result in significant monetary transmission differences; it is to the empirical evidence on the response of municipal bonds to monetary shocks the paper turns next.

To explore the effects of monetary shocks on municipal bond yields, I first use time series evidence from the yields on a series of S&P municipal bond indices. Yields on indices of all general obligation bonds in the United States increase 22bp in response to a 100bp monetary shock, with no difference between state bonds or local bonds. This response represents a far lower elasticity than that of corporate bonds.

Additionally, I find evidence of significant heterogeneity across states, which exhibit coefficients implying responses of 17bp to 40bp. While default risk does explain responses of municipal bonds to monetary shocks, as high risk indices respond more strongly, persistently low responses across the board imply a role for illiquidity in determining these responses.

To investigate the role of illiquidity in municipal bond responses to monetary shocks, I use transaction-level MSRB data of municipal bonds. An exercise as in Gilchrist, Yue, and Zakrajšek (2019), who are concerned with soveriegn bonds, suggest liquidity factors play a role in determining municipal yield responses. This is consistent with literature documenting search frictions, bargaining power, and transaction costs in the municipal bond market. Decomposing average spreads into liquidity and risk components also indicates that movements in spreads after a monetary shock tend to occur in the liquidity component. Finally, I connect the bond microdata for a subset of the sample with annual government finance data, yielding consistent, but not significant, evidence that smaller governments' borrowing costs respond more strongly to monetary shocks; this could be due to risk or liquidity factors.

Armed with empirical estimates of the average and range of municipal yield responses to monetary shocks, I calibrate the small open economy to respresent average U.S. localities, and study the magnitude of the municipal government channel of monetary policy. Using the main response coefficient from the empirical section, I show that including the average passthrough of monetary shocks to municipal yields in the model dampens monetary transmission by over half, relative to a case in which one assumes the local fiscal authority has access to borrowing at the risk-free interest rate. One immediate conclusion, then is that the exact nature of municipal financial markets are crucial in any model of state and local governments over the business cycle.

Additionally, realistic heterogeneity over municipal debt pricing results in meaningful differences in monetary transmission. Increasing a locality's borrowing cost responses from the low end of empirical estimates to the high end of estimates results in up to a 50% increase in monetary transmission to output and employment. Furthermore, the dispersion of peak monetary transmission implied by the state-level empirical estimates can account for 10-20% of observed dispersion of monetary transmission across U.S. localities, based on my own estimates and estimates from the literature. While localities in the U.S. differ on a multitude of dimensions affecting monetary transmission, the ability of monetary policy to influence their borrowing costs is a factor policymakers should consider.

Given that I have constructed a model of state and local governments that links financial markets and business cycles, it is natural to use the model to lend insight into other macroeconomic questions. I use the model to provide an explanation for why state and local government spending decreased in the wake of the 2008 recession, the first time it has ever done so. In my model, a real recession induces fiscal stimulus, but the same recession combined with a lockup of financial markets sees a fiscal contraction; this result suggests the financial nature of the 2008 recession placed heavy constraints on the ability of municipal governments to respond. I also consider a shock to the model representing the early days of the COVID-19 pandemic;

in the model, quick actions by the Fed to shore up municipal debt markets prevent a prolonged dampening of government spending.

Related Literature. This paper provides meaningful contributions to a number of important strands of economic literature. The baseline model is, in most ways, a canonical open economy New Keynesian model. In this vein, it adds to papers such as Galí and Monacelli (2008), Beetsma and Jensen (2005), Farhi and Werning (2017a), Farhi and Werning (2017b), Nakamura and Steinsson (2014a), and Chodorow-Reich (2019), which study monetary and fiscal policy in monetary unions, by highlighting that interest rates for the member governments of a monetary union may differ substantially in response to the same monetary policies. Similarly, in showing how monetary policy works *through* municipal fiscal policy, this paper merges the monetary union literature with the literature on monetary policy passthrough, exemplified by Bernanke, Gertler, and Gilchrist (1999), McKay, Nakamura, and Steinsson (2016), and Kaplan, Moll, and Violante (2018). More specifically, although focused on local governments, this paper contributes to a literature on international monetary transmission to small open economies, as in Auer et al. (2019) and Cesa-Bianchi, Thwaites, and Vicondoa (2016). In contrast to papers which study optimal fiscal policy for a member of a monetary union, I study the government as an agent in the model, whose behavior in response to monetary policy is taken as part of the passthrough effect.

By analyzing a model of a locality in the U.S., this paper enters in to the discussion on regional effects of monetary policy, and macroeconomic models with regions in general, as in Beraja et al. (2019) and Beraja, Hurst, and Ospina (2019). Other papers, such as Seegert (2015), Cashin et al. (2018), and Fisher and Wassmer (2014), examine the responses of state and local governments to significant macro events; I use the model in this paper as a playground to study the behavior of state and local governments in the aftermath of the financial crisis. Finally, the over-the-counter markets version of the model in Section 1.2.8 builds on the work of Duffie, Garleanu, and Pedersen (2005), who model OTC markets for financial assets, and Bethune, Sultanum, and Trachter (2019), who model the issuance side of OTC markets. I show in the paper how such a model can be connected to a DSGE macro model as a microfounded explanation for why borrowing costs may differ from the risk-free rate for local governments.

The results from the empirical section of this paper contribute to a number of strands of literature. First, and most obviously, this paper adds to recent work on the effect of U.S. monetary policy on various asset prices. Rosa (2014) does this for municipal bonds; I expand on his work by including a host of indices and exploiting a trade-level panel dataset to investigate potential determinants of muni responses to monetary shocks. Gilchrist, Yue, and Zakrajšek (2019) studies the response of international sovereign yields to U.S. monetary shocks; I perform similar exercises to their paper, but in the U.S. municipal bond market. Anderson and Cesa-Bianchi (2020) study how firm leverage affects corporate bond responses to monetary shocks.

By shedding light on the relationship between municipal bond yields and U.S. monetary policy, I also add to a robust literature on municipal bond pricing, especially as it relates to risk and liquidity; mine is the first paper to explore explicitly the roles of these channels in determining monetary policy responses. Two important papers, Schwert (2017) and Ang, Bhansali, and Xing (2014), debate the relative importance of risk and liquidity in municipal bond spreads, with the former emphasizing risk and the latter liquidity. Another strand of papers (Harris and Piwowar, 2006; Green, Hollifield, and Schürhoff, 2007a; Green, Hollifield, and Schürhoff, 2007b; Brancaccio, Li, and Schúrhoff, 2020; Garrett et al., 2018; Moldogaziev, 2018) highlights explicit frictions in the secondary market for municipal bonds, such as information asymmetries and market power, that result in price dispersion over given bonds. A number of other papers explore determinants of municipal bond prices, from state laws on bankruptcy (Yang, 2019) to climate change (Painter, 2020). Other relevant municipal bond pricing papers include Gao, Lee, and Murphy (2019), Grigoris (2019), Adelino, Cunha, and Ferreira (2017), and a host of others.

1.2 A Model of a Municipal Government

In order to examine the effects of the financial market for municipal bonds on municipal behavior and household welfare, I propose a quantitative heterogeneous agents DSGE model of municipalities in a monetary union, with municipal government debt sold on outside markets, subject to financial frictions. The economies in question are small open economies,[2] reflecting the tens of thousands of distinct municipal governments in the United States. In this model, households choose labor and purchases of local municipal bonds, financial markets buy and sell bonds, local governments choose spending debt issuance, and the central government chooses the risk-free in-

[2]The model builds on the canonical Calvo model found in Schmitt-Grohe and Uribe (2016), adding the municipal government sector and hand-to-mouth consumers.

terest rate. Local output is produced by monopolistically competitive firms using labor from households, resulting in standard New Keynesian features for prices.

This model is able to quantify the effects of financial frictions, which appear here as a wedge between the risk-free rate and municipal borrowing costs, and may potentially arise from an OTC framework, on local fiscal policy in response to macroeconomic shocks. As such, it serves as a contribution to the macro literature on passthrough of shocks, the fiscal policy literature, monetary unions models, models of the regional effects of macroeconomic shocks, and (in one possible application) models on OTC asset pricing. Note that, in this section, as well as in the main quantitative results, I focus on the problem of a single locality for simplicity.

1.2.1 Environment

The locality is modelled as a small open economy with a representative household and representative government, each of which maximizes the utility of the representative consumer. The household is made up of $1 - \kappa$ traditional consumers and κ "hand to mouth," or HTM, consumers. The traditional, or "Ricardian," optimizers in the household choose a consumption bundle comprised of two types of final goods, tradable (c^T) and non-tradable (c^N), as well as labor supply h and debt d^H. The government chooses government purchases of non-tradable goods g and municipal debt d^G, given a tax rate τ^G on the exogenous stream of tradable goods for the economy, y^T. HTM consumers simply choose labor h^H, and consume exactly their labor income in every period.

Tradable goods and bonds of both agents are traded with the rest of the world, where the locality is endowed with an exogenous income of the tradable good in every period. Non-tradable goods are produced using domestically supplied labor by monopolistically competitive firms within the region to satisfy demand for non-tradable consumption and government purchases of public goods. Inflation is induced by Calvo-style price setting on the part of these monopolistic competitors, who maximize expected future profits subject to household and government demand, as well as the expected constraints on price changes.

The aggregate risk-free interest rate in the economy is determined by a central authority and is exogenous with respect to local variables. Additionally, both the household and the government are subject to their own "proprietary" interest rates, r^H and r^G, which depend on the aggregate rate, deviation of debt from steady state, and parameters determining the relationship between monetary shocks and the actual interest rate paid by either the household or the government. Of particular interest in this project is the response of the *government's* borrowing costs to monetary shocks. The strength of this response will affect passthrough of interest rate shocks to households, and presents a potential source of heterogeneity of monetary passthrough in the U.S. economy.

1.2.2 Household

1.2.2.1 Basic Problem

The household in the small economy is made up of $1-\kappa$ "Ricardian" agents and κ "hand-to-mouth" agents. The representative Ricardian optimizer solves

$$\max_{\{c_t^T, c_t^N, h_t, d_{t+1}\}} \mathbb{E}_0 \beta^t \big[U(c_t) - V(h_t) + W(g_t)\big], \tag{1.1}$$

where

$$c_t = A(c_t^T, c_t^N) \tag{1.2}$$

$$c_t^T + p_t c_t^N + d_t = y_t^T(1-\tau^G) + w_t h_t + \frac{d_{t+1}}{1+r_t^H} + T_t. \tag{1.3}$$

Here, tradable consumption and debt are denominated in terms of the "national" price, which is normalized to unity. Prices p_t and w_t are prices–of non-tradables and labor, respectively–relative to the price of the tradable good. τ^G is the tax on exogenous tradable good allocation, and T_t is the lump-sum transfers to households from firm profits. Quantities are in per-person terms, such that total labor supply from optimizers is given by $(1-\kappa)h_t$.

1.2.2.2 Optimality Conditions

Accordingly, the household's first order conditions are given as follows:

$$\lambda_t = U'(c_t) A_1(c_t^T, c_t^N) \tag{1.4}$$

$$\frac{A_2(c_t^T, c_t^N)}{A_1(c_t^T, c_t^N)} = p_t \tag{1.5}$$

$$\frac{\lambda_t}{1 + r_t^H} = \beta \mathbb{E}_t \lambda_{t+1} \tag{1.6}$$

$$V'(h_t) = \lambda_t w_t. \tag{1.7}$$

Equation 1.4 is the first order condition for tradable goods consumption, defining the shadow value of income denominated in the tradable goods price. The condition for non-tradable consumption, when plugged into Equation 1.4, yields 1.5, which allocates consumption according to the relative price of the two goods. Equation 1.6 trades off the benefits of borrowing today with the costs of paying it back tomorrow, and Equation 1.7 equates marginal costs and benefits of labor supply.

1.2.2.3 Hand-to-Mouth Agents

The remaining κ agents in the economy behave in a hand-to-mouth manner. These agents supply per-capita labor h_t^H, and use labor income to consume tradable and non-tradable goods:

$$c_t^{T,H} + p_t c_t^{N,H} = w_t h_t^H. \tag{1.8}$$

Here, as above, consumption is aggregated according to $c_t^H = A(c_t^{T,H}, c_t^{N,H})$.

The first order conditions for these consumers mirror those of the traditional agents, without the intertemporal condition:

$$\lambda_t^H = U'(c_t^H) A_1(c_t^{T,H}, c_t^{N,H}) \tag{1.9}$$

$$\frac{A_2(c_t^{T,H}, c_t^{N,H})}{A_1(c_t^{T,H}, c_t^{N,H})} = p_t \tag{1.10}$$

$$V'(h_t^H) = \lambda_t^H w_t. \tag{1.11}$$

The total amount of non-tradable consumption in the economy, then, is the sum $(1-\kappa)c_t^N + \kappa c_t^{N,H}$, and likewise with tradable consumption $(1-\kappa)c_t^T + \kappa c_t^{T,H}$ and total labor $(1-\kappa)h_t + \kappa h_t^H$. This departure from Ricardian equivalence in the model allows a realistic local government spending multiplier to be calculated.

1.2.3 Local Government

The representative local government uses local fiscal policy to solve the problem 1.1[3] by choosing a borrowing level d_{t+1}^G subject to

$$g_t = \tau^G y_t^T + \frac{d_{t+1}^G}{1+r_t^G} - d_t, \tag{1.12}$$

Other than its choice variables, another key difference emerges for the local government: it exerts full market power over its debt. In other words, while the representa-

[3]Note that fiscal policy in this model is "passive" in the sense that policymakers are concerned primarily with the efficient provision of public goods. The local policymaker does not factor explicitly its stimulative effects on the economy. Beetsma and Jensen (2005) shows that such passive policies in a two-economy fiscal union result in welfare loss relative to centralized optimization or fiscal policy rules, and Carlino and Inman (2013) finds that indeed U.S. states have the tools to achieve stabilization policy through deficit spending. However, it is reasonable to think of small local economies in the U.S. as being more concerned with public goods provision than stabilization at the level of a fiscal union. In a case in which the local government was optimizing according to a fiscal rule, or if fiscal policy was coordinated at the union level, its borrowing costs would not matter for fiscal adjustments. Additionally, these adjustments would "lean against" monetary policy, which is contrary to the data on local government responses.

tive household is a price taker with respect to the interest rate, the local government can only be thought of as a singular agent and the only issuer of its asset. The government, therefore, must take into account the effect of its debt issuances on its borrowing costs, both in the current period and in the future. The first order condition with respect to debt purchases, then, is given by

$$W'(g_t)\frac{1 + r_t^G - d_{t+1}^G \frac{\partial r_t^G}{\partial d_{t+1}^G}}{(1 + r_t^G)^2} = \beta \mathbb{E} W'(g_{t+1})\Big(-1 + \frac{-d_{t+2}^G \frac{\partial r_{t+1}^G}{\partial d_{t+1}^G}}{(1 + r_{t+1}^G)^2}\Big). \tag{1.13}$$

In an analog to the household's debt decision, the government uses debt issues to balance government increased spending now with decreased spending from debt obligations later. When the yields on municipal bonds move strongly with interest rate shocks, these fiscal policy responses will tend to be greater, while the opposite is true when responses of yields to monetary shocks are weak.

The condition that the government take into account its effect on interest rates is not an innocent one. It affects the response of debt to transitory shocks, but it also has an effect on the steady state of the model. Figure 1.2 presents this as an illustrative example using the functional forms and calibration from the quantitative section of the paper. Failing to account for the effects of debt on borrowing costs would result in massive increases in government debt, under the same parameterization.

Figure 1.2: Steady State Comparison, Local Government Debt

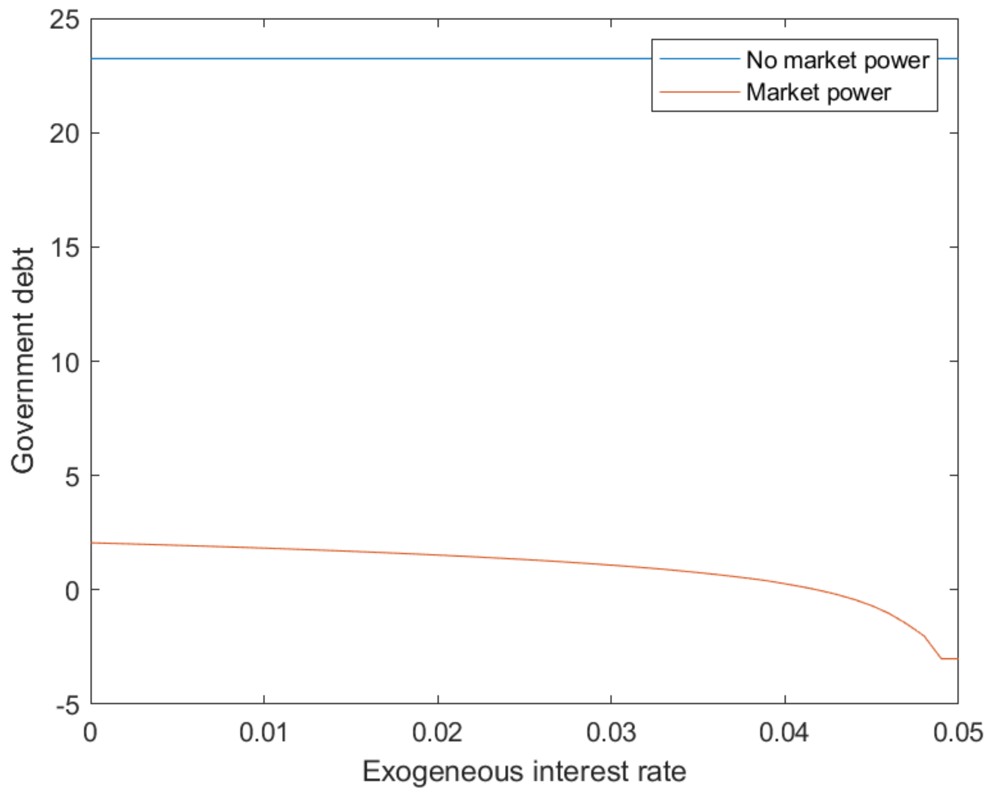

1.2.4 New Keynesian Production

1.2.4.1 Final Goods Production

The non-tradable good y_t^N is produced by a final goods producer which buys intermediate goods y_{it}^N from a continuum of intermediate goods producers in the local economy. Production of the final good from intermediates is determined by the aggregating equation

$$y_t^N = \Big(\int_0^1 (y_{it}^N)^{1-\frac{1}{\mu}} di \Big)^{\frac{1}{1-\frac{1}{\mu}}}, \tag{1.14}$$

and final goods firm profits are given by

$$P_t^N y_t^N - \int_0^1 P_{it}^N y_{it}^N di.$$

Profit maximization on the part of the final goods producer implies the demand equations for the monopolistically competitive intermediate goods producers

$$y_{it}^N = y_t^N \left(\frac{P_{it}^N}{P_t^N} \right)^{-\mu}.$$

Here, the domestically produced good is used both for consumption–by both types of agents–and government spending,

$$y_t^N = (1-\kappa) c_t^N + \kappa c_t^{N,H} + g_t, \tag{1.15}$$

so the relevant demand equations become

$$y_{it}^N = ((1-\kappa)c_t^N + \kappa c_t^{N,H} + g_t)\left(\frac{P_{it}^N}{P_t^N}\right)^{-\mu}. \tag{1.16}$$

P_t^N here is the price of final non-tradable goods, which is given by the aggregator $P_t^N = \left(\int_0^1 (P_{it}^N)^{1-\mu} di\right)^{\frac{1}{1-\mu}}$.

1.2.4.2 Intermediate Goods Producers

Intermediate goods firms exist on the continuum $[0,1]$, and produce differentiate inputs to the final non-tradable good using household labor:

$$y_{it}^N = h_{it}^\alpha, \ \alpha \in (0,1]. \tag{1.17}$$

The choice variable for these firms is the price for good i, P_{it}^N, which determines demand for the intermediate good as in Equation 1.16. Profit for an individual intermediate goods firm is given by $P_{it}^N y_{it}^N - (1-\frac{1}{\mu})W_t h_{it}$, where W_t is the raw wage and $(1-\frac{1}{\mu})$ is a labor subsidy meant to offset the distortions from monopolistic competition.

Plugging in the demand equations and production functions, period t profits for the intermediate goods firm are given as a function of the chosen price P_{it}^N:

$$P_{it}^N((1-\kappa)c_t^N + \kappa c_t^{N,H} + g_t)\left(\frac{P_{it}^N}{P_t^N}\right)^{-\mu} - (1-\frac{1}{\mu})W_t((1-\kappa)c_t^N + \kappa c_t^{N,H} + g_t)^{\frac{1}{\alpha}}\left(\frac{P_{it}^N}{P_t^N}\right)^{-\frac{\mu}{\alpha}}.$$

Prices are sticky according to a Calvo mechanism, i.e., intermediate goods firms may only change prices in each period with probability $(1-\theta)$. In Appendix A.1, I show that maximization of expected profits on the part of intermediate goods firms, since all price-adjusting firms choose the same price, result in choosing the flexible relative price $\tilde{p_t}^N$ to equate present value marginal costs and marginal revenues, $mr = mc$, where

$$mr_t = \frac{\mu - 1}{\mu} y_t^N p_t (\tilde{p_t}^N)^{1-\mu} + \beta\theta\mathbb{E}_t \frac{\lambda_{t+1}}{\lambda_t} \left(\frac{\tilde{p_t}^N}{\tilde{p_{t+1}}^N} \frac{1}{\pi_{t+1}^N} \right)^{1-\mu} mr_{t+1} \tag{1.18}$$

and

$$mc_t = \frac{1 - \frac{1}{\mu}}{\alpha} (y_t^N)^{\frac{1}{\alpha}} w_t (\tilde{p_t}^N)^{-\frac{\mu}{\alpha}} + \beta\theta\mathbb{E}_t \frac{\lambda_{t+1}}{\lambda_t} \left(\frac{\tilde{p_t}^N}{\tilde{p_{t+1}}^N} \frac{1}{\pi_{t+1}^N} \right)^{-\frac{\mu}{\alpha}} mc_{t+1}. \tag{1.19}$$

I also show in Appendix A.1 that inflation dynamics are given by

$$1 = \theta(\pi_t^N)^{\mu-1} + (1-\theta)(\tilde{p_t}^N)^{1-\mu}, \tag{1.20}$$

where $\pi_t^N = \frac{P_t^N}{P_{t-1}^N}$, and aggregate production is given by $y_t^N = s_t^{-\alpha} h_t^{\alpha}$, where $s_t = \int_0^1 \left(\frac{P_{it}^N}{P_t^N} \right)^{-\frac{\mu}{\alpha}}$ represents the amount of price dispersion in the intermediate goods sector that has a dampening effect on aggregate output. Price dispersion evolves according to

$$1 = \theta s_{t-1} (\pi_t^N)^{\mu/\alpha} + (1-\theta)(\tilde{p_t}^N)^{-\mu/\alpha}, \tag{1.21}$$

1.2.5 Financial Sector

As mentioned before, household and government debt are traded on an external financial market, resulting in two interest rates r_t^H and r_t^G. These interest rates reflect the aggregate interest rate r_t^*, debt stock/purchases d_t, d_{t+1}, d_t^G, and d_{t+1}^G, and monetary shocks. The interest rates can be thought of as being determined by the functions

$$r_t^H = f^H(d_t, d_{t+1}, r_t^*, m_t) \tag{1.22}$$

and

$$r_t^G = f^G(d_t^G, d_{t+1}^G, r_t^*, m_t). \tag{1.23}$$

The function $f^G(m_t)$ is of particular interest in this paper, as it will be a key determinant of the passthrough of monetary policy to local governments. The response of municipal government borrowing costs could vary based on a multitude of factors, including trading costs and illiquidity in the muni market, as I will show in the empirical section.

These pricing functions could arise from a number of financial market specifications. For example, a common formulation in the international literature is the debt-elastic interest rates of Schmitt-Grohe and Uribe (2003). Below, I will show another alternative to this formulation, which incorporates a standard over-the-counter asset market model into the model; such a feature serves as a microfoundation for the functions $f^H(m_t)$ and $f^G(m_t)$ and provides intuition into the mechanism working behind the empirical investigations.

1.2.6 Aggregation

Thus far I have focused on the problem of a single locality. In a full version of the model, there are a large number of localities subscripted by $s \in \{1, ..., S\}$. These economies exchange tradable goods and debt with each other, as well as with the rest of the world.[4] The nationwide interest rate r_t^* is set by the economy's monetary authority in response to aggregate output and inflation, which are made up of of local values:

$$r_t^* = R(\{y_{st}^T\}_s, \{y_{st}^T\}_s, \{\pi_{st}\}_s). \tag{1.24}$$

1.2.7 Equilibrium

A competitive equilibrium is a set of quantities y_t^N,c_t^T, c_t^N, h_t, g_t, d_t, d_t^G, λ_T, π_t^N, s_t, mr_t, mc_t and prices p_t, w_t, $\tilde{p}_t{}^N$, r_t^H, and r_t^G for each locality s satisfying:

(i) The optimizing household problem is solved by Equations (1.4), (1.5), (1.7), and (1.6)

(ii) Hand-to-mouth quantities satisfy Equations (1.11), (1.9), and (1.10)

(iii) The government's problem is solved by Equation (1.13)

(iv) Marginal revenue and marginal cost are given by Equations (1.18) and (1.19), where $mr = mc$

[4]The relaxation of market clearing at the union-wide level is essentially a preservation of the assumption of an external financial market with which households and governments trade. This allows the monetary authority to easily set r_t^*. Such an approach is consistent with the regional model found in Beraja et al. (2019).

(v) Aggregate production satisfies $y_t^N = s_t^{-\alpha}((1-\kappa)h_t + \kappa h_t^H)^\alpha$

(vi) Inflation and price dispersion evolve according to Equations (1.20) and (1.21)

(vii) Inflation is defined by $\pi_t^N = \frac{p_t}{p_{t-1}}$

(viii) Market clearing in tradable goods implies $(1-\kappa)c_t^T + \kappa c_t^{T,H} + d_t = y_t^T + \frac{d_{t+1}}{1+r_t^H}$

(ix) Interest rates satisfy Equations (1.22) and (1.23)

(x) The risk free rate is set by the monetary authority according to Equation 1.24

given the exogenous processes $y_t^T$and initial conditions s_{-1}, d_0, and d_0^G.

1.2.8 An Over-the-Counter Markets Model for Debt Pricing

While the borrowing cost functions in the model above are presented in a reduced form way, the finance literature provides a path to a microfounded relationship between monetary shocks and municipal yields. Specifically, I consider the class of models for which Duffie, Garleanu, and Pedersen (2005) was the seminal work, in which the trading of assets on secondary OTC markets is modeled carefully. In summary, municipal bonds are bought and sold to risk-neutral financial firms on primary markets, then sold on secondary markets to buyers who value the bonds highly but are subject to trading costs or incomplete market power. This friction in the secondary market dampens the response of the present value of the asset for financial firms to changes in the aggregate interest rate, thereby muting the primary market price response to monetary policy.

Every period, the municipal government makes debt issues x_t. Municipal bonds pay coupon rate c and mature with probability ν. Governments buy and sell from risk-neutral financial firms at competitive prices. A mass α of municipal buyers purchase bonds from financial firms; these buyers value the asset above its present value, at value v^H. This high valuation can be thought of as reflecting the tax advantage in municipals or warm-glow utility from supporting projects in one's community.[5]

The value of the bond for the financial firm, v^F, then, is the present value of the bond at time t, discounted by the expected path of aggregate interest rates r_t:

$$v_t^F = \mathbb{E}_t \sum_{s}^{\infty} \left[c(1-\nu)(1-p_{t+s}^{sell}) + (1-\nu)P_{t+s}p_{t+s}^{sell} + \nu \right] \prod_{k=1}^{s} \frac{1}{1+r_k}(1-p_{t+k})(1-\nu)). \tag{1.25}$$

The price P_t is the price of the bond on the secondary market, which is determined by Nash bargaining, as in the OTC literature:

$$P_t = \theta v_t^F + (1-\theta)v^H. \tag{1.26}$$

θ is a key parameter in the OTC model: it captures the financial frictions in the market resulting from trading costs or asymmetric information, which can broadly be described as contributing to illiquidity in munis. Additionally, p_t is the probability that a given muni held by the financial firm–that does not mature–is sold in period

[5] In support of both of these motivations, Pirinsky and Wang (2011) shows a great deal of market segmentation in the muni market, wherein household buyers tend to buy munis primarily from their own geographic areas.

t, which is given by the system

$$p_t^{sell} = \max\left\{\frac{\alpha - (1-\nu)B_t^H}{(1-\nu)B_t^F}, 1\right\}$$

$$B_{t+1}^F = (1 - p_t^{sell})(1-\nu)B_t^F + x_t$$

$$B_{t+1}^H = (1-\nu)B_t^H + p_t^{sell}(1-\nu)B_t^F,$$

where B_t^F and B_t^H are the bond holdings of financial firms and buyers, respectively. At any time t, we have $D_t = B_t^F + B_t^H$.

For simplicity, assume municipal governments face the competitive price–the financial firm's valuation of the bonds–for their issuances and purchases of municipal bonds, v^F. We can now transform this model into the structure of the full model, where $r_t^G = f^H(d_t, d_{t+1}, r_t^*, m_t)$. In the model, the government's net income from debt purchases is given by $y_t^G = \frac{d_{t+1}^G}{1+r_t^G} - d_t^G$. In the context of an OTC model, this income is given by $y_t^G = v_t^F x_t - (\nu + (1-\nu)c)d_t^G$. By setting these terms equal to each other, we get that the *effective* interest rate at time t for the municipal government is given by

$$r_t^G = \frac{d_{t+1}^G}{v_t^F d_{t+1}^G + (1-\nu)(1-c-v_t^F)d_t^G} - 1, \tag{1.27}$$

where v_t^F is defined as above. Here, monetary shocks work through the term v_t^F, as they affect the future path of aggregate interest rates r_t.

For this formulation of debt pricing to make sense, we need both that $\frac{\partial r_t^G}{\partial d_{t+1}^G} > 0$ and $\frac{\partial r_t^G}{\partial r_t^*} > 0$.

1.2.9 Elasticities of Interest

A few key elasticities are important for understanding the passthrough of monetary policy through municipal public finance and its potential heterogeneity using this model. First, we need to know the effect of borrowing costs on government spending. Results in Table 1.2 suggest these effects could be quite sizeable, though these results are merely suggestive and not the main focus of the paper. Additionally, it is important to know the local government spending multiplier: a helpful review in Chodorow-Reich (2019) suggests a point estimate of 1.8, suggesting a meaningful role of fiscal policy at the local level.

Of course, knowing the effects of an aggregate interest rate shock on the borrowing costs of municipal governments is a crucial piece of studying the potential size of this channel. Furthermore, any heterogeneity in these yield elasticities will result in differential fiscal policy responses across municipalities, and therefore different passthroughs in different localities. In the following empirical section, I investigate in depth the effect of monetary shocks on municipal bond yields, identifying significant heterogeneity in responses across government and highlighting a few possible sources.

1.3 The Effect of U.S. Monetary Shocks on Municipal Bonds

The key elasticity in the model outlined above is the effect of interest rate shocks on the borrowing costs of state and local governments in the U.S. The extent to which these borrowing costs respond to monetary shocks will determine the strength of the municipal public finance channel of monetary policy; my baseline estimate of the response coefficient is 0.22. Additionally, heterogeneity across municipalities in the responses of their borrowing costs to monetary shocks will imply heterogeneity in the transmission of monetary policy to households, as shocks will induce dispersion of borrowing costs. This section documents the average effects of monetary shocks on municipal borrowing costs, as well as the variance of such effects across state and local governments. I also investigate whether we can identify root causes of response heterogeneity, such as liquidity or risk in the muni market. I find evidence for both possibilities; a transaction-level dataset suggests illiquidity as a key factor, while high-yield (lower rated) indices respond twice as strongly to monetary shocks than the most highly rated indices.

I begin the section with a brief description of the municipal bond market, explaining the similarities and differences between munis and treasuries. Next I present summary statistics from the muni market, highlight its behavior during the financial crisis, and discuss the monetary shock identification strategy. I also provide some brief evidence that municipal bond yields on the secondary market are valid representations of municipal borrowing costs; in short, secondary market yields both reflect primary market prices and have a statistically significant effect on municipal government behavior. After this, I move on to the main empirical exercises.

The first set of exercises investigate the time series evidence on the effect of monetary shocks on a set of muni indices; this section is in the spirit of Rosa (2014), who looks at munis, and Gilchrist, Yue, and Zakrajšek (2019), who study foreign bond responses to U.S. interest rate shocks. I find that an index of General Obligation

(GO) municipal bonds responds to a 100bp monetary shock by an average of 22bp; furthermore, I find evidence of substantial heterogeneity by state, as responses vary from 17bp to 40bp across space, despite little difference between state and local bonds. High-risk indices respond more strongly to monetary shocks, as one might expect, but the coefficients remain persistently low, suggesting a role for illiquidity in dampening these coefficients.[6]

The second set of exercises exploits a trade-level dataset of the municipal bond market, in which liquidity and risk can be explored further as possible drivers of monetary transmission (or lack thereof). First, I perform an exercise from Gilchrist, Yue, and Zakrajšek (2019), finding joint significance for some possible indicators of illiquidity, but not risk. Second, I replicate the main strategy of Schwert (2017), who decomposes the (tax-adjusted) municipal spread into liquidity and risk components, and compute the response of each of these spread components to monetary shocks, finding that the liquidity component of spreads may exhibit a transitory response to monetary shocks. Finally, I take a sample of munis from the largest state and local governments, for which I can match annual finance data from the Census of Governments, finding that correlates of government size tend to lower the response coefficient, though without much statistical significance. In any case, my overall results suggest dampened (relative to, say, corporate bonds) but heterogeneous responses of municipal borrowing costs to U.S. interest rate shocks, with which I can use the municipal open economy model to evaluate the characteristics of monetary

[6]The tax-free nature of municipal bonds, in the absence of risk or illiquidity, would imply coefficients of $1 - \tau$, where τ is the relevant tax rate. Coefficients lower than this, as I find here, suggest the presence of illiquidity.

passthrough via state and local governments.

1.3.1 The Municipal Bond Market

State and local governments in the U.S. rely heavily on debt markets to finance a wide range of activities, from covering budget shortfalls to infrastructure investment. When a government decides to raise funds through a debt issue, it issues municipal bonds through a financial broker, one of a number of bidders for the rights to issue the bonds. At the time of issue, the broker sells the bonds on a "primary market" investors and other broker-dealers.

After the primary market sale, municipal bonds are traded in Over-the-Counter (OTC) markets, rather than on a central exchange. The OTC nature of the municipal secondary market requires a specific buyer-seller match for a sale to take place. Harris and Piwowar (2006) and Green, Hollifield, and Schürhoff (2007b) show that buyer characteristics, namely size, influence prices on the secondary market, such that there is price dispersion even on the same day for a given bond; some investors may have better information and market power than others. Municipal bonds are not like treasuries in that they are tax-free and not entirely risk-free; additionally, the OTC market induces and amount of transaction costs into this market. The average municipal bond is traded every 10 days, suggesting that price adjustment may be slower in this market than in other markets.

There is a wide dispersion on the yields of municipal bonds in the MSRB data discussed below, reflecting a high degree of heterogeneity in municipal borrowing costs. Table 1.1 presents basic summary statistics on these yields, and Figures 1.3

and 1.4 shows the dispersion of yields conditional on time to maturity for the years just before and during the financial crisis. As mentioned in the introduction, there are three main avenues by which yields on municipal bonds might differ from yields on U.S. treasuries of similar maturities, as well as yields of other municipal bonds: taxes, risk, and liquidity. There are some differences in marginal tax rates across locations in the U.S., which may be driving a small portion of these differences, but as Figure 1.5 shows, the variance of muni spreads varies at a higher frequency than changes in tax rates. Most importantly, all three of these figures suggest a massive change in yield dispersion in the wake of the recent financial crisis[7].

Returning to Figure 1.5, note the systematic difference in spread behavior in the aftermath of the financial crisis. In addition to the increase in the variance of spreads, note that spreads move from negative to positive on average. Despite the tax advantage of municipals, yields were higher during and after the crisis than comparable treasuries; in other words, yields on municipals did not decrease as rapidly as yields on treasuries during a time of unprecedented monetary expansion. For further motivation of the importance of market inefficiencies in the behavior of muni yields, Figure 1.6 shows average daily yields for the same time period, broken out by the IRC illiquidity measure used in Schwert (2017) and described in Appendix A.2.3. The bonds recorded as "illiquid" on this measure, i.e., those with high average imputed trading costs, saw a bigger increase in spreads during the crisis, despite similar behavior beforehand. This movement is suggestive that macroeconomic shocks may have heterogeneous effects on municipal borrowing costs; I return to an experiment

[7] The yields plotted here are simply the raw trade-level yields in the MSRB data, described in Appendix A.2.1.

Table 1.1: Summary Statistics, Daily Municipal Bond Yields

	Mean	Variance	25th	50th	75th
Yield, 2005-2019	2.760	2.077	1.630	2.743	3.758
Yield, 2008-2014	2.712	2.494	1.452	2.600	3.754
Spread, 2005-2019	0.4216	1.856	-0.535	0.155	1.147
Spread, 2008-2014	0.9137	2.171	-0.076	0.767	1.774

Note: Each variable is the median of daily trades for each municipal bonds. The sample does not include a bond on days in which it is not traded. Spreads are calculated as the difference between a bond's yield and an interpolated hypothetical U.S. treasury bond of the same length to maturity in days.

with a financial crisis later in the paper.

Heterogeneity in yields will be important for the upcoming framework, as I will show that there is also a heterogeneity over the degree to which muni prices respond to monetary shocks. This heterogeneity in response implies that different types of governments not only have varying borrowing costs over the long run, but will also encounter differing changes in borrowing costs after short-run shocks. As a result, the transmission of shocks through municipal borrowing costs should be expected to have both level and distributional effects, as governments are affected differentially by shocks, depending on factors such as trading costs and liquidity.

1.3.2 The Importance of the Secondary Market for Municipal Finance

So far I have described the secondary market for state and local government debt, i.e., the market in which the yields on bonds which have already been issued are determined. Of course, the going yield on bond A does not affect the payments of its issuer, which simply continues to make the predetermined coupon payments to

Figure 1.3: Muni Yields on FOMC Dates, 2006

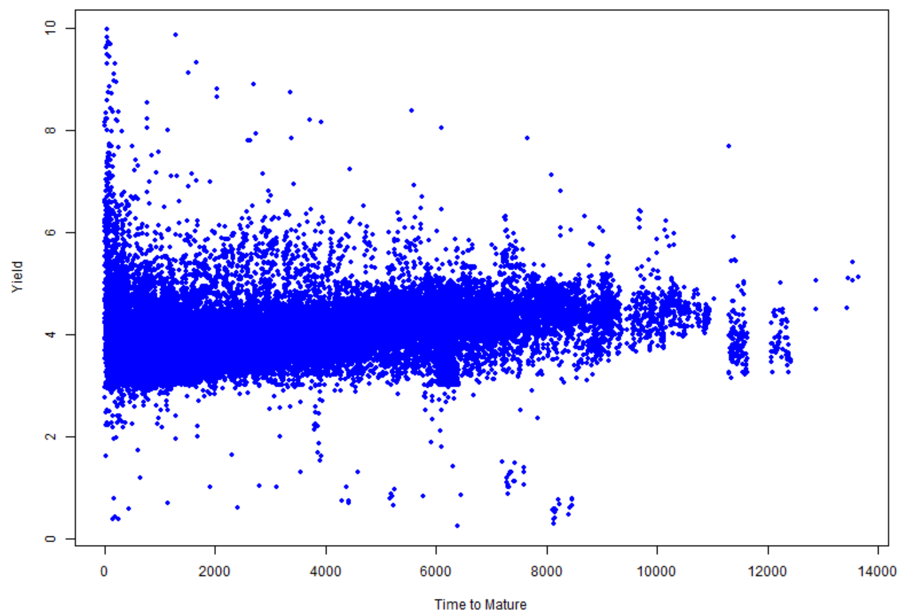

Figure 1.4: Muni Yields on FOMC Dates, 2008

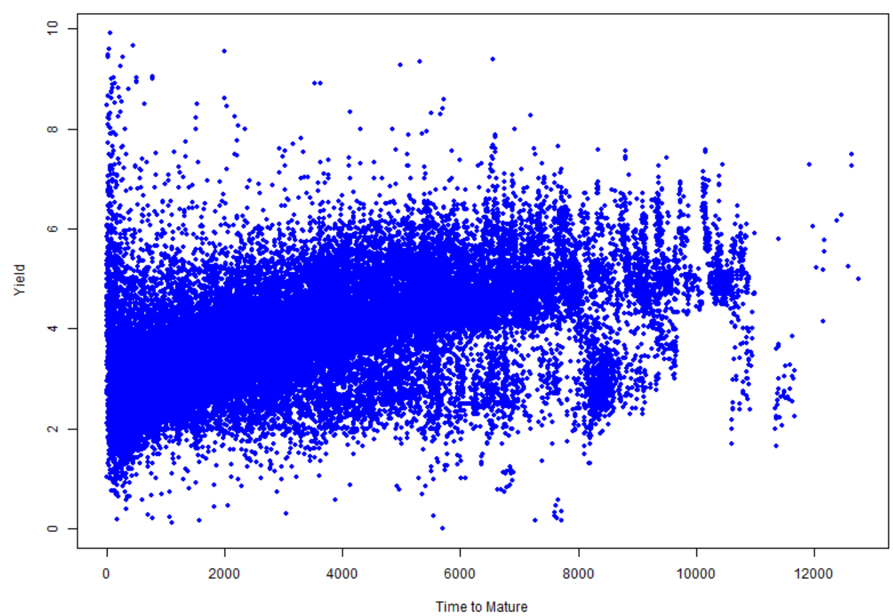

Figure 1.5: Muni Spreads Over Time

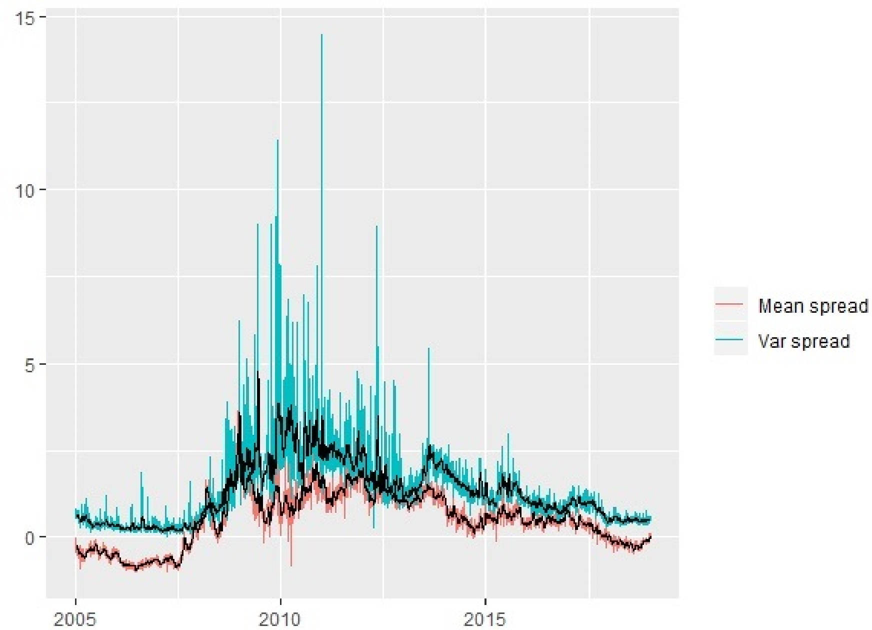

whomever is in possession of bond A at the time. In order to argue that yields affect borrowing costs, one must know that the yield of bond A has an effect on the cost of issuing bond B in the future. In this section I argue that secondary market yields not only affect borrowing costs, but the real behavior of state and local governments.

First, consider the association between the secondary market and the "primary" market for municipal bonds. I define a bond's "unexplained yield" as the residual ξ_{it} on the regression

$$R_{it} = 1 + \beta_1 T_{it} + \beta_2 i_{it} + \xi_{it}, \tag{1.28}$$

where R_{it} is the yield on muni i on day t, T_{it} is its time to maturity, and i_{it} is a benchmark treasury rate. This simple distinction yields a high correlation between

Figure 1.6: Muni Spreads, by Illiquidity

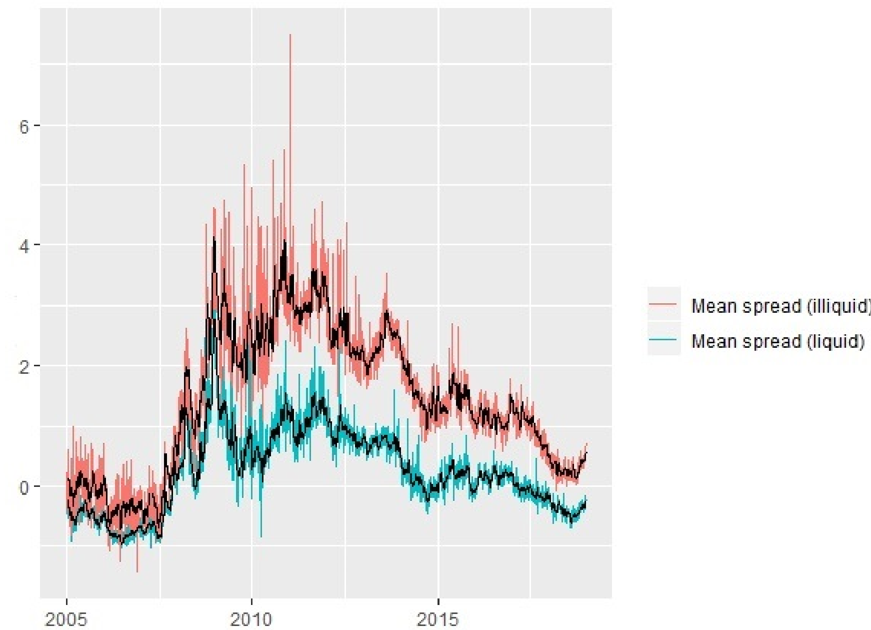

an issuer's (defined as a 6-digit CUSIP code) unexplained yields on the debt issued before day t and debt issued on date t: 0.68. This is a strong correlation, suggesting a strong relationship between the primary and secondary markets.

Secondary markets may be strongly related to primary markets and borrowing costs, but do they actually *affect* state and local government behavior? To explore this question, I connect governments with revenues more than $50 million from the Census Bureau's annual survey of state and local governments with CUSIP-6 issuer codes in the MSRB data. I use this dataset to estimate the regression equation

$$\log G_{it} = 1 + \beta R_{it} + \Gamma X_{it} + \varepsilon_{it}, \qquad (1.29)$$

Table 1.2: Government Responses to Secondary Market Muni Yields (IV)

	log(Debt Issues)	log(Current Exp)	log(Capital Exp)
Yield ($100bp = 1$)	-5.814	1.122	-0.6605
	(2.476)	(0.4715)	(0.4342)

Note: An observation is a municipality-year pair. The sample includes all municipalities in the Census Bureau's Annual Survey of State and Local Government Finances for which average revenues are greater than $500,000 and bonds could be found on the Bloomberg database from 2005 to 2012. Control variables include GDP, municipal revenues, and treasury rates. The explanatory variable is instrumented using summed monetary shocks as described below.

where G_{it} represents a few categories of government spending in a year, while R_{it} is the average yield of that government's debt on the secondary market and X_{it} is a vector of controls, including the average treasury yield. These yields are instrumented with an annualized version of the monetary shocks used later in the paper. Results are summarized in Table 1.2.

A couple of key suggestive results emerge from this exercise. The first, and most striking, is the apparent massive effect of secondary yields on new debt issues. An decrease of average annual yields on a government's debt of 100 basis points results in a *sixfold* increase in its new debt issues. Municipalities seem to respond in powerful ways to borrowing costs. Additionally, note that higher borrowing costs seem to shift the *composition* of municipal spending away from debt-financed capital projects to current expenditures. The secondary market for municipal debt clearly influences states and localities, both in terms of debt issuance and spending composition.[8]

1.3.2.1 Monetary Shocks

For the monetary shocks used in the exercise of this section, I use the strategy of Bu, Rogers, and Wu (2019) to identify monetary shocks on FOMC announcement dates.

[8]For aggregate time series evidence, see Appendix A.3.1.

I give a brief overview of the identification strategy here, with a full presentation in Appendix A.2.2. The BRW method uses the movements of prices of zero-coupon U.S. Treasury bonds with maturities $i \in \{1, 2, ..., 30\}$ on FOMC announcement dates to back out the implied monetary shocks on each date. This is accomplished in a Fama-Macbeth-style (Fama and MacBeth, 1973) procedure, which starts by making a standardizing assumption that defines the monetary shock as having a one-to-one effect on the five-year Treasury yield:

$$\Delta R_t^5 = \alpha + m_t + \eta_t. \tag{1.30}$$

Here ΔR_t^5 is the one-day movement in the five-year treasury yield, α is a constant, m_t is the monetary shock at date t, and η_t catches all other factors affecting the yield. Armed with this equation, we can use the change in the five-year Treasury yield on FOMC dates as a proxy to back out the monetary shocks m_t.

First, we need to estimate the sequence of 30 time series equations

$$\Delta R_{it} = \alpha_i + \beta_i \Delta R_t^5 + \varepsilon_{it} \tag{1.31}$$

for each zero-coupon bond $i \in \{1, 2, ..., 30\}$.[9] Armed with estimates $\hat{\beta}_1, \hat{\beta}_2, ..., \hat{\beta}_{30}$, shocks can be backed out by regressing the daily yield changes on the set of estimated $\hat{\beta}_i$ on each FOMC announcement date. This requires estimating the following

[9] In practice, these are estimated with a heteroskedasticity-based IV approach a la Rigobon and Sack (2004). Details are found in Appendix A.2.2.

equation for each $t \in \{1, 2, ..., T\}$:

$$\Delta R_{it} = \alpha_t + \hat{m}_t \hat{\beta}_i + \varepsilon_{it}. \tag{1.32}$$

The set of estimated $\hat{m}_t$ from each of these estimations is taken to be the monetary shock series.

The Bu, Rogers, and Wu (2019) monetary shock series has a number of desirable characteristics. First, it relies completely on publicly available data; the treasury yield is taken from the Federal Reserve website, and the estimated zero-coupon yields as estimated by Gürkaynak, Sack, and Wright (2006) are found at https://www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html. The publicly available nature of the data allows the shocks to be constructed at no cost, and the series is quite easy to update through the current date. Furthermore, the authors argue in the paper that this series is robust to the information critique of Nakamura and Steinsson (2018), but is nevertheless highly correlated with existing estimates of monetary policy shocks. Finally, these shocks are able to incorporate well the unconventional nature of monetary policy in the aftermath of the financial crisis; this is an especially important feature for this paper, as my sample only begins in 2005.

1.3.3 Time Series Evidence

The main set of results I use for the calibration of the model are time series estimates of the response of municipal yields to monetary shocks based on the yields of muni indices constructed by S&P.[10] The exercise is in the spirit of Rosa (2014), who studies the effect of monetary shocks on indices of AAA and AA bonds exclusively. I expand on Rosa's work by considering indices representing a broader set of munis, as well as specific geographic and sectoral indices, reminiscent of Gilchrist, Yue, and Zakrajšek (2019), who study the effects of U.S. interest rate shocks on the sovereign bond yields of several small open economies. While my results for highly rated bonds are similar to what Rosa finds, my other results provide a fuller picture of the effect of monetary shocks on the municipal bond market, especially with regard to potential heterogeneity.

S&P constructs these indices using a broad selection of bonds issued by state, local, and regional government entities in the U.S., which are not subject to income taxes. The bonds must have been issued in 2010 or later, and must have a minimum of 2 million U.S. dollars par value on the market. The indices are constructed as value weighted averages of the constituent bonds.[11]

For each of the indices in question, I am interested in estimating the equation

$$\Delta y_t = \beta_0 + \beta_1 m_t + \varepsilon_t, \tag{1.33}$$

where m_t is the Bu, Rogers, and Wu (2019) monetary shock series and $\Delta y_t = y_{t+1} - y_{t-1}$ is the two-day change in the yield to maturity of the asset around the FOMC

[10] All of these indices are available for download at https://www.spglobal.com/spdji/en/index-family/fixed-income/.

[11] More information is available from S&P for each index; for example, the methodology for the baseline index is found at https://www.spglobal.com/spdji/en/documents/methodologies/methodology-sp-municipal-bond.pdf.

Table 1.3: Baseline Time Series Results

	All GO	State GO	Local GO	S&P 500	All GO	State GO	Local GO	S&P 500
Monetary shock	0.22	0.24	0.22	0.50	0.23	0.23	0.28	0.65
	(0.08)	(0.08)	(0.08)	(0.14)	(0.14)	(0.16)	(0.13)	(0.21)
Horizon	2 days	2 days	2 days	2 days	6 days	6 days	6 days	6 days
N	2147	2147	2147	2147	2139	2139	2139	2139

Note: An observation corresponds to one day, around which a window is constructed from the previous day's price and the price at a given horizon. Each column refers to a separate time series regression of an index on monetary shocks. Heteroskedasticity-robust standard errors are reported in parentheses.

meeting date.[12] The coefficient $\hat{\beta}_1$ reflects the number of basis points the muni index yield should increase for every 1bp monetary shock, and is estimated with heteroskedasticity-robust standard errors.

1.3.3.1 Baseline Results

This section summarizes the average response of municipal bond yields to monetary shocks. Table 1.3 gives the estimated coefficients of monetary shocks on yields for three indices of general obligation municipal bonds. These indices group all municipals, state governments, and local governments, respectively, with an index for bonds from S&P 500 firms for comparison[13]. I choose to focus on general obligation bonds, which are backed by the full taxing power of the issuer rather than a specific revenue source, in order to more closely match the budget situation of the government in the model. These bonds can be reasonably thought of as representing the borrowing cost situations of their local governments.

[12] I follow GYZ in choosing a 2 day window; their alternative window of 6 days produces similar results.

[13] Suppose that the tax free nature of municipals were the only difference between munis and corporate bonds. In this case, the yield on municipal bonds y^m would simply be the after-tax return of corporate bonds, $(1-\tau)y^c$. The link between munis and treasuries, then, would be the coefficient on corporates discounted by the tax rate, $\frac{\partial y^m}{\partial r} = (1-\tau)\frac{\partial y^c}{\partial r}$. To the extent the coefficient on munis is lower than this, I say that the response of munis is *dampened* relative to corporates, for reasons other than their tax-free nature.

This exercise closely mirrors that of Rosa (2014), though my results exhibit slightly higher responses to monetary shocks than his; the reasons for this will be explained in part when I break these bonds out by rating. Nevertheless, the coefficients are quite dampened relative to models in which governments can borrow at the risk-free rate: municipal bond yields only increase (decrease) by 22 basis points in response to a 100 point change in the risk free rate, which is less than half of the response of corporate bonds, and far less than treasuries. This dampened response cannot be fully explained by the tax-free nature of municipals, and must be composed of illiquidity and/or risk effects. Models that do not take this dampened response into account will tend to overestimate the effects of monetary policy on local fiscal policy.

Interestingly, there does not seem to be much difference, on average, between the responses of state bonds and local bonds to monetary shocks.[14] This result is somewhat surprising, given the quite different tax and spending obligations between these two types of governments. Instead, it seems heterogeneity shows up in other ways, which I show in the next sections.

1.3.3.2 Heterogeneity by State

A natural place to look for heterogeneity across localities is in the presence of *geographic* variation. For this section, I estimate Equation 1.33 separately for indices of GO bonds originating in U.S. states, for which these indices exist.[15] Figure 1.7

[14] Additionally, in Appendix A.3.2 I investigate some alternative specifications. There is no significant difference between positive and negative monetary shocks, and including controls for the stock market does not affect the results.

[15] A number of states prohibit or limit the use of GO bonds.

Figure 1.7: Responses of Muni Yields to Monetary Shocks

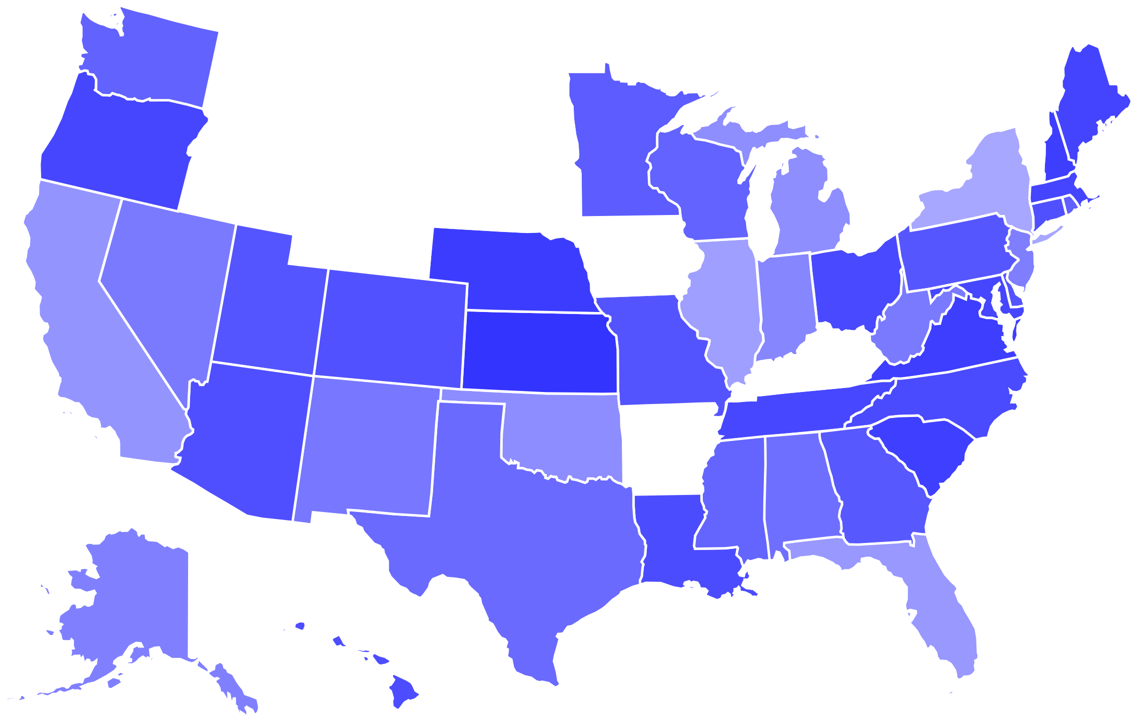

Note: Darker color indicates a stronger response of municipal yields issued in the state to monetary shocks. Missing states are those for which an index of GO bonds is not available. Maximum value: 39.81 (KS); minimum value: 17.32 (NY).

summarizes the estimates, highlighting visually the heterogeneity of responses across the U.S.

A few notes on these results are worth highlighting. First, the range of coefficients runs from 17.32 for New York to 39.81 for Kansas.[16] Second, no immediately obvious patterns emerge, save the apparent tendency for high-population areas to have muni bonds which response *more weakly* to monetary shocks; this tendency is suggested,

[16]We can reject that New York's response is 39.81 or higher, though the standard error bands for the two estimates do overlap. Even though sample size reduces power in this exercise, I interpret these results as evidence of heterogeneity across location given the remarkable consistency of coefficients for most of the sector-level indices and for state vs. local bonds.

Table 1.4: Time Series Results by S&P Rating

	AAA	AA	A	BBB Band	BB Band	NR
Monetary shock	0.087	0.062	0.021	0.194	0.218	0.372
	(0.078)	(0.077)	(0.090)	(0.099)	(0.674)	(0.170)

Note: An observation corresponds to one day, around which a window is constructed from the previous day's price and the price at a given horizon. Each column refers to a separate time series regression of an index on monetary shocks. Heteroskedasticity-robust standard errors are reported in parentheses.

though not proven by any means, later in Section 1.3.4.4. In any case, the results suggest that heterogeneity exists across space in the U.S. in the response of municipal bond yields to monetary shocks, and monetary policy may affect different areas of the U.S. differently.

1.3.3.3 Heterogeneity by Risk and Sector

In this section, I investigate heterogeneity in yield responses for muni indices broken out by rating and sector. Lower rated and unrated bonds seem to respond more strongly to monetary shocks than highly rated bonds, providing an explanation for why the magnitudes in this paper might differ somewhat from Rosa (2014). There does not seem to be much explanatory power in examining differences in municipal bonds broken out by sector.

Table 1.4 shows the coefficient estimates of Equation 1.33 run separately for indices of bonds in various S&P rating categories, at the 2-day horizon. Although some of the coefficient estimates are not statistically significant, there is a clear upward trajectory, i.e., the coefficients on riskier bonds are higher. This is consistent with a world in which expansionary monetary policy–for example–lowers default risk for risky bonds.[17] Furthermore, the coefficients on AAA and AA bonds are consistent with the findings of Rosa (2014), who only looks at low-risk indices. My baseline estimates, therefore, are higher than his in part because the GO index is not made up entirely of AAA and AA bonds.

Another way in which municipal bonds may differ from each other is the sector for which the bond was issued, especially for revenue bonds. Many local government entities may issue debt: schools, utility authorities, etc. While not directly entering a local government's general fund, these bonds do contribute to the overall burden of debt for local governments in a given places. Perhaps surprisingly, there is not much evidence of heterogeneity in the responses of these indices by sector, whose results can be found in Appendix A.3.3. Most of the coefficients, with a small number of exceptions, are close to and slightly lower than the baseline estimate of around 0.22.

Ultimately, this section has documented the presence of heterogeneity in municipal bond responses to monetary shocks, which are dampened on average relative to U.S. treasuries and corporate bonds. The type of government does not seem to make much difference, though it is possible that smaller and riskier governments have greater responses to monetary shocks. In the next section, I move from time series to panel data in order to obtain more evidence on the potential sources of this heterogeneity, including illiquidity, which is suggested by the consistently low values of these coefficients.

[17] It is worth noting that the majority of munis fall in the investment-grade category.

1.3.4 Panel Evidence: MSRB Data

Clearly, monetary shocks have a dampened effect on municipal bonds, with a fair amount of heterogeneity across different types of bonds. There is some evidence that risk may be driving some of this, but time series evidence cannot say whether liquidity and transaction costs may also be a contributing factor. Additionally, there may be information about specific local governments that influence the responses of munis to monetary shocks. Knowing that there is heterogeneity in these responses is sufficient for the quantitative exercise below, but knowing the *reasons* for this heterogeneity may be necessary to begin drawing out policy implications from the model.

To investigate more fully the sources of heterogeneity in monetary passthrough, I use a transaction-level dataset of municipal bond trades from the Municipal Securities Rulemaking Board, hereafter MSRB. These data are available through Wharton Research Data Services, and are available from 2005 onward. I end the sample on December 31, 2019 to avoid entanglements with the tumultuous nature of the muni market in early 2020. I restrict the sample to general obligation (GO) bonds issued by general governments (as defined by Bloomberg). Appendix A.2.1 provides more details on the dataset construction.

In this section–with the exception of the Schwert procedure–I estimate the equation

$$\Delta y_{it} = \beta_0 + \beta_1 m_t + \Gamma_0 X_{it} + \Gamma_1 X_{it} m_t + \varepsilon_{it}, \tag{1.34}$$

where m_t is the same monetary shock as before and X_{it} reflect bond-specific char-

Table 1.5: Baseline Panel Estimates

	Yield	Yield	Spread	Spread
Monetary shock	0.44	0.63	-0.01	0.16
	(0.27)	(0.28)	(0.17)	(0.24)
N	22758	22758	22699	22699
Time to Maturity Controls?	N	Y	N	Y

Note: An observation corresponds to an FOMC date-bond pair. Each column refers to a separate regression. Standard errors are reported in parentheses, and are clustered at the date level.

acteristics that might influence a bond's response to monetary shocks. Here, I allow a longer adjustment period for yields (and spreads) y_{it}, owing to the sparse nature of municipal bond trades. My baseline time period of adjustment is two weeks[18]; furthermore, I assign a bond's yield as its most recent daily yield, provided the trade happened within the last week. Standard errors are clustered at the time level, owing to the grouped nature of the shock m_t.

1.3.4.1 Average Results

Before moving on to other dimensions of heterogeneity, it may be helpful to benchmark baseline estimates in the panel data. In Table 1.5, I present coefficient estimates of β_1 from Equation 1.34. These estimates represent four regressions of monetary shocks on muni yields and spreads, and varying the inclusion of controls for time to maturity.

Of particular note here is the coefficients on the response of muni yields to monetary shocks. Why are these estimates higher than the baseline estimate in the time series section? The main reason is the selection inherent in this exercise: in order to

[18]Anderson and Cesa-Bianchi (2020) use one week for corporate bonds, and I double this window to allow for more trading to occur.

properly impute a price to bonds in this data, I only assign prices for trades within the last week. Trading rates of municipal bonds are quite low, so *conditional on trade*, we should expect individual-level responses to be higher. If we include those bonds which trade before the shock but not after, coding them as $\Delta y_{it} = 0$, the estimates (on yields) are quite close to the baseline 0.22 from the time series section. In any case, these estimates, should be kept in mind as a baseline as we move through the rest of the section.

1.3.4.2 GYZ Method

In addition to time series evidence on the effect of monetary policy on sovereign bond yields, Gilchrist, Yue, and Zakrajšek (2019) also perform an experiment to evaluate the effects of risk and liquidity on these responses. Specifically, in the context of Equation 1.34, X_{it} includes an indicator for whether a bond is investment grade or not (S&P rating BBB- or above), as well as a series of basic characteristics that the authors argue may *influence* liquidity: par value $\log PAR_i$, age $\log(1 + AGE_{it})$, time to maturity $\log T2M_{it}$, and coupon $\log(1 + COUP_{it})$. Furthermore, the response variable is the change in the muni spread rather than the yield.

Table 1.6 reports results from this estimation exercise, along with a joint test for the significance of the liquidity variables together. While none of the individual interactions are significant, the interactions of the four liquidity variables are significant for determining yields, suggesting a role for bond characteristics in the responses of borrowing costs. Finally, while the coefficient on risk is noisy–in contrast to the GYZ results for sovereign bonds–the sign is consistent with theory and time series

Table 1.6: Heterogeneous Responses: GYZ Method

	Yield	Spread
Monetary Shock	1.48	1.11
	(1.01)	(1.06)
Investment Grade = 1	0.01	0.01
	(0.02)	(0.02)
$\log PAR_i$	0.03	0.03
	(0.00)	(0.00)
$\log(1 + AGE_{it})$	-0.00	0.00
	(0.00)	(0.00)
$\log T2M_{it}$	-0.00	-0.02
	(0.01)	(0.01)
$\log(1 + COUP_{it})$	-0.00	0.00
	(0.03)	(0.03)
Monetary Shock * Investment Grade = 1	-0.14	-0.16
	(0.21)	(0.15)
Monetary Shock * $\log PAR_i$	-0.02	-0.02
	(0.04)	(0.04)
Monetary Shock * $\log(1 + AGE_{it})$	-0.01	-0.01
	(0.03)	(0.02)
Monetary Shock * $\log T2M_{it}$	-0.14	-0.11
	(0.11)	(0.12)
Monetary Shock * $\log(1 + COUP_{it})$	0.21	0.05
	(0.31)	(0.29)
P-value, liquidity interactions	0.051	0.228
N	22758	22699

Note: An observation corresponds to an FOMC date-bond pair. Each column refers to a separate regression. Standard errors are reported in parentheses, and are clustered at the date level.

evidence, in which less risky bonds exhibit lower responses to monetary shocks.

1.3.4.3 Schwert Spread Components

While GYK deals with sovereign bonds, Schwert (2017) is a paper at the frontier of the municipal bond pricing literature. The main exercise in the paper exploits the microstructure of the MSRB data to examine the portions of tax adjusted muni spreads that are accounted for by risk and illiquidity concerns. The basic procedure assumes that yields on municipal bond trades are determined according to

$$y_{it} = (1 - \tau)(r_t + \gamma_{it} + \psi_{it}),$$

where τ is the marginal tax rate, r_t is the risk-free rate, γ_{it} is the risk premium, and ψ_{it} reflects illiquidity, perhaps in the form of trading costs or asymmetric information.

To estimate the liquidity component ψ_{it}, the method first constructs λ_{it}, an average of several (standardized) illiquidity measurements from the literature. I describe these illiquidity measures in more detail in Appendix A.2.3; I follow Schwert closely, dropping a measure that requires more observations in order to extend to a larger sample of municipal bonds from smaller governments. The following equation for (tax-adjusted) spreads is then estimated at each time t:

$$\frac{y_{it}}{1 - \tau_{it}} - r_t = \beta_0 + \beta_t \lambda_{it} + \beta_t^R Rating_{it} + \varepsilon_{it}. \tag{1.35}$$

Here, y_{it} is the daily yield of a bond, τ_{it} is an imputed tax rate,[19] r_t is the zero coupon

[19] In my estimation, this is the same for all bonds, since I want to use this procedure on the full

U.S. treasury rate of similar maturity, and $Rating_{it}$ is a factor variables describing the S&P bond rating, if one exists. Armed with a series of betas on each day, the series of liquidity spread components is computed according to

$$\psi_{it} = \beta_t(\lambda_{it} - \lambda_{1t}), \tag{1.36}$$

where λ_{1t} is the first percentile of the liquidity measure. The risk component, γ_{it}, is simply computed as the portion of the tax-adjusted spread unexplained by the liquidity component.

Because the individual measures of these components are noisy, Schwert aggregates them into time series variables, using the four-month rolling average of daily cross-sectional mean spread components. This results in the time series γ_t and ψ_t, which are plotted in Figure 1.8. Note that my estimates of the relative magnitudes of these components differ substantially from Schwert's, which put the majority of the weight on risk; this is because I am using a more extensive sample of municipal bonds, whereas he uses only the largest state and local governments. This suggests a difference in spread makeup between smaller and larger state governments, and merits investigation in future research.

I estimate the effects of monetary shocks on these series, as in the time series results above, and present results in Table 1.7. Not much of significance stands out here, although there is some weak evidence of a transitory effect of monetary shocks on the liquidity component of spreads. Overall, it does not seem, to the extent monetary policy affects borrowing costs in a heterogeneous way, that the effect is

range of the data and cannot match all bonds to a geographic area.

Figure 1.8: Muni Spread Components

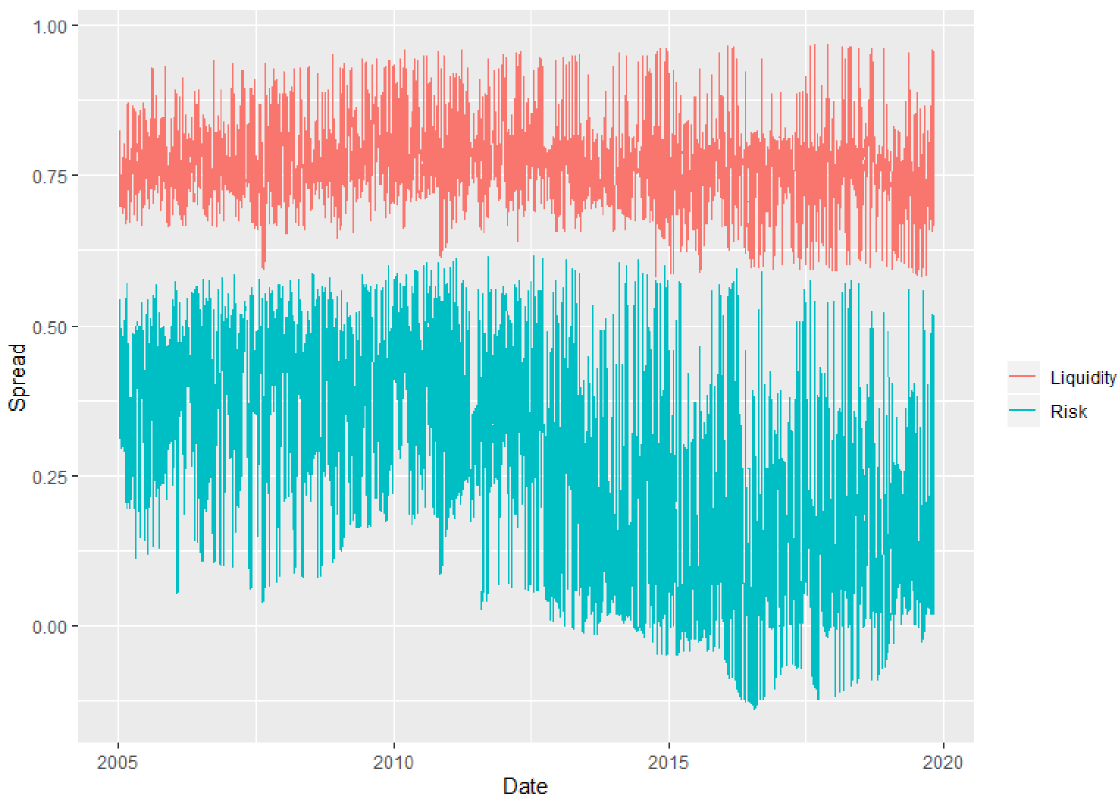

Table 1.7: Effect of Monetary Shocks on Spread Components

	Default spread	Default spread	Liquidity spread	Liquidity spread
Monetary shock	-0.11	-0.10	0.10	0.02
	(0.27)	(0.17)	(0.05)	(0.08)
N	2952	2944	3360	3352
Horizon	2 days	6 days	2 days	6 days

Note: An observation corresponds to one day. Each column refers to a separate regression. Heteroskedasticity-robust standard errors are reported in parentheses.

working through altering the components of risk and liquidity on spreads.

1.3.4.4 Government Finances Data

The last margin of heterogeneity I investigate involves a series of finance variables for municipalities. To obtain these data, I use the Census/Survey of Governments data from the U.S. Census. This survey obtains hundreds of balance sheet variables for state and local governments in the U.S., taking a representative sample annually and a full population sample every five years. Following Schwert (2017), I select the local governments in the U.S. with annual revenues of over $50 million. I then obtain the 6-digit CUSIP codes for these government issuers from the Bloomberg Terminal, and connect them to my panel dataset.

I estimate a series of regressions using these financial variables, and results are given in Tables A.4, A.5, and A.6 in Appendix A.3.4. None of the coefficients are precisely estimated. However, a clear pattern emerges: larger values seem to depress the response of bond yields to monetary policy. While not a smoking gun, this pattern is consistent with the state-level index coefficients on monetary shocks, which seem to imply that bonds from more populous states respond less to monetary shocks. This is a conjecture, however, and requires a more powerful identification strategy. Unfortunately, these finance statistics only vary at the annual level, and such data is not universally available for all muni issuers in all periods.

1.3.5 Empirical Takeaways

This empirical section has studied, from a number of angles, the effect of monetary shocks on municipal bond markets. On average, the yields on an index of general obligation municipal bonds respond by 22 basis points to a 100 basis point shock to the risk-free rate. This dampened effect is consistent with the lower volatility of municipal yields in general relative to treasuries. Furthermore, while the "level" of the issuer does not seem to matter, the location of the issuer does. In other words, I document heterogeneity across the U.S. in the response of municipal yields to monetary shocks. This heterogeneity maps to the model from Section 2.2.

I then investigate potential sources of this heterogeneity. There is some limited evidence that the heterogeneity may arise from differences in bond ratings or liquidity considerations, which map to a model of over-the-counter debt pricing. There may be an association as well between the size of a locality and its borrowing costs' response to monetary shocks, but the data on government finances is limited, especially for smaller governments. In the main quantitative results below, I do not take a stance on the source of heterogeneous responses, but the various options may carry different policy implications.

1.4 Quantitative Results

This section shows quantitatively the importance of the municipal public finance channel of monetary transmission; significantly, the estimated heterogeneity in muni yield responses to monetary shocks implies heterogeneous effects of monetary policy on localities, depending on local financial frictions. Increasing the response of local government borrowing costs from the lowest state-level estimates to the highest implies an increase of up to a half percent of household utility for an expansionary shock of 25bp. Furthermore, the nature of municipal bonds and frictions in muni markets significantly dampen this channel: allowing a municipal government to borrow at the aggregate risk-free rate more than doubles the output response to a monetary shock. Finally, although expansionary monetary policy is stimulative in early periods, the increase in municipal debt finance becomes a drag on local economies far into the future.

In this section, I focus on the response of single small open economies, abstracting from a full model of the U.S. I take this approach for ease of exposition. In Appendix A.4.1, I show that the results on heterogeneity of borrowing costs indeed carry over to a model in which there are multiple localities at one time. The only difference from the perspective of the locality is the process of the risk-free rate; when I consider the "partial equilibrium" with only one locality, I take the vector $[y_t^T \; r_t^*]$ as an exogeneous process, as is common in the international literature. I specify its process below.

1.4.1 Calibration and Solution

Before moving on to the quantitative results, it is necessary to discuss specifics on the calibration of the model for proper interpretation, which in this case is a U.S. municipality as a small open economy. I first discuss functional form choices, then parameters which are taken from the literature or estimated. I also calibrate a set of parameters to match some average statistics on state and local spending, revenues, consumption, and debt, followed by an investigation into the effects of local openness on the local fiscal multiplier. Finally, I briefly discuss the solution method.

1.4.1.1 Functional Forms

In the body of the paper, I use the model outlined above, opting for an *ad hoc* version of financial markets, in the vein of models with external debt elastic interest rates in the open economy literature. I use a simpler formulation for simplicity and ease of mapping the empirical response coefficients. Furthermore, while the sources of heterogeneity in borrowing costs have been investigated in multiple ways, this formulation allows an abstraction from taking a stand on their relative magnitudes for the time being. To begin, assume the exogenous tradable endowment y^T reflects shocks in the aggregate economy. Then let the aggregate risk-free r_t^* be determined by the system

$$z_t = Bz_{t-1} + \varepsilon_t^z, \tag{1.37}$$

where $z_t = \left[\log \frac{y_t^T}{y^{T*}} \;\; \log \frac{1+r_t^*}{1+r^*}\right]'$, and y^{T*}, r^* are parameters reflecting steady state values. $\varepsilon_t^z = [\varepsilon_t^y \;\; m_t]'$ reflect the exogenous shocks, with m_t being our shock of interest. Household interest rates are given by a standard debt-elastic interest rate formulation,

$$r_t^H = r_t^* + \phi^H(\exp(d_t^H - \bar{d^H}) - 1). \tag{1.38}$$

The benefit of this formulation is the ability to set any arbitrary steady state debt level; in the baseline calibration, I make the representative local household a *saver*, i.e., $\bar{d^H} < 0$.[20] I set the government's borrowing costs in a similar fashion, but with a friction on the adjustment to the treasury rate:

$$r_t^G = r^* + \theta^G(r_t^* - r^*) + \phi^G(\exp(d_t^G - \bar{d^G}) - 1). \tag{1.39}$$

First, note the imperfect response of actual borrowing costs to the risk-free rate, governed by θ^G, which will be the key parameter of interest capturing the response of muni yields to aggregate interest rate shocks. These pricing functions will be used in the main results of this paper, and similar results for the OTC model can be found in Appendix A.4.2.

Utility over consumption in this model is CRRA, with log utility over leisure and public goods consumption:

$$U(c) = \frac{c^{1-\sigma} - 1}{1 - \sigma} \tag{1.40}$$

$$V(h) = \Phi \log(\bar{h} - h) \tag{1.41}$$

$$W(g) = \gamma \log(g). \tag{1.42}$$

Furthermore, the consumption aggregator is CES:

$$A(c^T, c^N) = \left(A(c^T)^{1-\frac{1}{\xi}} + (1 - A)(c^N)^{1-\frac{1}{\xi}}\right)^{\frac{1}{1-\frac{1}{\xi}}}, \tag{1.43}$$

where ξ determines the substitution elasticity across tradable and non-tradable con-

[20] This imposition captures two features of the real world. First, it is more reasonable to assume that a household saves at the risk-free rate than that it borrows at this rate, and I wish to abstract from the market for household debt in this project. Second, it provides a "demand side" for the external financial market; while not strictly necessary here, it may be helpful for the reader to be able to think of having borrowers and savers in the model.

Table 1.8: Fixed Parameters

Parameter	Description	Strategy	Value
σ	CRRA utility parameter	Literature	2
α	Labor share of production	Data/Literature	0.6
μ	Elasticity of substitution in production	Literature (Gali and Monacelli, markup target 20 percent)	6
β	Discount rate	Literature, imply s.s. interest rate of 0.03	0.9694
θ	Calvo parameter	Data/Literature (target average 10 mos between price changes)	0.7
$\bar{h}$, h	Labor endowment and steady state	Literature, labor supply = 1/3	3, 1
Φ	Leisure utility	Set to solve $mr = mc$ in steady state	
γ	Government utility	Normalize to log utility	1
ξ	Elasticity of substitution	Literature, set to $1/\sigma$	1/2

sumption, and A represents the "openness" of the economy, which will be a key factor in the size of the local fiscal multiplier–more on this below.

1.4.1.2 Parameters

The parameters in the model are set using a combination of data and literature. Table 1.8 gives the basic parameters from the model which are standard in the literature. Nothing of extreme note is here, other than to note the decreasing returns to scale in intermediate production, which are consistent with the labor share of production in the U.S.

The exogeneous process for y_t^T and r_t^* is defined by the coefficient matrix B, which amounts to a VAR process with a lag of one period. For the baseline results, I simply estimate the matrix B as a bivariate VAR on the series $[log(Y_t)\ log(r_t)]'$, where Y_t is real GDP and r_t is a treasury rate, in this case, the U.S. 10-year. This estimation results in the (quarterly) coefficient matrix

$$\hat{B} = \begin{bmatrix} 0.985 & -0.0004 \\ 0.012 & 0.96 \end{bmatrix}.$$

Table 1.9: Calibration Targets

Target	Value
State and local government consumption and investment / GDP	0.11
State and local government own revenues / GDP	0.09
State and local government debt / GDP	0.15
Household savings / GDP	0.05
Imports / total shipments, from CFS ($>$ 50 miles)	0.66
Municipal bond yield	2.75
Local government spending multiplier (Chodorow-Reich)	1.8

The baseline elasticity of municipal borrowing costs to the exogenous component of the aggregate interest rate is taken directly from the main time series result for all general obligation bonds in the empirical section, $\theta^G = 0.22$. In the section studying the effects of heterogeneity in this elasticity, I use the distribution of the state-level estimates as plausible *lower bounds* on the heterogeneity, since they themselves represent averages to some extent. The low and high end of these estimates are used to define "low" and "high" elasticities, respectively.

The remaining six parameters,[21] y^{T*}, A, τ^G, $\bar{d^H}$, $\bar{d^G}$, ϕ^G and κ are calibrated to a set of moments that represent averages for state and local governments in the U.S. economy, in addition to a selected point estimate from the literature on local spending multipliers. These targets and their values–approximated for simplicity–are given in Table 1.9. At this point, a discussion is necessary on the exact interpretation of this small open economy. Is it a state government or a local government, or something else? The issue, of course, is that households in the U.S. are under the jurisdiction of

[21]I also set $\phi^H = 1$. This matches fairly well the persistence of the household debt response to a monetary shock in the Christiano, Eichenbaum, and Evans (1996) style procedure mentioned previously, as well as the magnitude of the economy's output response.

multiple tiers of governments, each of which exerts its own sphere of responsibility. Why are state *and* local expenditures being used for calibration? A robust literature exists describing the determination of public policy in federalist systems, but it is not the goal of this paper to enter in to that exciting discussion. Suffice it to say, for now, that the "government" imagined in this model is some sufficiently small combination of government roles that can be thought to be representative of its constituents' value functions. In the baseline calibration, I set a tradable consumption to total consumption ratio of 0.66, matching the proportion of shipments in the Commodity Flow Survey that travel further than 50 miles. This yields a value for A of 0.2683.

The degree of openness, given in this paper by the parameter A, can also be thought of as defining the "size" of the locality in question. As the locality's area increases, a higher proportion of household consumption is produced within the locality; for example, a good produced in San Francisco but consumed in Oakland is considered an import if the locality is defined by the Oakland city limits, but as a domestic good if we define the locality as the state of California or the more nebulous "Bay Area." To see how the choice of openness affects the multiplier, Figure 1.9 plots the on-impact government multiplier as a function of A, holding the rest of the calibration constant.[22]

In this figure, the vertical line represents the baseline calibration, which is chosen to match the preferred local government spending multiplier from the literature survey in Chodorow-Reich (2019), which puts forward 1.8 as a good value. Clearly, the impact multiplier depends crucially on the definition of a locality, or its openness.

[22] ϕ is dependent on the choice of A, for market clearing.

Figure 1.9: Local Government Multiplier and Openness

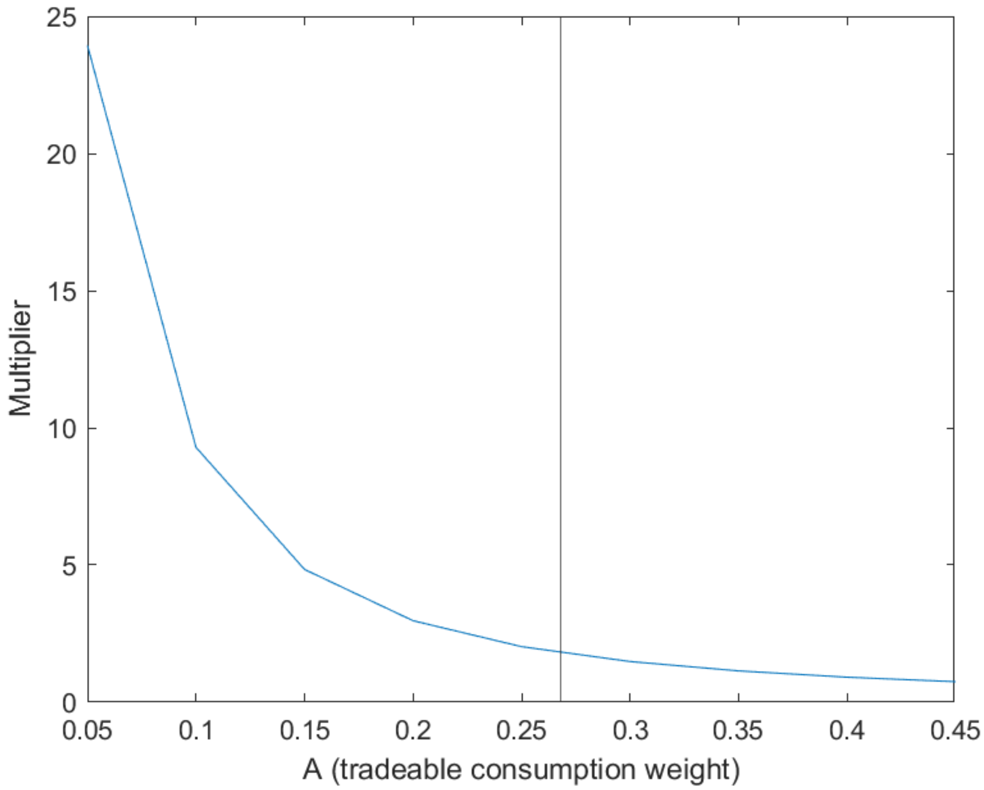

Alternative definitions of a locality, such as a state, in which tradable consumption is less important, will result in larger multipliers. As the economy becomes increasingly closed, the multiplier increases, as we approach the closed economy case. The multiplier,in addition to openness, also depends crucially on the proportion of non-Ricardian agents in the household, κ; these two determinants stand in agreement with the papers of Chodorow-Reich (2019) and Farhi and Werning (2017a), which analyze fiscal multipliers in monetary unions. A resulting implication is that if openness A decreases, fewer non-Ricardian agents will be required to match the preferred multiplier. Finally, this impact multiplier does not take into account the dynamic effects of local fiscal policy; the next section presents these effects in fuller detail.

1.4.1.3 Solution Method

For the results in this paper, I solve for the impulse response functions of the model by simulating the model's response to a one-time, unexpected shock to the exogenous portion of the interest rate. I assume that the economy is in steady state before the shock, and returns to steady state after 300 periods. While more computationally intensive than a perturbation strategy, this method allows me to later extend the model by including explicit constraints on debt issuance by the local government.

For the calibration of the hand-to-mouth share κ to match the local government spending multiplier of 1.8, I iterate over solutions of a stripped-down version of the model with exogenous government spending.[23] Because this calibration procedure

[23]In Appendix A.4.5 I further explore the effects of fiscal shocks in this model. I find that, for a shock to federal government spending, the effects on the local economy depend crucially on whether the federal government is purchasing local output or output from elsewhere. Federal spending on public goods crowds out local spending on public goods, and is only stimulative when that spending

requires potentially many evaluations of the response of the economy to an exogenous government spending shock, I compute these responses with a second-order perturbation in the `Dynare` package for `Matlab`. The quantitative responses from perturbation methods are quantitatively similar to those obtained from the more computationally intensive "MIT shock" method.

1.4.2 Results of an Interest Rate Shock

The response of this calibrated economy to a 25bp expansionary monetary shock is shown in Figures 1.10 and 1.11. The figure shows the percent deviations from the steady state in response to the monetary shock. Notice the logic of the channel shown in the second figure: borrowing costs decrease, increasing government debt and spending, resulting in output stimulus. In the baseline calibration, a 25bp decrease in the risk-free rate results in over a 3 tenths percent increase in output on impact.

Note, however, there is a long-run effects of the government debt buildup. Because the only margins of fiscal adjustment for the government are debt and spending–revenues are exogeneous–the government will have to reverse its debt accumulation through costly decreases in government spending later on. One immediate consequence of this result is the importance of thinking about the long-term consequences of stimulating debt-financed spending. Expansionary monetary policy allows local governments to shift spending from far in the future to the present, stimulating output in the short run but depressing it in the long run, even after interest rates have returned to steady state.

occurs locally.

Figure 1.10: IRFs, 25bp Expansionary Shock

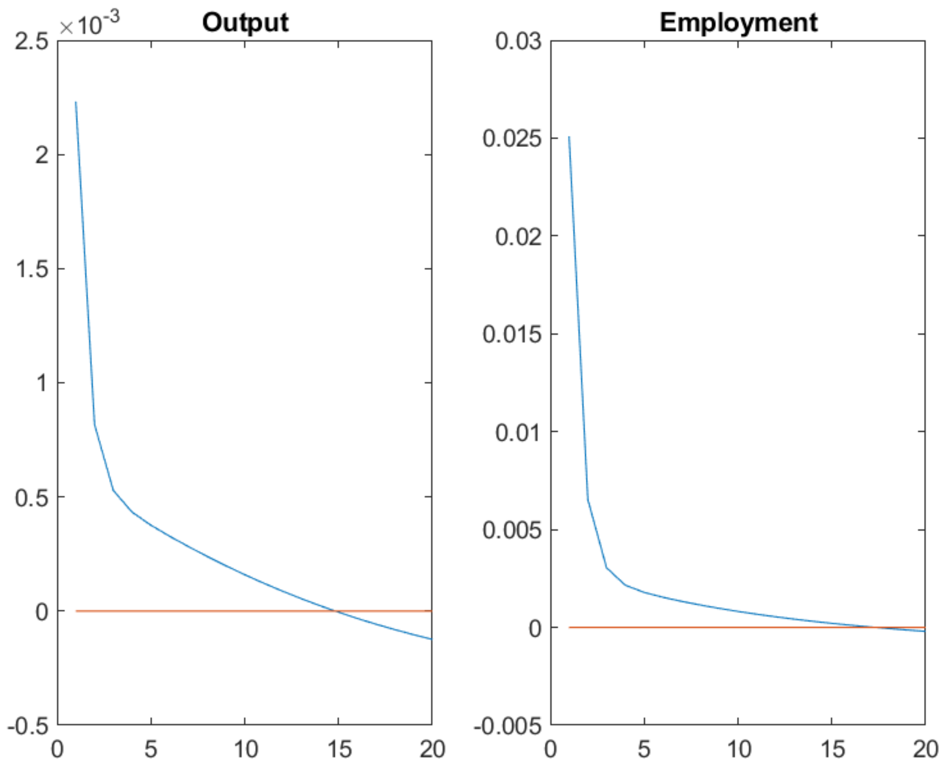

Figure 1.11: IRFs, 25bp Expansionary Shock

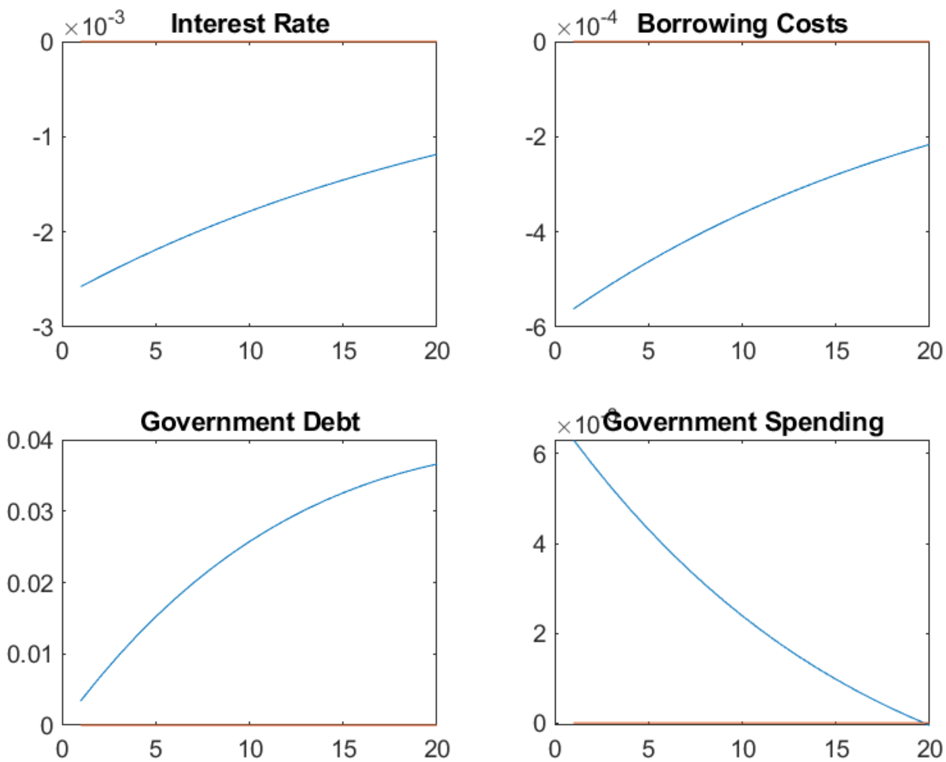

Figure 1.12: IRFs Compared to Standard Case

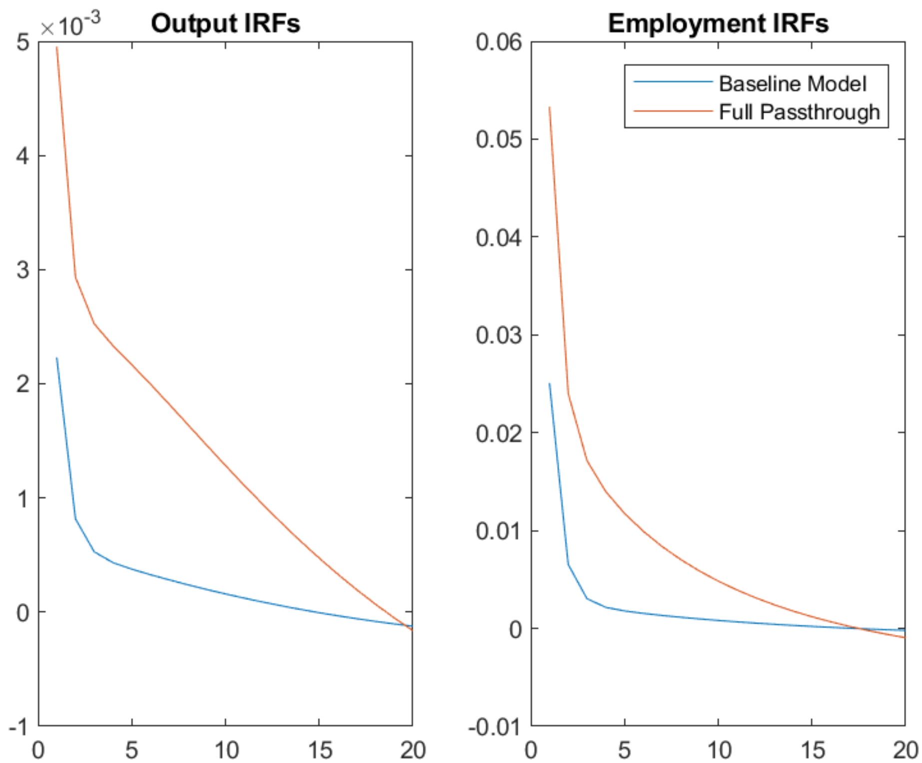

In the empirical section, I noted that the responses of municipal yields to monetary shocks were muted relative to what one might expect. As a result, the stimulative or depressive effects of monetary policy through local public finance are lower in a model which takes these features of muni markets into account, relative to a model which assumes local governments can borrow at the risk-free rate. To see the magnitude of the difference between the calibrated model and a model with borrowing costs at the risk-free rate, see Figure 1.12.

In the model that does not take the muted response of local borrowing costs to monetary shocks into account, the stimulative effects of monetary policy are more than double the model with realistic borrowing costs. When local government borrowing costs are tied more closely to the risk-free rate, local government debt increases by almost five percent relative to the steady state, and the reaction of spending, output, and labor are much higher than before. This stark difference suggests that models which assume municipal governments have access to borrowing at the aggregate risk-free rate will significantly overstate the stimulative effects of monetary policy on local economies; such models will also *understate* the possibility for heterogeneity of stimulus across localities.[24] I take up the extent of the heterogeneity in the next section.

In Figure A.5, I plot an estimated response of output to a monetary shock according to the Christiano, Eichenbaum, and Evans (1996) methodology mentioned earlier. The peak response of output to a one standard deviation monetary shock matches very well the magnitude of the peak response in the calibrated model, suggesting that the magnitudes of response in this model are reflective of the real world. Across the board, this type of "canonical" open economy DSGE model fails to generate the hump-shaped responses of economic variables to monetary shocks observed in VARs. One could imagine a number of features to supplement the model which might better match these hump shapes, but these features are not the focus of this paper.

[24]In Appendix A.4.3, I show that the inclusion of an explicit limit on debt issuances dampens the effect even further, resulting in anther potential source of heterogeneity, depending on the distribution of such limits in practice.

1.4.3 Heterogeneity Over Monetary Responses

A model that assumes a one-for-one relationship between municipal borrowing costs and national interest rates eliminates the possibility that borrowing costs might respond differently to monetary shocks in different localities. As a result, such a model will eliminate an important source of heterogeneity in the passthrough of monetary shocks across regions and localities in the U.S. In Figure 1.13, I plot the same impulse response functions as in the previous section for the baseline economy $\theta^G = 0.22$, and two additional economies: $\theta^G = 0.1$ and $\theta^G = 0.34$. These two additional economies are symmetric deviations from the baseline estimate of muni yield responses to monetary shocks, roughly corresponding to the low-end and and high-end of the state-level responses described in Section 1.3.3.2.

Realistic differences in the response of local government borrowing costs to monetary shocks have quantitatively important implications for heterogeneity in monetary transmission. For example, Francis, Owyang, and Sekhposyan (2012) estimate the responses of several U.S. cities to monetary shocks, grouping them according to region, and finding that the difference between the smallest regional peak employment response and the largest is about tenfold. In the model here, the difference is just under twofold. Furthermore, the dispersion of peak employment responses to monetary shocks in the full model of Appendix A.4.1 is about 9% of estimated dispersion in the responses obtained from a series of simple VARs[25]. These two comparisons

[25]The experiment here is the following. For each U.S. state, I estimate a VAR of state employment, national GDP, and the federal funds rate with a lag of four quarters, calculating the peak employment response to a one standard deviation shock to the federal funds rate for each state. I define dispersion of the responses as the standard deviation of percent deviations from the mean peak response. Dispersion in the model is 0.0746, versus 0.8421 in the data

Figure 1.13: IRFs by Response Elasticity

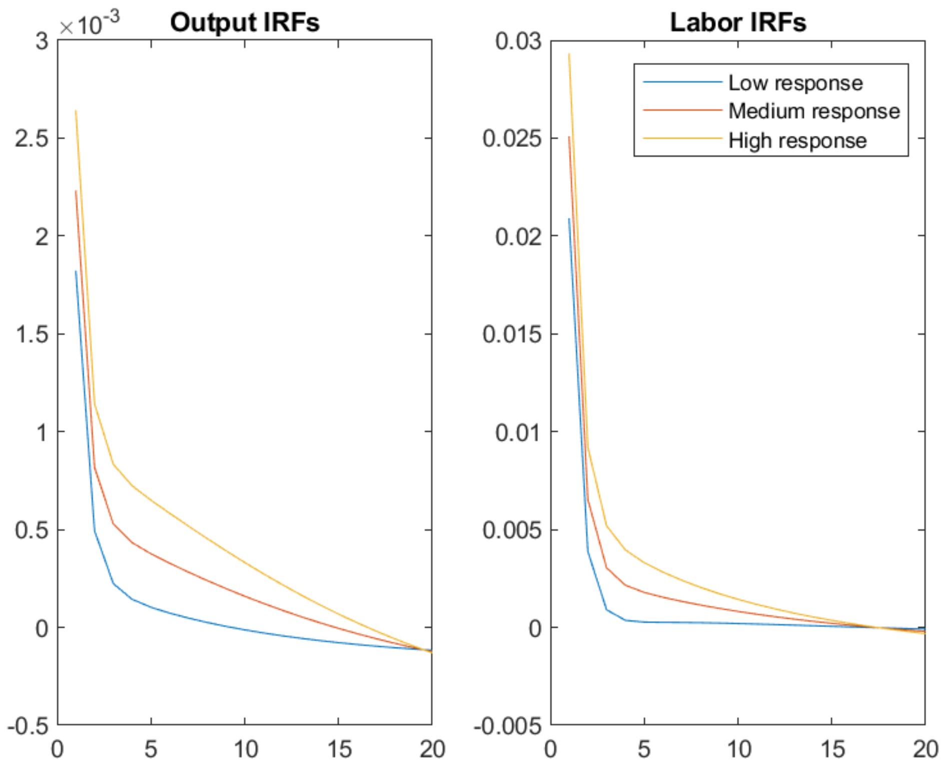

combined suggest that differences in borrowing cost responses to monetary shocks account for 10-20% of the variation in monetary transmission across U.S. states.

The municipal bond market is the only difference[26] between the economies in Figure 1.13, yet monetary transmission is markedly different between the economies. A government whose borrowing costs fall more after an expansionary shock borrows more, spends more, and sees greater output and labor increases on impact; the impact effects on output and labor are almost 50 percent greater in the high-response economy than the low-response economy. In the long run, of course, these effects will be flipped: the high-response governments have more debt to pay down, and thus a bigger future recession. The magnitude of response of municipal bond prices to monetary shocks is a key parameter, then, in determining both the *size* and *path* of local economic outcomes.

Why is this effect important? First, such differences in passthrough may affect the desirability of a given central bank policy in a monetary union. The results here suggest that, for example, the Federal Reserve should take into account heterogeneous effects on municipal governments across the U.S. when it is considering a given policy. Additionally, to the extent that a given policy can strengthen the relationship between treasuries and munis, that policy can increase the short-term output effects of monetary policy on the economy as a whole. Finally, since we have a calibrated model of a local government in a monetary union, in which the muni market plays an important role, we can examine the implications of this model for the two most

[26] Appendix A.4.4 gives one example of another difference between localities which might affect transmission: the steady state level of government spending. Unsurprisingly, the municipal spending channel of monetary policy is stronger when steady state government spending is higher.

recent U.S. crises, providing insights into local government behavior during each of them. It is to these crises that the next section now turns.

1.5 Application: Two Crises

In the previous sections, I have specified and calibrated an open-economy model of a local government. The open-economy model highlights the key role of financial markets in municipal government decisions over the business cycle. Such a model, then, is a good candidate for studying the effects of some recent economic crises on state and local governments.[27] First, I use the model to show that the effect of a recession combined with a financial crisis, as in 2008, dampens the ability of local governments to respond; this corresponds with the true observed behavior of state and local spending following the Great Recession. Finally, I am able to generate some predictions going forward from the current COVID-19 crisis.

1.5.1 State and Local Government Spending During the Financial Crisis

The Great Recession and its subsequent recovery were unique in myriad ways, and the behavior of state and local governments is no exception. For the three recessions

[27] It should be noted that the results in this section are meant to be more illustrative that quantitative. The model of the paper is built for more conventional study of monetary policy transmission. Specifically, abstracting from the default decision in the model is generally acceptable for steady state, given the extremely low rates of municipal default. Crises of unprecedented size, as these are, might render this abstraction unrealistic. Nevertheless, the experiments here help to illustrate some of the issues at hand; it may be helpful to think of the government as a state government, since these governments in general do not default on their debt.

Figure 1.14: State and Local Government Spending After Recessions

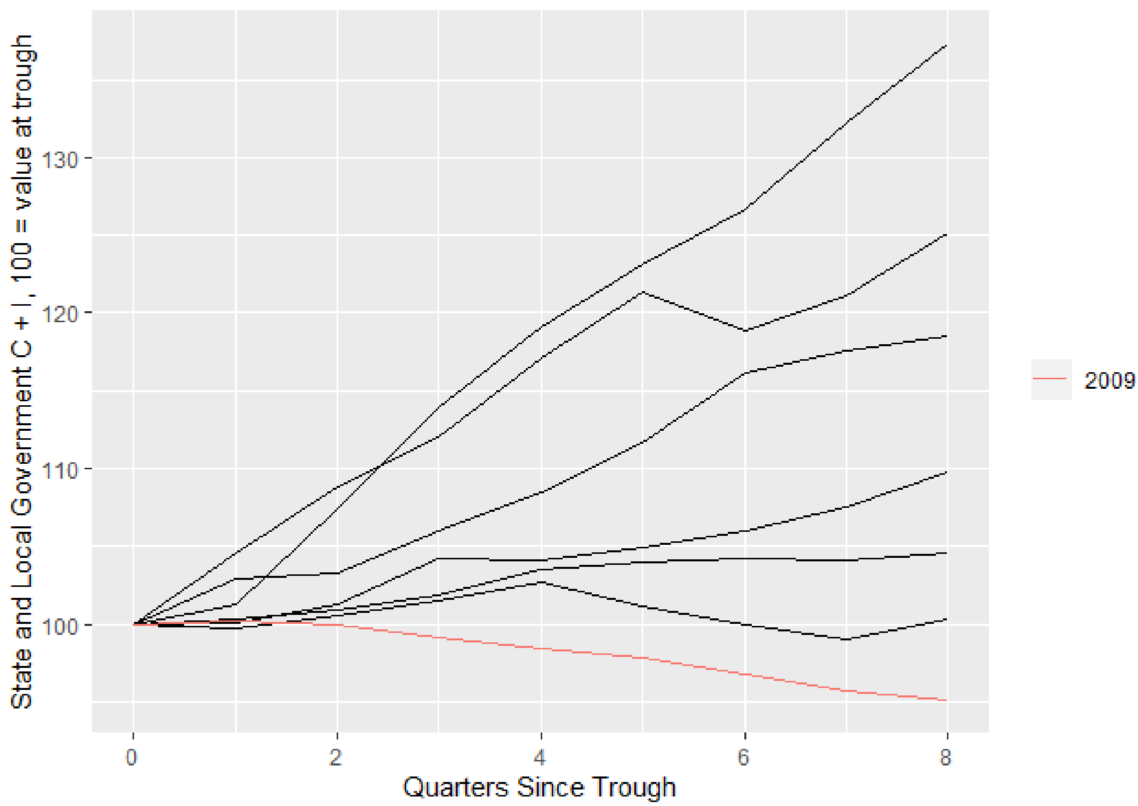

leading up to 2008, state and local government expenditures increased during the immediate recovery. In 2009, however, state and local governments *decreased* their spending, representing a break from previous recoveries. Figure 1.14 shows these recoveries.

One key difference in 2008, aside from the severity of the recession, was the associated crisis in financial markets. The model outlined in this paper allows us to examine the interaction between financial markets and state and local fiscal policy during recessions. Table 1.1 shows that, while U.S. treasury yields decreased quite dramatically during 2008-2014, municipal bond yields did not move much at all, especially at the high end of yields. This suggests either a decrease in liquidity in muni markets, or an increased perception of risk due to financial conditions.

Figure 1.15 plots the response of government spending in the open-economy model to two types of models. The first is a simple decrease in the tradable endowment y^T; in other words, an external crisis. The second combines an external crisis with a *financial* crisis; following the muni market during and after the Great Recession, I define this as a negative shock to θ^G, i.e., a dampening of the ability of municipal bond yields to decrease during the recession, and an increase in ϕ^G, meaning that inceasing debt becomes costlier.

Clearly, the financial crisis dampens the fiscal response of state and local governments to an external crisis. While a *decrease* is not induced for long, the financial crisis does cut out much of the government's fiscal response. A number of factors go in to the fiscal decisions of these governments in response to crises, including political considerations and budget rules, the non-response of borrowing costs during the Great Recession is likely an important factor in the lack of fiscal response by state and local governments.

1.5.2 Municipal Budgets and COVID-19

In March 2020, the COVID-19 pandemic caused widespread economic shutdowns in the U.S. Additionally, municipal bond markets went haywire, precipitating unprecedented action by the Federal Reserve to stabilize prices, including a mechanism by which the Fed would purchase munis, in addition to other assets. Despite the stabilization, it is widely thought that state and local governments will now be faced

Figure 1.15: Government Spending After Two Recessions

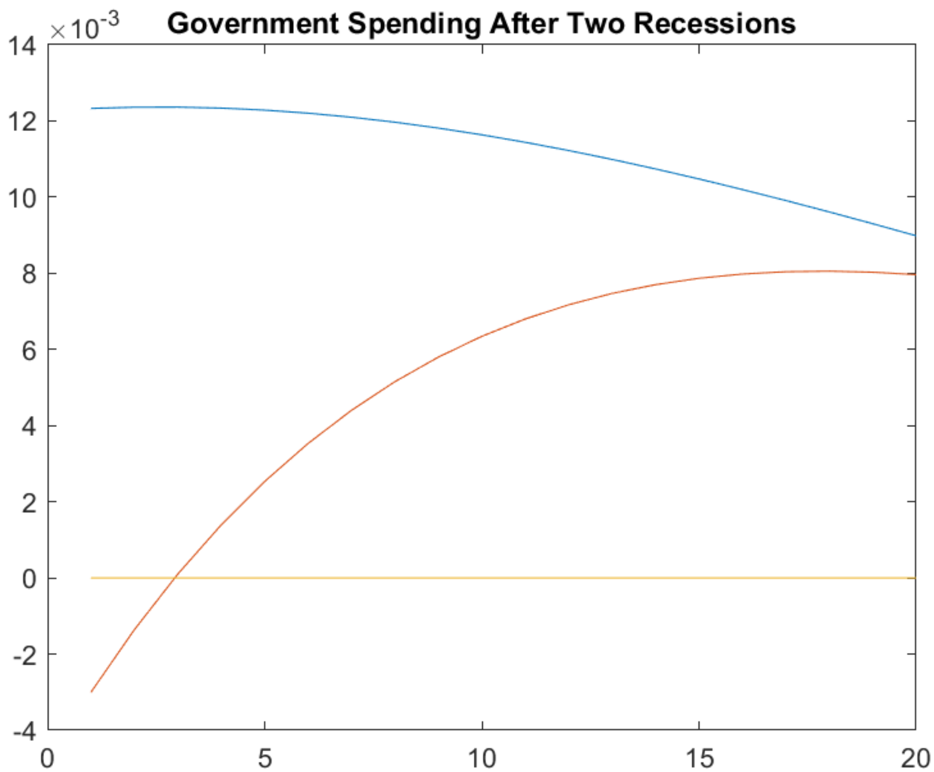

Figure 1.16: IRFs to a Pandemic Shock by Fed Response Time

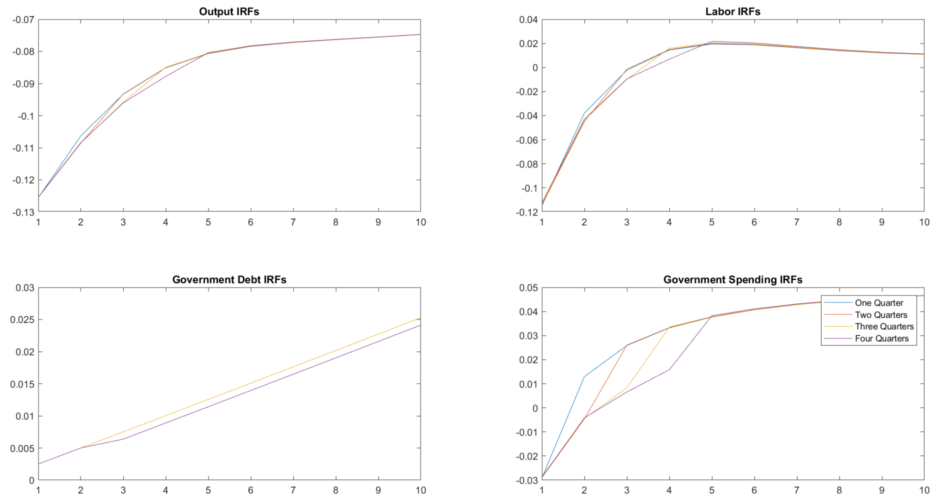

with many quarters of low revenues and difficult fiscal pressure. How does this crisis show up in the context of the model of this paper, and what might it tell us about the future path of state and local government spending?

While this model in many ways is not designed to capture the current pandemic crisis, I interpret the events thus far in three ways. First, as an external crisis like in Section 1.5.1, with a decrease in yT, reflecting the contraction in the entire U.S. Second, I shock $\bar{h}$, which in this model reflects an exogenous shock to unemployment, corresponding to the pandemic-induced lockdowns. I include a limit on debt increases to 2.5%, refelcting real debt issuance constraints faced by governments in times of crisis. Finally, I insert a transitory positive shock to borrowing costs, reflecting the chaos in the muni market which resulted in yields increasing by over 100bp at one point, relative to the baseline levels to which the Fed helped return them.

Figure 1.16 shows four hypothetical pandemic crises, each imposing a different persistence of the shock to borrowing costs for the local government. These can be thought of as hypothetical scenarios reflecting the speed of response by the Fed; in other words, what might have happened had the Fed delayed in its response to shore up liquidity in the municipal bond markets. In each case, the temporary financial market shock dampens government spending to a significant degree. Clearly, time of response has a moderate effect on the response of state and local governments, preventing governments from spending as much as they'd like to support households. These effects are transitory, however, as government spending and output quickly catch up to their preferred levels when interest rates return to "normal" levels. This simple experiment supports the rationale for the Fed's quick actions to prop up state and local government debt and spending in the early phases of the crisis.

1.6 Conclusion

This paper has provided a framework for understanding the passthrough of monetary policy to localities in the U.S. through state and local government spending. In an open economy model of a small U.S. region, the financial market underlying municipal borrowing costs affect the local government's ability to borrow and spend on fiscal policy in response to a change in the national risk-free rate. Municipal bond yields in the data exhibit dampened but heterogeneous responses to monetary shocks. These responses may be affected by liquidity in the over-the-counter municipal market, default risk perceptions, or some combination of the two.

Realistic heterogeneity in the response of municipal borrowing costs to monetary shocks implies differences of over 20 percent in output and employment responses to monetary shocks in the calibrated small open economy model. The financial market is important: failing to take into account the dampened response of municipal yields to monetary shocks would overstate the local stimulative effects of monetary policy by more than double. The importance of borrowing costs in determining local fiscal policy provides a playground which may give some insight into local fiscal policies in response to recent economic crises.

CHAPTER II

Optimal Minimum Wage Setting in a Federal System[1]

2.1 Introduction

After Kansas City, and St. Louis, Missouri set their own minimum wages above the state level in 2017, the Missouri state legislature prohibited any city from setting its own policy. While there is likely a political motive for a red state to overturn a progressive policy in its blue cities, there are also potential economic benefits from centralized redistribution. The central government can more efficiently implement such policies when workers are mobile by internalizing spillovers, and a state may not want one of its cities to set a higher minimum wage because of the externality imposed on the rest of the state. This may justify minimum wage preemption laws in Missouri and 24 other states. However, since the state is constrained to setting a

[1] Joint work with Andrew Simon.

uniform policy, as in Oates' (1972) decentralization theorem, then local policy setting may be preferred.

We examine optimal minimum wage setting in a federation with mobile workers to understand the relative benefits of centralized, decentralized, and joint policy setting with interregional spillovers. Which level of government should set minimum wage policies? Although our substantive focus is on the minimum wage, our framework has implications for a broad set of policies where local governments supplement federal decisions, like the EITC, income taxes, and Medicaid. Understanding the appropriate level of government decision making is increasingly important as the U.S. federal government has continued to shift more responsibility for these policies to states (Baicker, Clemens, and Singhal, 2012).

Our work most directly builds on Lee and Saez (2012), who examine the role of the minimum wage in the competitive labor market of a single jurisdiction, where workers are implicitly immobile. They find that a binding policy is desirable, even when non-linear taxes are available, if the newly unemployed have zero surplus from working and the government values redistribution to low-wage workers. In their model, workers can migrate between low and high wage jobs, though they face different costs from working in each. When a worker's sector is fixed so that labor supply can only respond on the extensive margin, the minimum wage is second-best Pareto inefficient.[2] With only one jurisdiction, their framework is unable to consider the relative merits of having different levels of government set policy as well as the ex-

[2] Gerritsen and Jacobs (2020) allow for educational investment in the Lee and Saez framework and find that a minimum wage is only optimal in the presence of non-linear taxes if the gains from more education outweigh higher unemployment.

ternalities that result from horizontal and vertical government competition. Vertical competition in a federation is particularly important for the minimum wage because it is a price floor policy; only the higher one matters. If the central government that is restricted to a uniform policy sets a binding minimum wage in the lower wage region, then it lowers the costs associated with horizontal competition.

We adapt the Lee and Saez (2012) model to a two-jurisdiction framework with mobile agents, regional governments, and a federal government to study the trade-off presented in the decentralization theorem. Workers are not mobile across sectors, but the high-skilled are mobile across jurisdictions.[3] Our goal is to understand the conditions under which setting a binding minimum wage is optimal for each type of government and analyze the welfare implications of different policy setting authorities. Local governments compete for high-skilled workers, which hinders their ability to redistribute through the minimum wage, but the central government is restricted to a uniform policy. Our stylized framework only includes two inputs for tractability, but the high-skilled can be thought of more broadly as mobile factors of production that may respond to increases in the minimum wage, like capital and firm location. In addition to the "tiered" U.S. model of a federal uniform price floor in which states may "top off," or raise the floor, we consider federal uniform, decentralized, and federal non-uniform policy setting.

We first provide theoretical results to understand the extensive margin decision of implementing a binding minimum wage when governments maximize social welfare.

[3] Blundell and MaCurdy (1999), and Lee and Saez (2012) both note that the extensive margin working decision where workers are not mobile across sectors is the most important and relevant in practice.

The local governments will set a binding policy if the welfare loss from unemployment is zero on the margin, and the government values redistribution to low-skilled workers more than emigration. Our conditions are the same as in Lee and Saez, except governments are concerned with policy induced migration. The federal government only cares about migration insofar as it affects total output and moving costs, leading it to have stronger preferences for a minimum wage than local jurisdictions.[4] Since the theoretical model only gives predictions on the extensive margin, we proceed to calibrate a two region model to the aggregate U.S. economy to illustrate how mobility and regional heterogeneity in productivity impact the relative benefits of centralization, decentralization, a hybrid system as in the U.S., and centralized non-uniform policy in general equilibrium.

The key insight from the model is that different levels of government are strategic complements in policymaking and the extent of this complementarity depends on input mobility and regional heterogeneity in productivity. Higher levels of mobility increase the costs associated with horizontal competition, while higher levels of regional heterogeneity decrease the effectiveness of uniform central policy. When jurisdictions are identical and inputs are mobile, centralized authority leads to greater social welfare while additionally allowing local governments to enact policy has no effect. However, when jurisdictions differ but inputs are immobile and therefore there are no interregional spillovers, decentralized authority is preferred. Simultaneously allowing the central government to also set a minimum wage in this case does not

[4] We also discuss joint policy setting by the local and federal government as well as federal non-uniform policies and several alternative models, such as rent-seeking governments, economy-wide aggregate production functions, and additional mobile factors of production, to understand how the model assumptions impact the sufficient conditions for a minimum wage.

change the equilibrium. For the more realistic cases with heterogeneous regions and imperfectly mobile inputs, a hybrid system improves welfare since central uniform policy reduces horizontal competition and decentralized policy allows for a different minimum wage in each jurisdiction.

Regional heterogeneity in redistributive preferences may also drive differences in the minimum wage, as reflected in the Missouri example. Urban areas with high minimum wages tend to be both more productive and more progressive. We find that progressivity has a nonlinear effect on optimal policy since the minimum wage redistributes to the low-skilled from the high-skilled and the newly unemployed. A government with a very progressive social welfare function would not implement a binding minimum wage because of the negative effects of additional unemployment, while a government at the other extreme would not implement a binding policy because it does not value redistribution. Local heterogeneity in preferences can also lead to higher minimum wages. If one local government enacts a minimum wage because of its redistributive preferences, this can lead the other local government to set a binding policy, since a higher minimum wage in one region decreases the migration externality in the other.

To understand the welfare implications of tiered and centralized policy setting, we calibrate an economic geography model of the continental United States to match regional heterogeneity in productivity, high-skilled location decisions, employment, and the federal government's optimal policy. Consistent with our earlier findings, this model predicts that tiered policy setting yields a small welfare gain over centralization alone, as it allows the federal government to more effectively redistribute from high-skilled to low-skilled workers. Even though states are heterogeneous in productivity, and high-skilled workers are fairly mobile, we also find that tiered minimum wage setting closely approximates the social planner's optimal policies. Similar models have been used to study place based policies in general (Kline and Moretti, 2014), corporate tax cuts (Suárez Serrato and Zidar, 2016), and income taxes (Colas and Hutchinson, 2020).

While we focus on the trade-offs of different minimum wage setting authorities, the previous normative literature focused on different tax systems. Stiglitz (1982), Allen (1987), and Guesnerie and Roberts (1987) find that the minimum wage is not optimal when non-linear taxes are available, however, Guesnerie and Roberts find they may be desirable under linear taxes. The literature has also considered more sophisticated minimum wage policies in a single jurisdiction where government competition is implicitly absent, like a graduated minimum wage tied to firm size (Danziger and Danziger, 2018), in-kind redistribution (Economides and Moutos, 2017), monitoring job search (Boadway and Cuff, 2001), and bargaining power (Hungerb uhler and Lehmann, 2009).[5] Previous work on optimal tax that incorporates government competition in a federation (e.g. Wilson, 1982; Gordon, 1983; Hamilton and Pestieau, 2005; Gordon and Cullen, 2012; Dvorkin, 2017) has not simultaneously allowed governments implement minimum wages. We connect these two previous literatures by understanding optimal minimum wage policies as a tool for redistribution in a federation.[6]

[5]We examine the implications in a competitive labor market, but other work has focused on the optimality when there are search frictions, as in search-and-matching model. See Flinn (2006), Dube, Lester, and Reich (2016), and Lavecchia (2019), for example.

[6]There is also a large empirical literature on the effects of the minimum wage. See Neumark

By putting the optimal minimum wage analysis into a federation, our work relates to the literature on tax competition. In our model, when a jurisdiction increases its minimum wage, high-skilled workers migrate to the other region, creating a horizontal externality that is increasing in mobility. Similar to models like Wilson (1986), and Zodrow and Mieszkowski (1986) on tax competition, agent mobility will lead to low local minimum wages. By adding a federal government to the Wilson-Zodrow-Mieszkowski set-up, Keen and Kotsogiannis (2002, Mar. 2004) find vertical externalities leave state taxes too high, which is in contrast with our result that federal policy decreases the cost of local minimum wage setting. Our results are therefore more similar to Janeba and Wilson (2011) who examine local public good provision with both horizontal and vertical externalities. In their model, local provision is too low due to competition, but central provision is inefficient because it is determined by a winning coalition in the legislature. In other political economy models, like Lockwood (2002) and Besley and Coate (2003), central provision is inefficient for similar reasons, whereas it is inefficient in our framework because the federal government is restricted to a uniform policy. Our assumption is based on the U.S. system, but this restriction does not need to hold. In May 2019, European Commission Vice-President Frans Timmermans called for each member of the European Union to set a minimum wage of 60% of its median salary.

Although not centered on the minimum wage, there is a complementary literature on the interaction of mobility and redistribution, beginning with Stigler (1957) and Oates (1972). Our focus in this work is on the implications of different levels

(2018) for a recent review.

of government undertaking redistribution as in Oates (1977), Ladd and Doolittle (1982), Brown and Oates (1987), and Dixit and Londregan (1998). Our approach to understanding this mechanism is similar to Epple and Romer (1991), who also use a computational model to capture general equilibrium effects. The importance of mobility for policy making in a federation has also been studied in the context of public goods (Epple and Platt, 1998; Caplan, Cornes, and Silva, 2000; Calabrese, Epple, and Romano, 2012; Simon, 2020),[7] income redistribution (Wildasin, 1997; Armenter and Ortega, 2011; Gordon and Cullen, 2012), and higher education (Wildasin, 2000).[8]

2.2 Model

We consider a simple two-region, two-factor model to emphasize the role of regional heterogeneity and factor mobility in determining the relative effectiveness of centralized and decentralized minimum wage setting. The economy, I, consists of two regions, $i \in \{1, 2\}$, low-skilled and high-skilled workers, local firms, and local and federal governments. We consider each in turn.

Each region i has a unit mass of immobile low-skilled labor. Low skilled labor may either be employed, l, or unemployed, u, so that $u_i + l_i = 1$. If employed, they pay a cost θ of working, which is distributed according to a known distribution $\mu(\theta)$. There are $\bar{h}$ high-skilled workers who are (im)perfectly mobile across regions, so that $h_1 + h_2 = \bar{h}$. High-skilled workers face moving costs of $\xi \sim \zeta(\xi)$, but do not

[7]There is also a relatively large literature on the value of centralized environmental policies with spillovers. See Banzhaf and Chupp (2012) for a recent example.

[8]There is also a growing empirical literature on income taxation and mobility. See Kleven et al. (2019) for a review.

face a working cost.[9] We interpret this as a normalization on working costs so that the low-skilled costs are relative to those of the high-skilled. It also ensures that all high-skilled workers are in the labor force for any reasonable minimum wage policy. Workers can only be employed in their region of residence. Consumption, c_j, for each worker j is given by her wage minus her cost of working and moving costs paid, if applicable.

Our mobility assumptions are based on the high correlation between education level and mobility, and the fact that low-skilled workers sluggishly respond to employment shocks.[10] To better understand the implications of these restrictions on workers, we consider how low-skilled mobility affects our theoretical results in the next section. Then, in section 2.4, our calibration exercise varies $\zeta(\xi)$ to describe the relationship between the high-skilled migration elasticity and optimal minimum wage setting authority. Although the model and calibration are based solely on low and high-skilled workers for tractability, the high-skilled represent input factor mobility more broadly. For example, firms may respond to a minimum wage in their capital investment or location decisions and local governments compete for these resources. Our choice of using the high-skilled as the mobile factor allows us to naturally incorporate the mobile input into a social welfare calculation. We discuss how additional mobile inputs affect our results later in this section.

Each region has a single perfectly competitive firm with production function

[9] In our simplified framework, we abstract away from congestion costs, which are likely important for high-skilled migration (Moretti, 2013).

[10] (Amior, 2019) documents this relationship using the CPS. Our assumptions are similar to those in (Wildasin, 2000) who examines the public provision of human capital in a federation, but we allow for variable high-skilled mobility to understand the comparative statics. We discuss this assumption more explicitly later in this section.

$\gamma_i f(l,h)$. Without loss of generality, we assume $\gamma_1 = 1$ and $\gamma_2 = \gamma \leq 1$. Workers are paid their marginal product, and if the high-skilled are perfectly mobile, they migrate so that $w_1^h = w_2^h$. Relative low-skilled competitive wages across regions depends on γ and the relative stock of high-skilled workers. We assume that the optimal minimum wage is never binding for high-skilled workers, so that they are also always employed. Let Θ_{ie} be the set of θ's of the low-skilled employed workers in i and Θ_{iu} the unemployed. For a given wage, workers with relatively large values of θ are unemployed because they choose to exit the labor force. If there is a binding minimum wage, others become unemployed due to the change in low-skilled labor demand by the firm. Similarly, define Θ_{ih} as the set of high-skilled workers in i.

There is one central government, and each region has its own local government. Both the local and central governments have perfect information about workers' working and moving costs and regional productivity, and we explore the consequences of this assumption in Section 2.4.3. Governments care about the consumption of their constituents, according to the concave function $G(c)$, where $G' > 0, G'' < 0$, and $G'(0) < \infty$. Total utility in r is given by $\int_{j \in r} G(c_j)dj$, $r = \{1, 2, I\}$, i.e. each region maximizes the total utility of its own residents, while the central government maximizes the utility of both regions. We assume that $G(c_h) > 0$ for any feasible policy so that high-skilled workers are valued and local governments compete for them. Federal and local governments are either endowed or not with the ability to set a minimum wage. There are four potential authority structures: (1) only regional governments; (2) only central government uniform; (3) both central uniform and regional governments; and (4) central non-uniform. Governments do not need to raise revenue.

Equilibrium is determined in two stages. In the first stage, no government has minimum wage setting authority; this is the competitive equilibrium. This gives each worker a residence. In the second stage, some governments are endowed with minimum wage setting power and optimally set their policy, given the policies of other governments, if applicable, and the responses by firms and workers.

Definition 1. The economy is in equilibrium after the second stage for a given authority structure if:

1. Each government with minimum wage setting authority optimally sets its minimum wage to maximize the welfare of its residents, given the policies of the other governments with authority.

2. Firms maximize profits, taking the residencies of workers and government policies as given.

3. Each worker optimally chooses whether to enter the labor force to maximize utility.

4. No high skilled workers want to move.

2.2.1 Decentralized Minimum Wage

We consider the conditions under which it is optimal for local governments to set a binding minimum wage. Each jurisdiction i's social welfare function is defined as

$$SW_i = (1 - l_i)G(c_u) + \int\limits_{\Theta_e} \mu(\theta)G(c_l)d\theta + \int\limits_{\xi(\Theta_{ih})} \zeta(\xi)G(c_{h_i})d\xi \tag{2.1}$$

where $c_j = w_j - \theta_j$ for low-skilled employed individual of type θ_j, $c_h = w_h$ for high-skilled individuals, and c_u is the consumption of the unemployed. Without taxes and transfers, $c_u = 0$. Since moving costs are only paid if residents leave, they do not directly enter the regional government's social welfare function. As in Lee and Saez (2012), let $g_u = G'(c_0)/\lambda; g_l = \int_{\Theta_e} \mu(\theta)G'(c_l)d\theta/\lambda$ and $g_h = G'(c_h)/\lambda$, where λ is the Lagrange multiplier on the Social Planner's budget constraint.[11]

For a given level of the minimum wage in the other jurisdiction, $\bar{w}_{-i}$, i implements a binding minimum wage $\bar{w}_i$ if and only if:

$$0 < \frac{\partial SW_i}{\partial \bar{w}_i}\Big|_{\bar{w}_{-i}} \iff \tag{2.2}$$

[11]The Social Planner's budget constraint is $(1 - l_i)c_u + l_i c_l + h_i c_h \leq l_i w_l + h_i w_h$. Since the governments in our model do not redistribute and only set a minimum wage, they do not have a budget constraint. Also note that the construction of the g terms are analogous for the regional and central governments, although their respective populations are different.

$$0 < \frac{-\partial l_i}{\partial \bar{w}_i} G(0) + \frac{\partial l_i}{\partial \bar{w}_i} \int\limits_{\Delta_{\bar{w}_i}\Theta_{iu}} \mu(\theta) G(c_l) d\theta$$
$$+ \int\limits_{\Theta_{ie}(\bar{w}_i)} \mu(\theta) G'(c_l) d\theta \frac{\partial w_i^l}{\partial \bar{w}_i} + h_i G'(c_h) \frac{\partial w_i^h}{\partial \bar{w}_i}$$
$$+ \frac{\partial h_i}{\partial \bar{w}_i} G(c_h)$$

The first line represents the change in welfare induced by newly unemployed low-skilled workers. $\Delta_{\bar{w}_i}\Theta_{iu}$ denotes the set of workers who become unemployed due to the minimum wage $\bar{w}_i$, which depends on how workers are separated from the firm after the policy change. Following Lee and Saez (2012), we assume efficient rationing, i.e. workers with the lowest surplus from working involuntarily lose their jobs first. Under the assumption, $\Delta_{\bar{w}_i}\Theta_{iu} = [\gamma_i f_l(l,h), \theta(\bar{w})]$, where $\theta(\bar{w})$ is the largest θ worker who becomes unemployed. $\theta(\bar{w})$ depends on the new policy level, the relative production functions in each jurisdiction, $-i$'s minimum wage, and the amount of migration. For a very small binding policy, $\theta(\bar{w}) = \gamma_i f_l(l,h)$. Therefore, the marginal low-skilled worker who loses her job has $c_j = 0$ before and after the policy change. The second line represents the change in welfare from the still-employed in region i and the third line is the change in welfare due to emigration.

First note that $\frac{\partial w_i^l}{\partial \bar{w}_i} > 0$ by construction and $\frac{\partial l_i}{\partial \bar{w}_i} < 0$ since it is less profitable to employ low-skilled workers at a higher wage. Increasing the minimum wage also impacts the high-skilled workers since $\frac{\partial w_i^h}{\partial \bar{w}_i} < 0$ from the zero-profit condition.[12] As the marginal product of labor for high-skilled workers falls, some high-skilled

[12] With perfectly competitive markets and constant returns to scale, the zero-profit condition implies that $ldw^l + hdw^h = 0$, giving the relationship between the minimum wage and high skilled wages. See Lee and Saez (2012) Appendix A.1. for more information.

workers leave, giving $\frac{\partial h_i}{\partial \bar{w}_i} < 0$ as well. For this reason, the third line of the first order condition only involves emigration and not immigration. Since the local government only cares about its residents, i.e. natives less those who move out, it does not internalize individual migration costs.

Proposition II.1. *If (1) there is efficient rationing; (2) the demand elasticity is finite; (3) the supply elasticity of low-skilled workers is positive; (4) the government values additional wages to low-skilled workers more than the loss of wages to high-skilled (redistribution) and outmigration, then it is total welfare improving for a jurisdiction to impose a (small) binding minimum wage.*

Proof. Our assumptions imply that the first two terms sum to 0 and the magnitude of the third term is larger than the sum of the magnitude of the fourth and fifth terms. More explicitly, using the above expressions and the zero profit condition, we can rewrite the last three terms as $l_i \lambda [g_l - g_h] + \frac{\partial h_i}{\partial \bar{w}_i} G(c_h)$. This is identical to proposition 1 in Lee and Saez (2012) except jurisdictions are now also affected by potential outmigration, weakening the incentive to enact a binding minimum wage. □

Under these assumptions, this extensive margin decision to adopt a minimum wage depends on γ, $\bar{w}_{-i}$, the elasticity of substitution between low and high-skilled workers, and the mobility of the high-skilled. In section 2.3.2, we conduct a calibration exercise to better understand how these parameters impact equilibrium, on both the extensive and intensive margins.

We now consider how changes to the assumptions of Proposition II.1 affect the desirability of a binding minimum wage. Without efficient rationing, the first two terms do not sum to zero for a small increase in the minimum wage because workers with a large surplus from working lose their jobs. If workers were fired randomly after the increase in the minimum wage, then social welfare would be lower in expectation as workers with positive value of employment are separated rather than the indifferent marginal employee. However, even if the lowest cost workers were separated first, a government may still implement a binding minimum wage if the value of redistribution is sufficiently large.[13] This reasoning leads to a corollary to Proposition II.1.

Corollary II.2. *Efficient rationing leads local governments to weakly prefer the highest minimum wages, all else equal, relative to any other separation assumption.*

Proof. Efficient rationing maximizes the sum of the first two terms, while not affecting any other terms. Otherwise, there is a first order welfare loss from unemployment. The condition is weak because a government may not desire a binding minimum wage for any separation policy if the other conditions of Proposition II.1 do not hold. □

If conditions (2) or (3) do not hold, then the employment of low-skilled workers will change too dramatically for the first two terms to sum to zero. Condition (4) may fail if governments care too much about outmigration relative to redistribution, if they do not care about redistribution, or if they only care about redistribution to the unemployed. As a special case, consider when high and low-skilled labor are

[13]Lee and Saez (2012) note that while this is the most favorable assumption for optimal policy, it may not be realistic or may be costly to reach through queuing or search costs. However, there is some empirical literature supporting this assumption. Neumark and Wascher (2007) find that the minimum wage has larger impacts on teenagers and secondary earners who are more likely to have more elastic labor supply and Luttmer (2007) found that reservations wages do not increase with the minimum wage.

perfect substitutes. Then $\frac{\partial w_i^h}{\partial \bar{w}_i} = 0$ and there is no outmigration. In this case, the minimum wage is optimal if there is efficient rationing as the government trades off the unemployment of marginal workers with the additional earnings of all others. As the two types of labor become less substitutable, increasing the minimum wage will lead to lower high-skilled wages and more outmigration. Decreasing the elasticity of substitution will weaken the incentive to have a minimum wage, all else equal.

To further understand the migration externality, consider the impact of raising the minimum wage on the other jurisdiction. If high-skilled workers migrate from jurisdiction i because moving costs are sufficiently small on the margin, the marginal product of labor of high-skilled workers in jurisdiction $-i$ will fall. Government $-i$ is made relatively better off, all else equal, as there is a small change in high-skilled consumption, but a larger change in welfare due to the additional high-skilled workers. This effect lowers the cost for $-i$ to implement a binding minimum wage as fewer high-skilled workers will want to migrate compared to the case where i does not have a binding minimum wage. In our model, the marginal product of low-skilled workers increases when high-skilled workers immigrate, and wages and employment both increase, leading unambiguously to higher utility for those workers.

2.2.1.1 Mobility and the Minimum Wage

Throughout the model we assume that low-skilled workers are perfectly immobile for tractability. While these workers tend to have relatively low mobility rates overall, Monras (2019) finds that after a minimum wage increase, the relative share of low-skilled workers decreases because of changes in low-skilled in-migration. This effect is potentially driven by immigration from other countries. Cadena (2014) provides evidence that low-skilled immigrants prefer states with unchanged minimum wages relative to those experiencing increases because of the disemployment effects. Since our static framework considers a fixed point in time without population change, including from immigration, or shocks to workers that induce migration besides directly from the minimum wage, this concern is not particularly relevant in our highly stylized setting. Monras additionally finds no effect on out-migration. Taken together, these effects on migration are consistent with moving low-skilled not choosing to live in minimum wage areas, rather than migration due to changes in the minimum wage itself. In a dynamic setting where households may move at any given time for any reason, the immobility assumption becomes less plausible and more important.

We now consider the implications of relaxing the perfect immobility assumption.[14] Low-skilled workers initially sort in stage 1 so that w^l is equal across jurisdictions. After a minimum wage increase in i, the firm does not hire any additional low-skilled and some are separated. Therefore, no low-skilled in $-i$ have an incentive to immigrate to i after the policy change. The unemployed in i emigrate if the new wage is higher than the sum of working and moving costs. After a minimum wage increase, low-skilled employment weakly increases in $-i$ from high-skilled migration, but the model does not specify who is hired. If the firm in $-i$ only hires initially local low-skilled workers, then there is no incentive to move and the immobility assumption does not impact our results.[15] If the firm hires workers with the lowest working

[14]Fukumura and Yamagishi (2020) study government competition with low-skilled migration responses from minimum wage changes. However, they focus on an economy with only minimum-wage workers and identical regions.

[15]Phillips (2018) finds that low wage employers discriminate against applicants with a long com-

cost, then the incentive to move depends on γ. Absent any binding minimum wage policies, the marginal worker in jurisdiction A will have a weakly higher working cost than the marginal worker in B since $\gamma \leq 1$. Therefore, a minimum wage in A will not lead to low-skilled migration, but one in B will if moving costs are sufficiently small. In general, adding in low-skilled mobility has an ambiguous effect on the optimal minimum wage policy that depends on the value of $G(0)$ and whether the newly unemployed would be hired if they moved.[16]

The main migration mechanism in the model is the link between regions created from high-skilled. The extent to which high-skilled workers respond to local minimum wage policies, $\frac{\partial h_i}{\partial \bar{w}_i}$, determines the effectiveness of local government policy relative to centralization. If the derivative is zero, because moving costs are high or the marginal product of low and high-skilled workers are unrelated, then our framework reduces to the single region case as studied by Lee and Saez (2012). Monras (2019) estimates the relationship between wages and high-skilled out-migration, finding negative or zero effects; however, they are not statistically different from zero. Cadena (2014) also provides some suggestive, although mixed, evidence that this derivative is in fact negative. Taking the point estimates across all specifications at face value, he finds a negative relationship between the count of immigrants with at least some college and the minimum wage in a given state.[17] More generally, the high-skilled migration

mute, which provides suggestive evidence for this assumption.

[16]If $G(0) < 0$, then local governments can set high minimum wages to induce the unemployed to migrate and raise total social welfare in the jurisdiction.

[17]When controlling for state-specific trends in the main results, the point estimate becomes small in magnitude and imprecise. The unweighted results are slightly larger in magnitude, but still imprecise. Although he interprets this as the minimum wage having no effect on high-skilled migration, the result is likely underpowered.

externality represents the impact on endogenous business locations of low-skilled employers which have been shown to be somewhat responsive to the minimum wage (Rohlin, 2011 and Aaronson et al., 2018). Since our model treats the low and high-skilled as linked through production, we think of the high-skilled as the managers of those firms, with the very high-skilled not captured in our framework.

2.2.2 Extensions of the Decentralized Model

In our baseline model, the results and conditions for a binding optimal minimum wage are driven by the high-skilled migration externality: when a local government raises its minimum wage some high-skilled workers move out. The size of the externality depends on the choice of social welfare function and the production technology. We extend the model to see how changes in these assumptions alter the conclusions. We consider a social welfare function where governments only care about natives or maximize housing values, when there is a single firm that operates in both regions, and the introduction of additional inputs to production.

2.2.2.1 Social Welfare Function

In our framework, we assume that the government only cares about the total utility of its final residents in the second stage equilibrium. In the first part of this section, we consider a government that only cares about natives, regardless of whether they move.[18] When governments maximize utility of their natives, the derivative of

[18]See Wilson (2015) for an overview of different assumptions on government maximization with resident mobility and Cremer and Pestieau (2004) for a review of the literature on government maximization with factor mobility.

the social welfare function shown in equation 2.2 no longer includes the final term: $\frac{\partial h_i}{\partial \bar{w}_i} G(c_h)$. However, they internalize migration costs. The government still trades off lower consumption of high-skilled workers with higher consumption of the low-skilled through the minimum wage, but no longer faces a first order cost from emigration. The implied minimum wages are higher compared to our baseline model.

Proposition II.3. *If conditions 1-3 from Proposition II.1 hold, the government cares only about their initial residents, and it values additional wages to low-skilled workers more than the loss of wages to high-skilled (redistribution) and moving costs, then it is total welfare improving for a jurisdiction to impose a (small) binding minimum wage.*

Proof. This is immediate from the first order condition, where the last term now concerns moving costs and not outmigration. Under this new assumption on the social welfare function, h_i is fixed.

$$0 < \frac{-\partial l_i}{\partial \bar{w}_i} G(0) + \frac{\partial l_i}{\partial \bar{w}_i} \int\limits_{\Delta_{\bar{w}_i}\Theta_{iu}} \mu(\theta) G(c_l) d\theta + \int\limits_{\Theta_{ie}(\bar{w}_i)} \mu(\theta) G'(c_l) d\theta \frac{\partial w_i^l}{\partial \bar{w}_i} + \int\limits_{\xi(\Theta_{ih})} G'(c_h) \left(\frac{\partial w_i^h}{\partial \bar{w}_i} - \frac{\partial \xi}{\partial \bar{w}_i} \right) d\xi$$

□

Corollary II.4. *Local governments that care about their initial residents set weakly higher minimum wage policies than those that care about their final residents.*

Proof. Comparing the first order conditions for a small binding minimum wage, this requires $\frac{\partial h_i}{\partial \bar{w}_i} G(c_h) \leq -\int_{\xi(\Theta_{ih})} G'(c_h) \frac{\partial \xi}{\partial \bar{w}_i} d\xi$. By construction, $\frac{\partial \xi}{\partial \bar{w}_i} = \xi$ if the worker moves and 0 otherwise, allowing us to rewrite the inequality as $\frac{\partial h_i}{\partial \bar{w}_i} G(c_h) \leq -\frac{\partial h_i}{\partial \bar{w}_i} G'(c_h) \mathbb{E}_{\Delta\xi(\Theta_{ih})}[\xi]$. First consider when $-i$ does not have a binding minimum wage. Then $\frac{\partial h_i}{\partial \bar{w}_i} < 0$, and the inequality holds since $G(c_h), G'(\cdot), \xi > 0$ by assumption. If $-i$ has a binding minimum wage and its initial residents return after i puts on a binding minimum wage, then both sides of the inequality are 0, and it holds as well. If i residents migrate instead, then the argument from the first case applies. □

The preceding analysis assumed that governments are utility maximizers, but we now consider a model where rent-seeking governments maximize property values.[19] In the most general application of this approach to our framework, all households and firms demand land. Define $H^d(\bar{w})$ to be the housing or land demand under minimum wage policy $\bar{w}$ and $\Delta_j^{j'} H^d(\bar{w})$ be the change for workers who transition from type j to j' due to the policy change, or for firm f. We assume that housing demand is nondecreasing in income, and that the elasticity of housing supply is finite. Then, for jurisdiction i, a binding minimum wage is optimal if and only if

$$\frac{\partial H^d}{\partial \bar{w}} = \Delta_u^u H^d(\bar{w}) + \Delta_l^u H^d(\bar{w}) + \Delta_l^l H^d(\bar{w}) + \Delta_{h_s}^{h_s} H^d(\bar{w}) + \Delta_{h_s}^{h_{s'}} H^d(\bar{w}) + \Delta_f H^d(\bar{w}) > 0 \tag{2.3}$$

It is immediate that $\Delta_l^l H^d(\bar{w}) \geq 0$. Workers who remain employed after the wage increase will have more income and demand more housing. Similarly, workers that lose their jobs will demand weakly less housing, giving $\Delta_l^u H^d(\bar{w}) \leq 0$. The magnitude of this term importantly depends on how we treat the working cost. If θ is a real

[19]See Epple and Nechyba (2004) for a survey of these two different assumptions on the government's objective function.

cost, then for a small binding minimum wage, the workers who lose their jobs under efficient rationing will have zero net income both before and after the policy change. If it is a utility cost, they will have less income and demand strictly less housing. In the welfare maximization framework, this distinction did not affect optimal policy. In the case where θ is a utility cost, the minimum wage is less attractive because we no longer have an assumption akin to efficient rationing. The always unemployed also have no change in income, implying $\Delta_u^u H^d(\bar{w}) = 0$.

Similar to our analysis in the utility maximization case, there is weakly less demand from high-skilled workers when the minimum wage increases. $\Delta_{h_s}^{h_s} H^d(\bar{w}) \leq 0$ since their wages weakly decrease and $\Delta_{h_s}^{h_s'} H^d(\bar{w}) \leq 0$ due to outmigration from falling wages. The changes are zero in the extreme case where low and high-skilled labor are perfect substitutes.

Finally, consider firms' demand for land. We can decompose the effect on firms as: $\Delta_f H^d(\bar{w}) = \frac{\partial H_f^d}{\partial L(l,h)} \frac{\partial L(l,h)}{\partial \bar{w}}$, where $L(l,h)$ is the total amount of land used by low and high-skilled workers for production. We showed that $\frac{\partial L(l,h)}{\partial \bar{w}} < 0$ as low-skilled workers become unemployed and the high-skilled migrate. If we assume that more workers requires more land, then $\frac{\partial H_f^d}{\partial L(l,h)} > 0$, and firm demand for land decreases. However, if the firm engages in more land-intensive activity when it has fewer workers, then $\frac{\partial H_f^d}{\partial L(l,h)} < 0$, and firm demand for land increases.[20] Summarizing the above analysis gives us the following proposition.

Proposition II.5. *A rent-seeking local government will implement a small binding minimum wage if conditions 1-3 from Proposition II.1 hold and if the housing supply*

[20]In Section 2.2.2.3 we derive the general effect of additional inputs on the optimal minimum wage.

elasticity is finite, housing demand is nondecreasing in after working-cost net income, and the net change in housing demand from the low-skilled employed and firms is positive and larger than the decrease in demand from the high-skilled.

Proof. Under the assumptions, the condition in equation 2.3 is satisfied as described above since the first two terms are 0, and the sum of the remaining are positive. Note that the final term can be either positive or negative. □

The optimal minimum wage trades off more housing demand from the still employed low-skilled with lower demand from the high-skilled due to wage decreases and outmigration. We discussed a comparable set of assumptions to the welfare maximizing case, (i.e. efficient rationing with no change in housing demand for the newly unemployed, and no land demand by firms), that leads to similar policies. We specify utility maximization in our main theoretical and quantitative analysis because the results are theoretically similar and it does not require us to also specify a specific housing market system. This allows us to focus our analysis on the high-skilled migration externality.

2.2.2.2 Regional Economic Linkages

When each jurisdiction has its own production function, asymmetric minimum wage policies lead high-skilled workers to migrate as the policies affect their relative marginal products across locations. If there is a binding minimum wage, employed low-skilled workers are paid that wage, but the amount of unemployment depends on the features of the production function. We now consider alternative assumptions on the production technology to understand how changes in the linkages across states affect the desirability of a minimum wage. We first consider a model where there is only one aggregate production function in the economy. This is motivated by the importance of large firms that operate in multiple states and choose where to hire their labor. For a company that operates in many states, the marginal product of labor of any worker may not only depend on the number of workers in her place of work, but in the entire company.

Proposition II.6. *If conditions 1-3 from Proposition II.1 hold, there is a single aggregate production function, and the government values additional wages to low-skilled workers more than the loss of wages to high-skilled (redistribution), then it is total welfare improving for a jurisdiction to impose a (small) binding minimum wage.*

Proof. With one aggregate production function, the marginal product of high-skilled workers will be the same in both jurisdictions, regardless of the minimum wage policies. High-skilled workers have no desire to migrate and the final term in equation 2.2 equals zero[21]. We are then in the Lee and Saez (2012) model where workers' locations and skill sectors are fixed. □

Proposition II.6 replaces the outmigration condition with the single production function assumption. All else equal, this leads to higher minimum wages than in the baseline model. High-skilled workers are perfectly inelastic to all governments and now bear more of the burden of minimum wage policies, allowing the government to engage in more redistribution to the low-skilled. Since low-skilled workers are

[21]Note that this also implies that there is no difference between a government that cares about natives or residents.

also always paid their marginal product, and marginal products do not differ across locations, when a jurisdiction raises its minimum wage, some low-skilled workers will become unemployed in each jurisdiction due to efficient rationing. Under our assumptions, each jurisdiction has low-skilled employed workers earning zero surplus from working in the competitive equilibrium. This negative externality imposed on the other jurisdiction will lead minimum wage policies to be too high from an efficiency perspective as governments engage in behavior akin to tax exporting.

We now briefly consider two other production technologies that affect optimal policy setting– trade linkages and agglomeration. In the baseline model above, we assume that production is consumed exclusively by local residents. In reality, of course, regions in a federation engage in trade with each other, such that households consume both “home” and “foreign” goods, with relative prices being affected by the schedule of minimum wage policies. If, instead, agents consume goods from both regions the externalities of local minimum wage setting will be smaller.

When the home region sets a binding minimum wage, it restricts the supply of its output, putting upward pressure on its relative price. Without trade, the foreign good becomes relatively more attractive at the margin and agents move in response are able to take better advantage of the lower price of their now home good. However, trade mitigates this effect as workers at home can purchase the relatively less costly foreign good without moving.

Agglomeration would instead amplify the externalities. Since a binding minimum wage in a region leads to high-skilled out-migration, agglomeration effects would imply that relative productivity would decrease. More high-skilled leave because of the implied change in wages.

2.2.2.3 Capital and Other Inputs

With two inputs to production and zero-profits, high-skilled wages decrease with the minimum wage. We introduce elastically supplied capital, k, with required return, $\bar{r}$, so that this is not necessarily the case. With three production inputs, the zero-profit condition implies $\frac{dw^h}{d\bar{w}} = \frac{-l}{h} - \frac{k}{h}\frac{d\bar{r}}{d\bar{w}}$. If $\frac{d\bar{r}}{d\bar{w}}$ is relatively large in magnitude and negative, the marginal product of high-skilled workers increases with the minimum wage. The increased marginal product will induce migration into the jurisdiction as high-skilled workers move for higher wages. In this case, it is Pareto efficient[22] to implement a very small binding minimum wage since the marginal low-skilled worker is indifferent between working and not due to efficient rationing, and the wages of all other still employed low-skilled and all high-skilled workers increase. Raising the minimum wage further will not be Pareto efficient, and the government will trade off additional consumption from the employed low-skilled and high-skilled, and migration into the region from high-skilled workers with reduced consumption from the newly unemployed.

From the zero-profit condition, we can generalize $\frac{dw^h}{d\bar{w}}$ when there are additional inputs, indexed by ι with price p^ι. Then,

$$\frac{dw^h}{d\bar{w}} = \frac{-l}{h} - \sum_{\iota \neq h,l} \frac{\iota}{h}\frac{dp^\iota}{d\bar{w}} \tag{2.4}$$

[22] We are implicitly assuming that capital is paid its marginal product and always earns zero profit so we do not need to worry about rents to capital owners.

Proposition II.7. *When there are additional inputs to production besides high and low-skilled labor, a small binding minimum wage is Pareto Efficient if conditions 1-3 from Proposition II.1 hold and* $\frac{dw^h}{d\bar{w}} \geq 0$

Proof. When $\frac{dw^h}{d\bar{w}} \geq 0$, there is no outmigration, and by efficient rationing all workers are weakly better off. □

2.2.3 Alternative Configurations for Minimum Wage Setting Authority

Using our baseline model, we consider the implications of alternative minimum wage setting authorities. As in the decentralized case, we develop conditions under which each type of government finds it optimal to implement a binding minimum wage.

2.2.3.1 Only Federal Government

In this section, we consider optimal minimum wage setting when only the federal government can set a minimum wage. The government's social welfare function is defined as

$$SW_f = \sum_i (1 - l_i)G(0) + \int_{\Theta_{ie}} \mu(\theta)G(c_l)d\theta + \int_{\xi(\Theta_{ih})} \zeta(\xi)G(c_h)d\xi \qquad (2.5)$$

Note that in the federal government's optimization problem, it cares about the utility of all high skilled workers regardless of residence. Unlike in the regional government's problem, when high skilled workers move, their utility still matters for the federal government. The government also cares about the migration costs paid by individual workers.[23]

2.2.3.2 Federal Non-Uniform Policy

When the federal government can set a different policy in each region, the policies are given by:

$$(\bar{w}_1, \bar{w}_2) = \arg\max SW_f(\bar{w}_1, \bar{w}_2) \tag{2.6}$$

The FOC for $\bar{w}_1$ imply that it will be binding in jurisdiction 1 if:

$$\begin{aligned}
0 < & -\frac{\partial l_1}{\partial \bar{w}_1} G(0) + \frac{\partial l_1}{\partial \bar{w}_1} \int\limits_{\Delta_{\bar{w}_1}\Theta_{1u}} \mu(\theta) G(c_l) d\theta \\
& -\frac{\partial l_2}{\partial \bar{w}_1} G(0) + \frac{\partial l_2}{\partial \bar{w}_1} \int\limits_{\Delta_{\bar{w}_1}\Theta_{2u}} \mu(\theta) G(c_l) d\theta \\
& + \int\limits_{\Theta_{1e}(\bar{w}_1)} \mu(\theta) G'(c_l^1) \frac{\partial w_l^1}{\partial \bar{w}_1} + \int\limits_{\Theta_{2e}(\bar{w}_1)} \mu(\theta) G'(c_l^2) \frac{\partial w_l^2}{\partial \bar{w}_1} \\
& + G'(c_h) \left(\bar{h} \frac{\partial w^h}{\partial \bar{w}_1} - \int\limits_{\xi(\Theta_{ih})} \frac{\partial \xi}{\partial \bar{w}_1} \zeta(\xi) d\xi \right)
\end{aligned}$$

The expression above is simplified by noting that $\frac{\partial h_1}{\partial \bar{w}_1} = -\frac{\partial h_2}{\partial \bar{w}_1}; h_1 + h_2 = \bar{h};$ and $\frac{\partial w_h^1}{\partial \bar{w}_1} = \frac{\partial w_h^2}{\partial \bar{w}_1}$.[24] We also use the fact that moving costs are randomly assigned after households choose a residence in the competitive equilibrium and that under our assumptions,

[23] The final term of the government's social welfare function therefore involves integrating over the types of households in each jurisdiction.

[24] The last equality only holds if the lower bound on the moving costs are zero, which we maintain throughout, and for a small minimum wage.

there will be net migration to jurisdiction 2. The FOC for $\bar{w}_2$ is defined analogously.

The first row of the FOC gives the welfare loss from newly unemployed low-skilled workers in jurisdiction 1. The second row gives the welfare gain from new low-skilled employment in jurisdiction 2. As high-skilled move to jurisdiction 2, the marginal product of low-skilled labor increases, leading workers to enter the labor market. The third row gives the welfare gain to the still-employed in both jurisdictions. The wage, conditional on employment, increases in jurisdiction 1 due to the new binding minimum wage, while the wage in jurisdiction 2 increases because of induced high-skilled immigration. The final row gives the welfare loss to high-skilled workers, due to both the change in wages and moving costs paid by those who move to jurisdiction 2.

Proposition II.8. *If (1) there is efficient rationing; (2) the demand elasticity is finite; (3) the supply elasticity of low-skilled workers is positive; (4) the government values additional wages to low-skilled workers more than the loss of wages to high-skilled; and (5) the first worker to move to the other jurisdiction has zero moving costs, $\xi = 0$, then it is total welfare improving for the federal government to impose a (small) binding minimum wage.*

Proposition II.8 is the same as Proposition II.1, except the central government only cares about migration insofar as it affects efficiency and moving costs paid. Condition (5) is similar to efficient rationing, but for moving, so that moving costs are zeros for a small binding minimum wage. It implies that the lower bound on the moving cost distribution is zero.[25] Even without condition (5), the federal govern-

[25]This assumption on the moving cost distribution is not realistic and is chosen to provide suf-

ment may implement a binding minimum wage if it values redistribution enough. An increase in the federal minimum wage leads to lower high-skilled wages, regardless of moving costs, and is therefore not Pareto improving. Comparing Propositions II.1 and II.8, we expect the federal government to prefer higher minimum wages than regional governments, all else equal.

The case in which the federal government can set a different minimum wage in each region is synonymous with the social optimum. The social optimum is defined in our context by the sum of total welfare for all agents in all regions. This is exactly the function maximized by the federal government. When its policy is allowed to be asymmetric, it has the flexibility to reach the socially optimal minimum wage in each region.

2.2.3.3 Federal Uniform Policy

In this case, we restrict the central government to set a single uniform policy. The minimum wage is given by:

$$\bar{w} = \arg\max SW_f(\bar{w}) \tag{2.7}$$

The FOC is quite similar to the non-uniform case, where the responses are given with respect to $\bar{w}$ instead of $(\bar{w}_1, \bar{w}_2)$. When $\gamma = 1$, the solution will be the same in the non-uniform and uniform cases. If $\gamma < 1$, then the competitive equilibrium

ficient, but not necessary, conditions for a federal minimum wage. However, in a long-run version of our model where workers do not have a fixed residence, the moving costs can be interpreted as preferences for one region compared to the other. In that case, it is plausible that some workers are indifferent, all else equal.

low-skilled wages in jurisdiction 1 will be greater than those in 2. Therefore, when considering a small binding minimum wage, it will only be binding in 2. Under the assumptions of Proposition II.8, the Social Planner will always want a binding minimum wage in each jurisdiction, however it may be the case that even when the assumptions hold, the federal government does not set a binding minimum wage if constrained to only set a uniform policy. When the Social Planner sets a binding policy in each jurisdiction, it can mitigate the welfare loss due to migration, both moving costs and the suboptimal allocation of individuals across jurisdictions. When the federal government can only set one policy and it is only binding in a jurisdiction, then there is an externality imposed on the other. This is seen by noting that all of the components of the government's first order conditions depend on the enforced minimum wage in both jurisdictions. For $\gamma < 1$, a small binding uniform policy is equivalent to the Social Planner only setting a minimum wage in the low productivity region, which the previous section showed is suboptimal.

2.2.3.4 Tiered System

The U.S. has a tiered system in which the federal and local government both set minimum wages, and the larger of the two is enforced. The federal government is restricted to set a uniform policy. Based on our previous analysis, under the assumptions of Proposition II.1, each local government sets a binding minimum wage, and under the assumptions of Proposition II.8, the federal government also sets a binding policy for at least the less productive region. As a result, there will be binding policy enacted in each. In the symmetric case, the federal minimum wage will be weakly higher than both jurisdictions' preferred policies because of the differential effects of outmigration and moving costs. The effect of moving costs must be less than outmigration by construction.

2.3 A Two Region Example

2.3.1 Parameterization

In the interest of exploring the relative welfare consequences of different configurations and the relative advantages of decentralized and centralized policy setting, we consider a parameterized version of the model. The calibrated model illustrates the mechanisms highlighted in the theory section, particularly regional heterogeneity and competition. In the previous sections, we are unable to give theoretical predictions for each policymaker beyond the extensive margin conditional on the other policies. The calibrated simulations in the next section on optimal policy provide results for the intensive margin decisions when governments compete while factoring in the general equilibrium effects of their actions.

We specify the following functional forms. We assume f is a constant elasticity of substitution (CES) production function with elasticity of substitution between high and low-skilled labor $\sigma = \frac{1}{1-\rho}$ and weight α on the low-skilled.

$$f(l,h) = (\alpha l^{\rho} + (1-\alpha)h^{\rho})^{\frac{1}{\rho}} \tag{2.8}$$

With this production function, $\rho \leq 1$. If $\rho = 1$ the two types of labor are perfect substitutes and if $\rho = -\infty$ then the two types are perfect complements. Worker heterogeneity is determined by the distributions on θ and ξ, which we assume to be $\theta \sim \mathcal{U}(0, \bar{\theta})$ and $\xi \sim \mathcal{U}(0, \bar{\xi})$. The sufficient conditions from propositions II.1 and II.8 require that the lower bounds are zero. Alternative distributional assumption on the moving costs do not qualitatively affect our results.[26]

We also assume that the social welfare aggregator is defined as:

$$G(c_j) = \int_{j \in r} \log(c_j + \nu) dj \tag{2.9}$$

where ν is a utility function shifter. Since some low-skilled workers are always unemployed with $c_u = 0$, ν must be positive. $\log(\nu)$ is therefore the value of unemployment to the government. Smaller values of ν do not affect labor supply, but increase the social cost of unemployment. We choose individual utility to be net, after migration and disutility from working, consumption so that there are no income effects.

2.3.2 Calibration

We calibrate our model to match aggregate moments of the U.S. economy. We have seven parameters: α, γ, and ρ for the production function; $\bar{h}, \bar{\theta}$, and $\bar{\xi}$ to describe the worker population; and ν for the social welfare function. We vary γ to see how

[26]For example, modeling the moving costs as a type I extreme value idiosyncratic preference for region 1, similar to our quantitative model in Section 2.5, does not affect the mechanisms highlighted here. In both cases, low ξ high-skilled workers in i will move after an increase in the minimum wage in i. However, the number of movers, and therefore optimal policy, depends on the distributional assumptions. We specify that the moving costs are normally distributed so that all draws of ξ are non-negative for all values of γ. Under the type I extreme value assumption with asymmetric regions, some individuals have negative moving costs in the baseline. For the purposes of this exercise, we do not allow for that source of misallocation.

different levels of regional heterogeneity impacts our results; for the calibration, we set $\gamma = 1$, corresponding to the symmetric case. Additionally, we let $\bar{h} = 1$ to reflect the fact that about one-third of the population has a Bachelor's degree (CPS), and $\rho = 0.286$, based on the estimate of the elasticity of substitution between high and low-skilled workers in Katz and Murphy (1992).[27]

Based on the timing assumptions of our model, we calibrate the remaining parameters in two stages. In the first stage of the calibration, based on the competitive equilibrium, we set α and $\bar{\theta}$. We calibrate α to match the relative wages across skill types. Using the Current Population Survey Outgoing Rotation Group microdata, the Economic Policy Initiative reports hourly wages in 2017 by highest degree received. Based on their report, we set $\frac{w^h}{w^l} = 1.82$. We calibrate $\bar{\theta}$ to match the 2017 labor force participation rate of high school graduates. Based on estimates from the Bureau of Labor Statistics in 2017, the value for high school graduates is 57.7%, while the value for those with only a Bachelor's degree is 73.3%. Since our model assumes that high skilled workers are always in the labor force, we scale up the high school rate to $\frac{57.7}{73.3} = 78.7$. ν is normalized to 1 so that the utility of the unemployed is 0. It also ensures that $G(c_h) > 0$

In the second stage, after all workers have established residencies from the competitive equilibrium, we calibrate the upper bound on the moving cost distribution, $\bar{\xi}$ to match the elasticity of migration with respect to wages of 0.5 from Kennan and Walker (2011). For a 10 percent decrease in γ, $\bar{\xi}$ is found so that there is an increase

[27]For the purposes of the calibration, we set $\bar{h}$ using college graduates to match the definition of high-skilled in Katz and Murphy.

in the stock of high-skilled workers in jurisdiction 1 of 5 percent.[28]

The parameter estimates are presented in Table 2.1. The estimate column gives the values used in the welfare calculations and the residuals note the distance between the model implied value and the target, when relevant.

Table 2.1: Calibration Results

Parameters	Estimate	Residual
α	.43	10^{-9}
$\bar{h}$	1	-
$\bar{\theta}$	.46	10^{-9}
ρ	0.29	-
$\bar{\xi}$	1.31	10^{-13}
ν	1	-

From the baseline calibrated model, Figure 2.1 plots total social welfare in the economy as a function of the minimum wage. In the left panel, the minimum wage is effective in both regions and so this figure represents the problem of the central government restricted to a uniform policy. A small binding minimum wage increases welfare since the government values redistribution, the marginal separated low-skilled worker has 0 surplus from working, and no high-skilled workers move since the regions are identical and a uniform policy is enforced. As the minimum wage increase further to 102.5 percent of the competitive wage, welfare also increases. After this point, the higher price floor causes relatively too much additional unemployment and lost high-skilled wages. When the minimum wage is 5.6 percent above the competitive

[28] We interpret the Kennan and Walker (2011) partial equilibrium results as a productivity shock in our general equilibrium model. The shock in one region affects high-skilled wages directly and through its effect on low-skilled labor supply, while the general equilibrium migration effect moves in the opposite direction. If instead we shock just high-skilled wages in one state holding everything else fixed, then $\bar{\xi}$ would barely change to 1.3140 instead of 1.3134.

level, the policy is worse than no policy at all.

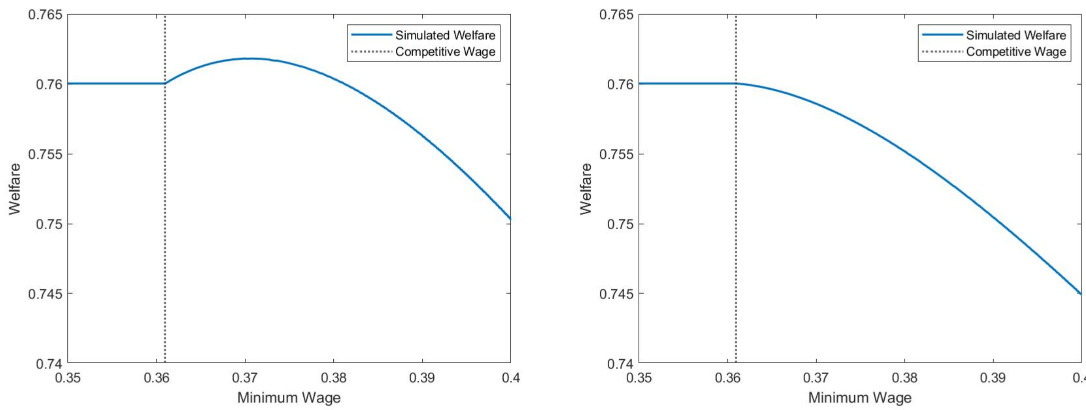

(a) Single Minimum Wage in Both Regions (b) Minimum Wage in Region 1 Only

Figure 2.1: Welfare and the Minimum Wage

Note: Both panels of the figure plot the total social welfare in the economy for various levels of the minimum wage. The left panel plots welfare as a function of a single minimum wage applied in both regions while the minimum wage only applies in region 1 in the right panel. Since the regions are identical, the implications from a minimum wage in only 1 or 2 are the same. The vertical line denotes the competitive low-skilled wage in the absence of any policy.

The right panel of the figure plots total welfare as a function of the minimum wage in region 1 with no minimum wage in the other. If region 2's government does not set a binding policy, this figure depicts the decentralized problem for region 1. For a small binding policy, total social welfare decreases in region 1 because the loss of high-skilled workers is more than the value from additional wages to the low-skilled. Welfare in region 2 increases from the inflow of new high-skilled workers, which increases the total population and low-skilled wages but at the expense of high-skilled wages. The increase in region 2 is not enough to offset the loss in region 1 and so a small binding minimum wage in only one region decreases total welfare in the economy. The migration externality is too large relatively to the benefits from redistribution.

2.4 Optimal Minimum Wage Setting

To supplement the theoretical work on the extensive margin, we use the calibrated model to trace out the relative benefits of centralization verse decentralization as a function of mobility and regional heterogeneity for optimal minimum wage setting. Our results illustrate the trade-off captured by the decentralization theorem, where the federal government can only set a single uniform policy and local governments can better adjust to their own needs. We find that centralization of minimum wage setting authority is always weakly preferred to decentralization in our baseline calibrated model, even when regions are heterogeneous, because of interregional spillovers and the corresponding competition for high-skilled workers. Local governments never want to set binding minimum wages, while the central government loses its incentive to set a uniform minimum wage as heterogeneity across regions increases.

Panel A of Table 2.2 presents results for the symmetric case, i.e., when relative productivity $\gamma = 1$. The five rows correspond to the five different types of government structure that we allow in our context. The top row is the competitive equilibrium with no government, and the bottom row is the case of a benevolent central government setting non-uniform policy, which is synonymous with a social planner solution. The middle three rows correspond to cases with local governments only, a central government restricted to uniform policy only, and a U.S.-style hybrid system with both of these types, respectively. Columns list the values for unemployment, wages for high-skilled and low-skilled workers, and the fraction of (mass 1) high skill workers that sort to region 1. If the wages of a particular case match the wages in the competitive case, then there is no binding minimum wage. The far right column lists the welfare loss relative to the federal non-uniform or social planner.

In the symmetric case, local governments do not want to set a binding minimum wage because the resulting outmigration is too large to be offset by the increase in wages for the low-skilled. Condition 4 of Proposition II.1 does not hold. In contrast, the central government is not affected by this race-to-the-bottom behavior to keep high-skilled workers because migration does not change the set of agents in the central government's welfare function. In the symmetric case, the central government exactly matches the social planner's policy, since the social planner sets identical minimum wages in each region.

In our calibrated equilibrium, we find that the social planner and central government set the optimal minimum wage to 102.5% of the competitive wage. The binding minimum wage leads to a small total welfare increase over the competitive and decentralized cases of 0.23%. With higher low-skilled wages, the always employed low-skilled have a welfare gain at least 2.14% that depends on their working costs. The gain comes at the expense of the newly unemployed, whose welfare decreases by 100% as consumption drops to 0, and the high-skilled, whose welfare falls by 1.65%. As more low-skilled workers become unemployed, the marginal product of the high-skilled workers falls, leading to lower consumption. The always unemployed are unaffected by changed in the policy.

With heterogeneous regions, the central government's ability to match the social planner's solution disappears. Panel B of Table 2.2 presents the outcomes when γ is reduced to .995. With this very small productivity difference, the central government desires a binding minimum wage that is between the social planner's two policies: slightly too high for the less productive region, but too low for the more productive area. The centralized and tiered environments both result in a positive welfare loss relative to the social planner, though they are still preferred to the decentralized case where local governments do not implement a binding policy.

Finally, we present results for $\gamma = 0.99$ in Panel C of Table 2.2. In this case, the local and central uniform governments do set a binding minimum wage, while the social planner's solution sets a different binding policy in each region. As in the previous cases, the migration incentive is too strong for the local governments to set a binding policy, and now that is also the case for the central uniform as well. A small binding central uniform policy only affects the less productive jurisdiction, which causes high-skilled workers to migrate. The migration externality decreases welfare as total output in the economy shrinks, even more low-skilled workers in region 2 become unemployed, and some high-skilled pay migration costs. In the $\gamma = 0.995$ case, the jurisdictions were similar enough that the same policy implemented in both did not cause too large of an externality. With large amounts of heterogeneity, like in the U.S., central uniform policies are inefficient. The optimality of a central uniform minimum wage is decreasing in regional heterogeneity. With two policies to set, the social planner is able to set the binding minimum wage in each jurisdiction so that the populations remain the same as in the competitive equilibrium.

Table 2.2: Baseline Results

Authority	u	w_1^l	w_2^l	avg w^h	h_1	SW Loss
Panel A: $\gamma = 1$						
Competitive	.426	.361	.361	.657	.5	.23%
Decentralized	.426	.361	.361	.657	.5	.23%
Federal Uniform	.523	.370	.370	.643	.5	0
U.S.	.523	.370	.370	.643	.5	0
Federal Non-Uniform	.523	.370	.370	.643	.5	-
Panel B: $\gamma = .995$						
Competitive	.429	.362	.358	.655	.506	.23%
Decentralized	.429	.362	.358	.655	.506	.23%
Federal Uniform	.533	.370	.370	.640	.509	.16%
U.S.	.533	.370	.370	.640	.509	.16%
Federal Non-Uniform	.533	.372	.368	.640	.507	-
Panel C: $\gamma = .99$						
Competitive	.432	.364	.356	.653	.513	.23%
Decentralized	.432	.364	.356	.653	.513	.23%
Federal Uniform	.432	.364	.356	.653	.513	.23%
U.S.	.432	.364	.356	.653	.513	.23%
Federal Non-Uniform	.522	.372	.364	.640	.513	-

Note: This table presents the simulated equilibria for our baseline calibrated model different minimum wage setting authorities and levels of regional heterogeneity. The competitive equilibrium is provided for comparison. Social Welfare (SW) Loss is defined relative to the Federal Non-Uniform, or Social Planner.

2.4.1 Variable Mobility

The relative differences in government competition of mobile high-skilled agents drives our main results in the previous section. Local governments are averse to minimum wage induced outmigration, as it creates a first order welfare loss. The central government faces a much smaller cost from migration, especially since the first workers who move have no migration costs. However, both types of governments bear the burden of decreased total output as households sort across jurisdictions. We now explicitly vary the mobility of the high-skilled to better understand this mechanism. With high migration costs, the model captures the short-run effects as workers are unable to move in response to the policy. As migration costs decrease, the model better captures the long-run equilibrium effects.

When migration costs increase, minimum wages become more desirable for regional governments; when high-skilled workers are immobile, regional governments enact exactly the social planner policy. Without migration, there are no interregional spillovers and the local governments' problems exactly match that of the social planner because the two economies are additively separable in the social welfare function. From the baseline symmetric model, if the minimum of the moving cost distribution is increased slightly to 0.0025 from 0, or 0.3% of the competitive high-skilled wage, then the decentralized outcome would match the social planner as well. The lowest moving cost is large enough that the high-skilled are effectively immobile in equilibrium. As long as the minimum moving cost is more than 0.03% of the high-skilled competitive wage, regional governments will enact binding policy in the baseline model. Small changes in the moving cost have large implications for government competition.

To understand the interaction between productivity heterogeneity and mobility, we set γ to .95 and the minimum of the migration cost distribution to 0.0025. Table 2.3 presents the results. As before, the social planner implements a different binding minimum wage in each jurisdiction. In the decentralized case, migration costs lessen horizontal competition enough so that both local governments set binding policies, but not enough to replicate the social planner's solution. The central uniform policy alone leads to greater welfare loss than the decentralized policy setting because there is considerable heterogeneity, but the government only has one instrument. However, the hybrid system captures the social optimum. When the central government sets its policy equal to that of the social planner in jurisdiction 2, it now becomes optimal for jurisdiction 1 to raise its policy. The local government faces a smaller threat from outmigration than the decentralized case where the regional government in jurisdiction 2 sets its policy too low. In this more realistic case, where jurisdictions are heterogeneous and all migration is at least a little costly, the hybrid minimum wage setting authority is welfare improving over either by itself. Centralized and decentralized policy setting are strategic complements. This justifies the policies we see in Seattle, for example, while not supporting the policies by the Missouri state government of prohibiting its cities from setting higher minimum wages or in Europe where many countries set a single national policy.

At the other extreme, we consider perfectly mobile high-skilled workers. In this case, the outmigration response to a minimum wage will be even stronger than in the baseline case, such that local governments will have even less of an incentive to implement binding minimum wages. For the central government, the effect is ambiguous. The lack of moving costs means that migration is less costly, since what matters to the central government is the costs paid by individuals rather than the loss of population. However, the central government also cares about the total output

Table 2.3: $\gamma = .95$ and $\xi \sim \mathcal{U}[0.0025, \bar{\xi}]$

Authority	u	w_1^l	w_2^l	w^h avg	h_1	SW Loss
Competitive	.456	.374	.334	.638	.566	.21%
Decentralized	.497	.378	.338	.632	.566	.08%
Federal Uniform	.466	.374	.336	.367	.566	.19%
U.S.	.542	.383	.343	.626	.566	0%
Federal Non-Uniform	.542	.383	.343	.626	.566	-

Note: This table presents the simulated equilibria for the case where $\gamma = .95$ and $\xi \sim \mathcal{U}[0.0025, \bar{\xi}]$ for different minimum wage setting authorities. The competitive equilibrium is provided for comparison. Social Welfare (SW) Loss is defined relative to the Federal Non-Uniform, or Social Planner.

in the economy, which is decreasing with migration. When $\gamma = 1$, no households move as a result of federal policy, and the calibrated migration and perfect mobility equilibria are identical. When $\gamma = 0.95$, the federal non-uniform optimal policies are 0.379 in jurisdiction 1 and 0.340 in jurisdiction 2, compared to 0.383 and 0.343 in the baseline calibrated case. Higher levels of mobility push down optimal federal policy.

2.4.2 Alternative Social Welfare Functions

The previous quantitative results rely on a single specification of the social welfare function, namely, $U(c) = \log(c + \nu)$, where we set $\nu = 1$. ν is a shifter on the progressivity of the welfare function, with lower values corresponding to a more progressive desire for redistribution. The right panel of Figure 2.2 presents the equilibrium minimum wages desired by local governments and by the social planner as a function of the ν for the symmetric case holding all other structural parameters fixed. Since our theory requires that $G(c_h) > 0$, ν cannot be too small in our baseline model. To fully investigate the importance of progressivity, the left panel of the figure uses the social welfare function $\log(2(c+\nu))$ so that $G(c_h) > 0$ for all $\nu > 0$ since in the symmetric baseline equilibrium, $c_h > .6$ for any feasible minimum wage policy.

The optimal minimum wages are an inverted-U shape in progressivity. As ν approaches zero, the welfare of the unemployed quickly heads toward negative infinity, such that adding to unemployment becomes increasingly costly from a social welfare perspective. Minimum wages distribute from the high-skilled and the newly unemployed to the low-skilled employed. If ν were 0, then the social welfare would be negative infinity for all policies, including no binding minimum wage, since there will always be some unemployed workers in our model. In that case, the optimal minimum wage would not be defined for any government. The pattern is identical when $\gamma < 1$.

We now consider a utilitarian social welfare function. Since our agents maximize consumption, this social welfare function is equivalent to maximizing the sum of wages minus any working costs or moving costs paid. In the symmetric $\gamma = 1$ case, no binding minimum wages are optimal for any government. With a relatively large amount of regional heterogeneity, for example when $\gamma = 0.95$, the social planner desires a small binding minimum wage of about 0.3% over the competitive wage in each jurisdiction. Under our baseline social welfare function, when $\gamma = 0.95$, the social planner would set optimal minimum wages of about 2.5% above the competitive.

In addition to the productivity heterogeneity explored previously, regions may also differ in their progressivity. St. Louis and Kansas City, Missouri as blue cities in a red state may have set higher minimum wages because they value redistribution more. To test this implication, Figure 2.3 varies ν for only jurisdiction 1, holding

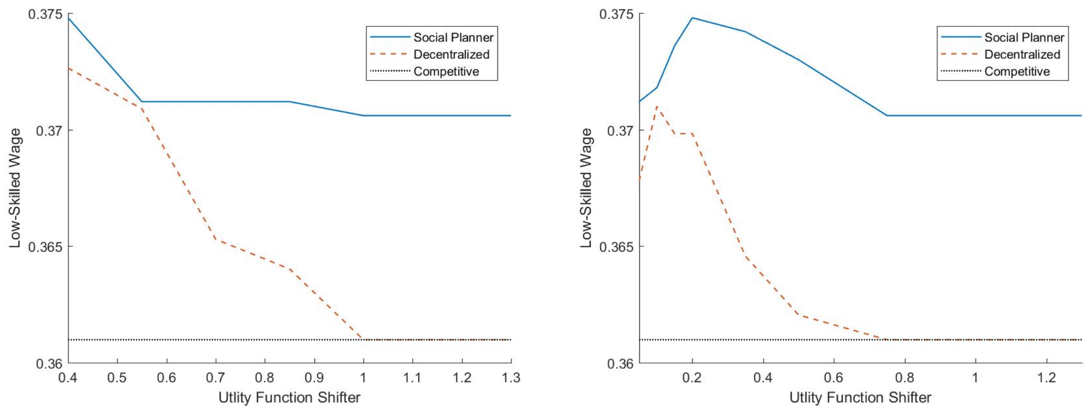

(a) Baseline model

(b) Utility given by: $\log(2(c + \nu))$

Figure 2.2: Optimal Minimum Wage as a Function of ν

Note: In the right panel, we vary ν from 0.4 to 1.3 using the baseline calibrated model. In the left panel, we vary ν from 0.05 to 1.3 changing the social welfare function to $\log(2(c + \nu))$ The low-skilled wage in a jurisdiction is the maximum of the competitive wage and that government's minimum wage policy. The Social Planner's optimal policy (solid line) equals the Central Uniform and Tiered optimal policies since jurisdictions are symmetric.

the value for jurisdiction 2 fixed at 1, and plots the equilibrium decentralized policies and the competitive outcome when jurisdictions are otherwise identical. When ν is slightly less than 1, jurisdiction 1 values redistribution sufficiently to set a binding minimum wage. Unlike in the baseline, condition 4 of Proposition II.1 is now satisfied. Although the welfare cost of setting a minimum wage for the government in jurisdiction 2 is now lower because the binding minimum wage in the other region reduces the value of migration for the high-skilled, it remains at the corner of no binding policy. It is not until ν in jurisdiction 1 falls below 0.7 that jurisdiction 2 sets a binding minimum wage. As ν in jurisdiction 1 falls further, jurisdiction 2 is induced to set a binding minimum wage. In this region of ν, governments engage in race to the top. When ν decreases even further, the benefit of setting a high minimum wage also decreases and the policies begin to level off. Even when regions are equally productive, differences in the value of redistribution can lead competing governments to set binding decentralized policies.

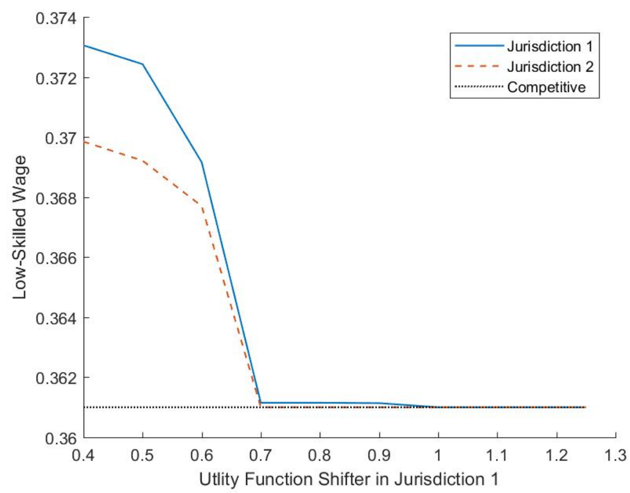

Figure 2.3: Optimal Minimum Wage as a Function of ν in Jurisdiction 1

Note: We vary ν in jurisdiction 1 from 0.4 to 1.3 keeping ν in jurisdiction 2 fixed at 1. The low-skilled wage in a jurisdiction is the maximum of the competitive wage and that government's minimum wage policy.

2.4.3 Information

In addition to mobility and regional competition, the relative information quality of different levels of government affects the optimal minimum wage setting authority. Decentralized policy setting may be preferred because local governments have better knowledge of the labor supply, labor demand, and migration elasticities or local productivity. These parameters jointly determine how employment changes after a minimum wage increase. In a 2019 report, the CBO documented a large amount of variation in estimates of the employment elasticities from the recent minimum wage literature from 0.4 in Cengiz et al. (2019) to -1.7 in Clemens and Wither (2019). Which estimate should a government use when determining its policy? In the context of our model, if a government thought the employment elasticity were positive, then it would believe that a binding minimum wage is Pareto improving and only faces a welfare trade-off when the estimate is negative. Therefore different estimates lead to very different "optimal" policies. Local governments may additionally have better information about the market structure, which we do not explore in this paper, or political beliefs leading them to favor one estimate over another; for example, if it believes there is significant monopsony power in the region, it may place more weight on positive estimates. If any government has noisy or incorrect information about the regions or these parameters, then it will not be able to set policy optimally.

While the previous analysis assumed that both levels had perfect information about workers and regions, we now explore the consequences of misspecified policy.[29] Panel (a) of Figure 2.4 presents the welfare for different minimum wage policies in each region for the baseline model normalized by the welfare in the competitive case with no binding policy. In each row, the minimum wage varies from 0 to 10 percent above the competitive wage in region 1. The columns vary the price floor in region 2. The grid therefore constructs a discretized version of the social planner's problem and it confirms that the optimal minimum wage is about 2.5

[29] We do not make explicit assumptions about why the minimum wage differs from the optimal policies found in the previous subsection but instead simulate how deviations affect welfare. These errors may be due to a lack of information.

percent above the competitive. Similarly, looking down the main diagonal gives the optimal central uniform policy. Based on these simulations, we see that very large uniform minimum wages can lead to lower total welfare than having no policy, which is also the decentralized case. Importantly for policy-making, these results suggest that when regions are symmetric, a minimum wage that is too high is better than one that is too low. Specifically, we find that a minimum wage 0.5 percentage points too high yields higher welfare than one that is 0.5 percentage points too low. The same is true when we compare policies that are 1.5 and 2.5 percentage points off in either direction.[30]

Asymmetric policies for which the difference is more than 2 percentage points also lead to lower welfare than having no policies at all. When the enforced minimum wage is too different in otherwise identical regions, the migration externality overpowers the benefits from redistribution. Since symmetric policies do not induce migration, relatively more of the burden falls on high-skilled workers that the progressive government values less. However, large differences in the enforced low-skilled wage across equally productive regions leads to a large amount of migration. This exacerbates the disemployment effects of the minimum wage in the region with a higher policy. In our simulations, we find that minimum wages of 101.5 and 103.5 percent of the competitive performs worse than a uniform policy of 106 percent.

To understand how errors in optimal policy setting from a lack of information vary with productivity and mobility, Panel (b) displays the relative welfare based on the simulations shown in Table 2.3 where $\gamma = .95$ and $\min(\xi) = 0.0025$. With

[30]We cannot compare more than 2.5 percentage points off of the optimal because then the low minimum wage would not be binding.

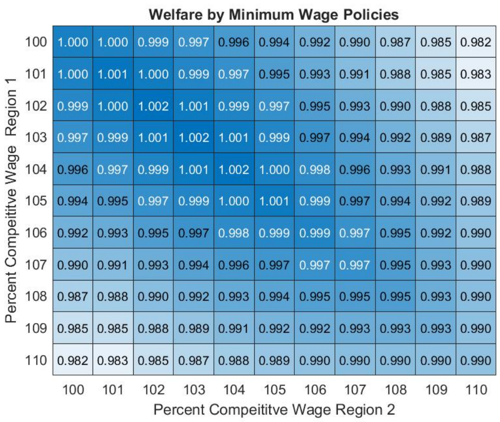

(a) $\gamma = 1$, $\min(\xi) = 0$

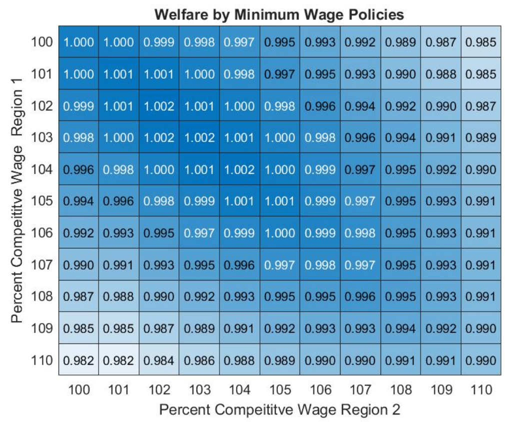

(b) $\gamma = .95$, $\min(\xi) = 0.0025$

Figure 2.4: Welfare and the Minimum Wage

Note: Panel (a) presents the total social welfare in the economy for different minimum wage policies in both regions ranging from the competitive wage to 110 percent of the competitive. Welfare is normalized so that it is 1 in the competitive case where neither region has a binding minimum wage. Panel (b) conducts the same simulations but for the case where $\gamma = .95$ and $\min(\xi) = 0.0025$ as in Table 2.3. Minimum wage combinations with higher total welfare are shaded darker.

perfect information, the optimal uniform central policy is less than 101 percent of the competitive wage in region 2 and not binding in region 1 while the optimal non-uniform policy is about 102.5 percent above each region's respective competitive wage.[31] Since moving costs are positive on the margin, both regions set binding policies of about 1 percent above the competitive in the decentralized case as well. The patterns are very similar to the baseline example except the welfare matrix is no longer symmetric. Based on the asymmetry, we want to determine if the government should be relatively more cautious in the high or low productivity region. Our results indicate that setting too high a minimum wage in the low-productivity region while setting too low a minimum wage by the same percentage points in the high-productivity region is better than the reverse.

2.4.4 Taxes and Transfers

In the previous sections, we study the optimal minimum wage without taxes, which is an "optimal suboptimal policy." When the federal government has access to linear income taxes (a single income tax rate and a demogrant), it will always prefer linear taxes to a combination of a minimum wage and linear taxes in our calibrated model. The proof of this numerical result in our framework is presented in Appendix B.1. Lee and Saez (2012) prove in a model without migration, when workers are not mobile across skill levels and non-linear taxes and transfers are available, then it is not optimal to also have a binding minimum wage. This is analogous to the no mobility

[31] The diagonal of this figure does not give the optimal central uniform policy because regions have different competitive wages. Since productivity in region 2 is much lower than 1, the optimal uniform minimum wage is found by looking across the top row.

case of our model discussed above. In a search-and-matching framework, Lavecchia (2019) finds that both a minimum wage and optimal taxes are preferred only if the government has very strong redistributive preferences.[32] As shown in the previous sections, the migration externality from any policy that restricts low-skilled labor supply and induces migration is worse for the local governments than the federal, and income taxes are no exception. This is intuitive in light of the literature which tends to find that redistribution is better handled by the federal government (e.g. see Gordon and Cullen, 2012).

If the government taxes income, then low-skilled workers will restrict their labor supply, impacting the marginal product of all workers and lowering total output in the economy. High-skilled labor supply is not perfectly inelastic to each jurisdiction, but it is perfectly inelastic to the economy and the federal government. The government trades off redistribution with efficiency. A binding minimum wage will increase low-skill labor supply, but decrease labor demand, while an income tax will decrease labor supply and have no effect on demand, potentially allowing the policies to be complements. While the optimality of jointly specifying a minimum wage and linear income tax is uncertain, we show that under the assumptions of our calibrated model, the federal government would never implement a binding minimum wage if linear taxes are available.

Although it is not optimal to have both a binding minimum wage and a federal

[32]Redistributive policies may also have an insurance component that high-skilled workers care about (Hoynes and Luttmer, 2012). If a high-skilled worker was separated from her job due to a shock, and then was able to find employment at the minimum wage rather than the low-skilled competitive one, the insurance value would be positive. However, in our static model with fixed worker skill levels, the high-skilled are made weakly worse off from a minimum wage increase.

income tax, we incorporate this into our model and recompute the equilibrium to compare the policies. Our simulations confirm that the government does not want to set a binding minimum wage when it sets the income tax optimally. In our baseline symmetric calibrated model, we find that federal government would choose to set a tax rate of 24% and no binding minimum wage in both jurisdictions, as expected. The policy dramatically increases unemployment to 35.5% of the low-skilled workers from 26.2% under the optimal minimum wage. The restricted labor supply leads low-skilled wages to rise to 0.390, and high-skilled wages to fall to 0.616. The government trades off its redistribution incentive with decreased low-skilled labor supply that leads to increased unemployment and falling high-skilled wages. As the income tax increases, the marginal (zero consumption) low skilled worker exits the labor force. This raises the marginal product of all the still employed low-skilled workers, thus acting very much like a minimum wage, except now the government is also able to redistribute some of the additional income to the unemployed.

The optimal linear income tax increases total welfare over the optimal minimum wage by 2%. Compared with the optimal minimum wage model, where the utility of the unemployed is 0 by construction, the utility of unemployed is 0.126. The low-skilled worker with the lowest working cost sees her utility increase by 16.2% under the optimal linear income tax compared with a minimum wage, which is a lower bound on the welfare increase for low-skilled workers. High-skilled workers see their welfare fall by 5.1%, as the government redistributes total wealth. If the federal government were able to set a different tax rate on each skill type, it would tax low-skilled at 73% and the high-skilled at 100%, leading to a welfare gain of

8.5% over the optimal minimum wage. This result is driven by the fact that the high-skilled workers supply labor inelastically to the economy. Since the government can use that revenue to lessen the harm of unemployment, it then sets a much higher tax rate on the low-skilled as well.

Throughout this section, we assumed that the governments tax income. However, if the tax is on consumption, equal to income minus cost of working and migration, then there is no labor supply response, and a federal government with a concave social welfare function would optimally set a leveling tax so that all workers have the same utility. With a leveling tax, a minimum wage is suboptimal as the induced unemployment lowers total output in the economy that can be redistributed. A utilitarian government would have no desire to redistribute through a consumption tax since the total output consumed would not change.

2.5 A Quantitative Model of Minimum Wage Setting

Our two region model demonstrated that decentralized and centralized minimum wage setting are strategic complements and that the extent to which they are depends on regional heterogeneity and mobility. While this approach highlights important mechanisms, we now extend the model to the case of U.S. states to quantify the welfare gain from the tiered policy compared to other minimum wage setting authorities.

2.5.1 Model

We consider a spatial equilibrium model of local labor markets as in Rosen (1979) and Roback (1982) with two imperfectly substitutable types of labor, and state and federal governments. High-skilled workers in state j have indirect utility given by

$$V_j^h = w_j^h + \phi_j + \kappa \epsilon_j \tag{2.10}$$

where w_j^h is the high-skilled wage rate in j, ϕ_j is a state j parameter that captures amenities and costs of living, and ϵ_j is an i.i.d. type I extreme value idiosyncratic preference for j, with standard deviation κ. High-skilled workers pick the location that maximizes their utility, which implies that the probability a worker lives in j is

$$p_j = \frac{\exp\left(\frac{w_j^h + \phi_j}{\kappa}\right)}{\sum_m \exp\left(\frac{w_m^h + \phi_m}{\kappa}\right)} \tag{2.11}$$

Immobile low-skilled workers can be employed or unemployed. If a worker i is employed in state j, she pays working cost $\theta_{ij} \sim \mathcal{U}(0, \bar{\theta}_j)$ and receives wages w_j^l. Therefore, her indirect utility is

$$V_j^l = \begin{cases} \phi_j & \text{if unemployed} \\ w_j^l - \theta_{ij} + \phi_j & \text{if employed} \end{cases} \tag{2.12}$$

Low and high-skilled workers in the same state face the same cost of living and receive the same amenities.

Each state has one firm that hires low and high-skilled workers with production function

$$f_j(l,h) = \gamma_j(\alpha_j l^\rho + (1-\alpha_j)h^\rho)^{1/\rho} \tag{2.13}$$

This specification differs slightly from Equation 2.8 in that it allows both γ and α to vary across markets. The government's social welfare function is given in Equation 2.9.

2.5.2 Data and Calibration

We calibrate the model to the continental U.S. states in 2015. In 2013 and 2014, Congress considered the Minimum Wage Fairness Act, which would have raised the federal price floor in 2015 for the first time since 2009. After this failed, President Obama continued to push for an increase through 2015, although Congress did not act. Since the federal government considered changing the minimum wage, but chose not to, we take the 2015 policy to accurately reflect the government's preferences for redistribution. This allows us to back out ν for the federal government from the observed equilibrium.[33] While there has been a continued discussion of increasing the minimum wage, the federal government has not considered changing it since. We

[33] More specifically, ν is calibrated so that \$7.25 is the welfare maximizing central minimum wage policy, holding the state policies fixed at their 2015 levels. This implies a Nash equilibrium where the federal government does not find it beneficial to deviate given the strategies of the states. However, the governments may be more sophisticated and know that changes in federal policy will also impact the state's optimal policies. We address calibrating ν for all states in Section 2.5.5.

therefore use the 2015 American Community Survey (ACS) to estimate the model parameters.

From the ACS, we observe individual level wages by education and labor force attachment in each state. This gives the number of the low (high school equivalent) and high-skilled (college graduates) employed workers, as well as their average wages, the number of unemployed, and the number not in the labor force. Since many high school equivalent workers earn more than the minimum wage, the model implied high-skilled wage equals the state's minimum wage times the high-skilled/low-skilled wage ratio in the data. $\bar{\theta}_j$ in each state is picked to match the state's high school equivalent labor force participation rate in the data.

Taking log of Equation 2.11, we first estimate κ given the observed population shares and high-skilled wages in the data from the following regression equation:

$$\log(p_j) = \beta_0 + \beta_1 w_j^h + \upsilon_j \tag{2.14}$$

where $\beta_1 = 1/\kappa$. However, endogeneity from the correlation of amenities and wages as well as measurement error will lead to a biased estimate of κ. We therefore instrument for the high-skilled wage with the number of low-skilled residents. The model provides intuition for it as a valid instrument. First, the marginal product of the high-skilled depends on the number of low-skilled employed in a state, and so it should be correlated with the total stock of low-skilled. We confirm this in the middle column of Table 2.4. Additionally, low-skilled workers are immobile, and so the total number is exogenous in the model. We find that κ is 1.048. As expected, the estimate is slightly larger than the value of 0.717 from Suárez Serrato and Zidar

Table 2.4: High-Skilled Migration Elasticity

	OLS $\log(p_j)$	First-Stage High-skilled Wage	IV $\log(p_j)$
High-skilled Wage	0.313***		0.954***
	(0.060)		(0.255)
Low-skilled Workers/100000		3.385***	
		(0.681)	
Observations	48	48	48

Note: The leftmost column (OLS) reports the estimated effect of average high-skilled wages in a state on the log of the percent of high-skilled workers in that state, using data from the 2015 American Community Survey. The middle column reports the first-stage for the IV regression in the rightmost column. Robust standard errors are reported in all columns. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

(2016), who conduct a similar analysis on county groups, a smaller unit of geography, using a Bartik instrument approach. They similarly find an OLS/IV ratio for κ of just above 3.

After estimating κ, we can recover ϕ_j up to a normalization that $\min \phi_j = 0$. Since workers' indirect utility is a function of ϕ_j, the normalization together with our estimate of ν determine the concavity of the social welfare function. Under the functional form assumptions, this does not affect high-skilled location choices and is chosen so that all workers have non-negative utility. Since only $\phi_j + \nu$ is identified, an alternative normalization would only affect the magnitude of ν and not the results.

The calibration of the state specific production function follows Section 2.3.2. First, we set $\rho = 0.286$ following Katz and Murphy (1992). We calculate α_j to match the wage gap between high and low-skilled workers and γ_j so the marginal product of the low-skilled employed matches the state's minimum wage. Given the model's parametrization, the values perfectly match the targeted moments. [34]

After recovering the other parameters, we finally back out the federal government's redistribution preferences ν so that the observed policy of \$7.25 maximizes total social welfare in the economy. We aim to minimize $|7.25 - w^*(\nu)|$, where $w^*(\nu)$ is the optimal minimum wage as a function of the concavity of the social welfare function. Table 2.5 summarizes the parameter estimates. The model is unable to perfectly match the current federal minimum wage for any progressivity shifter, although the residual is only about 1 cent. In the context of the model where states are heterogeneous, the federal government and 21 states should not coordinate on the same policy. However, \$7.25 may be salient at a national level as a reference point, or the federal government may understate regional heterogeneity because of a lack of information.

Table 2.5: U.S. Model Calibration Results

National Parameters	Estimate	Residual
ρ	0.29	-
ν	4.05	0.012
State Parameters	Mean	Std. Dev.
$\bar{\theta}_j$	14.65	1.28
α_j	0.43	0.06
γ_j	21.44	2.72

[34] For this calculation, we calculate the state's wage gap as the ratio of college graduates and high school equivalent or less in the ACS for employed workers between ages 16 and 70. Since the minimum wage is a floor, the ratio of high-skilled wages to the minimum wage would overstate their relative marginal product in the model.

2.5.3 Welfare and U.S. Minimum Wage Setting Authority

Does joint state and the federal tiered minimum wage setting in the U.S. increase aggregate social welfare over centralized policy alone? While our previous theoretical results suggest there should be positive gains, in the application, the sign is ambiguous since state are heterogeneous in their preferences for redistribution as well as their productivities. This implies that the most productive states do not always set the highest minimum wages. The correlation between the 2015 state policies and γ_j is only 0.53.[35] If a relatively low productivity state sets too high a minimum wage relative to what is binding in others, too many high-skill workers move. Since the state is of relatively low productivity, this disproportionately harms its low-skilled workers by increasing unemployment. Given this tension created from differences in preferences for redistribution, it is possible for optimal centralized policy to lead to higher total welfare from the perspective of the federal government than the tiered system.

To measure the potential gain or loss, Panel (a) of Figure 2.5 plots total welfare in the economy under both minimum wage setting authorities. The solid line shows how welfare in the tiered system changes as a function of the federal minimum wage. Based on the calibration, when the states' minimum wages are held at their 2015 levels, the federal government optimally sets a minimum wage of nearly \$7.25, which confirms our calibration of ν is successful. The dashed line in the figure shows

[35] This calculation may overstate heterogeneity in preferences because states whose optimal policies are below \$7.25 have no incentive to set their own. Using the 2015 effective minimum wages, the correlation is 0.85, but for similar reasons this likely overstates heterogeneity. Heterogeneity in other production parameters matters as well. The correlation between α_j and the state policies is -0.36.

welfare as a function of centralized policy. Without the state policies, the federal government optimally chooses a lower minimum wage of \$6.90. Panel (a) of Figure 2.6 shows the relationship between tiered minimum wages and low-skilled wages under centralization.[36] This result is in line with the previous findings that state and federal polices are strategic complements. Higher state minimum wages reduce the high-skilled migration externality and allow the federal government to engage in more redistribution by setting a higher wage floor itself.

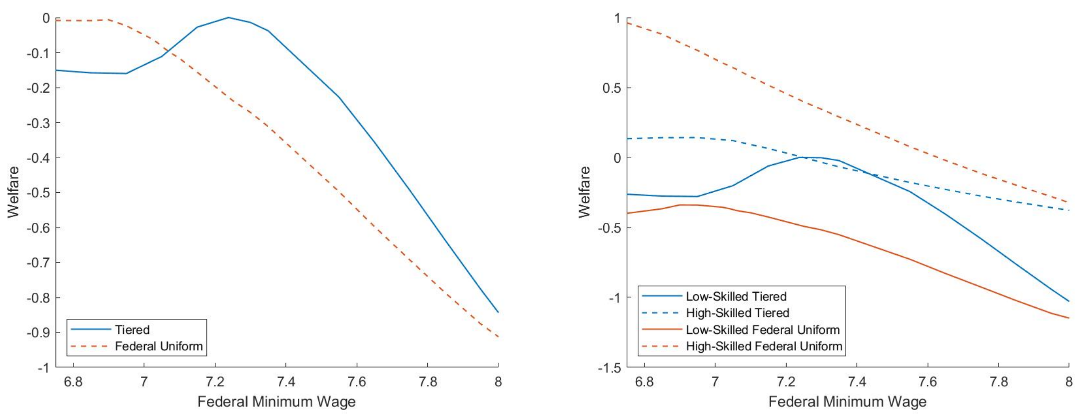

(a) Aggregate Welfare (b) Welfare by Worker Type

Figure 2.5: U.S. Welfare and the Federal Minimum Wage

Note: Panel (a) plots total social welfare as a function of the federal minimum wage under tiered and federal uniform minimum wage setting authorities. Welfare is normalized to be 0 under optimal tiered policy setting with a federal minimum wage. Panel (b) splits welfare into the contributions from low and high-skilled workers separately. Welfare is normalized to be 0 for each worker type under optimal tiered policy setting.

Optimal centralized policy has a welfare loss of about 0.006 percent compared to

[36]Based on our calibration, the tiered minimum wages are binding in all states, while our simulations predict that the optimal federal uniform minimum wage is only binding in 17 states.

the tiered system.[37] Although the aggregate welfare change is small, state and federal minimum wages together are better able to redistribute to the low-skilled employed, as shown in Panel (b) of Figure 2.5. The figures also show that suboptimal policy in the form of low wages, for example due to information frictions discussed earlier, has a higher social welfare cost under tiered policy setting. This is because many states set their own policies at or above the federal level. High state policies become more costly when the federal price floor falls. The opposite is true when the federal governments sets too high of a minimum wage.[38] However, as the federal policy increases, the tiered and central uniform systems become more similar since only the higher of the state and national policy is implemented. If the federal policy were above \$9.47, the two would be equivalent.

While the current U.S. minimum wage setting authority yields higher welfare than centralized policy alone, it may be far from the social planner's solution. The 48 states are heterogeneous in their productivity and low-skilled working costs, but 21 have effective minimum wage policies of \$7.25. This is in part because states face a different trade-off than the federal government as emphasized in equations 2.2 and 2.5, and because some states may be more or less progressive. Panel (b) of Figure 2.6 shows the relationship between the observed tiered policies, and those under federal non-uniform. States below the 45 degree line have too high minimum wages under tiered policy setting. When not limited to a single policy, the government can better differentiate among states. Those subject to the \$7.25 minimum wage in the tiered

[37]The competitive equilibrium has a welfare loss of 0.021 percent.

[38]When interpreting the value of the tiered system across different minimum wage levels, it is important to note that the analysis in this section holds state policies fixed. If the federal minimum wage were to fall (rise), states should lower (raise) theirs as well.

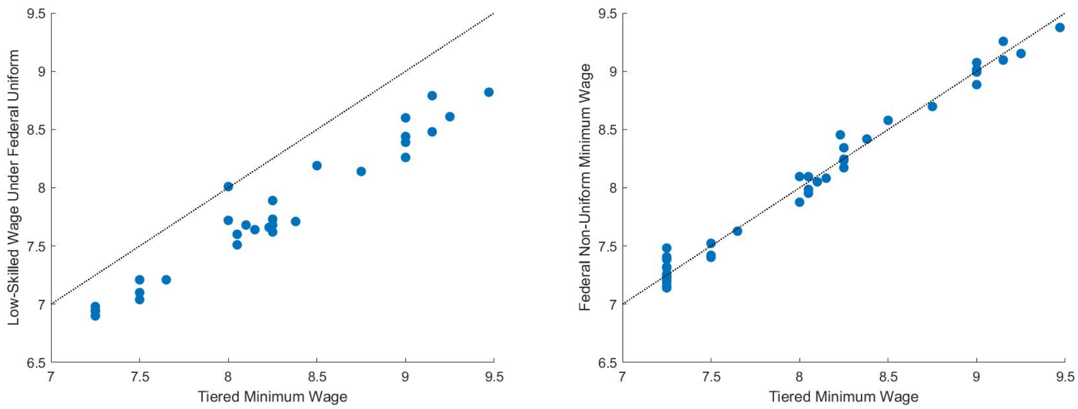

(a) Centralization (b) Federal Non-Uniform

Figure 2.6: Comparing U.S. Minimum Wage Setting Authorities

Note: Panel (a) plots the low-skilled wages in each state under tiered and federal uniform policy setting as well as a dotted 45 degree line for reference. Panel (b) similarly shows the relationship between tiered policy and the optimal federal non-uniform.

system have federal non-uniform policies that range from \$7.14 to \$7.53. The federal government also sets slightly higher minimum wages for states with initially high policies in this case since it is able to better internalize the spillovers from high-skilled migration. Overall, current U.S. policy closely matches the social optimum predicted in this model with only a 0.05 percent welfare loss.

2.5.4 State Heterogeneity and Regional Redistribution

The minimum wage setting authority potentially impacts redistribution across states in addition to skill groups. Under centralization, the federal government has one policy lever that only directly affects the lowest productivity states. Higher federal uniform policies will therefore redistribute to more productive states. To some extent, this can improve total welfare because the more productive states are larger and have more low wage workers and unemployed, but a progressive government does not value transferring wealth from low to high income areas. Moving from centralization to the tiered system, there is an ambiguous effect on activity across states that depends on the relative progressivity of different governments.

To measure how states would be differentially affected under centralization compared to the current tiered system, the left three panels of Figure 2.7 show low-skilled employment, and high-skilled location decisions and wages change under these two regimes. Under centralization, Panel (a) shows that low-skilled employment is higher in most places. It mostly, but not exclusively, increases in states with low tiered minimum wages. The states experiencing the biggest decreases from centralization, represented by points far below the 45 degree line, are all relatively small, which reflects the government's trade off of redistributing to large workforces and low wage places. Without state minimum wages to reduce the high-skilled outmigration externality, the low-skilled workers in the lowest productivity areas are hit the hardest.

There is very little state heterogeneity in the impact on high-skilled wages because these workers are imperfectly mobile. If states became too different, some workers will move. Panel (c) shows that wages are close to the 45 degree line, but nearly uniformly above. With only a single low federal minimum wage of $6.90, aggregate low-skilled employment increases in most states and drives up high-skilled wages. Centralization alone has only a small effect on the high-skilled migration externality and therefore allows for less redistribution from high to low-skilled workers, leading to higher average wages. However, since low-skilled employment decreases in some states, high-skilled workers' marginal products fall and they leave those affected low productivity areas. For this reason, Panels (a) and (e) look similar.

The regional distributional consequences of federal non-uniform over tiered policy setting are also ambiguous because states' preferences may be very different from the federal government. However, the optimal tiered minimum wages closely match the federal non-uniform, and so the right three panels of Figure 2.7 also show very little difference in low-skilled employment, and high-skilled location choices and wages.

2.5.5 A Brief Note on Decentralized Minimum Wage Setting

Competition between states for mobile high-skilled workers leads to low minimum wages under decentralized policy setting. Local governments are therefore less able to redistribute income from high to low-skilled workers. The amount of redistribution under decentralization compared with other systems importantly depends on regional heterogeneity in ν. Following the previous calibration for the federal government, we aim to estimate $\nu_j, j = \{1, ..., 48\}$, such that the observed 2015 minimum wage in state j maximizes j's welfare holding the other minimum wages fixed, so that the 2015 tiered policies are a Nash equilibrium. Since only 27 states set a binding minimum wage above \$7.25, ν_j is not point identified for every state.

In the calibrated model, there does not exist ν_j that replicates the state's 2015 policies. Under our assumptions and calibration, each state would prefer a smaller minimum wage, holding all other states' and the federal government's policies constant. That is, there does not exist a progressivity level such that the current state minimum wages are best responses. There are several possible reasons why the model

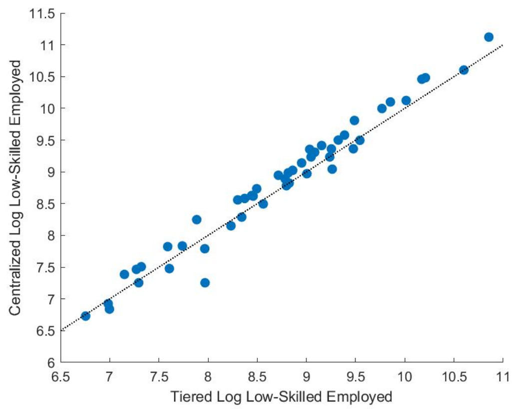

(a) Centralized Low-Skilled Employment

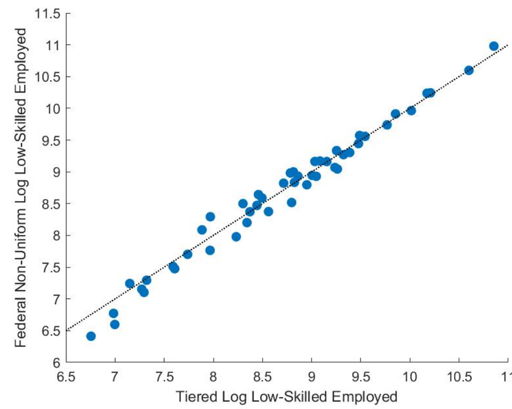

(b) Federal Non-Uniform Low-Skilled Employment

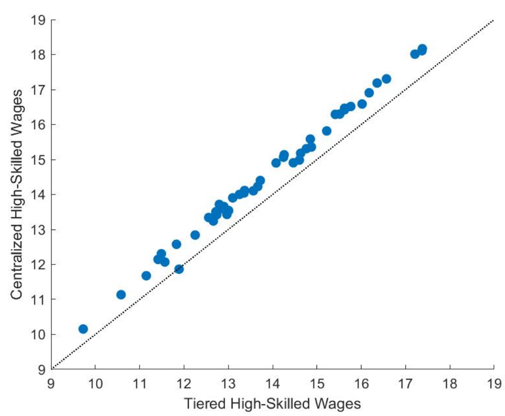

(c) Centralized High-Skilled Wages

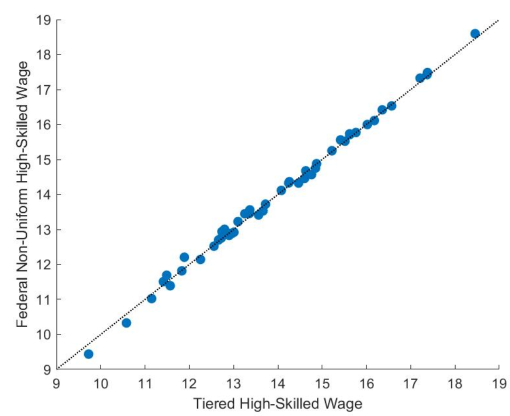

(d) Federal Non-Uniform High-Skilled Wages

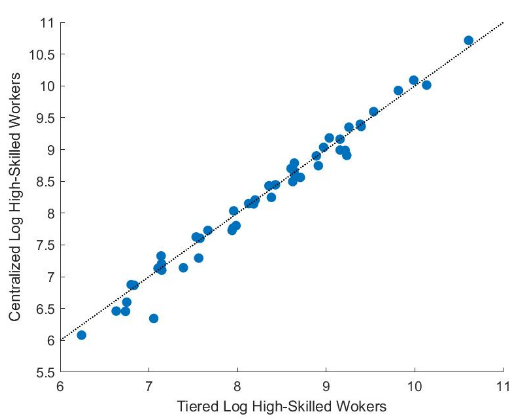

(e) Centralized High-Skilled Residencies

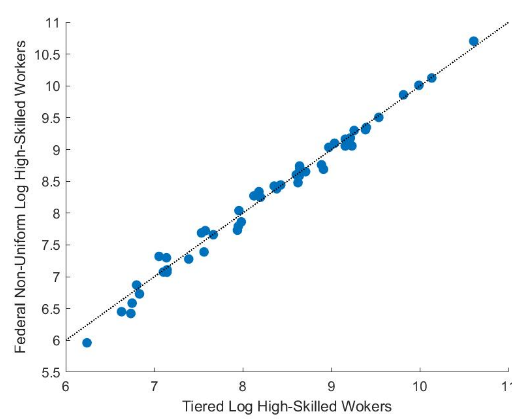

(f) Federal Non-Uniform High-Skilled Residencies

Figure 2.7: The Impacts of Minimum Wage Setting Authorities Across States

Note: The left three panels (a, c and e) compare tiered and centralized policy setting, while the right three (b, d, and f) compare the tiered and federal non-uniform systems. All figures contain a a dotted 45 degree line for reference.

is unable to match the data. First, governments may be more sophisticated. If a states believes that others will set lower minimum wages too when it lowers its own policy, then there are smaller gains from doing so. The main advantage of a lower minimum wage is to attract more high-skilled workers, but the amount is a function of low-skilled wages everywhere. Second, states could face political frictions. They may not be able to set their policies on a continuum but are instead restricted to a small grid. The observed policies may be optimal on the grid but not globally. Finally, the model is potentially misspecified or, at least, states may have different beliefs about the implied elasticities. In particular, the model implications for the effect of the minimum wage on low-skilled unemployment and high-skilled wages may not be true, for example if workers have monoposony power or the high-skilled have larger moving costs. More work is needed in the future to better understand state competition and minimum wage setting.

While we are unable to compare the decentralized and tiered equilibria explicitly, we aim to measure the relative trade-off of local and federal policy setting, holding preferences for redistribution fixed. Only 23 of the 48 states set binding minimum wages under decentralization when all governments have the same value of ν equal to 4.05. Without a binding policy in many states from the federal government, the high-skilled migration externality is large and states optimally set lower minimum wages. Figure 2.8 shows that the low-skilled wages in every state are lower under decentralized policy setting compared to the current U.S. policies. However, total welfare under this version of decentralization is only 0.009 percent lower than the tiered system, or slightly worse than the centralized policy.

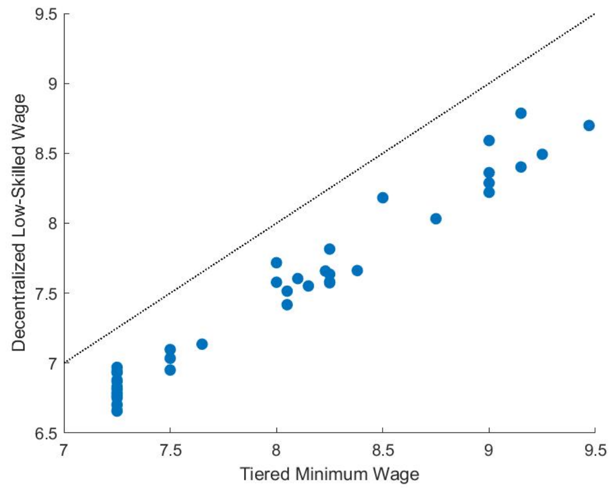

Figure 2.8: Decentralized Minimum Wage Setting

2.6 Conclusion

This paper provides a framework to think about the interaction of minimum wage setting and federalism. Previous optimal minimum wage research focuses on one jurisdiction, which misses the implications of government competition. We first build on the theoretical model of Lee and Saez (2012), extending it to a context in which high-skilled labor is mobile across regions. We present sufficient conditions for the desirability of a binding minimum wage, from the perspective of both a local and central government. Though minimum wages are not optimal in the presence of even a linear tax system in our model, their prevalence in the real world warrants consideration from an optimal policy perspective.

We calibrate a two-region model of a federation and compare minimum wage policies across four types of government structures: decentralization, centralization with uniform policy, a U.S.-style combination of both, and a social planner. In our baseline model, the desirability of centralized policy setting is increasing in mobility; when mobility is shut down completely, decentralization obtains the social planner's solution. However, when all movers face positive moving costs, and regions are heterogeneous, then the decentralized outcome leads to higher total welfare than federal uniform policy. Furthermore, suboptimal policy takes the form of minimum wages that are too low, such that the tiered system we observe in the U.S. is weakly preferred to either decentralization or centralization exclusively. In the tiered system, the central uniform policy lessens the externalities from horizontal competition by local governments leading both types of governments to set policies more closely aligned with the social planner. Decentralized and centralized policy setting are strategic complements. When we extend our framework to U.S. states, we find that the tiered system closely matches the social planner's solution. Our results indicate that higher levels of government should not forbid lower levels from setting their own minimum wages. These results are consistent across social welfare functions, although the effect of heterogeneity may depend on the progressivity of the social welfare function.

Though the U.S. system performs well relative to other potential systems, its ability to efficiently redistribute income still falls short of a simple system of linear taxes. From the perspective of our calibrated model, the existence of minimum wages and redistributive taxes in the U.S. remains a puzzle. Further research remains to be done on the types of models, economic or political in nature, which might give rise to the joint optimality of both types of policies. It is likely that the income tax is not set optimally due to political frictions, which lead to demand for other redistributive policies. For now, our results suggest the policy makers should pay attention to the mobility implications of minimum wage policies, as well as their interactions with the existing tax system.

CHAPTER III

Federalism, Fiscal Savings, and Information Asymmetries

3.1 Introduction

The United States is an important example of fiscal federalism and fiscal unionization. One of the main purposes of fiscal unionization is thought to be the ability of a fiscal union to smooth idiosyncratic risks among its member jurisdictions (Farhi and Werning, 2017b). The U.S. is a strong version of a fiscal union: the central government enjoys full tax and spending power. One might expect that such a fiscal authority would fully insure subnational governments against risk; however, this may not be the case if the central government itself faces significant frictions. The behavior of state governments and the federal government in the U.S. provides a useful case study to examine how fiscal policy is conducted by members of a strong federation.

Fiscal policy at the state level is an important component of the economic environment in the U.S. Spending by subnational governments totaled 2.85 trillion dollars in 2016, representing almost 15 percent of GDP; state government spending alone made up 28 percent of total government spending.[2] In addition, the makeup of state government spending is fundamentally different than that of the federal government. While federal government spending primarily constitutes defense, social security, and interest payments, state governments tend to spend primarily on education and public welfare (this includes Medicaid). Constraints faced by state governments differ substantially from the federal government, as well; state governments interact with different tax bases, and most face deficit limits of varying strengths.

In light of these observations, policy analysts should not expect state-level fiscal policy to behave like federal policy. One major way in which state governments differ from the U.S. government is in the use of Budget Stabilization Funds, or "Rainy Day Funds", to improve funding for public programs in times of fiscal distress. All fifty state make use of rainy day funds; in 2017, Montana became the fiftieth state to establish such a fund.[3] The median balance of these funds in 2016 was 477 million dollars, having built up significantly since their depletion after the Great Recession. The presence and size of these funds indicates that fiscal savings may be an important way in which U.S. states can enact fiscal policy over the business cycle, even in the presence of (sometimes strict) deficit limits. The savings motive for states motivates

[2]BEA, 2018.

[3]https://www.pewtrusts.org/en/research-and-analysis/articles/2018/08/29/states-make-more-progress-rebuilding-rainy-day-funds

the question of this paper: *How do U.S. states' savings respond to business cycles and federal fiscal policy, and what kinds of models might help us understand this behavior?*

3.1.1 Overview of Results

The broad results are as follows. I present three stylized facts about state public finances over the business cycle: states engage in precautionary savings, transfers from the federal government respond less to local shocks than national shocks, and states whose cycles are less correlated with the national cycle tend to save more. I obtain these facts using data on state government finances from two different sources, under four possible definitions of state government savings; the results are robust to the choice of definition for state government savings.

I interpret these facts as evidence that state government savings behavior is driven both by balanced budget rules and by the transfer policies of the federal government. To illustrate, I propose a model of a small open endowment economy in a fiscal union; i.e., an economy under both a regional government and a central government, much like a U.S. state. The regional government faces a debt limit, but the central government observes the state of the world with a noisy signal. The central government, therefore, is not able to fully insure the regional household against adverse shocks, and the regional government must build up a stock of savings.

I calibrate the model to U.S. data, and find that it fits qualitative features of the data quite well. The implied noise shock in the central government's signal is almost three times as large as the real economic shocks, indicating significant frictions to optimal policy making at the centralized level. The baseline calibration of the model implies a 2.1 percent welfare loss relative to the social planner solution, indicating a sizeable influence of the frictions in the model on household utility. Two variations to the information structure of the central government yield very similar results to the baseline model. I conclude that information (or other political) frictions at the centralized level of policy making may create large deviations from the socially optimal policy; this implies states should consider actively pursuing countercyclical fiscal policies. Additionally, the results suggest that significant barriers to cross-region risk sharing may still exist, even in the presence of a strong fiscal union.

3.1.2 Literature and Outline

This paper connects to a number of distinct strands of literature in economics research. On the empirical side, Hines (2010) characterizes the behavior of state-level spending over the business cycle, arguing that small and large states behave differently in response to macroeconomic conditions. Nakamura and Steinsson (2014b) estimate government spending multipliers for U.S. states, using military spending shocks from the central government. Owyang and Zubairy (2008) find heterogeneous effects of fiscal stimulus in different states and regions of the union, depending on regional makeup. In addition to these, Chodorow-Reich (2017) summarizes findings from empirical literature on the effects of fiscal policy at subnational levels.

In the public finance literature, economies with fiscal federalism have been much studied (Oates, 2008). This literature tends to compare public goods provision at the local level to that at the central level. These models are static, however, and do not say much about cyclical policy. Furthermore, they tend to compare two methods of public spending rather than having both local and central governments spending at the same time. For the purposes of this model, I abstract from the choice of public goods and study the optimal taxation behavior of states subject to an exogenous stream of public goods. This is the approach taken by optimal fiscal policy papers such as Schmitt-Grohe and Uribe (2004), Chari and Kehoe (1999), and Bhandari et al. (2017).

This paper also draws on the precautionary savings under credit constraints literature. Aiyagari (1994) is a key early example of this literature. Additionally, papers such as Bianchi and Mendoza (2018) and Durdu, Mendoza, and Terrones (2009) analyze savings decisions and credit contraints explicitly in small open economy models; I take this approach, though the optimizing agents of interest in my model are governments, while households are passive. Bhandari et al. (2017) consider explicitly fiscal policy for a credit constrained government with access to imperfect markets. I contribute to this literature by considering the problem of a government conducting fiscal policy under debt limits, but in a small open economy when there is a "higher" level of government which is also conducting policy, i.e., in a fiscal union or federation.

Other papers have studied fiscal policy in the context of a union of economies. Beetsma and Jensen (2005), Debrun (2000), and Ferrero (2009) consider optimal fiscal policy for economies in a monetary union. Luque, Morelli, and Tavares (2014) and Farhi and Werning (2017b) consider the desirability of explicit fiscal unionization. This paper considers the strongest version of a fiscal union, in which the central government has full power to tax and spend in addition to the regional governments; this is the system in place in the U.S.

Models in which a central government might not have the same access to information as local governments have been explored in other contexts. Bordignon, Manasse, and Tabellini (2001), for example, consider optimal redistribution policy when information is asymmetric. Silva and Cornes (2000) examine information asymmetries in the context of interregional transfers and public goods provision. In his survey of the future of the fiscal federalism literature, (Oates, 2005) also mentions information asymmetries between the different levels of government as a feature of federal systems. This paper contributes to the information asymmetries literature in fiscal federalism by applying the idea to a dynamic model of fiscal policy.

The rest of the paper proceeds as follows. Section 3.2 presents the main stylized facts of the paper. Section 3.3 introduces the information model with which the stylized facts are interpreted. Section 3.4 calibrates, analyzes, and interprets the model. Section 3.5 concludes.

3.2 State Government Savings: Three Stylized Facts

This section lays out three stylized facts apparent in the data on U.S. state government savings. First, state governments overwhelmingly engage in precautionary savings: savings are positive and procyclical. Second, transfer receipts from the federal government are countercyclical, but depend more on the aggregate U.S. economy than on a state's idiosyncratic business cycle. Finally, states whose business cycles are less correlated with the national business cycle tend to save more than states experiencing fluctuations more in step with the aggregate cycle. This section first describes the data sources and definitions; the second subsection presents the three facts.

3.2.1 Data and Descriptions

3.2.1.1 Data Sources

Several sources are used to assemble the data for this part of the paper. Data on rainy day funds and end-period balances for state governments are obtained from the National Association of State Budget Officers' - hereafter, NASBO - "Fiscal Survey of the States." I use the spring edition of this semiannual report from 1979 to 2017 to obtain data from previous years which is self-reported by states and collected by NASBO. Due to heterogeneity in the structure of BSFs, some state governments do not report BSF balances separately from end-year balances, rendering analysis of rainy day funds alone a bit hairy; I discuss this below when considering all possible definitions of "savings."[4]

Data on state government revenues, spending, and debt holdings comes from the U.S. Census Bureau's *Census of Governments*. While the full sample of local governments is only administered every five years, all state governments are included in the limited survey taken every year, such that yearly observations from 1970 to 2012 are available for every state. Other state variables of interest are provided at

[4]An appendix containing a full explanation of the data collection from NASBO reports is available upon request.

the yearly level on the website of the University of Kentucky's Center for Poverty Research. I estimate state-level recession dates using the Philadelphia Fed's state coincident index. National annual price level indices are obtained from the OECD.

3.2.1.2 Definition of Savings

In order to study the cyclical behavior of state government savings, some definition of "savings" is naturally required. Four potential definitions are available in the data; I choose to focus on a couple of them for ease of exposition. The first obvious definition of state government savings is the balance of the state's rainy day fund as reported to NASBO. While some amount of heterogeneity exists across funds, and not all states report their RDF balance separately from their general fund, budget stabilization funds are a useful metric due to their explicit purpose of preparation for adverse shocks. A second, and slightly more expansive, definition includes all end-year balances in a state's general fund; while such a measure will include unplanned revenue and spending shocks, it captures all rainy day fund activity and provides a consistent measure across states.

While the first two potential measures are taken from the NASBO reports, the other two are found in the U.S. Census Bureau's annual Census of Governments dataset. The third potential measure of state government savings is a state's net assets–cash and securities less debt outstanding–not including assets set aside for insurance purposes (pensions, etc.). The fourth measure is all of a state's net assets, including those in insurance-type funds.

My preferred measures of state government savings are measures two and three.

These measures, total balances in general funds (including rainy day funds) and net non-insurance assets, provide a nice balance between the ideal features of a savings measurement. They are consistent across states, relatively general, and include a good deal of long-term savings components. Importantly, however, the qualitative results are not altered by the choice of savings measure.

3.2.2 Three Stylized Facts

3.2.2.1 Fact 1: State Savings Are Positive and Procyclical

The first stylized fact I identify is the presence of positive and procyclical savings behavior on the part of state governments. Regardless of which measure of savings measure is observed, U.S. states mostly run positive balances. This is not in itself a surprising result; in fact, it is exactly what one might expect given the balanced budget requirement imposed on 49 of the 50 U.S. states.[5] Table 3.1 presents summary statistics for the savings measures of interest, both as a fraction of gross state product and as a fraction of general current state government expenditures. Clearly, states run positive levels of savings–0.5 percent or 4.3 percent, depending on the definition–on average, although some observations do record negative savings levels.

In addition to being overwhelmingly positive, state government savings also move with the business cycle. Figures 3.1, 3.2, 3.3, and 3.4 show how various percentiles of the distribution of savings across states move over the business cycle for the four measures of state savings, where the shaded regions indicate NBER recession dates. Clearly, savings balances build up in economic expansions and spend down

[5] NCSL, 2010.

Table 3.1: Measures of State Government Savings: Summary Statistics

Savings measure	Mean	Variance	Percentiles		
			50th	10th	90th
BSF over GSP	0.0038	0.0003	0.0012	0	0.0040
Gen. fund balance over GSP	0.0051	0.0003	0.0026	0.0002	0.0083
Net noninsurance assets over GSP	0.0427	0.0139	0.0175	-0.0258	0.1060
Net total assets over GSP	0.1763	0.0208	0.1600	0.0594	0.2823
BSF over expenditures	0.0228	0.0066	0.0102	0	0.0371
Gen. fund balance over expenditures	0.0379	0.0063	0.0242	0.0013	0.0776
Net noninsurance assets over expenditures	0.3367	0.5174	0.1604	-0.2329	0.9355
Net total assets over expenditures	1.5531	0.7419	1.4838	0.6304	2.4441

Note: Moments reported here are over all state-year observations. Data on budget stabilization funds and general fund balance come from the NASBO fiscal survey of the states, data on net assets come from the Census of Governments, and gross state products are obtained from UKCPR. Sample periods are as follows: BSFs from 1985-2016, balances from 1979-2016, both net assets series from 1981-2012.

in recessions; this is consistent with the stated purpose of RDFs, which are included in these measures. I interpret this behavior as being indicative of a precautionary savings motive on the part of state governments, induced by the presence of balanced budget rules and the desire of policy makers to smooth expenditures over the cycle.

As a supplemental example, I also plot the series of Kentucky's balances over GSP alongside its HP-filtered log GSP series in Figure 3.5. The cyclical behavior of Kentucky's balances seems acyclical in the 1980s; however, in the early 1990s they come more into line with what would be expected under a precautionary savings motive, building up in state level expansions and spending down during contractions. Notably, the fund doesn't simply respond to U.S. level decreases in output relative to trend; it experiences a decrease in the mid-1990s and recently in the mid-2010s, both corresponding to downturns in gross state product. Furthermore, note that the mid-2000s recession seems to begin earlier in the state, and the state's balances begin to respond accordingly before the U.S. as a whole fell into recession. Furthermore, over the entire sample the correlation coefficient of the two series is 0.40, further

Figure 3.1: Real State Government Rainy Day Funds

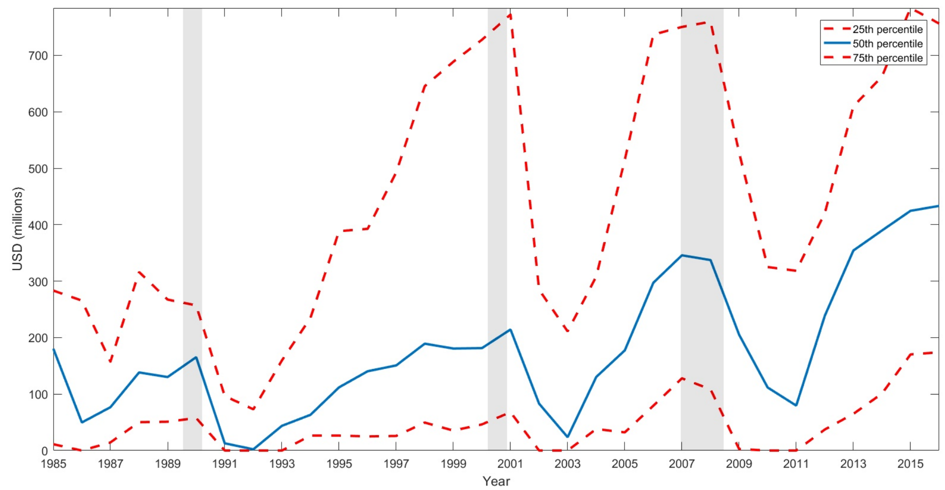

Figure 3.2: Real State Government Total Balances

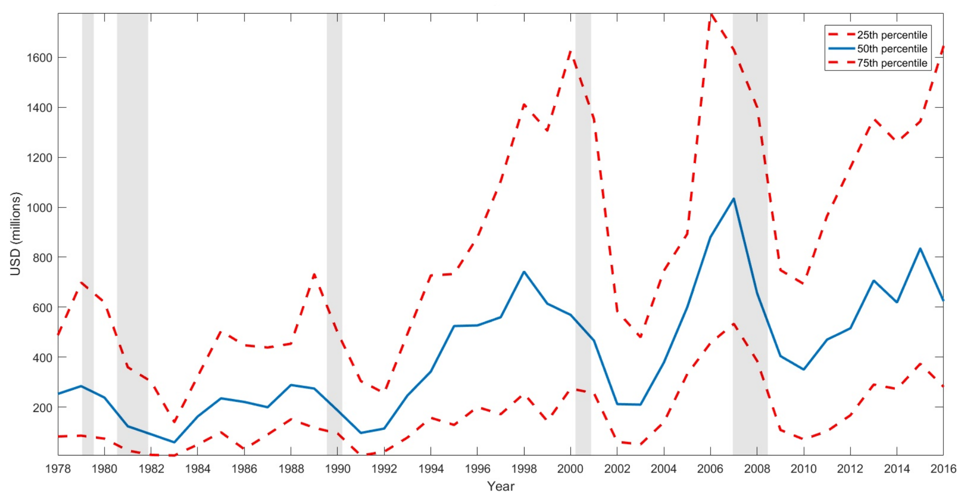

Figure 3.3: Real State Government Net Financial Assets (Not Incl. Insurance Funds)

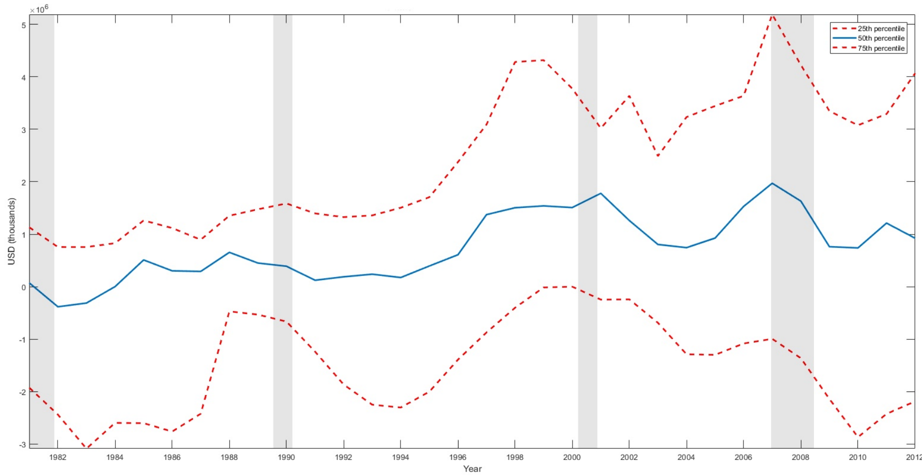

Figure 3.4: Real State Government Net Financial Assets

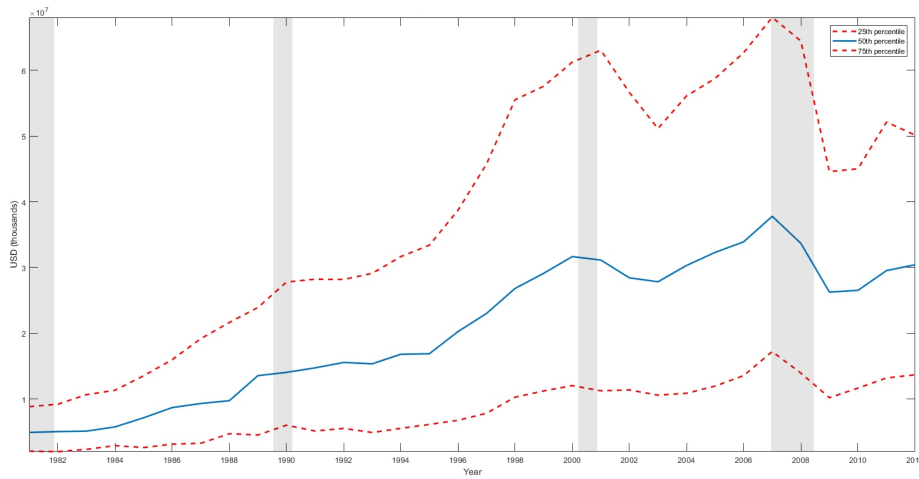

Figure 3.5: Kentucky State Government Balances and Business Cycle

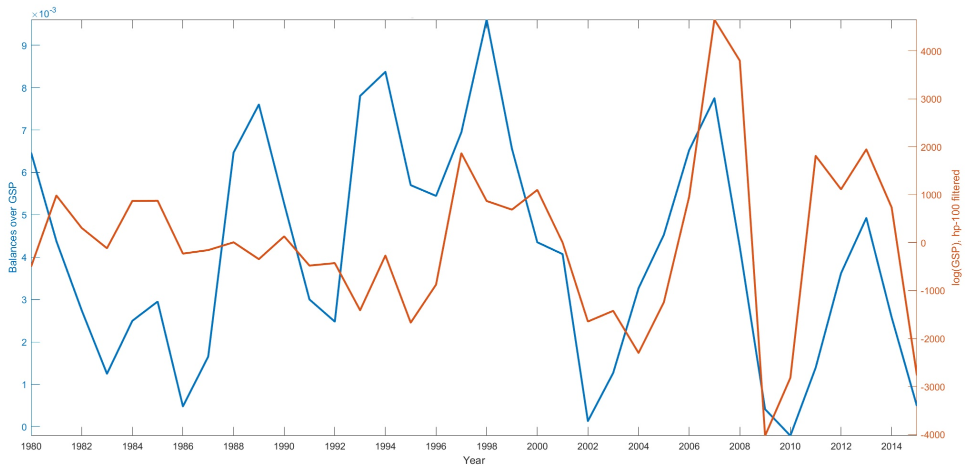

indication of significant precautionary savings behavior in Kentucky; similar behavior is observed in other states.

3.2.2.2 Fact 2: Transfer Receipts Are Countercyclical and Respond Heavily to National Cycle

The second stylized fact describes the behavior of state governments' transfer receipts from the federal government. These transfer payments are countercyclical, as one might expect, but respond quite differently to aggregate and idiosyncratic fluctuations. Specifically, a state government's transfer receipts from the federal government respond more strongly to the condition of the U.S. economy as a whole than to economic conditions within a state, running counter to the idea that fiscal unions serve to smooth idiosyncratic risk across members. In other words, Michigan might expect an increase in transfers when the rest of the country goes into recession, even if Michigan is expanding; conversely, Michigan may not expect as much revenue from the federal government when it is contracting, if the rest of the country is doing well.

In order to observe the relationship between business cycles and transfers from the federal government to state governments, I estimate the following equation:

$$\log(T_{it}) = 1 + \beta_1 y_{it} + \beta_2 y_{-i,t} + \Gamma X_{it} + \varepsilon_{it}. \tag{3.1}$$

In this equation, T_{it} represents the amount of transfers state i receives from the federal government in year t. y_{it} and $y_{-i,t}$ represent the cyclical components of state i's output and the sum of the other states' outputs, respectively. X_{it} is a vector of controls, including population and state fixed effects. There is, of course, a potential source of endogeneity in this regression: transfers from the federal government to state government i may have an effect on the local business cycle in state i, biasing the estimate of β_1 toward zero. In Appendix C.2, I instrument for y_{it} and $y_{-i,t}$ using a measure of monetary shocks, which I argue are plausibly exogeneous to state-level transfers, and show that the main result of this section is not affected; if anything, the difference is more stark.

Table 3.2 presents the output from a regression of state receipts from the federal government on population, cyclical GSP, and the cyclical component of the sum of the GSP of the other states; this corresponds to an estimation of Equation 3.1 with a population control and standard errors clustered at the state level. The response of federal transfers to a state respond more strongly to the cyclical component of the *aggregate* economy (less GSP of the state itself) than to the *idiosyncratic* cycle of the

Table 3.2: Determinants of State Government Receipts from Federal Government

Variable	log(pop)	log(GSP_i), cyclical	log(GSP_{-i}), cyclical
Coefficient	0.0061	-0.2083**	-0.5068***
(s.e.)	(0.0021)	(0.1016)	(0.1111)

Note: Results from a fixed-effects regression estimating Equation 3.1 of receipts on the explanatory variables with standard errors clustered at the state level. Observations include 46 states for which holes do not exist in the Census of Governments data from 1981 to 2012. * $p \leq 0.10$, ** $p \leq 0.05$, *** $p \leq 0.01$.

individual state. A one percent decrease in a state's own GSP relative to trend results in a 0.2 percent increase in transfer receipts; at the same time, a similar decrease in other states' GSP would yield a 0.5 percent increase in the state's receipts.

For a bit more insight into the composition of these transfer receipts, consider Figure 3.6. Clearly, the most significant component of federal government transfers to states is the "public welfare" category; this category includes funding for a wide range of public assistance programs that are administered at the state level, Medicaid being the largest among them. the second biggest category is for education funding, and the third for highways. The previous result about the response of federal transfers to the business cycle of the aggregate economy seems to be driven by the two biggest categories, public welfare and education, as both of these transfer categories exhibit the same pattern of greater response to the aggregate business cycle.

That public welfare drives a large part of the main result in this section might raise concern upon a first glance. The public welfare portion mainly consists of Medicaid payments, unemployment insurance, etc.; these transfers are typically formulaic, simply responding to claims made by low income individuals. Thus, it might seem strange to claim, as the model below does, that the result displayed in Table 3.2 could

Figure 3.6: Composition of Federal Transfers to States

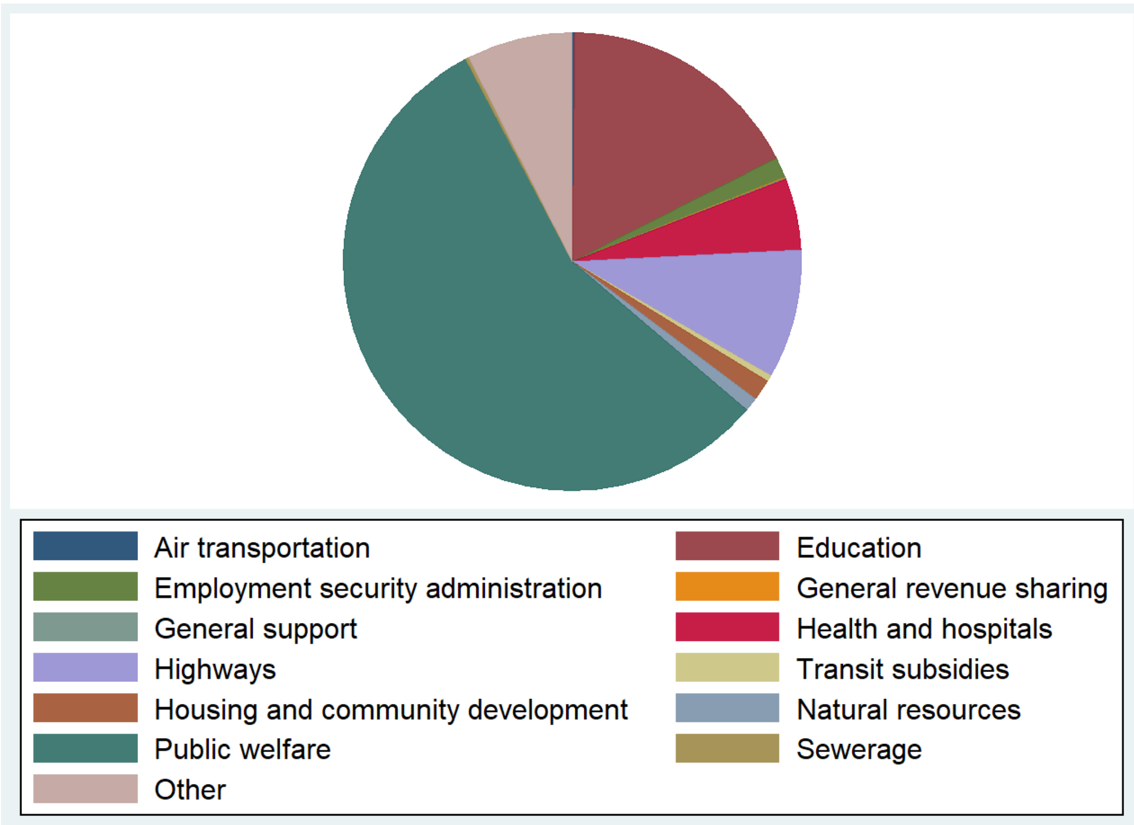

reflect some information or political friction in the fiscal policy system. However, while the business-as-usual operations of these transfers involve responses to demand, the federal government has shown that it is willing to step in and increase these transfers in a *discretionary* way in response to large business cycle shocks. The most recent example of such an action is in the American Reinvestment and Recovery Act (ARRA) of early 2009, which was the largest fiscal stimulus response to the Great Recession in the U.S. The ARRA included an increase in Medicaid payments amounting to about $87 billion dollars to be paid to states[6], over and above what would normally have been paid, for the years 2009 and 2010.

Such a discretionary increase in these payments is exactly in line with the results found in this section; the federal government only responded with large increases in Medicaid transfers when *the whole country* was in recession, rather than in response to idiosyncratic state-level downturns. This only matters, of course, if state-level business cycles are substantially different from the aggregate business cycle. I argue here that state-level recessions are indeed quite distinct from the aggregate business cycle, and occur with greater frequency.

To identify state-level recessions, I follow Brown (2017) by using the Philadelphia Fed's monthly coincident indices, which are available both for states and for the U.S. as a whole. I identify the turning points of the cyclical component of this index using the modified BBQ (MBBQ) algorithm from Engel[7], and identify recessions for the U.S. and all fifty states therein. While only 4 recessions are identified for the U.S. over

[6] https://www.healthaffairs.org/do/10.1377/hpb20100715.109669/full/

[7] Engel modifies the BBQ algorithm of Harding and Pagan (2002), and provides code at http://www.ncer.edu.au/data.

the sample period (1979-2017), most states experienced more than four recessions; furthermore, states were in recession for an average of 73 months, compared to just 46 months for the U.S. nationally. These recession differences suggest that U.S. states experience heterogeneous idiosyncratic shocks apart from the U.S. business cycle as a whole.

That states' business cycles differ substantially from the U.S. business cycle suggests the transfer responses identified here will have significant implications for risk in state government budgets. Notably, a state whose business cycle moves more independently from the rest of the country might be exposed to more risk because of the federal transfer system than an otherwise equal state whose cycle moved more instep with the U.S. cycle. Having already identified possible precautionary savings behavior by state governments, one might expect these independent states' governments to save more relative to other states; indeed, that is what these policy makers do, as I show in the following section.

3.2.2.3 Fact 3: Less Correlated States Save More

If the transfer policy of the U.S. federal government to U.S. state governments doesn't respond as much to idiosyncratic fluctuations, then states whose cycles are less correlated with the rest of the country might be expected to run higher balances of government savings. To evaluate this prediction, I develop a measure of a state's correlation with the business cycle of the rest of the country. For each state, I apply an HP-100 filter to two annual time series: the state's annual real GSP series and real U.S. GDP less the state's GSP. The long-run correlation of the cyclical component of each of these time series yields the correlation of a state's business cycle with that of the other 49 states. In this section, I show that this "correlation" measure is negatively associated with precautionary savings behavior on the part of U.S. state governments.

Figures 3.7, 3.8, 3.9, and 3.10 show the time path of state government savings for the five most correlated and least correlated U.S. states with GDP.[8] Clearly, states whose business cycles are least correlated with the U.S. business cycle run higher levels of government savings as a percentage of GSP than those which are most correlated. The most stark example is Figure 3.9, in which the most correlated states on average run slightly negative net assets (not including insurance funds). For a flavor of how correlations vary across the U.S., see Figure 3.11, in which states whose cyclical GSP is more correlated with U.S. GDP are highlighted.

Of course, there are a multitude of factors determining how correlated a state is with the rest of the country, some of which may also affect a state government's level of savings. To further demonstrate the relationship between the correlation measure and a state's government savings, I estimate the equation

$$s_{it} = 1 + \beta_1 y_{it} + \beta_2 y_{it}^{cycle} + \beta_3 \rho_i + \beta_4 \rho_i y_{it}^{cycle} + \Gamma X_{it} + \varepsilon_{it}. \tag{3.2}$$

Here s_{it} is some measure of state government savings, y_{it} is gross state product, y_{it}^{cycle} is its cyclical component, and ρ_i is the main explanatory variable of interest, namely, a state's correlation with the business cycle of the rest of the country. Because ρ_i

[8]For the remainder of the paper, I disregard Alaska. Alaska's reserve funds are massive in comparison to the other states, and it is the least correlated with the rest of the U.S. The case of Alaska certainly supports my conclusions, but I want to prevent it from driving the results entirely.

Figure 3.7: Average BSF Over GSP

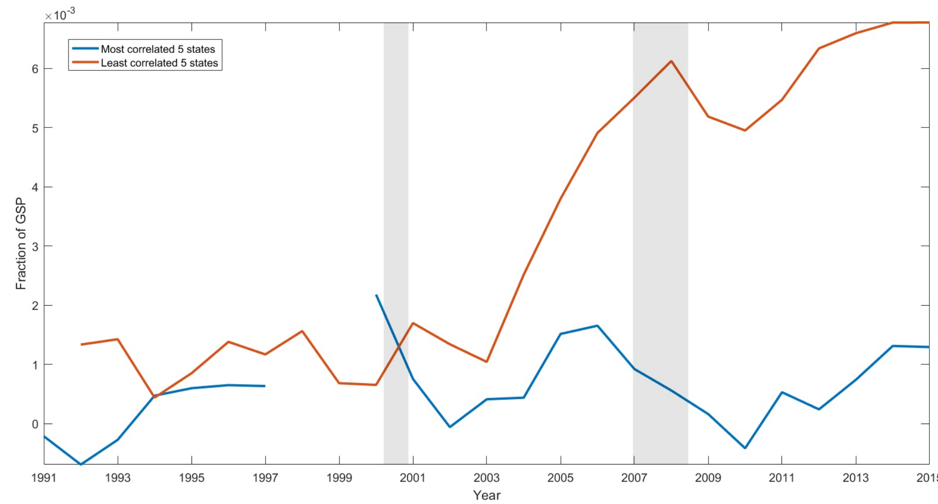

Figure 3.8: Average State Government Total Balances Over GSP

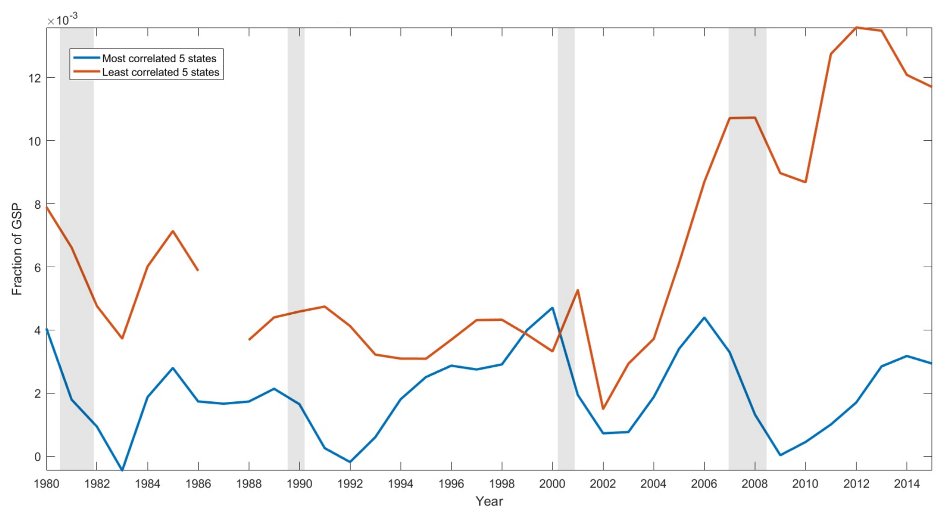

Figure 3.9: Average State Government Assets Over GSP (Not Incl. Insurance Funds)

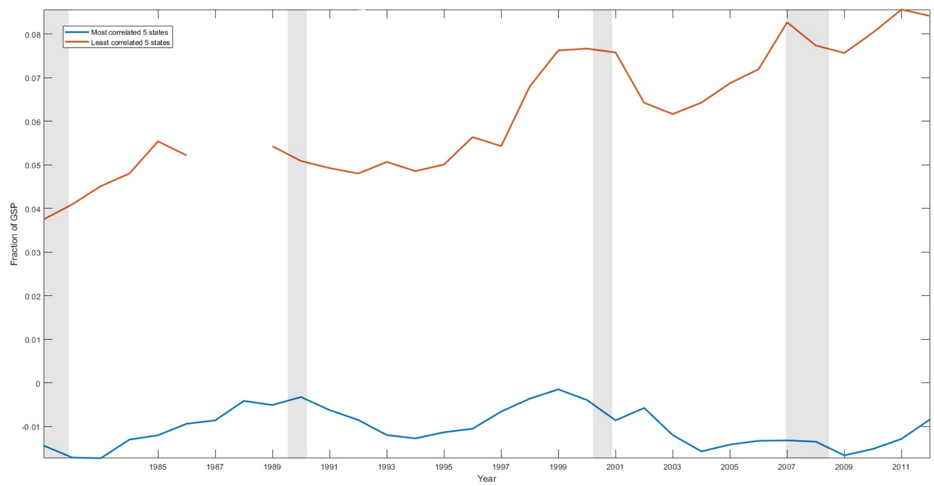

Figure 3.10: Average State Government Assets Over GSP

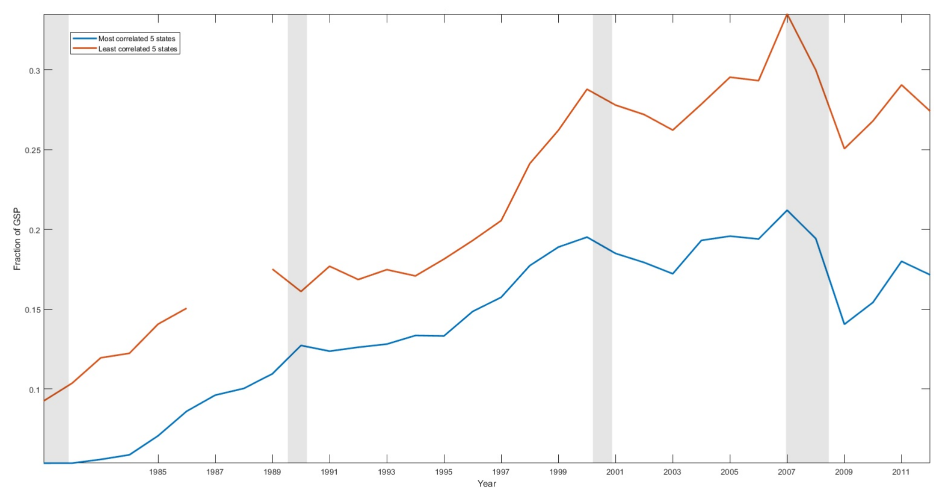

Figure 3.11: State Correlations with the National Business Cycle

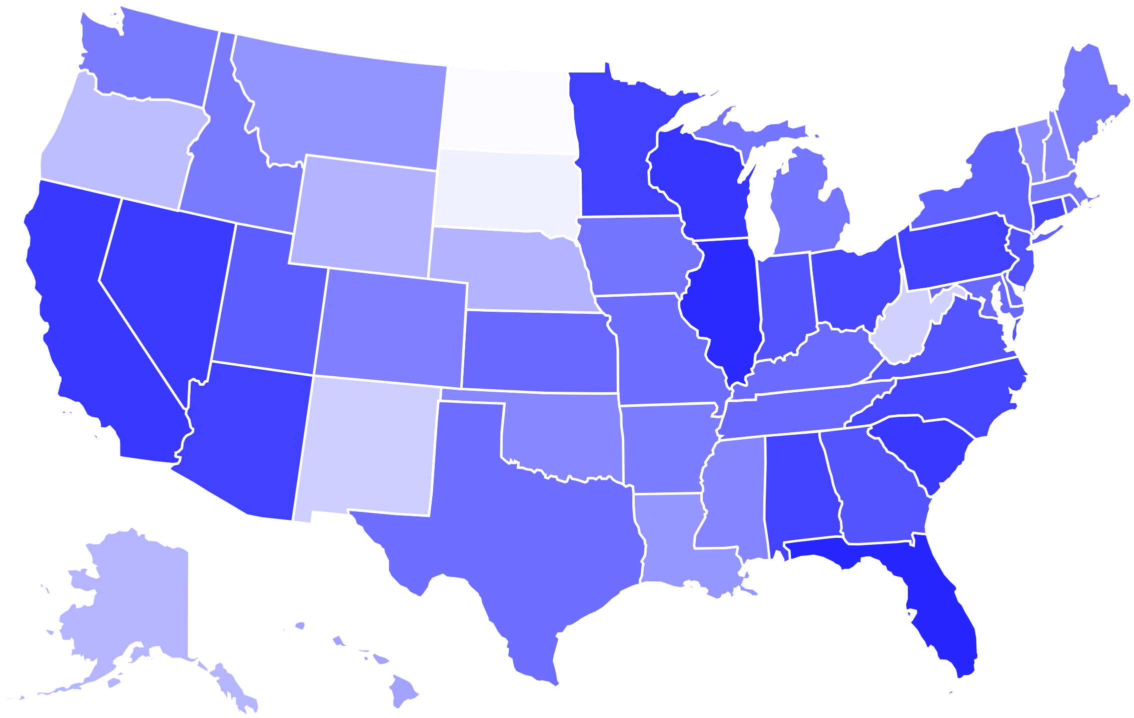

Note: Darker shade indicates a state's business cycle is more correlated with the business cycles of the other 49 states. Highest value is Illinois, with a correlation of 0.83.

Table 3.3: Determinants of State Government Balances

Dependent variable	log(∼ Real balances)			Real balances / GSP			Real balances / Expenditures		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Business cycle correlation	-0.0920**	-0.0825*	-0.0847*	-0.0054***	-0.0038**	-0.0038**	-0.0412***	-0.0262**	-0.0262**
	(0.0412)	(0.0466)	(0.0468)	(0.0016)	(0.0015)	(0.0015)	(0.0144)	(0.0128)	(0.0127)
Log(GSP), cyclical	0.5505***	0.4131**	0.1138	0.0225***	0.0177***	0.0167***	0.3036***	0.2536***	0.2487***
	(0.1886)	(0.1885)	(0.0789)	(0.0042)	(0.0050)	(0.0063)	(0.0377)	(0.0454)	(0.0547)
Log(GSP), cyclical * High correlation	-	-	0.9854*	-	-	0.0031	-	-	0.0155
			(0.5528)			(0.0068)			(0.0652)
Log(GSP)	0.0537***	0.0570***	0.0577***	-	-	-	-	-	-
	(0.0124)	(0.0131)	(0.0132)						
Controls	N	Y	Y	N	Y	Y	N	Y	Y

Note: This table reports results estimating Equation 3.2 for state balances. The sample for these regressions is 49 U.S. states (Alaska not included) for the years 1981-2012. The regressions are according to a random effects model. Standard errors are clustered at the state level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 3.4: Determinants of State Government Net Assets (Not Incl. insurance funds)

Dependent variable	log(∼ Real net assets)			Real net assets / GSP			Real net assets / Expenditures		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Business cycle correlation	-1.2630***	-1.2365***	-1.2320***	-0.1168***	-0.0605**	-0.0607**	-1.0787***	-0.5528***	-0.5541***
	(0.3325)	(0.3857)	(0.3845)	(0.0393)	(0.0267)	(0.0267)	(0.3393)	(0.2115)	(0.2118)
Log(GSP), cyclical	0.9278**	0.5754	0.0538	0.0513*	0.0581*	0.0202	0.9424***	1.0980***	0.9055***
	(0.3416)	(0.6129)	(0.6406)	(0.0274)	(0.0328)	(0.0381)	(0.1349)	(0.1655)	(0.2011)
Log(GSP), cyclical * High correlation	-	-	1.8011**	-	-	0.1196***	-	-	0.6059**
			(0.7729)			(0.0413)			(0.2714)
Log(GSP)	0.1993**	0.3001***	0.2979***	-	-	-	-	-	-
	(0.0834)	(0.0857)	(0.0852)						
Controls	N	Y	Y	N	Y	Y	N	Y	Y

Note: This table reports results estimating Equation 3.2 for non-insurance assets. The sample for these regressions is 49 U.S. states (Alaska not included) for the years 1981-2012. The regressions are according to a random effects model. Standard errors are clustered at the state level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

is a time-invariant object at the state level, I estimate a random-effects model[9] and cluster standard errors at the state level. Tables C.2 and C.3 give the output from a selection of these regressions for the two preferred definitions of state government savings.

Clearly, the relationship between "correlation" and state government savings is negative and significant in all specifications of the estimation model. The interpretation is exactly the stylized fact highlighted in this section: states whose business cycles are less correlated with the rest of the U.S. run higher levels of government

[9] Wooldridge, 2010.

savings. Evidence of the procyclicality of these savings balances is also seen in most specifications. Qualitatively, these results are robust to any of the aforementioned measures of state government savings.[10]

Among the controls, the variance of a state's cyclical GSP is sometimes positively related with savings levels, lending more evidence to the idea that these savings measures capture precautionary savings behavior.[11] There may also be weak evidence that the less correlated governments are more hesitant to spend out of their savings in bad times, perhaps out of desire to have funds in case worse conditions hit later. Regarding other control variables of interest, states with Democrat governors tend to run lower levels of savings, and the strength of a state's balanced budget rule seems not to be associated with higher savings levels. I include a fuller version of some of these regressions in Appendix C.3.

When combined with stylized facts 1 and 2, this third fact hints at an important feature of the effect of the federalist structure of the U.S. on state-level fiscal policy. It seems that, by exposing less correlated state governments to more risk by not insuring them against downturns as much as the more correlated states, the federal government creates stronger incentives for these state governments to engage in precautionary savings behavior. The rest of the paper attempts to expand on this story by putting forth a quantitative model of federalism and government policy over the business cycle that are able to reproduce these facts and others related to public finance over the business cycle. The model will allow counterfactual analysis of poli-

[10]The stylized fact also remains when the regressions are weighted by GSP.

[11]Seegert (2017) documents the rise in volatility of state tax revenues over time; this may also explain some of the rise in the savings measure over time for many states.

cies like balanced budget rules at the state level, shed light on the broader debate about the role for states in pursuing countercyclical fiscal policy, and contribute to the discussions aboutfiscal unionization. For example, counter to Oates (1972), if real frictions to centralized fiscal policy are sizeable enough, it may be optimal for lower levels of government to engage in stimulus policy.

3.3 An Information Model

To begin thinking about the cyclical behavior of state government savings in a quantitative sense, I put forward a model of multi-tiered governments and imformation asymmetries. The basic structure of the model is as follows: a state is modeled as a small endowment economy in a fiscal federation, or a "region," to avoid confusing usage of the word "state." There are two levels of benevolent governments, regional and central, each with its own advantages and disadvantages. The regional government must provide a certain level of some public good, but is subject to strict borrowing limits. The central government is not subject to borrowing restrictions and makes tax and transfer policy to help regional governments smooth consumption, but only observes the state of the world with a noisy signal.

The balanced budget restriction at the regional level versus the information friction at the central level is the key trade-off in the allocation of fiscal policy. In a model without frictions, there would be no difference between the provision of financing for the public good at the local or central level, as either government would be able to perfectly smooth the representative household's consumption over the cycle.

Furthermore, in the presence of any sort of balanced budget requirement or a similar disadvantage in smoothing consumption at the local level, it would be optimal for the central government to take over all countercyclical fiscal policy, leaving the local government to simply levy a tax exactly equal to public goods spending.[12] This is one conventional wisdom on federalism and fiscal policy articulated in Oates (1972).

Such specialization in fiscal policy is not observed in the data, however. As noted above, U.S. states engage in precautionary savings behavior. Therefore, it is likely not the case that centralized fiscal policy is strictly preferred to decentralized policy; there must be some trade-off between policies at regional and central levels. I propose to model this trade-off by way of an information friction on the part of the central government. While the central government has an advantage in its ability to smooth consumption through borrowing and/or compulsory transfers, it does not observe the state of the economy exactly, receiving a noisy signal about the endowment shock. Because of this, the central government will not respond perfectly to region-specific shocks, requiring the regional government to save up funds in order to smooth.

This way of motivating the trade-off is consistent with discussions about the advantages of state and local governments *vis-a-vis* central governments. For example, the CBO references differences in information about citizens' situations and preferences as a reason why local government action might be preferred in some cases,[13]

[12] Goodspeed (2016) studies the trade-off between federal transfers and rainy-day funds, where the relative effectiveness of each government's fiscal policies may differ structurally. In this case, the central government may not want to fully take over.

[13] CBO (2013) captures the basic tradeoff: "[Federal grants] can increase economic efficiency in instances when state and local governments have localized knowledge that would allow them to implement a program more efficiently and effectively than the federal government could, but they lack an incentive or funding to provide as much of a good or service as would be desirable from a national perspective..."

and Bordignon, Manasse, and Tabellini (2001) examines the behavior of optimal transfers when local governments have better information about their tax bases than the federal government. The "signal" method of modeling an information friction provided here can be interpreted in a number of ways. The most obvious interpretation is that of a central fiscal authority having imperfect measurement of indicators the regional economy[14]; however, it could also be that a far-away central authority, although receiving accurate measurements, is not as "tuned in" as local authorities with the effects on local citizens of the observed shocks.

Furthermore, an even more reduced form interpretation of the information friction is that it captures other factors which might dampen the ability of a central fiscal authority to respond to local shocks, including political economy frictions like slow or biased legislatures; below I elaborate on the relationship between the information and political economy models. The idea that local governments have better ability to know and match the preferences of their constituents goes all the way back to Tiebout (1956), and the political economy friction imposed by a legislature in centralized provision is found most notably in (Besley and Coate, 2003). That the federal government insures national shocks more consistently than local shocks suggests using some sort of friction in a model involving federal transfers; while I choose in this paper to explore an information interpretation, the political economy interpretation is also a helpful framework. The information angle serves as an example of how information economics can be applied to policy problems.

[14]Such imperfect information is likely not a lack of data but a lack of more informal information local policymakers might have access to via conversations with business leaders and workers.

3.3.1 Model Environment

3.3.1.1 Endowment Process

Income for the household in region i is allocated exogenously according to an endowment process. The household is passive; it does not engage in any behavior to affect its consumption.[15] The endowment in period t for region i is given by the following:

$$y_{it} = \bar{y} + \gamma z_{it} + \epsilon_{it}, \tag{3.3}$$

where $\bar{y}$ is the long-run mean of income, z_{it} is an aggregate component, ϵ_{it} is an idiosyncratic component, and γ multiplies the aggregate component. Both components follow an AR(1) process, for example, the process for ϵ_{it} is given by

$$\epsilon_{it} = \rho_\epsilon \epsilon_{i,t-1} + \xi_{it}^\epsilon, \tag{3.4}$$

where $\xi^\epsilon \sim N(0, \sigma_\epsilon^2)$.

Note that regions here are modeled as a continuum of *ex-ante* identical islands. The lack of ability to trade with other regions is an extreme assumption, but it highlights the lack of complete insurance available to the regional agents. It also eliminates the complications that might be introduced by regional governments competing strategically with one another. While inter-region games are no doubt interesting,

[15]The household may be thought to engage in one action, namely, the election of a regional government. In the framework here of a representative household, a government whose preferences exactly align with the household's is elected; this is exactly the type of regional government I consider. Equating the household and government's preferences and decisions for consumption and savings is a strategy used in Aguiar and Amador (2011).

the strategic interactions of interest in this paper are those between the central government and the subnational governments.

3.3.1.2 Regional Government

The problem of the government in region i is to choose a stream of taxes and savings, τ_{it} and s_{it}, to finance an exogenous stream of government purchases g_{it}, which generate no utility. The preferences of the government are exactly aligned with those of the representative household, which has utility over consumption in every period:

$$W\Big(\{c_t\}_{t=0}^{\infty}\Big) = \mathbb{E}\Bigg[\sum_{t=0}^{\infty} \beta^t U(c_{it})\Bigg]. \tag{3.5}$$

The government is subject to the following constraints:

$$c_{it} = y_{it} - \tau_{it}$$

$$g_{it} + s_{it} = \tau_{it} + T_{it} + (1 + r)s_{i,t-1}$$

$$s_{it} \geq \phi,$$

where τ_{it} is the lump sump tax chosen in period t, s_{it} is the savings (negative debt) of the regional government, y_{it} is the endowment, g_{it} is required government spending, ϕ is the per-period borrowing constraint, and T_{it} is the fiscal transfer from the federal government, which will be taken as given from the perspective of the regional government.

The transfers will result from the policy function of the central government,

$T(s_t, z_t, \theta_t, f_t)$, where θ is a noisy signal of ϵ_t which is unobserved by the regional government, and f_t is its prior belief about ϵ_t; both of these will be explained shortly. Given the observed state variables, the regional government cannot predict transfers T exactly, and must choose its policy to maximize *expected* utility over the possible realizations of the transfer function, such that its dynamic programming problem can be described by the following equation:

$$\begin{aligned} V^R(s_t, z_t, \epsilon_t) = \max_{s_{t+1} \geq \phi} \mathbb{E}_t[c_t] + \beta \mathbb{E}_t\big[V^R(s_{t+1}, z_{t+1}, \epsilon_{t+1})\big] \\ \texttt{s.t. } c_t + s_{t+1} = y_t + T(s_t, z_t, \theta_t, f_t) + (1+r)s_t - g_t. \end{aligned} \tag{3.6}$$

Note the presence in this problem of an expectation operator on current period consumption; in this model, the local government chooses next period savings before observing the realization of transfers T from the central government, and must use taxes and subsidies to balance the budget at the end of the period.

3.3.1.3 Central Government

In addition to the regional government, the central government features as a second optimizing agent in the model. Its inclusion reflects the fact that, in the context of the U.S., the federal government is not a passive agent with regard to fiscal policy. While much federal spending, such as the Social Security program, is indeed formulaic, the federal government also engages in a large amount of discretionary spending, up to a third of which is a direct transfer to the states. The central government is benevolent, so its optimization problem at first glance is almost equivalent to that of the regional government:

$$V^C(s_t, z_t, \theta_t, f_t) = \max_{T_t} \mathbb{E}_t[c_t] + \beta \mathbb{E}_t\big[V^C(s_{t+1}, z_{t+1}, \theta_{t+1}, f_{t+1})\big]$$
$$\texttt{s.t. } c_t + s(s_t, z_t, \epsilon_t) = y_t + T_t + (1+r)s_t - g_t. \tag{3.7}$$

Note, however, that the differences between the two decision problems are not inconsequential; the model generates a wedge between the decision rules even in the presence of identical welfare functions.[16] First, the central government's choice variable is T_t rather than s_t. While the fiscal balance carried over into the next period is chosen by the regional governments, the central government chooses how much to give to (or take from) the regional governments in the form of transfers in each period. Furthermore, there is no period budget balance necessarily required on the part of the central government. I assume that the central government can perfectly observe the aggregate shock z_t; as a result, it is able to perfectly insure the regions against aggregate shocks. Therefore, for the remainder of the paper I abstract from the aggregate shocks, and consider a model with only idiosyncratic shocks and the recasted policy functions $s(s_t, \epsilon_t, f_t)$ and $T(s_t, \theta_t, f_t)$. Given this recasting, two mechanisms can be employed to discipline the financial behavior of the central government and prevent it from accumulating debt indefinitely.

The first mechanism that can be employed is budget balance over the infinite horizon, i.e., a no-Ponzi-game condition on the central government's assets. In this

[16]This paper abstracts from the possibility that the regional government could simply tell the central government what its state is. The information friction might be thought of as the noise introduced as information is passed from lower levels of government to higher levels. Of course, differences in objective functions would exacerbate the differences in behavior, but the information mechanism alone is of sufficient interest for the purposes of this paper.

framework, there is a finite number of regions, and the central government has a stock of assets out of which positive transfers are paid and into which negative transfers are deposited. The no-Ponzi-game condition, then, is given by

$$\lim_{t \to \infty} \mathbb{E}_0 \Big[\frac{A_t}{(1+r)^t} = 0 \Big], \tag{3.8}$$

where A_t is the stock of assets held by the central government for the purposes of transfers to the regional government.

The second method is to have period budget balance and an infinite number of regional governments. In this setup, the central government in every period takes from some regions and gives to others, such that total transfers (for idiosyncratic shocks) net out to zero:

$$\int T_{it} di = 0. \tag{3.9}$$

For the simple case in which all regions are identical, solving for the transfer function in this case is analytically equivalent to the solution for the first mechanism. Appendix C.4 shows that both mechanisms result in the simple budget condition $\mathbb{E}_0 \big[T(s_t, \theta_t, f_t) \big] = 0$.

In addition to the differences in choice variables, note that the state variables are also different for the central government. The central government observes s_t, but receives a noisy signal θ_t about the state variable ϵ_t:

$$\begin{aligned} \theta_t &= \epsilon_t + \xi_t^\theta \\ \xi_t^\theta &\sim N(0, \sigma^2_{\xi^\theta}). \end{aligned} \tag{3.10}$$

Here, ξ_t^θ is the noise component of the signal, and σ_{ξ^θ} reflects the relative noisiness of the signal. Finally, the central government brings into the period a prior belief on the distribution of the idiosyncratic component ϵ: $f_t = N(\mu_t, \sigma^2_{\mu,t})$. After observing the signal, the central government updates this prior to form a posterior with which it forms its expectation for the choice of transfer, then projects this posterior forward into the next period using the law of motion for ϵ. This process is described in further detail in the next section.

3.3.2 Bayesian Updating

The central government begins time period t with a prior belief f_t on the distribution of ϵ_t: $f_t = N(\mu_t, \sigma^2_{\mu,t})$. Upon observing the noisy signal θ_t, the central government updates its belief to $\hat{f}_t = N(\hat{\mu}_t, \hat{\sigma}^2_{\mu,t})$ according to the following rules, which mimic the classic signal extraction problem put forth in Lucas (1973):

$$\hat{\mu}_t = \mu_t + \frac{\sigma^2_{\mu,t}}{\sigma^2_{\mu,t} + \sigma^2_{\xi^\theta}}(\theta_t - \mu_t) \tag{3.11}$$

$$\hat{\sigma}^2_{\mu,t} = \frac{\sigma^2_{\mu,t}\sigma^2_{\xi^\theta}}{\sigma^2_{\mu,t} + \sigma^2_{\xi^\theta}}. \tag{3.12}$$

It is this distribution $\hat{f}_t$ that the central government uses to form its expectations when solving for its optimal policy. The extent to which the belief about the mean is updated after observing the signal is determined by the relative variance of the noisy portion of the signal. The noisier the signal, the less weight is attached to it in the process of forming beliefs about the region's endowment.

At the end of the period, the central government must form its belief about ϵ_{t+1}, which is the prior distribution it will bring into the next period as a state variable. These priors for period $t+1$ are formed from applying the known $AR(1)$ process to the posteriors formed in period t:

$$\mu_{t+1} = \rho_\epsilon \hat{\mu}_t \tag{3.13}$$

$$\sigma^2_{\mu,t+1} = \rho^2_\epsilon \hat{\sigma}^2_{\mu,t} + \sigma^2_{\xi^\epsilon}. \tag{3.14}$$

Given these laws of motion, the posterior variance $\hat{\sigma}^2_{\mu,t}$ is bounded in the long run, and under certain conditions converges to a single value. In Appendix C.5, I show that the fixed point is

$$\hat{\sigma}^2_{\mu,\infty} = \frac{(\rho^2_\epsilon - 1)\sigma^2_{\xi^\theta} + \sigma^2_{\xi^\epsilon} + \sqrt{\left[(\rho^2_\epsilon - 1)\sigma^2_{\xi^\theta} + \sigma^2_{\xi^\epsilon}\right]^2 + 4\rho^2_\epsilon \sigma^2_{\xi^\theta}\sigma^2_{\xi^\epsilon}}}{2\rho^2_\epsilon}. \tag{3.15}$$

In solving the model, I assume that the central government has already reached this value for the posterior variance. This eliminates another state variable and allows the belief about the distribution of ϵ to be characterized by movements in μ_t.

3.3.3 Timing and Equilibrium

The timing of events in the model is as follows. In every period t,

1. All shocks ξ are realized.

2. Regional governments observe the true shock to their endowment ξ_ϵ, but not

the private signal θ_t. The central governments observes the noisy signal θ_t.

3. The central government forms its update belief $\hat{f}_t$ from the prior belief f and the signal θ_t.

4. Transfers and next period savings are chosen and committed to simultaneously by the central government and regional governments, respectively.

5. Regional taxes adjust to satisfy the choice of s_{t+1}, given the realization of T_t.

6. The central government uses $\hat{f}_t$ to form f_{t+1}, the prior belief going into the next period.

Definition: *A Markov Perfect Equilibrium is a set of policy functions* $\{s(s_t, \epsilon_t, f_t), T(s_t, \theta_t, f_t)\}$ *such that, given exogenous processes for* ϵ_t, θ_t, *and* g_t,

1. $s(s_t, \epsilon_t, f_t)$ *solves the regional government's problem given* $T(s_t, \theta_t, f_t)$, *and*

2. $T(s_t, \theta_t, f_t)$ *solves the central government's problem given* $s(s_t, \epsilon_t, f_t)$.

Here I define the equilibrium with one region, but in principle there could be many of these regions, each with its own equilibrium with respect to the central government. Since these regional governments are islands in the model, the solutions are separable, and it is helpful to cast the problem in terms of one region only.

Note that the simultaneity is an important assumption in the model. If, upon realization of shocks, one government moved first, then its move would reveal what its information was. A sequential game, therefore, would include signaling elements into the strategic decision-making process. While signaling may be an important factor in reality, whether the local or federal government is the more realistic first mover within a period is unclear. Additionally, given that, in the real world, policy-making is a continuously ongoing process, the simultaneity assumption is not particularly incredible.

3.3.4 Relation to a Political Economy Model

Thus far, I have specified the main friction to central government policy making as a matter of imperfect information and learning on the part of the central government. The other compelling source of inefficiency in centralized policy is political in nature. Discretionary transfers to regions/states and their agencies originate from budgetary decisions made in a legislative body; while noisy signals are likely at play here, political processes and voting are a major determinant of transfers. The federal government may not respond to idiosyncratic shocks simply because others vote down extra transfers to places in recession.

While the political story is somewhat different, I propose that the information frictions modeled above can be thought of as including political frictions, as well. To see this, consider a simplified version of the information model, in which there is no persistence for the idiosyncratic component ($\rho_\epsilon = 0$) for simplicity. Suppose the noise shock ξ^θ can take on one of two values: $\xi^\theta \in [-\bar{\xi^\theta}, \bar{\xi^\theta}]$, each with probability 1/2, where $\bar{\xi^\theta} > 0$. From the perspective of the regional government, which knows the true ϵ as well as the transfer function $T(s, \theta)$, it could receive one of two values for the transfer. If $\xi^\theta = -\bar{\xi^\theta}$, transfer T will be higher than it would be in a frictionless model, and vice versa if $\xi^\theta = \bar{\xi^\theta}$. So the regional government forms its expectations and policy knowing that, given s and ϵ, its transfer will be either T_{high} or T_{low}, each with probability 1/2.

Now consider a slightly different model, in which the friction to centralized policy making is political, i.e., transfers are voted on by a legislature for the central government. Following the political economy setup of Besley and Coate (2003), suppose there are two regions in the fiscal federation, and utility spillovers of the following type:

$$U_i(c_i, c_{-i}) = (1-\kappa)u(c_i) + \kappa u(c_{-i}), \kappa \in (0, 1/2).$$

Transfers are chosen by the legislature, which is modeled using the *minimum winning coalition* strategy. In every period, each region has probability 1/2 of being "in power," i.e., casting the median vote on transfer policy. If a region ends up in power, it receives a higher-than-efficient transfer, T_{high}, and if it is the minority, it receives a lower-than-efficient transfer T_{low}.[17] So, just as in the simple model of information, this framework requires the regional government to set policy knowing that it will receive T_{high} or T_{low} with probability 1/2, and we can expect its behavior to be similar to that in the simple information model. Given that a well known and widely used model of political frictions can be mapped into a similar version of the information model presented here, I think it reasonable to think of the "information" friction as potentially including political factors, as well.

[17]Here, as in Besley and Coate's model, the spillover term κ ensures that the deciding voter does not completely disregard the utility of the other region.

3.4 Quantitative Analysis

3.4.1 Estimation and Calibration

In order to examine the properties of the model, it is sufficient to consider the case of a single region. There are no interactions in this model across sectors, and the interactions of interest are between the regional governments and the central government. Given this strategy, I attempt to get some results roughly corresponding to the 'median' U.S. state. I estimate an AR(1) model for HP-100 filtered $\log(GSP)$ in all 50 states, and set ρ_ϵ and σ_ϵ to be the respective medians of the AR(1) parameter estimates. I then calibrate σ_{ξ^θ} to match the median of $corr(T_t, y_t)$, the correlation of transfer receipts from the federal government with output, at the state level.

Parameters for the baseline case are given in Table 3.5. I normalize $\bar{y} = 1$ and set $g_t = g = 0.05$ to roughly approximate data on U.S. state government spending. Utility is CRRA: $u(c) = \frac{c^{1-\nu}-1}{1-\nu}$, and I let $\nu = 2$. I choose an annual interest rate of 0.04, and set the discount rate such that $\beta < \frac{1}{1+r}$ to keep the region from wanting to increase savings indefinitely.[18] I set $\phi = 0$ to reflect the balanced budget constraints that are present in most U.S. states, and choose a realistic upper bound for regional government savings of 0.2, which is not binding in the baseline case. Later, I study the potential welfare effects of lowering ϕ, but 0 is an intuitive choice for the baseline case. I also restrict the transfer policy to respond linearly to the central government's signal, given its prior beliefs. This does not alter its optimal policy much, but it greatly eases the computation burden involved in solving the problem.

[18] Aiyagari, 1994.

Table 3.5: Baseline Parameters, Information Model

ρ_ϵ	ξ^ϵ	ξ^θ	$\bar{y}$	g	ν	ϕ	β	r
0.5095	0.0280	0.0671	1	0.05	2	0	0.961	0.04

I solve for the equilibrium policy functions by the use of an "inner loop, outer loop" strategy. The "inner loop" refers to the process of solving for each policy function given the policy function of the other government. The regional government policy is solved by value function iteration, and the central government policy is a static optimization problem, since its choices do not affect its future value function. The 'outer loop,' then, repeats this process, updating each policy function until both have converged.

3.4.2 Results

To assess the performance of the model compared to U.S. data, I simulate the calibrated baseline region for 10000 periods and observe its behavior. Table 3.6 presents some basic moments for some of the key variables of interest. Of the moments which are not explicitly targeted, $corr(y_t, s_t)$, the co-movement of local government savings and output matches remarkably well. The autocorrelation of transfer receipts is a bit low, but in an acceptable qualitative range. The variability of savings and transfers are low relative to the data, but their relative magnitudes to each other seem to make sense.

The one outlier, of course, is the behavior of savings. In the model, savings is more persistent and less variable than it is in the data. One reason for this may be

Table 3.6: Business Cycle Moments

	Model	Data (median of log hp-100 filter at state level)
$corr(y_t, T_t)$	-0.1457	-0.1536
$corr(y_t, s_t)$	0.2674	0.2862
$corr(y_t, y_{t-1})$	0.5205	0.5430
$corr(T_t, T_{t-1})$	0.1670	0.3531
$corr(s_t, s_{t-1})$	0.9466	0.1072
$sd(y_t)$	0.0366	0.0281
$sd(T_t)$	0.0170	0.0694
$sd(s_t)$	0.0549	1.16

Note: Moments reported here are for the cyclical components of HP filters of these variables.

the following: the moments I report in the data are in the extreme long-run case of a static economy. Savings here doesn't grow over time: once it reaches a desired level, it stays there and simply fluctuates around that level. It may be that, in the real world, U.S. states have not yet 'settled' into their desired long-run levels of savings, thus exhibiting more unpredictability.

I also compute the impulse response functions of savings and transfers to idiosyncratic shocks. Figures 3.12 and 3.13 give the predictions from the model, while Figures 3.14 and 3.15 display results from the data. The IRFs from the data are computed from running a VAR with savings (or transfer), state GSP, and U.S. GDP for all 50 states; I plot the median of the estimated IRFs and the median of the confidence bands. The model predicts a hump-shaped response of savings to an idiosyncratic regional shock, and the data seems to present some weak evidence in favor of this prediction. The complete lack of response of transfers to the regional shock is also consistent with the model, which predicts a response that is stunningly low in magnitude.

Figure 3.12: IRF of Savings to GSP Shock, Model

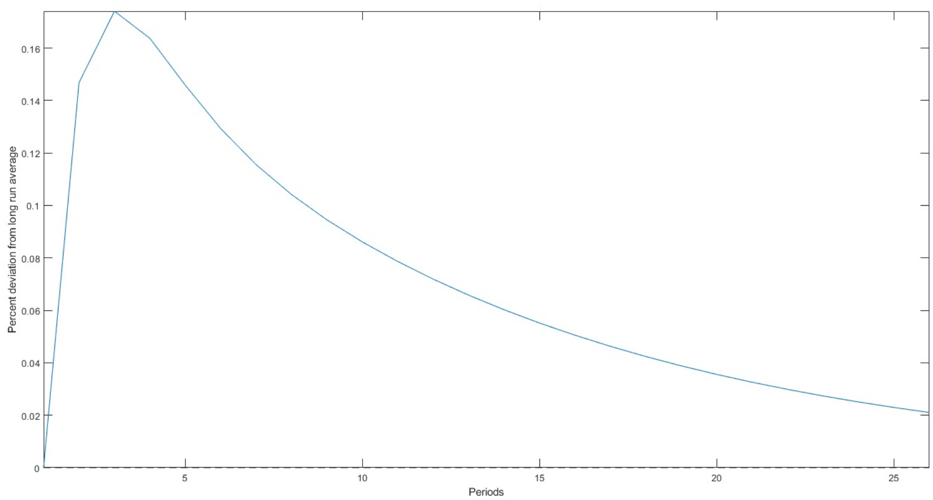

Figure 3.13: IRF of Transfers to GSP Shock, Model

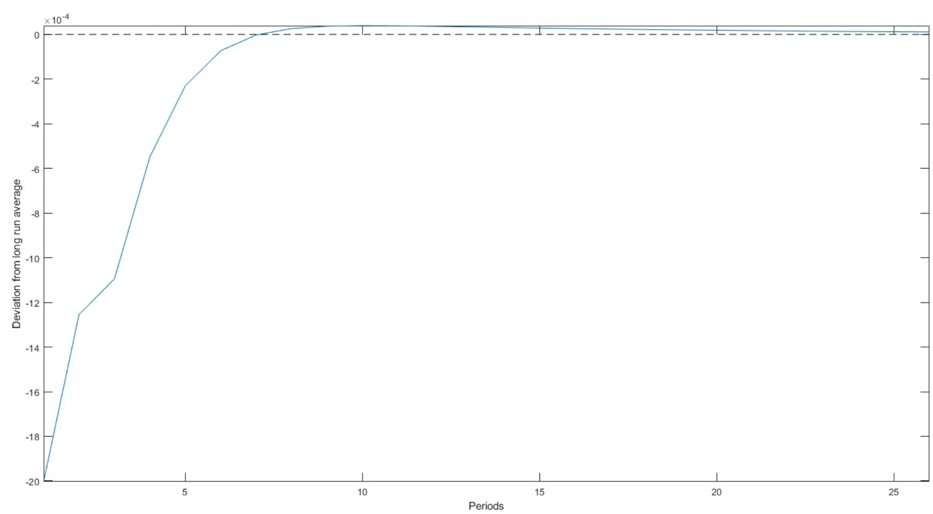

Figure 3.14: IRF of Savings to GSP Shock, Data

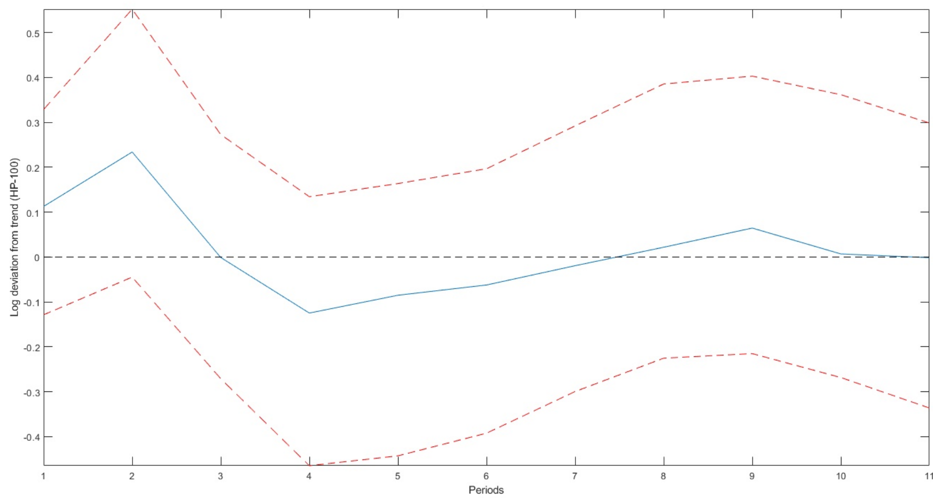

Figure 3.15: IRF of Transfers to GSP Shock, Data

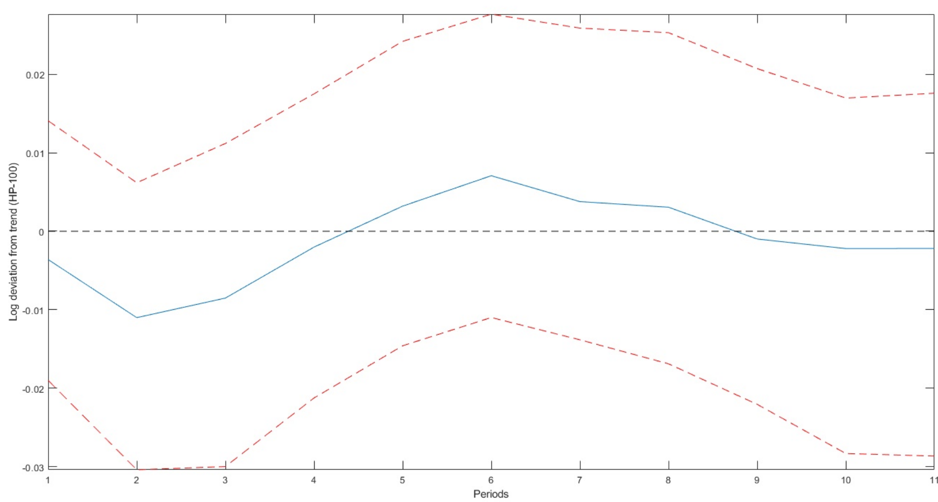

The lack of much meaningful transfer response in the model is consistent with the data, and is driven by the massive information cost implied by the baseline calibration. To generate a realistically low correlation between local fluctuations and transfer receipts from the federal government, the variability of the noise component of the central government's signal has to be almost *three times as large* as that of the real idiosyncratic shock. Even though such a sizeable noise shock is necessary to match the data relatively well, information costs that large seem almost incredible. Certainly, in the real world, the frictions at the level of centralized policy making are more diverse; for example, political dynamics are likely an important part of a central government's inability (or unwillingness) to insure regions against adverse shocks. In the context of this model, all such frictions are being captured by the single noise parameter; nevertheless, its size implies a central government that is quite weak in responding to economic shocks at the regional level.

3.4.3 Model Extensions

The baseline model makes an assumption on the nature of the central government's knowledge set; specifically, that the central government observes a noisy signal on the idiosyncratic state and beginning-of-period savings, but is not able to back out the previous period's idiosyncratic shock. This information structure is in the middle, so to speak, of two other possible setups, for which I present results in this section. I show that both model extensions behave in a qualitatively similar way to the baseline model, with only a couple key differences. These similarities suggest that the presence of this noisy signal is indeed the key driver of the quantitative results.

Table 3.7: Business Cycle Moments

	Baseline Model	Data	Extension 1 (observes lag perfectly)	Extension 2 (does not observe savings)
$corr(y_t, T_t)$	-0.1457	-0.1536	-0.1999	-0.1485
$corr(y_t, s_t)$	0.2674	0.2862	0.2643	0.2588
$corr(y_t, y_{t-1})$	0.5205	0.5430	0.5319	0.512
$corr(T_t, T_{t-1})$	0.1670	0.3531	0.0420	0.1779
$corr(s_t, s_{t-1})$	0.9466	0.1072	0.9590	0.9405
$sd(y_t)$	0.0366	0.0281	0.0377	0.0387
$sd(T_t)$	0.0170	0.0694	0.0173	0.0134
$sd(s_t)$	0.0549	1.16	0.0619	0.0568

Note: Moments reported here are for the cyclical components of HP filters of these variables.

Model extension 1 allows the central government to perfectly back out the previous period's idiosyncratic state ϵ_{t-1}. This could either come from its observation of s_t or a lagged ability to obtain data on local economic conditions, and either interpretation is fine for the solution of the model. While extension 1 represents an improvement on the central government's ability to obtain information relative to the baseline model, extension 2 is a tighter restriction. In extension 2, the central government is not able to observe the region's savings account whatsoever; it must form an expectation about local savings based on its posterior beliefs and its knowledge of the regional government's policy function.

Moments from the simulations of the two model extensions, along with the baseline model and the data, are given in Table 3.7. For each of these extensions, I use the parameterization from the baseline model for direct comparison with the baseline model. The model extensions behave in many respects very similarly to the baseline model, especially extension 2, suggesting that the exact nature of the central government's relationship with information on local savings is not the key driver of the model. The main difference occurs in model extension 1, in which transfers are more correlated with regional business cycles, with a much lower autocorrelation.

This is to be expected, given that the central government's prior is more accurate in this extension. The impulse response functions in these extensions unsurprisingly behave similarly in the extensions, with extension 1 displaying a stronger response of transfers to shocks, unsurprisingly with a one-period delay for the peak response, given that perfect information is revealed one period later.

3.4.4 Welfare Analysis

The precautionary savings behavior in the model depends critically on the presence of a strict limit on deficits on the part of the regional governments. This is a realistic feature of the model, given the widespread use of such balanced budget rules in the real world. In this model, such rules are not a result of optimizing behavior, but external parameters imposed on the agents. There are no other frictions on policy at the local level; therefore, were balanced budget rules to be sufficiently relaxed, regional governments would be able to achieve the social planner solution of full consumption smoothing for the household. By studying the effects of removing this constraint, I can say something about the welfare loss imposed by the baseline model relative to the social planner, as well as examine the potential effects of removing a balanced budget constraint.

To compute social welfare in the region, I solve the model and then simulate it starting at $s_0 = 0$, $\epsilon_0 = 0$, with time-0 social welfare being given by the sum of discounted utilities. As ϕ is lowered, the model approaches the social planner outcome quite quickly; when $\phi \leq -0.10$, the model essentially matches the social optimum. The implied welfare loss of the balanced budget rule (combined with the frictions on the central government, of course), then, is about 2.1 percent of welfare in the social planner case. This is, of course, a sizeable number. All of this loss occurs in the early periods, when the regional government increases taxes in order to build up its stock of savings to its desired long-run level. From this point on, the household is *better off* than it would be in the no-balanced-budget case, due to the extra interest income for its government, but these gains are far outweighed by the losses in early periods.

3.4.5 Implications for Fiscal Policy

The model presented in this paper implies that significant structural frictions to optimal countercyclical fiscal policy making may exist at centralized levels of government. The magnitude of these frictions suggests that a role exists for U.S. states to participate in active fiscal policy over the business cycle. Clearly, as the data show, U.S. state governments do in fact save in order to manage public finances during downturns. The most visible vehicles for such savings are rainy day funds, but states have other avenues, as well.

I have avoided using the word "stimulus" thus far, as Keynesian-type stimulus does not appear in my model. However, it is likely that the presence of these frictions in the making of fiscal policy at the federal level imply that it may be optimal for U.S. states to engage in stimulus policy during recessions; this conclusion runs counter to a conventional wisdom going back at least as far as Oates (1972), though Gramlich (1997) does find a stimulus role for states. More work should be done to explicitly model the implications of information and other political frictions for state-level fiscal policy over the business cycle.

Of note as well is the analysis of the effects of balanced budget rules in this model. In the information model presented here, removing the balanced budget rules completely eliminates welfare losses from baseline model, allowing regional governments to smooth completely over the business cycle. Of course, the assumption is that there are no real costs to borrowing, no default risk, etc. It may be that, in the real world, the balanced budget rules for U.S. are optimal responses to real costs of debt finance. In that case, the policy implications of loosening balanced budget rules would be completely the opposite: a welfare loss instead of a welfare gain.

Finally, these results may have something to say about the debate over fiscal policy in Europe. The results in this paper serve as a caution to potential efforts to establish a European fiscal union. While the ability of a such a union to finance spending with a deficit would be an advantage, there may be significant frictions to effective and timely fiscal stimulus along the lines of the frictions identified in this paper. Such frictions may prevent the effective smoothing of risks across countries, working against the stated purpose of unionization. In the presence of significant information or political frictions to optimal policy, a fiscal union in Europe may have trouble responding to localized shocks, especially as diverse an economic environment as the Eurozone.

3.5 Conclusion

This paper identified three key facts about the public finances of U.S. state governments over the business cycle, in the context of the role of the U.S. as a strong fiscal union. First, state governments engage in precautionary savings, in large part due to balanced budget requirements. Second, transfer payments from the federal government tend to respond more strongly to the aggregate business cycle than to state-level economic cycles. Third, states whose business cycles are less correlated with the national cycle tend to save more relative to other states. In light of the first two facts, I interpret the third fact as an indication that federal transfers (or lack thereof) influence state government savings behavior.

To give structure to this interpretation, I turn to a modeling framework in which I interpret a U.S. state as a small open endowment economy in a fiscal federation. Both levels of government, regional and central, may conduct fiscal policy, but each is faced with a different friction. Regional governments face borrowing limits, but central governments are faced with an information friction (which may also be interpreted as a political friction). The information friction prevents the central government from perfectly smoothing over the cycle, and thus regional governments must engage in precautionary savings.

I find that the model fits many qualitative features of the data well, and conclude that it is a useful framework in which to begin thinking about state government finances. The implied information friction is almost three times as large as real volatility; this is a formidable friction implying a central government with little ability to smooth over idiosyncratic cycles for states. Two variations on the structure of the information available to the central government yield quantitatively similar results. The baseline calibration also implies a long-run welfare loss of two percent relative to the social optimum. These results imply that frictions to policy making at the central level may be a significant factor in fiscal policy at the state level.

I conclude, contrary to some conventional wisdom, that space exists for states to actively pursue robust fiscal policies over the business cycle; furthermore, states that save are behaving optimally given their constraints. To the extent that central governments are constrained by a lack of information or stymied by politics, they may not be able to perfectly implement a first-best countercyclical fiscal policy. Of course, more research is needed into the size and nature of these frictions to policy, especially as they relate to recessions and expansions. Federalism, its complexities and mysteries notwithstanding, remains an important vehicle through which policy makers can insure citizens against adverse outcomes.

APPENDICES

APPENDIX A

Chapter 1 Appendix

A.1 New Keynesian Model Details

This appendix section provides further details on some of the equations in the New Keynesian portion of the model. First, consider again the period t profits of an intermediate goods producer,

$$P_{it}^N(c_t^N + g_t)\left(\frac{P_{it}^N}{P_t^N}\right)^{-\mu} - (1 - \frac{1}{\mu})W_t(c_t^N + g_t)^{\frac{1}{\alpha}}\left(\frac{P_{it}^N}{P_t^N}\right)^{-\frac{\mu}{\alpha}},$$

where P_{it}^N is the firm's price. When a firm is given the opportunity to set its price, it maximizes the present value of these per-period profits, taking into account the probability θ that its chosen price will continue on to the next period:

$$\mathbb{E}_t \sum_{s=0}^{\infty} Q_{t,t+s}\theta^s \left[\tilde{P}_{it}^N y_{t+s}^N \left(\frac{\tilde{P}_{it}^N}{P_{t+s}^N}\right)^{-\mu} - (1 - \frac{1}{\mu})W_{t+s}(y_{t+s}^N)^{\frac{1}{\alpha}}\left(\frac{\tilde{P}_{it}^N}{P_{t+s}^N}\right)^{-\frac{\mu}{\alpha}}\right].$$

Here $\tilde{P}_{it}^{N}$ is the chosen price of the firm, W_t is the raw wage in time t, and $Q_{t,t+s}$ is the nominal discount factor that converts income in $t+s$ to payments t; this discount factor is based on β and λ_t. The first order condition associated with this problem is given by

$$\mathbb{E}_t \sum_{s=0}^{\infty} Q_{t,t+s}\theta^s \left(\frac{\tilde{P}_{it}^{N}}{P_{t+s}^{N}}\right)^{-\mu} \left\{ \frac{\mu-1}{\mu}\tilde{P}_{it}^{N} - \frac{1}{\alpha}(1-\frac{1}{\mu})W_{t+s}\left[y_{t+s}^{N}\left(\frac{\tilde{P}_{it}^{N}}{P_{t+s}^{N}}\right)^{-\mu}\right]^{\frac{1-\alpha}{\alpha}} \right\} = 0.$$

The first term of the bracketed piece is the marginal revenue in each period, while the second term is the marginal cost. Separating the two terms results in a present value marginal revenue and a present value marginal cost, and recognizing that all adjusting firms will choose the same price $\tilde{P}_{t}^{N}$, marginal costs and revenues can be written recursively as

$$mr_t = \frac{\mu-1}{\mu} y_t^N \tilde{P}_t^N \left(\frac{\tilde{P}_t^N}{P_t^N}\right)^{-\mu} + \theta\mathbb{E}_t Q_{t,t+1} mr_{t+1}$$

and

$$mc_t = -\frac{1}{\mu}(1-\frac{1}{\mu})(y_t^N)^{\frac{1}{\mu}} W_t \left(\frac{\tilde{P}_t^N}{P_t^N}\right)^{-\frac{\mu}{\alpha}} + \theta\mathbb{E}_t Q_{t,t+1} mc_{t+1}.$$

Converting to relative variables $w_t = \frac{W_t}{P_t^N}$, $\tilde{p}_t^N = \frac{\tilde{P}_t^N}{P_t^N}$, and using $\pi_t^N = \frac{P_t^N}{P_{t-1}^N}$ and $Q_{t,t+1} = \beta\frac{\lambda_{t+1}}{\lambda_t}$ results in Equations 1.18 and 1.19 from the text.

The nontradable price index is defined by

$$P_t^N = \int_0^1 (P_{it}^N)^{1-\mu} di^{\frac{1}{1-\mu}}.$$

Again using the fact that all prices set in period t will be the same, we see that

$$(P_t^N)^{1-\mu} = \theta(P_{t-1}^N)^{1-\mu} + (1-\theta)(\tilde{P}_t^N)^{1-\mu}.$$

Dividing both sides by $(P_t^N)^{1-\mu}$ gives rise to Equation 1.20 in the text.

Similarly, to obtain Equation 1.21, consider the aggregation of total hours worked, along with the definition of production and demand equations:

$$h_t = \int_0^1 h_{it} di = \int_0^1 (y_{it}^N)^{\frac{1}{\alpha}} di = \int_0^1 \left(y_t^N \left(\frac{P_{it}^N}{P_t^N} \right)^{-\mu} \right)^{\frac{1}{\alpha}} di = y_t^{N\frac{1}{\alpha}} \int_0^1 \left(\frac{P_{it}^N}{P_t^N} \right)^{-\frac{\mu}{\alpha}} di.$$

Define price dispersion

$$s_t = \int_0^1 \left(\frac{P_{it}^N}{P_t^N} \right)^{-\frac{\mu}{\alpha}} di,$$

such that now $y_t^N = s_t^{-\alpha} h_t^{\alpha}$. Again making use of the symmetry of price decisions, we get

$$s_t = \theta \left(\frac{P_{t-1}^N}{P_t^N} \right)^{-\frac{\mu}{\alpha}} + (1-\theta) \left(\frac{\tilde{P}_t^N}{P_t^N} \right)^{-\frac{\mu}{\alpha}},$$

which simplifying gives Equation 1.21.

A.2 Empirical Methodology Appendix

A.2.1 Sample Selection and Data Cleaning

This appendix describes the process for selecting and cleaning the trade-level municipal bond data for use in the paper. First, I use the Bloomberg terminal to obtain CUSIP codes for all General Obligation (GO) bonds issued by general governments at any time up to present day. Depending on download limits, it may be necessary to break up the downloads into blocks of 5000 bonds or fewer.

Bloomberg provides CUSIP codes for each bond at the 8-digit level, consisting of a 6-digit issuer code followed by a 2-digit issue-specific code. MSRB, however, reports CUSIPs at the 9-digit level. The 9th digit in any CUSIP code is an automatically generated character according to the following algorithm:

1. Assign each character of the 8-digit code a numeric value x_i, with numeric characters being assigned their own value $0 - 9$, and alphabetic characters assigned numeric values beginning with 10: $A = 10$, $B = 11$, ... $Z = 35$

2. Construct a sum $S = \sum_{i=1}^{8}(1+I(i))x_i$, where $I(i) = 1$ if i is even, and $I(i) = 0$ if i is odd. In short, every other x_i is multiplied by 2.

3. Let s be the last (ones) digit of the sum S

4. Assign the 9th digit of the CUSIP code to be the *complement* of s, i.e., $10 - s$, or simply s if $s = 0$.

With the full 9-digit CUSIP codes in hand, I then request the full trade-level MSRB dataset from WRDS, which includes info on every brokered trade of a muni bond included in the list of CUSIPs I provide. This data begins in 2005, when brokers in the municipal bond market were required to provide real-time transaction information to the MSRB, and continues to the present day. The MSRB data includes bond characteristics such as coupon, dated date, and maturity date, and trade characteristics like par value, price, yield, time and date, and whether the trade was a purchase from or sale to a customer or if it was an inter-dealer trade.

I broadly follow Schwert's conditions for cleaning this trade-level dataset to remove potential errors in the data. This includes all bonds with coupons greater than 20% and times to maturity over 100 years in the future. It also drops individual trades with a yield to maturity of 0, a price outside the range $[50, 150]$, or a recorded trade date after the maturity date. This results in a dataset of 1,587,426 trades from 2005 to 2019.

The dataset used in the main estimation procedure takes the monetary shocks series described below and merges with a dataset of daily yields and spreads. Daily yields and spreads are assumed to be the median value for a bond-day pair. An observation in the resulting data is an FOMC decision day-muni bond pair, with two sets of yields and spreads. The first is the most recent daily price as of the FOMC day, and the second is the most recent daily price as of the day two weeks after the FOMC day; for a bond that has not been traded in two weeks, these two values may be the same.[1]

[1] Though as I note in the body of the paper, such observations are ultimately dropped

A.2.2 Identification of Monetary Shocks

As mentioned in the body of the paper, I employ the method of Bu, Rogers, and Wu, 2019 to identify monetary shocks at the FOMC date frequency. The BRW method uses a Fama-MacBeth two-step procedure to extract monetary shocks from a series of U.S. treasury yields. The procedure normalizes monetary shocks m_t such that they enter one-to-one into daily changes in the 5-year treasury yield[2]:

$$\Delta R_t^5 = \alpha + m_t + \eta_t.$$

The method then takes the zero-coupon treasury series, representing years to maturity i from $i = 1$ to $i = 30$. Each of these yields are assumed to respond to monetary shocks on FOMC dates according to

$$\Delta R_t^5 = \tilde{\alpha}_i + \beta_i m_t + \eta_{it}.$$

The first step of the procedure seeks to estimate the series of 30 parameters β_i. Since m_t is unobserved, the method uses the normalization to R_t^5, allowing us to instead plug in and estimate the equation

$$\Delta R_{it} = \alpha_i + \beta_i \Delta R_t^5 + \varepsilon_{it},$$

where $\alpha_i = \tilde{\alpha}_i + \beta_i \alpha$ and $\varepsilon_{it} = \eta_{it} + \beta_i \eta_t$.

An immediate problem arises in this estimation: ε_{it} is correlated with R_t^5 through

[2]The choice of maturity is not crucial to the procedure, a 1- or 2- year bond would work, as well.

η_t, resulting in a biased OLS estimate. To deal with this issue, the BRW method estimates each β_i using a Rigobon and Sack (2004) instrumental variables method. In short, an estimate for β_i can be obtained from the equation

$$[\Delta R_{it}] = \alpha_i + \beta_i[\Delta R_t^5] + \mu_{it},$$

where $[\Delta R_t^5] = (\Delta R_t^5, \Delta R_t^{5*})'$ and $[\Delta R_{it}] = (\Delta R_{it}, \Delta R_{it}^*)'$. Variables with a * represent a one-day movement in the corresponding rate one week before the FOMC date. The instrumental variable for this estimation is $[\Delta R_t^{IV}] = (\Delta R_t^5, -\Delta R_t^{5*})'$. The procedure relies on the assumption that the variance of *non-monetary* news does not change from week to week.

The second step, armed with the IV estimates $\hat{\beta}_i$, then estimates the equation

$$\Delta R_{it} = \alpha_t + \hat{m}_t\hat{\beta}_i + \varepsilon_{it}$$

on each day t, recovering the estimated monetary shocks m_t as the resulting coefficients.

A.2.3 Schwert Illiquidity Measures

In Section 1.3.4.3 I describe the method for decomposing municipal spreads into risk and illiquidity components, as in Schwert (2017). To construct the illiquidity measure ψ_{it}, I standardize three measures of illiquidity used by Schwert, and construct ψ_{it} as the monthly average of these three measures. The monthly average is used due to the paucity of munis with multiple trades on a given day, which is required for the daily measures.

The first measure originated in Feldhutter (2012), and is intended to explicitly capture the transaction costs introduced by the over-the-counter nature of bond markets, in which bonds might trade at multiple prices at the same time. This is the "Imputed Round-Trip Cost" measure of illiquidity, and is measured as follows:

$$IRC_{its} = \frac{P_{its}^{max} - P_{its}^{min}}{P_{its}^{min}}, \tag{A.1}$$

where P is the price, i is a CUSIP code, t is a given day, and s is a trade size. The idea of this measure is to capture the common occurence in which a dealer matches a buyer with a seller, with the difference in the prices representing the costs of finding and making the transaction. In the data, trades of the same bond on the same day of the same size are coded as round-trip trades, and the daily illiquidity measure is the average of round-trip trades on that day.

Another measure of the transaction costs element of liquidity is the "Price Dispersion" measure from Jankowitsch, Nashikkar, and Subrahmanyam (2011). This measure is similar to the first measure, but uses all prices on a given day. This measure of illiquidity represents the average dispersion around the "market consensus" price, or the average price on a given day:

$$DISP_{it} = \sqrt{\frac{1}{\sum_j Q_j} \sum_j (P_{ij} - M_{it})^2 Q_j}, \tag{A.2}$$

where j represents a trade of bond i on day t, P_{ij} is the price of trade j, Q_j is the par value of the trade, and $M_{it} = \frac{\sum_j P_{ij} Q_j}{\sum_j Q_j}$. If a bond's prices are highly dispersed on a given day, it could reflect high transaction costs or inventory risks for dealers, among other sources of illiquidity in bond markets.

The third measure, from Amihud (2002), is meant to capture the price impact of trades for a municipal bond. This is related to the market depth component of liquidity, i.e., the ability of a bond to sustain large trades without large movements in price. If, on a day for which there are multiple price changes for a muni, the average price change relative to trade size is large, then trades are having an impact on prices, and market depth is low. The Amihud (2002) measure, then, is given by

$$DEPTH_{it} = \frac{1}{N_{it}} \sum_{j=2}^{N_{it}} \frac{\left| \frac{P_{ij} - P_{i,j-1}}{P_{i,j-1}} \right|}{Q_j}, \tag{A.3}$$

where notation is the same as above, and N_{it} is the number of trades of bond i on day t. Note that this measure begins with the second trade on a given day, since intraday price changes are the object of interest here.

A.3 Additional Empirical Results

A.3.1 VAR Evidence of Government Behavior

This section provides aggregate time series evidence on the repsonse of state and local government fiscal policy to monetary shocks. I follow closely the strategy of Christiano, Eichenbaum, and Evans (1996), by estimating the VAR equation

$$Y_t = A + B_1 Y_{t-1} + ... + B_4 Y_{t-4} + \varepsilon_t, \tag{A.4}$$

where t corresponds to one quarter. The vector Y includes variables in the following order:

$$Y = [\log GDP \;\; \log C \;\; \log P \;\; \log I \;\; \log X \;\; \log WL \;\; R \;\; \log \Pi \;\; \Delta M],$$

where GDP is GDP, C is Personal Consumption Expenditures, P is the GDP deflator, I is private investment, WL is earnings, R is the federal funds rate, Π is profits, and ΔM is the change in M2 from the previous period. t represents a quarter in the U.S.; for monthly variables I use the first month of the quarter.

X here is the response variable of interest, corresponding to state and local government total debt, consumption expenditures, investment, or consumption + investment. The expenditures are reported at the quarterly level as a part of NIPA. Debt is included in the Federal Reserve's flow of funds data; due to a definitional change in 2004, I adjust pre-2004 values to match the post-2004 series, imputing the 2004Q1 growth rate to 2003Q4. Figures A.1, A.2, A.3, and A.4 here show the effect of an expansionary shock to the federal funds rate on the variables of interest. Furthermore, Figure A.5 gives the estimated response of output in the VAR with debt.

Figure A.1: CEE Impulse Response to Fed Funds Shock

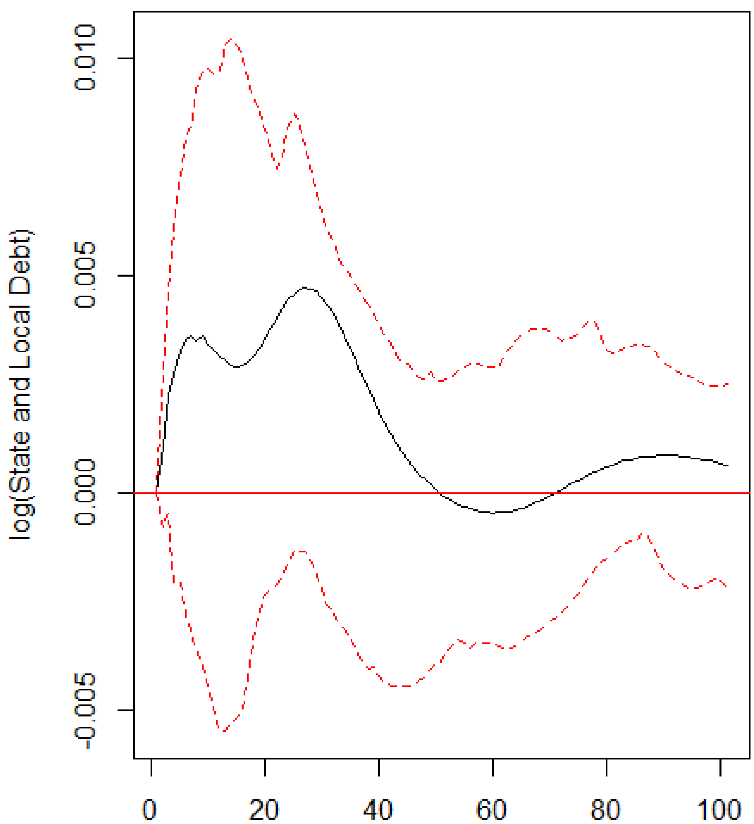

Figure A.2: CEE Impulse Response to Fed Funds Shock

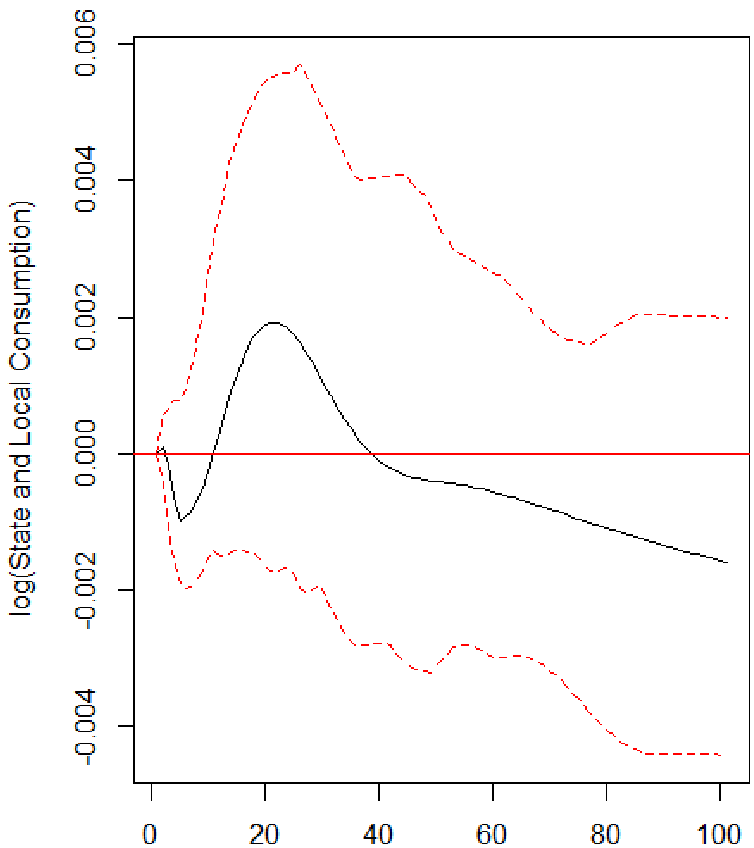

Figure A.3: CEE Impulse Response to Fed Funds Shock

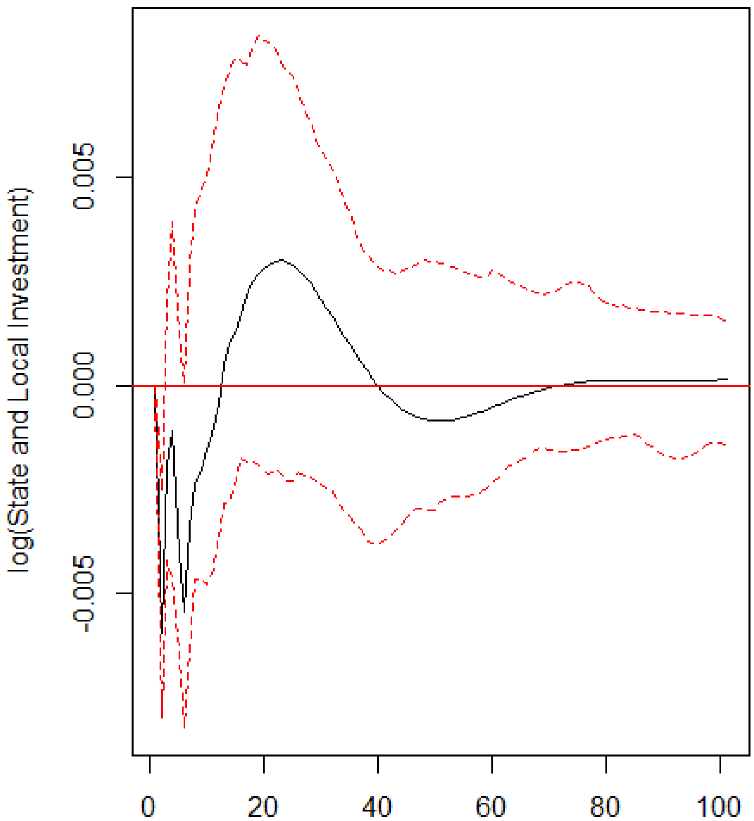

Figure A.4: CEE Impulse Response to Fed Funds Shock

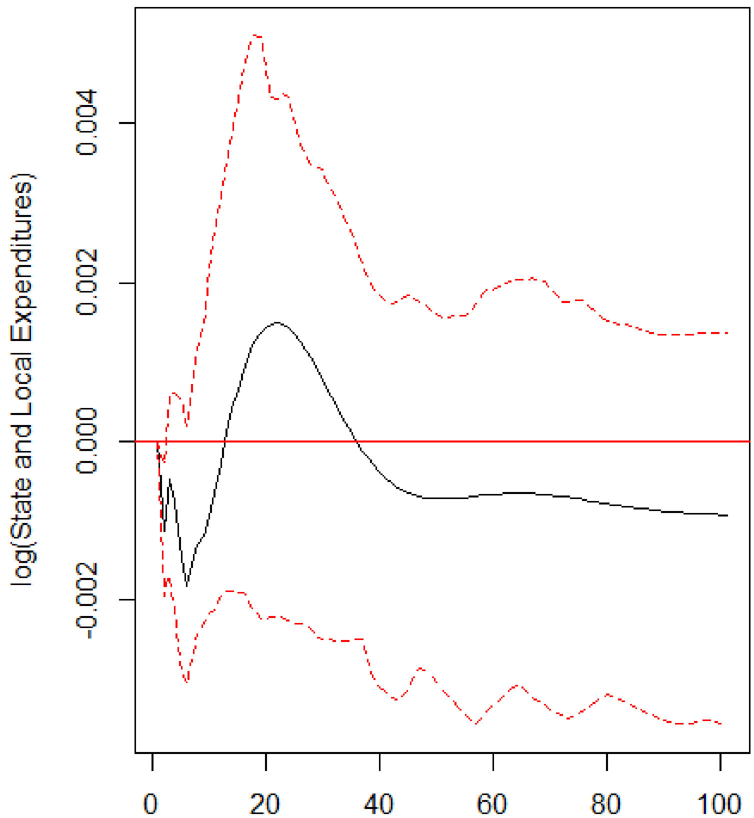

Figure A.5: CEE Impulse Response to Fed Funds Shock

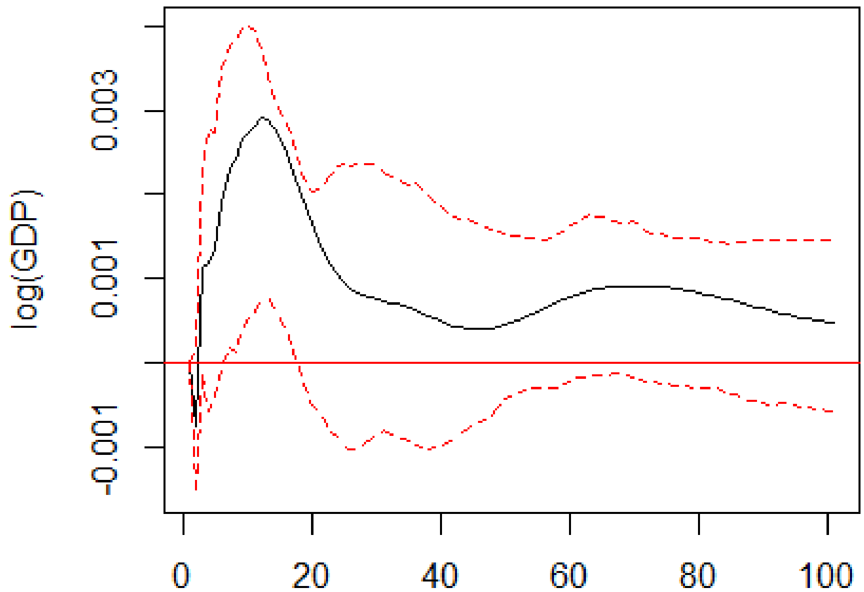

While there is an initial decrease in expenditures in the short run, due to conventional leaning-against-the-wind factors, expenditures do seem to rise in the medium run. This increase corresponds with the peak of the debt buildup response. As such, it is consistent with the borrowing costs channel put forward in this paper.

A.3.2 Additional Specifications

In this section, I investigate two additional specifications of the baseline results in Table 1.3. In Table A.1, I allow the response of municipal yields to monetary shocks to differ based on whether or not the monetary shock is positive, i.e., in include the indicator $\mathbb{I}[Shock > 0]$. Table A.2 is the same as in the main body of the text,

Table A.1: Time Series Results, Shock Direction

	All GO	State GO	Local GO	SP 500	All GO	State GO	Local GO	SP 500
Monetary shock	0.17	0.18	0.18	0.56	0.17	0.12	0.26	0.78
	(0.10)	(0.10)	(0.10)	(0.17)	(0.18)	(0.23)	(0.16)	(0.25)
Monetary shock * $\mathbb{I}[Shock > 0]$	0.09	0.10	0.09	-0.11	0.12	0.22	0.03	-0.23
	(0.16)	(0.16)	(0.16)	(0.27)	(0.29)	(0.32)	(0.26)	(0.42)
Horizon	2 days	2 days	2 days	2 days	6 days	6 days	6 days	6 days
N	2147	2147	2147	2147	2139	2139	2139	2139

Note: An observation corresponds to one day, around which a window is constructed from the previous day's price and the price at a given horizon. Each column refers to a separate time series regression of an index on monetary shocks. Heteroskedasticity-robust standard errors are reported in parentheses.

Table A.2: Time Series Results, Controlling for S&P 500 Index

	All GO	State GO	Local GO	SP 500	All GO	State GO	Local GO	SP 500
Monetary shock	0.22	0.24	0.23	0.50	0.24	0.24	0.28	0.66
	(0.08)	(0.08)	(0.08)	(0.14)	(0.14)	(0.16)	(0.13)	(0.21)
Horizon	2 days	2 days	2 days	2 days	6 days	6 days	6 days	6 days
N	2144	2144	2144	2144	2136	2136	2136	2136

Note: An observation corresponds to one day, around which a window is constructed from the previous day's price and the price at a given horizon. Each column refers to a separate time series regression of an index on monetary shocks. Heteroskedasticity-robust standard errors are reported in parentheses.

but controlling for the S&P 500 index. In each of these specifications, I find no significant differences from the baseline results reported in the paper. While the sign of the coefficients in Table A.1 indicate the upward shocks might have larger effects, the standard errors are large; additionally, the coefficients in Table A.2 are virtually indistinguishable to those in Table 1.3.

A.3.3 Sector-Level Time Series

In Table A.3, I summarize estimates of Equation 1.33, computed separately for sector-level indices.

Table A.3: Sector-level Time Series Estimates

Sector	Coefficient (s.e.)
Airport	0.175 (0.078)
Education	0.205 (0.083)
Health Care	0.170 (0.075)
Higher Education	0.195 (0.083)
Infrastructure	0.206 (0.116)
Land Backed	1.291 (0.684)
Lifecare	0.160 (0.071)
Multifamily	0.150 (0.078)
Nursing	0.569 (0.169)
Port	0.178 (0.079)
Public Power	0.210 (0.291)
Single Family	0.090 (0.058)
Student Loan	0.205 (0.086)
Tobacco	0.230 (0.145)
Toll Road	0.194 (0.085)
Transportation	0.197 (0.084)
Utility	0.197 (0.122)
Water and Sewer	0.191 (0.080)

Note: An observation corresponds to one day, around which a window is constructed from the previous day's price and the price at a given horizon. Each column refers to a separate time series regression of an index on monetary shocks. Heteroskedasticity-robust standard errors are reported in parentheses.

A.3.4 Government Finance Variables

In Tables A.4, A.5, and A.6, I summarize estimates of Equation 1.34, computed separately for a number of government finance statistics. The coefficient estimates in these tables represent the interaction between the variable and the monetary shock, in separately estimated regressions. The tables correspond to revenues, expenditures, and other categories, respectively.

Table A.4: Panel Estimates: Interactions with Public Finance Revenue Variables

Interaction Variable (log)	Interaction Coefficient (s.e.)
Total	-0.204
	(0.138)
General	-0.199
	(0.136)
Total Tax	-0.190
	(0.127)
Property Tax	-0.031
	(0.017)
Sales Tax	-0.094
	(0.052)
Income Tax	-0.064
	(0.045)
Intergovernmental	-0.192
	(0.128)
Miscellaneous	-0.180
	(0.133)
Interest	-0.150
	(0.122)
Utilities	-0.074
	(0.057)

Note: An observation corresponds to an FOMC date-muni bond pair. Each *row* corresponds to a regression with the specified variable interacted with monetary shocks, as in Equation 1.34. The reported coefficient is the coefficient on the interaction term. Standard errors are reported in parentheses, and are clustered at the date level. $N = 12100$ for all regressions.

Table A.5: Panel Estimates: Interactions with Public Finance Expenditure Variables

Interaction Variable (log)	Interaction Coefficient (s.e.)
Total	-0.198
	(0.135)
Current	-0.192
	(0.138)
Capital	-0.192
	(0.138)
Construction	-0.175
	(0.114)
Interest	-0.187
	(0.127)
Benefits	-0.097
	(0.062)
Wage	-0.199
	(0.138)
General	-0.196
	(0.134)
Education	-0.085
	(0.056)
Health	-0.104
	(0.068)
Highway	-0.156
	(0.102)
Housing	-0.122
	(0.096)
Public Welfare	-0.088
	(0.062)
Utility	-0.095
	(0.076)
Retirement	-0.095
	(0.063)

Note: An observation corresponds to an FOMC date-muni bond pair. Each *row* corresponds to a regression with the specified variable interacted with monetary shocks, as in Equation 1.34. The reported coefficient is the coefficient on the interaction term. Standard errors are reported in parentheses, and are clustered at the date level. $N = 12100$ for all regressions.

Table A.6: Panel Estimates: Interactions with Other Public Finance Variables

Interaction Variable (log)	Interaction Coefficient (s.e.)
Population	-0.236
	(0.148)
Total Long Term Debt	-0.173
	(0.119)
Total LTD Issued	-0.173
	(0.119)
Total Cash and Securities	-0.169
	(0.117)
Total Cash and Securities (non-insurance)	-0.194
	(0.136)

Note: An observation corresponds to an FOMC date-muni bond pair. Each *row* corresponds to a regression with the specified variable interacted with monetary shocks, as in Equation 1.34. The reported coefficient is the coefficient on the interaction term. Standard errors are reported in parentheses, and are clustered at the date level. $N = 12100$ for all regressions.

A.3.5 Government Type

In addition to the heterogeneity in the body of the paper, and in order to provide a more comprehensive view of heterogeneity in municipal bond spreads, I investigate the response of muni yields and spreads to monetary shocks on two important dimensions. The first is government type. Insofar as bonds issued by state governments are different in terms of liquidity (or risk) than their city and county counterparts, these bonds might respond differently to monetary shocks. In particular, if these bonds are on average more liquid than bonds of smaller governments, we might expect their yields to respond more strongly to monetary shocks. One piece of suggestive evidence in this direction is the fact that state governments are disproportionately represented in the sample of bonds which actually record a trade in the window following a monetary shock.

Table A.7: 2-Week Response of Muni Yields to Monetary Shocks, by Issuer Type

	d yield	d yield ($\neq 0$)	d spread	d spread ($\neq 0$)
Big government	0.01***	0.06***	0.01**	0.05***
	(0.00)	(0.01)	(0.00)	(0.01)
Monetary shock	0.24***	0.66***	0.05	0.10
	(0.06)	(0.15)	(0.07)	(0.16)
Monetary shock * Big government	-0.02	-0.21*	-0.13*	-0.22**
	(0.09)	(0.13)	(0.07)	(0.11)
N	74338	36747	74286	36685

Note: An observation corresponds to an FOMC date-bond pair. Each column refers to a separate regression. Standard errors are reported in parentheses, and are clustered at the date level.

To identify what type of government issued a particular bond, I first retrieve the bond issuer names from the Bloomberg Terminal. I then keep those issuers which include the words "state" or "commonwealth," and label these the "big" government issuers.[3] The estimate of interest, then, is the differential magnitude of response to monetary shocks for these "big" governments *vis-à-vis* other types of governments. Table A.7 repeats the estimation of Equation 1.34, reporting the interaction between the estimated shock and government type.

The results lend some further insight into the pattern observed earlier, in which the average response of muni yields to monetary shocks is made up of some bonds that adjust and some that do not. Interestingly, while there does not seem to be any differential response *on average* between bonds from big governments and smaller governments, there is a difference conditional on adjustment. Among the bonds which adjust price in response to a monetary shock, the bonds for state governments adjust 45 basis points in accordance with a 100 bp shock, whereas all other bonds

[3] I can also break out the smaller governments by the words "city," "county," "town," "village," etc., but the most important distinction seems to be the fifty states versus all other governments. Surely there is heterogeneity within cities and towns; this is a potential direction for future research.

adjust 66 basis points on average. Of course, the bonds for these bigger governments are a larger proportion of the "responding" sample than the full sample; in other words, their yields are more likely to respond to a monetary shock than those of smaller governments. Because these bonds trade at a higher frequency, their price adjustments in response to a given shock, conditional on adjustment, are smaller than those bonds which trade at a lower frequency. The pattern for spreads exhibits a similar pattern, in which the coefficient on the interaction term is more negative and more precisely estimated in the reduced sample; the average response remains zero, as before.

The second additional margin of heterogeneity involves an attempt at a comprehensive measure of unexplained spreads, which is made up of liquidity, risk, and tax components. I residualize the implied spread on every transaction in the original data by regressing out time to maturity and fixed effects at the month level. For each bond, I take the average of its residualized (actual minus predicted) spread to compute a time-invariant measure of unexplained spreads for each muni. If average residuals are above zero, I code the bond as "high spread;" similarly, I code as "low spread" those bonds for which average residuals are below zero. Of course these unexplained spreads include liquidity, risk, and tax components, but they represent a simple and intuitive margin of heterogeneity that does not require dropping observations yet carries some important information about the desirability of certain types of bonds.

Table A.8 presents the results of analogous regressions, in which the more liquid (and less risky) bonds are expected to be in the "low unexplained spread" category.

Table A.8: 2-Week Response of Muni Yields to Monetary Shocks, by High/Low Spreads

	d yield	d yield ($\neq 0$)	d spread	d spread ($\neq 0$)
Low spread	0.02***	0.05***	0.02***	0.05***
	(0.00)	(0.01)	(0.00)	(0.01)
Monetary shock	0.22***	0.56***	-0.00	-0.02
	(0.09)	(0.21)	(0.08)	(0.18)
Monetary shock * Low spread	0.03	0.05	0.03	0.06
	(0.06)	(0.12)	(0.06)	(0.12)
N	74143	36651	74080	36596

Note: An observation corresponds to an FOMC date-bond pair. Each column refers to a separate regression. Standard errors are reported in parentheses, and are clustered at the date level.

The "average unexplained spreads" dimension of heterogeneity seems not to have much of an effect on the response of muni yields to monetary shocks, in either the conditional or unconditional specifications. Note that these measures are time invariant and the bond level; while a time-varying measure of excess spread would be helpful, many of these bonds simply aren't traded at a high enough frequency to obtain a meaningful measure.

While the average spread differential doesn't reveal a systematic response in the same way that government type does, it may be the case that a differential response is revealed *within* certain types of governments. To finish the investigation into heterogeneity of responses, I include both the government issuer's type and unexplained spreads in the regression specifications. Results are given in Table A.9.

As before, state governments respond less strongly to monetary shocks, especially conditional on adjustment, and low unexplained spreads have zero effect on the response of yields to a monetary shock. Note, however, the triple-difference coefficient in these specifications. Big governments with low excess spreads do not exhibit a

Table A.9: 2-Week Response of Muni Yields to Monetary Shocks, by Issuer Type and High/Low Spreads

	d yield	d yield ($\neq 0$)	d spread	d spread ($\neq 0$)
Big Government	0.01	0.06***	0.01	0.06***
	(0.01)	(0.01)	(0.01)	(0.01)
Low spread	0.02***	0.05***	0.02***	0.05***
	(0.00)	(0.01)	(0.00)	(0.01)
Monetary shock	0.28***	0.76***	0.08	0.17
	(0.08)	(0.20)	(0.08)	(0.18)
Monetary shock * Big government	-0.21	-0.55**	-0.29**	-0.50***
	(0.14)	(0.23)	(0.12)	(0.20)
Monetary shock * Low spread	-0.06	-0.15	-0.04	-0.10
	(0.06)	(0.14)	(0.07)	(0.16)
Big government * Low spread	0.01	-0.01	0.01	-0.01
	(0.01)	(0.01)	(0.01)	(0.01)
Monetary shock * Big government * Low spread	0.29**	0.51**	0.24*	0.41*
	(0.13)	(0.22)	(0.12)	(0.22)
N	74338	36743	74264	36681

Note: An observation corresponds to an FOMC date-bond pair. Each column refers to a separate regression. Standard errors are reported in parentheses, and are clustered at the date level.

lower response to monetary shocks, even conditional on adjustment; this coefficient almost completely negates the negative coefficient on being a state government. In these specifications, the bonds which exhibit a lower response to monetary shocks are only those which are issued by state governments with high excess spreads.

The results of this section, to the extent they say anything systematic about heterogeneity in the response of municipal bond yields to monetary shocks, may be summarized as follows. Bonds issued from state governments ("big" governments) are traded more often, and therefore are more likely to experience a price change as a result of a monetary shock. Unsurprisingly, the magnitude of their responses are smaller on average. This lower response is mainly driven by state bonds with high excess spreads, reflecting higher illiquidity or potentially higher risk premia than other bonds from similar issuers.

A.4 Additional Quantitative Results

A.4.1 Full Model Results

In the body of the paper, I show the effects of changing the coefficient θ^G on the monetary transmission to each locality. I focused on one locality at a time for ease of exposition. Here, I consider the full model with a large number of localities, and confirm that monetary shocks transmit in a heterogeneous manner when localities' financial markets differ.

In the data, I identified the municipal yield responses to monetary shocks of 41 U.S. states. In this specification, then, I set $S = 41$, and assign each state a θ^G corresponding to one of the empirical estimates, while all other parameters remain the same. Aggregate output is defined as total output, and aggregate inflation as average inflation of non-tradables across localities. The risk-free rate is set according to a Taylor rule which mimics the interest rate process in the body of the paper, replacing y^T with total output and adding in a response of inflation, where the risk-free rate responds to a basis point of lagged inflation with a 1.5 basis point rate increase.

Figure A.6 shows the responses of percentiles of the full economy to a 25bp shock to the risk-free rate set in the Taylor Rule. Consistent with the results in the paper, realistic heterogeneity in monetary passthrough to munis results in significant transmission to U.S. localities. Here, also, we see that the difference between the 10th percentile and the median is much greater than that between the 90th and median. This is reflective of the asymmetry of the distribution of coefficient estimates in the

Figure A.6: IRFs Under Muni Heterogeneity

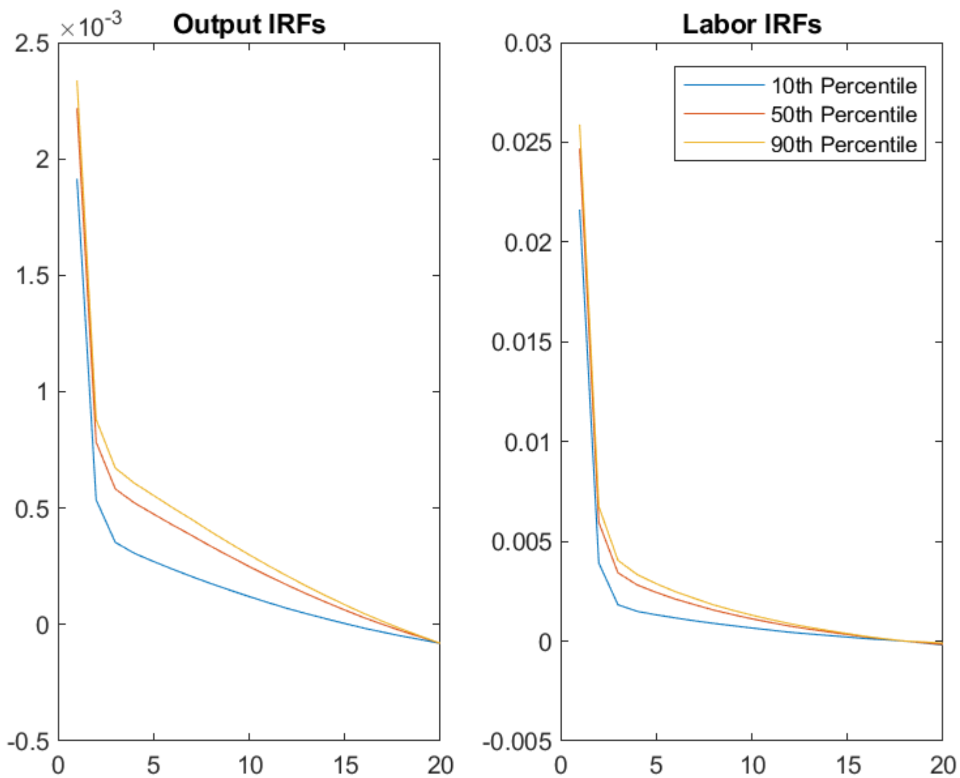

data.

A.4.2 OTC Model Results

Here, I present some brief results from the OTC version of the model described in Section 1.2.8 of the main paper. Calibration moves forward similarly as in the baseline model, with a few additions. First, for simplicity I assume that financial firms are immediately able to sell bonds on the secondary market, i.e., $p^{sell} = 1$. A fuller examination of heterogeneity in muni pricing would require a more specific specification of the thickness of this secondary market, but in this section I will show simply the basic results. I calibrate v^H and ν to match the steady state d^G and r^G from the baseline model, and set $\theta = 0.5$ for simplicity.[4]

Figure A.7: IRFs, 25bp Expansionary Shock

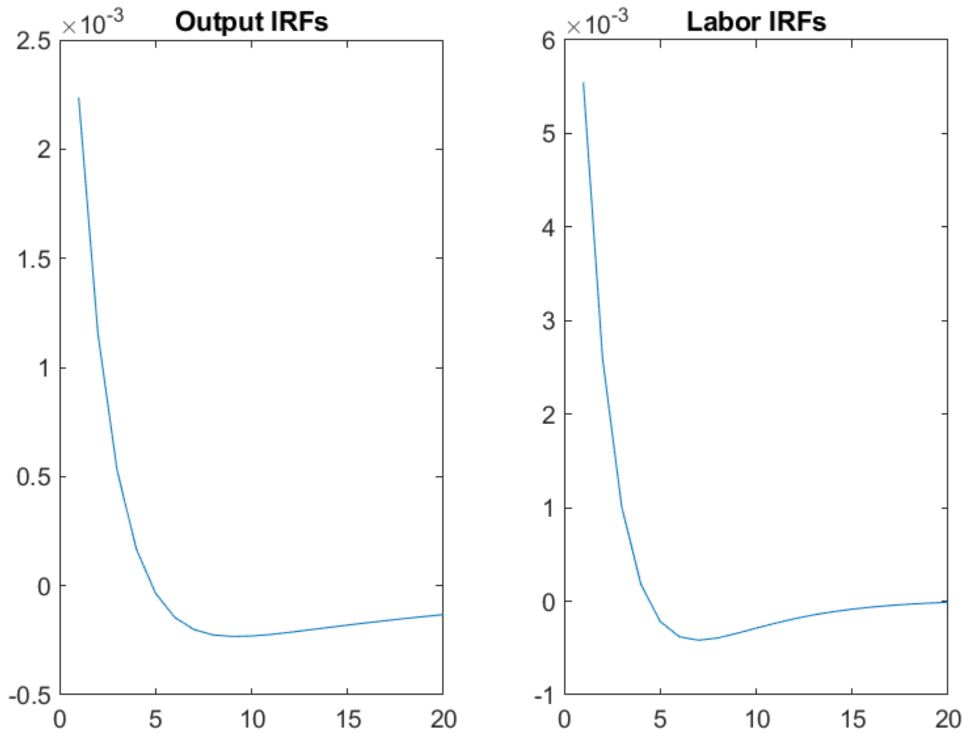

[4]Monetary responses do vary with θ, but without variation in α cannot match the observed heterogeneity in the data.

Figure A.8: IRFs, 25bp Expansionary Shock

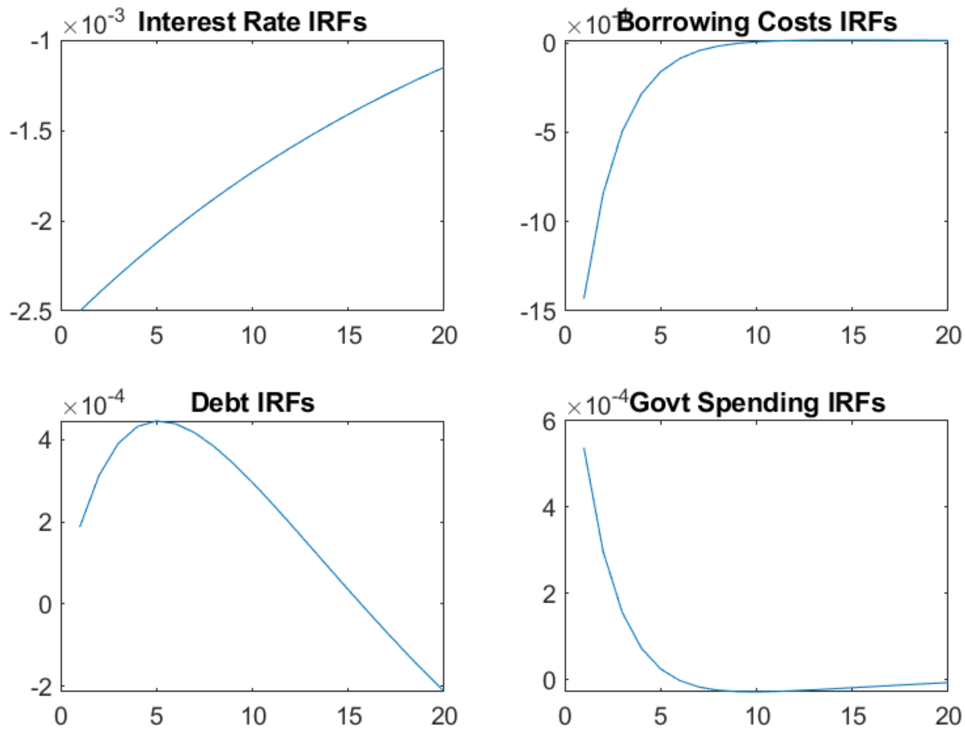

This model behaves similarly to the baseline model in terms of the directions of the IRFs. Note, however, in this special case of the model, the response of government debt and spending is much lower than the baseline model. The reason for this is the costs involved in issuing debt: the explicit inclusion of coupon payments and debt retirements puts upward pressure on their borrowing costs. As a result, $\frac{\partial r^G}{\partial d^G}$ is quite large, dampening significantly the responses to monetary shocks. In the baseline model, this translates to a larger value for ϕ^G.

A.4.3 Explicit Debt Constraints

This paper examines in detail the effect of municipal bond markets' response to monetary shocks on the size and potential heterogeneity in monetary policy transmission. Another fiscal dimension on which state and local governments differ is the stringency of balanced-budget rules, which vary across governments. Most governments have some sort of balanced budget requirement on the books; the rules surrounding these requirements likely result in an effective *politically imposed* on the amount of debt a government can issue.

Figure A.9: The Effect of Debt Constraints

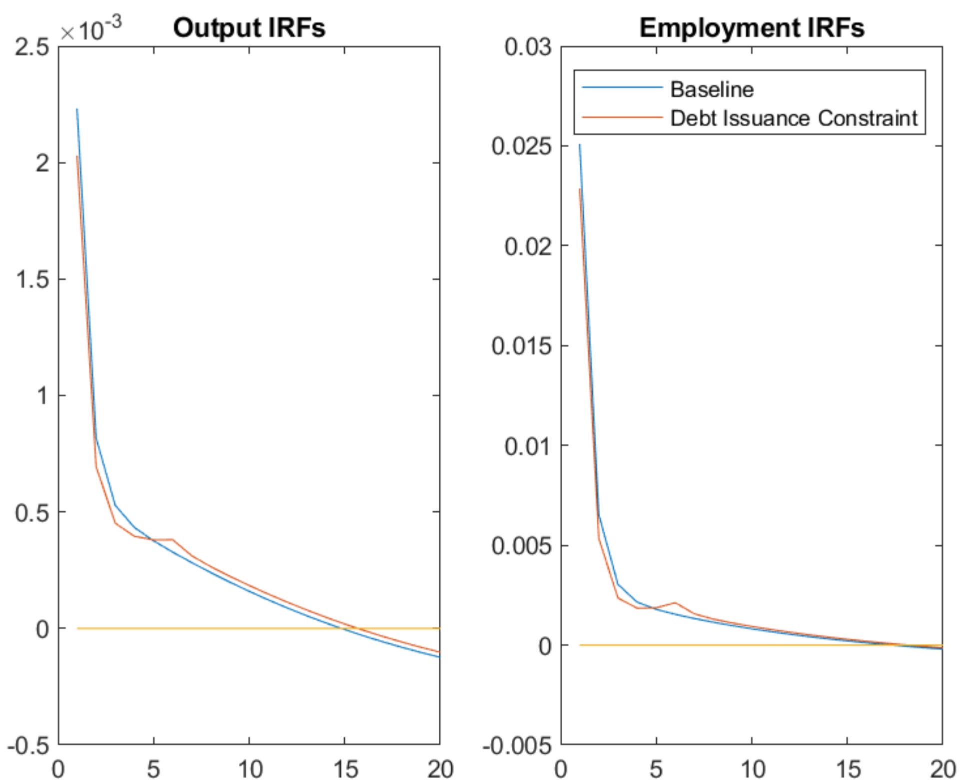

Figure A.10: IRFs by Response Elasticity, Debt Issue Constrained

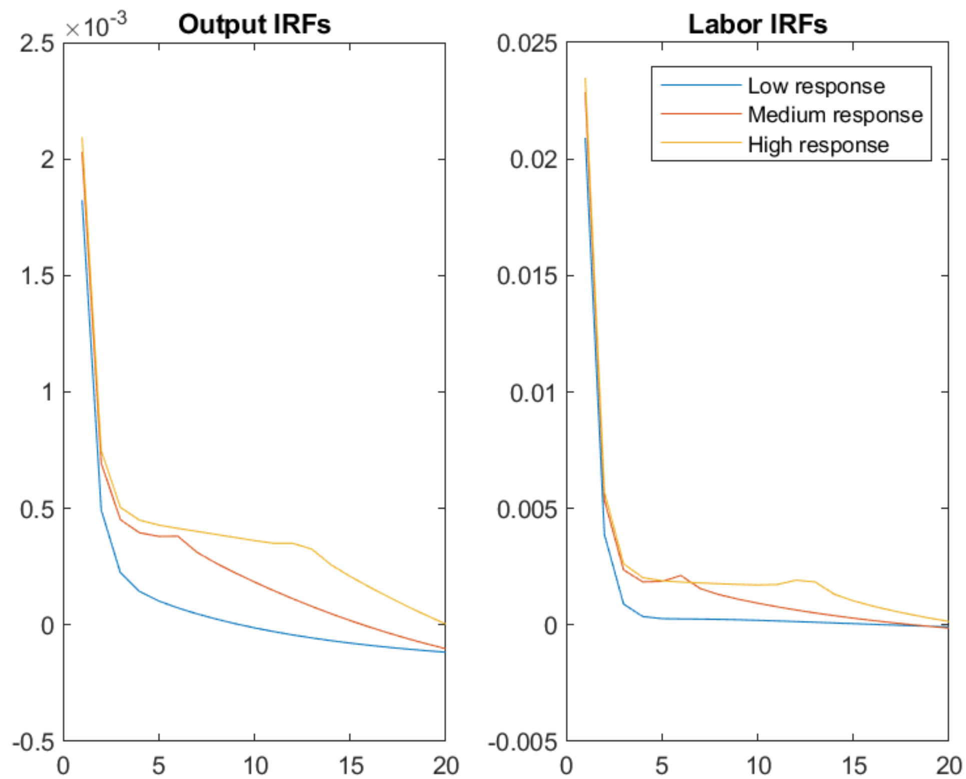

Figures A.9 and A.10 show how the results in the paper are affected by an additional constraint on debt issued by a local government: $d^G_{it} \leq 1.025 d^G_{i,t-1}$. First, note that the debt constraint dampens the transmission of monetary policy, even relative to the already-dampened baseline case, on the order of about 10 percent in the baseline case. This dampening seems to be stronger for the economies with more responsive debt prices (Figure A.10), making the medium- and high- response economies almost equal. Just as the response of borrowing costs to monetary shocks affects monetary transmission through local fiscal policy, so too will the debt issuance constraints placed on these localities through various budgeting laws.

A.4.4 Steady State Government Debt

Figure A.11 shows the positive relationship between steady state local government debt and the transmission of monetary policy in the baseline calibrated model. Here I recalibrate the parameters of the model in the same way as the body of the paper for a number of values for steady state local government debt as a fraction of output, from 0.05 to 0.15. I record the on-impact transmission of monetary policy for each of these economies in the figure.

Clearly, there is a positive correlation between the steady state level of government debt and monetary transmission. The same percentage increase in government spending will be more stimulative for a government which spends more in steady state. The steady state level of government spending is one of many possible dimensions on which these governments may differ, and which may contribute to the transmission of monetary policy.

A.4.5 Fiscal Policy Shocks

In the body of the paper, I explore the interactions between national monetary policy and local fiscal policy. A natural related question is how national fiscal policy might enter into the small open economy model, and what effects it might have given the problems faced by the local government. Figure A.12 shows the response of the baseline economy to two different federal government spending shocks. In the first, public goods spending exogenously increases, but the public goods come from outside

Figure A.11: Monetary Transmission and Government Debt

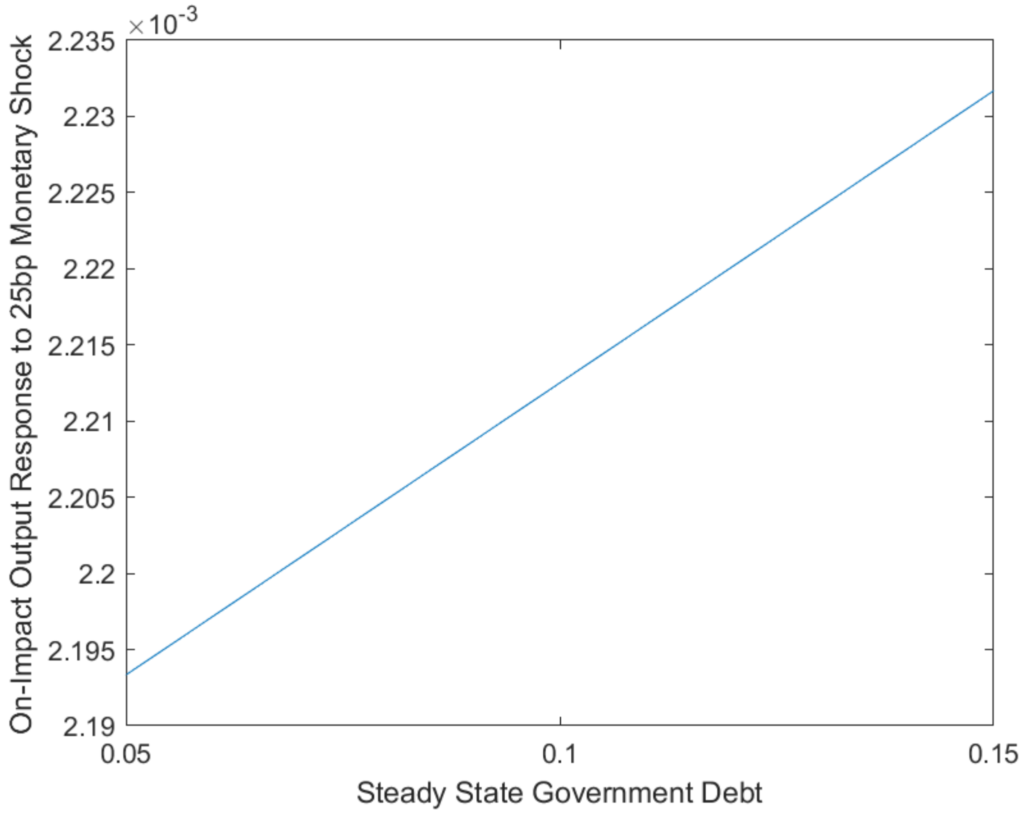

Figure A.12: Federal Government Spending Shock (5 Percent Increase)

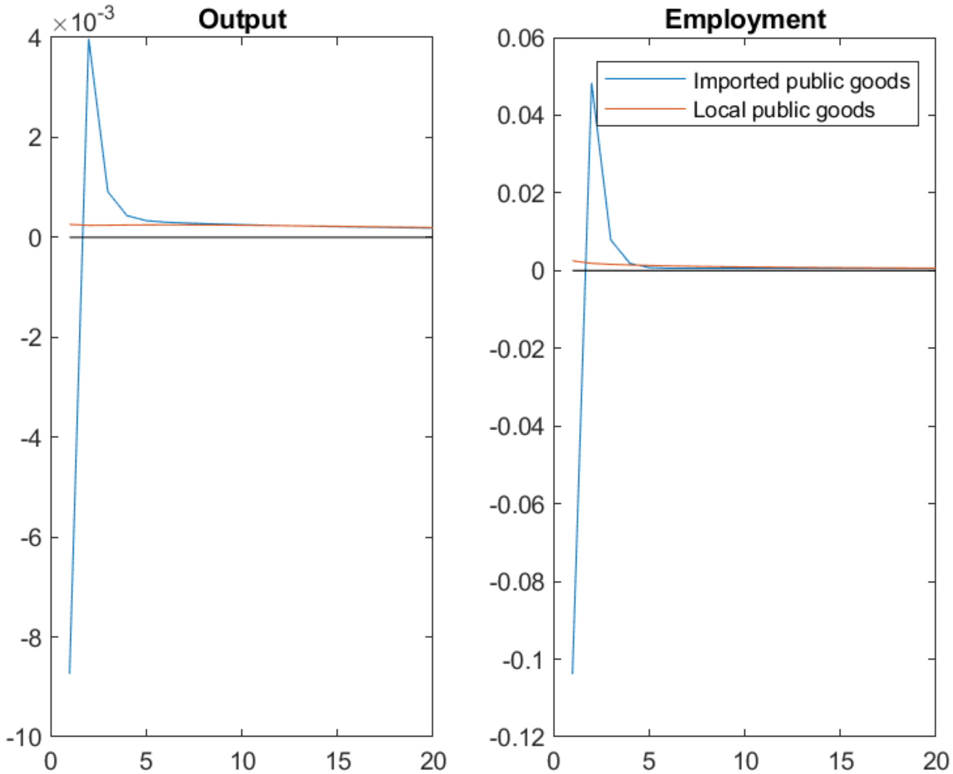

the local economy; i.e., the federal government does not purchase g from local goods producers. In the second case, the federal government purchases public goods from the local nontradable sector.

The fiscal shock in each case corresponds to a five percent (transitory) increase over the steady state level of government spending. In both cases, federal spending crowds out local spending as the local government simply wants to provide efficient levels of public goods given the household's preferences. When federal spending is imported from elsewhere in the economy, local output decreases markedly as the local government decreases its spending, lowering aggregate demand and employment. This lack of need to spend allows the local government to pay down debt, allowing for increased spending in the future, and an expansion of local aggregate demand. In the second case, the federal spending represents an increase in aggregate demand on impact, as public goods are now being bought by the federal government. Subsequently, local spending increases as before; the subsequent spending is muted in this case because of the inflationary effect of the initial federal shock, which is not present in the first case.

APPENDIX B

Chapter 2 Appendix

B.1 Optimal Linear Taxes and Minimum Wage

When the federal government has access to linear income taxes, it will always prefer linear taxes to a combination of a minimum wage and linear taxes under the assumptions of our calibrated model.

Consider a binding minimum wage and tax policy $\bar{w}, t \geq 0$. Under the assumptions of our model, mainly a perfectly competitive firm, this set of policy will induce additional unemployment over the competitive equilibrium so that $l(\bar{w}, t)$ low-skilled workers are employed. The minimum wage reduced labor demand while the tax reduced labor supply. The workers earn $\bar{w}$ and high-skilled workers earn $w_h(l(\bar{w}, t))$. The high-skilled wage depends on our structural parameters and the stock of low-skilled employed workers. Now consider $t^\star$ such that $l(0, t^\star) = l(\bar{w}, t)$. It must be the case that $t^\star \geq t$, as a binding minimum wage always reduces low-skilled employment, as long as $l(\bar{w}, t) > 0$. Since l is the same, w_l, w_h, and total output are equal across the policies. However, $t^\star \geq t$ gives that the tax alone raises more revenue. The additional tax revenue is, on net, redistributed to lower consumption agents through the demogrant, increasing total welfare since the social welfare function is concave. As the tax increases, low-skilled labor supply drops, increasing the low-skilled wage and decreasing the high-skilled wage; when it reaches 70%, only 34% of the low-skilled workers are employed, and the low-skilled wage is higher. We may be worried that since the low-skilled face a working cost, they can be high-income but low-consumption and the additional redistribution is harmful for social welfare. However, at that point, income tax collected is redistributed on net to the unemployed from the employed. Since the consumption of the unemployed is always bounded above by the consumption of the employed, even with this tax and demogrant, the tax policy always redistributes to the lower consumption agents.

We need to show that $t^\star$ exists. As long as the low-skilled labor supply is not perfectly inelastic, it will exist because as $l \to 0$, the marginal product goes to infinity, but increasing the tax rate will weakly lower labor supply. When the labor supply is perfectly inelastic (because the low-skilled workers do not face costs from working), then the government can implement a levelling tax, since high-skilled labor supply is also perfectly inelastic to the economy. The value of the levelling tax is maximized when output and labor demand is maximized, i.e. when there is no binding minimum wage.

So the government will always prefer only a linear income tax when both a linear income tax and minimum wage are available in our model.

APPENDIX C

Chapter 3 Appendix

C.1 Construction of Balances Data

I obtain the data on state balances and rainy day funds from "The Fiscal Survey of the States," a semiannual report published by the National Association of State Budget Officers (NASBO). When available, I use the spring edition of the report for consistency, and I report the "actual" numbers from the previous year, not the "estimated" number for the current year. This report surveys budget officers in U.S. states in order to detail the fiscal health of states across various dimensions. I utilize this specific series because it provides states with the option to report both general fund balances and rainy day fund balances separately.

Unfortunately, the data reported in the Fiscal Survey of the States is not available in electronic format, and must be recorded by hand for each year. The report is released in PDF format every year, and each report gives an end-year balance for the general fund and the rainy day fund for all 50 states. Recording the balances data for the states involves locating the relevant table in each document and pulling the data for each state into a common spreadsheet for analysis. The format tends to be similar across years, although it does change from time to time, with greater frequency of format changes in earlier years.

The basic strategy is to record two figures: the "ending balance" for each state, which is the surplus (or deficit) of the state's general fund, and the rainy day fund balance. Then "total balances" are reported as the sum of these two numbers. In more recent documents, NASBO constructs and reports total balances but not ending balances, so the data collection involves the collection of total balances and rainy day fund balances. Figures C.1, C.2, and C.3 show the difference between the 2017 report and the 1996 report. In 2017, total balances are reported and therefore do not need to be constructed; in 1996, however, ending balances are reported and total balances need to be constructed.

Note also that certain states report the rainy day fund as a part of their general fund. For these states, I am not able to back out the value of the rainy day fund.[1] For this reason, and because a few states don't report anything about rainy day funds, I include total balances as a balance of interest for the analysis.

So the overall procedure for collecting the balances data is as follows. First, collect the data for rainy day funds balances for each state/year observation, not including the values reported in footnotes. Second, collect the ending balances, some

[1]Sometimes there is a footnote with the value of the rainy day fund. I include these figures with the rainy day fund data when possible.

Figure C.1: NASBO Data Format, Balances 2017

TABLE 28

Total Balances and Total Balances as a Percentage of Expenditures, Fiscal 2016 to Fiscal 2018

	Total Balances ($ in Millions)			Total Balances as a Percent of Expenditures		
State	Fiscal 2016	Fiscal 2017	Fiscal 2018	Fiscal 2016	Fiscal 2017	Fiscal 2018
Alabama	$715	$814	$918	9.2%	9.8%	11.0%
Alaska*	7,120	7,033	6,310	130.9	158.4	145.6
Arizona	745	582	485	7.8	6.0	5.0
Arkansas	0	0	0	0.0	0.0	0.0
California*	8,553	7,741	10,404	7.5	6.3	8.5
Colorado***	513	588	676	5.0	5.6	6.1
Connecticut*	236	259	261	1.3	1.4	1.4
Delaware***	568	428	456	14.5	10.5	11.1
Florida	3,246	2,939	2,761	11.1	9.6	8.9
Georgia***	2,131	2,131	2,131	9.7	9.3	9.0
Hawaii	1,129	832	743	16.4	10.8	10.0
Idaho	310	374	361	10.2	11.4	10.4
Illinois*	534	133	133	2.0	0.4	0.4
Indiana	2,244	1,731	1,834	15.0	11.3	11.8
Iowa	773	608	709	10.7	8.3	9.7
Kansas	37	100	217	0.6	1.6	3.5
Kentucky	490	351	179	4.8	3.2	1.6
Louisiana	45	261	286	0.5	2.7	3.0
Maine	193	243	199	5.8	7.1	5.8
Maryland	1,217	926	960	7.5	5.4	5.6
Massachusetts***	1,482	1,345	1,410	3.7	3.2	3.3
Michigan	1,217	1,075	1,012	12.6	10.6	10.0
Minnesota***	3,102	2,720	2,328	15.4	12.6	10.3
Mississippi	356	337	557	6.2	5.8	9.7
Missouri	445	476	414	4.9	5.2	4.3
Montana	255	123	163	11.0	5.2	6.9
Nebraska	1,262	809	707	30.1	18.6	16.1
Nevada	418	484	298	11.6	12.5	7.2
New Hampshire	182	100	97	13.1	6.9	6.5
New Jersey	473	491	493	1.4	1.4	1.4
New Mexico***	146	-67	11	2.3	-1.1	0.2
New York***	8,934	7,232	5,917	13.1	10.4	8.2
North Carolina	2,155	1,794	1,787	10.2	8.1	7.6
North Dakota	983	11	57	32.7	0.4	2.5
Ohio	3,198	2,728	2,200	9.3	7.7	6.5
Oklahoma*	241	N/A	N/A	3.9	N/A	N/A
Oregon	811	1,119	1,021	9.0	12.3	10.6
Pennsylvania	2	-606	4	0.0	-1.9	0.0
Rhode Island	359	272	196	10.1	7.4	5.2
South Carolina***	1,131	856	890	15.8	10.9	11.6
South Dakota	157	160	160	10.8	10.1	9.9
Tennessee	1,958	1,838	801	15.5	13.4	5.5
Texas	14,047	11,783	11,639	26.6	22.3	22.5
Utah	658	501	502	10.4	7.8	7.6
Vermont	78	93	124	5.3	6.1	8.1
Virginia	501	626	297	2.6	3.1	1.5
Washington	1,922	1,951	1,592	10.6	10.0	7.8
West Virginia	1,150	871	853	27.5	19.8	19.2
Wisconsin	612	735	600	3.9	4.3	3.5
Wyoming	1,811	1,481	1,481	109.7	103.0	101.9
Total**	**$80,845**	**$69,409**	**$67,633**	**10.3%**	**8.5%**	**8.2%**
			Median	**9.9%**	**7.7%**	**7.6%**

*NOTES: Total balances include both the ending balance and Rainy Day Funds. Fiscal 2016 are actual figures, fiscal 2017 are estimated figures, and fiscal 2018 are recommended figures. N/A indicates data not available. *See notes to Table 28 on page 65. **Fiscal 2017 and Fiscal 2018 total balance amount and total balances as percentage of expenditures exclude Oklahoma, as complete data for these states was not available for this year. ***Ending Balance includes Rainy Day Fund.*

Figure C.2: NASBO Data Format, RDFs 2017

TABLE 29

Rainy Day Fund Balances and Rainy Day Fund Balances as a Percentage of Expenditures, Fiscal 2016 to Fiscal 2018

	Rainy Day Fund Balances ($ in Millions)**			Rainy Day Fund Balances as a Percent of Expenditures		
State	Fiscal 2016	Fiscal 2017	Fiscal 2018	Fiscal 2016	Fiscal 2017	Fiscal 2018
Alabama	$530	$766	$786	6.8%	9.3%	9.4%
Alaska*	7,120	7,033	6,310	130.9	158.4	145.6
Arizona	461	463	468	4.8	4.8	4.8
Arkansas	N/A	N/A	N/A	N/A	N/A	N/A
California*	7,573	6,761	9,424	6.6	5.5	7.7
Colorado	513	588	676	5.0	5.6	6.1
Connecticut*	236	259	261	1.3	1.4	1.4
Delaware	215	221	221	5.5	5.4	5.4
Florida	1,354	1,384	1,417	4.6	4.5	4.6
Georgia*	2,033	N/A	N/A	9.3	N/A	N/A
Hawaii	101	311	317	1.5	4.0	4.3
Idaho	259	259	293	8.5	7.9	8.5
Illinois	275	0	0	1.0	0.0	0.0
Indiana	1,468	1,472	1,476	9.8	9.6	9.5
Iowa	729	607	604	10.1	8.3	8.3
Kansas	N/A	N/A	N/A	N/A	N/A	N/A
Kentucky	209	236	179	2.0	2.1	1.6
Louisiana	359	261	286	4.1	2.7	3.0
Maine	122	168	173	3.7	4.9	5.1
Maryland	832	833	860	5.2	4.9	5.0
Massachusetts	1,292	1,303	1,401	3.2	3.1	3.2
Michigan	612	709	1,004	6.3	7.0	9.9
Minnesota	1,947	1,953	1,953	9.7	9.0	8.6
Mississippi	350	337	440	6.1	5.8	7.7
Missouri	291	294	309	3.2	3.2	3.2
Montana	N/A	N/A	N/A	N/A	N/A	N/A
Nebraska	731	546	503	17.4	12.6	11.4
Nevada	0	64	62	0.0	1.7	1.5
New Hampshire	93	100	100	6.7	6.9	6.7
New Jersey	0	0	0	0.0	0.0	0.0
New Mexico	146	-67	0	2.3	-1.1	0.0
New York	1,798	1,798	1,948	2.6	2.6	2.7
North Carolina	1,575	1,474	1,787	7.4	6.6	7.6
North Dakota	573	0	0	19.0	0.0	0.0
Ohio	2,005	2,034	2,034	5.8	5.7	6.0
Oklahoma*	241	N/A	N/A	3.9	N/A	N/A
Oregon	550	771	952	6.1	8.5	9.9
Pennsylvania	0	0	0	0.0	0.0	0.0
Rhode Island	192	194	196	5.4	5.3	5.2
South Carolina	505	487	509	7.0	6.2	6.6
South Dakota	143	157	160	9.8	9.9	9.9
Tennessee	568	668	800	4.5	4.9	5.5
Texas	9,715	10,254	10,972	18.4	19.4	21.2
Utah	493	493	493	7.8	7.7	7.5
Vermont	78	93	124	5.3	6.1	8.1
Virginia	236	549	281	1.2	2.7	1.4
Washington	550	1,340	1,350	3.0	6.9	6.6
West Virginia	779	635	617	18.7	14.4	13.9
Wisconsin	281	282	302	1.8	1.7	1.8
Wyoming	1,811	1,481	1,481	109.7	103.0	101.9
Total**	**$51,942**	**$49,572**	**$53,530**	**6.6%**	**6.3%**	**6.7%**
Median				**5.4%**	**5.5%**	**6.0%**

*NOTES: N/A indicates data not available. Fiscal 2016 are actual figures, fiscal 2017 are estimated figures, and fiscal 2018 are recommended figures. *See Notes to Table 29 on page 65. **Total Rainy day fund balances for fiscal 2017 and fiscal 2018 exclude Georgia and Oklahoma, as data were unavailable for these years.*

Figure C.3: NASBO Data Format, 1996

TABLE A-1

Fiscal 1995 State General Fund, Actual (Millions)

Region/State	*Beginning Balance*	*Revenues*	*Adjustments*	*Resources*	*Expenditures*	*Adjustments*	*Ending Balance*	*Budget Stabilization Fund*
NEW ENGLAND								
Connecticut*	$ 0	$ 8,480		$ 8,460	$ 8,399		$ 61	$ 61
Maine*	4	1,672	$ 37	1,712	1,687	$ 26	4	10
Massachusetts*	125	15,798		15,923	15,705		179	425
New Hampshire	12	963		975	971	4	0	24
Rhode Island*	4	1,643		1,646	1,641		5	45
Vermont*	0	673	3	675	690		-15	0
MID-ATLANTIC								
Delaware*	313	1,602		1,915	1,541		374	*
Maryland	65	7,068		7,133	7,000		133	286
New Jersey*	1,204	14,898		16,138	14,947	240	952	*
New York*	399	32,296	862	33,557	33,399		158	*
Pennsylvania*	302	15,765	148	16,215	15,732	-54	429	66
GREAT LAKES								
Illinois*	230	17,302	-300	17,232	17,201	-300	331	0
Indiana*	90	7,307	-30	7,367	6,332	356	679	419
Michigan*	0	7,995	44	8,040	8,041	-2	0	1,003
Ohio*	300	15,711		16,011	14,979	962	70	828
Wisconsin*	282	7,946		8,228	7,827		401	*
PLAINS								
Iowa*	0	3,907		3,907	3,616		292	116
Kansas*	454	3,219	4	3,677	3,310		367	5
Minnesota*	903	8,759		9,662	8,605		1,057	*
Missouri	275	5,459		5,734	5,261		473	24
Nebraska*	152	1,706	1	1,858	1,683		176	21
North Dakota*	28	632		660	629		31	0
South Dakota*	0	580	41	622	589	33	0	11
SOUTHEAST								
Alabama	128	4,078		4,206	4,151		54	0
Arkansas*	0	2,400	52	2,453	2,453		0	0
Florida	198	14,179		14,377	14,248		129	282
Georgia*	120	9,625		9,745	9,500		224	288
Kentucky	98	5,188	125	5,411	5,006	144	261	100
Louisiana*	0	4,784	10	4,794	4,729	-132	146	0
Mississippi*	166	2,624		2,790	2,675		115	268
North Carolina*	888	9,972		10,860	10,034		892	*
South Carolina*	407	4,234		4,641	4,051		589	*
Tennessee*	173	5,076	90	5,339	5,174	27	138	*
Virginia*	334	7,174		7,507	7,490		17	80
West Virginia*	69	2,309	3	2,380	2,210	43	127	64
SOUTHWEST								
Arizona	229	4,466		4,695	4,425		270	223
New Mexico*	156	2,692	-60	2,788	2,714	15	0	59
Oklahoma	118	3,513		3,631	3,436		195	45
Texas*	1,929	20,563		22,492	20,640		1,852	9
ROCKY MOUNTAIN								
Colorado*	405	3,996		4,402	3,914	4	484	*
Idaho*	38	1,288	-55	1,271	1,268		3	33
Montana*	50	938	7	995	948		47	NA
Utah	37	2,365		2,402	2,341		61	66
Wyoming*	22	445	35	502	476		26	55
FAR WEST								
Alaska*	0	2,489	83	2,572	2,572		0	2,136
California*	109	42,710	-175	42,644	41,961		683	*
Hawaii	291	2,969		3,259	3,169		90	0
Nevada*	129	1,206	165	1,500	1,103	295	102	100
Oregon*	439	3,390		3,829	3,333		496	*
Washington*	402	8,534	107	9,043	8,484		559	0
TERRITORIES								
Puerto Rico	255	5,211		5,466	5,340		126	82
Total	**$12,077**	**$354,584**	--	**$367,893**	**$352,291**	--	**$13,736**	**$7,171**

NOTE: NA indicates data are not available.

*See Notes to Table A-1.

of which include rainy day fund balances. Third, for those observations whose ending balances include a rainy day fund balance which is reported in a footnote, subtract this value from the ending balance data and add it to the rainy day fund series. Next, construct total balances by adding the ending (less rainy day) and rainy day series together. Finally, replace values in the total balances series with reported total balances from NASBO in years for which these are reported. This is a labor intensive process, but the only way to obtain these data, given the format in which they are published.

C.2 Determination of Transfer Equation Using IV

Equation 3.1 relates the transfers from the federal government received by state government i to, among other things, gross state product of state i and the sum of the gross state products of other states $-i$. A clear potential source of endogeneity in this equation is the effect of transfers on output, especially since these data are only available at the annual frequency. As a robustness check, I instrument $y_{i,t}$ and $y_{-i,t}$ with monetary policy shocks at the annual level. I use the annual shocks from Weiland and Yang (2019), which correspond to the shocks in Romer and Romer (2004), updated through 2007. These shocks are publicly available online at Johannes Weiland's webpage.[2]

Table C.1 shows the results of the instrumental variables regression. The response of transfers to the business cycle is much greater for $y_{-i,t}$, and the sign on $y_{i,t}$ flips, suggesting that perhaps the main results underestimate the true difference

[2] https://sites.google.com/site/johannesfwieland/

Table C.1: Determinants of State Government Receipts from Federal Government, IV

Variable	$\log(pop)$	$\log(GSP_i)$, cyclical	$\log(GSP_{-i})$, cyclical
Coefficient	0.1728*	3.314*	-8.290**
(s.e.)	(0.0915)	(1.937)	(4.128)

Note: Results from a fixed-effects instrumental variables regression with standard errors clustered at the state level. The instruments are monetary shocks described in the text. Observations include 46 states for which holes do not exist in the Census of Governments data from 1981 to 2012. * $p \leq 0.10$, ** $p \leq 0.05$, *** $p \leq 0.01$.

in reponses. Furthermore, the qualitative relationship between the two equations holds: fiscal transfers from the federal government respond more strongly to aggregate business cycle conditions than local conditions, even though the transfers are to specific state governments. These regressions should be taken with a grain of salt, however, given that the instruments may not be incredibly strong.

C.3 Extended Regression Results

This section provides the full results from the regressions provided in the main body of the paper. The extra variables in these tables are listed and explained below:

- Log(debt): The log of a state government's real debt
- Debt / GSP: The stock of a state government's debt as a share of its gross state product.
- Governor is democrat: = 1 if the state's governor is a member of the Democratic party, reported in the Center for Poverty Research data.

- Variance of cyclical GSP: the variance over time of the state's HP-filtered gross state product, a time-invariant variable for each state.

- Balanced budget strictness: measure of the strictness of the state's balanced budget rules, constructed from Hou and Smith (2006), which classify 9 possible restrictions on budgets. Sums the rules each state has adopted, weighting the 'stricter' rules more heavily.

Table C.2: Determinants of State Government Balances

Dependent variable	log(∼ Real balances)			Real balances / GSP			Real balances / Expenditures		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Business cycle correlation	-0.0920**	-0.0825*	-0.0847*	-0.0054***	-0.0038**	-0.0038**	-0.0412***	-0.0262**	-0.0262**
	(0.0412)	(0.0466)	(0.0468)	(0.0016)	(0.0015)	(0.0015)	(0.0144)	(0.0128)	(0.0127)
Log(GSP), cyclical	0.5505***	0.4131**	0.1138	0.0225***	0.0177***	0.0167***	0.3036***	0.2536***	0.2487***
	(0.1886)	(0.1885)	(0.0789)	(0.0042)	(0.0050)	(0.0063)	(0.0377)	(0.0454)	(0.0547)
Log(GSP), cyclical * High correlation	-	-	0.9854*	-	-	0.0031	-	-	0.0155
			(0.5528)			(0.0068)			(0.0652)
Log(GSP)	0.0537***	0.0570***	0.0577***	-	-	-	-	-	-
	(0.0124)	(0.0131)	(0.0132)						
Log(debt)	-	0.0041	0.0020	-	-	-	-	-	-
		(0.0065)	(0.0072)						
Debt / GSP	-	-	-	-	6.29e-6	6.30e-6	-	2.76e-5	2.76e-5
					(4.46e-6)	(4.47e-6)		(4.47e-5)	(4.49e-5)
Governor is Democrat	-	-0.132*	-0.0120*	-	-0.0002	-0.0002	-	-0.0027	-0.0027
		(0.0068)	(0.0069)		(0.0003)	(0.0015)		(0.0028)	(0.0028)
Variance of cyclical GSP	-	9.335**	9.572**	-	0.8860***	0.8870***	-	9.260***	9.265***
		(4.308)	(4.240)		(0.2214)	(0.2215)		(2.199)	(2.199)
Balanced budget strictness	-	0.0001	3.61e-5	-	-5.23e-6	-5.27e-6	-	5.16e-5	5.14e-5
		(0.0011)	(0.0010)		(4.51e-5)	(4.52e-5)		(0.0005)	(0.0005)

Note: The sample for these regressions is 49 U.S. states (Alaska not included) for the years 1981-2012. The regressions are according to a random effects model. Standard errors are clustered at the state level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table C.3: Determinants of State Government Net Assets (Not Incl. Insurance Funds)

Dependent variable	log(~ Real net assets)			Real net assets / GSP			Real net assets / Expenditures		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Business cycle correlation	-1.2630***	-1.2365***	-1.2320***	-0.1168***	-0.0605**	-0.0607**	-1.0787***	-0.5528***	-0.5541***
	(0.3325)	(0.3857)	(0.3845)	(0.0393)	(0.0267)	(0.0267)	(0.3393)	(0.2115)	(0.2118)
Log(GSP), cyclical	0.9278**	0.5754	0.0538	0.0513*	0.0581*	0.0202	0.9424***	1.0980***	0.9055***
	(0.3416)	(0.6129)	(0.6406)	(0.0274)	(0.0328)	(0.0381)	(0.1349)	(0.1655)	(0.2011)
Log(GSP), cyclical * High correlation	-	-	1.8011**	-	-	0.1196***	-	-	0.6059**
			(0.7729)			(0.0413)			(0.2714)
Log(GSP)	0.1993**	0.3001***	0.2979***	-	-	-	-	-	-
	(0.0834)	(0.0857)	(0.0852)						
Log(debt)	-	-0.0761	-0.801	-	-	-	-	-	
	-	(0.0582)	(0.0578)						
Debt / GSP	-	-	-	-	-5.42e-6	-4.91e-6	-	-0.0003	-0.0002
	-				(2.31e-5)	(2.36e-5)		(0.0002)	(0.0002)
Governor is Democrat	-	-0.0152	-0.0135	-	-0.0031	-0.0029	-	-0.0205	-0.0197
		(0.0482)	(0.0478)		(0.0025)	(0.0025)		(0.0272)	(0.0272)
Variance of cyclical GSP	-	41.87	41.80	-	25.64***	25.68***	-	237.8***	238.0***
		(43.14)	(42.86)		(5.479)	(5.486)		(60.83)	(60.82)
Balanced budget strictness	-	-0.0190	-0.0191	-	-0.0018***	-0.0018***	-	-0.0173***	-0.173***
		(0.0119)	(0.0118)		(0.0007)	(0.0007)		(0.0066)	(0.0066)

Note: The sample for these regressions is 49 U.S. states (Alaska not included) for the years 1981-2012. The regressions are according to a random effects model. Standard errors are clustered at the state level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

C.4 Equivalence of Central Government Budget Models

In this section I show that the two alternatives for formulating the budget constraint of the central government with respect to smoothing idiosyncratic shocks result in equivalent transfer policy rules. The first alternative is the no-Ponzi-game condition from Equation 3.8:

$$\lim_{t\to\infty} \mathbb{E}_0\Big[\frac{A_t}{(1+r)^t} = 0\Big].$$

This constraint prevents the central government from accumulating debt indefinitely. This constraint is valid for an economy with any number of regions.

The second alternative assumes that transfers to smooth for idiosyncratic shocks must be paid for in the current period. If regions are indexed by i, this constraint takes the form $\sum_i T_{it} = 0$ for all t, and for an infinite number of regions it is expressed as in Equation 3.9:

$$\int T_{it} di = 0.$$

Theorem C.1. *If all regions are ex ante identical, then the no-Ponzi-game budget constraint is equivalent to the period budget constraint for a continuum of regions, as both result in the condition* $\mathbb{E}_0\big[T(s_t, \theta_t, f_t)\big] = 0$.

Proof. Consider first the period budget constraint with an infinite number of *ex ante* identical, but heterogeneous, regions. By the law of large numbers, in every period the distribution over the state variables will yield densities equal to the long-run probabilities of each state. Since the transfer policy function for idiosyncratic shocks will be the same for every region, its distribution will also equal the long run distribution. This results in $\int T(s_{it}, \theta_{it}, f_{it}) di = \mathbb{E}_0\Big[T(s_{jt}, \theta_{jt}, f_{jt})\Big]$ for any region j, which, when plugged into the budget constraint yields $\mathbb{E}_0\Big[T(s_{jt}, \theta_{jt}, f_{jt})\Big] = 0$ for any region j.

Now observe the no-Ponzi-game budget constraint. This allows the central government to hold unlimited assets A_t for the purposes of smoothing idiosyncratic shocks. If regions are *ex ante* identical, then it is without loss of generality to consider a separate fund A_{it} for each region for the purposes of solving for the transfer function $T(s_{it}, \theta_{it}, f_{it})$, which will be the same in each region. The law of motion for A_{it} is given by $A_{it+1} = A_{it}(1 + r) + (-T(s_{it}, \theta_{it}, f_{it}))$, which can be expanded and solved to yield $A_{it+1} = -\sum_{j=0}^{t} T(s_{it-j}, \theta_{it-j}, f_{it-j})(1+r)^j$. Plugging in to the budget constraint, the condition now becomes

$$\lim_{t\to\infty} \mathbb{E}_0 \left[-\sum_{j=0}^{t-1} T(s_{it-(1+j)}, \theta_{it-(1+j)}, f_{it-(1+j)})(1+r)^{(j-t)} \right] = 0$$

$$\Rightarrow \lim_{t\to\infty} -\sum_{j=0}^{t-1} \mathbb{E}_0 \Big[T(s_{it-(1+j)}, \theta_{it-(1+j)}, f_{it-(1+j)}) \Big] (1+r)^{(j-t)} = 0.$$

This can only be satisfied if $\mathbb{E}_0 \Big[T(s_{it}, \theta_{it}, f_{it}) \Big] = 0$ for all regions i. □

C.5 Long-Run Posterior Variance

In this section I show that a long-run stable value for the posterior variance exists and solve for its value. First, I derive the law of motion for the posterior variance. Beginning with posterior $\hat{\sigma}^2_{\mu,t}$ in time t, the prior for next period is formed as in Equation 3.14: $\sigma^2_{\mu,t+1} = \rho^2_\epsilon \hat{\sigma}^2_{\mu,t} + \sigma^2_{\xi^\epsilon}$. Next period's posterior is then formed according to Equation 3.12: $\hat{\sigma}^2_{\mu,t+1} = \frac{\sigma^2_{\mu,t+1}\sigma^2_{\xi^\theta}}{\sigma^2_{\mu,t+1}+\sigma^2_{\xi^\theta}}$. Combining these two yields the law of motion for the posterior variance:

$$\hat{\sigma}^2_{\mu,t+1} = \frac{(\rho^2_\epsilon \hat{\sigma}^2_{\mu,t} + \sigma^2_{\xi^\epsilon})\sigma^2_{\xi^\theta}}{\rho^2_\epsilon \hat{\sigma}^2_{\mu,t} + \sigma^2_{\xi^\epsilon} + \sigma^2_{\xi^\theta}}. \tag{C.1}$$

Defining the function $h(x) = \frac{(\rho^2_\epsilon x + \sigma^2_{\xi^\epsilon})\sigma^2_{\xi^\theta}}{\rho^2_\epsilon x + \sigma^2_{\xi^\epsilon} + \sigma^2_{\xi^\theta}}$ and taking the limit as $x \to \infty$, it is clear that $h(x)$ converges to a finite value: $\lim_{x\to\infty} h(x) = \sigma^2_{\xi^\theta}$. Therefore, we can argue that $\hat{\sigma}^2_{\mu,t+1} = h(\hat{\sigma}^2_{\mu,t})$ is bounded on $\mathbb{R}^+$. Now assume that $\hat{\sigma}^2_{\mu,0}$ is initialized in a region such that it converges to the fixed point of $h(x)$; this fixed point then defines the long run value, and solving for it is straightforward:

$$\hat{\sigma}^2_{\mu,\infty} = \frac{(\rho^2_\epsilon \hat{\sigma}^2_{\mu,\infty} + \sigma^2_{\xi^\epsilon})\sigma^2_{\xi^\theta}}{\rho^2_\epsilon \hat{\sigma}^2_{\mu,\infty} + \sigma^2_{\xi^\epsilon} + \sigma^2_{\xi^\theta}}$$

$$\rho^2_\epsilon (\hat{\sigma}^2_{\mu,\infty})^2 + (\sigma^2_{\xi^\epsilon} + \sigma^2_{\xi^\theta} - \sigma^2_{\xi^\theta}\rho^2_\epsilon)\hat{\sigma}^2_{\mu,\infty} - \sigma^2_{\xi^\theta}\sigma^2_{\xi^\epsilon} = 0$$

Solving for the positive solution of the quadratic formula yields the result in Equation 3.15:

$$\hat{\sigma}^2_{\mu,\infty} = \frac{(\rho^2_\epsilon - 1)\sigma^2_{\xi^\theta} + \sigma^2_{\xi^\epsilon} + \sqrt{\left[(\rho^2_\epsilon - 1)\sigma^2_{\xi^\theta} + \sigma^2_{\xi^\epsilon}\right]^2 + 4\rho^2_\epsilon \sigma^2_{\xi^\theta}\sigma^2_{\xi^\epsilon}}}{2\rho^2_\epsilon}.$$

BIBLIOGRAPHY

BIBLIOGRAPHY

Aaronson, Daniel et al. "Industry dynamics and the minimum wage: a puttyclay approach". In: *International Economic Review* 59.1 (2018), pp. 51–84.

Adelino, Manuel, Igor Cunha, and Miguel A. Ferreira. "The Economic Effects of Public Financing: Evidence from Municipal Bond Ratings Recalibration". In: *The Review of Financial Studies* 30.9 (2017), pp. 3223–3268.

Aguiar, Mark and Manuel Amador. "Growth in the Shadow of Expropriation". In: *The Quarterly Journal of Economics* 126.2 (2011), pp. 651–697.

Aiyagari, S. Rao. "Uninsured indiosyncratic risk and aggregate saving". In: *The Quarterly Journal of Economics* (1994).

Allen, Stephen P. "Taxes, redistribution, and the minimum wage: a theoretical analysis". In: *The Quarterly Journal of Economics* 102.3 (1987), pp. 477–489.

Amihud, Yakov. "Illiquidity and stock returns: Cross-section and time series effects". In: *Journal of Financial Markets* 5 (2002), pp. 31–56.

Amior, Michael. "Education and geographical mobility: the role of job surplus". In: *CEP Discussion Paper* (2019).

Anderson, Gareth and Ambrogio Cesa-Bianchi. "Corssing the Credit Channel: Credit Spreads and Firm Heterogeneity". In: *Bank of England Working Paper* 854 (2020).

Ang, Andrew, Vineer Bhansali, and Yuhang Xing. "The Muni Bond Spread: Credit, Liquidity, and Tax". In: *Columbia Business School Research Paper* 14-37 (2014).

Armenter, Roc and Francesc Ortega. "Credible redistribution policy and skilled migration". In: *European Economic Review* 55.2 (2011), pp. 228–245.

Auer, Simone et al. "International monetary policy transmission through banks in small open economies". In: *Journal of International Money and Finance* 90 (2019), pp. 34–53.

Baicker, Katherine, Jeffrey Clemens, and Monica Singhal. "The rise of the states: US fiscal decentralization in the postwar period". In: *Journal of Public Economics* 96.11-12 (2012), pp. 1079–1091.

Banzhaf, H. Spencer and B. Andrew Chupp. "Fiscal federalism and interjurisdictional externalities: New results and an application to US Air pollution". In: *Journal of Public Economics* 96.5-6 (2012), pp. 449–464.

BEA. *National Income and Product Accounts*. 2018. URL: https://apps.bea.gov/iTable/iTable.cfm?isuri=1&reqid=19&step=2&0=survey.

Beetsma, Roel M.W.J. and Henrik Jensen. "Monetary and fiscal policy interactions in a micro-founded model of a monetary union". In: *Journal of International Economics* 67 (2005), pp. 320–352.

Beraja, Martin, Erik Hurst, and Juan Ospina. "The Aggregate Implications of Regional Business Cycles". In: *Econometrica* 87.6 (2019), pp. 1789–1833.

Beraja, Martin et al. "Regional Heterogeneity and the Refinancing Channel of Monetary Policy". In: *The Quarterly Journal of Economics* 134.1 (2019), pp. 109–183.

Bernanke, Ben S., Mark Gertler, and Simon Gilchrist. "The Financial Accelerator in a Quantitative Business Cycle Framework". In: *Handbook of Macroeconomics* 1 (1999), pp. 1341–1393.

Besley, Timothy and Stephen Coate. "Centralized versus decentralized provision of local public goods: a political economy approach". In: *Journal of public economics* 87.12 (2003), pp. 2611–2637.

Bethune, Zachary, Bruno Sultanum, and Nicholas Trachter. "Asset issuance in over-the-counter markets". In: *Review of Economic Dynamics* 33 (2019), pp. 4–29.

Bhandari, Anmol et al. "Fiscal policy and debt management with incomplete markets". In: *Quarterly Journal of Economics* (2017), pp. 617–663.

Bianchi, Javier and Enrique G. Mendoza. "Optimal Time-Consistent Macroprudential Policy". In: *Journal of Political Economy* 126.2 (2018).

Blundell, Richard and Thomas MaCurdy. "Labor supply: A review of alternative approaches". In: *Handbook of labor economics* 3 (1999), pp. 1559–1695.

Boadway, Robin and Katherine Cuff. "A minimum wage can be welfare-improving and employment-enhancing". In: *European Economic Review* 45.3 (2001), pp. 553–576.

Bordignon, Massimo, Paolo Manasse, and Guido Tabellini. "Optimal Regional Redistribution Under Asymmetric Information". In: *American Economic Review* 91.3 (2001), pp. 709–723.

Brancaccio, Giulia, Dan Li, and Norman Schúrhoff. "Learning by Trading: The Case of the US Market for Municipal Bonds". In: *Working paper* (2020).

Brown, Charles C. and Wallace E. Oates. "Assistance to the Poor in a Federal System". In: *Journal of Public Economics* 32 (1987), pp. 307–330.

Brown, Jason P. *Identifying State-Level Recessions*. Economic Review. Federal Reserve Bank of Kansas City, 2017.

Bu, Chunya, John H. Rogers, and Wenbin Wu. "A Unified Measure of Fed Monetary Policy Shocks". In: *FEDS Working Paper* 2019-043 (2 2019).

Bureau, U.S. Census. *Census of Governments, 1970-2016*. 2012.

Cadena, Brian C. "Recent immigrants as labor market arbitrageurs: evidence from the minimumm wage". In: *Journal of Urban Economics* 80 (2014), pp. 1–12.

Calabrese, Stephen M., Dennis Epple, and Richard Romano. "Inefficiencies from Metropolitan Political and Fiscal Decentralization: Failures of Tiebout Competition". In: *The Review of Economic Studies* 79.3 (2012), pp. 1081–1111.

Caplan, Arthur J., Richard C. Cornes, and Emilson CD Silva. "Pure public goods and income redistribution in a federation with decentralized leadership and imperfect labor mobility". In: *Journal of Public Economics* 77.2 (2000), pp. 265–284.

Carlino, Gerald A. and Robert P. Inman. "Local deficits and local jobs: Can US states stabilize their own economies?" In: *Journal of Monetary Economics* 60.5 (2013), pp. 517–530.

Cashin, David et al. "Fiscal policy and aggregate demand in the USA before, during, and following the Great Recession". In: *International Tax and Public Finance* 25 (2018), pp. 1519–1558.

CBO. *Federal Grants to State and Local Governments*. Congrassional Budget Office, 2013.

Cengiz, Doruk et al. "The effect of minimum wages on low-wage jobs". In: *The Quarterly Journal of Economics* 134.3 (2019), pp. 1405–1454.

Cesa-Bianchi, Ambrogio, Gregory Thwaites, and Alejandro Vicondoa. "Monetary policy transmission in an open economy: new data and evidence from the United Kingdom". In: *Bank of England Working Paper* 615 (2016).

Chari, V.V. and Patrick J. Kehoe. "Optimal fiscal and monetary policy". In: *Handbook of Macroeconomics, Volume 1*. 1999.

Chodorow-Reich, Gabriel. "Geographic cross-sectional fiscal spending multipliers: what have we learned?" In: *NBER Working Paper Series* (2017).

— "Geographic Cross-Sectional Fiscal Spending Multipliers: What Have We Learned?" In: *American Economic Journal: Economic Policy* 11.2 (2019), pp. 1–34.

Christiano, Lawrence J., Martin Eichenbaum, and Charles Evans. "The Effects of Monetary Policy Shocks: Evidence from the Flow of Funds". In: *The Review of Economics and Statistics* 78.1 (1996), pp. 16–34.

Clemens, Jeffrey and Michael Wither. "The minimum wage and the Great Recession: evidence of effects on the employment and income trajectories of low-skilled workers". In: *Journal of Public Economics* 170 (2019), pp. 53–67.

Colas, Mark and Kevin Hutchinson. "Heterogeneous Workers and Federal Income Taxes in a Spatial Equilibrium". In: (2020).

Cremer, Helmuth and Pierre Pestieau. "Factor mobility and redistribution". In: *Handbook of Regional and Urban Economics* 4 (2004), pp. 2529–2560.

Danziger, Eliav and Leif Danziger. "The Optimal Graduated Minimum Wage and Social Welfare". In: *Transitions through the Labor Market: Work, Occupation, Earnings and Retirement*. Emerald Publishing Limited, 2018, pp. 55–72.

Debrun, Xavier. "Fiscal Rules in a Monetary Union: A Short-Run Analysis". In: *Open economies review* 11 (2000), pp. 323–358.

Dixit, Avinash and John Londregan. "Fiscal federalism and redistributive politics". In: *Journal of Public Economics* 68.2 (1998), pp. 153–180.

Dube, Arindrajit, T. William Lester, and Michael Reich. "Minimum wage shocks, employment flows, and labor market frictions". In: *Journal of Labor Economics* 34.3 (2016), pp. 663–704.

Duffie, Darrell, Icolae Garleanu, and Lasse Heje Pedersen. "Over-the-counter markets". In: *Econometrica* 73.6 (2005), pp. 1815–1847.

Durdu, Ceyhun Bora, Enrique G. Mendoza, and Marco E. Terrones. "Precautionary demand for foreign assets in Sudden Stop economies: An assessment of the New Mercantilism". In: *Journal of Development Economics* 89 (2009), pp. 194–209.

Dvorkin, Maximiliano A. "Fiscal Federalism and Optimal Income Taxes". In: *Federal Reserve Bank of St. Louis Review* (2017), pp. 319–336.

Economides, George and Thomas Moutos. "Minimum wages in the Presence of in-kind Redistribution". In: *Working paper* (2017).

Epple, Dennis and Thomas Nechyba. "Fiscal decentralization". In: *Handbook of Regional and Urban Economics* 4 (2004), pp. 2423–2480.

Epple, Dennis and Glenn J. Platt. "Equilibrium and local redistribution in an urban economy when households differ in both preferences and incomes". In: *Journal of Urban Economics* 43.1 (1998), pp. 23–51.

Epple, Dennis and Thomas Romer. "Mobility and redistribution". In: *Journal of Political Economy* 99.4 (1991).

Fama, Eugene F. and James D. MacBeth. "Risk, Return, and Equilibrium: Empirical Tests". In: *Journal of Political Economy* 81.3 (1973), pp. 607–636.

Farhi, Emmanuel and Iván Werning. "Fiscal Multipliers: Liquidity Traps and Currency Unions". In: *Handbook of Macroeconomics* 2 (2017), pp. 2417–2492.

Farhi, Emmanuel and Ivan Werning. "Fiscal Unions". In: *American Economic Review* 107.12 (2017).

Feldhutter, Peter. "The same bond at different prices: Identifying search frictions and selling pressures". In: *Review of Financial Studies* 25.1 (2012), pp. 1155–1206.

Ferrero, Andrea. "Fiscal and monetary rules for a currency union". In: *Journal of International Economics* 77.1 (2009), pp. 1–10.
Fisher, Ronald C. and Robert W. Wassmer. "The Issuance of State and Local Debt During the United States Great Recession". In: *National Tax Journal* 67.1 (2014), pp. 113–150.
Flinn, Christopher J. "Minimum wage effects on labor market outcomes under search, matching and endogeneous contact rates". In: *Econometrica* 74.4 (2006), pp. 1013–1062.
Francis, Neville, Michael T. Owyang, and Tatevik Sekhposyan. "The Local Effects of Monetary Policy". In: *The B.E. Journal of Macroecnomics* 12.2 (2012).
Fukumura, Koichi and Atsushi Yamagishi. "Minimum wage competition". In: *International Tax and Public Finance* 27.6 (2020), pp. 1557–1581.
Galí, Jordi and Tommaso Monacelli. "Optimal monetary and fiscal policy in a currency union". In: *Journal of International Economics* 76 (2008), pp. 116–132.
Gao, Pengjie, Chang Lee, and Dermot Murphy. "Municipal borrowing costs and state policies for distressed municipalities". In: *Journal of Financial Economics* 132.2 (2019), pp. 404–426.
Garrett, Daniel et al. "Tax Advantages and Imperfect Competition in Auctions for Municipal Bonds". In: *NBER Working Paper* 23473 (2018).
Gerritsen, Aart and Bas Jacobs. "Is a Minimum wage an Appropriate Instrument for Redistribution?" In: *Economica* 87 (2020), pp. 611–637.
Gilchrist, Simon, Vivian Yue, and Egon Zakrajšek. "U.S. Monetary Policy and International Bond Markets". In: *Journal of Money, Credit, and Banking* 51 (S1 2019), pp. 127–161.
Goodspeed, Timothy. "Absorbing Shocks: National Rainy-Day Funds and Cross-Country Transfers in a Fiscal Union". In: *FinanzArchiv: Public Finance Analysis* 72.4 (2016), pp. 407–420.
Gordon, Roger H. "An optimal taxation approach to fiscal federalism". In: *The Quarterly Journal of Economics* 98.4 (1983), pp. 567–586.
Gordon, Roger H. and Julie Berry Cullen. "Income redistribution in a federal system of governments". In: *Journal of Public Economics* 96.11-12 (2012), pp. 1100–1109.
Gramlich, Edward M. "Subnational fiscal policy". In: *Perspectives on local public finance and public policy* 3.3-27 (1997).
Green, Richard C., Burton Hollifield, and Norman Schürhoff. "Dealer intermediation and price behavior in the aftermarket for new bond issues". In: *Journal of Financial Economics* 86.3 (2007), pp. 643–682.
— "Financial Intermediation and the Costs of Trading in an Opaque Market". In: *The Review of Financial Studies* 20.2 (2007), pp. 275–314.

Grigoris, Fotis. "The term structure of municipal bond yields, local economic conditions, and local stock returns". In: *Working paper* (2019).

Guesnerie, Roger and Kevin Roberts. "Minimum wage legislation as a second best policy". In: *European Economic Review* 31.1-2 (1987), pp. 490–498.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright. "The U.S. Treasury Yield Curve: 1961 to the Present". In: *FEDS Working Paper* 2006-28 (2 2006).

Hamilton, Jonathan and Pierre Pestieau. "Optimal income taxation and the ability distribution: Implications for migration equilibria". In: *International Tax and Public Finance* 12.1 (2005), pp. 29–45.

Harding, Don and Adrian Pagan. "Dissecting the cycle: a methodological investigation". In: *Journal of Monetary Economics* 49 (2002), pp. 365–381.

Harris, Lawrence E. and Michael S. Piwowar. "Secondary Trading Costs in the Municipal Bond Market". In: *The Journal of Finance* 61.3 (2006), pp. 1361–1397.

Hines Jr, James R. "State fiscal policies and transitory income fluctuations". In: *Brookings Papers on Economic Activity* 2010.2 (2010), pp. 313–337.

Hou, Yilin and Daniel L. Smith. "A Framework for Understanding State Balanced Budget Requirement Systems: Reexamining Distinctive Features and an Operational Definition". In: *Public Budgeting and Finance* 26.3 (2006), pp. 22–45.

Hoynes, Hilary W. and Erzo FP Luttmer. "The insurance value of state tax-and-transfer programs". In: *Journal of Public Economics* 96.11-12 (2012), pp. 1110–1128.

Hungerb uhler, Mathias and Etienne Lehmann. "On the optimality of a minimum wage: New insights from optimal tax theory". In: *Journal of Public Economics* 93.3-4 (2009), pp. 464–481.

Janeba, Eckhard and John Douglas Wilson. "Optimal fiscal federalism in the presence of tax competition". In: *Journal of Public Economics* 95.11-12 (2011), pp. 1302–1311.

Jankowitsch, Rainer, Amrut Nashikkar, and Marti G. Subrahmanyam. "Price dispersion in OTC markets: A new measure of liquidity". In: *Journal of Banking and Finance* 35 (2011), pp. 343–357.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. "Monetary Policy According to HANK". In: *American Economic Review* 108.3 (2018), pp. 697–743.

Katz, Lawrence F. and Kevin M. Murphy. "Changes in relative wages, 1963-1987: supply and demand factors". In: *The Quarterly Journal of Economics* 107.1 (1992), pp. 35–78.

Keen, Michael J. and Christos Kotsogiannis. "Does federalism lead to excessively high taxes?" In: *American Economic Review* 92.1 (2002), pp. 363–370.

Keen, Michael J. and Christos Kotsogiannis. "Tax competition in federations and the welfare consequences of decentralization". In: *Journal of Urban Economics* 56 (Mar. 2004), pp. 397–407.

Kennan, John and James R. Walker. "The effect of expected income on individual migration decisions". In: *Econometrica* 79.1 (2011), pp. 211–251.

Kleven, Henrik et al. "Taxation and Migration: Evidence and Policy Implications". In: *NBER working paper* w25740 (2019).

Kline, Patrick and Enrico Moretti. "People, Places, and Public Policy: some Simple Welfare Economics of Local Economic Development Programs". In: *Annual Review of Economics* 6.1 (2014), pp. 629–662.

Ladd, Helen F. and Fred C. Doolittle. "Which Level of Government Should Assist the Poor?" In: *National Tax Journal* 35 (1982), pp. 323–336.

Lavecchia, Adam. "Minimum Wage Policy with Optimal Taxes and Unemployment". In: *McMaster University Department of Economics Working Paper* (2019).

Lee, David and Emmanuel Saez. "Optimal minimum wage policy in competitive labor markets". In: *Journal of Public Economics* 96.9-10 (2012), pp. 739–749.

Lockwood, Ben. "Distributive politics and the costs of centralisation." In: *Review of Economic Studies* 69 (2002), pp. 313–338.

Lucas, Robert E. "Some International Evidence on Output-Inflation Tradeoffs". In: *The American Economic Review* 63.3 (1973), pp. 326–334.

Luque, Jaime, Massimo Morelli, and Jose Tavares. "A volatility-based theory of fiscal union desirability". In: *Journal of Public Economics* 112 (2014), pp. 1–11.

Luttmer, Erzo FP. "Does the minimum wage cause inefficient rationing?" In: *The BE Journal of Economic Analysis & Policy* 7.1 (2007).

McKay, Alisdair, Emi Nakamura, and Jón Steinsson. "The Power of Forward Guidance Revisited". In: *American Economic Review* 106.10 (2016), pp. 3133–58.

Moldogaziev, Tima T. "The Market(s) Left Behind: Price Formation in Municipal secondary Trades". In: *Working paper* (2018).

Monras, Joan. "Minimum Wages and Spatial Equilibrium: Theory and Evidence". In: *Journal of Labor Economics* 37.3 (2019).

Moretti, Enrico. "Real wage inequality". In: *American Economic Journal: Applied Economics* 5.1 (2013), pp. 65–103.

Nakamura, Emi and Jón Steinsson. "Fiscal Stimulus in a Monetary Union: Evidence from US Regions". In: *American Economic Review* 104.3 (2014), pp. 753–92.

— "Fiscal stimulus in a monetary union: Evidence from US regions". In: *The American Economic Review* 104.3 (2014), pp. 753–792.

— "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect". In: *The Quarterly Journal of Economics* 133.3 (2018), pp. 1283–1330.

NASBO. *Fiscal Survey of the States, 1979-2017*. 2017.

NCSL. "NCSL Fiscal Brief: State Balanced Budget Provisions". In: *National Conference of State Legislatures* (2010).

Neumark, David. "The Econometrics and Economics of the Employment Effects of Minimum Wages: Getting from Known Unknowns to Known Knowns". In: *NBER Working Paper* (2018).

Neumark, David and William L. Wascher. "Minimum wages and employment". In: *Foundations and Trends in Microeconomics* 3.12 (2007), pp. 1–182.

Oates, Wallace E. "An economist perspective on Fiscal federalism". In: *The Political Economy of Fiscal Federalism*. Ed. by WE Oates. 1977, pp. 3–19.

— *Fiscal Federalism*. 1972.

Oates, Wallace E. "On the evolution of fiscal federalism: theory and institutions". In: *National Tax Journal* 61.2 (2008), pp. 313–334.

Oates, Wallace E. "Toward A Second-Generation Theory of Fiscal Federalism". In: *International Tax and Public Finance* 12 (2005), pp. 349–373.

OECD. *Main Economic Indicators - complete database*. 2017.

Owyang, Michael T. and Sarah Zubairy. "The Regional Variation in the Response to Government Spending Shocks". In: *Working paper* (2008).

Painter, Marcus. "An inconvenient cost: The effects of climate change on municipal bonds". In: *Journal of Financial Economics* 135.2 (2020), pp. 468–482.

Philadelphia, Federal Reserve Bank of. *State Coincident Indices*. 2017.

Phillips, David C. "Do low-wage employers discriminate against applicants with long commutes? Evidence from a correspondence experiment". In: *Journal of Human Resources* (2018), 1016–8327R.

Pirinsky, Christo A. and Qinghai Wang. "Market Segmentation and the Cost of Capital in a Domestic Market: Evidence from Municipal Bonds". In: *Financial Management* 40.2 (2011), pp. 455–481.

Poverty Research, University of Kentucky Center for. *UKCPR National Welfare Data, 1980-2016*. 2017.

Rigobon, Roberto and Brian Sack. "The impact o fmonetary policy on asset prices". In: *Journal of Monetary Economics* 51.8 (2004), pp. 1553–1575.

Roback, Jennifer. "Wages, rents, and the quality of life". In: *Journal of Political Economy* 90.6 (1982), pp. 1257–1278.

Rohlin, Shawn M. "State minimum wages and business location: Evidence from a refined border approach". In: *Journal of Urban Economics* 69.1 (2011), pp. 103–117.

Romer, Christina D and David H Romer. “A New Measure of Monetary Shocks: Derivation and Implications”. In: *American Economic Review* 94.4 (2004), pp. 1055–1084.

Rosa, Carlo. “Municipal Bonds and Monetary Policy: Evidence from the Fed Funds Futures Market”. In: *The Journal of Futures Markets* 34.5 (2014), pp. 434–450.

Rosen, Sherwin. *Wage-based indexes of urban quality of life*. Baltimore: Johns Hopkins University, 1979.

Schmitt-Grohe, Stephanie and Martin Uribe. “Closing small open economy models”. In: *Journal of International Economics* 61.2 (2003), pp. 163–185.

— “Nominal Rigidity, Exchange Rates, and Unemployment”. In: *Open Economy Macroeconomics* 2 (2016).

— “Optimal fiscal and monetary policy under sticky prices”. In: *Journal of Economic Theory* 114 (2004), pp. 198–230.

Schwert, Michael. “Municipal Bond Liquidity and Default Risk”. In: *The Journal of Finance* 72.4 (2017), pp. 1683–1722.

Seegert, Nathan. “Optimal Taxation with Volatility: A Theoretical and Empirical Decomposition”. In: *SSRN working paper* (2017).

— “The Performance of State Tax Portfolios During and After the Great Recession”. In: *National Tax Journal* 68.4 (2015), pp. 901–918.

Silva, Emilson C.D. and Richard C. Cornes. “Local Public Goods, Risk Sharing, and Private Information in Federal Systems”. In: *Journal of Urban Economics* 47 (2000), pp. 39–60.

Simon, Andrew. “Public Good Spillovers and Fiscal Centralization: Evidence from Community College Expansions”. In: (2020).

Stigler, George J. “The Tenable Range of Functions of Local Government”. In: (1957).

Stiglitz, Joseph E. “Self-selection and Pareto efficient taxation”. In: *Journal of Public Economics* 17.2 (1982), pp. 213–240.

Suárez Serrato, Juan Carlos and Owen Zidar. “Who benefits from state corporate tax cuts? A local labor markets approach with heterogeneous firms”. In: *American Economic Review* 106.9 (2016), pp. 2582–2624.

Tiebout, Charles M. “A Pure Theory of Local Expenditures”. In: *The Journal of Policitcal Economy* 64.5 (1956), pp. 416–424.

Weiland, Johannes and Mu-Jeung Yang. “Financial Dampening”. In: *Working paper* (2019).

Wildasin, David E. “Income distribution and redistribution within federations”. In: *Annales d’Economie et de Statistique* (1997), pp. 291–313.

— “Labor-market integration, investment in risky human capital, and fiscal competition”. In: *American Economic Review* 90.1 (2000), pp. 73–95.

Wilson, John D. “A theory of interregional tax competition”. In: *Journal of Urban Economics* 19.3 (1986), pp. 296–315.

— “Optimal linear income taxation in the presence of emigration”. In: *Journal of Public Economics* 18.3 (1982), pp. 363–379.

— “Tax competition in a federal setting”. In: *Handbook of Multilevel Finance* (2015), p. 264.

Wooldridge, Jeffrey M. *Econometric Analysis of Cross Section and Panel Data*. 2010.

Yang, Lang. “Not All State Authorizations for Municipal Bankruptcy Are Equal: Impact on State Borrowing Costs”. In: *National Tax Journal* 72.2 (2019), pp. 435–464.

Zodrow, George R. and Peter Mieszkowski. “Pigou, Tiebout, property taxation, and the underprovision of local public goods”. In: *Journal of Urban Economics* 19.3 (1986), pp. 356–370.