NBER WORKING PAPER SERIES

THE AGGREGATE IMPLICATIONS OF REGIONAL BUSINESS CYCLES

Martin Beraja
Erik Hurst
Juan Ospina

Working Paper 21956
http://www.nber.org/papers/w21956

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2016

First draft: May 2014. A previous version of this paper circulated as "The Regional Evolution of Prices and Wages During the Great Recession". We thank Mark Aguiar, Manuel Amador, David Argente, Mark Bils, Elisa Giannone, Adam Guren, Simon Gilchrist, Paul Gomme, Bob Hall, Loukas Karabarbounis, Pat Kehoe, Virgiliu Midrigan, Elena Pastorino, Harald Uhlig and Joe Vavra for their very helpful comments and suggestions. Finally, we thank seminar participants at Berkeley, Boston University, Brown, Chicago, Chicago Federal Reserve, Columbia, Duke, Harvard, Michigan, Minneapolis Federal Reserve, Minnesota Workshop in Macroeconomic Theory, MIT, NBER's Summer Institute EF&G, NBER's Summer Institute Prices Program, Northwestern, Princeton, Rochester, St. Louis Federal Reserve, UCLA, Yale's Cowles Conference on Macroeconomics, the Board of Governors of the Federal Reserve, and the Bank of England. Any remaining errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Aggregate Implications of Regional Business Cycles
Martin Beraja, Erik Hurst, and Juan Ospina
NBER Working Paper No. 21956
February 2016
JEL No. E24,E31,E32,R12,R23

ABSTRACT

We argue that it is difficult to make inferences about the drivers of aggregate business cycles using regional variation alone because (i) the local and aggregate elasticities to the same type of shock are quantitatively different and (ii) purely aggregate shocks are differenced out when using cross-region variation. We highlight the importance of these confounding factors by contrasting the behavior of U.S. aggregate time-series and cross-state patterns during the Great Recession. In particular, using household and scanner data for the US, we document a strong relationship across states between local employment growth and local nominal and real wage growth. These relationships are much weaker in US aggregates. In order to identify the shocks driving aggregate (and regional) business cycles we develop a semi-structural methodology that combines regional and aggregate data within a model of a monetary union. The methodology uses theoretical restrictions implied by a wage setting equation with nominal wage rigidities. Taking this methodology to the data, we find that a combination of both "demand" and "supply" shocks are necessary to account for the joint dynamics of aggregate prices, wages and employment during the 2007-2012 period in the US while only "demand" shocks are necessary to explain most of the observed cross-state variation. We conclude that the wage stickiness necessary to get demand shocks to be the primary cause of aggregate employment declines during the Great Recession is inconsistent with the flexibility of wages estimated from cross-state variation.

Martin Beraja
University of Chicago
martinberaja@gmail.com

Juan Ospina
University of Chicago
juan.jose.ospina@gmail.com

Erik Hurst
Booth School of Business
University of Chicago
Harper Center
Chicago, IL 60637
and NBER
erik.hurst@chicagobooth.edu

A data appendix is available at http://www.nber.org/data-appendix/w21956

1 Introduction

There is a large and growing literature using regional variation to learn about the determinants of aggregate economic variables.[1] However, making inferences about the aggregate economy using only regional variation is complicated by two issues. First, we show that in a model of a monetary union, local and aggregate elasticities to the same type of shock are quantitatively different because of both factor mobility and general equilibrium forces. This implies that it is problematic to use local shock elasticities estimated from regional data to ascertain the importance of a given aggregate shock. Second, certain types of shocks get differenced out when using cross-region variation. It is then impossible to learn anything about these aggregate shocks by exploiting variation across locations. Thus, in general we cannot expect to fully understand the joint evolution of aggregate variables by using only cross-sectional variation.

In this paper, we present a methodology that uses regional data along with aggregate data to identify aggregate (and regional) shocks driving business cycles. The methodology uses theoretical restrictions implied by a wage setting equation that holds in models of a monetary union with wage rigidities. Under certain conditions, we are able to use cross-state variation in wages, prices, and employment to estimate the aggregate degree of wage rigidity. The extent to which aggregate wages are rigid is a key restriction that places structure on the type of shocks driving aggregate fluctuations.

Using household and scanner data for the US, we document a strong relationship across states between local employment growth and local nominal and real wage growth. This underlying variation drives our estimates of the parameters of the wage setting equation. Our estimates suggest that wages are relatively flexible, limiting the contribution of “demand" shocks to the aggregate employment decline during the Great Recession.[2] We find that a combination of “demand" and other shocks are necessary to account for the joint dynamics of aggregate prices, wages and employment during the 2007- 2012 period within the U.S.. We argue that the reason why aggregate wages in the time series appear to be relatively stable when compared to regional wages is not because wages are sticky in the aggregate, but because different aggregate shocks have relatively offsetting effects on aggregate wages. We conclude that the wage stickiness necessary to get demand shocks to be the primary cause of aggregate employment declines during the Great Recession is inconsistent with the flexibility of wages estimated from cross-state variation.

The paper is organized as follows. We begin the paper by documenting a series of new facts about the variation in prices and wages across U.S. states during the Great Recession. To do this, we use data from Nielsen’s *Retail Scanner Database* to create price indices for each U.S. state. As we discuss in detail below, the *Retail Scanner Database* includes prices and quantities for given UPC

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

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

codes at over 40,000 stores at a weekly frequency from 2006 through 2012. We show that an aggregate price index created with this data matches the BLS's Food CPI nearly identically. While the price indices we create from this data are based mostly on consumer packaged goods, we show how under certain assumptions the indices can be scaled to be representative of a composite local consumption good. Using this data, we show that there is a strong correlation between local prices and local economic activity; locations with the largest declines in real economic activity between 2007 and 2010 had the smallest three year inflation rate.

Using data from the *2000 U.S. Census* and the 2000 - 2012 *American Community Surveys* (ACS), we then make composition adjusted nominal wage indices for each U.S. state during the 2000 to 2012 period. We focus on a sample of wage measures for full time workers with a strong attachment to the labor force. We further adjust our measures for the fact that there are observable changes in the composition of the labor force over the business cycle. Using these indices, we show that states that experienced larger employment declines between 2007 and 2010 had significantly lower nominal wage growth during the same time period. Using the local prices variation that we estimate, we can further make measures of real wage growth at the state level. Our estimates suggest that real wages also vary significantly with local measures of employment at the state level.

These cross region patterns stand in sharp contrast with the well documented aggregate time series patterns for prices and wages during the same time period. As both aggregate output and employment contracted sharply within the U.S. during the 2007-2012 period, aggregate nominal wage growth remained robust and real wage growth did not break trend. The robust growth in nominal wages during the recession is viewed as a puzzle for those that believe that the lack of aggregate demand was the primary cause of the Great Recession.[3] The key point we wish to make with these new facts is that while aggregate wages appear sticky during the Great Recession using time series variation, local wages were strongly correlated with measures of local employment growth using cross-state variation. Any model that is calibrated to match the relationship between wage growth and employment growth at the aggregate level needs to confront the fact that there is a strong relationship between local employment growth and local wage growth at business cycle frequencies.

In Section 4, we present a simple model of a monetary union that we use for two purposes. First, the model specifies a wage setting equation for local and the aggregate economies. For a given set of preferences, there is a direct mapping between the parameters of the local and aggregate wage setting equations. These equations will be important to our procedure that identifies the underlying aggregate and local shocks. Second, the calibrated model allows us to quantify the differences between aggregate and local elasticities to a given shock.

The model has many islands linked by trade in intermediate goods which are used in the production of a non-tradable final consumption good and by trade in a risk-free asset. The nominal interest rate on this asset follows a rule that endogenously responds to aggregate variables and is

[3]For example, this point was made by Krugman in a recent New York Times article ("Wages, Yellen and Intellectual Honesty", NYTimes 8/25/14).

set at the union level. Labor is the only other input in production, which is not mobile across islands. We assume that nominal wages are only partially flexible. This is the only nominal rigidity in the model. Finally, the model includes a series of shocks: a shock to the household's discount rate, shocks to non-tradable and tradable productivity/mark-up, a shock to the household's labor supply, and a monetary policy shock. All shocks, aside from the monetary policy shock, have both local and aggregate components where by definition the weighted average of the local shocks sum to zero. We show that, under relatively few assumptions, the log-linearized economy aggregates. This allows us to study the aggregate and local behavior separately, a property that we will exploit when estimating the aggregate and regional shocks through our methodology.

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

In Section 5, we turn to the estimation of the local and aggregate shocks. We consider a broader class of models than the simple model outlined above. The broader class of models, however, nests our simple model. In particular, we show that the aggregate and local equilibria can be represented as vector autoregression (VAR) in prices, nominal wages, and employment with three shocks. We refer to the three shocks as a "demand shock" (which is a combination of the discount rate and monetary policy shock), a productivity/mark-up shock (which is a combination of the productivity/mark-up shocks in the tradable and non tradable sectors) and the labor supply shock. To back out the aggregate (local) shocks, we estimate the aggregate (local) VAR. Our estimation of the VAR is semi-structural in that we impose the aggregate (local) wage setting equation as an additional restriction to help identify the VAR.[5] The aggregate (local) wage setting equation implies a series of particular linear restrictions linking the reduced form errors to the underlying structural shocks. When we impose these linear restrictions, the VAR becomes identified. In essence, our procedure uses some elements of theory to help identify the underlying economic shocks.[6]

The shock identification procedure requires parameterizing the structural wage setting equation. We argue that the regional data on prices, wages and employment during the 2006-2012 period can

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

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

[6] Our estimating procedure is similar in spirit to the procedure developed in Hamilton (2015).

be used to estimate the Frisch elasticity of labor supply and the amount of wage stickiness which are the only parameters of the wage setting equation in our base specification. Using a variety of econometric techniques, including instrumenting for local labor demand shocks, we use regional data to estimate parameters of the wage setting equations. Across a variety of our specifications and identification procedures, we estimate only a modest amount of wage stickiness. The amount of wage flexibility we estimate using the local variation is much greater than estimates of wage flexibility obtained using only aggregate time series data.

With the parameterized aggregate wage setting equation, we use the VAR methodology described above to estimate the shocks driving aggregate employment, prices, and wages during the Great Recession. The results suggest that during the early part of the recession (2008-2009) roughly 30 percent of the aggregate employment decline can be traced to the "demand shock" (the discount rate plus the monetary policy shocks). The leisure shock explains roughly 30 percent of the decline in aggregate employment while the productivity/mark-up shock explaining the remaining 40 percent. Over a longer period (2008-2012), the demand shock cannot explain any of the persistence in the employment decline. Instead, it is the productivity/mark-up and labor supply shock explaining why employment remained low from 2010-2012. While the demand shock may have been important in the early part of the recession, it had little effect on explaining the low levels of employment in the U.S. after 2009.[7] Furthermore, we show that an aggregate labor supply shock is needed to explain why aggregate wages did not fall.

The regional data in our paper serve two purposes in our estimation. First, the regional data is needed to estimate the amount of wage stickiness which is a parameter of the aggregate wage setting equation. Second, the regional data is needed to estimate the local VAR. We use a similar procedure to estimate the shocks driving the local economies. Our results suggest that the "demand" shock is driving most of the cross-region variation in employment during the Great Recession. This is why price and wage growth are very positively correlated with employment growth at the local level.

Our paper contributes to many literatures. First, our work contributes to the recent surge in papers that have exploited regional variation to highlight mechanisms of importance to aggregate fluctuations. For example, Mian and Sufi (2011 and 2014), Mian, Rao, and Sufi (2013) and Midrigan and Philippon (2011) have exploited regional variation within the U.S. to explore the extent to which household leverage has contributed to the Great Recession.[8] Nakamura and Steinsson (2014) use sub-national U.S. variation to inform the size of local government spending multipliers. Blanchard and Katz (1991), Autor et al. (2013), and Charles et al. (2015) use regional variation to measure the responsiveness of labor markets to labor demand shocks. Our work contributes to this literature on

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

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

two fronts. First, we show that local prices and wages also respond to local changes in economic conditions at business cycle frequencies. Second, we provide a procedure where local variation can be combined with aggregate data to infer something about the nature and importance of certain mechanisms for aggregate fluctuations. With respect to the latter innovation, our paper is similar in spirit to Nakamura and Steinsson (2014).

Second, our paper contributes to the recent literature trying to determine the causes of the Great Recession. In many respects, our model is more stylized than others in this literature in that we include a broad set of shocks without trying to uncover the underlying micro foundation for these shocks. However, the shocks we chose to focus on were designed to proxy for many of the popular theories about the drivers of the Great Recession. For example, our discount rate shock can be thought of as reduced form representation of tightening of household borrowing limits. For example, such shocks have been proposed by Eggertsson and Krugman (2012), Guerrieri and Lorenzoni (2011) and Mian and Sufi (2014) as an explanation of the 2008 recession. Likewise, our productivity/mark-up shock can be interpreted as anything that changes firms' demand for labor. In a reduced form sense, credit supply shocks to firms, such as those proposed by Gilchrist et al (2014), would be similar to our productivity/mark-up shock. Finally, our labor supply shock can be seen as a proxy for increasing distortions within the labor market due to changes in government policy (e.g., Mulligan (2012) or as a reduced form representation of a skill mismatch story within the labor market (e.g., Charles et al. (2013, 2015)). It is these shocks which have been proposed to explain the well documented decline in labor force participation during the last decade. So while our specification only allows for broad reduced form shocks, we think these shocks nest many of the prominent stories about the underlying causes of employment declines during the Great Recession. As we show, the fact that prices and wages move with economic conditions at the local level help discipline how aggregate prices and wages should have moved in response to different types of shocks.

2 Creating State Level Price And Wage Indices

2.1 Local Price Indices

2.1.1 Price Data

Local price indices are necessary to make measures of local real wages. To construct state level price indices we use the *Retail Scanner Database* collected by AC Nielsen and made available at The University of Chicago Booth School of Business.[9] The Retail Scanner data consists of weekly pricing, volume, and store environment information generated by point-of-sale systems for about 90

[9] The data is made available through the Marketing Data Center at the University of Chicago Booth School of Business. Information on availability and access to the data can be found at http://research.chicagobooth.edu/nielsen/. Contemporaneously, Coibion et al. (2015), Kaplan and Menzio (2015) and Stroebel and Vavra (2014) also use local scanner data/household price data to estimate that local prices vary with local economic conditions at business cycle frequencies. Our paper complements this literature by actually making price indices using the Nielsen scanner data for each state at the monthly frequency and using those price indices to estimate structural parameters of the local wage setting equation.

participating retail chains across all US markets between January 2006 and December 2012. When a retail chain agrees to share their data, all of their stores enter the database. As a result, the database includes roughly 40,000 individual stores. Each entry includes a store identifier and a store-chain identifier so a given store can be tracked over time and can be linked to a specific chain. While each chain has a unique identifier, no information is provided that directly links the chain identifier to the name of the chain. Most sales in the data, about 97 percent, come from food, drug and mass merchandising stores.

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

Finally, the geographic coverage of the database is outstanding and is one of its most attractive features. It includes stores from all states except for Alaska and Hawaii (but including the District of Columbia). Likewise, it covers stores from 361 Metropolitan Statistical Areas (MSA) and 2,500 counties. The data comes with both zip code and FIPS codes for the store's county, MSA, and state. Over the seven year period, the data set includes total sales across all retail establishments worth over $1.5 trillion. In this paper, we aggregate data to the level of U.S. states and compute state level scanner data price indices. Online Appendix Table O1 shows summary statistics for the scanner data for each year between 2006 and 2012 and for the sample as a whole.[10]

2.1.2 A Scanner Data Price Index

Our goal is to construct regional price indices from the scanner data that are similar in spirit to how the BLS constructs the CPI.[11] While we briefly outline the price index construction in this subsection, the full details of the procedure are discussed in the Online Appendix that accompanies our paper. Our scanner price indices are built in two stages. In the first stage, we aggregate the prices of goods within the roughly 1,000 categories described above. For our base indices, a good is a given store-UPC pair such that a UPC in store A is treated as a different good than the same UPC sold in store B. We do this to allow for the possibility that prices may change as households substitute

[10]The online appendix is at http://faculty.chicagobooth.edu/erik.hurst/research/regional_online_appendix.pdf

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

from a high cost store (that provides a different shopping experience) to a low cost store when local economic conditions deteriorate. We then compute for each good the average price and total quantity sold for the month within a given level of geography. Next, we find the quantity weighted average price for all goods within each detailed category (sugar granulated, sugar powdered, etc.) for a given month within a given location. We aggregate our index to the monthly level to reduce the number of missing values.

Specifically, for each category, we compute:

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

where $P_{j,t,y,k}$ is category level price index for category j, in period t, with a base year y, in geography k. For our analysis, geographies will either be U.S. states or the country as a whole. $p_{i,t,k}$ is the price at time t of the specific good i (from category j) in geography k and $\bar{q}_{i,t-1,k}$ is the average monthly quantity sold of good i in the prior year in location k. By fixing quantities at their prior year's level, we are holding fixed household's consumption patterns as prices change. We update the basket of goods each year, and chain the resulting indices to produce one chained index for each category in each geography. Fixing quantities at a lagged level implies that the price changes we document below with changing local economic conditions is not the result of changing household consumption patterns.

The second stage of our price indices also roughly follows the BLS procedure in that we aggregate the category-level price indices into an aggregate index for each location k. The inputs are the category-level prices and the total expenditures of each category. Specifically, for each state we compute:

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

where $\bar{S}^{t}_{j,k}$ is the share of expenditure of category j in month t in location k averaged over the year. For the purposes of this paper, we make our baseline specification one that allows the weights of each category to be updated monthly.[12]

To benchmark our scanner price index, we compare our scanner price index for the aggregate U.S. to the BLS's CPI for food. We chose the BLS Food CPI as a benchmark given that most of the goods in our database can be classified as food.[13] Figure 1 shows that our scanner price index

[12]One issue discussed in greater depth within the Online Appendix is how we deal with missing data when computing the price indices. Seasonal goods, the introduction of new goods, and the phasing out of existing goods means that missing data on month to month price changes occur. When computing our price indices, we restrict our sample to only include (1) goods that had positive sales in the prior year and (2) goods that had positive sales in every month of the current year. Online Appendix Table R1 shows the percent of sales included within the price index for each sample year.

[13]Not all of our goods are food products. About 13 percent of our goods (expenditure weighted) are health and beauty products (including drugs). About 6 percent of our goods (expenditure weighted) are alcoholic beverages. About 13 percent are non-food grocery items (e.g., paper products, disposable diapers, laundry detergents, and household cleaning supplies). Finally, about 7 percent of our goods (expenditure weighted) are non-food, non-health and beauty,

matches nearly identically the BLS's Food CPI.

2.1.3 Computing Regional Inflation Rates Using Retail Data

One natural question is how to extend the spatial variation in inflation rates based on the goods in our sample to spatial variation in inflation rates for a composite basket of consumer goods. Most of the goods in our sample are produced outside the local market and are simultaneously sold to many local markets. These production costs represent the traded portion of local retail prices. If there were no additional local distribution costs, one would expect little variation in retail prices across regions if retail goods were purely tradable; the law of one price would hold. However, there are local costs associated with retail distribution. These costs include the wages of workers in the retail establishments, the rent of the retail facility, and expenses associated with local warehousing and transportation.

Assuming that these non-tradable shares are constant across regions and identical for all firms in the retail industry within our sample, we can express local retail grocery prices (P^r) in region k during period t as:

$$P^r_{t,k} = (P^T_t)^{1-\alpha_r}(P^{NT}_{t,k})^{\alpha_r}$$

where P^T_t is the tradable component of local retail grocery prices in period t and does not vary across regions and $P^{NT}_{t,k}$ is the non-tradable component of local retail grocery prices in period t and potentially does vary across regions. α_r represents the share of non-tradable prices in the total price for the retail goods in our sample.

What we are interested in is the traded and non-traded component of the typical good in the household's consumption basket as opposed to just the regional variation in grocery prices. Suppose that the composite good in a region can be expressed such that:

$$P_{t,k} = (P^T_t)^{1-\bar{\alpha}}(P^{NT}_{t,k})^{\bar{\alpha}}$$

The retail grocery sector is only one sector within a household's local consumption bundle. For example, one could imagine sectors where the non-tradable share is much larger than in the grocery sector. Many local services primarily use local labor and local land in the production of their retail activities (e.g., dry-cleaners, haircuts, education services, and restaurants). Conversely, for other sectors, the traded component of costs could be large relative to the local factors used to sell the good (e.g., auto dealerships). $\bar{\alpha}$ is the non-tradable share for the composite consumption good in the local economy. We also assume that $\bar{\alpha}$ is constant across all regions.

Given these assumptions, we can transform the variation in the grocery sector prices that we identify into variation in the broader consumption basket across regions. Taking logs and differencing across regions we get that the variation in log-prices of the composite good between two regions

and non alcohol and tobacco products. This latter group includes goods such as batteries, cutlery, pots and pans, candles, cameras, small consumer electronics, office supplies, and small household appliances. The remaining items are food.

k, k' $(\Delta \ln P_{t,k,k'})$ is proportional to the variation in log-grocery retail prices across those same regions $\left(\Delta \ln P^r_{t,k,k'}\right)$. Formally,

$$\Delta \ln P_{t,k,k'} = \left(\frac{\bar{\alpha}}{\alpha_r}\right) \Delta \ln P^r_{t,k,k'}$$

With knowledge of α_r and $\bar{\alpha}$ we can make such an adjustment. Below, we discuss empirical methodologies to discipline α_r and $\bar{\alpha}$ which allow us to scale variation in retail grocery prices across regions into variation in a broader composite consumption good. Throughout, we will discuss the robustness of our results to different measures of $\bar{\alpha}/\alpha_r$.

2.2 Local Wage Indices

To make nominal wages at the state level, we use data from the 2000 *Census* and the 2001-2012 *American Community Surveys* (ACS). The 2000 *Census* includes 5 percent of the U.S. population while the 2001-2012 ACS's includes around 600,000 respondents per year between 2001 and 2004 and around 2 million respondents per year between 2005-2012. The large sample sizes allows us to compute detailed labor market statistics at the state level. Within each year of the *Census/ACS* data, we make hourly nominal wages for each individual in the sample. To do so, we restrict our sample to only males between the ages of 21 and 55, who are currently employed, who report usually working at least 30 hours per week, and who worked at least 48 weeks during the prior 12 months. These restrictions result in our sample being comprised of males with a strong attachment to the labor force. For each individual in the resulting sample, we divide total labor income earned during the prior 12 months by a measure of annual hours worked during prior 12 months.[14]

Despite our restriction to prime age males with a strong attachment to the labor force, the composition of workers on other dimensions may still differ across states and within a state over time. The changing composition of workers could still explain some of the variation in nominal wages across states over time. To account for this, we run the following regression on the ACS data to create a composition adjusted wage measure (at least based on observables):

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

where $\ln(w_{ikt})$ is log nominal wages for household i in period t residing in state k and X_{it} is a vector of household specific controls. The vector of controls include a series of dummy variables for usual hours worked (with 40-49 hours per week being the omitted group), a series of five year age dummies (with 40-44 being the omitted group), 4 educational attainment dummies (with some college being the omitted group), three citizenship dummies (with native born being the omitted group), and a series of race dummies (with white being the omitted group). We run these regressions separately for each year such that both the constant, γ_t, and the vector of coefficients on the controls,

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

Γ_t, can differ for each year. We then take the residuals from these regressions, η_{itk}, and add back the constant, γ_t. Adding back the constant from the regression preserves differences over time in average log wages. To compute average wages within a state holding composition fixed, we average $e^{\eta_{itk}+\gamma_t}$ across all individuals in state k. We refer to this measure as the demographic adjusted nominal wage in time t in state k. This is the series we use to exploit cross region variation in wages during the Great Recession.

The benefit of the *Census/ACS* data is that it is large enough to compute detailed labor market statistics at state levels. However, one drawback of the *Census/ACS* data is that it not available at an annual frequency prior to 2000. To complement our analysis, we use data from the March Supplement of the *Current Population Survey* (CPS) to examine longer run trends in both aggregate nominal and real wages between 1978 and 2012. These long run trends are an input into our aggregate shock decomposition discussed below. We compute wages within the CPS analogously to the way we computed wages within the ACS data. In particular, we compute hourly wages for men 21-55 with a strong attachment to the labor force (those currently working at least 30 hours a week and those who worked at least 48 weeks during the prior year). Again, like for the ACS data, we adjust the wages to account for a changing vector of observables over time.[15] A full discussion of the way we calculate composition adjusted wages in the CPS can be found in the Online Appendix that accompanies the paper. For the 2000-2012 period, we can compare the time series variation in aggregate nominal wages using the ACS data and the CPS data. The two series during this time period have a correlation of 0.99.

Figure 2 shows aggregate nominal and real composition adjusted log wages during the 2000-2012 period by taking the weighted average of individual composition adjusted wages from the CPS. To get aggregate composition adjusted real wages, we deflate the aggregate nominal adjusted wages from the CPS by the aggregate June CPI-U with 2000 as the base year. Between 2007 and 2010, average composition adjusted nominal wages within the U.S. increased by roughly 4 percent. The patterns in our data replicate the aggregate nominal wage growth patterns documented by many others in the literature.[16] Given that consumer prices increased by 5 percent during the same period, aggregate real wages in the U.S. fell by roughly 1 percent between 2007 and 2010. This was similar to the trend in real wages prior to the start of the recent recession. As seen from Figure 2, nominal wages increased slightly and real wage growth did not seem to break trend during the Great Recession. The aggregate puzzle has been why aggregate wages did not decline relative to trend despite the very weak aggregate labor market.

[15] The X_{it} of controls to adjust the CPS wage data is identical to the controls used to adjust the ACS wage data except for the inclusion of citizenship dummies. The CPS does not measure citizenship status consistently over time. Additionally, we excluded anyone with business earnings from the sample. Finally, we estimate the regressions separately for the periods 1977-1995, 1996-2000, and 2000-2015 to account for the change in the earnings questions in 1995 and for the expansion of the CPS sample size after 2000.

[16] See, for example, Daly and Hobijn (2015).

3 Regional Variation in Prices and Wages During the 2000s

3.1 Regional Variation in Prices During the 2000s

Figure 3 and Table 1 explore the extent to which our regional scanner price index is correlated with measures of local economic activity. Specifically, Figure 3 plots the percentage point change in the state's average unemployment rate between 2007 and 2010 against the percent change in the state's scanner price index between 2007 and 2010.[17] Figure 3 shows the variation in P^r. In other words, the results in this Figure are not adjusted for the fact that the tradable share of the goods in our sample differs from the tradable share in the composite consumption good. The unemployment rate data come from the BLS's Local Area Unemployment Statistics. Each observation represents a U.S. state (excluding Alaska and Hawaii). The size of the circle in the figure represents the size of the U.S. state measured by their 2006 population (as reported by the BLS) while the line in the figure represents the weighted OLS regression line. In particular, we regress:

$$\ln\left(\frac{P^r_{2010,k}}{P^r_{2007,k}}\right) = \beta_0 + \beta_1 \Delta X_{k,07-10} + \varepsilon_k$$

where $\Delta X_{k,07-10}$ is our measure of the change in economic activity within the state between 2007 and 2010. For Figure 3, $\Delta X_{k,07-10}$ equals the percentage point change in the state unemployment rate between 2007 and 2010.

Figure 3 shows that there is a negative relationship between the change in the state's unemployment rate between 2007 and 2010 and the change in the state's price level between 2007 and 2010. The estimate of β_1 for this specification is -0.35 (standard error = 0.13). This implies that cumulative retail price inflation between 2007 and 2010 was 1.4 percentage points higher in states with a change in the unemployment rate of 6 percentage points during that same time period relative to states with an unemployment rate of 2 percentage points.

Column 1 of Table 1 shows different estimates of β_1 from the above regression with different measures of changing local economic activity ($\Delta X_{k,07-10}$). Each row in Table 1 has a different measure of the changing economic conditions within the state. For example, the first row has the change in the BLS unemployment rate in the state as our measure of $\Delta X_{k,07-10}$ (analogous to the results in Figure 3). Other local economic measures in the subsequent rows include the percent change in state per-capita nominal GDP, the percent change in state per-capita total hours worked, the percent change in state housing prices, and the percent change in the state employment rate.[18] Additionally, in some of the empirical work below, we isolate movements in local employment that were correlated with local housing price changes. The last two rows of Table 1 isolate the relationship between local price growth and local unemployment changes (row 6) and local employment changes (row 7)

[17] Our scanner index is monthly. When computing annual price indices for a given state, we simply take the arithmetic mean of the monthly price indices over the year.

[18] The information on state GDP comes from the U.S.'s Bureau of Economic Analysis (BEA). State population and state total employment comes from the BLS. State total hours worked were computed by the authors using micro data from the American Community Survey. State house price data is from the FHFA's repeat sales indices.

that are correlated with changes in local house price growth. As seen from the results in Table 1, all measures of the change in economic activity are correlated with the change in local prices. As local economic conditions deteriorated during the Great Recession (higher change in the unemployment, lower growth rate in the employment rate, lower house price growth, lower change in hours and GDP per capita), the lower the price inflation during Great Recession.

In column 2 of Table 1 we use an estimate of $\bar{\alpha}/\alpha_r$ to translate the local price variation from the grocery sector into local price variation for the composite local consumption good. For our preferred method, we exploit the relationship between local price inflation and local unemployment rate changes using BLS metro area price indices. These indices are only available for 27 MSAs at varying degrees of frequency (monthly, bi-monthly, semi-annually).[19] The BLS cautions that there is a fair bit of noise in these measures but we believe that it is still interesting to assess the patterns between unemployment rate changes and price changes within this data. In particular, the BLS creates both local food price indices and a price index for the total local consumption basket. Using this data, we regress the 3-year inflation rate (either for food or total CPI) at the MSA level on the 3 year change in the unemployment rate during the 2007-2010 period. The results of these simple regressions using the BLS data are strikingly similar to our results using the scanner data. Within the BLS data, we find that a 1 percentage point increase in the local unemployment rate is associated with a 0.34 percentage point decline in the local food inflation rate (standard error = 0.22). Additionally, we find that the relationship between the inflation rate for all goods (composite local price index) with the change in the unemployment rate is -0.47 (standard error = 0.15). The fact that the coefficient on the inflation rate for the composite local consumption bundle is larger in magnitude than the coefficient on the inflation rate for food inflation is consistent with our belief that the tradable share of food is higher than the tradable share of the local composite consumption good. Given these coefficients, the BLS data suggests a measure of $\bar{\alpha}/\alpha_r$ of 1.4 (-0.47/-0.34). We will use this as our base adjustment factor throughout the paper. But, as we highlight throughout, using this adjustment factor has little effect on the quantitative results of the paper. As seen in column 2 of Table 1, the scaling of prices only modestly increases the responsiveness of local prices to changes in local real activity.

3.2 Regional Variation in Nominal and Real Wages During the 2000s

Figure 4 shows the cross state variation in log demographic adjusted nominal wages from the ACS data between 2007 and 2010 against the change in the state's unemployment rate during the same time period. As seen from the figure, nominal wage growth is also strongly correlated with changes in the unemployment rate during the 2007-2010 period. A simple linear regression through the

[19]In the online appendix that accompanies this paper, we discuss the BLS local price indices in greater depth. Additionally, Fitzgerald and Nicolini (2014) use the BLS MSA level price indices to examine the relationship between local prices and local unemployment rates over a much longer time period. They find that over the period of 1976-2010, a 1 percentage point increase in the local unemployment rate is associated with a 0.3 percentage point decline in the local annual inflation rate. These estimates are in line with the results we highlight in the scanner data for the full sample of U.S. states.

data (weighted by the state's 2006 labor force) suggests that a 1 percentage point change in the state unemployment rate is associated with a 1.2 percentage point decline in nominal wage growth (standard error = 0.2). In Table 2 (column 1) we show that the growth in local nominal wages was highly correlated with changes in many measures of state economic activity during the 2007-2010 period. For example, lower GDP growth, lower employment growth, lower hours growth and lower house price growth were all strongly correlated with lower nominal wage growth during the recent recession.

The second and third columns of Table 2 show the coefficient on the change in local economic activity between 2007 and 2010 from a regression of real wage growth in a given state during that time period on the change in local economic activity. In column 2, we compute local real wages by deflating local nominal wage growth by the growth in the local scanner price index. In column 3, we compute local real wages by deflating local nominal wage growth with the growth in the prices of a composite local consumption good. As discussed above, we scale the growth in the scanner price index by a factor of 1.4 to account for the fact that grocery/mass merchandising goods have a higher tradable share than the composite consumption good. Not surprising, the coefficients in column 2 of Table 2 are roughly equal to the coefficients from column 1 of Table 2 less the coefficients from columns 1 and 2 of Table 1. In all specifications, real wages fell as measures of local economic conditions worsened. For example, a 1 percentage point increase in the unemployment rate was associated with a 0.75 and 0.89 percentage point decline in real wage growth during the 2007 to 2010 period (depending on the scaling factor).

The patterns we document in Table 2 and Figure 4 also show up in other wage series. While there are no government data sets that produce broad based composition adjusted wage series at the local level, the Bureau of Labor Statistics's Quarterly Census of Employment and Wages (QEW) collects firm level data on employment counts and total payroll at local levels. These measures are broad based in that the underlying data are collected as part of the state and federal unemployment insurance programs and covers roughly 98 percent of workers in the U.S.. Using this data, yearly earnings-per-worker can be computed at the state level. This measure is an imperfect measure of wages in that it is not adjusted for cyclical movements in hours worked. Additionally, the measure does not adjust for changes in the composition of workers over the business cycle. Finally, the earnings measures reported include wages and salary as well as bonuses, stock options, and in some states, contributions to deferred compensation plans. These latter measures are not included in the ACS wage indices.

Despite these differences, the cross state correlation between growth in our composition adjusted wage index from the ACS and the growth in earnings-per-worker from QEW is quite high. Appendix Figure A1 shows the simple scatter plot of the growth in the ACS and QEW wage measures between 2007 and 2010. If we fit a line through the scatter plot, the slope coefficient is 0.72 (standard error = 0.20). The correlation between the two measures is about 0.5. Appendix Figure A2 shows that even within the QEW data, there is a strong relationship between employment growth and earnings-per-worker growth during the Great Recession. The x-axis of Appendix Figure A2 is QEW employment growth between 2007 and 2010. QEW employment growth is essentially perfectly correlated with the employment growth measure we use from the BLS in Tables 1 and 2. The y-axis of Figure A2 is QEW nominal earnings per worker growth between 2007 and 2010. As seen from the figure, places with lower employment growth had lower nominal wage growth. The slope coefficient from the line in the scatter plot is 0.45 (standard error = 0.07). This is very similar to the estimated relationship between employment growth and ACS nominal growth during the same time period as show in Table 2.

The fact that nominal wages are moving in response to local employment changes during the Great Recession seems to be a robust finding across different data series and can be used to inform wage stickiness at the aggregate level. This finding is also consistent with the extensive literature in labor economics and public finance showing that local labor demand shocks cause both employment and wages to vary together in the short to medium run. For example, Blanchard and Katz (1991), Autor, Dorn and Hanson (2013) and Charles, Hurst and Notowidigdo (2013) all find that negative local labor demand shocks cause substantial declines in local wages over the three to five year horizon. These results also suggest that wages are fairly flexible in response to labor demand shocks at the local level.

3.3 Comparing Cross-State Patterns to Aggregate Time Series Patterns

One natural question is whether the cross-region patterns during the Great Recession were substantively different from the aggregate time series patterns during the Great Recession. Table 3 sheds light on this issue. The top two rows of Table 3 show the 2007 and 2010 aggregate level of the June unemployment rate (column 1), composition adjusted nominal wages from the CPS (column 2), the June CPI (column 3), and composition adjusted real wages where the June CPI is used to deflate the data (column 4). The third row of Table 3 shows the actual change in these aggregate variables between 2007 and 2010. Given that aggregate prices and wages have been trending over time, one would want an "expected" change in these variables to which the actual change can be compared. Although somewhat ad hoc, we compute a linear trend in these variables between 2000 and 2007 and use that linear trend to predict what wages and prices would have done between the 2007 and 2010 period.[20]

Row 5 takes the difference between the actual change in nominal wages, prices and real wages between 2007 and 2010 and the predicted change during that same time period. Given our rough estimates, nominal wages grew 1.7 percent less during the Great Recession relative to trend. Prices, on the other hand, grew 2.9 percent less than trend during the 2007-2010 time period. The two facts together imply that real wages actually grew 1.4 percent more relative to trend during the recession. To compare these to our cross region elasticities from Tables 1 and Table 2, we divide the aggregate deviation from predicted growth (row 5) by the aggregate change in the unemployment rate during the 2007 to 2010 period. We refer to this as the aggregate semi elasticity of prices, nominal wages

[20]We thank Bob Hall for giving us this idea. We based Table 3 on the analysis he did as part of his discussion of our paper at the 2015 NBER summer EFG program meeting.

and real wages to a 1 percentage point change in the unemployment rate.

The key results from this analysis is that the response of wages to changes in unemployment were *much* stronger at the local level during the Great Recession than at the aggregate level. At the aggregate level, a 1 percentage point increase in the unemployment rate was associated with only a 0.33 percent decline in nominal wage growth over three years. As discussed above, using cross region variation, a 1 percentage point increase in the local unemployment rate was associated with a 1.24 percent decline in nominal wage growth over three years. This is consistent with visual evidence shown in Figures 2 and 4. The price-unemployment relationship at the local level, however, is nearly identical to the price-unemployment relationship at the aggregate level during the 2007-2010 period.[21] Given this, the relationship between real wage growth (relative to trend) and unemployment changes at the aggregate level was actually positive during the Great Recession. However, at the local level, real wage growth plummeted as unemployment increased. It is these differences in the relationships between prices, wages and employment at the local level and at the aggregate level that forms the basis of the remainder of this paper. Why did local wages adjust so much when local employment conditions deteriorated during the Great Recession while aggregate wages hardly responded at all despite a sharp deterioration in aggregate employment conditions? We turn to answering that question next.

4 A Model of a Monetary Union

To help explain the difference between the local and aggregate economies, we build a model of local economies that aggregate. The goals of the model we develop are fivefold. First, as just noted, the model shows conditions under which local economies can aggregate. Second, within the model, we discuss our assumptions on how wages are set. The nominal wage rigidity we specify will be essential to our identification strategy in later parts of the paper. Third, the model allows us to calibrate local and aggregate employment, price and wage elasticities to a variety of different shocks. While it is well known that aggregate and local employment elasticities to the same underlying shock can differ, there is little known about the quantitative differences between aggregate and local elasticities to a given shock. A calibrated version of our model allows us to quantify these differences. This could provide guidance to researchers who want to take an estimated local elasticity to a given shock and apply it to the aggregate economy. Fourth, the model provides us with theoretical co-movement between variables that helps us in interpreting some of the shocks we recover. Finally, the model provides an example of a given economy that is nested in our VAR procedure in Section 5 of the paper. The VAR approach will allow us to estimate shocks for a broader set of models than the one we write down in this section. However, as we have just argued, it is useful to have an example economy that maps into the broader VAR specification.

[21] As we discuss later, there was some deviation between prices at the local level and prices at the aggregate level during the 2010-2012 period as aggregate price growth returned to trend despite the labor market remaining very weak.

Formally, our model economy is composed of many islands inhabited by infinitely lived households and firms in two distinct sectors that produce a final consumption good and intermediates that go into its production. The only asset in the economy is a one-period nominal bond in zero net supply where the nominal interest rate is set by a monetary authority. We assume intermediate goods are traded across islands but the consumption good is non-tradable.[22] Finally, we assume labor is mobile across sectors but not across islands. We explore the issue of labor mobility during the Great Recession when we take the model to the data. Throughout we assume that parameters governing preferences and production are identical across islands and the islands only differ, potentially, in the shocks that hit them.

4.1 Firms and Households

Producers of tradable intermediates x in island k use local labor N_k^x and face nominal wages W_k (equalized across sectors) and prices Q (equalized across islands k). Their profits are

$$\max_{N_k^x} Qe^{z_k^x}(N_k^x)^\theta - W_k N_k^x$$

where z_k^x is a tradable productivity shock in island k and $\theta < 1$ is the labor share in the production of tradables. Final (retail) goods y producers face prices P_k and obtain profits

$$\max_{N_k^y, X_k} P_k e^{z_k^y}(N_k^y)^\alpha (X_k)^\beta - W_k N_k^y - QX_k$$

where z_k^y is a final good (retail) productivity shock and $(\alpha, \beta) : \alpha + \beta < 1$ are the labor and intermediates shares. Unlike the tradable goods prices, final good prices (P_k) vary across islands.[23]

Households preferences are given by

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

where C_{kt} is consumption of the final good, N_{kt} is labor, δ_{kt} and ϵ_{kt} are exogenous processes driving the household's discount factor and the disutility of labor, respectively. Our base preferences abstract from income effects on labor supply. However, as we show in section 7.4, relaxing this assumption does not quantitatively change the conclusions of the paper.

Households are able to spend their labor income $W_{kt}N_{kt}$ plus profits accruing from firms Π_{kt}, financial income $B_{kt}i_t$ and transfers from the government T_t, where B_{kt} are nominal bond holdings

[22] The final good can be thought of as being retail: restaurants, barbershops and stores; and the intermediate sector providing physical goods: food ingredients, scissors and cellphones.

[23] It is worth noting that all model shocks will generate endogenous variation in markups given our assumption of decreasing returns to scale. Additionally, what we call a "productivity shock" is isomorphic to any shifter of unit labor costs and, hence, labor demand schedules. Later we will refer to it as the productivity/markup shock. We do not attempt to distinguish between the different interpretations of this shock in this paper.

at the beginning of the period and i_t is the nominal interest (equalized across islands given our assumption of a monetary union where the bonds are freely traded) on consumption goods (C_{kt}) and savings ($B_{kt+1} - B_{kt}$). Thus, they face the period-by-period budget constraint

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

A well known issue in the international macroeconomics literature is that under market incompleteness of the type we just described there is no stationary distribution for bond holdings across islands in the log-linearized economy; and all other island variables in the model have unit roots. This is problematic for reasons both theoretical (we will like to study log-deviations from a deterministic steady state) and empirical (regional data for the US does not suggest the presence of such unit roots). We follow Schmitt-Grohe and Uribe (2003) and let ρ_{kt} be the endogenous component of the discount factor that satisfies $\rho_{kt+1} = \rho_{kt} + \Phi(.)$ for some function $\Phi(.)$ of the average per capita variables in an island. As such, agents do not internalize this dependence when making their choices. This modification induces stationarity for an appropriately chosen function $\Phi(.)$. Schmitt-Grohe and Uribe (2003) show that alternative stationary inducing modifications (a specification with internalization, a debt-elastic interest rate or convex portfolio adjustment costs) all deliver similar quantitative results in the context of a small open economy real business cycle model.

4.2 Sticky wages

We allow for the possibility that nominal wages are rigid and use a partial-adjustment model where a fraction λ of the gap between the actual and frictionless wage is closed every period. Formally:

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

Given our assumption on household preferences, $P_{kt}e^{\epsilon_{kt}}(N_{kt})^{\frac{1}{\phi}}$ is the marginal rate of substitution between labor and consumption and the parameter λ measures the degree of nominal wage stickiness. In particular, when $\lambda = 1$ wages are fully flexible and when $\lambda = 0$ they are fixed. This implies that workers will be off their labor supply curves whenever $\lambda < 1$. A similar specification has been used by Shimer (2010) and, more recently, by Midrigan and Philippon (2011). Shimer (2010) argues that in labor market search models there is typically an interval of wages that both the workers are willing to accept and firms willing to pay. To resolve this wage indeterminacy he considers a wage setting rule that is a weighted average of a target wage and the past wage. The target wage in our case is the value of the marginal rate of substitution.

Popular alternatives in the literature include the wage bargaining model in the spirit of Hall and Milgrom (2008) as in Christiano, Eichenbaum and Trabandt (2015b); and the monopsonistic competition model where unions representing workers set wages period by period as in Gali (2009). The key difference with the partial adjustment model is that both alternatives result in a forward looking component in the wage setting rule that is absent in our specification. In fact, this wage setting rule can be derived from the monopsonistic competition setup in the case where agents are myopic about the future; or the labor market search setup in the special case where firms make take it or leave it offers and the probability of being employed in the future is independent of the current employment status. While there is no forward looking component in the reset wage in our base specification, we consider the implications of including forward looking behavior in Section 7.4 below.

4.3 Equilibrium

An equilibrium is a collection of prices $\{P_{kt}, W_{kt}, Q_t\}$ and quantities $\{C_{kt}, N_{kt}, B_{kt}, N^x_{kt}, N^y_{kt}, X_{kt}\}$ for each island k and time t such that, for an interest rate rule $i_t = i(.)e^{\mu_t}$ and given exogenous processes $\{z^x_{kt}, z^y_{kt}, \epsilon_{kt}, \delta_{kt}, \mu_t\}$, they are consistent with household utility maximization and firm profit maximization and such that the following market clearing conditions hold:

$$
\begin{aligned}
C_{kt} &= e^{z^y_{kt}}(N^y_{kt})^\alpha X^\beta_{kt} \\
N_{kt} &= N^y_{kt} + N^x_{kt} \\
\sum_k X_{kt} &= \sum_k e^{z^x_{kt}}(N^x_{kt})^\theta \\
\sum_k B_{kt} &= 0
\end{aligned}
$$

4.4 Shocks

We assume the exogenous shocks follow an AR(1) process, with an identical autoregressive coefficient across islands (and sectors in the case of productivity), and that the innovations are iid, mean zero, random variables with an aggregate and island specific component. First, define $\gamma_{kt} \equiv \delta_{kt} - \delta_{kt-1} - \mu_t$. This is a combination of the discount rate shock and the monetary policy shock. γ_{kt} will show up as a wedge in the Euler equation. Then,

$$
\begin{aligned}
z^y_{kt} &= \rho_z z^y_{kt-1} + \sigma_z u^y_t + \tilde{\sigma}_y v^y_{kt} \\
z^x_{kt} &= \rho_z z^x_{kt-1} + \sigma_z u^x_t + \tilde{\sigma}_x v^x_{kt} \\
\gamma_{kt} &= \rho_\gamma \gamma_{kt-1} + \sigma_\gamma u^\gamma_t + \tilde{\sigma}_\gamma v^\gamma_{kt} \\
\epsilon_{kt} &= \rho_\epsilon \epsilon_{kt-1} + \sigma_\epsilon u^\epsilon_t + \tilde{\sigma}_\epsilon v^\epsilon_{kt}
\end{aligned}
$$

with $\sum_k v^y_{kt} = \sum_k v^x_{kt} = \sum_k v^\gamma_{kt} = \sum_k v^\epsilon_{kt} = 0$. By assumption, we assume the weighted average of the regional shocks sum to zero in all periods.

Let $u^z_t \equiv u^y_t + \beta u^x_t$. We will call u^z_t, u^γ_t and u^ϵ_t the aggregate *Productivity/Markup, Discount rate* and *Leisure shocks* respectively. These are the innovations that the econometric procedure aims to identify. Analogously, $v^y_{kt}, v^x_{kt}, v^\gamma_{kt}, v^\epsilon_{kt}$ are the *Regional shocks*. The interpretation of the Leisure and Productivity/Markup shocks is relatively straightforward given our model environment. They are shifters of households and firms' labor supply (wage setting) and labor demand schedules respectively. On the other hand, what we identify as a "discount rate shock" (γ_{kt}) is really the combination of two more fundamental shocks. First, an innovation to the marginal rate of substitution between consumption in consecutive periods. Second, an innovation in the nominal interest rate rule set by the monetary authority. Our procedure is unable to distinguish between the two given that they both show up in the household's Euler equation and, hence, we treat it as a single shock.

4.5 Aggregation

Our first key assumption for aggregation is that all islands are identical with respect to their underlying production parameters (α, β, and θ), their underlying utility parameters (σ and ϕ) and the degree of wage stickiness (λ).[24] Our second assumption is that the islands are identical in the steady state and that price and wage inflation are zero. The last assumption is that the joint distribution of island-specific shocks is such that its cross-sectional summation is zero. If K, the number of islands, is large this holds in the limit because of the law of large numbers. We log-linearize the model around this steady state and show that it aggregates up to a representative economy where all aggregate variables are independent of any cross-sectional considerations to a first order approximation.[25] We denote with lowercase letters a variable's log-deviation from its steady state. Also, variables without a k subscript represent aggregates. For example, $n_{kt} \equiv log\left(\frac{N_{kt}}{N}\right)$and $n_t \equiv \sum_k \frac{1}{K} n_{kt}$. We assume that the monetary authority announces the nominal interest rate rule in log-linearized form: $i_{t+1} = \varphi_\pi \mathbb{E}_t[\pi_{t+1}] + \varphi_y(y_t - y_t^*) + \mu_{t+1}$ where π_t is the aggregate inflation rate and $y_t - y_t^*$ is the output gap; defined as the difference between output and the flexible wage equilibrium output for the same realization of shocks. Finally, we assume that the endogenous component of the discount factor is $\Phi(.) = \Phi_0\,(c_{kt} - c_t)$.[26]

The following lemmas present a useful aggregation result and show that we can write the island level equilibrium in deviations from these aggregates. Let w_t^r be real wage growth and π_t^w be nominal wage growth. Formally, $w_t^r \equiv log\left(\frac{W_t/P_t}{W/P}\right)$, $\pi_t^w = w_t^r - w_{t-1}^r + \pi_t$ and $z_t \equiv z_t^y + \beta z_t^x$.

Lemma 1 *The behavior of π_t^w, w_t^r, n_t in the log-linearized economy is identical to that of a representative economy with only a final goods sector with labor share in production $\alpha + \theta\beta$, no endogenous discount factor and only 3 exogenous processes $\{z_t, \epsilon_t, \gamma_t\}$.*

[24]When implementing our procedure empirically using data on US states, we discuss the plausibility of this assumption. Given that the broad industrial compositon at the state level does not differ much across states, the assumption that productivity parameters and wage stickiness are roughly similar across states is not dramatically at odds with the data. As a robustness exercise, we estimate our key equations with industry fixed effects and show that our key cross section estimates are unchanged.

[25]The model we presented has many islands subject to idiosyncratic shocks that cannot be fully hedged because asset markets are incomplete. By log-linearizing the equilibrium we gain in tractability, but ignore these considerations and the aggregate consequences of heterogeneity. As usual, the approximation will be a good one as long as the underlying volatility of the idiosyncratic shocks is not too large. If our unit of study was an individual, as for example in the precautionary savings literature with incomplete markets, the use of linear approximations would likely not be appropriate. However, since our unit of study is an island the size of a small country or a state we believe this is not too egregious of an assumption. The volatilities of key economic variables of interest at the state or country level are orders of magnitude smaller than the corresponding variables at the individual level.

[26]When $\Phi_0 > 0$ this will be enough to induce stationary of island level variables in log-deviations from the aggregate. At the same time, since $\Phi(.)$ depends only on these deviations, the aggregate equilibrium will feature a constant endogenous discount factor ρ.

Denote variables $\tilde{x}_t \equiv x_{kt} - x_t$ as island k log-deviation from aggregates at time t, where the subscript k is dropped for notational simplicity.

Lemma 2 *For given $\{\tilde{z}^y_t, \tilde{z}^x_t, \tilde{\gamma}_t, \tilde{\epsilon}_t\}$, the behavior of $\{\tilde{p}_t, \tilde{w}_t, \tilde{n}^y_t, \tilde{n}^x_t\}$ in the log-linearized economy for each island in deviations from aggregates is identical to that of a small open economy where the price of intermediates and the nominal interest rate are at their steady state levels, i.e. $q_t = i_t = 0\ \forall t$.*

Proof. See Appendix A for a proof of Lemma 1 and 2. ■

4.6 Aggregate vs. local shock elasticities

With the model in hand, we can now explore the extent to which aggregate employment, price and wage elasticities to a given shock differ from local employment, price and wage elasticities to the same shock. Many researchers use clever identification strategies exploiting regional variation to estimate local elasticities to a given shock. For example, Mian and Sufi (2014), use debt variation across U.S. metropolitan areas to isolate the extent to which something akin to a local discount rate shock affects local employment. It is our goal in this sub-section to show that the local employment (price, wage) elasticity to a given discount rate shock (productivity shock, leisure shock) is not the same as the aggregate elasticity to the same shock. Moreover, we want to calibrate the model to quantify the difference.

To gain some intuition as to the difference between local and aggregate elasticities in our model, we first consider the special case where there is an endowment of the tradable good and no labor is used in its production, i.e. $\theta = 0$. Focusing on a discount rate shock in this special case makes the comparison very transparent. We let $\zeta^{agg}_0 \equiv \frac{dn_0}{d\gamma_0}$ and $\zeta^{reg}_0 \equiv \frac{d\tilde{n}_0}{d\tilde{\gamma}_0}$ be the employment elasticities to the discount rate shock on impact. By solving for the recursive laws of motion in equilibrium we obtain,

$$\zeta^{agg}_0 = \frac{(1-\lambda)}{(1-\alpha+\frac{\lambda}{\phi})(\varphi_p - 1) + \left(\varphi_y \alpha - (\varphi_p - 1)(1-\alpha)\right)\frac{1-\lambda}{\rho_\gamma}}$$

$$\zeta^{reg}_0 = \frac{d\tilde{n}_0}{d\tilde{\gamma}_0} = \frac{(1-\lambda+\lambda\beta)}{\left(1-\alpha+\frac{\lambda}{\phi}+\left(\frac{\sigma(1-\lambda(\alpha-\frac{1}{\phi}))}{1-\frac{\alpha\phi}{1+\phi}} - (1+\frac{\lambda}{\phi})\right)\beta\right)\left(\frac{1+r}{\rho_\gamma} - 1\right)}$$

These expressions help understand the general equilibrium forces that make local and aggregate elasticities different. From the perspective of the closed economy, the endogenous response of the nominal interest rate rule $\{\varphi_p$ and $\varphi_y\}$ reduces the aggregate employment impact elasticity to an unanticipated discount rate shock. A negative discount rate shock puts downward pressure on employment and prices. The monetary authority can lower interest rates to offset such a shock. These parameters of the interest rate rule are entirely absent in the expression for the regional elasticity. Hence, from these forces, the aggregate employment elasticity to a discount rate shock is inherently smaller than the local employment elasticity to a local discount rate shock.

From the local perspective, since island level economies in deviations from the aggregate are small open economies, there are two extra margins of adjustment that are absent in the aggregate closed economy. First, the possibility to substitute labor for intermediate goods in the production of final consumption goods ($\beta > 0$) which will decrease the regional employment elasticity to the shock (as long as the term $\left(\frac{\sigma(1-\lambda(\alpha-\frac{1}{\phi}))}{1-\frac{\phi}{1+\phi}\alpha} - 1\right)$ is positive). Second, the possibility to transfer resources intertemporally through saving/borrowing at the interest rate r, as seen in the term $\left(\frac{1+r}{\rho_\gamma} - 1\right)$, decreases the regional employment elasticity. Theoretically, therefore, the aggregate employment elasticity to an aggregate discount rate shock can be either greater or smaller than the local employment elasticity to a local discount rate shock.

It is also interesting to compare how these discount rate elasticities change with the degree of nominal wage stickiness. Our identification procedure allows us to do this exercise when we estimate the impulse response to a discount rate shock. When $\varphi_p > 1$, both elasticities are decreasing in λ. In particular, employment does not respond to discount rate shocks at all in the limit when wages are perfectly flexible ($\lambda \to 1$). This is the standard intuition in New Keynesian models where some sort of nominal or real rigidity is necessary to get real effects to demand shocks.

While it is generally understood that local and aggregate elasticities can differ, there has been little quantitative work assessing the potential size of these differences. A parameterized version of our model can allow us to directly compute the local and aggregate employment elasticities to different types of shocks. To this end, Table 5 quantifies the employment impact elasticities to each of the shocks in the full model. Table 4 presents and explains the parameterization of our model. Most of the parameters are standard from the literature or are chosen to match the labor share in the tradable and non-tradable sectors. The Online Appendix has an extended discussion of our baseline parameter choice. For our base specification, we use estimates of λ and ϕ of 2 and 0.7, respectively. These are the parameters that show up in the aggregate and local wage setting equations. The value of these parameters are the ones that we estimate using local variation in Section 6.

Column 1 of Table 5 shows our base estimates of the local and aggregate employment elasticities. In columns 2 - 8 of Table 5, we show how the elasticities change across alternate parameterization. Specifically, in column 2, we re-compute the elasticities reducing the Frisch elasticity of labor supply (ϕ) from 2 to 1. In column 3, we make wages more sticky by reducing λ from 0.7 to 0.5 (returning the Frisch elasticity to our base parameterization). In column 4, we set $\beta = 0$ shutting down the possibility to substitute labor for intermediate goods in the production of final goods. In the next two columns, we shut down the endogenous feedback in the nominal interest rate to changes in the employment gap such that φ_y is set to zero. In this first of those two columns, we leave the response of the nominal interest rate to the inflation target (φ_p) at its base parameterization. In the second of those two columns, we lower φ_p such that the local and aggregate responses to a discount rate shock are the same on impact. Finally, in the last two columns, we explore how the elasticities change as the persistence of the demand shock changes.

In our base specification, we find that the regional employment elasticity to a discount rate shock is 2.3 times larger than the aggregate employment elasticity to a discount rate shock. This implies that the growing literature that uses cross region variation to estimate local employment elasticities to demand shocks dramatically overstates employment responses when those local elasticities are applied to the aggregate. Across the different parameterizations of the wage setting rule shown in columns 2-3, the conclusion remains unchanged. Local employment elasticities to discount rate shocks are always two to three times larger than the aggregate employment elasticities. In columns 4-6, we see the importance of the general equilibrium forces. As we shut down the ability to substitute labor for intermediate goods ($\beta = 0$), the gap between the regional and aggregate elasticities gets larger. The ability to trade intermediates across regions dampens the local employment elasticity to discount rate (demand) shocks. In columns 5 and 6, we see that the endogenous monetary policy response also dramatically dampens the aggregate response to a discount rate shock. This suggests that in periods where the economy is at the zero lower bound, aggregate and local employment elasticities to a demand shock are more similar, a point also made in Nakamura and Steinsson (2014). The last column explores the robustness to changes in the persistence of the discount rate shock. The more temporary is the local discount rate shock, the smaller the local employment elasticity because the regions can borrow and lend with each other.

Given that the literature has tried to estimate local employment responses to well identified local demand shocks, we have focused our discussion so far on the employment responses to discount rate shocks. In general, the local employment elasticities to a local demand shock are much larger than aggregate employment elasticities to an aggregate demand shock. This is because, at the aggregate level, the monetary authority can change interest rates to offset the discount rate shock. Table 5 also shows the local and aggregate employment response to local and aggregate productivity/markup and labor supply shocks. For these other shocks, the local employment elasticities are usually *smaller* than their aggregate counterparts. Much of this has to do with the specification of the nominal interest rate rule. What we want to stress, however, is that the quantitative difference between aggregate and local employment elasticities depends on the underlying shock and can be quite large.

Tables 6 and 7 summarize the aggregate and regional impulse responses, respectively, for all variables and shocks in our benchmark calibration. We get a sense of the impulse response by showing the results upon impact (the short run elasticities) and after 5 years (the long run elasticities). These tables allow one to assess the model predictions. We use the same parameterization as in Table 4. The short run responses in Columns 1 of Table 6 and Table 7 just restate the employment elasticities in column 1 of Table 5. The remainder of the tables show the estimates for the price, nominal wage and real wage elasticities to all the underlying shocks in the model upon impact. As seen from Table 6, an aggregate negative discount rate shock (households become less patient) lowers aggregate employment, lowers aggregate prices, and lowers (slightly) aggregate real wages. Conversely, an aggregate negative productivity shock lowers aggregate employment, raises aggregate prices, and raises aggregate real wages. We will compare these to the estimated impulse responses in the vector-autoregression model we estimate in Section 5.

5 A Semi-structural Approach to Estimating Aggregate and Local Shocks

The above model was designed to (1) link the aggregate and regional economies, (2) specify the local and aggregate wage setting equations, (3) provide a quantitative assessment of aggregate and local elasticities to the shocks embedded in the model, and (4) guide our interpretation of some of the shocks that we recover using the procedure in this section.[27] However, the simple model abstracts from many features which could be important in quantifying the underlying aggregate and local shocks. In this section, we develop a procedure that allows us to estimate the aggregate and local shocks under a broader class of models. We can be agnostic with respect to the other components of the model aside from the wage setting equation and still identify the underlying shocks. To do so, we estimate a semi-structural VAR where the structure comes from imposing theoretical restrictions to help identify the underlying shocks. As an overview, the structural shocks and impulse response matrix are identified by using a set of theoretical restrictions on the joint response of variables in the VAR to each type of shock. We will derive these restrictions from the parameterized wage setting equation developed above. Beraja (2015) discusses this semi-structural identification methodology in detail, as well as its application to more general VARs and theoretical models than the ones in this paper.

5.1 Estimating the Aggregate shocks

The recursive solution to the equilibrium system of equations in Lemma 1 can be written in reduced form as a VAR(∞) in $\{\pi_t, \pi_t^w, n_t\}$.[28]

$$(I - \rho(L)) \begin{bmatrix} \pi_t \\ \pi_t^w \\ n_t \end{bmatrix} = \Lambda \begin{bmatrix} u_t^\epsilon \\ u_t^z \\ u_t^\gamma \end{bmatrix}$$

With knowledge of $\rho(L)$ and an invertible matrix Λ together with aggregate data on consumer price indices, nominal wages and employment it is possible to recover the structural shocks. Hence, identification of the shocks implies identification of these matrices.

The benefit of estimating VARs to uncover underlying shocks is that VARs nest a larger class of models. However, additional restrictions are needed to identify the shocks. Some identification schemes rely on ordering of the shocks. These identification schemes are often made without

[27]We also performed a business cycle accounting exercise by solving the model and using the data to recover the exogenous stochastic processes. By doing so, we learned that different "wedges" are required to explain the joint dynamics of employment prices and wages. We find that the labor wedge is quite important in explaining variation in employment in the early stages of the Great Recession. Like our VAR results, the Euler equation wedge only explained less than half of the employment decline during the early part of the recession and explained essentially none of the persistence. However, unlike the VAR, the business cycle accounting does not allow us to recover the fundamental shocks and, therefore, we do not report more specific results here. As has been shown in Buera and Moll (2015), slight changes in model specification can alter the mapping between underlying structural shocks and corresponding aggregate wedges.

[28]The exogenous processes are AR(1) and the system of equations characterizing the equilibrium is of first order. When written in matrix form it is easy to show that there is a reduced form representation as a VARMA(1,2) and hence as a VAR(∞) if the moving average part of the process is invertible.

theoretical justification. Other identification schemes rely on long run theoretical restrictions. Our approach only requires the use of one component of the theory developed above to identify the VAR. In particular, we will use the wage setting equation. We do not need to take a stance on the rest of the underlying model (aside from the linearity and the orthogonality of the underlying shocks, which are also made in the other VAR identification schemes) and, as a result, our VAR will be consistent with other models. For example, in other models there may be price stickiness in the firm's setting of prices or habits in an individual's consumption decision. Our VAR will be consistent with these and other models as long as the following wage setting equation in log-linearized form (discussed in the prior section) holds:

$$\pi_t^w = \lambda(\pi_t + \epsilon_t - \epsilon_{t-1} + \frac{1}{\phi}(n_t - n_{t-1})) + (1-\lambda)\pi_{t-1}^w$$

We show that if this one equation is specified, we can identify the aggregate shocks while being agnostic about the remaining structural equations describing the economy. This is because this single structural equation imposes several linear constraints that the reduced form errors must satisfy. At the end of Section 7, we show the robustness of our estimation procedure to other wage setting equations.

The first step in our procedure consists of estimating the reduced form VAR to obtain the autoregressive matrix $\rho(L)$ and the reduced form errors covariance matrix V. In practice we will truncate $\rho(L)$ to be of finite order as it is typically done in the literature.

We now derive the identification restrictions that will allow us to estimate Λ and the shocks. Applying the conditional expectation operator $\mathbb{E}_{t-1}(.)$ on both sides of the above wage setting equation and constructing the reduced form expectational errors we obtain:

$$\begin{bmatrix} \lambda & -1 & \frac{\lambda}{\phi} \end{bmatrix} \Lambda \begin{bmatrix} u_t^\epsilon \\ u_t^z \\ u_t^\gamma \end{bmatrix} + \lambda\sigma_\epsilon u_t^\epsilon = 0 \tag{3}$$

Similarly, constructing $\mathbb{E}_{t-1}(.) - \mathbb{E}_{t-2}(.)$, we obtain:

$$\left([\begin{array}{ccc} \lambda & -1 & \frac{\lambda}{\phi} \end{array}]\rho_1 + [\begin{array}{ccc} 0 & 1-\lambda & 0 \end{array}] \right) \Lambda \begin{bmatrix} u_{t-1}^\epsilon \\ u_{t-1}^z \\ u_{t-1}^\gamma \end{bmatrix} + \lambda(\rho_\epsilon - 1)\sigma_\epsilon u_{t-1}^\epsilon = 0 \tag{4}$$

where ρ_1 is the matrix collecting the first order autoregressive coefficients in the reduced form VAR. If the VAR includes two lags, we can construct $\mathbb{E}_{t-1}(.) - \mathbb{E}_{t-2}(.)$.

The above equations (3) and (4) have to hold for all realizations of the shocks. In particular, equation (3) gives us two linear restrictions in the elements of Λ for given parameters in the wage setting equation when there are either contemporaneous discount rate or productivity/markup shocks. Moreover, from equation (4), we obtain two extra linear restrictions that hold when there is a lagged discount rate shock or a lagged productivity/markup shock. These restrictions, together with the six restrictions coming from the orthogonalization of the shocks, are sufficient to identify all nine elements in the Λ matrix. Intuitively, (3) allows us to "separate" the leisure shock from the discount rate and productivity/markup shocks combined; and (4) "separates" the discount rate from the productivity/markup shock.

It is worth noting that there is a sense in which the shocks are not completely identified. The issue arises because the above procedure does not allow us to label which of the shocks that do not appear in the wage setting equation corresponds to u_t^z or u_t^γ. Basically, the linear restrictions from equation (3) and (4) are identical for both discount rate and productivity/markup shocks. A solution to this labeling problem is to use the theoretical co-movement on impact of employment, wages and prices from our model after a u_t^γ and u_t^y shock, respectively, and label the estimated shocks accordingly. To label the shocks we search over all linear combinations $\psi \in [0,1]$ of the independent restrictions coming from equation (4) such that a discount rate (productivity shock/markup) shock: (i) moves prices and employment in the same (opposite) direction on impact; and (ii) moves real wages and employment in opposite (same) direction on impact. If more than one linear combination of the restrictions satisfy these, we pick the one that is closer to giving equal weighting to both restrictions. The qualitative co-movement of these variables that we require to label the shocks is consistent with the signs of the impact elasticities in our model from Section 4. This is the approach we will follow when we apply the procedure to identify the shocks that hit the US economy during the Great Recession.

For completeness, the matrix Λ solves the system:

$$
\begin{bmatrix} \lambda & -1 & \frac{\lambda}{\phi} \end{bmatrix} \Lambda \begin{bmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 0 & 0 \end{bmatrix}
$$

$$
\left(\begin{bmatrix} \lambda & -1 & \frac{\lambda}{\phi} \end{bmatrix} \rho_1 + \begin{bmatrix} 0 & 1-\lambda & 0 \end{bmatrix} \right) \Lambda \begin{bmatrix} 0 \\ \psi \\ 1-\psi \end{bmatrix} = 0
$$

$$
\Lambda \Lambda' = V
$$

It is worth noting that, when $\lambda = 1$, this procedure cannot identify all columns in the impulse response matrix Λ because the system above is underdetermined. This is because equation (4) implies linear restrictions that are merely linear combinations of the restrictions implied by equation (3). Thus, some degree of wage rigidity is key for identification of the shocks through this semi-structural procedure.

5.2 Estimating the Regional Shocks

The procedure for estimating regional shocks and elasticities follows a similar logic than the one for aggregate shocks. The recursive solution to the equilibrium system of equations in Lemma 2 can be written in reduced form as a VAR(∞) in $\{\tilde{p}_t, \tilde{w}_t, \tilde{n}^y_t\}$ when $\tilde{v}^\epsilon_t = 0$. Given that there are four shocks at the local level, we need one further identification restriction. We set $\tilde{v}^\epsilon_t = 0$. We provide some evidence for this choice in the next section. The regional VAR, therefore, can be expressed as follows:

$$(I - \tilde{\rho}(L)) \begin{bmatrix} \tilde{p}_t \\ \tilde{w}_t \\ \tilde{n}^y_t \end{bmatrix} = \tilde{\Lambda} \begin{bmatrix} \tilde{v}^y_t \\ \tilde{v}^x_t \\ \tilde{v}^\gamma_t \end{bmatrix}$$

Again, this reduced form vector autoregression representation is consistent with a much more general class of models than the one characterized in Lemma 2. From here on we will be working with a subset of these general class of models such that the regional wage setting equation in log-linearized form holds,

$$\tilde{w}_t = \lambda \left(\tilde{p}_t + \frac{N^y}{N\phi} \tilde{n}^y_t + \frac{N - N^y}{N\phi(1-\theta)} (\tilde{z}^x_t - \tilde{w}_t) \right) + (1-\lambda)\tilde{w}_{t-1}$$

This is obtained from replacing the tradable goods labor demand and labor market clearing condition into the wage setting equation.

The first step in the procedure consists in estimating the reduced form VAR via OLS to obtain the autoregressive matrix $\tilde{\rho}(L)$ and the reduced form errors covariance matrix $\tilde{V}$. In practice we will truncate $\tilde{\rho}(L)$ to be of finite order.

We now derive the identification restrictions that will allow us to estimate $\tilde{\Lambda}$ and the shocks. Applying the conditional expectation operator $\mathbb{E}_{t-1}(.)$ on both sides of the wage setting equation and constructing the reduced form expectational errors we obtain,

$$\begin{bmatrix} 1 & -\left(\frac{1}{\lambda} + \frac{N-N^y}{N\phi(1-\theta)}\right) & \frac{N^y}{N\phi} \end{bmatrix} \Lambda \begin{bmatrix} \tilde{v}^y_t \\ \tilde{v}^x_t \\ \tilde{v}^\gamma_t \end{bmatrix} + \frac{N - N^y}{N\phi(1-\theta)} \sigma^x \tilde{v}^x_t = 0 \tag{5}$$

Similarly, constructing $\mathbb{E}_{t-1}(.) - \mathbb{E}_{t-2}(.)$, obtain

$$\left(\begin{bmatrix} 0 & \frac{1-\lambda}{\lambda} & 0 \end{bmatrix} + \begin{bmatrix} 1 & -\left(\frac{1}{\lambda} + \frac{N-N^y}{N\phi(1-\theta)}\right) & \frac{N^y}{N\phi} \end{bmatrix} \tilde{\rho}_1 \right) \Lambda \begin{bmatrix} \tilde{v}^y_{t-1} \\ \tilde{v}^x_{t-1} \\ \tilde{v}^\gamma_{t-1} \end{bmatrix} - \frac{1}{\phi} \frac{N - N^y}{N(1-\theta)} \rho^x \sigma^x \tilde{v}^x_{t-1} = 0$$

where $\tilde{\rho}_1$ is the matrix collecting the first order autoregressive coefficients in the reduced form VAR. As with the procedure for identifying aggregate shocks, we can identify the impulse matrix $\tilde{\Lambda}$ with these extra linear restrictions and we can label the shocks accordingly using the same restrictions on the model-implied co-movement on impact of the variables in the VAR.

6 Estimating Parameters of the Wage Setting Equation Using Regional Data

In this section, we discuss how we estimate λ and ϕ which are necessary for our shock decomposition procedure. Given the above assumptions, the aggregate and local wage setting equations can be expressed as:

$$\pi_t^w = \lambda(\pi_t + \frac{1}{\phi}(n_t - n_{t-1}) + (1-\lambda)\pi_{t-1}^w + \lambda(u_t^\epsilon - (1-\rho_\epsilon)\epsilon_{t-1})$$

$$\pi_{kt}^w = \lambda(\pi_{kt} + \frac{1}{\phi}(n_{kt} - n_{kt-1})) + (1-\lambda)\pi_{kt-1}^w + \lambda(u_t^\epsilon - (1-\rho_\epsilon)\epsilon_{t-1}) + \lambda v_{kt}^\epsilon$$

The aggregate and local wage setting curves are functions of the Frisch elasticity of labor supply (ϕ) and the wage stickiness parameter (λ). There is a literature on estimating micro and macro labor supply elasticities. However, it is hard to estimate the amount of wage stickiness using aggregate data given the small degrees of freedom inherent in aggregate data and given that at the aggregate level it is hard to isolate movements in employment growth and price growth that are arguably uncorrelated with the aggregate labor supply shock (u_t^ϵ). In some instances, regional data can be used to estimate these parameters.

In order for regional data to be used to estimate λ and ϕ, one of the following must hold: (1) the labor supply shock has no regional component ($v_{kt}^\epsilon = 0$) or (2) the regional component of the labor supply shock must be uncorrelated with changes in local economic activity (i.e., $cov(v_{kt}^\epsilon, (n_{kt} - n_{kt-1})) = 0$ and $cov(v_{kt}^\epsilon, \pi_{kt}) = 0$). The latter condition holds if a valid instrument can be found that isolates movement in $n_{kt} - n_{kt-1}$ and π_{kt} that is orthogonal to v_{kt}^ϵ. In this section, we estimate λ and ϕ using the regional data on prices, wages and employment growth during the Great Recession. We argue that local labor supply shocks were small during the Great Recession allowing us to estimate λ and ϕ using OLS. Additionally, we use local house price variation during the early part of the Great Recession as an instrument to isolate movements in $n_{kt} - n_{kt-1}$ and π_{kt} that are orthogonal to local labor supply shocks. Both procedures yield estimates of λ and ϕ that are fairly similar.

6.1 Estimating Equation and Identification Assumptions

Formally, we estimate the following specification using our regional data:

$$\pi_{kt}^w = b_0 + b_1\pi_{kt} + b_2(n_{kt} - n_{kt-1}) + b_3\pi_{kt-1}^w + \Psi D_t + \Gamma X_k + e_{kt}$$

where $b_1 = \lambda$, $b_2 = \lambda/\phi$, $b_3 = (1-\lambda)$, and $b_0 = \lambda(u_t^\epsilon - (1-\rho_\epsilon)\epsilon_{t-1})$. Any aggregate labor supply shocks are embedded in the constant term. The local error term includes λv_{kt}^ϵ as well as measurement error for the local economic variables. We estimate this equation pooling together all annual employment, price and wage data for years between 2007 and 2011. When estimating the above regression, we include year fixed effects (D_t). This ensures that we are only using the cross-region variation to estimate the parameters. We estimate this equation annually because we only have annual measures of wages at the state level. Our annual nominal wage measures at the state level are the composition adjusted nominal log wages computed from the American Community Survey discussed above. π^w_{kt}, therefore, is just the log-growth rate in adjusted nominal wages within the state between year t and $t-1$. Our measure of employment growth at the state level is calculated using data from the U.S. Bureau of Labor Statistics. The BLS reports annual employment counts and population numbers for each state in each year. We divide the employment counts by population to make an annual employment rate measure for each state. $n_{kt} - n_{kt-1}$ is the log-change in the employment rate between year t and $t-1$. π_{kt} is log-change in the average price index in each state i within year t. In our base specification, we use the scanner data local inflation rate scaled to account for the fact that the local non-tradable share in the grocery sector may differ from the composite consumption good. However, in alternate specifications, we use the raw inflation rate from the scanner data as our measure of local inflation. Finally, in some specifications we include controls for the state's industry mix in 2007. This allows for the potential that local labor supply shocks, to the extent that they exist, may be correlated with the state's industry structure. Given that we have observations on 48 states for 4 years of growth rate data, our estimating equation includes 192 observations in our base specification. We also show results using data from 2007-2009 before the large changes in unemployment benefits extension starting in 2010.

Two additional comments are needed about our estimating equation. First, the theory developed above implies that $b_1 + b_3 = 1$. We impose this condition when estimating the cross region regression. Second, we believe our measures of local wage growth and price growth are measured with error. The measurement error, if classical, will attenuate our estimates of b_1 and b_3. Additionally, because we are regressing wage growth on lagged wage growth, any classical measurement error in wages in year t will cause a negative relationship between wage growth today and lagged wage growth. We take these measurement error concerns seriously when estimating the above regression. Specifically, given the large sample sizes on which our wage indices (price indices) are based, we can split the sample in each year and compute two measures of wage indices (price indices) for each state within each year. For example, if we have 1 million observations in the 2007 American Community Survey, we split the sample into two distinct samples with 500,000 observations each. Within each sample, we can compute a wage measure for each state. The wage measures within each sub-sample, will be measured with error. We can use the growth rates in wages in one half of the samples as an instrument for growth rate in wages in the other half of the samples. We discuss these procedure in detail in the Online Appendix that accompanies the paper. As we show in that appendix, the procedure dramatically corrects the attenuation bias from measurement error in our estimates.

In order to recover unbiased estimates of λ and ϕ via OLS, we must assume $v^{\epsilon}_{kt} = 0$. The assumption that there are no local labor supply shocks cannot generically be true. However, in the Online Appendix, we provide some evidence suggesting that this assumption may be roughly valid during the 2007-2011 period. We show that many potential labor supply shocks highlighted in the literature had large aggregate components but very little cross-state variation. For example, the decline in routine jobs (Jaimovich and Siu (2014), Charles et al (2013, 2015)) was dramatic at the aggregate level during the 2007-2011 period, but occurred in all U.S. states with roughly equal propensity. We show these results in Online Appendix Figure O5. Likewise, some have argued that the expansion of government policies acted like a labor supply shock that discouraged work (Mulligan (2012)). We show that many of these government policies - like the expansion of the Supplemental Nutrition Assistance Program (SNAP) - was large at the aggregate level but had little cross state variation. For example, in Online Appendix Figure O2, we show that SNAP benefits per recipient increased by roughly 30 percent between 2007 and 2011. Because the increase in per recipient benefit occurred at the Federal level, there was statutorily no variation in per recipient benefits across U.S. states during this time period.[29]

One policy that has received considerable attention in its potential to act as a labor supply shock is the differential extension of the duration of unemployment benefits across states during the Great Recession. By law in 2010, weeks of unemployment benefits were tied to the state's unemployment rate. As of 2010, most U.S. states had a duration of unemployment benefits that was close to the maximum of 99 weeks. These states with at least 86 weeks of unemployment benefits represent roughly 90 percent of the U.S. population. However some smaller states, mostly in the Plains region of the U.S., had small employment declines and only had an extension of unemployment benefits from 60-85 weeks.[30] While some smaller states had benefit extensions that were much lower than other states, it should be noted that by 2010 almost all states had unemployment benefit extensions that were close to the maximum. Despite this, we still perform two additional robustness exercises in our OLS specifications to account for the fact that there were policy differences across states during this time period that could discourage labor supply. First, when using our full time period, we exclude any state that had less than 85 weeks of unemployment benefit extensions leaving us with a sample of states that had essentially no variation in unemployment benefits. Given that the excluded states were small in population terms, such exclusion had essentially no effect on our OLS estimates. Additionally, we estimate our key parameters using only data prior to 2010. Prior to 2010, the duration of extended unemployment benefits were the same across all states.

While we try to defend that OLS estimation of the above equation yields unbiased estimates of λ and ϕ using cross state variation during the Great Recession, it is impossible to completely rule out that labor supply shocks are causing some of the variation in state business cycles during this period. To further explore the robustness of our results, we also estimate IV specifications of the above equation. Following the work of many recent papers including Mian and Sufi (2014),

[29] In the Online Appendix, we also show that there was no systematic variation in state labor income tax rates during the 2007-2010 period. Additionally, we show that there was little state variation in Federal programs to help underwater homeowners (like HAMP) that occurred during the 2007-2010 period. The reason that there was little state variation is that take up rates of the program during this time period were very low (with take up rates being essentially zero prior to 2010).

[30] States also had some discretion as to whether they opted into the program. This explains why some states did not have the maximum weeks of unemployment benefits even when their unemployment rate was higher. We discuss these policies and how they varied across states in detail in the Online Appendix.

we use contemporaneous and lagged variation in local house prices as our instruments for local employment and price growth. The argument is that local house price variation during the 2007-2011 period (in our base specification) or during the 2007-2009 period (in our restricted specification) is orthogonal to movements in local labor supply shocks. This seems like a plausible assumption for the 2007-2009 period as state policy changes did not occur prior to 2009. In the Online Appendix, we discuss the IV procedure in detail. We also show that contemporaneous housing price growth strongly predicts contemporaneous employment growth and lagged measures of housing growth predicts price growth.

Before turning to the estimation, it is also worth discussing the no cross-state migration assumption that we have imposed throughout. If individuals were more likely to migrate out of poor performing states and into better performing states, our estimated labor supply elasticities from the state regression may be larger than the aggregate labor supply elasticity. While theoretically interstate migration could be problematic for our results, empirically it is just not the case. Using data from the 2010 American Community Survey, we can compute both the in-migrants and the out-migrants to and from each state. Given this data, we can compute a net-migration rate for each state. As documented by others, we find that the net migration rate was very low during the Great Recession (Yagan 2014). This can be seen from Appendix Figure A3. Both the low level of interstate migration and the fact that it is uncorrelated with employment growth during this period makes us confident that our estimated parameters of our local wage setting curve can be applied to the aggregate.

6.2 Estimates of λ and ϕ

Column 1 of Table 8 shows the estimates of our base OLS specification where we use all data from 2007-2011 and do not include any additional controls. Our base estimates are $b_1 = 0.69$ (standard error = 0.13) and $b_2 = 0.31$ (standard error = 0.08). As noted above, b_1 is λ and b_2 is λ/ϕ. Given our base estimates, the cross sectional variation in prices and wages implies a labor supply elasticity of 2.2. Standard macro models imply a labor supply elasticity of 2 to 4 based on time series variation. The estimates from the cross-section of states are in-line with these macro time series estimates. Our base estimate of λ = 0.69 suggests only a modest amount of wage stickiness. Perfectly flexible wages imply $\lambda = 1$ while perfectly sticky wages imply $\lambda = 0$. In other words, lagged wage growth predicts current wage growth conditional on current employment and price growth, but the effect is much less than one for one. As we discuss below, similar regressions run on aggregate date yield much smaller estimates of λ implying a greater amount of wage stickiness.

Columns 2-3 of Table 8 show a variety of robustness checks for our base estimates. In column 2 we include industry controls. In particular, we include the share of workers in 2007 working in manufacturing occupations or in routine occupations. This allows us to proxy for different amounts of wage stickiness or different potential labor supply shocks that are correlated with industrial mix. In column 3, we use the actual change in retail prices as opposed to scaling the local retail price difference for the fact that retail grocery sector is more tradable than the composite local consumption good. Both including controls for local industry mix and changing the scaling on local retail price variation does not effect our estimates of λ and ϕ in any meaningful way. In columns 4 and 5, we re-estimate our base specification with and without industry controls using only data from 2007-2009 prior to the changes in national policy extending unemployment benefit duration. Again, our estimate λ and ϕ remain 0.73 and 1.9, respectively.[31] Finally, in columns 6 and 7, we show our IV estimates for the 2007-2011 period and the 2007-2009 period where we instrument local employment growth and local price growth with contemporaneous and one lag of local house price growth. Our estimates of λ and λ/ϕ are 0.77 (standard error = 0.13) and 0.76 (standard error = 0.17) implying an estimated Frisch elasticity of 1.0.

Regardless of our specification we estimate labor supply elasticities of between roughly 1.0 and 2.0. More importantly, all of our estimates imply a fair degree of wage flexibility with our estimates of λ ranging from about 0.7 to 0.8. This is consistent with the patterns shown in Figure 4 where local wages moved quite a bit with local economic conditions during the Great Recession. The estimation that wages are fairly flexible is a key insight that is important for the interpretation of our results and has broader implications for the literature. Linking our estimates back to the prior section, it is hard to get aggregate demand shocks to be the primary shock driving economic conditions during the Great Recession if wages are fairly flexible. Put another way, if wages were sticky enough in the aggregate to have demand shocks be the primary driver of aggregate employment declines during the recent recession, we would not observe wages moving as much as they did in the cross section during the same time period.

To show the stark difference between local estimates of wage stickiness and aggregate estimates of wage stickiness, we used aggregate data on prices, nominal wages and employment between 1976 and 2012. This is the same data that we will use in our VAR estimation below. Given the short time series sample, power was in issue. However, across all specifications we explored, estimates of λ using aggregate data ranged from about 0.4 to 0.6. These estimates are well below the estimates of 0.7 to 0.8 using local variation. If aggregate labor supply shocks occur along with shocks that shift labor demand, wages will appear sticky in the aggregate time series. The assumption of no aggregate labor supply shocks is a common one when estimating wage stickiness using aggregate data (see, for example, Christiano et al. (2015b)).

7 An Application to the US Great Recession

The cross sectional facts presented above represent a puzzle. At the aggregate level, nominal wages did not appear to respond much (relative to trend) as aggregate employment fell during the Great Recession. However, exploiting variation across regions, there appears to be a significant negative relationship between nominal wages and local employment. Why did aggregate wages respond so little during the Great Recession while there was a strong relationship at the regional level?

[31] Additionally, we estimated our base specification excluding CA, NV, AZ, and FL. In both cases, our estimates were nearly identical to our base specification in column 1 of Table 3.

One potential explanation is that a series of shocks hit the aggregate economy during this period - some putting downward pressure on prices and wages and others putting upward pressure on prices and wages. If some of those shocks had only aggregate effects they would be differenced out in the cross region variation during the recession. Our econometric procedure allows us to quantify the relative magnitudes of these shocks and to assess their contributions to the behavior of prices, wages and employment during this period at both the aggregate and local level.

7.1 Findings in the aggregate

We follow the procedure described in Section 5.1. We first estimate the VAR with two lags in aggregate employment growth, price growth and nominal wage growth via OLS equation by equation using annual data from 1976 to 2012. From the reduced form errors U we obtain sample estimators of the covariance matrix $\hat{V} = \frac{UU'}{Years - \#Variables \# Lags}$.

The aggregate variables we construct are comparable to our regional measures. Given that our cross-sectional equations are estimated using annual data, we analogously define our aggregate data at annual frequencies. We use data from the CPI-U to create our measure of aggregate prices. Specifically, we take log-change in the CPI's between the second quarter of year t and $t-1$ for our measure of p_t. For n, we use BLS data on the aggregate employment to population rate of all males 25-54. We choose this age range so as to abstract from the downward trend in employment rates due to the aging of the population over the last 30 years.[32] Finally, we use data from the Current Population Survey (CPS), discussed above, to construct our aggregate composition adjusted wage measure. As with the CPI, we take the log-change in this wage measure between t and $t-1$ for our measure of w_t. For all data, we use years between 1976 and 2012.

Figures 5, 6, and 7 report the impulse response of aggregate employment, nominal wages and price growth to each of the shocks when we use our benchmark estimates for λ and ϕ reported in column 1 of Table 8 ($\lambda = 0.69$ and $\phi = 2.2$). Figure 5 shows their behavior after an initial discount rate/monetary (γ) shock of the same magnitude and sign as in 2008. Qualitatively, after a negative discount rate shock both prices and employment fall sharply relative to trend while real wages decline slightly relative to trend. These results are identical to the theoretical predictions shown in Table 6. Figure 6 shows the impulse responses to a 2008 productivity/markup (z) shock. Prices relative to trend increase on impact while employment growth falls sharply. Nominal wages, however, only decline slightly (with a lag). Again, these predictions match the predictions of our simple theoretical model shown in Table 6. While both negative γ and z shocks reduce employment, the γ shock puts downward pressure on prices while the z shock puts upward pressure on prices. Figure 7 shows the impulse response of employment, prices, and nominal wages to the leisure shock. Upon impact, the leisure shock reduces employment and prices while it increases wages.

[32] We detrend all data when estimating the VAR. Specifically, we allow for a linear trend in the employment to population ratio between 1978 and 2007. For the price inflation rate and the nominal wage inflation rate, we use an HP filter (with a smoothing parameter of 100). Given that we detrend the data, our results are essentially unchanged when we use the employment to population ratio for all individuals as opposed to using it just for prime age males.

Again, these predictions match the predictions from the simple model. It is the leisure shock that is putting upward pressure on nominal wages in the aggregate.

We turn now to the cumulative response of each individual variable when we feed the VAR with the sequence of shocks between 2008 and 2012 one at a time.[33] The nature of the counterfactuals aims at quantifying the contribution of each shock during the Great Recession in explaining the behavior of the aggregate US economy. Figure 9 presents the counterfactual employment response. During the Great Recession employment fell in the US by more than 4 percent between 2008-2009 (relative to trend) and remained at the low level thereafter. The counterfactual exercise shows that the productivity/markup and discount rate shocks contributed about the same amount to the initial decline during the 2008-2009 period (each explaining roughly 40 percent of the aggregate employment decline). However, the discount rate/monetary does not explain any of the persistence in the employment decline post 2009. Instead, it is the productivity/mark-up shock that explains most of the sluggish response of employment post 2009. The leisure shock contributed only a modest amount to the observed employment decline both on impact or over the longer 2008-2012 period.

Figures 10 and 11 provide insight into the shock decomposition as well as to helping understand the "missing deflation puzzle"and "missing nominal wage decline puzzle".[34] Figure 10 shows the counterfactual price response to each of the shocks. With respect to the data, aggregate prices fell relative to trend between 2008 and 2009. However, prices quickly stabilized relative to trend. This is the sense that there was "missing deflation". Despite the weak employment situation post 2009, prices were growing relative to trend. Both the discount rate/monetary shock and the leisure shock put downward pressure on aggregate prices. However, the productivity/markup shock put upward pressure on aggregate prices. The counterfactual analysis shows that if the economy was only hit with the productivity/mark-up shock, prices would have be rising (by upwards of 2 percent) relative to trend during the Great Recession. According to our procedure, it is this countervailing productivity/markup shock that arises as the explanation for the missing deflation puzzle - particularly post 2009. This finding is consistent with the results of Christiano et al. (2015a).

Figure 11 shows the cumulative nominal wage response to each of the shocks. Again, the figure shows the missing nominal wage puzzle during the Great Recession. Throughout the early part of the recession and the entire recession, nominal wage growth was close to zero (relative to trend). However, our model shows - like conventional wisdom - if the economy only experienced the discount rate/monetary shock, nominal wages should have fallen by roughly 2.5 percent relative to trend by 2010 and remained well below trend in 2012. The leisure and productivity/markup shocks, however, are necessary to explain why nominal wages did not fall during the Great Recession. Some of the exit in the labor force due to the leisure shock put upward pressure on nominal wages. Likewise, the productivity shock also put upward pressure on nominal wages. In summary, our

[33]For the interested reader, the actual realizations of the shocks we estimate can be seen in Figure 8.

[34]The robust growth in nominal wages and consumer prices during the recession is viewed as a puzzle for those that believe that the lack of aggregate demand was the primary cause of the Great Recession. For discussions of the "missing deflation", see Hall (2011), Ball and Mazumder (2011), Stock and Watson (2012), and Del Negro et al. (2015).

methodology suggests that demand shocks cannot be solely responsible for the employment decline during the Great Recession. If demand shocks were solely responsible, price inflation would have been lower - particularly post 2009 and nominal wages would have fallen. Our method suggests that a combination of productivity shocks and labor supply shocks are needed to explain the missing deflation and nominal wage declines. Moreover, our methodology allows us to quantify the relative importance of each shock in explaining the aggregate movement in wages, prices, and employment during the Great Recession.

7.2 Robustness to Estimated Parameters λ and ϕ

How do our estimated parameters affect our employment, price and wage counterfactuals? In Table 9, we report the contribution of each shock to the explanation of aggregate employment declines implied by different combinations of $\{\phi, \lambda\}$. We do this for both the initial years of the recession (2008 to 2009) and over the longer period encompassing the recovery (2008 to 2012). The table shows that the qualitative conclusions of the previous section still hold for the range of $\{\phi, \lambda\}$ estimates using alternative specifications that we report in Table 8. These go from roughly 0.7 to 0.8 for λ and from roughly 1.0 to 2.5 for ϕ. When reading Table 9, each cell shows the decomposition of how much of the employment change during the time period can be attributed to the discount rate/monetary shock (γ) and how much can be explained by the productivity mark up shock (z). The sum of all three shocks should sum to 100 percent. So, the difference between the sum of the γ and z contributions and 100 percent is attributed to the labor supply shock (ε).

Table 9 offers several further results worth discussing. First, we observe that the relative importance of the labor supply shock vis a vis the discount rate and productivity/markup shocks combined is governed by the Frisch labor supply elasticity ϕ. We estimate a relatively large elasticity; in the range of that used to calibrate standard macro models.[35] However, suppose we used a much lower elasticity $\phi = 0.5$ instead which is in line with some microeconomic estimates in the literature. In this case, the labor supply shock would account for a much larger fraction of the employment decline in the Great Recession. This comes out of the VAR estimation. However, the intuition for this result is straightforward and in line with the simple theory we wrote down in Section 4. If labor supply is fairly elastic, large movements in employment can be rationalized without the need of large labor supply shocks given the relatively small movements in real wages in the data.

The intuition for the decomposition between discount rate and productivity/markup shocks is more subtle but also in line with the simple theory we wrote down. We find that the degree of wage flexibility λ affects the relative importance of one vis a vis the other within the remaining

[35]This result may be of independent interest to the reader familiar with the macro v. micro labor supply elasticities (see Chetty, Guren, Manoli, and Weber (2011)). Using cross-sectional data (same as in most of the micro labor-supply elasticity literature) we arrive at an estimate similar to the macro elasticity (estimated from aggregate time series data). We believe this is because the regional variation in employment rates that we use to estimate this elasticity only incorporates the extensive margin adjustment in the labor supply; which is the same margin that is most important in accounting for aggregate fluctuations in total hours over the business cycle.

unexplained part by the leisure shock. For example, suppose we increased the degree of wage flexibility λ.[36] Then the productivity/markup shock would account for a much larger fraction of the employment decline in the Great Recession. Theoretically, it is clear that when λ is large, the discount rate shocks do not matter much for the determination of employment. To see this, consider the extreme case where wages are perfectly flexible and the demand shock is only composed of the monetary shock. Then the equilibrium in the simple theoretical model satisfies monetary neutrality. We formalized this point in Section 4.6 when we derived the model's implied elasticity of aggregate employment to a discount rate shock. Conversely, when wages are very rigid ($\lambda = 0.1$), our procedure is suggesting that demand shocks can explain essentially all of the decline in the early part of the recession and much of the persistence in employment declines during the 2008-2012 period. Given that all of our estimates of λ are around 0.7, the demand shock cannot have been the sole cause of the recession because if it did nominal wages would have fallen and prices would have declined relative to trend after 2009.

7.3 Findings at the regional level

We follow the procedure in Section 5.2. We first estimate the VAR with two lags in state-level non-tradable employment, price and nominal wage growth via OLS equation by equation.[37] All variables are expressed in log-deviations from their average weighted by population in 2006. We pool all data between 2006 and 2011, and estimate common autoregressive coefficients and reduced form errors covariance matrix for all states. In our benchmark specification, we set $\{\phi, \lambda\}$ equal to our estimates from Section 6.2. We set $\theta = 0.55$ to match the labor share in the manufacturing sector in the US and $\frac{N^y}{N} = 0.85$ to match the share of total employment in the service sector plus self-employed/family workers as reported in the BLS.

Table 10 summarizes the contribution of the discount rate shock and the combined productivity/markup shocks to non-tradable employment, wages and prices. We define shock j's contribution to the change in variable y between 2007 and 2010 as

$$\xi_y^j = \frac{\frac{1}{\sum_k \omega_k (\Delta^j \tilde{y}_k - \Delta \tilde{y}_k)^2}}{\sum_j \frac{1}{\sum_k \omega_k (\Delta^j \tilde{y}_k - \Delta \tilde{y}_k)^2}}$$

where ω_k are population weights in 2006; $\Delta \tilde{y}_k$ is the change in variable y in state k between 2007 and 2010; and $\Delta^j \tilde{y}_k$ is the counterfactual change if only shock j would have occurred. Note that ξ_y^j is always in $[0, 1]$ and increases when the counterfactual and actual changes in variable y are close to each other.

For our benchmark $(\lambda, \phi) = (0.7, 2)$ we find that the discount rate shock contributed 79 percent to

[36] It is worth mentioning that for large values of λ and small values of ϕ, the results in Table 9 become rather sensitive to small variations in parameters. This is because our shock identification procedure needs a certain degree of wage rigidities, as explained in the Section 5.1. For example, for values of λ around 0.95 it is not possible to identify the productivity/markup and discount rate shocks.

[37] We define non-tradable employment in a state as the employment rate in the service and retail sectors combined.

the change in non-tradable employment between 2007 and 2010 across all states; 41 percent to local price changes and 24 percent to local wage changes. We obtain similar numbers for $(\lambda, \phi) = (0.5, 1)$. We conclude that the discount rate shock was the main driver of regional variation in non-tradable employment.

Table 11 summarizes a second set of results from the regional counterfactuals. We characterize the joint distribution of cumulative growth rates between 2007 and 2010 for each variable across states with two statistics: the variance and the correlation with each other. We compare the actual statistics from our data in Section 3 with the counterfactuals obtained when simulating one shock at a time. Again, we find that the discount rate shock alone can generate 98 percent of the cross-state variance of non-tradable employment growth; 69 percent of the price growth variance; and 47 percent of the nominal wage growth variance. Moreover, it reproduces the right sign for the cross-state correlations of price growth and non-tradable employment growth; nominal wage growth and non-tradable employment growth; and nominal wage growth and price growth. Although, quantitatively, it generates a larger correlation between prices and non-tradable employment than in the data. Both productivity/markup shocks combined can explain only 35 percent of the non-tradable employment growth variance across states. They do as good a job as the discount rate in explaining 85 percent of the variation in price growth and a worse job in explaining 26 percent of nominal wage growth variance. However, they imply negative correlations between price/wage growth and non-tradable employment growth. The opposite is observed in the data. We conclude that the discount rate shock alone does a fairly good job in reproducing the regional patterns that we documented in Section 3. To improve the estimated fit, nontradable and tradable productivity/markup shocks are needed. However, their contribution to aggregate employment changes are small.

7.4 Robustness to Alternate Identifying Assumptions

The above results are based on the particular functional form of our wage setting equation:

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

This particular wage setting equation is based on our assumption of preferences as well as our assumption on the reset wage. In our baseline model, we assumed a utility function with no income effects on labor supply. Likewise, our wage setting equation assumed that no forward looking behavior is used when wages are reset. Both of these assumptions were made for tractability. In this sub-section, we explore the robustness of our results to relaxing both of these assumptions.

In Appendix B1, we derive the aggregate and local wage setting equations under a broad set of utility functions where consumption and leisure are non-separable. This class of utility functions allows for arbitrarily large income and substitution effects. As we show in the appendix, the use of local consumption data allows us to estimate the extent of wage stickiness as well as to estimate the parameters that encompass both the income and substitution effects on labor supply. In particular, we can estimate the following equation using local data:

$$\pi^w_{kt} = \tilde{b}_t + \tilde{b}_1 \pi_{kt} + \tilde{b}_2 (n_{kt} - n_{kt-1} + \tilde{b}_3 \pi^w_{kt-1} + \tilde{b}_4 (c_{kt} - c_{kt-1}) + \Psi D_t + \Gamma X_k + e_{kt}.$$

This equation is identical to our estimating equation above aside from the addition of local consumption growth. As outlined in Appendix B1, the coefficients $\tilde{b}_1$ and $\tilde{b}_3$ should sum to 1 even under the broader preference specification. We impose this restriction when estimating the modified equation. For our measure of local real consumption growth, we use the change in real retail expenditures at the state level computed within the Nielsen sample. To get real expenditures, we deflate nominal expenditures by our measure of local prices.[38] For our base specification, our estimates of $\tilde{b}_1$, $\tilde{b}_2$, and $\tilde{b}_4$ are, respectively, 0.72 (standard error = 0.12), 0.25 (standard error = 0.08), and 0.16 (standard error = 0.06). The coefficient on real consumption growth ($\tilde{b}_4$) is positive and significant saying that there is an estimated income effect on labor supply. Controlling for the income effect, our estimate of λ ($\tilde{b}_1$) is slightly higher than our base specification without allowing for an income effect.

For the aggregate wage setting equation, we can substitute out consumption growth using the model definition ($c_t = w_t + n_t - p_t$). Substituting out consumption growth, Appendix B1 shows that the aggregate wage setting equation still takes the following form:

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

where ω is a parameter that represents the strength of the income effect on labor supply (and maps directly to $\tilde{b}_4$ from the above local labor supply regression, see equation (6) in Appendix B1). Aside from the coefficient scaling the aggregate labor supply shock, this equation is identical to the identification restriction we imposed when estimating the aggregate VAR. The only difference is that there is no longer a one-to-one mapping between λ and λ/ϕ in the above aggregate restriction on the VAR and the reduced-form parameters $\tilde{b}_1$ and $\tilde{b}_2$ from the local wage setting equation. However, as shown in Appendix B1, there is still a specific mapping between the parameters we estimate from the local regression ($\tilde{b}_1$, $\tilde{b}_2$, and $\tilde{b}_4$) and the aggregate parameters we need to identify the VAR (λ and ϕ). With the correctly specified λ and ϕ, we can just use the matrix in Table 9 to read off the decomposition of shocks during the Great Recession. While λ and ϕ are no longer structural parameters (instead being combinations of structural parameters), knowing them still fully identifies the aggregate VAR. Using our estimates of $\tilde{b}_1$, $\tilde{b}_2$, and $\tilde{b}_4$ and the procedure developed in Appendix B1, we find that our new estimates of λ and ϕ which allow for income effects on labor supply to be 0.68 and 2.0, respectively. These parameters are nearly identical to our base specification without income effects. Having local consumption data allows us to control for income effects on labor supply in our local wage setting equations. The take away from this robustness exercise is that abstracting from preferences that allow for an income effect on labor supply is not biasing our

[38]This measure of real expenditures is (1) highly correlated with measures of local employment and (2) highly correlated with the BEA's recent state level personal expenditures measure. Our results are similar if we use the BEA's local consumption measure. However, we prefer our measure given that much of the BEA's local consumption measure is imputed (where the imputation uses local employment measures).

decomposition of the shocks driving aggregate employment declines during the Great Recession in any meaningful way.

In Appendix B2, we re-specify our wage setting equation allowing for forward looking behavior when wages are reset. In this Appendix, we show that ignoring forward looking wage setting behavior biases up our estimates of wage flexibility. That is, the amount of wage flexibility that we estimate from our local equation is too large relative to the true amount of wage flexibility. We show that the amount of the bias is dependent on two parameters: (1) the extent to which firms put weight on forward looking behavior when setting wages (which we call κ in Appendix B) and (2) the underlying persistence process of local wages (which we call $\bar{\rho}_w$). Imposing that at the aggregate level the monetary authority wants to stabilize expected nominal wage growth, we can quantify the extent to which our estimates of λ from the local regressions are biased upwards as a function of these two parameters. Under a range of plausible parameter estimates for κ and $\bar{\rho}_w$, we show that the bias is quite small. For our base specification, where our estimated $\lambda = 0.69$, we show that $\kappa(1 - \bar{\rho}_w)$ must exceed 2.5 for the true λ to be lower than 0.4. This is an order of magnitude larger than any plausible parametrization for either κ or $\bar{\rho}_w$. Moreover, as seen in Table 9, our estimates of the role of demand shocks in explaining employment decline during the Great Recession are broadly similar for values of λ between 0.4 and 0.69. These results suggest that our abstraction from including expectations in our wage setting equation is not quantitatively altering the paper's conclusions.

8 Conclusion

Regional business cycles during the Great Recession in the US were strikingly different than their aggregate counterpart. This is the cornerstone observation on which we built this paper. Yet, the aggregate US economy is just a collection of these regions connected by trade of goods and assets. We argued that their aggregation cannot be arbitrary and that regional business cycle patterns have interesting implications for aggregate business cycles: regions and their connections place restrictions on the structure of the economy and, thus, that structure helps to identify the underlying shocks driving aggregate fluctuations.

In this paper, we show that making inferences about the aggregate economy using regional variation is complicated by two issues. First, the local elasticity to a given shock may differ from the aggregate elasticity to the same shock because of general equilibrium effects. Second, the type of shocks driving most of the regional variation may be different than the shocks driving most of the aggregate variation. We document that both of these issues are quantitatively important using local and aggregate data for the U.S. during the Great Recession.

There are a few key takeaways from the paper. First, the relationship between prices, wages, and employment in the aggregate time series during the 2006-2011 period are very different than the cross-region relationship between prices, wages, and employment during the same time period. For example, while aggregate wages appeared to be very sticky despite aggregate employment falling sharply, both nominal and real wages co-varied strongly with local employment growth in cross-section of U.S. states. Both the documentation of the cross-region facts and the creation of the underlying local price and wage indices are the first innovations of the paper.

The second take-away is that wages are only modestly sticky using cross-region data. The amount of wage stickiness is often a key parameter in many macro models. Despite its importance, there are not many estimates of the frequency with which wages adjust (particularly relative to estimates of price adjustments). We develop a procedure to estimate the amount of wage stickiness using cross-region variation. The wage stickiness parameter is key to our empirical methodology to estimate the underlying shocks and elasticities at the aggregate and local level. The fact that we estimate that wages are only modestly sticky limits the importance of demand shocks at the aggregate level in explaining the Great Recession. If wages are only modestly sticky, aggregate demand shocks should have resulted in falling wages. This is not something that was observed in the aggregate time series. Regardless of the use of this parameter in our empirical work, we think that our estimate of wage stickiness could be of independent interest to researchers.

The third takeaway from this paper is the development of a new econometric procedure that allows us to estimate both the aggregate and local shocks as well as aggregate and local elasticities to a given shock. This procedure is a hybrid method that merges very few restrictions imposed by a theoretical model with aggregate and cross-sectional data when estimating a VAR and identifying the corresponding shocks. We show that if our assumption about the form of the aggregate wage setting equation is true, a parameterized version of that equation is enough to identify the aggregate and local VARs without any additional assumptions (aside from the usual orthogonalization conditions). We view this as a contribution to the growing literature that uses model-based structure to estimate VARs.

Finally, the fourth takeaway is the most important for the goals of the paper. Using our various empirical components, we show that a combination of both "demand" and "supply" shocks are necessary to account for the joint dynamics of aggregate prices, wages and employment during the 2007-2012 period within the U.S.. In contrast with the aggregate results, we find that discount rate shocks explain most of the observed employment, price and wage dynamics across states. These results suggest that only using cross-region variation to explain aggregate fluctuations is insufficient when some shocks do not have a substantive regional component. The reason that aggregate prices and wages are not falling is not because wages and prices are completely sticky. The reason aggregate prices and wages are not falling is that the series of shocks experienced by the aggregate economy are such that some shocks are putting downward pressure on prices and wages (discount rate shocks) while other shocks are putting upward pressure on prices and wages (productivity/mark-up and leisure shocks). In the cross-section, however, the discount rate shocks are causing prices, wages and employment to covary positively. Lastly, we quantify that the local employment elasticity to a local discount rate shock is substantially larger than the aggregate employment elasticity to a similarly sized aggregate discount rate shock. These results suggest that even when the aggregate and regional shocks are the same, it is hard to draw inferences about the aggregate economy using regional variation. Collectively these results suggest that researchers should be cautious when extrapolating variation across regions to make statements about aggregate dynamics.

References

An, Sungbae and Frank Schorfheide, 2007. “Bayesian Analysis of DGSE Models”, *Econometric Reviews*, 24(2-4), 113-73.

Autor, David, David Dorn, and Gordon Hanson, 2013. "The China Syndrome: Local Labor Market Effects of Import Competitions in the United States", *American Economic Review*, 103(6), 2121-2168.

Ball, Laurence and Sandeep Mazumder, 2011. “Inflation Dynamics and the Great Recession”, *Brookings Papers on Economic Activity*, 42(1), 337-405.

Beraja, Martin, 2015. “A Semi-Structural Methodology for Policy Counterfactuals and an Application to Fiscal Unions”, *University of Chicago Working Paper*.

Berger, David and Joe Vavra, 2015. “Dynamics of the U.S. Price Distribution”, *Working Paper*.

Blanchard, Olivier and Lawrence Katz, 1991. "Regional Evolutions", *Brookings Papers on Economic Activity*, 1992:1, 1-75.

Cage, Robert, John Greenless and Patrick Jackman, 2003. “Introducing the Chained Consumer Price Index,” U.S. Bureau of Labor Statistics.

Carriero, Andrea and Raffaella Giacomini, 2011. “How Useful are No-Arbitrage Restrictions for Forecasting the Term Structure of Interest Rates”, *Journal of Econometrics*, 164(1), 21-34.

Charles, Kerwin, Erik Hurst, and Matthew Notowidigdo, 2015. "Masking and Employment Changes During the 2000s: Housing Booms and Manufacturing Decline”, *University of Chicago Working Paper*.

Charles, Kerwin, Erik Hurst, and Matthew Notowidigdo, 2013. "Manufacturing Decline, Housing Booms and Non-Employment", *NBER Working Paper #18989*.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber, 2011. “Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins”, *American Economic Review Papers and Proceedings*, 471-5.

Christiano, Lawrence, Martin Eichenbaum, and Mathias Trabandt, 2015a. “Understanding the Great Recession” *American Economic Journal - Macro*, January, 110-67.

Christiano, Lawrence, Martin Eichenbaum, and Mathias Trabandt, 2015b. "Unemployment and Business Cycles", *Econometrica (forthcoming)*.

Coibion, Olivier, Yuriy Gorodnichenko, and Gee Hee Hong, 2015. "The Cyclicality of Sales, Regular and Effective Prices: Business Cycle and Policy Implications", *American Economic Review*, 105(3), 993-1029.

Del Negro, Marco and Frank Schorfheide, 2004. "Priors from General Equilibrium Models for VARs", *International Economic Review*, 45, 643-73.

Del Negro, Marco, Marc Giannoni, and Frank Schorfheide, 2015. "Inflation in the Great Recession and New Keynesian Models", *American Economic Journal: Macroeconomics, 7(1): 168-96*.

Diewert, Erwin W., 1976. "Exact and Superlative Index Numbers," *Journal of Econometrics* 46, 115–45.

Diewert, Erwin W., Lorraine Ivancic and Kevin J. Fox, 2011. "Scanner data, time aggregation and the construction of price indexes," *Journal of Econometrics* 161, 24-35.

Eggertsson, Gauti and Paul Krugman, 2012. "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach", *Quarterly Journal of Economics*, 127(3), 1469-513.

Fitzgerald, Terry and Juan Pablo Nicolini, 2014. "Is there a Stable Relationship Between Unemployment and Future Inflation? Evidence from U.S. Cities", *Federal Reserve Bank of Minneapolis Working Paper 713*.

Gali, Jordi, 2011. "Monetary Policy and Unemployment", Handbook of Monetary Economics, vol. 3A, Elsevier B.V., 487-546

Giroud, Xavier and Holger Mueller, 2015. "Firm Leverage and Unemployment during the Great Recession", *Working Paper*

Gilchrist, Simon, Raphael Schoenle, Jae Sim, and Egon Zakrajsek, 2014. "Inflation Dynamics During the Financial Crisis", working paper.

Guerrieri, Veronica and Guido Lorenzoni, 2011. "Credit Crises, Precautionary Savings and the Liquidity Trap", University of Chicago Working Paper.

Hagerdon, Marcus, Fatih Karahan, Iourii Manovskii, and Kurt Mitman, 2015. "Unemployment Benefits and Unemployment in the Great Recession: The Role of Macro Effects" NBER Working Paper 19499.

Hall, Robert, 2011. "The Long Slump", *American Economic Review*, 101(2), 431-469.

Hall, Robert and Paul Milgrom, 2008. "The Limited Influence of Unemployment on the Wage Bargain", *American Economic Review*, 98(4), 1653-74.

Baumeister, Christiane and James Hamilton (2015). "Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information", *Econometrica*, 83, 1963-99.

Jaimovich, Nir and Henry Siu, 2014. "The Trend is the Cycle: Job Polarization and Jobless Recoveries", Duke University Working Paper.

Kaplan, Greg and Guido Menzio, 2015. "The Morphology of Price Dispersion", *International Economic Review, forthcoming*.

Mehrotra, Neil and Dmitriy Sergeyev, 2015. "Financial Shocks and Job Flows", Brown University Working Paper.

Melser, Daniel, 2011. "Constructing High Frequency Price Indexes Using Scanner Data," Working Paper, 2011

Mian, Atif and Amir Sufi, 2014. "What Explains the High Unemployment? The Aggregate Demand Channel", *Econometrica* (forthcoming)

Mian, Atif, Kamalesh Rao, and Amir Sufi, 2013. "Household Balance Sheets, Consumption and The Economic Slump", *Quarterly Journal of Economics*, 1687-1726.

Mian, Atif and Amir Sufi, 2011. "House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis", *American Economic Review*, 101(4), 2132-56.

Midrigan, Virgiliu and Thomas Philippon, 2011. "Household Leverage and the Recession", *NYU Working Paper*.

Mondragron, John, 2015. "Household Credit and Employment in the Great Recession", *University of California at Berkeley Working Paper.*

Mulligan, Casey, 2012. "The Redistribution Recession: How Labor Market Distortions Contracted the Economy*", Oxford University Press*.

Nakamura, E. and J. Steinsson (2014). "Fiscal Stimulus in a Monetary Union: Evidence from US Regions," *The American Economic Review*, 104, 753–792.

Ruggles, Steven, J. Trent Alexander, Katie Genadek, Ronald Goeken, Matthew B. Schroeder, and Mathew Sobek, 2010. *Integrated Public Use Microdata Series: Version 5.0*. Minneapolis: University of Minnesota.

Schmitt-Grohé, S. and M. Uribe, 2003: "Closing small open economy models," *Journal of International Economics*, 61, 163–185.

Schorfheide, Frank, 2000. "Loss Function-based Evaluation of DSGE Models", *Journal of Applied Econometrics*, 15(6), 645-70.

Shimer, Robert, 2010. *Labor Markets and Business Cycles*, Princeton University Pres.

Stock, James and Mark Watson, 2012. "Disentangling the Channels of the 2007-2009 Recession", *Brookings Papers on Economic Activity, Spring (2012).*

Stroebel, Johannes and Joseph Vavra, 2014. "House Prices, Local Demand, and Retail Prices", *NBER Working Paper, 20710*.

Vavra, Joseph, 2014. "Inflation Dynamics and Time Varying Volatility: New Evidence and an Ss Interpretation", *Quarterly Journal of Economics*, 129(1).

Yagan, Danny, 2014. "Moving to Opportunity? Migratory Insurance Over the Great Recession", Working Paper.

A Proof of Lemma 1 and 2

The following equations characterize the log-linearized equilibrium

$$
\begin{aligned}
w^r_{kt} &= \lambda(\epsilon_{kt} + \frac{1}{\phi} n_{kt}) + (1-\lambda)(w^r_{kt-1} - \pi_{kt}) \\
w^r_{kt} &= -(1-(\alpha+\theta\beta))n^y_{kt} - \beta(1-\theta)(n^x_{kt} - n^y_{kt}) + z^y_{kt} + \beta z^x_{kt} \\
0 &= \mathbb{E}_t\left(mu_{kt+1} - mu_{kt} - \pi_{kt+1} - \gamma_{kt+1} - \Phi_0(c_{kt} - c_t) + \varphi_p \pi_t + \varphi_y(c_t - c^*_t)\right) \\
mu_{kt+1} &= -\frac{\sigma}{C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}} \left(C c_{kt+1} - N^{\frac{1+\phi}{\phi}} (\frac{1+\phi}{\phi} \epsilon_{kt+1} + n_{kt+1}) \right) \\
N n_{kt} &= N^x n^x_{kt} + N^y n^y_{kt} \\
c_{kt} &= w^r_{kt} + n^y_{kt} \\
b_{kt} &= (1+r)(b_{kt-1} + i_t) + \frac{X}{B}(z^x_{kt} + \theta n^x_{kt} - x_{kt}) - r\tau_t \\
0 &= z^x_{kt} - (w_{kt} - q_t) - (1-\theta) n^x_{kt} \\
x_{kt} &= n^y_{kt} + (w_{kt} - q_t) \\
\sum_k x_{kt} &= \sum_k (z^x_{kt} + \theta n^x_{kt})
\end{aligned}
$$

From the last 3 equations, after adding up, it holds that $n^x_t = n^y_t$. Then the aggregate log-linearized equilibrium evolution of $\{\pi^w_t, w^r_t, n_t\}$ is characterized by

$$
\begin{aligned}
0 &= \mathbb{E}_t(mu_{t+1} - mu_t - \pi_{t+1} - \gamma_{t+1}) + \varphi_p \mathbb{E}_t[\pi_{t+1}] + \varphi_y(w^r_t + n_t) \\
\pi^w_t &= \frac{\lambda}{1-\lambda}(\epsilon_t + \frac{1}{\phi} n_t - w^r_t) \\
w^r_t &= -(1-(\alpha+\theta\beta))n_t + z_t \\
mu_{t+1} &\equiv -\frac{\sigma}{C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}} \left(C(w^r_{t+1} + n_{t+1}) - N^{\frac{1+\phi}{\phi}} (\frac{1+\phi}{\phi} \epsilon_{t+1} + n_{t+1}) \right) \\
\pi_{t+1} &\equiv \pi^w_{t+1} - (w^r_{t+1} - w^r_t)
\end{aligned}
$$

which is equivalent to the system of equations characterizing the log-linearized equilibrium in a representative agent economy with a production technology that utilizes labor alone with an elasticity of $\alpha + \theta\beta$, no endogenous discounting and only 3 exogenous processes $\{z_t, \epsilon_t, \gamma_t\}$. The top equation is the aggregate Euler equation. The second equation is the aggregate wage setting equation. The third equation is effectively the aggregate labor demand curve.

To prove Lemma 2, just take log-deviations from the aggregate in the original system. This results in the system characterizing the evolution of $\{\tilde{p}_t, \tilde{w}_t, \tilde{n}_t^y, \tilde{n}_t^x\}$ for given $\{\tilde{z}_t^y, \tilde{z}_t^x, \tilde{\gamma}_t, \tilde{\epsilon}_t\}$,

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

This system is identical to the original where we have set $i_t = q_t = 0$ and dropped the market clearing condition in the intermediate goods market.

Finally, the recursive solution to the aggregate system is $[\pi_t^w, w_t^r, n_t]' = Q[\gamma_t, z_t, \epsilon_t]'$ and the solution to the regional system is $[\tilde{w}_t, \tilde{b}_t, \tilde{p}_t, \tilde{n}_t^y]' = \tilde{P}[\tilde{w}_{t-1}, \tilde{b}_{t-1}] + \tilde{Q}[\tilde{\gamma}_t, \tilde{z}_t^y, \tilde{z}_t^x, \epsilon_t]'$. The matrices characterizing the solution can be found using the method of undetermined coefficients.

B Alternative wage setting specifications

B.1 Preferences with wealth effects in labor supply

In our benchmark specification for the wage setting equation we assumed that the marginal rate of substitution between consumption and hours worked is independent of consumption (as is the case with GHH preferences). In this section we explore the consequences of moving away from this assumption for our econometric procedure in Section 5.1

For a general set of preferences represented by $u(c,n)$, we can write the marginal rate of substitution in log-deviations from steady state as,

$$
\begin{aligned}
mrs_{kt} &= \left(\frac{u_{cn}c}{u_n} - \frac{u_{cc}c}{u_c} \right) c_{kt} + \left(\frac{u_{nn}n}{u_n} - \frac{u_{nc}}{u_n} c \right) n_{kt} \\
&\equiv \omega c_{kt} + \left(\omega + \frac{1}{\hat{\phi}} \right) n_{kt}
\end{aligned}
$$

which nest the special case with no wealth effects ($\omega = 0$) and so we obtain the marginal rate of substitution from our benchmark specification. The aggregate and state level wage setting equations become,

$$
\begin{aligned}
\pi_t^w &= \hat{\lambda}(\pi_t + \epsilon_t - \epsilon_{t-1} + (\frac{1}{\hat{\phi}} + \omega)(n_t - n_{t-1}) + \omega(c_t - c_{t-1})) + (1-\hat{\lambda})\pi_{t-1}^w \\
\tilde{w}_{kt} &= \hat{\lambda}(\tilde{p}_{kt} + \epsilon_{kt} + (\frac{1}{\hat{\phi}} + \omega)\tilde{n}_{kt} + \omega\tilde{c}_{kt}) + (1-\hat{\lambda})\tilde{w}_{kt-1}
\end{aligned}
$$

Replacing aggregate consumption with the model implied $w_t + n_t - p_t$ we obtain

$$\pi_t^w = \lambda\pi_t + \frac{\lambda}{\phi}(n_t - n_{t-1}) + (1-\lambda)\pi_{t-1}^w + \frac{\lambda}{1-\omega}(\epsilon_t - \epsilon_{t-1})$$

where $\lambda \equiv \frac{\hat{\lambda}(1-\omega)}{1-\lambda\omega}$ and $\frac{1}{\phi} \equiv \frac{\frac{1}{\phi}+2\omega}{1-\omega}$. Also, we can re-write the state level equation as

$$\tilde{w}_{kt} = \lambda\tilde{p}_{kt} + \frac{\lambda}{\phi}\tilde{n}_{kt} + (1-\lambda)\tilde{w}_{kt-1} + \frac{\lambda\omega}{1-\omega}(\tilde{p}_{kt} + \tilde{c}_{kt} - (\tilde{w}_{kt} + \tilde{n}_{kt})) + \frac{\lambda}{1-\omega}\tilde{\epsilon}_{kt} \tag{6}$$

Since state level economies are open economies, in general, the term $\tilde{p}_{kt} + \tilde{c}_{kt} - (\tilde{w}_{kt} + \tilde{n}_{kt})$ will be different from zero. By omitting it in our cross-sectional regressions we could be obtaining biased estimates of λ, ϕ.

B.2 Forward looking wages

In our benchmark specification for the wage setting equation we assumed that there was no forward looking term in the target wage. In this section we explore the consequences of having a forward looking component in the wage setting equation for our econometric procedure in Section 5.1

In particular, consider the aggregate and state level wage setting equations

$$\begin{aligned}
{\pi^w}_t &= \lambda(\pi_t + \epsilon_t - \epsilon_{t-1} + \frac{1}{\phi}(n_t - n_{t-1})) + \lambda\kappa\mathbb{E}_t[\pi_{t+1}^w] + (1-\lambda){\pi^w}_{t-1} \\
\tilde{w}_{kt} &= \lambda(\tilde{p}_{kt} + \tilde{\epsilon}_{kt} + \frac{1}{\phi}\tilde{n}_{kt}) + \lambda\kappa\mathbb{E}_t[\tilde{w}_{kt+1} - \tilde{w}_{kt}] + (1-\lambda)\tilde{w}_{kt-1}
\end{aligned}$$

where κ parametrizes the importance of the forward looking term.

Also, lets consider the case where local wages follow an AR(1) process in equilibrium with coefficient $\tilde{\rho}_w$ and aggregate expected wage inflation is zero. Our model from Section 4 , would imply this, for instance, when $\theta \to 1$ so that $\tilde{w}_{kt} = \tilde{z}_{kt}^x$ in equilibrium and $\tilde{\rho}_w = \rho_x$; and the monetary authority fully stabilizes expected aggregate nominal wage growth.

We obtain,

$$\begin{aligned}
{\pi^w}_t &= \lambda(\pi_t + \epsilon_t - \epsilon_{t-1} + \frac{1}{\phi}(n_t - n_{t-1})) + (1-\lambda){\pi^w}_{t-1} \\
\tilde{w}_{kt} &= \frac{\lambda}{1+\lambda\kappa(1-\tilde{\rho}_w)}(\tilde{p}_{kt} + \tilde{\epsilon}_{kt} + \frac{1}{\phi}\tilde{n}_{kt}) + \frac{1-\lambda}{1+\lambda\kappa(1-\tilde{\rho}_w)}\tilde{w}_{kt-1}
\end{aligned}$$

Then, we can write,

$$\lambda = \frac{1-\beta_w}{1+\beta_w\kappa(1-\tilde{\rho}_w)}$$

where $\beta_w \equiv \frac{1-\lambda}{1+\lambda\kappa(1-\tilde{\rho}_w)}$. From this expression we see that our estimates for λ using cross-state variation are upward biased. However, we can get a notion on the magnitude of the bias by asking what would $\kappa(1-\tilde{\rho}_w)$ have to be in order for λ to be less than some λ_0. We obtain,

$$\kappa(1-\tilde{\rho}_w) > \frac{1-\beta_w-\lambda_0}{\beta_w\lambda_0}$$

For example, given our lower estimate for $\beta_w = 0.5$, in order for λ to be below 0.1 we would need a $\kappa(1-\tilde{\rho}_w)$ larger than 8.

C Figures and Tables

Figure 1: Nielsen Retail Price Index vs. CPI Food Price Index

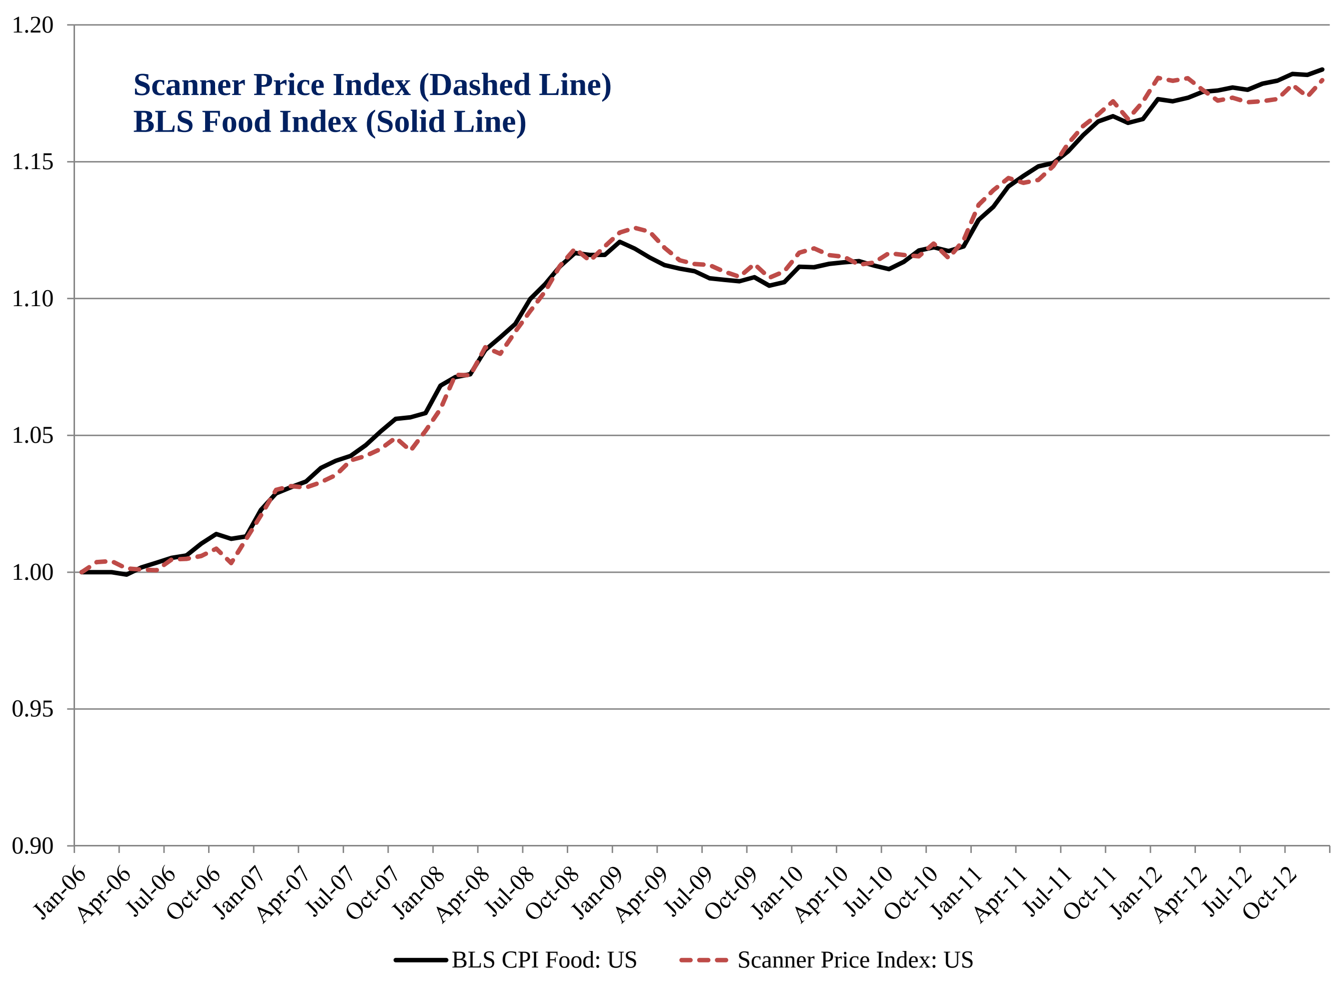

Note: In this figure, we compare our monthly retail scanner price index for the U.S. as a whole (dashed line) to the CPI's aggregate monthly food price index (solid line). Given that the goods in our price index come predominantly from grocery, pharmacy, and mass merchandising stores, we choose the CPI food index as our benchmark. We normalize both indices to 1 in January of 2006.

Figure 2: The Evolution of Aggregate Real and Nominal Composition Adjusted Wages

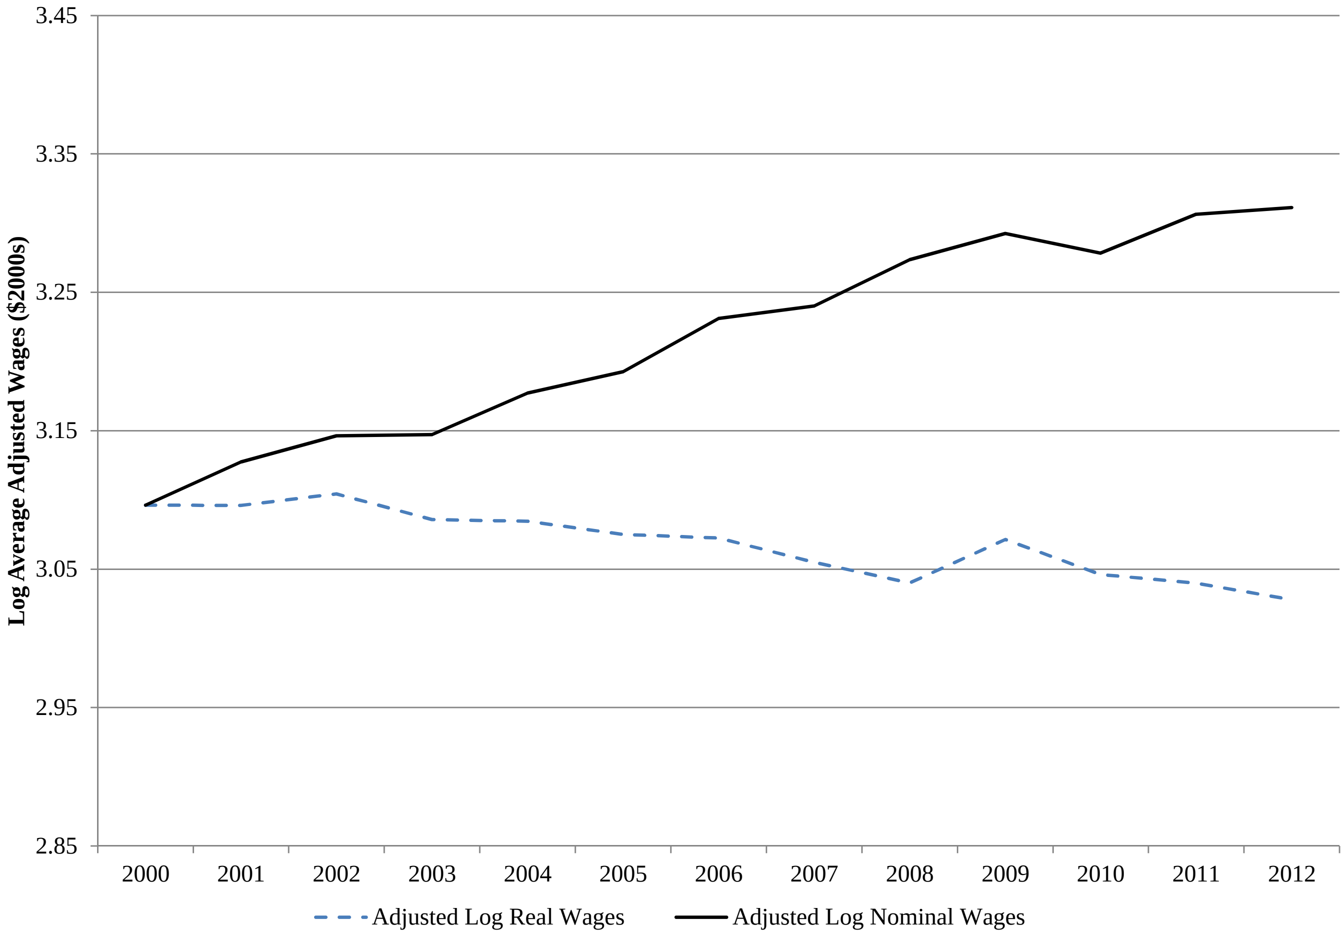

Note: Figure shows the evolution of aggregate real and nominal log wages within the U.S. between 2000 and 2012. Nominal wages were computed using data from the Current Population Survey. The sample is restricted to only males between the ages of 21 and 55, who are currently employed, who report usually working 30 hours per week, and who worked at least 48 weeks during the prior 12 months. Nominal wages are computed by dividing individual reports of labor earnings over the last 12 months by their hours worked over the last 12 months. Hours worked over the last 12 months are computed by multiplying weeks worked last year by the usual hours they currently report working. Survey measures from year t refer to wages earned in year t-1. The figure tracks when wages were earned. As discussed in the text, we adjust wages for the changing labor market condition over time by controlling for age, race, education, and usual hours worked. As computed, the wages are for a white male aged 40-44 who attended some college (but without a 4-yr degree) working 40 hours per week. We compute real wages by deflating our nominal wage index by the CPI-U of the corresponding year.

Figure 3: Change in State Unemployment Rate vs. Cumulative State Retail Price Inflation, 2007-2010

Note: Figure shows a simple scatter plot of the percentage point change in the BLS unemployment rate in the state between 2007 and 2010 against the cumulative percent change in our retail price index based on the Nielsen scanner data during the same period. The size of the underlying state is represented by the size of the circle in the figure. The line represents a weighted regression line from the bi-variate regression.

Figure 4: Change in State Unemployment Rate vs. State Nominal Wage Growth, 2007-2010

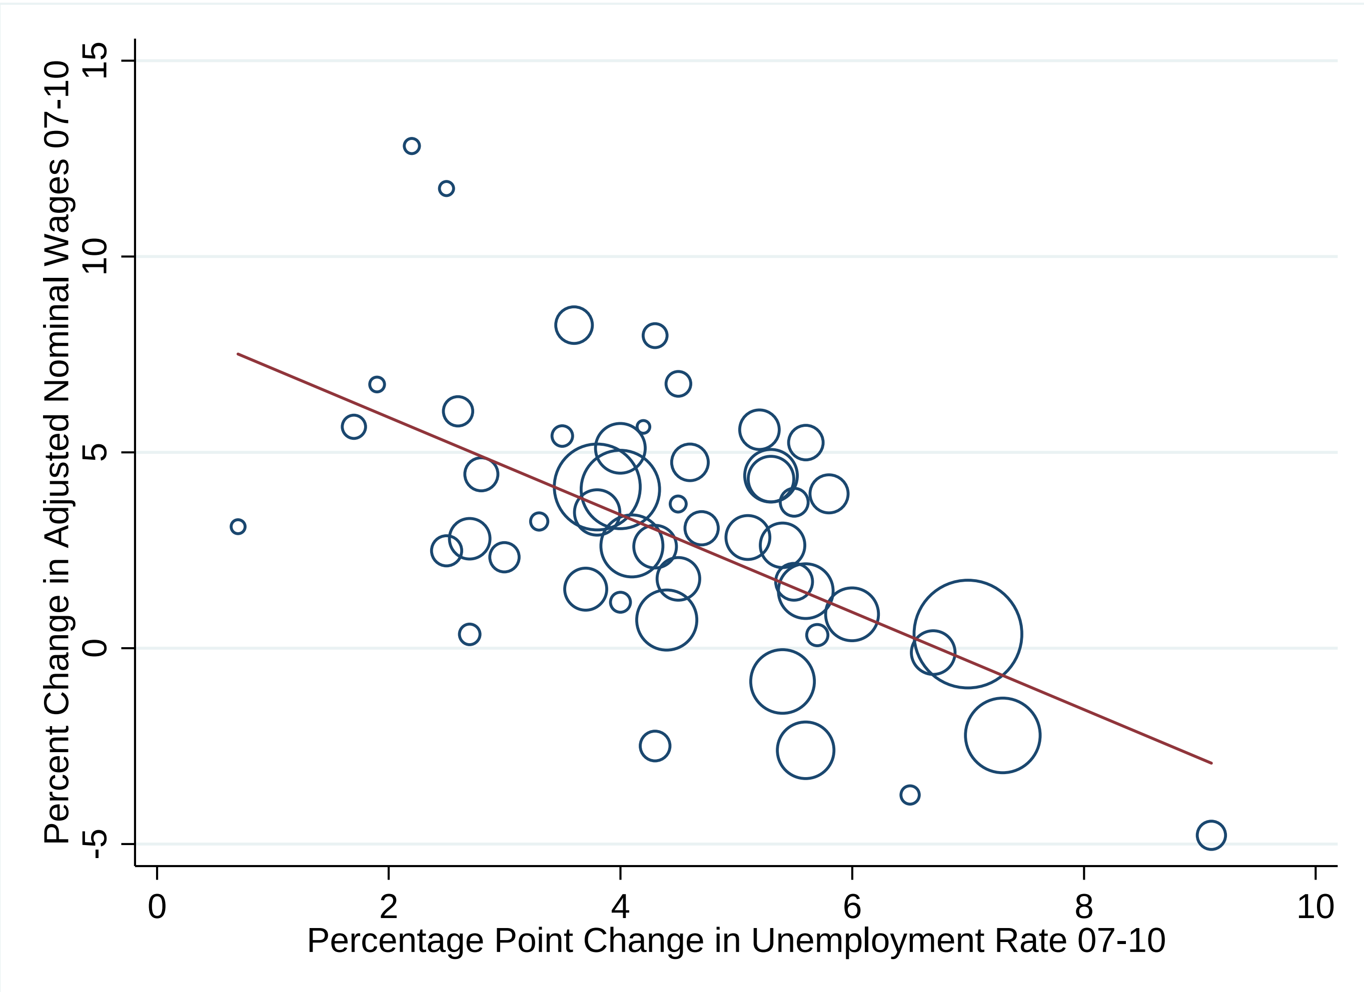

Note: Figure shows a simple scatter plot of the percentage point change in the BLS unemployment rate in the state between 2007 and 2010 against nominal wage growth during the same period. Nominal wages are computed from the ACS and are adjusted for the changing labor market composition of workers within each state over time. We restrict wage measures to a sample of men between the ages of 21 and 55 with a strong attachment to the labor market. Our composition adjustment controls for age, education, race, nativity and usual hours worked. See text for details. The size of the underlying state is represented by the size of the circle in the figure. The line represents a weighted regression line from the bi-variate regression.

Figure 5: Impulse Response to a Discount rate Shock

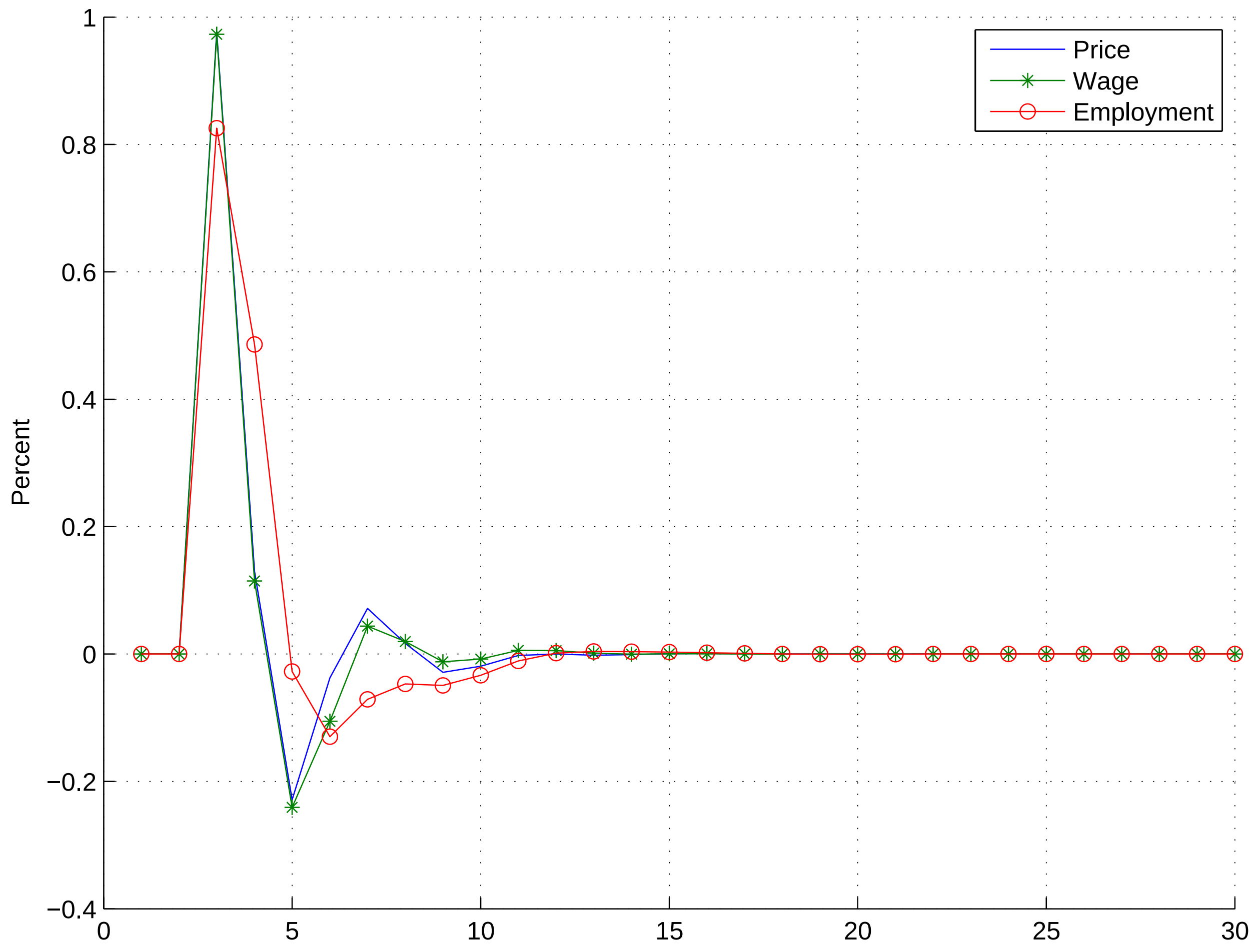

Note: Figure shows the impulse response to a one standard deviation discount rate shock. The horizontal axis are years after the shock.

Figure 6: Impulse Response to a Productivity / Markup Shock

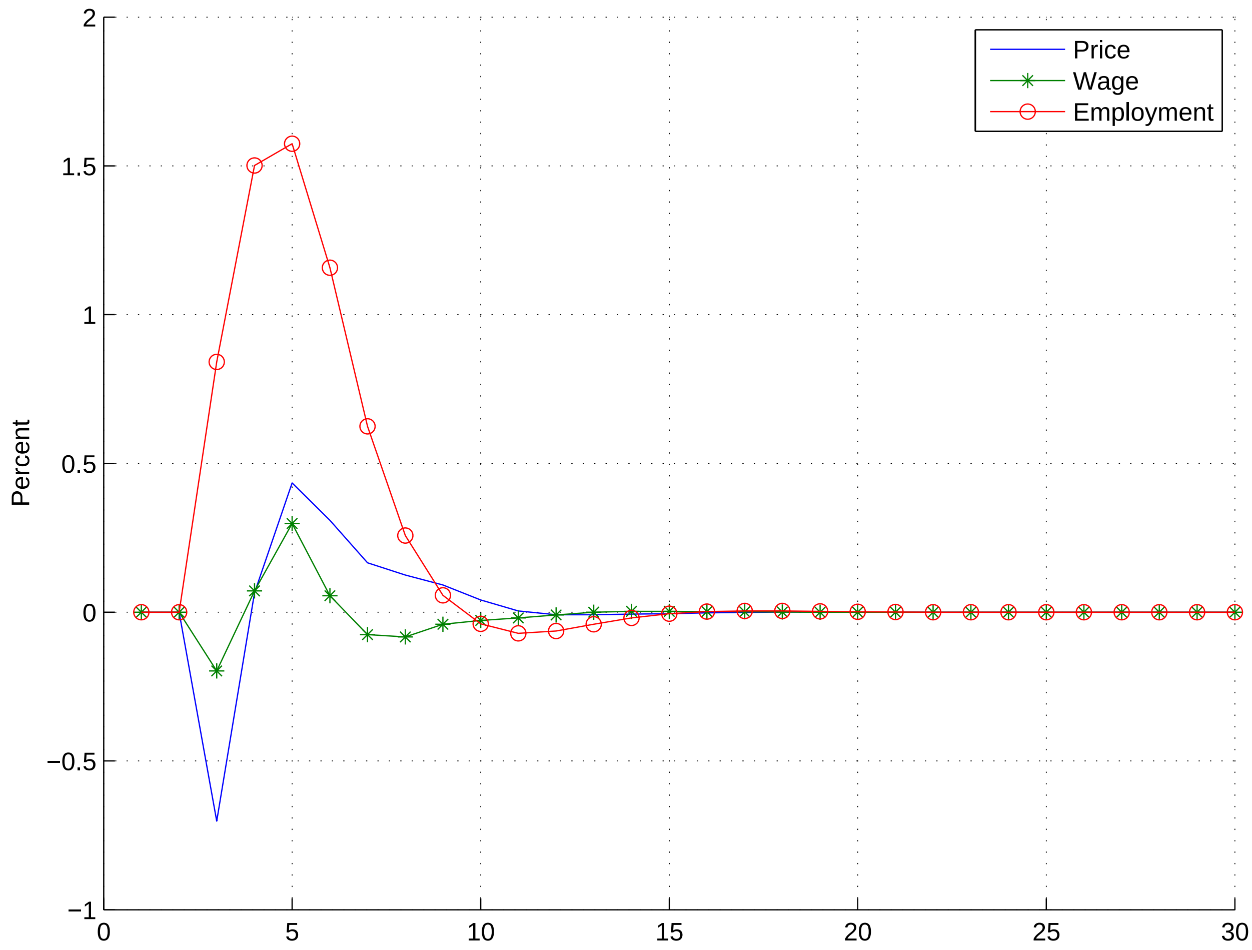

Note: Figure shows the impulse response to a one standard deviation productivity/markup shock. The horizontal axis are years after the shock.

Figure 7: Impulse Response: Leisure Shock

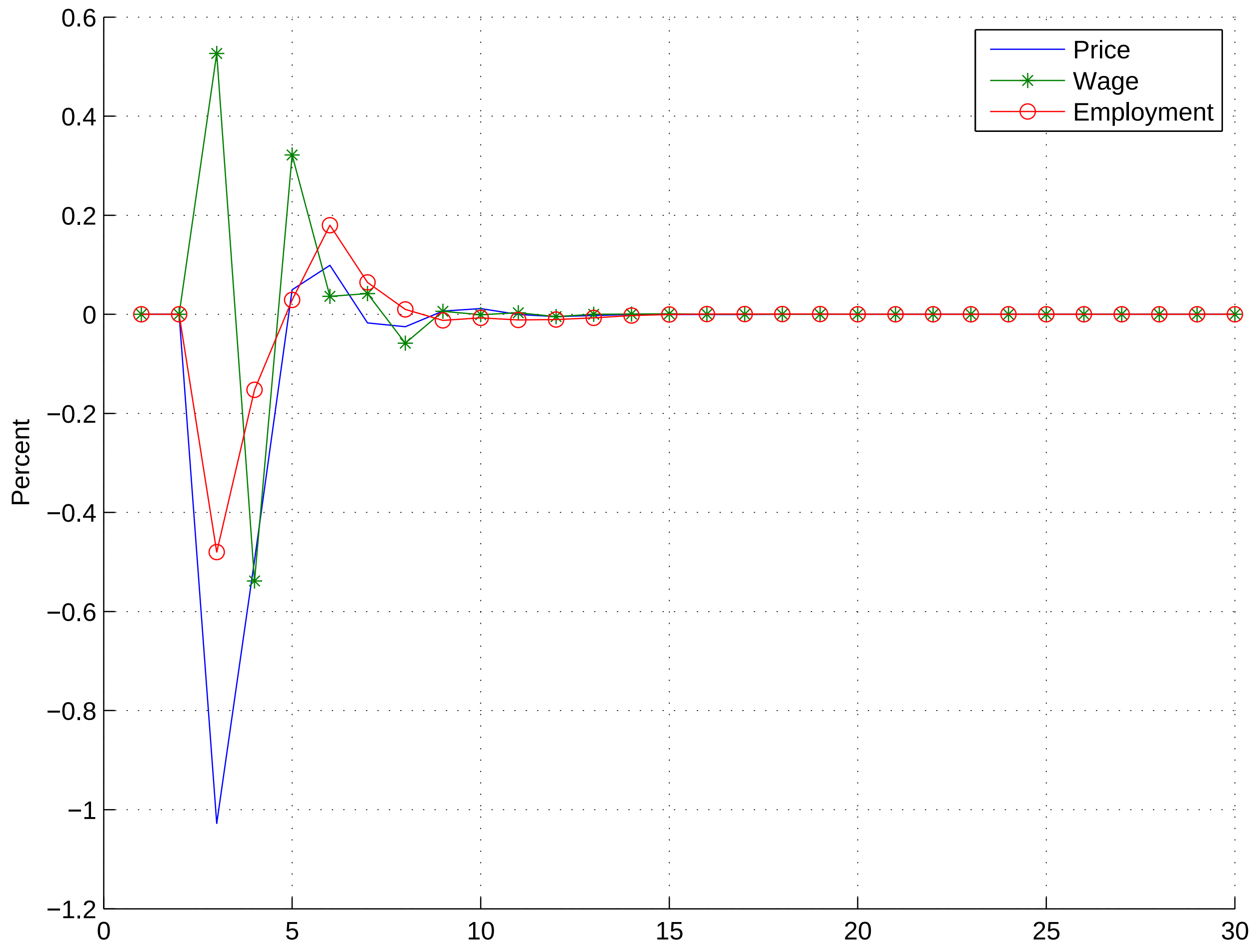

Note: Figure shows the impulse response to a one standard deviation leisure shock . The horizontal axis are years after the shock.

Figure 8: Shock Time Series

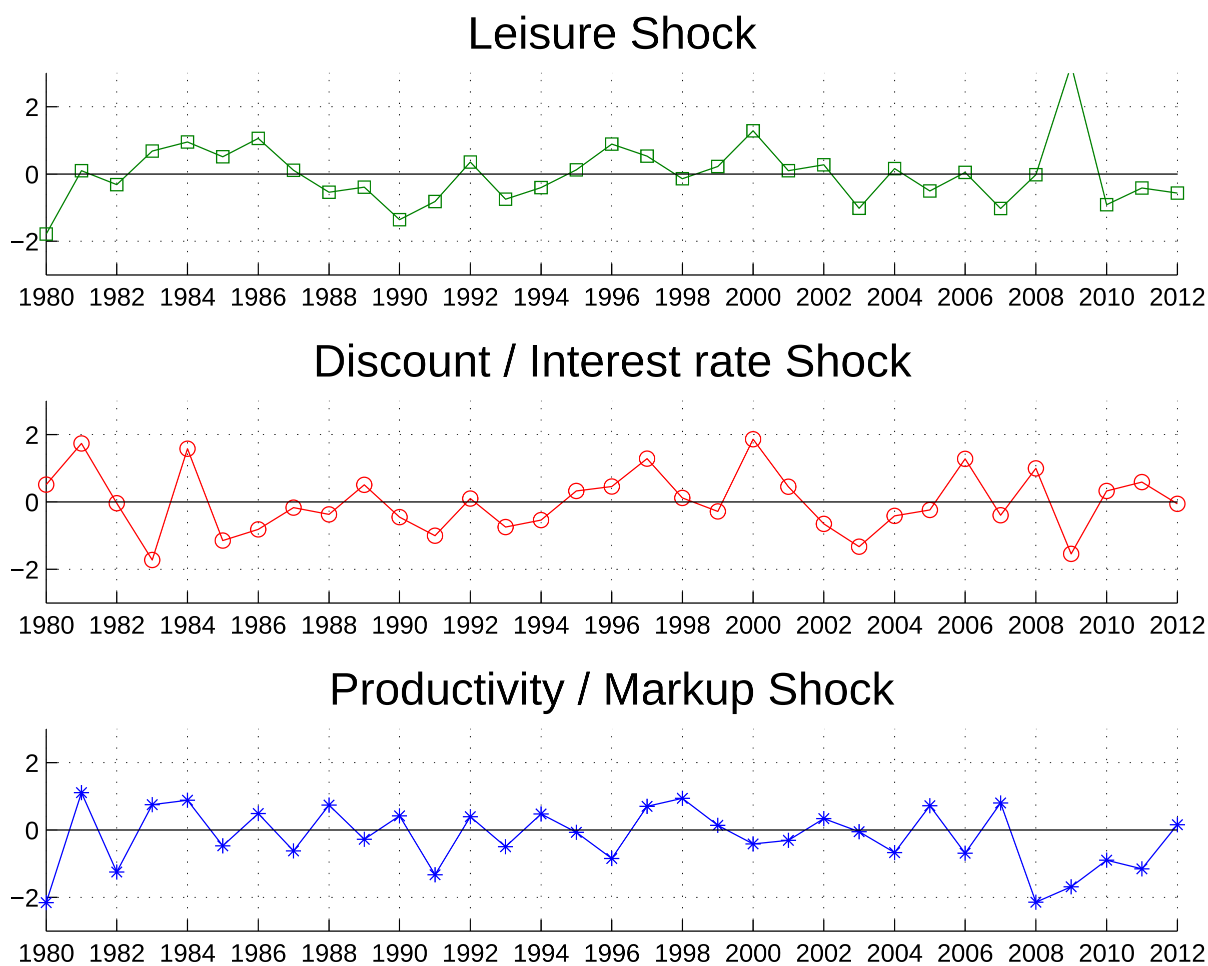

Figure 9: Counterfactual Employment Response

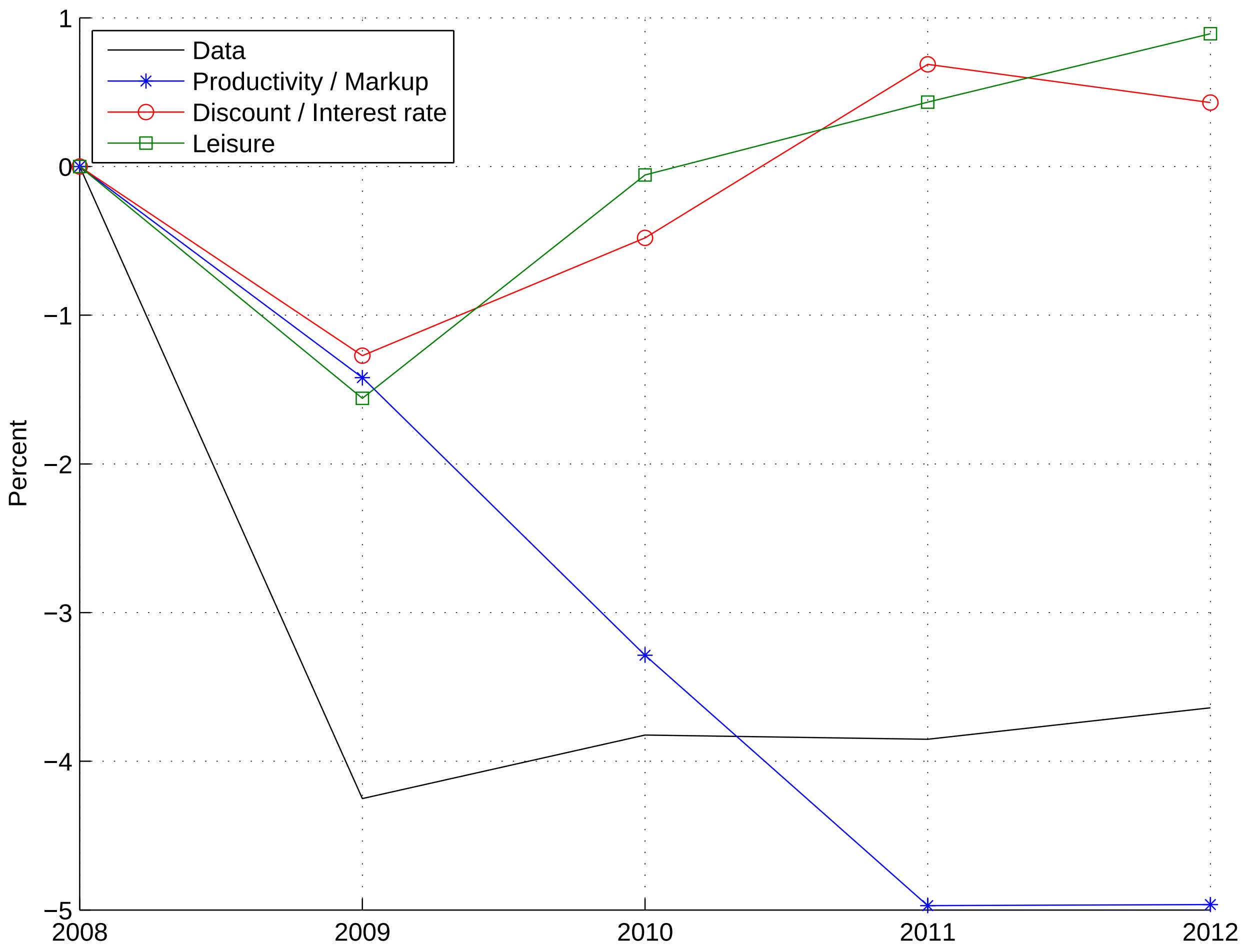

Note: Figure shows the cumulative response of employment when we feed the VAR with the sequence of shocks between 2008 and 2012; one at a time.

Figure 10: Counterfactual Price Response

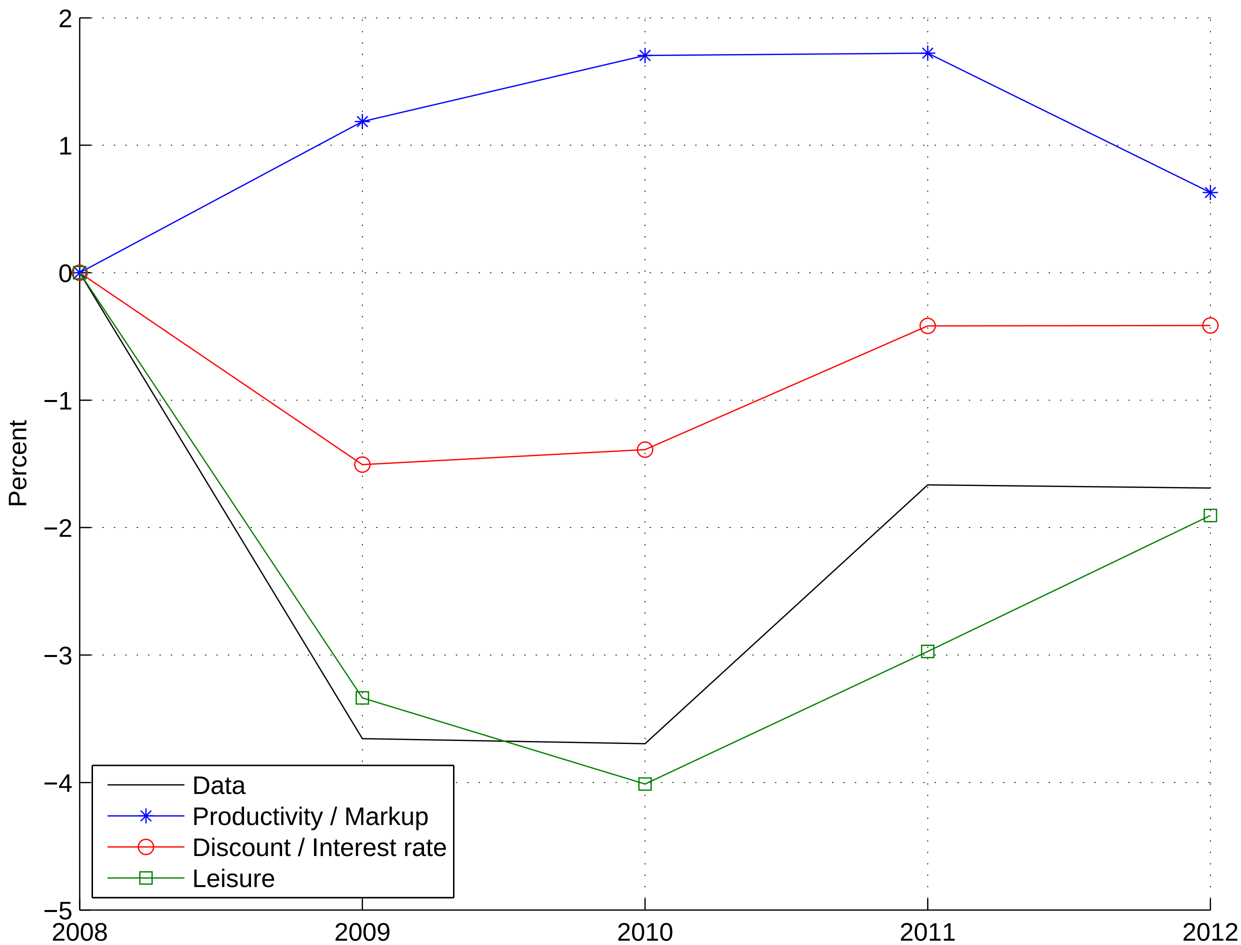

Note: Figure shows the cumulative response of Prices when we feed the VAR with the sequence of shocks between 2008 and 2012; one at a time.

Figure 11: Counterfactual Wage Response

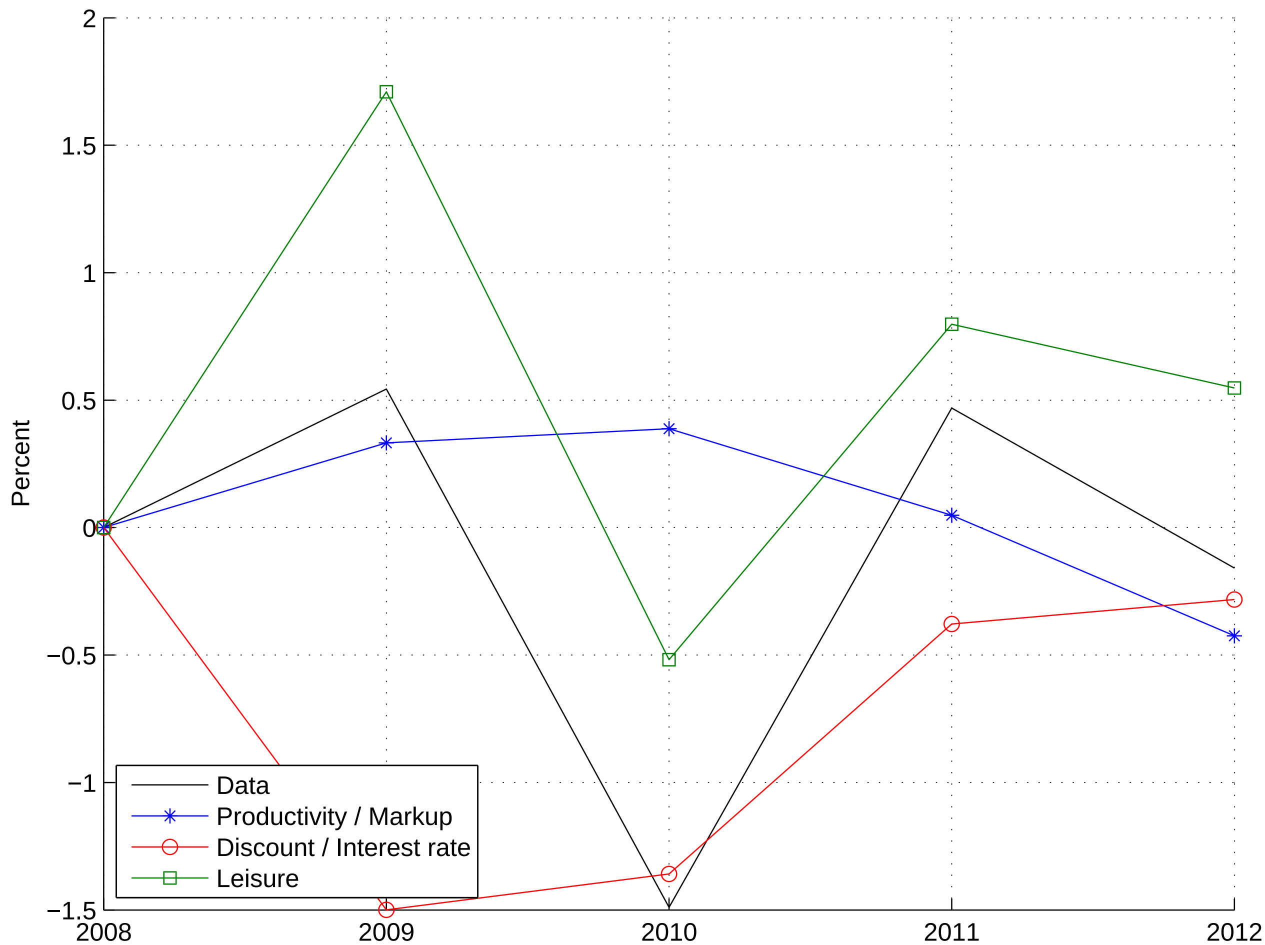

Note: Figure shows the cumulative response of Wages when we feed the VAR with the sequence of shocks between 2008 and 2012; one at a time.

Table 1: The Relationship between Regional Price Changes and Changes in Regional Economic Activity: 2007-2010

	Specification	
	(1)	(2)
	No Scaling	Scaling Based on BLS Regional Variation
Change in Unemployment Rate (Percentage Point)	-0.354	-0.495
	(0.133)	(0.187)
Per-Capita GDP Growth (Percent)	0.087	0.122
	(0.047)	(0.065)
Per-Capita Hours Growth (Percent)	0.139	0.195
	(0.077)	(0.108)
House Price Growth (Percent)	0.033	0.045
	(0.012)	(0.017)
Employment Rate Growth (Percent)	0.070	0.097
	(0.073)	(0.102)
IV: Change in Unemployment Rate (Percentage Point)	-0.432	-0.604
	(0.162)	(0.227)
IV: Employment Growth	0.292	0.409
	(0.127)	(0.178)
Scaling Factor	1.0	1.4

Note: Table shows the results of a bi-variate regression of the inflation rate in a given state between 2007 and 2010 against changing measures of real activity within the state between 2007 and 2010. Column 1 uses the underlying data from our sample to compute the price indices (Pr). Column 2 scales the price index from the scanner data (Pr) by estimates of the nontradable share in the retail grocery industry and the nontradable share for a composite local consumption good (a factor of 1.4). See text for additional details. Standard errors are in parenthesis. Each regression is weighted by the state's 2006 population.

Table 2: The Relationship between Regional Wage Changes and Changes in Regional Economic Activity: 2007-2010

	Dependent Variable		
	Nominal Wage Growth	Real Wage Growth (No Scaling of Price Index)	Real Wage Growth (Base Scaling of Price Index)
Change in Unemployment Rate (Percentage Point)	-1.244	-0.892	-0.751
	(0.205)	(0.264)	(0.300)
Per-Capita GDP Growth (Percent)	0.487	0.399	0.365
	(0.060)	(0.080)	(0.093)
Per-Capita Hours Growth (Percent)	0.653	0.520	0.464
	(0.120)	(0.146)	(0.165)
House Price Growth (Percent)	0.113	0.080	0.067
	(0.019)	(0.024)	(0.027)
Employment Rate Growth (Percent)	0.618	0.546	0.518
	(0.109)	(0.129)	(0.146)
IV: Change in Unemployment Rate (Percentage Point)	-1.500	-1.063	-0.890
	(0.253)	(0.321)	(0.364)
IV: Employment Growth (Percent)	1.011	0.719	0.602
	(0.195)	(0.209)	(0.232)
Scaling Factor For Regional Price Differences		1.0	1.4

Note: Table shows the results of a bi-variate regression of nominal wage growth (column 1) or real wage growth (columns 2 and 3) in a given state between 2007 and 2010 against changing measures of real activity within the state between 2007 and 2010. Wages are measured using the American Community Survey and are adjusted for the changing composition of the workforce. When computing real wages, we adjust nominal wages the unscaled retail price index (column 2) and by the retail price index scaled by a factor of 1.4 (column 3). See the note to Table 1 for a discussion of the scaling of the retail scanner price index. Standard errors are in parenthesis. Each regression is weighted by the state's 2006 population.

Table 3: Comparison of Aggregate and Local Wage and Price to Unemployment Changes During the Great Recession

	Aggregate Variable			
	Unemployment Rate	Nominal Wage	Prices	Real Wage (in 2000$)
2007 Aggregate Level	4.6 percent	$25.5/hour	207.2	$21.2/hour
2010 Aggregate Level	9.6 percent	$26.5/hour	217.2	$21.0/hour
Actual Aggregate Change/Growth, 2007-2010	5.0 percentage points	3.8 percent	4.7 percent	-0.9%
Expected Aggregate Growth, 2007-2010 (Based on 2000-2007 Trend)		5.5 percent	7.6 percent	-2.1%
Aggregate Deviation from Expected Growth, 2007-2010		-1.7 percent	-2.9 percent	1.2 percent
Aggregate Semi Elasticity With Respect to 1 Percentage Point Change in Unemployment Rate, 2007-2010		-0.33	-0.57	0.24
Local Semi Elasticity With Respect to 1 Percentage Point Change in Unemployment Rate, 2007-2010		-1.24	-0.50	-0.75

Note: Table compares the aggregate time series relationship between nominal wage growth, price growth, and real wage growth with unemployment rate changes during the 2007 to 2010 period. The wage data is the adjusted wage series we created using data from the CPS. See text for details. The unemployment and price data come from the U.S. Bureau of labor statistics. For prices, we use the CPI-U. For both measures, we use data from June of the corresponding year. Real wages are our composition adjusted nominal wage deflated by the June CPI-U. To get predicted wage and price growth between 2007 and 2010, we take a simple linear prediction of the corresponding wage and price growth between the 2000 and 2007 period. Once we get the deviation between actual wage (price) growth and predicted wage (price) growth between 2007 and 2010, we divide that difference by the 5 percentage point decline in the unemployment rate between 2007 and 2010. The last row of the table just redisplays the corresponding local relationships between wages (prices) and unemployment documented in Tables 1 and 2.

Table 4: Calibration

	Value	**Target**
α	0.62	Aggregate labor share in the non-tradable sector in 2006
θ	0.55	Aggregate labor share in the tradable sector in 2006
β	0.16	Aggregate labor share in 2006 (0.61). See on-line appendix
σ	2	Intertemporal elastictiy of substitution
ϕ	2	Frisch elasticity. See Section 6.2 for estimation
λ	0.7	Wage stickiness. See Section 6.2 for estimation
φ_p	1.5	Taylor Rule from Galí (2011)
φ_y	0.5	Taylor Rule from Galí (2011)
Φ_0	0.01	Trade balance-output ratio volatility = 2 from Mendoza (1991).
R	0.03	Real interest rate
X	0.32	Intermediate inputs over output ratio US 2006
B	2.1	Median net worth to output ratio in the US 2006
ρ_γ	0.9	Persistence of discount rate shock
ρ_z	0.76	Persistence of productivity shock from the model. See on-line appendix
ρ_ϵ	0.66	Persistence of leisure shock from the model. See on-line appendix

Table 5: **Aggregate v. Regional Employment Impact Elasticities**

	ρ_γ	**0.9**			**0.9**	**0.9**		**0.6**	**0.1**
	β	**0.16**			**0**	**0.16**		**0.16**	
	φ_p	**1.5**			**1.5**	**1.5**	**1.31**	**1.5**	
	φ_y	**0.5**			**0.5**	**0**		**0.5**	
	(λ,ϕ)	**(0.7,2)**	**(0.7,1)**	**(0.5,2)**	**(0.7,2)**	**(0.7,2)**	**(0.7,2)**	**(0.7,2)**	**(0.7,2)**
Aggregate	γ	0.77	0.53	1.29	0.74	1.11	1.76	0.71	0.31
	z	0.29	0.20	-0.35	0.28	-0.24	-1.07	0.29	0.29
	ϵ	-0.65	-0.51	-0.28	-0.69	-0.26	0.41	-0.65	-0.65
Regional	γ	1.76	1.26	2.95	2.85	1.76	1.76	0.36	0.03
	z^y	0.44	0.38	-0.18	0.10	0.44	0.44	0.44	0.44
	z^x	0.08	-0.06	0.23	0.05	0.08	0.08	0.08	0.08
	ϵ	-0.64	-0.52	-0.16	-0.43	-0.64	-0.64	-0.64	-0.64

Note: The table summarizes the response of employment on impact to each of the model shocks under our base calibration (column 1) and alternate calibrations (columns 2-8). The rows represent the employment response to different aggregate shocks (top three rows) and different local shocks (bottom three rows). Columns 2 and 3, explore the elasticities under different calibrations of wage stickiness and the Frisch elasticity of labor supply. Column 4 examines the robustness to changes in the tradable share of the intermediate good. Columns 5 and 6 examine the results under alternate Taylor Rule parameters. The final two columns change the persistence of the demand shock. The units are percentage deviations from the steady state, in the case of aggregate employment, and percentage deviations from the aggregate in the case of regional employment.

Table 6: **Aggregate Elasticities (λ=0.7, ϕ=2)**

		n	π	π^w	w^r
Short run	γ	0.77	1.64	1.42	-0.22
	z	0.29	-2.73	-1.81	0.92
	ϵ	-0.65	0.94	1.13	0.19
Long run	γ	0.56	1.01	1.03	-0.16
	z	0.12	-0.66	-0.79	0.40
	ϵ	-0.18	0.35	0.32	0.05

Note: The table summarizes the response of each aggregate variable to the shocks in percentage deviations from the steady state. The "short" elasticity is the response at date $t = 0$. The "long" elasticity is the response after 5 years.

Table 7: **Regional Elasticities (λ=0.7, ϕ=2)**

		n	π	π^w	w^r
Short run	γ	1.76	2.54	2.40	-0.14
	z^y	0.44	-2.00	-1.25	0.75
	z^x	0.08	0.29	0.31	0.02
	ϵ	-0.64	0.71	0.97	0.26
Long run	γ	0.12	-0.28	-0.30	0.19
	z^y	0.63	0.31	0.23	0.21
	z^x	-0.06	-0.03	-0.04	0.04
	ϵ	-0.44	-0.12	-0.20	0.14

Note: The table summarizes the response of each island level variable to the shocks in percentage deviations from the steady state. The "short" elasticity is the response at date $t = 0$. The "long" elasticity is the response after 5 years.

Table 8: Estimates of λ and $\frac{\lambda}{\phi}$ using Cross- Region Data

	Specification						
	OLS					IV	
	2007-2011			2007-2009		2007-2011	2007-2009
	(1)	(2)	(3)	(5)	(6)	(7)	(8)
λ	0.69	0.69	0.75	0.73	0.73	0.77	0.79
	(0.13)	(0.13)	(0.13)	(0.17)	(0.17)	(0.13)	(0.18)
λ/ϕ	0.31	0.32	0.31	0.39	0.39	0.76	0.99
	(0.08)	(0.08)	(0.07)	(0.09)	(0.10)	(0.17)	(0.25)
Implied ϕ	2.2	2.2	2.4	1.9	1.9	1.0	0.8
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry Controls	No	Yes	Yes	No	Yes	Yes	Yes
Scaling Factor of Prices	1.4	1.4	1.0	1.4	1.4	1.4	1.4

Note: Table shows the estimates of λ and $\frac{\lambda}{\phi}$ from our base wage setting specification using the regional data. Each observation in the regression is state-year pair. Each column shows the results from different regressions. The regressions differ in the years covered and additional control variables added. The first three columns show the OLS results using all local data between 2007 and 2011. Columns 4 and 5 show OLS results using only data from 2007 through 2009. The final two columns show IV results for the different time periods. In the IV specifications, we instrument contemporaneous employment and price growth with contemporaneous and lagged house price growth. We adjust for measurement error in wage growth, lagged wage growth, and price growth using the split sample methodology discussed in the Online Data Appendix. All regressions included year fixed effects. All standard errors are clustered at the state level.

Table 9: Discount/Interest rate (γ) and Productivity/Markup (z) shocks' contribution to aggregate employment change

		2008 to 2009					2008 to 2012				
		ϕ					ϕ				
λ		**0.5**	**1**	**2**	**3**	**4**	**0.5**	**1**	**2**	**3**	**4**
0.1	γ	103	83	108	107	108	40	-34	46	57	52
	z	-3	22	-3	-1	-2	47	126	48	38	43
0.3	γ	47	66	51	29	101	-13	-25	-33	-20	9
	z	2	16	45	71	1	98	123	134	121	92
0.5	γ	8	36	31	13	92	-1	-13	-19	2	47
	z	6	16	48	74	0	94	123	133	111	65
0.7	γ	0	0	30	21	11	47	48	-12	-12	0
	z	3	33	33	53	69	53	72	136	135	122
0.9	γ	0	0	0	2	3	45	47	41	15	9
	z	-1	24	54	64	69	58	79	91	114	118

Note: Table shows the percent contribution of the demand and supply shocks to the aggregate employment change implied by our procedure for different combinations of the parameters. For a given pair $\{\phi, \lambda\}$, the 'γ' entry corresponds to the demand shock. The 'z' entry to the supply shock. The percent contribution of the leisure shock can be calculated by subtracting the sum of both entries from 100. Entries with $*$ are such that no decomposition of the shocks satisfy the identification restrictions for those parameter values.

Table 10: Discount rate (γ) and Productivity/Markup (z) shocks' contribution to change in regional variables

(λ, ϕ)		$\tilde{\mathbf{w}}$	$\tilde{\mathbf{p}}$	$\tilde{\mathbf{n}}^{\mathbf{y}}$
(0.7,2)	ξ^{γ}	24	41	79
	ξ^{z^y}	17	39	11
	ξ^{z^x}	60	20	11
(0.5,1)	ξ^{γ}	19	40	79
	ξ^{z^y}	14	41	11
	ξ^{z^x}	67	19	10

Note: The table shows the contribution of each regional shock to each regional wage, price and non-tradable employment change between 2007-2010. As defined in Section 7.3, for each variable y and shock j we compute $\xi^j_y = \frac{\frac{1}{\sum_k \omega_k (\Delta^j \tilde{y}_k - \Delta \tilde{y}_k)^2}}{\sum_j \frac{1}{\sum_k \omega_k (\Delta^j \tilde{y}_k - \Delta \tilde{y}_k)^2}}$.

Table 11: Regional counterfactual statistics

	Data	Discount rate	Productivity/Markup
$\sigma_n^2/(\sigma_n^{data})^2$	1	0.98	0.35
$\sigma_p^2/(\sigma_p^{data})^2$	1	0.69	0.62
$\sigma_w^2/(\sigma_w^{data})^2$	1	0.47	0.26
β_{p,n^y}	0.48	0.73	-1.82
β_{w,n^y}	0.51	0.47	-0.66
$\beta_{p,w}$	0.32	0.65	0.37

Note: The first three lines in the table show the counterfactual variance across states relative to the actual variance of the total percent change in each variable between 2007-2010. The last three lines show the population weighted OLS coefficient corresponding to each variable pair. For example, β_{p,n^y} is the coefficient in the regression of price growth between 2007-2010 onto employment growth in the non-tradable sector where each state is weighted by its population in 2006. The second column corresponds to the counterfactual with the γ shock alone. The third column corresponds to the counterfactual with both z^x, z^y shocks and no γ shock.

Figure A1: QEW Nominal Per Worker Earnings Growth vs. ACS Adjusted Nominal Wage Growth, 2007-2010

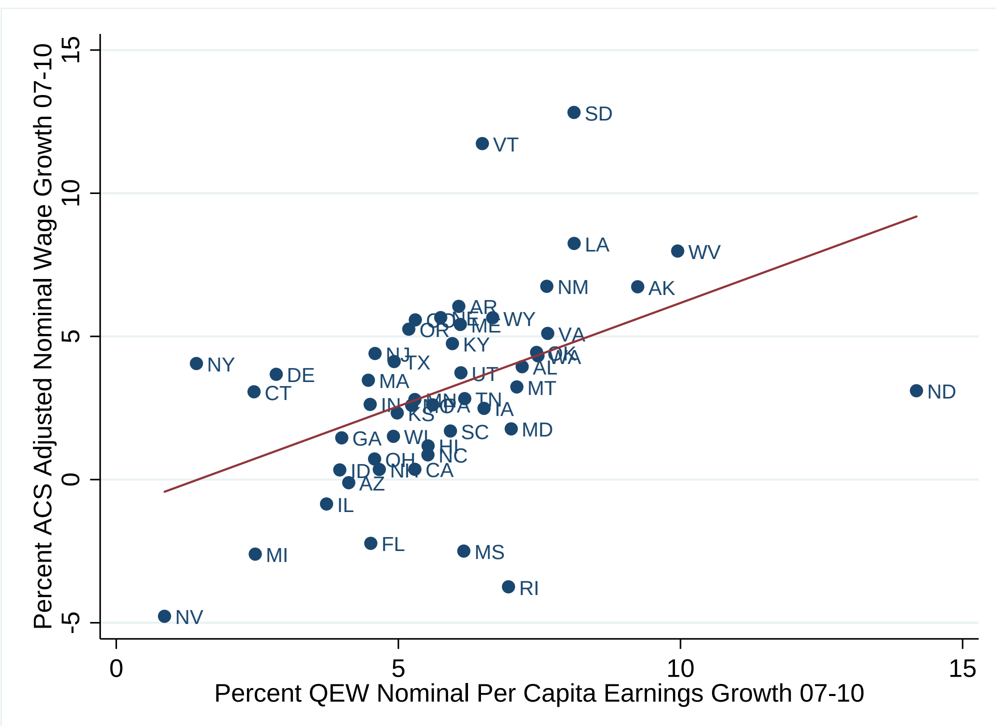

Note: Figure shows the simple scatter plot of the growth in our ACS adjusted nominal wage index from 2007-2010 against the nominal per worker earnings growth from 2007-2010 from the BLSÕs QEW. Each observation is a U.S. state. The slope of the regression line is 0.72 with a standard error of 0.20.

Figure A2: QEW Nominal Per Worker Earnings Growth vs. QEW Employment Growth, 2007-2010

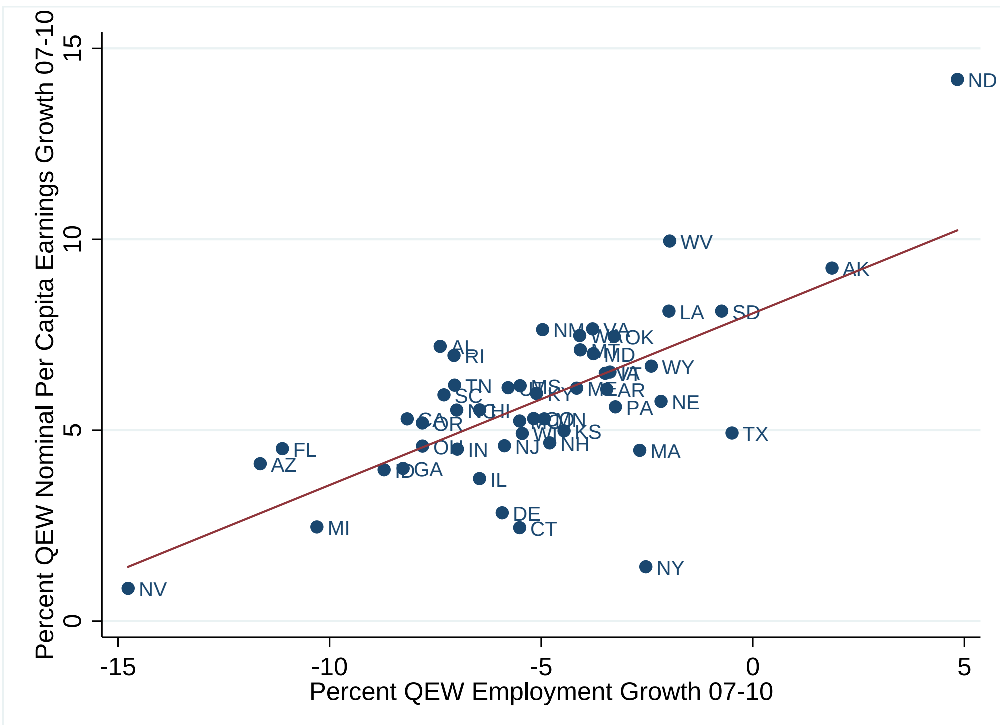

Note: Figure shows the simple scatter plot of the nominal per worker earnings growth from 2007-2010 from the BLSÕs QEW and employment growth from the QEW from 2007-2010. Each observation is a U.S. state. The slope of the regression line is 0.45 with a standard error of 0.07.

Figure A3: State Net Migration Rate 2009-2010 vs. State Employment Growth 2007-2010

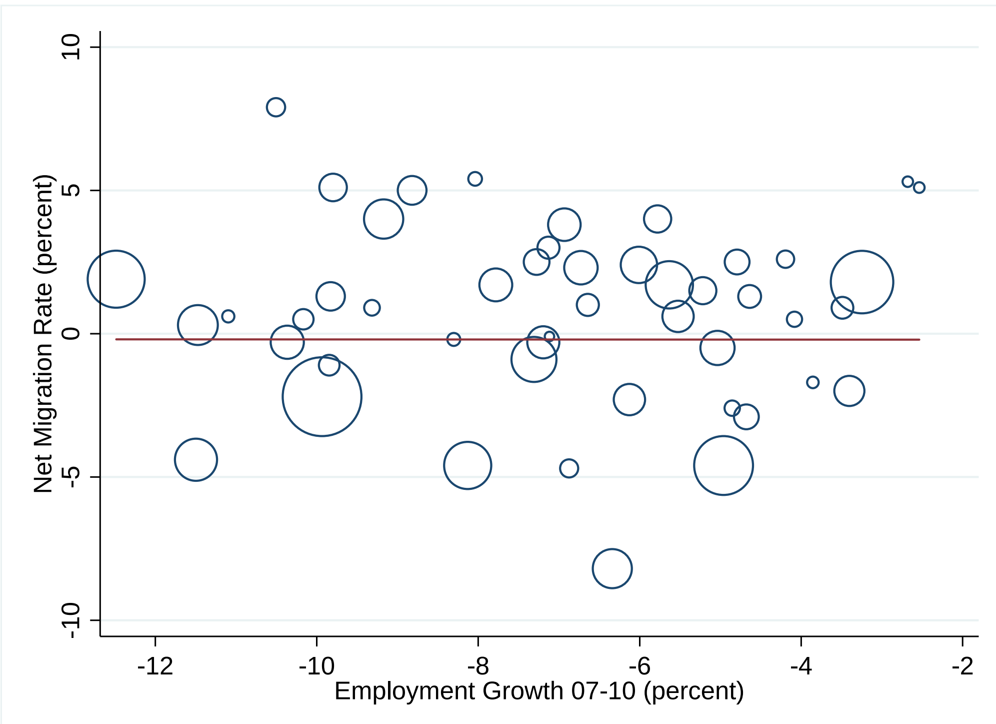

Note: Figure shows state net migration rate between 2009 and 2010 against employment growth in the state during 2007-2010. Employment growth comes from the BLS and is defined in the text. State net migration rates come from American Community Survey.