

The Aggregate Implications of Regional Business Cycles

Martin Beraja Erik Hurst Juan Ospina
University of Chicago

February 25, 2015

Preliminary

Abstract

During the last recession and its aftermath, aggregate price growth and aggregate nominal wage growth remained robust despite the high rate of unemployment. We document that U.S. states with weaker labor markets experienced lower price growth, lower nominal wage growth, and lower real wage growth during the Great Recession. Given that reliable measures of local prices do not exist, we create our own state price indices using scanner data from the Nielsen Retail Scanner Database. The main innovation of the paper is the development of an econometric procedure that combines cross-region data with aggregate time series data to infer the underlying shocks driving both aggregate and regional business cycles. Applying our procedure to the Great Recession, 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 contrast, we find that only demand shocks are needed to explain the cross-region variation. The results suggest that only using cross-region variation to explain aggregate fluctuations is insufficient when shocks do not have a substantive regional component.

1 Introduction

As both output and employment contracted sharply within the U.S. during the Great Recession, consumer price growth and nominal wage growth remained robust. Figure 1 plots the relationship between the CPI inflation rate and the aggregate unemployment rate between 2000 and 2013. Each observation is a quarter with the inflation rate from one quarter to the next being expressed at an annual rate. Aside from the fourth quarter of 2008 and the first quarter of 2009, consumer prices grew around 2 percent. What is striking is that as the unemployment rate continued to increase after the first quarter of 2009, consumer prices grew at a rate similar to the lower unemployment periods prior to 2007. Figure 2 shows nominal and real wage growth during the same period.¹ Between 2007 and 2010, average nominal wages within the U.S. increased by roughly 5 percent. Given that consumer prices increased by 5 percent during the same period, aggregate real wages in the U.S. were roughly constant between 2007 and 2010. This was similar to the trend in real wages prior to the start of the recent recession. 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.² Recently, a literature has emerged trying to explain the missing disinflation and the missing wage declines during this time period.³

In this paper, we document a series of new facts about the variation in prices and wages across U.S. states during the Great Recession. Specifically, we show that states that experienced smaller unemployment increases between 2007 and 2010 had much larger consumer price increases and had much larger nominal wage increases. While aggregate prices and wages did not appear to respond to weakening aggregate output and employment, we find a strong relationship between price growth/wage growth and employment estimated off of the cross-section of U.S. states during the Great Recession. The relatively flexible prices and wages at the local level stands in sharp contrast to the aggregate time series patterns. The goal of this paper is to combine the cross sectional relationships between local prices and wages and local real activity with the aggregate trends in consumer prices, nominal wages, and employment to infer the type of shocks experienced by both the aggregate and local economies during the Great Recession.

The paper is comprised of four distinct parts. In the first part we use data from Nielsen's *Retail Scanner Database* to compute price indices for each U.S. state. As far as we know, we are the first to use scanner data to compute local price indices. As a result, this portion of the paper provides an innovation to the literature in itself. The data comes directly from point of sale records of approximately 40,000 grocery, pharmacy, and mass

¹The wage data was computed by the authors using data from the Current Population Survey (CPS). Below, we provide a more extensive discussion of how we computed wages and how we adjusted for composition. To get the aggregate real wage series, we deflated the aggregate nominal wage series by the CPI.

²See, for example, Hall (2011), Ball and Mazumder (2011), and King and Watson (2012). This point was further made by Krugman in a recent New York Times article (Wages, Yellen and Intellectual Honesty, NYTimes 8/25/14).

³See, for example, Del Negro et al. (2014).

merchandising stores within the United States. We have access to the database for all weeks between the first week of 2006 and the last week of 2011. The database is massive in that it covers roughly 76 billion product*store*week transactions representing 200 billion dollars worth of transactions per year. The geographic coverage is also extensive. We have price information from 85 percent of the counties within the continental U.S.. The nature of the database allows us to make high frequency price indices for each U.S. state during the 2006-2011 period. We document that our retail scanner price index for the entire U.S. tracks the BLS's Food CPI for the entire U.S. nearly identically. Given that our retail goods have a large traded component, we also discuss how the variation in our retail good prices can be used to inform changes in a broader set of local prices.

In the second part of the paper, we show the key cross-sectional facts pertaining to the evolution of prices and wages in response to local economic activity during the Great Recession. Specifically, we show that during the recession, local consumer prices increased the most in places with the smallest employment declines. This may seem surprising in that we are primarily focusing on relatively tradable goods in our database. But, local distribution costs (rent paid by the retail establishment, wages of the retail workers, transportation and warehousing) are a nontrivial component of these retail goods. We estimate that a one percentage point increase in the unemployment rate between 2007 and 2010 was associated with a roughly 0.4 percentage points decrease in grocery/mass-merchandising prices and a roughly 0.8 percentage point decrease in prices for a composite local good.⁴ The difference between the grocery results and the results for the composite good stems from the fact that grocery prices have a larger tradable share. These findings show that the assumption that prices across U.S. sub-regions are invariant to local shocks is strongly rejected by the data.

Using data from the Current Population Survey and the Census/American Community Surveys, we also explore the variation in both nominal and real wages across U.S. states during the recession. Nominal wage growth varied substantially across U.S. states during the late 2000s. For example, California - which experienced a relatively large decline in employment during the recent recession - experienced no nominal wage growth between 2007 and 2010. Texas, on the other had, experienced a much smaller decline in employment during the recent recession and experienced nominal wage growth of 7 percent between 2007 and 2010. Pooling across all states, we find that a 1 percentage point increase in the state unemployment rate was associated with a 1.2 percentage point lower nominal wage growth. Using the nominal variation in wages and the documented regional variation in prices, we find that real wage growth was also negatively associated with declines in the state's employment rate.

In the third part of the paper, we lay out an econometric procedure that uses both the cross region and aggregate time series data to infer the underlying set of shocks driving both the aggregate and local economies. Our approach is semi-structural in that we infer the aggregate shocks using a VAR. However, we use theory to specify

⁴Below, we discuss the assumptions needed to translate cross-region variation in retail prices into cross-region variation in the price of a composite consumption good.

one structural equation within the VAR. We then show how the regional data can be used to estimate the key parameters of this structural equation. With one equation specified in the VAR, we use aggregate data to compute a variable that is correlated with one of the aggregate shocks and, by definition, orthogonal to the other shocks. These additional restrictions allows us to fully estimate the VAR without relying on ad hoc ordering assumptions or sign restrictions. The benefit of this procedure is that we can discipline the VAR using economic theory. However, we only need take a stance on one of the underlying structural equations. In this sense our procedure does not require us to rely on the complete model specification. With the estimates from the aggregate VAR and the regional micro data we then can infer the shocks that drive the cross-regional variation.

To provide context for our estimation procedure, we develop a simple model economy with many islands linked by trade in intermediate goods within a monetary union. Each island has a non-tradable sector that produces final consumption goods. Furthermore, we assume that nominal wages are only partially flexible, workers cannot move across islands and output is produced with only labor. Within the model, we allow for three broad types of shocks with each shock having an aggregate and local component. These broad types of shocks have been the focus of much recent work trying to explain the causes of the Great Recession. The first shock is akin to a standard “demand” shock. We model this as a shock to the household’s discount rate but it can be viewed as a proxy for the tightening of household borrowing limits, a decline in household wealth or a change in monetary policy. 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 current recession. The model also includes a shock to a firm’s productivity. We call it a “supply” shock. This shock is modeled as a pure productivity shock. However, this shock could also be interpreted as anything that changes firms’ demand for labor. For example, credit supply shocks to firms, such as those proposed by Gilchrist et al (2014), would be similar to our productivity shock. Finally, the model includes a preference shock for leisure relative to consumption. This “leisure” shock can be seen as a proxy for increasing distortions within the labor market due to changes in government policy. Mulligan (2012) has suggested such shocks were an important deterrent of economic activity during the Great Recession. Additionally, the leisure shock can be seen as a reduced form representation of the skill mismatch story put forth in Charles et al. (2013) for the recent decline in the U.S. employment rate.

Log-linearizing the above model yields something akin to a labor demand curve, a labor supply curve and an Euler equation that hold within each island which jointly determine local wages, local prices, and local employment. We also show that under relatively few assumptions, the local economy can be aggregated. The aggregate economy can then be represented with aggregate labor supply, labor demand, and Euler equations that have similar parameters as their local counterparts. Additionally, the aggregate economy can be represented as a reduced form VAR in prices, wages and employment. With the estimate of one of the structural equations, for example the ag-

aggregate labor supply (demand) curve, aggregate data can be used to infer a proxy for the aggregate labor supply (demand) shock. This information can then be used as an additional restriction when estimating the VAR. This procedure does not require one to take a stance on the true underlying model driving the remaining structural equations in the economy. With this additional restriction, the VAR is fully identified and the aggregate shocks to the discount rate, productivity, and the taste for leisure can be recovered.

Next, we show that under certain assumptions, regional data can be used to parameterize the chosen aggregate structural equation. Suppose the chosen equation was the aggregate labor supply (labor demand) curve. The key additional assumption in this case is that the regional component to the labor supply (labor demand) shock is orthogonal to the local movements in prices, wages and employment. Clearly, over different periods of time, the assumption that labor supply shocks and productivity shocks have no systematic regional component would be violated. However, during certain periods of time, it is potentially plausible to assume that either the labor supply shock or the productivity shock has no systematic regional component.

In the final part of the paper, we use both the regional and aggregate data to implement our procedure to estimate both the aggregate and local shocks during the Great Recession. We argue that during this period, regional data can be used to parameterize the aggregate labor supply curve. This amounts to assuming that the local shocks to labor supply are orthogonal to the shocks driving the regional variation in economic variables. To assess the plausibility of this assumption, we perform a variety of empirical tests illustrating that (1) there is very little regional variation in either federal or state policies that could affect labor supply during the Great Recession and (2) the variation that does exist is uncorrelated with local measures of economic activity. The policies we study include unemployment benefit extensions, welfare payments, state tax rates, and government programs to help homeowners.⁵ With the assumption of a stable labor supply curve across regions, we use annual variation across states in prices, wages and employment during the 2007 to 2010 period to estimate the parameters of the labor supply curve. We perform a variety of robustness exercises on our estimates. In particular, we explore whether differential migration patterns occurred across the states. The extent to which migration was important could imply that our labor supply estimates from the local economies would be larger than the parameters for the aggregate economy. We document that interstate migration rates were small, on average, during this period and that the extent to which they differed across states were uncorrelated with changes in local employment, prices, and wages.

Using our baseline estimates of the local labor supply curve to discipline the aggregate VAR, we find that demand shocks alone cannot explain the aggregate dynamics of prices, wages, and employment during the Great Recession. Quantitatively, we find that

⁵We do find some variation in unemployment benefit extensions across states that are correlated with local unemployment rates. This is not surprising given the extensions are triggered by the state unemployment rate. However, in 2010, most states were at 99 weeks. When estimating our local labor supply curve, we use data from 2007-2010 and 2007-2009. In 2009, the unemployment benefit extension policy had not been announced. We discuss this in greater depth below.

the demand shock was only responsible for roughly 50 percent of the fall in employment during the early portion of the recession. As the recovery continued, the demand shock contributed less and less to the low aggregate employment. If the economy only experienced the aggregate demand shock, the US economy would have experienced stagnant prices between 2007 and 2011. The aggregate supply shock, however, also explained roughly 50 percent of the decline in employment during the early part of the recession. As the recovery progressed, the supply shock was the primary cause of low employment growth during the recovery. We also show that it was the supply shock that prevented aggregate prices and wages from falling. We find that the leisure shock had relatively little effect on aggregate employment during either the recession or its aftermath. We did find that the leisure shock put modest downward pressure on aggregate prices and upward pressure on aggregate wages.

It is worth stressing that the fact that prices and wages adjusted so much at the local level helps to discipline our aggregate results. The rationale is that in the cross section of states (1) labor supply is fairly elastic and (2) wages are fairly (but not perfectly) flexible. Given this, it is not surprising that the demand shock is only explaining a portion of the price, wage, and employment dynamics during the Great Recession. In the cross section of states, there is a strong negative relationship between prices and economic activity as would be suggested by typical New-Keynesian demand shocks. However, aggregate prices and aggregate nominal wages did not fall. This suggests some other aggregate supply side shock must have occurred to put upward pressure on aggregate prices and wages. Our procedure quantifies how important these shocks were and distinguishes between the productivity and labor supply shock with the aid of our model.

Lastly, we perform many additional robustness exercises and extensions. Across our various robustness exercises, the qualitative conclusion that aggregate supply shocks were important in explaining the missing deflation in aggregate prices and wages. Additionally, we use the model structure, local data and the aggregate VAR to infer the local shocks. We find that while the supply shock was important for explaining the aggregate evolution of prices and employment, it is essentially only regional differences in demand shock that explain the cross region variation. Finally, we compare our VAR approach to other standard VAR approaches using aggregate data in the literature and show there are quantitative differences across the procedures.

We want to stress that a limitation of our analysis is that we do not point to what specific "demand" specific shock or "supply" specific shock drove the Great Recession. For example, we cannot distinguish between a tightening of household borrowing constraints vs. households wanting to deleverage. Despite that, we think our conclusions are important in the extent that we quantify the relative importance of broad types of shocks. One take away from our work is that something akin to a supply shock was at play in the aggregate economy during the Great Recession and that shock had very little variation across regions. This finding will hopefully guide researchers to focus on exploring the origins of these broad shocks in future research. The results also suggest that only using cross-region variation to explain aggregate fluctuations is insufficient

when the aggregate shocks do not have a substantive regional component.

Our paper contributes to many additional 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. 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. (2014) use regional variation to measure the responsiveness of labor markets to labor demand shocks. Our work contributes to this literature on two fronts. First, we show that local prices also respond to local changes in economic conditions. 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.

Second, our paper contributes to the recent literature highlighting that supply shocks were important for explaining aggregate fluctuations during the Great Recession. For example, Christiano et al (2014) estimate a New Keynesian model using data from the recent recession. Although their model is 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. Gilchrist et al. (2014) show that liquidity constraints facing firms can result in firms cutting employment and raising prices. Our work complements these papers by using regional variation to help to parameterize the aggregate economy. The fact that prices and wages move with economic conditions at the local level help to discipline how aggregate prices and wages should have moved if only demand shocks were driving aggregate fluctuations.

Third, there is some recent work using scanner data to explore the relationship between local economic conditions and prices.⁶ Contemporaneously, Coibion et al. (2014) use data from Symphony IRI to examine regional variation in prices during the 2000-2011 period. The main focus of that paper is to examine the nature of household shopping behavior in response to changes in local economic activity. Kaplan and Menzio (2014) use data from Nielsen's Homescan data to examine how the variance of prices paid change with economic conditions. They find that within a given market and a given time period, there is a large difference in prices paid for a given product. They

⁶There was an older literature that used scanner data to create price indices for a particular good. See, for example, Hawkes and Piotrowski (2003), Richardson (2003), and Lowe and Ruscher (2003) create scanner price indices for, respectively, ice cream, breakfast cereal, and televisions. Additionally, others have used scanner data to create price indices for different groups. For example, Aguiar and Hurst (2007) and Broda, Leibtag and Weinstein (2009) use scanner data to produce price indices for individuals of different ages and incomes, respectively. Broda and Weinstein (2010) and Handbury, Wantanabe, and Weinstein (2013) use scanner data to quantify biases in government provided price indices. Finally, Handbury and Weinstein (2011) use the Homescan data to examine persistent pricing differences across U.S. locations. In their analysis, they find that prices paid for a given good do not systematically differ across different regions. While it may be true that regions do not have persistently different prices on average, we document that local prices do move with business cycle frequencies.

conclude that only a small portion of the price variability in a market time period is due to some stores being persistently more expensive than others. Stroebel and Vavra (2014) use the IRI data to explore the relationship between house prices and retail prices. They conclude that increasing house prices cause retail prices to increase. They provide evidence that mark-ups change in response to changes in housing wealth. Finally, Fitzgerald and Nicolini (2014) use data from the 27 MSA level price indices published by the BLS to create MSA level Phillips curves. Consistent with our findings, they also show a negative relationship between inflation and unemployment at the MSA level that holds historically. Our paper complements this literature by actually making price indices using scanner data for each state at the monthly frequency for each state. We post these indices so other researchers can use them in their research going forward.

2 Creating State Level Price Indices

2.1 Data

We begin by constructing state-level price indexes using the *Retail Scanner Database* collected by AC Nielsen and made available at The University of Chicago Booth School of Business.⁷ The Retail Scanner data consists of weekly pricing, volume, and store environment information generated by point-of-sale systems for about 90 participating retail chains across all US markets between January 2006 and December 2011. 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. The stores in the database vary in terms of the channel they represent: food, drug, mass merchandising, liquor, and convenience stores. 97 percent of the sales in the data come from food, drug and mass merchandising stores. It is important, however, to point out that one of the most important low-price retailers, Walmart, has only recently agreed to share their data with Nielsen and therefore all of their stores are excluded from the database to which we have access.

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

⁷The 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/>.

since they do not have a UPC code assigned. Nielsen sorts the different UPCs into over one thousand narrowly defined "categories".⁸ 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 (defined below). We will first aggregate prices to a category level and then compute the price index aggregating across categories.

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 371 Metropolitan Statistical Areas. The data comes with both zip code and FIPS codes for the store's county, MSA, and state. In this paper, we aggregate data to the level of U.S. states and compute state level scanner data price indices. In future iterations, similar indices can be made at the MSA level.

Table 1 shows summary statistics for the scanner data for each year between 2006 and 2011 and for the sample as a whole. A few things are of particular note. The sample sizes - in terms of stores covered - increased from 31,642 stores (in 2006) to 35,645 stores (in 2011). Much of this is due to Nielsen creating additional contracts with more chains. As seen from the table, the number of chains in the sample increased from 73 chains in 2006 to 86 chains in 2011. Most of this increase occurred between 2006 and 2007. Second, notice the number of observations (store*week*UPC code) is massive. The database includes nearly *12 billion* unique observations. Third, during the entire sample, there is about 1.4 million unique UPC codes within the database. On average, each year contains roughly 760,000 UPC codes. Fourth, the geographic coverage of the database is substantial in that it includes stores for 81 percent of all counties within the United States. Moreover, the number of geographical units (zip codes, counties, MSAs, states) is very similar from year to year highlighting that the geographical coverage is consistent through time. Finally, the dataset includes between \$184 billion and \$236 billion of transactions within each year. For the time periods we study, this represents roughly 30 percent of total U.S. expenditures on food and beverages (purchased for off-premise consumption) and roughly 2 percent of total household consumption.⁹

2.2 A Scanner Data Price Index

Our goal is to construct regional price indices from the scanner data that is similar in spirit to how the BLS constructs the CPI.¹⁰ Our scanner price indices are built in

⁸The total number of categories in the database is 1092, however only 1075 have information for the entire time period 2006-2011.

⁹To make these calculations, we compare the total transaction value in the scanner data to BEA reports of total spending on food and beverages (purchased for off-premise consumption) and total household consumption. See NIPA Table 2.5.5. Additionally, the Online Appendix which accompanies this paper provides coverage statistics (volume, number of stores, number of UPC's) for every state by year.

¹⁰The BLS does publishes twenty-seven metro area price indices semi-annually. We explore the patterns in these data below and show that they are very similar to the patterns we document for a large number of U.S. sub-regions using our scanner data index.

two stages.¹¹ In the first stage, we aggregate prices of goods within the roughly 1,000 categories described above. For our base index, a good is either a UPC or a store-UPC pair. In the latter case, a two liter bottle of Coke sold in store A is treated as a different good than a two liter bottle of Coke sold in store B. We do this to allow for the possibility that prices may change as households substitute from a high cost store (that provides a different shopping experience) to a low cost store when local economic conditions deteriorate.¹² For each state, within each detailed category (sugar granulated, sugar powdered, etc.), we find the quantity weighted average price for all goods (UPC or UPC-store pair) within a given month. We then compute for each good the average price and total quantity sold for the month. We aggregate our index to the monthly level to reduce the number of missing values.

Formally, the first step is to produce a category-level Laspeyres price index which can be expressed as follows:

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

where $P_{j,t,y,k}^L$ is Laspeyres price index for category j , in year t , with base year y , in geography k . For our analysis, geographies will either be U.S. states or the country as a whole. $p_{i,t,k}$ is the price at time t of the specific good i in geography k and $\bar{q}_{i,t-1,k}$ is the average monthly quantity sold of good i in the prior year in location k . By fixing quantities at their prior year's level, we are holding fixed household's consumption patterns as prices change. We update the basket of goods each year, and chain the resulting indices to produce one chained index for each category in each geography, denoted by $P_{j,t,k}^L$. In this way, the index for months in 2007 uses the quantity weights defined using 2006 quantities and the index for months in 2008 uses the quantity weights defined using 2007 quantities. This implies that the price changes we document below with changing local economic conditions is not the result of changing household consumption patterns. Fixing the basket also minimizes the well documented chain drift problems of using scanner data to compute price indices (Dielwert et al. (2011)). Notice, this procedure is very similar to the way the BLS builds category-level first stage for their price indices.

When computing our monthly price indices, one issue we confront is how to deal with missing values from period to period. For example, a product that shows up in month m may not have a transacted price in month $m + 1$ making it impossible to

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

¹²In practice, controlling for store effects had little effect on our price indices. However, the possibility that store effects can move local prices was discussed prominently in Coibion et al (2012). For completeness, we constructed our price indices allowing for store effects in pricing.

compute the price change for that good between the two months. Missing values may be due to new products entering the market, old products withdrawing from the market, and seasonality in sales. Our results in the paper were robust to how we deal with missing values but clearly the price indices will generally differ depending on how one treats such data points. Although we could have used some ad hoc imputation methods like interpolation between observed prices or keeping a price fixed until a new observation appears, we chose to follow a more conservative approach. Looking at equation (1), we see that we can handle the missing values without imputation by restricting the goods that enter the basket to those that have positive sales over at least one month in the previous year and over the 12 months of the current year. This is what we do when creating our indices. For example, when computing the category prices in 2008 we use the reference basket for 2007. In doing so, we only take the goods that have $q_{i,2007,k} > 0$ and $q_{i,t,k} > 0$ for all $t \in 2008$.¹³ This ensures that for a given product in the price index during year t , we will have a weight for this product based on $t - 1$ data and we will have a non-missing transaction price in all months in which the price index is computed during that year.¹⁴ The bottom row of Table 1 includes the share of all expenditures (value weighted) that were included in our price index for a given year. In the four later years of the sample, our price index includes roughly two-thirds of all prices (value weighted).

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

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

where $\bar{s}_{j,k}^t$ is the share of expenditure of category j in month t in location k averaged over the year. We calculate the shares using total expenditure on all goods in each category, even though for the category-level indices some goods were not included due to missing data. For the purposes of this paper, we make our baseline specification one that fixes the weights of each category for a year in the same fashion as we did for the category-level indices. However, as a robustness specification, we allowed the weights in the second step to be updated monthly. The results using the two methods were nearly identical.

Figure 3 shows that our scanner price index matches nearly identically the BLS's Food CPI. We chose the BLS Food CPI as a benchmark given that most of the goods

¹³The database starts in 2006. As a result, our baseline specification of the 2006 price indices only includes products that have positive sales in all months of 2006.

¹⁴This procedure implies that we will miss products that are introduced within a given year. These products, however, will be incorporated in next year's basket as long as they have continuous sales during the subsequent calendar year.

in our database are food data.¹⁵ For ease of comparison, we normalize both our index and the BLS Food CPI to 1 in January of 2006. Notice that the inflation rate between January 2006 and January of 2009 is close to identical between our index and the BLS's food index at 12.0 percent and 12.1 percent, respectively. Prices in both indices fall through mid 2009 and then both indices show a rise in prices after that. The fact that our price index matches the BLS Food CPI so closely suggests that the underlying data in our database is broadly representative of the goods included in the BLS's Food CPI. This gives us confidence that we will be able to create meaningful CPI's at the local level for the grocery/mass-merchandizing products included in our data.¹⁶

2.3 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 distribution costs, one would expect little variation in retail prices across regions if retail goods were purely tradable. However, local distribution costs include the wages of workers in the retail establishments, the rent of the retail facility, and expenses associated with local warehousing. These local distribution costs represent the non-traded component of the retail goods in our sample.

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 prices (P^r) in region k during period t as:

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

where P_t^T is the tradable component of local retail prices in period t and does not vary across regions and $P_{t,k}^{NT}$ is the non-tradable component of local retail 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.

Analogously, we can define similar price indices for all sectors s within the economy.

¹⁵Not 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, 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.

¹⁶As noted above, our scanner price index, in practice, differs from the BLS's Food CPI given that we are allowing the quantity weights to be updated more frequently. This could explain some of the difference between the two surveys prior to 2009 and post 2011. Given the short time series of our sample, it is hard to evaluate this conjecture systematically. But, our results do suggest that such bias that may result from the BLS using lagged data to define the quantity weights is small.

Specifically, we assume that each sector within a region has a price index that combines tradable and local non-tradables using a different weight α_s such that:

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

Again, we assume that α_s is constant across all regions for all s . α_s differs from α_r to the extent that the share of tradables in the grocery/mass merchandizing retail sector differs from the share of tradables in sector s . For example, one could imagine that α_s is much larger than α_r for many local services that almost exclusively use local labor and local land in the production of their retail activities (e.g., dry-cleaners, haircuts, education services, and restaurants). For other sectors, α_s could be lower than α_r if the traded component is large relative to the local factors used to sell the good. This may be true, for example, for car dealerships. What we are interested in is the traded and non-traded component of the typical good in the household's consumption basket. Suppose that the composite good in a region can be expressed such that:

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

where $\bar{\alpha}$ is the non-tradable share for the composite good in a region. Given these assumptions, we can transform the variation in the grocery/mass-merchandising 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 k, k' ($\Delta \ln P_{t,k,k'}$) is proportional to the variation in log-grocery retail prices across those same regions ($\Delta \ln P_{t,k,k'}^r$). Formally,

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

With knowledge of α_r and $\bar{\alpha}$ we can make such an adjustment. Burstein, Neves and Rebelo (2003) document that distribution costs represent more than 40 percent of retail prices in the United States. Industry analysts report the grocery industry in the U.S. has a gross margin of 25-30 percent suggesting that local distribution costs are a significant component of costs. When converting the variation in local retail prices into local non-tradable prices, we use an estimate of $\alpha_r = 0.3$. This implies that 30 percent of retail prices are attributed to non-traded inputs. This is on the upper end of industry reports but lower than the findings of Burstein, Neves and Rebelo. For $\bar{\alpha}$, we looked for an estimate of the share of total local consumption at the state level that is imported from outside the state. Assuming that all housing consumption is locally consumed, our estimate of $\bar{\alpha}$ should exceed 0.2 (the share of housing services out of total consumption). Based on the work of Nakamura and Steinsson (2014), we use an estimate of 0.6. In that paper, Nakamura and Steinsson measure the fraction of output in a U.S. region that is imported from other U.S. regions.¹⁷ Putting the two estimates together, we adjust the

¹⁷The level of analysis in Nakamura and Steinsson (2014) is U.S. regions. They define 10 regions - the nine Census divisions where they segment the "South Atlantic" division into two regions. Their estimate

variation in the regional inflation rates computed using the goods in our database by a factor of 2 (0.6/0.3).¹⁸

3 Regional Variation in Prices and Wages During the 2000s

3.1 Regional Variation in Prices During the 2000s

Figure 4 and Table 2 explore the extent to which our regional scanner price index is correlated with measures of local economic activity. Specifically, Figure 4 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 2010.¹⁹ For the results in Figure 4, we use our price index where a good is a given UPC within a state (as opposed to a UPC-store pair). Additionally, Figure 4 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 represent 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 4, $\Delta X_{k,07-10}$ equals the percentage point change in the state unemployment rate between 2007 and 2010. In Table 2, we run this regression for a variety of different measures of changes in state economic activity.

Figure 4 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 -0.46 with a standard error of 0.14 (and an adjusted R-squared of 0.18). This implies that cumulative retail price inflation between 2007 and 2010 was 1.84 percentage points higher in states with a change in the unemployment rate of 6 percentage points during that same time

of $\bar{\alpha}$ is 0.69. Given their unit of analysis is larger than a U.S. state, their estimate should be seen as an upper bound on the nontraded share of local consumption at the state level. Given this, we choose our estimate of $\bar{\alpha} = 0.6$. We have redone all the results in the paper using an estimate of $\bar{\alpha} = 0.69$ and none of the results change in any meaningful way.

¹⁸The adjustment factor places a minimal role in our quantitative work below. Our base assumption is that the grocery/mass-merchandising sector has a larger tradable share than the average good. We have explored a variety of adjustment factors between 1.5 and 3 and the quantitative implication of our estimation procedure in Section 6 were relatively robust.

¹⁹Our 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.

period relative to states with an unemployment rate of 2 percentage points. Given our discussion above, the responsiveness of regional differences in retail prices for the grocery/mass-merchandising sector may be muted relative to the responsiveness of the composite local consumption good given the relatively high tradable share of costs in these sectors. Scaling the regional variation by our scaling factor, we find that a one percentage point increase in the state unemployment rate is associated with a fall in local prices -0.92 percent ($-0.46 * 2$).

Table 2 shows different estimates of β_1 from the above regression with different measures of changing local economic activity ($\Delta X_{k,07-10}$). For each measure, we show the results for our price index where a good is defined as UPC within a state (columns (1) and (3)) and for our price index where a good is defined as a UPC-store pair within a state (columns (2) and (4)). Panel A measures the variation for our retail grocery and mass-merchandising goods. Panel B shows the results for our composite good which is just a scaled version of the coefficients in Panel A. Each row in Table 2 is a different measure of the changing economic conditions within the state. For example, the first row is the change in the BLS unemployment rate in the state (analogous to the results in Figure 4). 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.²⁰ Additionally, to isolate the effect of housing price changes working through the change in the unemployment rate and the percent change in the employment rate, we estimate a version of the above regression instrumenting the change in the unemployment rate or the growth in the employment rate with the change in local house prices. Aside from isolating the part of the change in the unemployment (employment) rate that is associated with changing house prices, it also may change the coefficient by correcting for measurement error in the local economic measures.

As seen from the results in Table 2, 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. Defining goods at the UPC-store level (columns 2 and 4) only mitigates slightly the underlying relationships when we only define goods at the UPC level (columns 1 and 3). Isolating the part of the change in the unemployment rate due to changing housing market conditions does not alter at all the relationship between unemployment changes and price changes. However, isolating the part of the change in the employment rate that is due to changing housing market conditions strengthens the coefficients on the employment rate change.

Figure 5 allows for a comparison of the timing of the price changes within the states relative to when the unemployment rate change occurred. The results in Figure 4

²⁰The 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.

compared long differences in both the unemployment rate and the inflation rate. In Figure 5, we can exploit the monthly nature of our data. For ease of exposition, we group all states into three groups. The first group includes the top one third of states based on the change in the unemployment rate between 2007 and 2010. This group includes Nevada, California and Florida (among others). We refer to this group as the "high unemployment change states". The second group includes the bottom one third of states based on the change in the unemployment rate between 2007 and 2010. This group includes Texas and Massachusetts (among others). We refer to this group as the "low unemployment change states". The third group includes the remaining states.

In Figure 5, we plot two separate lines. The solid line is the unemployment difference between the low and high unemployment change groups of states (weighted by state population within each group). As seen from Figure 5, the unemployment rate between the low and high unemployment change states started opening up in mid 2007 and by mid 2009 had stabilized. After the recession ended, the relative unemployment rate between the high and low unemployment rate states remained relatively constant through 2011. The dashed line in Figure 5 is the difference between the average retail price level in the low unemployment change states relative to the high unemployment change states.²¹ The differences in retail prices between low and high unemployment change states is essentially the mirror image of the differences in the unemployment rates between the high and low unemployment change states. As the unemployment rates diverged, the prices also quickly diverged. When the unemployment differences stabilized, the price differences quickly stabilized. For example, during 2011, there was a slight narrowing of the unemployment gap across the two groups of states. At exactly that same time, there was a narrowing of the price levels between the two groups of states. The simple correlation between the two series in Figure 5 is -0.93. There is nothing in our procedure that ensured we would find such a correlation. The unemployment data comes from the BLS's Local Area Unemployment Statistics. The price level data was created using the scanner data.

The second thing to note from Figure 5 is that prices appear to adjust rather quickly as the unemployment rate changed. There does not seem to be a delay between when the unemployment rate changed and when the prices changed. Moreover, once the unemployment rate stabilized between the two groups of states, there was no further change in the price level differences. The speed at which prices adjusted will be in sharp contrast to the speed in which nominal wages adjusted. This will motivate our theoretical model where we focus on wage stickiness rather than price stickiness.²²

In the online appendix that accompanies the paper, we discuss how our price indices compare with the MSA level price indices produced by the BLS. For the most part,

²¹To get the difference in the composite consumption good between the low and high unemployment change states, we would just multiply the dashed line by 2.

²²The differences in the price inflation between high and low unemployment change states is highly statistically significant. To save space, we do not discuss the procedure by which we computed standard errors in the text. However, in the online appendix that accompanies this paper, we reproduce Figure 5 with standard errors and discuss the method by which they were computed.

the BLS does not put out sub-national price indices. However, for 27 metro areas price indices are reported monthly (3 MSA's), bimonthly (11 MSA's), and semi-annually (13 MSA's).²³ Overall, we concluded two things from our analysis of the 27 BLS MSA level price indices. First, the patterns we find in our scanner data broadly extend to the BLS data at the MSA level. MSAs (states) with weaker local economic conditions during the recession also had lower food price growth during that same time period. However, we needed to aggregate the BLS data into groups of MSAs based on the state's unemployment change to find statistically significant patterns. Relative to our data, there appears to be more noise in the BLS local price measures. Aside from having measures for all states, we feel the scanner data allows for reduced measurement error in the local price indices. Second, and potentially more importantly, the BLS data allowed us to examine whether the patterns we document for food data extend to broader measures of the CPI. Using the BLS data, we found that the relationship we document between local economic activity and food prices also show up between local economic activity and broader measures of the local CPI. This further gives us confidence that our results are not limited to the goods covered within the scanner database.

In summary, despite there not being any strong relationship between price growth and changes in economic activity at the aggregate level during the recession, there is a strong relationship between prices and real activity across U.S. states during this time period.

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

We now explore how both nominal and real wages evolved across states during the 2000s. To make nominal wages at the state level, we use data from the 2000 U.S. Census and the 2001-2012 American Community Surveys (ACS). The 2000 Census includes 5 percent of the U.S. population. The 2001-2012 ACS's include around 600,000 respondents between 2001-2004 and around 2 million respondents after 2004. The large sample sizes allows for detailed labor market information at the state level. We begin by using the data to make individual hourly nominal wages. We restrict our sample to only those individuals 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. For each individual, we divide total labor income earned during the prior 12 months by a measure of annual hours worked during prior 12 months.²⁴ The composition of

²³The 27 MSAs are (in order of reporting frequency): Chicago (CMSA), LA (CMSA), New York (CMSA), Atlanta (MSA), Boston (MSA), Cleveland (CMSA), Dallas-Fort Worth, Detroit (CMSA), Houston (CMSA), Miami (CMSA), Philadelphia (CMSA), San Francisco (CMSA), Seattle (CMSA), Washington (CMSA), Anchorage (MSA), Cincinnati (CMSA), Denver (CMSA), Honolulu (MSA), Kansas City (MSA), Milwaukee (CMSA), Minneapolis (MSA), Phoenix (MSA), Pittsburgh (MSA), Portland, OR (CMSA), St. Louis (MSA), San Diego (MSA), and Tampa (MSA). The Phoenix series only started in 2002. For that reason, we start our regional BLS CPI sample in 2002.

²⁴Total 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 months is the multiple of total weeks worked during the prior 12 months and the respondents report of their usual hours worked per week. In

workers differs across states and within a state over time which could explain some of the variation in nominal wages across states over time. To account for this, we run the following regression:

$$\ln(w_{itk}) = \gamma_t + \Gamma_t \tilde{X}_{itk} + \eta_{itk}$$

where $\ln(w_{itk})$ is log nominal wages for household i in period t residing in state k and \tilde{X}_{itk} is a vector of household specific controls. The vector of controls include a series of dummy variables for usual hours worked (30-39, 50-59, and 60+), 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, Γ_t , can differ for each year. We then take the residuals from these regressions for each individual, 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 .

Figure 6 shows the cross state variation in log demographic adjusted nominal wages 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 strongly correlated with changes in the unemployment rate during the 2007-2010 period. A simple linear regression through the 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 1.23 percentage point decline in nominal wage growth (standard error = 0.21). In Table 3 (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.

Figure 7 shows a demographic adjusted nominal wage index for high and low unemployment change states over the 2000 to 2012 period. The high and low unemployment change states are defined similarly as in Figures 4. The high and low unemployment change states had roughly similar growth in nominal wages between the 2000 and 2007 period. However, after 2007, the nominal wages started to diverge. One key insight from Figure 7 is that after 2010, the nominal wages between high and low unemployment states continued to diverge. This result occurred despite the fact that the difference in unemployment rates (and the growth in GDP per capita and employment per capita) stabilized during this period. This result will be important when calibrating our model below in that it is this feature of the data that will help discipline our estimates of wage

some years, bracketed reports are provided for the weeks worked during prior 12 months and the usual hours per week worked. In those cases, we take the mid point of the brackets.

stickiness.

Finally, the second column of Table 3 shows 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. We measure real wage growth in each state as the nominal wage growth less the scaled up inflation rate for the composite good (allowing for store effects). Not surprising, the coefficients in column 2 of Table 3 are roughly equal to the coefficients from column 1 of Table 3 less the coefficients from column 4 of Table 2. 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 0.5 percent decline in real wage growth during the 2007 to 2010 period. Unlike the results for prices and nominal wages, most of the coefficients in column 2 of Table 4 are only marginally significant at standard levels.

4 Using Regional Information to Identify Aggregate Shocks

As seen from the prior section, both wages and prices co-vary with economic conditions in the cross-section. Despite the relatively little movement in wages and prices at the aggregate level during the Great Recession, there is much interesting variation to exploit in the cross-section. In this section, we now turn to a more formal exploration of these cross-sectional patterns and develop a methodology for their use, together with aggregate time series data, to learn about the underlying shocks driving both aggregate and local business cycles. The methodology that we propose is a relatively general econometric procedure that allows one to incorporate cross-sectional information to a structural VAR describing the evolution of aggregate variables of interest and, under some assumptions, for the identification of structural shocks.

The rapid expansion in micro-data availability created many opportunities for researchers to engage macroeconomic questions where heterogeneity and distributional issues are important.²⁵ By and large, the way cross-sectional information is used, follows a model based approach: write down a model with a proposed mechanism and show patterns in the data that are indicative of the mechanism either because these patterns are in line with some of the model's implications and/or support certain assumptions. Finally, the model is calibrated/estimated using micro-data and simulations are conducted to discuss the aggregate/macroeconomic question of interest. Our methodology provides an alternative to this approach that places less restrictions in the aggregate behavior of the variables of interest than a fully specified theoretical model while still incorporating micro-data and theoretical restrictions in the cross-section to discipline it. The cost, however, is the range of questions that can be addressed and the use of linear approximations.

²⁵Leading examples include the literature on precautionary savings and incomplete markets, e.g. Aiyagari (1994), Krusell and Smith (1998); or nominal price rigidities, e.g. Bils and Klenow (2002), Alvarez, Beraja, Neumeyer and Gonzalez-Rozada (2014).

These limitations apply to structural VARs more generally, regardless of the identification scheme, e.g. ordering, sign-restriction, long-run restrictions, narratives, or theoretical impulse response matching. Nonetheless, our methodology has mainly two advantages in comparison: (i) it incorporates not only aggregate data, but regional data as well; and (ii) it places identification restrictions in the behavior of cross-sectional of variables and not the aggregate variables/shocks of interest. In this last sense, the methodology is related to the “hybrid” literature 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).

We start by providing a model of regional economies within a monetary union. In the model we outline, each local economy has something akin to a labor demand curve, a labor supply curve, and an Euler equation. The goal of the model section is twofold. First, it highlights the assumptions needed such that the local labor supply (demand) curves aggregate to an aggregate labor supply (demand) curve. Second, it shows that the aggregate economy can be represented as a VAR. In the subsequent section, we outline our semi-structural approach to estimating the VAR. We show that if we know the parameters of either the aggregate labor supply curve, the aggregate labor demand curve, or the aggregate Euler equation, the aggregate VAR can be fully identified with no additional restrictions. The key is that one needs to only know the parameters of one of the aggregate structural equations. We then show that under certain assumptions, the regional data can be used to parameterize either the aggregate labor supply equation or the aggregate labor demand equation. Finally, we will show how we can use the regional data and the results of the aggregate VAR to infer the underlying shocks driving the local economy.

4.1 A Model of a Monetary Union

The 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.²⁶ Finally, we assume labor is mobile across sectors but not across islands.²⁷ Throughout we assume that parameters governing preferences and production are identical across islands and these only differ, potentially, in the shocks that hit them.

²⁶We can think of the final good as being retail: restaurants, barbershops and stores; and the intermediate sector providing physical goods: food ingredients, scissors and cellphones.

²⁷We freely admit this is a rather extreme assumption. One that we make for tractability and ease of exposition which is the purpose of this model section: to illustrate a simple economy where our methodology can be applied. Whether it is easier for workers to move across industries/occupations than space is certainly debatable and depends on the time frame and type of worker that one might have in mind. For evidence, the interested reader should look into...

4.2 Firms and Households

Tradable intermediates x producers in island i use local labor N_i^x and face nominal wages W_i (equalized across sectors) and prices Q (equalized across islands k). Their profits are

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

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

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

where z_i^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.

Households preferences are given by

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

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. Households are able to spend their labor income $W_{kt} N_{kt}$ plus profits accruing from firms Π_{kt} and financial income $B_{kt} i_t$, where B_{kt} are nominal bond holdings at the beginning of the period and i_t is the nominal interest (equalized across islands given our assumption of a monetary union where the bonds are freely traded) on consumption goods (C_{kt}) and savings ($B_{kt+1} - B_{kt}$). Thus, they face the period-by-period budget constraint

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

4.3 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, typically, workers will be off their labor sup-

ply curves. A similar specification has been used recently, for example, in Midrigan and Philippon (2011). Popular alternatives in the literature include the wage bargaining model in the spirit of Hall and Milgrom (2008) as in Christiano, Eichenbaum and Trabandt (2013) or the monopsonistic competition model where unions representing workers set wages period by period as in Galí (2009). The key difference with the partial adjustment model is that both alternatives result in a forward looking component in the wage setting rule. In fact, the partial adjustment wage setting rule can be derived from the monopsonistic competition setup in the case where agents are myopic about the future.

4.4 Equilibrium

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

$$\begin{aligned} C_{kt} &= e^{z_{kt}^y} (N_{kt}^y)^\alpha X_{kt}^\beta \\ N_{kt} &= N_{kt}^y + N_{kt}^x \\ \sum_k X_{kt} &= \sum_k e^{z_{kt}^x} (N_{kt}^x)^\theta \\ \sum_k B_{kt} &= 0 \end{aligned}$$

4.5 Aggregation

Our first key assumption for aggregation is that that all islands are identical with respect to their underlying production parameters (α , β , and θ), their underlying utility parameters (σ , ϕ and ρ) and the degree of wage stickiness (λ).²⁸ Our second assumption is that the islands are identical in the steady state and that price and wage inflation are zero. 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.²⁹ We denote with low-

²⁸When implementing our procedure empirically using data on US states, we discuss the plausibility of this assumption. Given that the broad industrial composition at the state level does not differ much across states, the assumption that productivity parameters and wage stickiness are roughly similar across states is not dramatically at odds with the data. As a robustness exercise, we estimate our key equations with industry fixed effects and show that our key cross section estimates are unchanged. We are comfortable with the assumption that preference parameters are constant across states.

²⁹The 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,

ercase letters the log-growth rates of variables. Also, variables without an k subscript represent aggregates. For example, $n_{kt} \equiv \log \left(\frac{N_{kt}}{N_{kt-1}} \right)$ and $n_t \equiv \sum_k \frac{1}{K} n_{kt}$. We assume that the monetary authority announces a nominal interest rate rule which in log-linearized form can be written as: $i_{t+1} = \varphi_p p_t + \varphi_n n_t + \mu_{t+1}$ where μ_{t+1} is an exogenous stochastic process. The following lemmas present a useful aggregation result and show that we can write the island level equilibrium in deviations from these aggregates.

Lemma 1 *The behavior of p_t, w_t, 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 $\alpha + \theta\beta$ and only 3 exogenous processes $\{z_t, \epsilon_t, \gamma_t\}$ where $z_t = z_t^y + \beta z_t^x$ and $\gamma_t = \delta_t - \delta_{t-1} + \mu_t$.*

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. Also, with a slight abuse of notation, define $\tilde{z}_t^x \equiv \beta(z_{kt}^x - z_t^x)$.

Lemma 2 *For given $\{\tilde{z}_t^y, \tilde{z}_t^x, \tilde{\gamma}_t, \tilde{\epsilon}_t\}$, the behavior of $\{\tilde{p}_t, \tilde{w}_t, \tilde{n}_t, (\tilde{n}_t^x - \tilde{n}_t^y)\}$ 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. For a given sequence of aggregate variables $\{p_t, n_t, q_t\}$, the following equations characterize the log-linearized equilibrium

$$\begin{aligned}
w_{kt} &= \lambda(p_{kt} + \epsilon_{kt} - \epsilon_{kt-1} + \frac{1}{\phi} n_{kt}) + (1 - \lambda)w_{kt-1} \\
w_{kt} &= p_{kt} - (1 - (\alpha + \theta\beta))n_{kt}^y - \beta(1 - \theta)(n_{kt}^x - n_{kt}^y) + z_{kt}^y - z_{kt-1}^y + \beta(z_{kt}^x - z_{kt-1}^x) \\
0 &= \mathbb{E}_t (mu_{kt+1} + \delta_{kt+1} - \delta_{kt} - p_{kt+1} + \varphi_p p_t + \varphi_n n_t + \mu_{t+1}) \\
mu_{kt+1} &= -\frac{\sigma}{C - N^{\frac{1+\phi}{\phi}}} \left(Cc_{kt+1} + N(\epsilon_{kt+1} - \epsilon_{kt} + \frac{1+\phi}{\phi} n_{kt+1}) \right) \\
Nn_{kt} &= N^x n_{kt}^x + N^y n_{kt}^y \\
c_{kt} &= z_{kt}^y - z_{kt-1}^y + (\alpha + \beta) \left(\frac{N}{N^y} n_{kt} - \frac{N^x}{N^y} n_{kt}^x \right) - \beta(q_t - w_{kt}) \\
w_{kt} &= z_{kt}^x - z_{kt-1}^x + q_t - (1 - \theta)n_{kt}^x \\
x_{kt} &= n_{kt}^y + w_{kt} - q_t \\
\sum_k x_{kt} &= \sum_k (z_{kt}^x - z_{kt-1}^x + \theta n_{kt}^x)
\end{aligned}$$

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.

From the last 3 equations, after adding up, it holds that $n_t^x = n_t^y$. Then the aggregate log-linearized equilibrium evolution of $\{p_t, w_t, n_t\}$ is characterized by

$$\begin{aligned}
0 &= \mathbb{E}_t(mu_{t+1} - p_{t+1} + \gamma_{t+1}) + \varphi_p p_t + \varphi_n n_t \\
w_t &= \lambda(p_t + \epsilon_t - \epsilon_{t-1} + \frac{1}{\phi} n_t) + (1 - \lambda)w_{t-1} \\
w_t &= p_t - (1 - (\alpha + \theta\beta))n_t + z_t - z_{t-1} \\
mu_{t+1} &= -\frac{\sigma}{C - N^{\frac{1+\phi}{\phi}}} \left(C(z_{t+1} - z_t + (\alpha + \theta\beta)n_{t+1}) + N(\epsilon_{t+1} - \epsilon_t + \frac{1+\phi}{\phi}n_{t+1}) \right)
\end{aligned}$$

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 $\alpha + \theta\beta$ and only and only 3 exogenous processes $\{z_t, \epsilon_t, \gamma_t\}$. The top equation is the aggregate Euler equation. The second equation is effectively the aggregate labor supply curve. 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 $\{\tilde{p}_t, \tilde{w}_t, \tilde{n}_t, (\tilde{n}_t^x - \tilde{n}_t^y)\}$ for given $\{\tilde{z}_t^y, \tilde{z}_t^x, \tilde{\gamma}_t, \tilde{\epsilon}_t\}$,

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

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. ■

4.6 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. Formally,

$$\begin{aligned} z_{kt}^y &= \rho_z z_{kt-1}^y + u_t^y + v_{kt}^y \\ z_{kt}^x &= \rho_z z_{kt-1}^x + u_t^x + v_{kt}^x \\ \gamma_{kt} &= \rho_\gamma \gamma_{kt-1} + u_t^\gamma + v_{kt}^\gamma \\ \epsilon_{kt} &= \rho_e \epsilon_{kt-1} + u_t^\epsilon + v_{kt}^\epsilon \end{aligned}$$

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

Let $u_t^z \equiv u_t^y + \beta u_t^x$. We will call u_t^z , u_t^γ and u_t^ϵ the aggregate *Supply, Demand and Leisure shocks* respectively. These are the innovations that the econometric procedure aims to identify. Analogously, $v_{kt}^y, v_{kt}^x, v_{kt}^\gamma, v_{kt}^\epsilon$ are the *Regional shocks*. The interpretation of the Leisure and Supply shocks is relatively straightforward given our model environment. They are shifters of households and firms' labor supply and demand schedules respectively. On the other hand, what we identify as a Demand shock 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 and, hence, we treat it as a single shock.

5 Estimation Procedure

5.1 Estimating the Aggregate Shocks

The system of equations in Lemma 1 can be written in reduced form as a VAR(∞).³⁰

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

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 Supply, Demand and Leisure shocks. Hence, identification of the shocks is identification of these matrices.³¹

As is well known in the literature, estimation of the above VAR is not possible without additional identifying restrictions. For example, identification is often achieved via imposing ordering, sign, or long-run restrictions. Our identification approach is to use theory to specify one equation of the VAR. If the one equation is specified, we can identify the aggregate shocks while being agnostic about the remaining structural equations

³⁰The 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(2,2) and hence as a VAR(∞).

³¹When we implement our procedure we will truncate $\rho(L)$ to be of finite order as is typically done in the literature.

driving the economy.

To provide intuition, suppose we knew that either (1) the aggregate labor supply curve is specified according to the above model such that $w_t = \lambda(p_t + \frac{1}{\phi}n_t) + (1 - \lambda)w_{t-1} + \lambda e_t$ with $e_t \equiv u_t^\epsilon - (1 - \rho_\epsilon)\epsilon_{t-1}$ or (2) the aggregate labor demand is specified according to the above model such that $w_t = p_t - (1 - (\alpha + \theta\beta))n_t + \eta_t$ with $\eta_t \equiv u_t^z - (1 - \rho_z)z_{t-1}$. If we knew the specification of the above labor supply curve and if we had estimates of λ and ϕ , $\hat{\lambda}$ and $\hat{\phi}$, we can then use aggregate data and the estimated aggregate labor supply curve to compute an instrument for the aggregate labor supply shock.³² Note, $\hat{\lambda}$ is an estimate of the aggregate wage stickiness and $\hat{\phi}$ is an estimate of the aggregate labor supply elasticity. Specifically, we can use aggregate data to define a variable s_t such that:

$$\hat{s}_t = \frac{1}{\hat{\lambda}}(w_t - \hat{\lambda}(p_t + \frac{1}{\hat{\phi}}n_t) - (1 - \hat{\lambda})w_{t-1})$$

\hat{s}_t is an estimate of e_t . Demean it to obtain the instrument $s_t = \hat{s}_t - \text{Mean}(\hat{s}_t)$. Under the following identification and orthogonality assumptions it is true that:

$$\mathbb{E}(u_t^\epsilon) = \mathbb{E}(u_t^z) = \mathbb{E}(u_t^\gamma) = 0$$

$$\mathbb{E}(u_t^\epsilon u_{t-k}^\epsilon) = \mathbb{E}(u_t^z u_{t-k}^z) = \mathbb{E}(u_t^\gamma u_{t-k}^\gamma) = 0 \quad \forall k$$

$$\Sigma = \text{Var} \begin{bmatrix} u_t^\epsilon \\ u_t^z \\ u_t^\gamma \end{bmatrix} = I_{3,3}$$

$$\mathbb{E} \begin{bmatrix} s_t u_t^\epsilon \\ s_t u_t^z \\ s_t u_t^\gamma \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$$

The last set of equations allow us to fully estimate the aggregate VAR. Notice, we can still remain agnostic to the true form of the aggregate labor demand curve and the aggregate Euler equation in our estimation. At this point, the estimation procedure is rather standard. First, we estimate the reduced form aggregate VAR via OLS equation by equation and obtain the reduced form errors U . Then we obtain estimators for the

³²Going forward, we describe the intuition of our procedure using the aggregate labor supply curve. We do this because this is how we will proceed with our application to the Great Recession below. However, the same intuition will hold if we knew the parameterized aggregate labor demand curve.

covariance matrix Σ and the moments vector $\tilde{E} \equiv \mathbb{E}(sU)$. Since the variance of the shocks have been normalized to 1 in the model, the same is done to the moments vector to obtain $E = \frac{\tilde{E}}{\sqrt{\tilde{E}'\Sigma^{-1}\tilde{E}}}$.³³ Second, we solve the system of equations given by

$$\Sigma = \Lambda\Lambda' \quad (3)$$

$$\Lambda \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = E \quad (4)$$

This is a system of 9 independent equations in 9 unknowns to solve for the elements in Λ . Finally, obtain the structural shocks as $u = \Lambda^{-1}U$.

It is worth noting that there is a sense in which Λ , and hence u , is still unidentified. The issue arises because the above procedure does not allow us to label which of the remaining shocks, in this case u_t^z and u_t^γ , corresponds to the aggregate labor demand shock and the aggregate labor supply shock. Basically, the moment conditions in equation (4) identify the first column of Λ . This is the impulse response vector to a u_t^ϵ (leisure) shock. Intuitively, under our assumptions, the instrument provides exogenous variation in this shock that is uncorrelated with the other shocks and hence it allows for the identification of its impulse response vector. However, while the moment conditions in (3) orthogonalize the shocks and the procedure identifies one of the remaining columns in Λ with the impulse response vector of u_t^z and the other with the impulse response vector of u_t^γ , it does not allow one to tell which column is which. A solution to this labeling problem is to use the theoretical co-movement on impact of employment, wages and prices after an aggregate demand and aggregate supply shock, respectively, and label the estimated shocks accordingly. This is the approach we follow when we apply the procedure to identify the shocks that hit the US economy during the Great Recession. To label the shocks we assume that a positive aggregate demand shock does not decrease prices and moves prices and employment in the same direction.

5.2 Using Regional Data To Parameterize Structural Equations

Having regional data allows two advantages to our procedure. First, under certain conditions, the regional data can be used to estimate the structural parameters of one of the aggregate equations. Second, the estimates from the aggregate VAR can be combined with the regional data to infer the shocks driving the cross-region variation. In this subsection, we discuss the conditions needed to use the regional data to infer the structural parameters of one of the aggregate equations. In the next subsection, we discuss how the regional data can be combined with the results from the aggregate VAR to infer the local shocks.

Consider the model economy developed above. Under our assumptions, both the

³³This is the right normalization because the first element in \tilde{E} is not 1 as in the theoretical model, but σ_ϵ^2 . Then, we can obtain the normalization coefficient as $(\sigma_\epsilon^2)^2 = \tilde{E}'\Lambda'^{-1}\Lambda^{-1}\tilde{E} = \tilde{E}'\Sigma^{-1}\Lambda\Lambda^{-1}\tilde{E} = \tilde{E}'\Sigma^{-1}\tilde{E}$

local and aggregate labor supply curves share similar parameters. Likewise, the both the local and aggregate labor demand curves share similar parameters. Consider the local aggregate supply curve:

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

If $v_{kt}^\epsilon = 0$ or $cov(v_{kt}^\epsilon, n_{kt}) = 0$ and $cov(v_{kt}^\epsilon, p_{kt}) = 0$, regional data can be used to estimate ϕ and λ . Under the former assumption, there is no local component to the labor supply shock and the remaining variables in the regression are measured without error. The latter assumption allows for measurement error in the variable regressions or shocks to labor supply. However, in order for the estimates of ϕ and λ to be unbiased, the regional shocks to labor supply or the measurement error in the that are uncorrelated with local price and employment growth.³⁴ The intuition is straight-forward. If there is no systematic variation in the labor supply curve across regions, local movements in prices, employment and nominal wages will be driven by changes in local labor demand. This allows the parameters of the labor supply curve to be traced out.

Conversely, $v_{kt}^{z^i} = 0$ or $cov(v_{kt}^{z^i}, n_{kt}) = 0$ and $cov(v_{kt}^{z^i}, p_{kt}) = 0 \forall i = \{x, y\}$, regional data can be used to trace out the parameters of the local labor demand curve. In this case, all the variation in the labor market would be driven by shifts in labor supply and sticky wages. As a rule, it is probably unrealistic to assume that there is either no regional component to labor supply shocks or no regional component to productivity shocks. However, during certain periods, the assumption may be more plausible. It is in these periods that the regional data can be used to infer structural parameters that help to discipline the aggregate VAR.

Below, we argue that the assumption that the regional component of labor supply shocks was both small and uncorrelated with local price and employment growth during the Great Recession. Given that, we use the regional data to provide estimates of $\hat{\phi}$ and $\hat{\lambda}$. These estimates will be used to parameterize the aggregate labor supply curve when estimating our aggregate VAR.

5.3 Estimating the Regional Shocks (Preliminary)

The procedure for inferring regional shocks follows a similar logic than the one just described for inferring aggregate shocks. The objective is to construct a suitable instrument, this time using information about certain relationships in the aggregate together with local tradable versus non-tradable employment data, that identifies one impulse column-vector in the regional VAR. For exposition, we maintain the assumption that $v_{kt}^\epsilon = 0$. This will be the assumption we provide supporting evidence for below during the period of the Great Recession. From Lemma 2, we can write the reduced form re-

³⁴Both employment and price growth are endogenous to the local labor supply shock. However, $cov(v_{kt}^\epsilon, n_{kt})$ and $cov(v_{kt}^\epsilon, p_{kt})$ could still be equal to zero depending on how v_{kt}^ϵ is correlated with the regional shocks to the discount rate or productivity.

gional VAR in $\{\tilde{w}_t, \tilde{p}_t, \tilde{n}_t^y\}$ with the variables expressed as deviations from the aggregate:

$$(I - \tilde{\rho}(L)) \begin{bmatrix} \tilde{p}_t \\ \tilde{w}_t \\ \tilde{n}_t^y \end{bmatrix} = \tilde{\Lambda} \begin{bmatrix} \tilde{v}_t^x \\ \tilde{v}_t^y \\ \tilde{v}_t^\gamma \end{bmatrix}$$

Moreover, consider a structural relationship in Lemma 2 describing the relative allocation of employment in the tradable versus non-tradable sector as a function of the variables in the regional VAR.

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

When this relationship holds, we can infer that for a fixed vector of $\{\tilde{p}_t, \tilde{w}_t, \tilde{n}_t^y\}$ for all regions the extent to which a regional productivity shock (either in the tradable or non-tradable sector) occurred. The key moment that allows for this identification is relative employment in the tradable and non-tradable sectors $(\tilde{n}_t^x - \tilde{n}_t^y)$.

Parametrizing this equation and using regional data we can obtain the residual and demean it to obtain the instrument \tilde{s}_t . Under the maintained orthogonality assumptions, it satisfies the moment conditions:

$$\mathbb{E} \begin{bmatrix} \tilde{s}_t \tilde{v}_t^x \\ \tilde{s}_t \tilde{v}_t^y \\ \tilde{s}_t \tilde{v}_t^\gamma \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}$$

As before, we can then estimate the reduced form regional VAR and identify the impulse matrix $\tilde{\Lambda}$ with these extra moment conditions. The procedure uniquely identifies the demand shock. However, it does not allow us to label the productivity shocks. We know that one is the productivity shock in the tradable sector and the other one corresponds to the non-tradable sector, but cannot tell which one is which. In our decomposition below, we will assess the relative importance of the discount rate shock relative to the combined productivity shocks in the tradable and nontradable sectors.

6 An Application to the US Great Recession

The cross sectional facts presented above represent a puzzle. At the aggregate level, nominal wages and consumer prices did not appear to respond much (relative to trend) as aggregate unemployment increased during the 2007-2010 period. However, exploiting variation across regions, there appears to be a significant negative relationship between nominal wages and local employment and between non-tradable prices and local employment. Why did aggregate wages and prices respond so little during the Great Recession while there was a strong relationship between these variables at the regional level?

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.

6.1 Cross Sectional Estimation

As discussed above, if we know that the aggregate labor supply curve takes the following form:

$$w_t = \lambda(p_t + \frac{1}{\phi}n_t) + (1 - \lambda)w_{t-1} + \lambda(u_t^\epsilon - (1 - \rho_\epsilon)\epsilon_{t-1})$$

estimates of λ and ϕ allows us to create an instrument for the aggregate labor supply shock. From above, the cross region labor supply curve can be expressed as:

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

If either (1) the labor supply shock has no regional component ($v_{kt}^\epsilon = 0$) or (2) the regional component of the labor supply shock is uncorrelated with changes in local economic activity (i.e., $cov(v_{kt}^\epsilon, n_{kt}) = 0$ and $cov(v_{kt}^\epsilon, p_{kt}) = 0$), we can estimate λ and ϕ using cross region variation.

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

$$w_{kt} = b_t + b_1 p_{kt} + b_2 n_{kt} + b_3 w_{kt-1} + \Psi D_t + e_{kt}$$

where $b_1 = \lambda$, $b_2 = \lambda/\phi$, $b_3 = (1 - \lambda)$, and $b_t = \lambda(u_t^\epsilon - (1 - \rho_\epsilon)\epsilon_{t-1})$. The error term also includes 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 2010. When estimating the above regression, we include year fixed effects (D_t). This ensures that we are only using the cross-sectional 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 in Section 3.2. w_{kt} , therefore, is just the log-growth rate in nominal wages within the state between year t and $t + 1$. Our measure of employment at the state level is the employment rate for each state 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} is the log-change in the employment rate between year t and $t + 1$. Finally, p_{kt} is log-change in the the average monthly price index in each state i within year t . We use the price index scaled

to account for the fact that the local non-tradable share in the grocery sector may differ from the composite consumption good. Given that we have observations on 48 states for 4 years, our estimating equation includes 192 observations.

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 in which our wage (price) measures are based, we can split the sample in each year and compute two measures of wages (prices) 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 Estimation Appendix that accompanies the paper. As we show in that appendix, the procedure dramatically corrects the attenuation bias from measurement error in our estimates.

Column 1 of Table 4 shows the estimates of our base specification. Running the above specification using all years of data and including all states, we estimate $\lambda = 0.69$ (standard error = 0.13) and $\lambda/\phi = 0.31$ (standard error = 0.08). Note, from the above utility function, the Frisch elasticity of labor supply is ϕ . Given the above 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. The reason we are able to identify this in the cross section is that employment varied quite a bit across states despite the relatively small movement in real wages. Figure 7 gives the intuition for our wage stickiness findings. Even though employment growth differences across regions stabilize by 2010, nominal wages kept diverging throughout the 2012 period. Places that experienced the biggest employment declines (highest unemployment increases) had a growing wage gap relative to those places with the smallest employment declines continuously during the 2007-2012 period. The amount of stickiness is represented by $(1 - \lambda)$ and is identified of of the response of today's wage growth to lagged wage growth, conditional on today's employment and price growth.

6.1.1 Assessing the Identifying Assumption

Before we use the above estimates of λ and ϕ to help to identify our aggregate VAR, it is necessary to justify our assumption that labor supply shocks were small across regions or, to the extent that they exist, were uncorrelated with local measures of employment and price growth. To assess whether our assumption is plausible, we perform a variety of exercises. As discussed in the introduction, the labor supply shock is modeled as a shock to the preference for leisure relative to consumption. While to our knowledge no one has argued that actual preference parameters changed during the Great Recession, the preference parameter can be thought of as a proxy for either changes in government policy that altered the return to work (see, for example, Mulligan 2011) or for skill mismatch story put forth by Charles et al. (2013). Given this, we explored the regional variation in government policies and industrial mix during this time period. We begin by focusing our attention on four such government policies: state income tax rates, federal food assistance programs, federal programs to help underwater homeowners renegotiate their mortgage contract, and the extension of unemployment benefits. In the online appendix that accompanies the paper, we discuss the exact methodology and data used to explore the regional variation in policy. We briefly summarize the results here. Our analysis shows that for these major policies, our assumption that these policies varied little across regions and the extent to which they did vary was uncorrelated with local measures of economic activity is not at odds with the data.

Using data on statutory tax rates by state from the tax foundation and micro data on individual incomes by state from the American Community Survey, we compute the average marginal tax rate for each state in 2007 and 2010. State income tax changes were not that common during this time period. Roughly 90 percent of the states, population weighted, had essentially no change in their average marginal tax rate during this time period. As seen from Appendix Figure A1, the extent to which the average marginal tax rate changed between 2007 and 2010 was uncorrelated with state employment or price growth between 2007 and 2010.

Likewise, Figure A2 shows no correlation in the growth in benefits from the federal Supplemental Nutrition Assistance Program (SNAP) and employment growth across states during the 2007-2010 period. SNAP is the successor to the federal Food Stamps program. This program was expanded dramatically during the Great Recession. Given that the SNAP program is means tested, an expansion of the program can discourage work effort. This point is made by Mulligan (2012). Using summary data from the US Department of Agriculture, we measure the dollar per SNAP recipient for each state between 2007 and 2010. The average recipient received an increase about 33 percent during the 2007 to 2010 period. Yet, there was very little regional variation in the increase (a standard deviation of 4 percent across the states, population weighted). As seen from Figure A2, the variation that did occur across states was uncorrelated with state unemployment growth. So, while the increase in SNAP benefits may have discouraged labor supply at the aggregate level, there is very little variation across U.S. states.³⁵

³⁵We also explored whether the eligibility for SNAP differed across states in a way that is correlated

Another new federal program that was means tested and was argued to possibly discourage work effort was the Home Affordable Modification Program (HAMP). HAMP was designed to help homeowners who were underwater renegotiate their mortgage. The program was authorized in early 2009 but there were no significant modifications taking place until mid 2010. By the end of 2010, only about 0.5% of households had participated in the program. As seen from Appendix Figure A3, the extent to which households participated in the program varied slightly with the state's employment growth between 2007 and 2010. A one percent decline in employment growth was only associated with a 0.07 percentage point increase in HAMP take up (i.e., from 0.50% to 0.57%). This effect is very small. However, given that there was some relationship between HAMP take up and underlying economic conditions within the state, we perform a robustness specification in our cross sectional estimation that excludes 2009 and 2010 data and therefore only focuses on the periods before HAMP went into effect. Focusing only on the growth in employment, prices and wages between 2007 and 2008 and between 2008 and 2009 yields estimates of $\lambda = 0.73$ (standard error = 0.17) and $\lambda/\phi = 0.39$ (standard error = 0.09). These results are shown in column (2) of Table 4. The estimates are very similar to our baseline estimates. Additionally, in column (4) of Table 4, we exclude states with the highest amount of loan modifications (CA, FL, NV and AZ). As seen from this robustness specification, our results are essentially unchanged.

Finally, we explored the extent to which unemployment benefits were extended at the state level. During the 2007-2010 period, unemployment benefits increased from about 26 weeks per recipient to upwards of 99 weeks per recipient. Some researchers have argued that this large increase in the duration of unemployment benefits can explain only a small portion of the aggregate decline in employment (Rothstein (2012)) while others have argued that it can explain a substantial portion of the aggregate decline in employment (Hagedorn et al. (2013)). By law in 2010, weeks of unemployment benefits were tied to the state's unemployment rate. This implies that there will be a correlation between the total amount of unemployment benefit extension within the state and the change in the state employment rate between 2007 and 2010. Figure A4 shows this correlation. As of 2010, 70 percent of U.S. states had a duration of unemployment benefits that exceeded 86 weeks. These states represent roughly 90 percent of the U.S. population. However many smaller states, mostly in the Plains region of the U.S., had small employment declines and only an extension of unemployment benefits from 60-85 weeks.³⁶ A simple regression line through figure A4 shows that a 1 percent decline in employment growth was associated with an additional 1.5 weeks of unemployment benefits. Again, this is a tiny change. However, as discussed above, we can reestimate our model focusing only on data before the unemployment benefit extension occurred (i.e., prior to 2009). The results are nearly identical to our base

with the change in state economic conditions between 2007 and 2010. We found no evidence suggesting such a relationship.

³⁶States 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.

specification. Additionally, we can include only those states that had an increase in unemployment benefits of at least 86 weeks (using all years of our data).³⁷ These results are shown in column (3) of Table 4. In this specification, our estimates of $\lambda = 0.77$ (standard error = 0.15) and $\lambda/\phi = 0.27$ (standard error = 0.08) are very similar to our base estimates.

In column (5) of Table 4, we control for the industry mix of the state in year 2007. Charles et al. (2013) document that the secular decline in manufacturing depressed employment rates during the 2000s. Autor and Dorn (2013) show that declines in routine employment also depressed employment rates during the 2000s. Both of these prior papers exploit variation across either MSAs or commuting zones. However, as shown in Appendix Table A5, there is very little variation in routine occupation shares across U.S. states as of 2007 and the extent to which there is variation, it is uncorrelated with employment growth during the 2007-2010 period. We define routine occupations to include all manufacturing and administrative occupations. Although not shown, a similar pattern exists for just manufacturing employment as of 2007. This is consistent with the results in Charles et al. that most of the decline in manufacturing employment within the U.S. during the last 15 years took place in the early 2000s. As seen from column (5) of Table 4, adding in controls for the share of employment in routine occupations in 2007 does not alter our estimates in any meaningful way.

As one final robustness specification, we explore whether our results are sensitive to the scaling factor we used to go from local grocery prices to prices for a composite local good. As discussed above, we use a scaling factor of two implying that the composite local good has a non-tradable component that is twice as large as the goods in our retail scanner database. In column (6) of Table 4, we re-estimate our local labor supply equation using no scaling factor for the local inflation rates. The scaling factor is not important for our main results. With a scaling factor of two (our base case), our estimates are $\lambda = 0.69$ and $\lambda/\phi = 0.31$. With a scaling factor of one (no scaling), our estimates are $\lambda = 0.75$ and $\lambda/\phi = 0.32$. Results were similar using other scaling factors.

From this, we conclude that our estimates are very robust to alternate specifications. Throughout, our estimates of λ range from roughly 0.65 to 0.75 and our estimates of ϕ range from roughly 1.8 to 2.8. For our base estimation, we use the estimates from column (1) of Table 4. However, we then show the robustness of our results to alternate estimates of λ and ϕ .

6.1.2 Migration

One assumption we have imposed throughout is that there is no migration across states. However, 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

³⁷The excluded states are: Arkansas, Iowa, Louisiana, Maryland, Mississippi, Montana, Nebraska, New Hampshire, North Dakota, Oklahoma, South Dakota, Utah, and Wyoming.

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. Interstate migration rates have been declining for the last few decades (Kaplan and Schulhofer-Wohl (2013)). Additionally, the decline in the housing market may have further prevented households from migrating across states (Valletta (2013)). Most importantly for us, the net migration rate was uncorrelated with employment growth during this period. This can be seen from Appendix Figure A5. 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 assumption we can apply the parameters of our local labor supply curve to the aggregate labor supply curve.

6.2 Construction of Aggregate Instrument

As noted above, we can construct our instrument for the aggregate labor supply shock by combining parameters estimated from the cross section with aggregate data:

$$\hat{s}_t = \frac{1}{\hat{\lambda}}(w_t - \hat{\lambda}(p_t + \frac{1}{\hat{\phi}}n_t) - (1 - \hat{\lambda})w_{t-1})$$

For the aggregate data, we construct measures of p , n , and w that are comparable to our regional measures. Given that our cross-sectional equations are estimated using annual data, we analogously define our aggregate data at annual frequencies. We use data from the CPI-U to create our measure of aggregate prices. Specifically, we take log-change in the simple average of the monthly CPI's during year t for our measure of p_t .³⁸ For n , we use BLS data on the aggregate employment rate of all individuals of workers 25-54.³⁹ We choose this age range so as to abstract from the downward trend in employment rates due to the aging of the population.⁴⁰

Finally, we use data from the Current Population Survey (CPS) to construct our aggregate wage measure. Like with our regional wage measures, we attempt to control for the changing labor force composition over the business cycle. In the data appendix, we fully discuss how we calculated our measure of aggregate wages. Briefly, we pooled together all data from the March CPS between 1976 and 2013. Within each survey, we restrict our sample to men between the ages of 25 and 54 who currently work at least 30 hours per week and who worked at least 48 weeks during the prior year. We cre-

³⁸All results in the paper are nearly identical if we use annual measures of the PCE as our aggregate price measure.

³⁹We downloaded the monthly data on prime age employment rates directly from the Federal Reserve Economic Data site maintained by the St. Louis Federal Reserve.

⁴⁰Given that we filter all of our data, allowing for trend growth in the employment to population rate does not affect our results in anyway. Our results are essentially unchanged when we use the employment to population ratio for all groups as opposed to using it just for prime age workers.

ate wage measures by dividing annual earnings during the prior year by annual hours worked. Using the pooled data, we regress wage rates on the age, education, race, hours worked and citizenship controls.⁴¹ After running the regression, we take the residuals from this regression and average the residuals for each year. Given that income reports in the March CPS during year t refer to income earned during year $t-1$, we define our wage measures such that they refer to when the income was earned.⁴² Given these data and the estimated parameters, we estimate \hat{s}_t for all t between 1976 and 2012.

6.3 Construction of the Aggregate Shocks

We estimate the reduced form VAR with 2 lags by OLS equation by equation. From the reduced form errors U we obtain sample estimators of the covariance matrix $\hat{\Sigma} = \frac{UU'}{\text{Years} - \# \text{Variables} \# \text{Lags}}$ and moments $\widehat{\mathbb{E}(sU)} = \frac{U\hat{s}}{\text{Years}}$. We numerically solve for the system of equations in (5)-(6), obtain Λ and construct the shocks as $u = \Lambda^{-1}U$. Given (6), the first column in Λ corresponds to the impulse vector of the labor supply shock. We label the second column as the demand impulse vector and the third as the supply impulse vector. As can be seen in Figure 8 and Figure 9, the demand impulse vector has the elements corresponding to prices and employment of identical sign, while in the supply vector these elements have opposing signs. In other words, we label the demand shock as the one that causes prices and employment to move in the same direction. The supply shock, on the other hand, moves prices and employment in opposite directions.

6.4 Findings

We start by reporting the impulse response of aggregate employment, nominal wages and price growth to each of the shocks. Figure 8 shows their behavior after an initial demand (discount rate) shock of the same magnitude and sign as in 2008. All three variables experience a sharp drop on impact of approximately 1 percent. Qualitatively, after a negative demand shock both prices and employment fall and real wages remain constant. However, while both price and nominal wage inflation remain depressed for several years after the shock; employment growth recovers rather quickly. Figure 9 shows the impulse responses to a 2008 supply (productivity) shock. Price growth increases 0.3 percent on impact and employment growth falls approximately the same. Nominal wage growth moderately increases around 0.1 percent, but real wage growth decreases. It is worth noting that the employment response is rather persistent; in contrast with the response to a demand shock. The qualitative employment and price growth response is less clear cut in the case of a 2008 labor supply shock as shown in Figure 10. The employment growth response is rather small, while price growth declines 0.5 percent on impact

⁴¹To the extent possible, the controls in this regression were created to match the controls used to adjust our regional wage data from the ACS.

⁴²In the data appendix we also discuss how we deal with the change in income measurement that occurred in the CPS during 1994.

but recovers quickly; even turning positive after two years. Nominal wage growth, on the other hand, exhibits a steep (0.8 percent on impact) and persistent increase.

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.⁴³ 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 12 presents the counterfactual employment response. During the Great Recession employment fell in the US by more than 3 percent between 2007-2010 and started slowly recovering thereafter; although in 2012 was still 2.5 percent below its 2007 level. The counterfactual exercise shows that the supply and demand shocks contributed about the same to the initial decline during the 2007-2010 period. Nonetheless, absent the sequence of supply shocks, aggregate employment in the US would have recovered much faster. This in line with the difference in persistence in the response to each shock that we noted earlier on. The labor supply shock, however, barely contributed to the observed employment decline.

Figure 13 offers this paper's contribution towards understanding the "missing deflation puzzle". Prices increased around 4 percent overall during the 2007-2012, being stable only during the year 2008. This implies an annual inflation rate of over 1 percent every year, but 2008. The counterfactual cumulative price response to the sequence of Demand shocks alone shows that prices would have decreased by more than 1 percent in the absence of Leisure and Supply shocks over this time period. Such price behavior would not have been considered "puzzling" by most economists given the sharp employment decline. Notwithstanding, the counterfactual response to the sequence of Supply shocks alone shows that prices would have increased even more; approximately 5 percent. We note that, according to our procedure, it is this countervailing Supply shock that arises as an explanation for the missing deflation puzzle.

Finally, Figure 14 shows the cumulative nominal wage response. Similar to the price response, if the US economy would have been hit by the sequence of Demand shocks alone, nominal wages would have declined by approximately 1.5 percent between 2007 and 2012. The combination of the Leisure and Supply shocks was behind the actual 3 percent increase during this period.

6.5 Robustness and Discussion

What patterns across regions in the US are behind our particular decomposition? In Tables 5-7 we report the contribution of each shock to the aggregate employment, price and wage changes implied by different combinations of $\{\phi, \lambda\}$. We do this for both the initial years of the recession (2008 to 2010) and all years (2008 to 2012).

First, we observe that the relative importance of the Leisure shock vis a vis the Demand and Supply 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

⁴³For the interested reader, the actual realizations of the shocks we estimate can be seen in Figure 11.

macro models.⁴⁴ However, suppose we constructed \hat{s}_t using a much lower elasticity $\phi = 0.5$ instead; in line with microeconomic estimates in the literature. In this case, the leisure shock would account for a much larger fraction of the employment decline in the Great Recession. The theoretical intuition for this result is straightforward. Large movements in employment can be rationalized without the need of large leisure shocks given the relatively small movements in real wages in the data only if the labor supply elasticity is large enough. In terms of our procedure, by decreasing the labor supply elasticity we make \hat{s}_t more correlated with employment. Since \hat{s}_t provides exogenous variation in the leisure shock, this results in a larger estimated element in its impulse vector corresponding to employment.

The intuition for the decomposition between demand and supply shocks is more subtle. We find that the degree of wage flexibility λ affects the relative importance of demand vis a vis supply shocks within the remaining unexplained part by the leisure shock. For example, suppose we increase the degree of wage flexibility to $\lambda = 1$ (no wage rigidities) instead of using our estimated $\lambda = 0.66$ and constructed \hat{s}_t . Then the supply shock would account for most of the employment decline in the Great Recession. Hence, the fact that across regions nominal wage differences seem to be relatively persistent, opens the door for the demand shock to play a larger role in the aggregate. Theoretically, it is clear that when λ is large demand 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 model satisfies monetary neutrality.

Regarding our procedure, when λ increases, \hat{s}_t becomes more correlated with inflation and less so with nominal wages. Since we estimate that a leisure shock generates deflation (the element in its impulse vector corresponding to p_t is negative) and now this shock is more correlated with inflation, we need a larger negative supply shock (a shock that increases prices in our VAR) to account for the observed patterns in aggregate inflation. Finally, given a larger negative supply shock, a smaller negative demand shock is needed to account for the employment decline during the Great Recession. Hence, the relevance of demand shocks relative to supply shocks decreases when wages become more flexible.

⁴⁴This 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.

7 Estimating the Regional Shocks

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 their aggregation cannot be arbitrary. That particular regional patterns have interesting implications about aggregate business cycles by placing restrictions on the structure of the economy and, thus, the nature of the underlying shocks driving both regional and aggregate fluctuations.

We started by documenting several facts about the joint behavior of consumer prices, nominal wages and employment across states in the US; going to great lengths to construct state-level price indexes from Nielsen's *Retail Scanner Database* that addressed the main drawbacks in available measures (like the BLS metro-area indices): geographical coverage and measurement error. We found that states where employment fell the most between 2007 and 2010 had much larger nominal and real wage declines and, less so but still significant, price decreases. A pattern that contrasts with the puzzling "missing aggregate price and wage deflation".

We then showed how to use the information contained in regional business cycles to identify the structural shocks in a VAR describing the evolution of aggregate variables. Our methodology first places theoretical restrictions in the structure of regional shocks. We assume that at least one of the shocks is purely aggregate, and thus, the observed regional variation could not be driven by it. This allows us to estimate key parameters in certain relationships in the cross-section that apply in the aggregate as well (for example, a relationship akin to a labor supply schedule) to construct an instrument for the purely aggregate shock, and together with the usual orthogonality conditions for shocks in the VAR, identify the impulse response matrix and all the aggregate shocks. Moreover, we presented a model economy of a monetary union that illustrates the way in which regions aggregate, the identification assumptions and the broad types of shocks that can and cannot be inferred with our methodology.

In our preferred specification we assumed that regional business cycles were not driven by Leisure shocks and estimated a wage setting schedule using regional variation in wages, prices and employment. We estimated a Frisch labor supply elasticity of approximately 2, about the same magnitude used to calibrate macroeconomic models to match aggregate business cycle volatility, and a moderate degree of nominal wage rigidities.

We used this information in our procedure and found that the initial aggregate employment collapse during the recession years could be attributed equally to the aggregate Demand and Supply shocks; both explaining about half of the decline. However, we noted that the impulse response to a Demand shock is relatively short lived when compared to the Supply shock and found that the later sequence of shocks is behind the

sluggish employment during the recovery years. The Leisure shock contributed very little to the employment behavior overall. Finally, we observed that absent the Supply and, to a lower extent, Leisure shocks, the aggregate price level would have fallen more than 1 percent between 2007 and 2012. A similar picture arises when explaining the behavior of nominal wages (although the role of the Leisure shock is much more important in this case). From the exercise we conclude that not only Demand and Leisure shocks hitting US households during the Great Recession, but Supply shocks hitting firms as well, are necessary to jointly explain the lack of nominal wage and price inflation, and the slow employment recovery.

Future work could make use of our local price indexes to study questions where their regional variation is relevant, if only to deflate nominal quantities. For instance, Beraja, Fuster and Hurst (2014) study the implications of regional heterogeneity for monetary policy and need local deflators to come up with real measures of certain variables of interest. Furthermore, the new methodology we presented is relatively general and could be used in a wealth of other applications where cross-sectional information can be exploited to identify parameters and shocks in VARs. For example, Beraja (2014) studies stabilization in fiscal unions. He uses an extended version of our methodology to estimate regional shocks and certain key parameters to perform counterfactuals aimed at quantifying the regional stabilization benefits of federal fiscal rules in a union.

References

- Aguiar, Mark and Erik Hurst, 2007. "Lifecycle Prices and Production", *American Economic Review*, 97(5), 1533-59.
- Aiyagari, S. Rao, 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving", *The Quarterly Journal of Economics*, 109(3), 659-84.
- Alvarez, Fernando, Martin Gonzalez-Rozada, Andy Neumeyer, and Martin Beraja, 2014. "From Hyperinflation to Stable Prices: Argentina's Evidence on Menu Cost Models", working paper.
- 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.
- Benigno, Pierpaolo, 2004. "Optimal Monetary Policy in a Currency Area", *Journal of International Economics*, 63(2), 293-320.
- Bils, Mark and Pete Klenow, 2004. "Some Evidence of the Importance of Sticky Prices", *Journal of Political Economy*, 112(5), 947 -85.
- Blanchard, Olivier and Lawrence Katz, 1991. "Regional Evolutions", *Brookings Papers on Economic Activity*, 1992:1, 1-75.
- Broda, Christian and David Weinstein, 2010. "Product Creation and Destruction: Evidence and Price Implications". *American Economic Review*, 100(3), 691-723.
- Broda, Christian, Ephraim Leibtag, and David Weinstein, 2009. "The Role of Prices in Measuring the Poor's Living Standards", *Journal of Economic Perspectives*, 23(2), 77-97.
- Burstein, A. T., Neves, J. C., & Rebelo, S. (2003). Distribution costs and real exchange rate dynamics during exchange-rate-based stabilizations. *Journal of Monetary Economics*, 50(6), 1189-1214.
- Cage, Robert, John Greenless and Patrick Jackman, 2003. "Introducing the Chained Consumer Price Index," U.S. Bureau of Labor Statistics.
- Charles, Kerwin, Erik Hurst, and Matthew Notowidigdo, 2014. "Housing Booms, Labor Market Outcomes, and Educational Attainment", *University of Chicago Working Paper*.
- Charles, Kerwin, Erik Hurst, and Matthew Notowidigdo, 2013. "Manufacturing Decline, Housing Booms, and Non-Employment", *NBER Working Paper #18949*.

Christiano, Lawrence, Martin Eichenbaum, and Mathias Trabandt, 2014. "Understanding the Great Recession", working paper.

Christiano, Lawrence, Martin Eichenbaum, and Mathias Trabandt, 2013. "Unemployment and Business Cycles", *NBER Working Paper*, 19265.

Coibion, Olivier, Yuriy Gorodnichenko, and Gee Hee Hong, 2014. "The Cyclicity of Sales, Regular and Effective Prices: Business Cycle and Policy Implications", University of Texas Working Paper.

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, 2014. "Inflation in the Great Recession and New Keynesian Models", *Federal Reserve Bank of New York Staff Reports*, 618.

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.

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.

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

Handbury, Jessie, Tsutomu Wantanabe, and David Weinstein, 2013. "How Much Do Official Price Indexes Tell Us About Inflation?", *NBER Working Paper* #19405.

Handbury, Jessie and David Weinstein, 2011. "Is New Economic Geography Right? Evidence from Price Data", *NBER Working Paper* #17067.

Hawkes, William and Frank Piotrowski, 2003. "Using Scanner Data to Improve the Quality of Measurement in the Consumer Price Index", in Scanner Data and Price Indexes, ed. Robert Feenstra and Matthew Shapiro, University of Chicago Press, Chicago, IL.

King, Robert and Mark Watson, 2012. "Disentangling the Channels of the 2007-2009 Recession", *NBER Working Papers 18094*.

Krusell, Per and Anthony Smith, 1998. "Income and Wealth Heterogeneity in the Macroeconomy", *Journal of Political Economy*, 106(5), 867-896.

Lowe, Robin and Candace Ruscher, 2003. "Estimating Price Movements for Consumer Durables Using Electronic Retail Transaction Data", in Scanner Data and Price Indexes, ed. Robert Feenstra and Matthew Shapiro, University of Chicago Press, Chicago, IL.

Kaplan, Greg and Guido Menzio, 2014. "The Morphology of Price Dispersion", *Princeton 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*.

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

Nakamura, Emi and Jon Steinsson, 2014. "Fiscal Stimulus in a Monetary Union: Evidence from US Regions", *American Economic Review*, 104(3), 753-92.

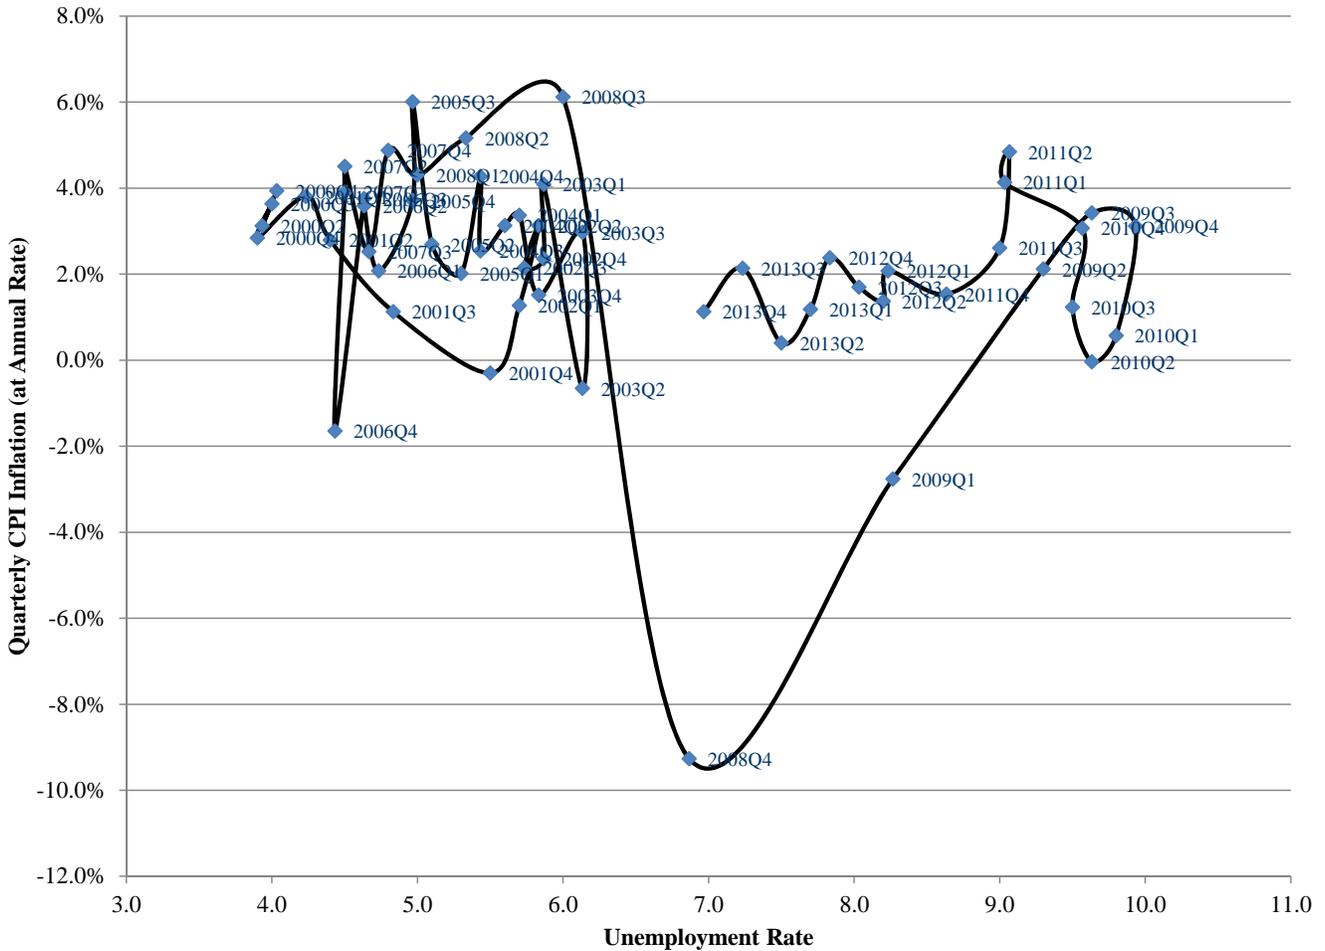
Richardson, David, 2003. "Scanner Indexes for the Consumer Price Index", in Scanner Data and Price Indexes, ed. Robert Feenstra and Matthew Shapiro, University of Chicago Press, Chicago, IL.

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.

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

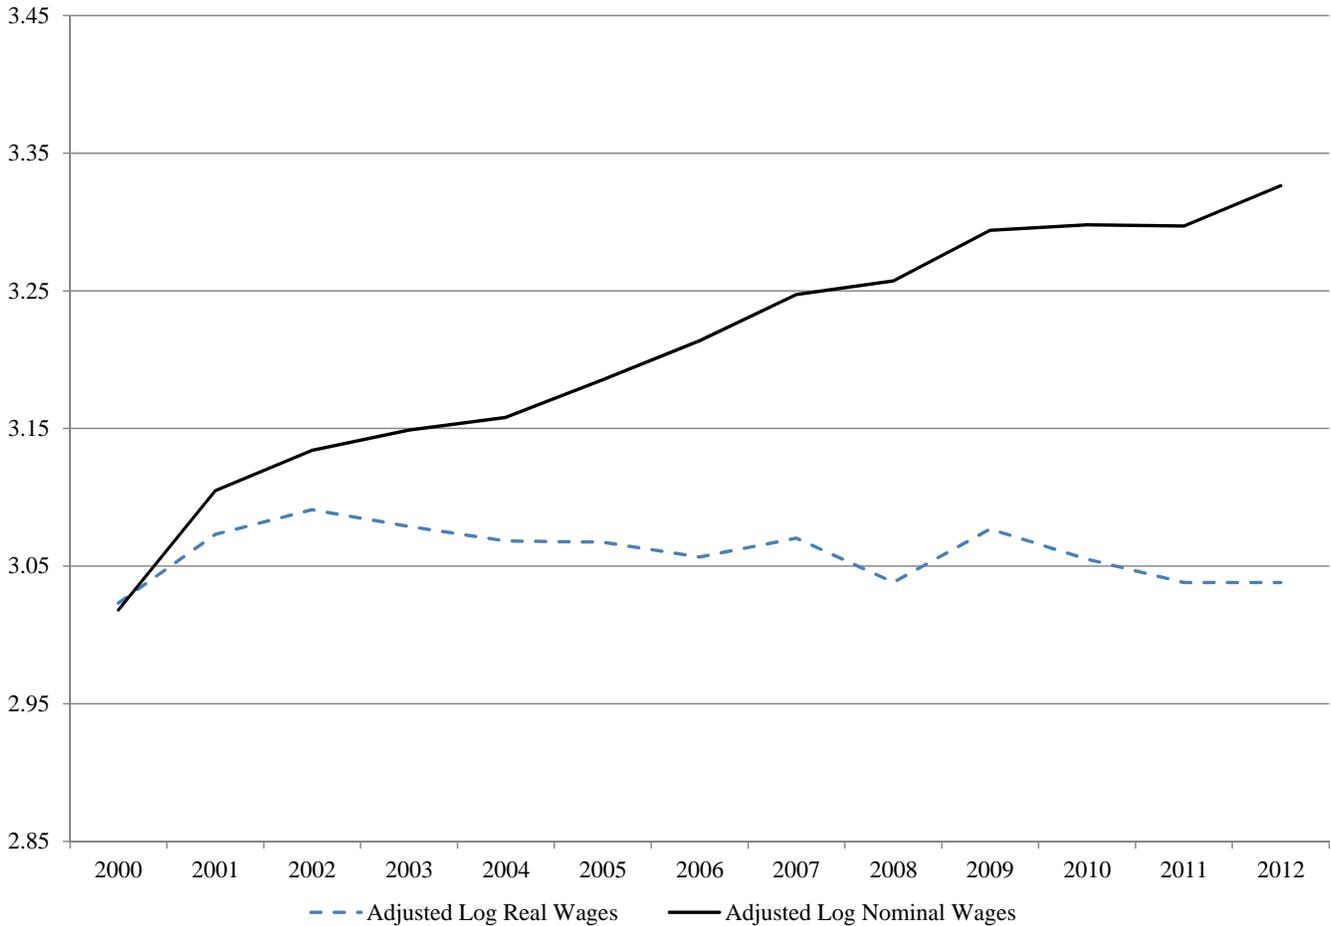
A Figures and Tables

Figure 1: Aggregate Inflation vs Aggregate Unemployment, Quarterly Data (2000Q1-2013Q4)



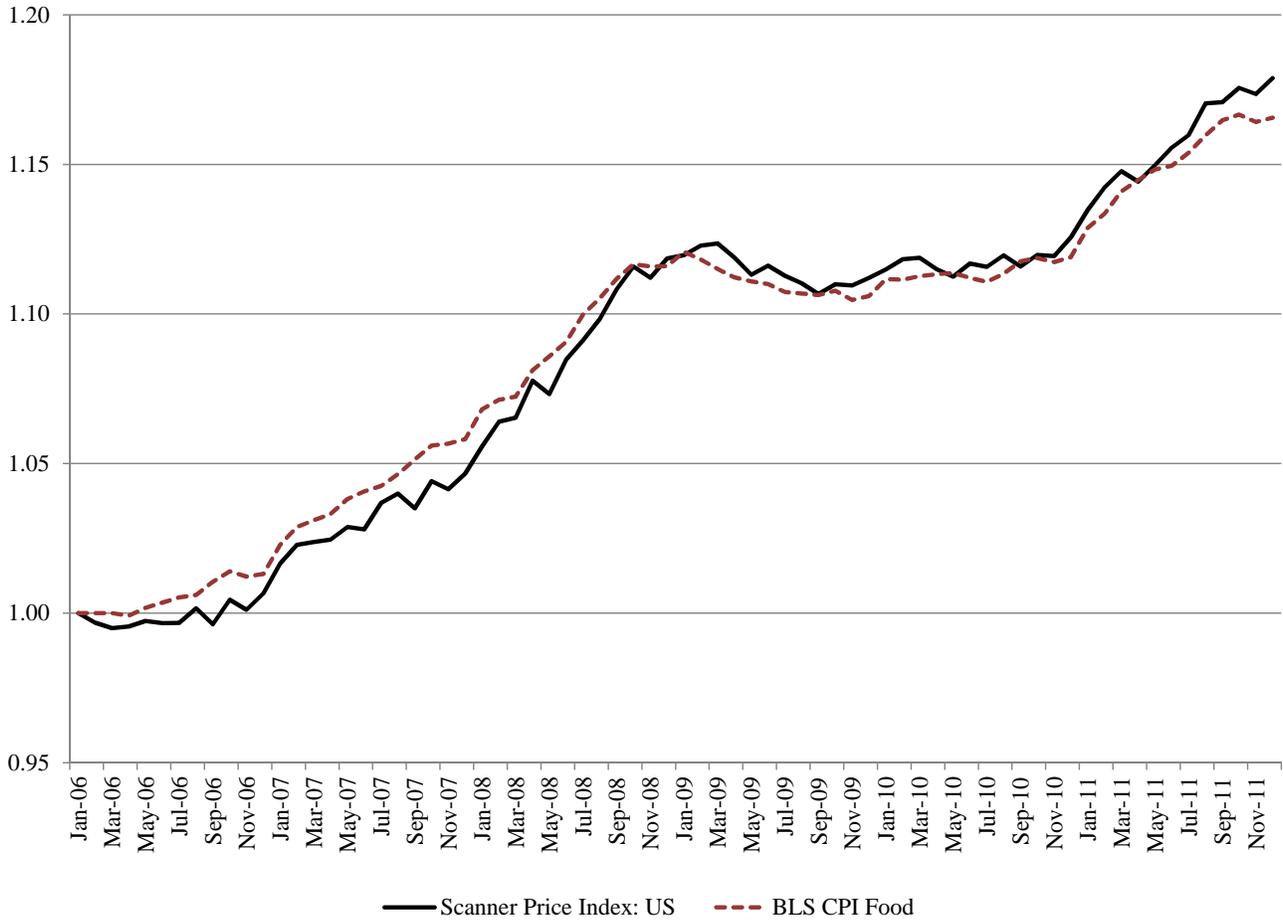
Note: Figure shows the relationship between the aggregate U.S. unemployment rate in quarter t (x-axis) against the annualized inflation rate from quarter t-1 to quarter 1 (y-axis). The line connects quarters from one period to the next. The unemployment rate comes from the Bureau of Labor Statistics. Our measure of inflation is based on the CPI-U.

Figure 2: The Evolution of Aggregate Real and Nominal Log Wages, Annual Data (2000-2012)



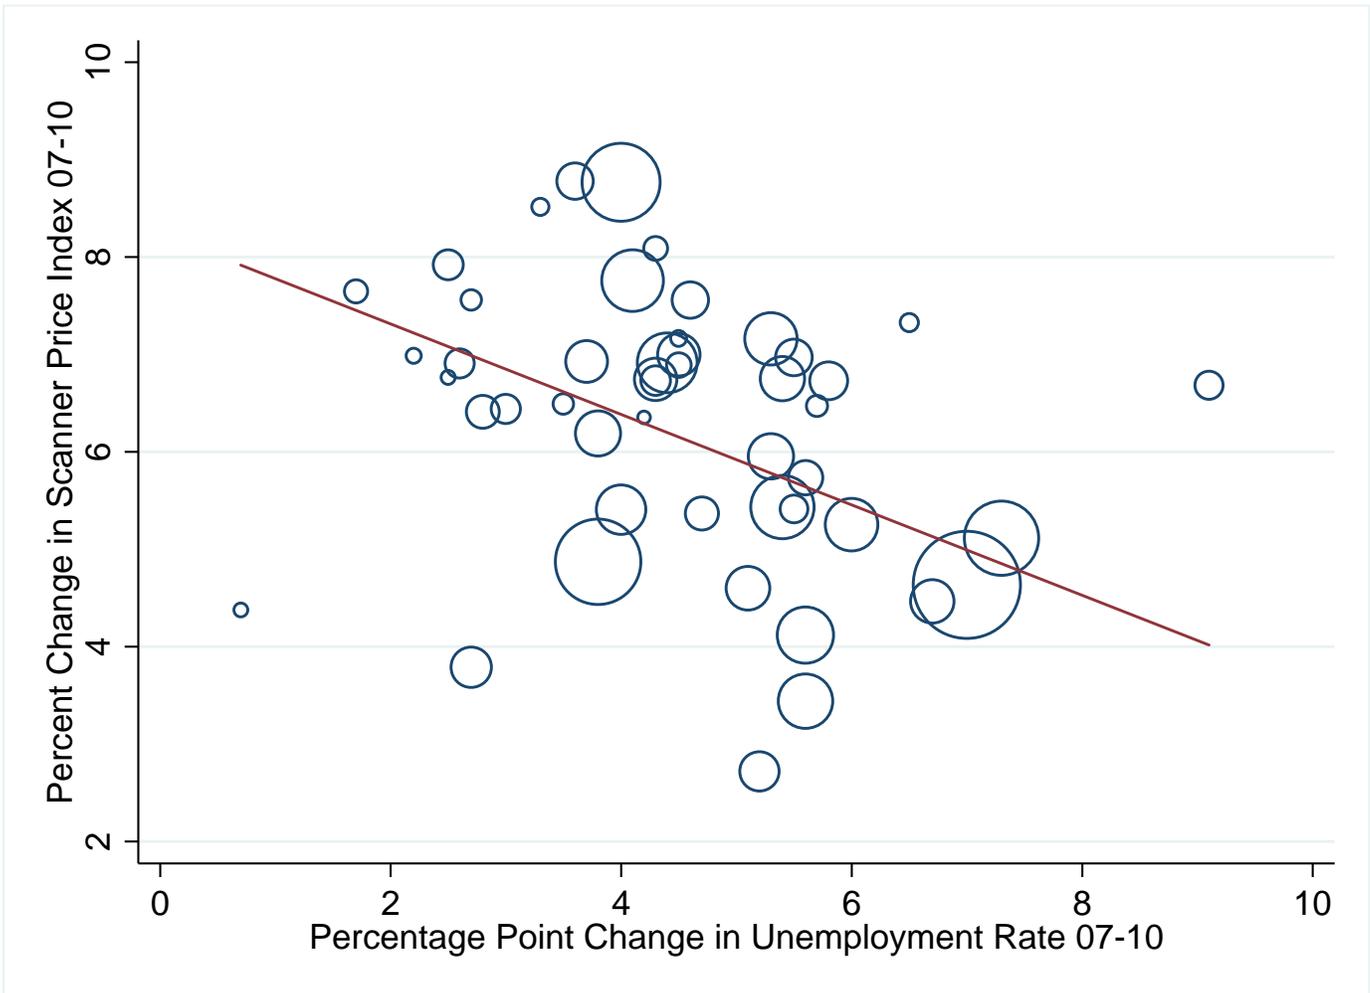
Note: Figure shows the evolution of aggregate real and nominal log wages within the U.S. between 2000 and 2012. Nominal wages were computed using data from the American Community Survey. The sample is restricted to only those individuals who are currently employed, who report usually working 30 hours 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. As discussed in the text, we adjust wages for the changing labor market condition over time. As computed, the wages are for a white male aged 40-44 who was born in the US having 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: Nielsen Retail Price Index vs. CPI Food Price Index, 2006M1 to 2011M12



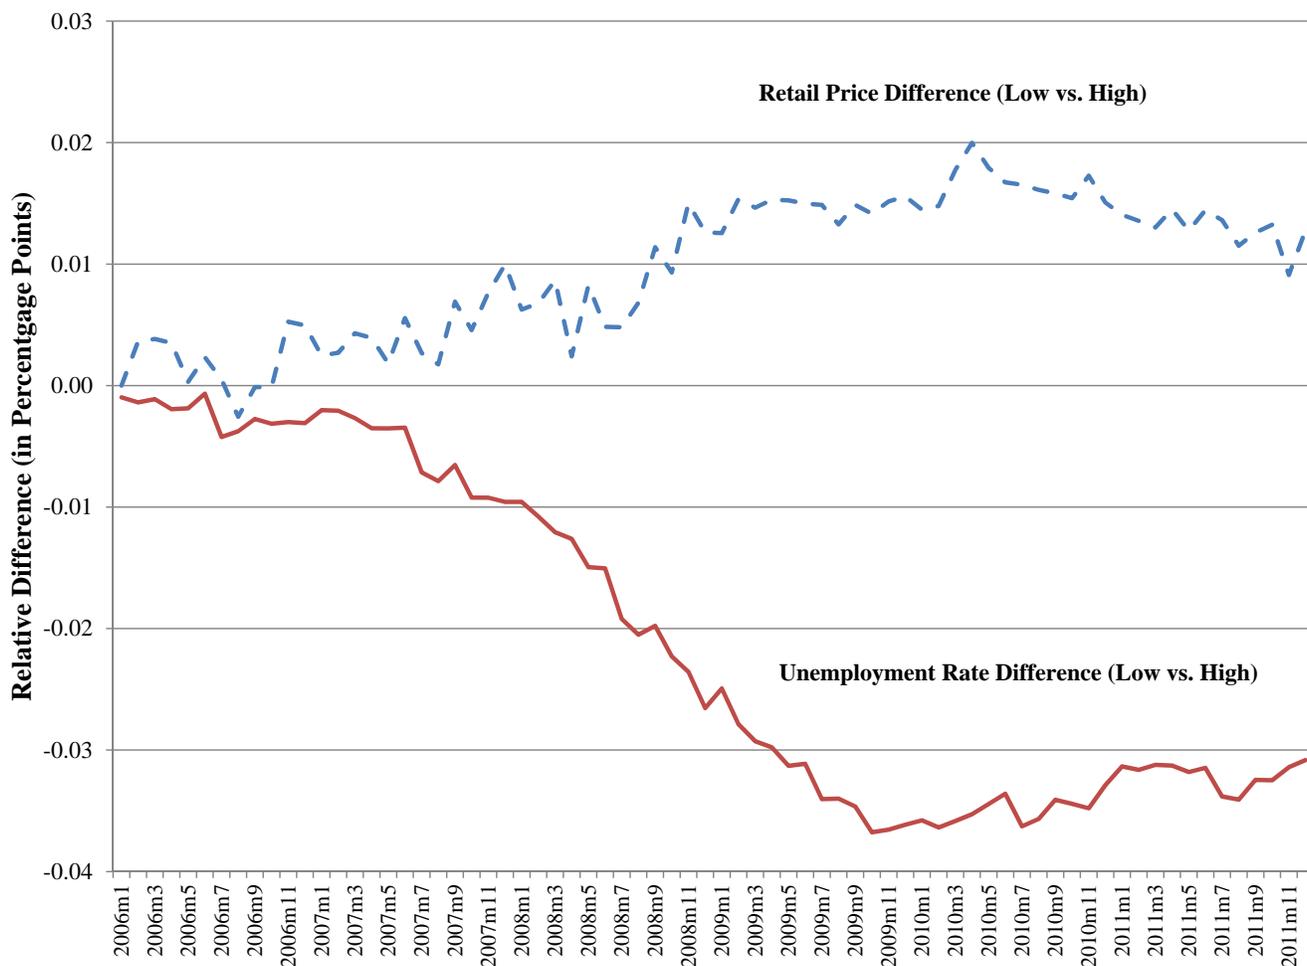
Note: In this figure, we compare our retail price index for the U.S. as whole to the CPI food price index. Given that the goods in our price index come predominantly from grocery, pharmacy, and mass merchandising stores, we thought the food CPI was an appropriate benchmark. For the Nielsen retail price index in this figure, we define a good as a UPC-Store pair. See text for additional details. We normalize both our index and the CPI Food index to 1 in January of 2006.

Figure 4: Change in State Unemployment Rate vs. Cumulative State Retail Price Inflation (2007-2010)



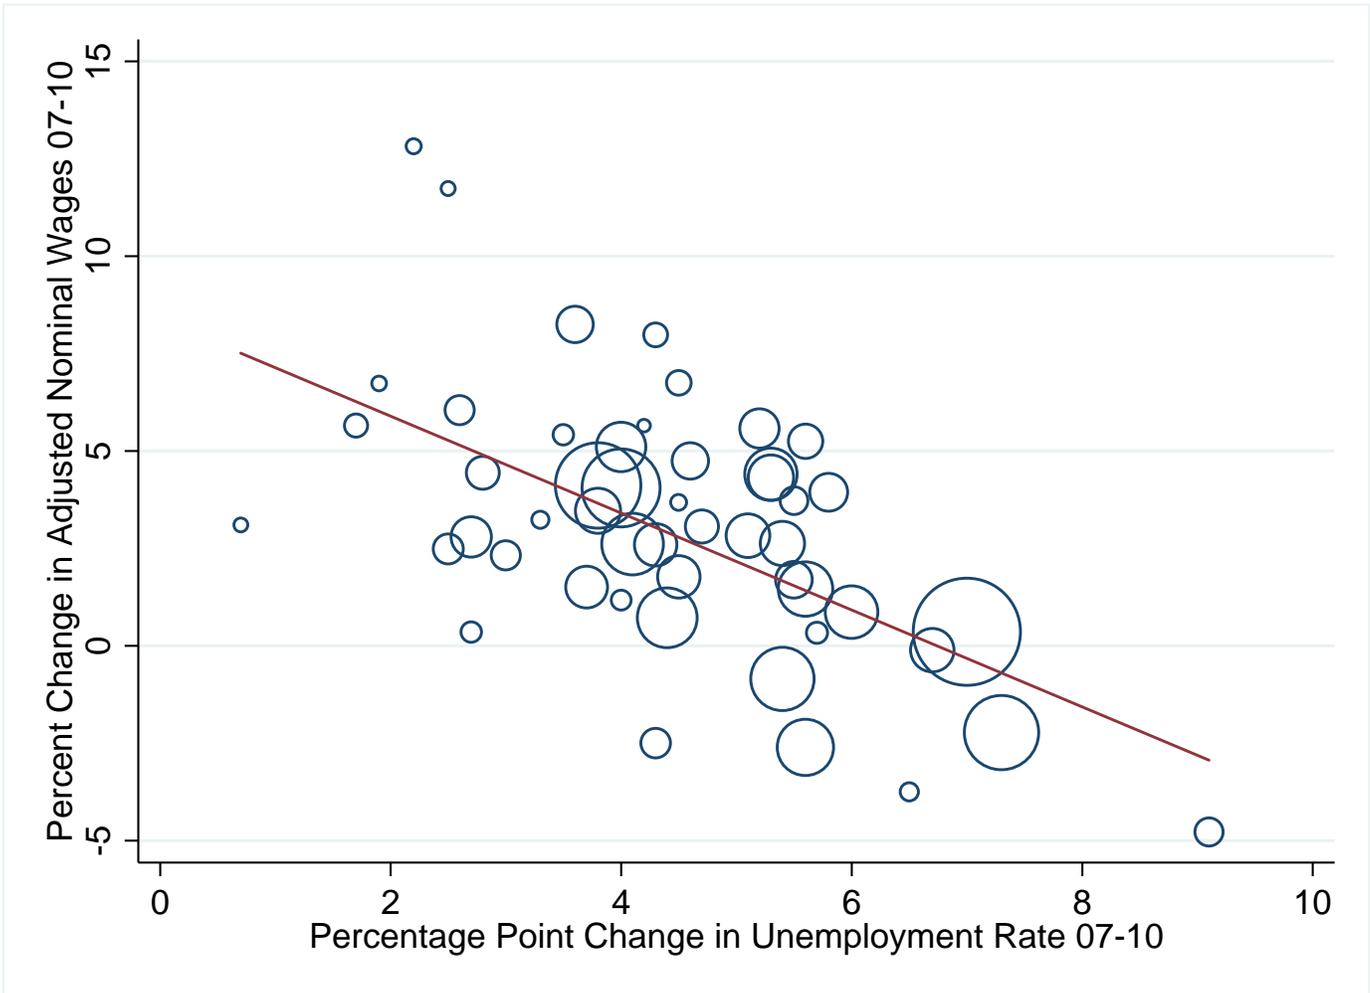
Note: Figure shows a simple scatter plot of the percentage point change in the BLS unemployment rate in the state between 2007 and 2010 against the cumulative percent change in our retail price index based on the Nielsen scanner data during the same period. The retail price index is computed where each good is based on a UPC within the state (as opposed to a UPC-store). 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: Differential Retail Prices and Unemployment between Low and High Unemployment Change States, 2006M1 to 2011M12



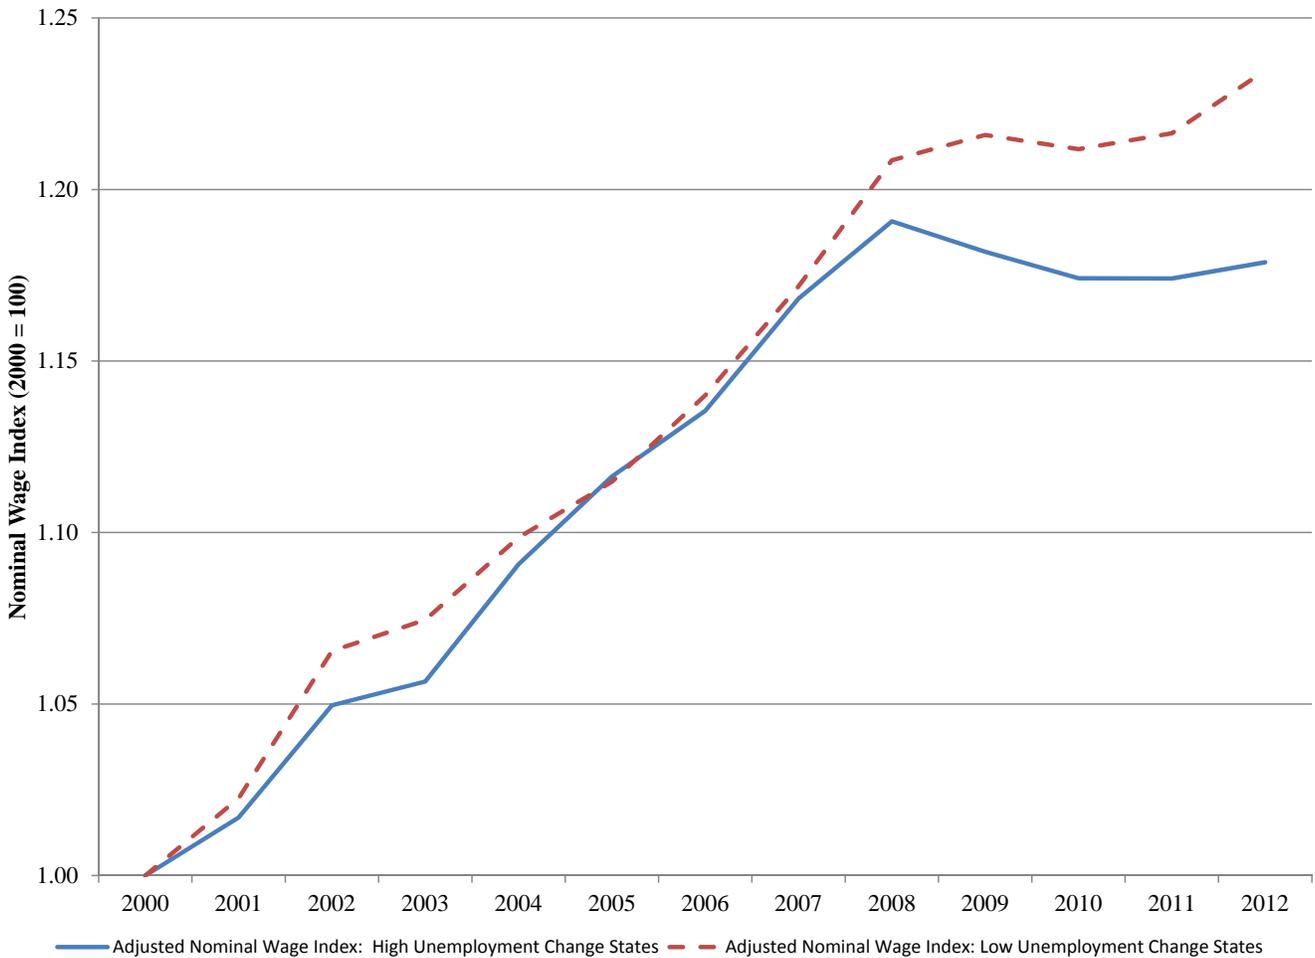
Note: Figure shows the trends in the relative monthly unemployment rate between low and high unemployment change states against the trends in the relative monthly retail price index between low and high unemployment change states. High unemployment change states are the top one-third of all states with respect to the change in unemployment between 2007 and 2010. Low unemployment change states are the bottom one-third of all states with respect to the change in unemployment between 2007 and 2010. Within each group of states for each month, we average the unemployment rate and price index across states weighting each state by their population.

Figure 6: Change in State Unemployment Rate vs. State Nominal Wage Growth (2007-2010)



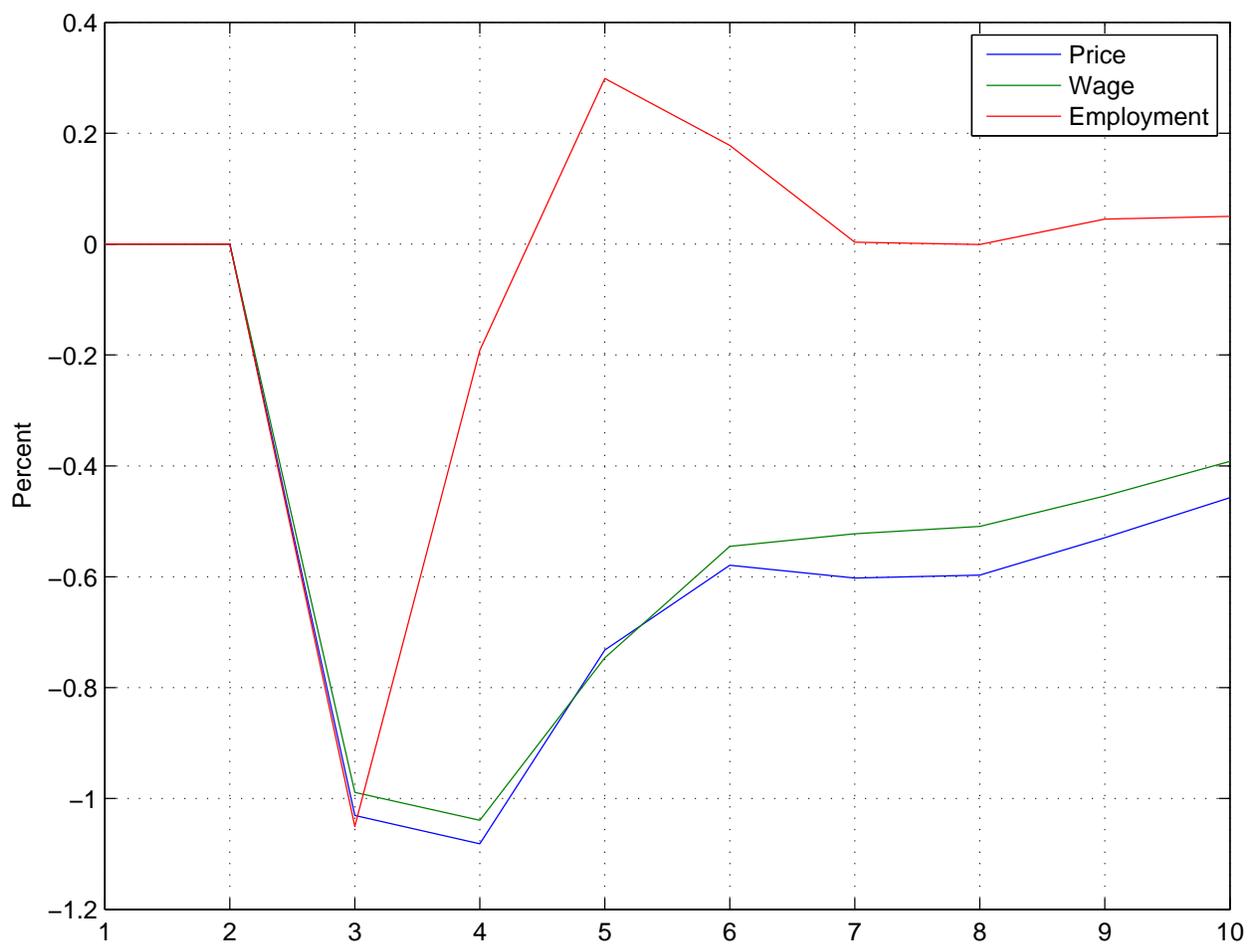
Note: Figure shows a simple scatter plot of the percentage point change in the BLS unemployment rate in the state between 2007 and 2010 against nominal wage growth during the same period. The retail price index is computed where each good is based on a UPC within the state (as opposed to a UPC-store). Nominal wages are adjusted for changing labor market composition within each state. 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 7: Nominal Wage Index in Low and High Unemployment Change States, 2000-2012 (2000=100)



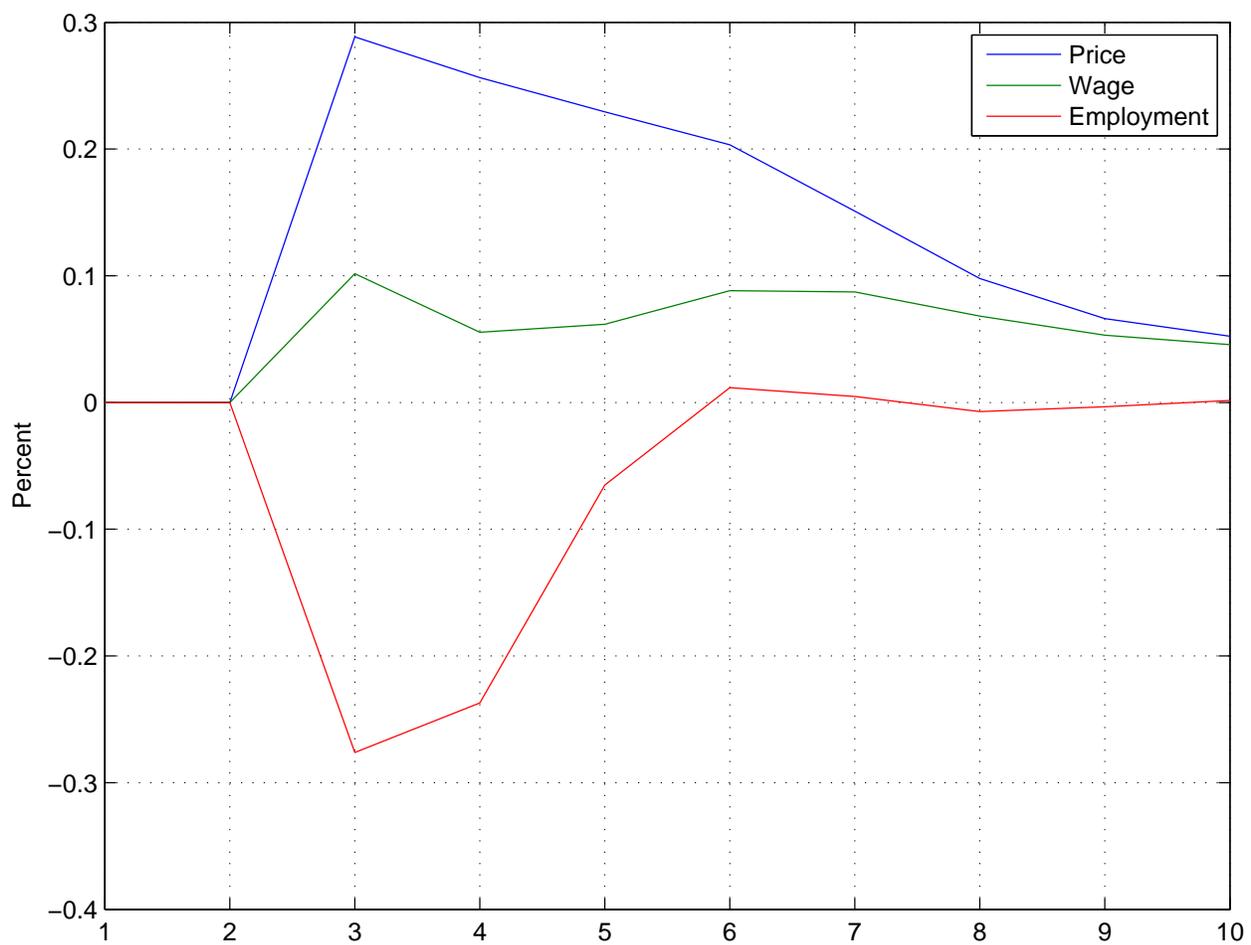
Note: Figure shows the trends in nominal wages between low and high unemployment change states from 2000-2012. As in Figure 6, nominal wages are adjusted for changing labor force composition within the state. High unemployment change states are the top one-third of all states with respect to the change in unemployment between 2007 and 2010. Low unemployment change states are the bottom one-third of all states with respect to the change in unemployment between 2007 and 2010. We normalize nominal wages to 1 in year 2000 for all states. Within each group of states, we average nominal wages weighting each state by their population.

Figure 8: Impulse Response to 2008 Discount / Interest rate Shock



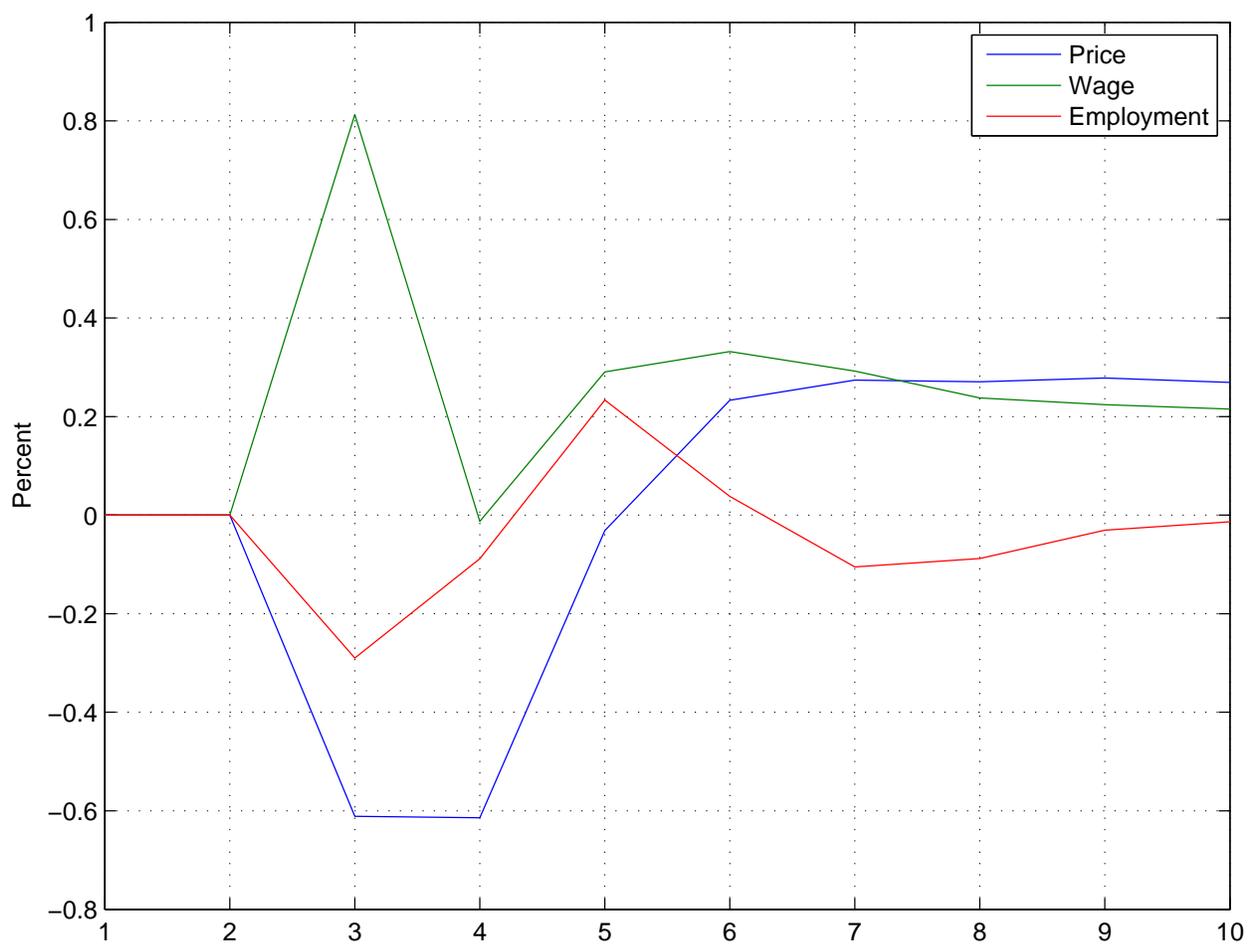
Note: Figure shows the impulse response to a Demand shock of the same sign and magnitude as the one we estimate for 2008. The horizontal axis are years after the shock.

Figure 9: Impulse Response to 2008 Productivity / Markup Shock



Note: Figure shows the impulse response to a Supply shock of the same sign and magnitude as the one we estimate for 2008. The horizontal axis are years after the shock.

Figure 10: Impulse Response to 2008 Leisure Shock



Note: Figure shows the impulse response to a Leisure shock of the same sign and magnitude as the one we estimate for 2008. The horizontal axis are years after the shock.

Figure 11: Shock Time Series

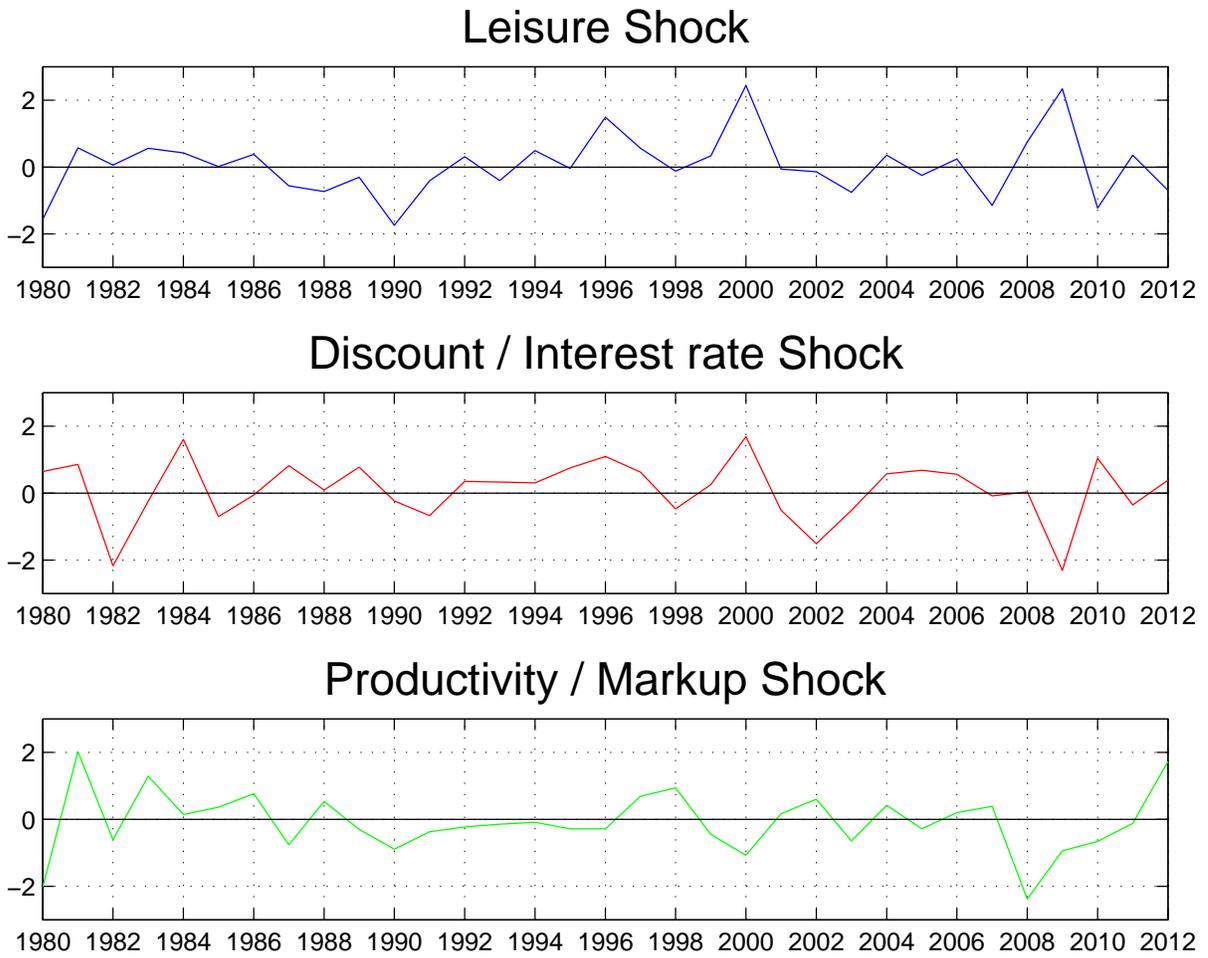
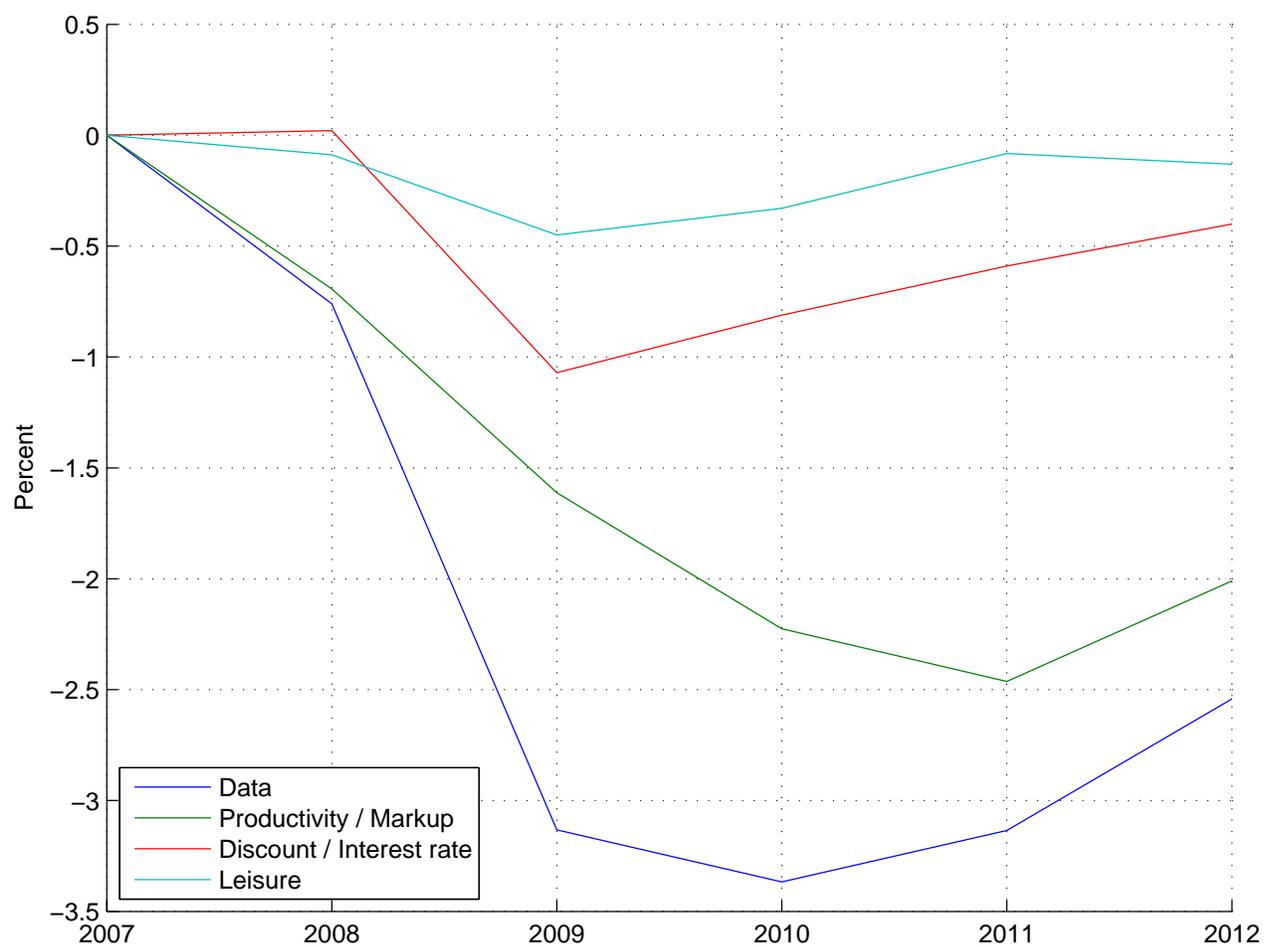
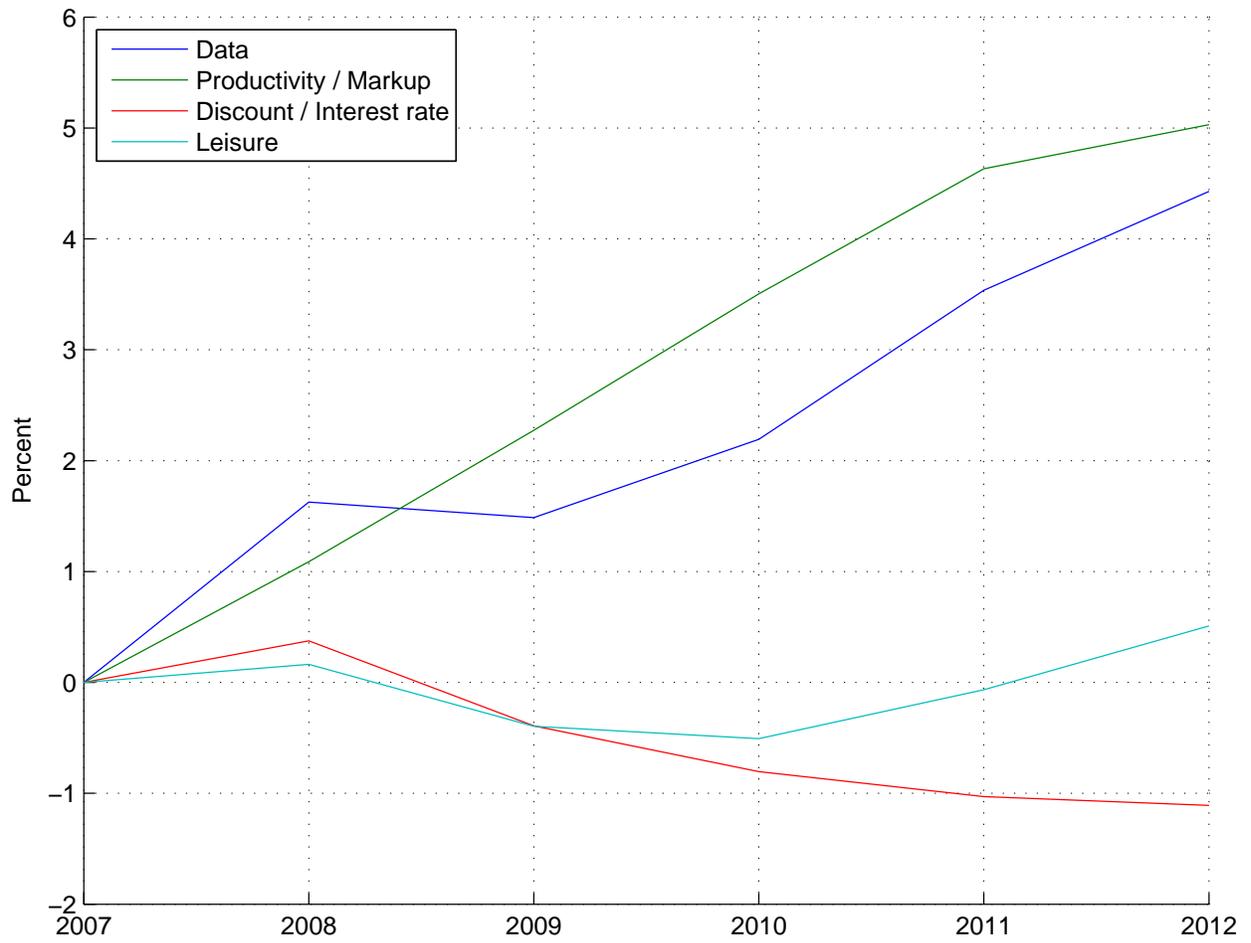


Figure 12: Counterfactual Employment Response



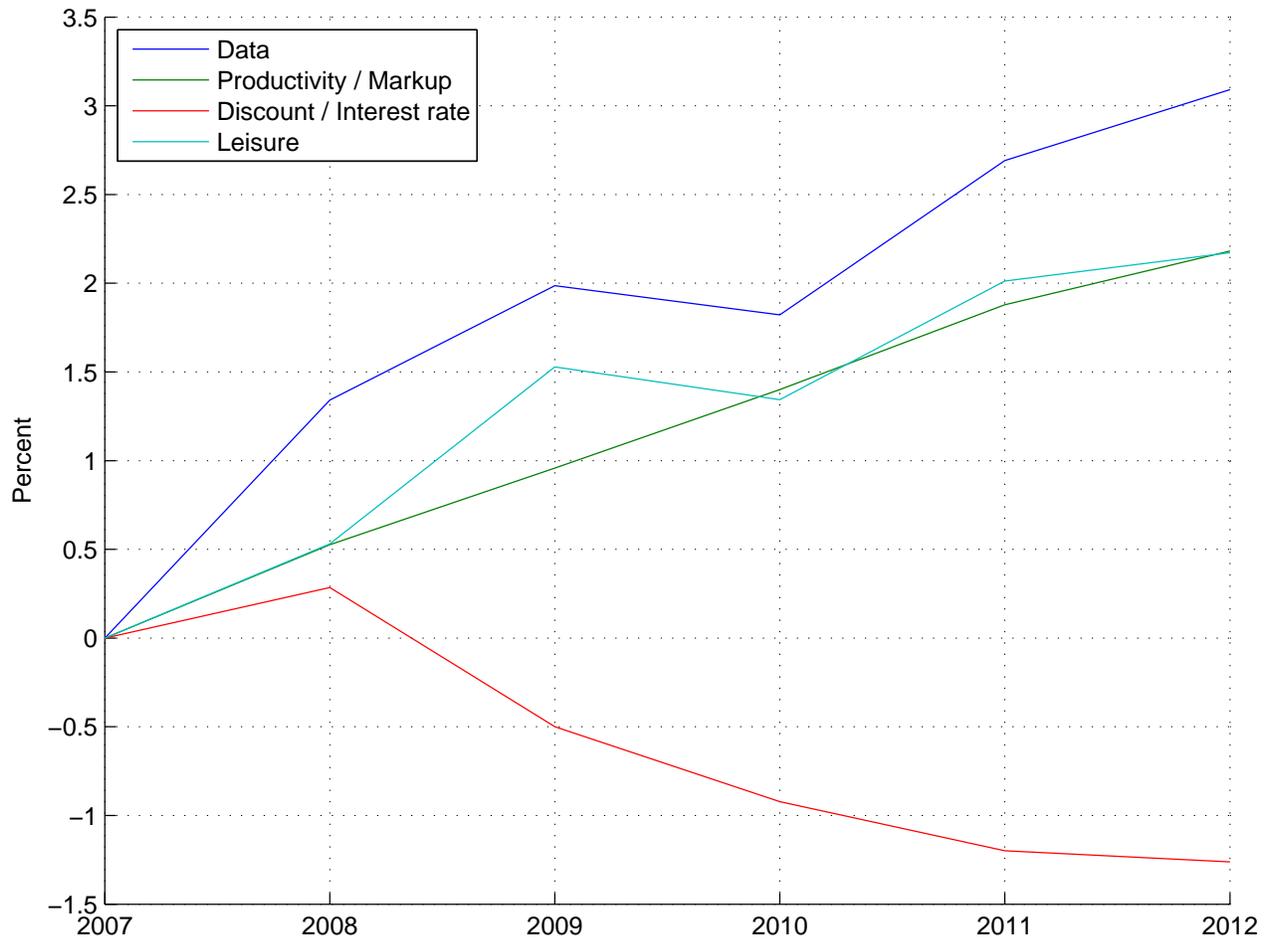
Note: Figure shows the cumulative response of Employment when we feed the VAR with the sequence of shocks between 2008 and 2012; one at a time.

Figure 13: 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 14: 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: AC Nielsen Database Description

Descriptor	Year						Total	Average
	2006	2007	2008	2009	2010	2011		
Number of Observations (million)	11,742.5	12,522.6	12,746.2	12,738.0	12,861.8	13,327.9	75,939.1	12,656.5
Number of UPCs	724,316	761,778	763,802	760,554	746,680	748,224	1,363,965	750,892.3
Number of Chains	73	85	86	86	86	86	88	83.7
Number of Stores	31,642	33,646	34,730	35,285	35,671	35,645	38,379	34,436.5
Number of Zip Codes	10,741	11,099	11,339	11,469	11,576	11,639	11,781	11,310.5
Number of Counties	2,365	2,464	2,496	2,507	2,518	2,526	2,545	2,479.3
Number of MSAs	371	371	371	371	371	371	371	371.0
Number of States	49	49	49	49	49	49	NA	49.0
Transaction Value (US billion)	184.3	204.0	215.1	220.1	223.9	235.9	1283.3	213.9
Pct. Value Used in Price Indexes	52.56	48.39	64.69	64.60	66.39	66.05	NA	60.45

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

	Measure of Price Inflation			
	A. Grocery/Mass Merchandising Price Inflation		B. Composite Price Inflation	
	(1)	(2)	(3)	(4)
Change in Unemployment Rate (Percentage Point)	-0.464 (0.140)	-0.375 (0.125)	-0.928 (0.279)	-0.751 (0.250)
Per-Capita GDP Growth (Percent)	0.170 (0.046)	0.085 (0.044)	0.341 (0.092)	0.170 (0.089)
Per-Capita Hours Growth (Percent)	0.300 (0.075)	0.151 (0.073)	0.591 (0.150)	0.302 (0.147)
House Price Growth (Percent)	0.036 (0.013)	0.031 (0.012)	0.072 (0.026)	0.062 (0.023)
Employment Rate Growth (Percent)	0.222 (0.073)	0.078 (0.069)	0.444 (0.146)	0.157 (0.139)
IV: Change in Unemployment Rate (Percentage Point)	-0.477 (0.169)	-0.413 (0.152)	-0.954 (0.338)	-0.825 (0.303)
IV: Employment Growth	0.323 (0.118)	0.279 (0.120)	0.646 (0.237)	0.558 (0.240)
Goods Defined as UPC-Store?	No	Yes	No	Yes

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. Panel A uses the underlying data from our sample to compute the price indices (Pr). Panel B scales up the variation in the grocery goods within our sample to estimate the variation in prices for a composite consumption good (P). Given our scaling factor, the coefficients in Panel B are two times the coefficients in Panel A. Columns (1) and (3) use our price index measures where we define a good in our price index without conditioning on the store it was sold. Columns (2) and (4) define a good a store-UPC pair. Standard errors are in parenthesis. Each regression is weighted by the state's 2006 population.

Table 3: The Relationship between Regional Wage Changes and Changes in Regional Economic Activity: 2007-2010

	Dependent Variable	
	Nominal Wage Growth	Real Wage Growth
Change in Unemployment Rate (Percentage Point)	-1.244 (0.205)	-0.495 (0.354)
Per-Capita GDP Growth (Percent)	0.487 (0.060)	0.316 (0.113)
Per-Capita Hours Growth (Percent)	0.653 (0.120)	0.357 (0.200)
House Price Growth (Percent)	0.113 (0.019)	0.050 (0.032)
Employment Rate Growth (Percent)	0.618 (0.109)	0.458 (0.173)
IV: Change in Unemployment Rate (Percentage Point)	-1.500 (0.253)	-0.669 (0.430)
IV: Employment Growth	1.011 (0.195)	0.453 (0.276)

Note: Table shows the results of a bi-variate regression of nominal wage growth (column 1) or real wage growth (column 2) rate 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 by the composite price index (i.e., the real price index scaled to account for different non-tradable shares). Standard errors are in parenthesis. Each regression is weighted by the state's 2006 population.

Table 4: Estimates of λ and $\frac{\lambda}{\phi}$ using Cross- Region Data

	Specification					
	(1)	(2)	(3)	(4)	(5)	(6)
λ	0.69 (0.13)	0.73 (0.17)	0.77 (0.15)	0.63 (0.11)	0.69 (0.13)	0.75 (0.13)
λ/ϕ	0.31 (0.08)	0.39 (0.09)	0.27 (0.08)	0.25 (0.08)	0.32 (0.08)	0.32 (0.07)
Price Adjustment	2	2	2	2	2	1
Time Period	2007-2010	2007-2008	2007-2010	2007-2010	2007-2010	2007-2010
Additional Restrictions	None	None	Exclude States That Did Not Extend Unemployment Benefits for at Least 85 weeks	Exclude CA, FL, NV, and AZ	Include Industry Controls	None

Note: Table shows the estimates of λ and ϕ 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, the states included, and additional control variables added. We adjust for measurement error in wage growth, lagged wage growth, and price growth using the split sample methodology discussed in the text. All regressions included year fixed effects. All standard errors are clustered at the state level.

Table 5: Discount/Interest rate (γ) and Productivity/Markup (z) shocks' contribution to employment change

λ	2008 to 2010						2008 to 2012				
	ϕ						ϕ				
		0.5	1	2	3	4	0.5	1	2	3	4
0.1	γ	74	74	73	72	72	41	40	39	38	38
	z	28	31	32	33	33	52	56	58	59	59
0.3	γ	41	60	64	64	64	20	30	30	30	30
	z	27	34	38	40	41	51	61	66	68	68
0.5	γ	8	35	49	52	53	6	17	22	23	23
	z	28	38	45	47	48	51	66	73	75	76
0.7	γ	-2	14	32	37	40	4	8	15	17	17
	z	27	43	51	53	54	50	70	79	81	82
0.9	γ	-4	3	18	24	27	6	4	9	12	13
	z	26	45	56	60	61	48	71	83	85	86
1	γ	-3	0	12	19	22	7	3	7	10	11
	z	25	45	59	62	64	47	72	84	87	88

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.

Table 6: Discount/Interest rate (γ) and Productivity/Markup (z) shocks' contribution to price level change

λ		2008 to 2010					2008 to 2012				
		ϕ					ϕ				
		0.5	1	2	3	4	0.5	1	2	3	4
0.1	γ	-111	-115	-115	-115	-115	-59	-56	-55	-54	-54
	z	140	145	147	148	148	94	98	100	101	101
0.3	γ	-39	-86	-101	-104	-105	-33	-55	-58	-58	-58
	z	139	151	155	156	156	93	104	108	110	111
0.5	γ	38	-26	-63	-74	-78	19	-25	-45	-49	-50
	z	136	153	157	158	158	90	107	113	115	116
0.7	γ	72	23	-20	-35	-43	50	7	-21	-30	-34
	z	131	150	153	152	151	86	105	113	115	116
0.9	γ	85	51	15	-1	-10	65	30	1	-10	-15
	z	127	146	145	142	141	83	102	109	110	111
1	γ	88	59	28	12	4	70	38	10	-1	-6
	z	125	145	141	137	135	82	101	106	108	108

Note: Table shows the percent contribution of the demand and supply shocks to the aggregate price level 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.

Table 7: Discount/Interest rate (γ) and Productivity/Markup (z) shocks' contribution to wage level change

λ		2008 to 2010					2008 to 2012				
		ϕ					ϕ				
		0.5	1	2	3	4	0.5	1	2	3	4
0.1	γ	-23	-12	-7	-5	-4	-54	-42	-36	-34	-33
	z	29	32	33	34	34	44	48	50	51	51
0.3	γ	-33	-37	-28	-24	-22	-53	-65	-58	-53	-51
	z	30	36	40	42	43	45	55	61	63	64
0.5	γ	16	-27	-34	-33	-31	17	-43	-58	-58	-57
	z	27	37	46	49	51	40	57	68	72	75
0.7	γ	56	3	-23	-27	-29	68	-1	-37	-45	-47
	z	23	35	47	52	55	35	53	69	75	78
0.9	γ	77	31	-4	-13	-17	95	35	-10	-23	-29
	z	21	31	44	50	54	31	47	64	72	76
1	γ	84	41	6	-5	-11	104	48	3	-12	-19
	z	21	29	42	49	53	30	45	61	69	73

Note: Table shows the percent contribution of the demand and supply shocks to the aggregate wage level 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.

Table 8: From λ to aggregate nominal wage growth volatility σ_w

	λ					
	0.1	0.3	0.5	0.7	0.9	1
$\frac{\sigma_w^{\gamma \text{ only}}(\lambda)}{\sigma_\gamma}$	5	18	31	41	48	50
$\frac{\sigma_w^{z \text{ only}}(\lambda)}{\sigma_z}$	2	7	11	12	12	12

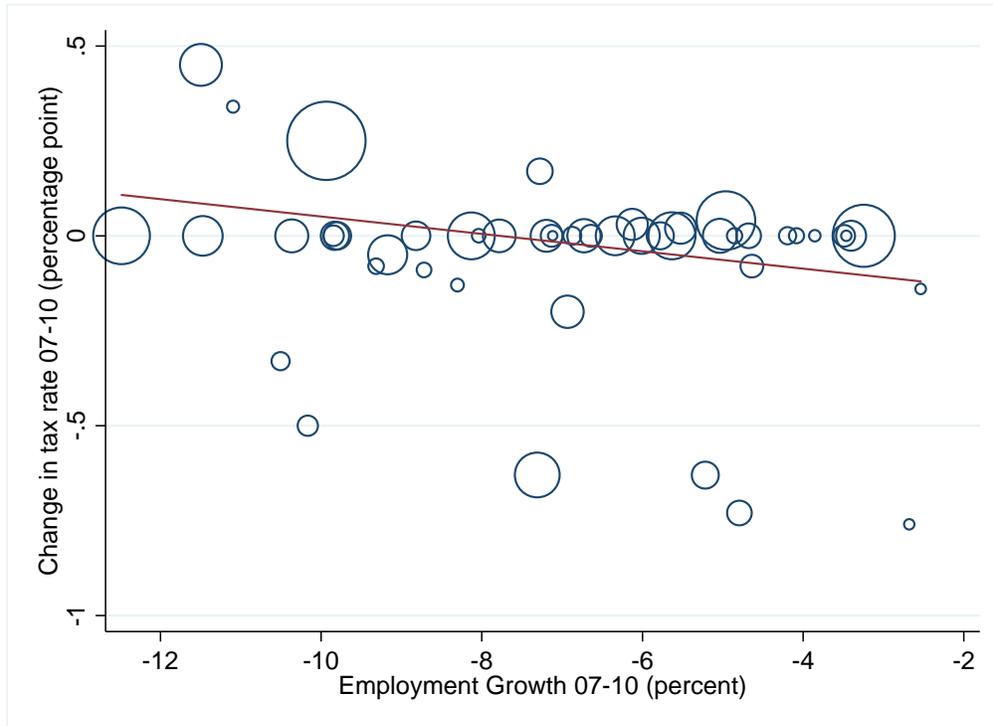
Note: Table shows the stationary aggregate nominal wage growth volatility when the only driving process is γ or z expressed as basis points from the stationary volatility of the γ and z processes respectively. This is done for different degrees of nominal wage rigidity λ .

Table 9: Regional counterfactuals

	Data	Productivity/Markup	Discount/Interest rate
$\sigma_n^2/(\sigma_n^{data})^2$	1	0.05	0.93
$\sigma_p^2/(\sigma_p^{data})^2$	1	0.78	0.27
$\sigma_w^2/(\sigma_w^{data})^2$	1	0.98	0.05
β_{p,n^y}	0.46	-3.29	0.54
β_{w,n^y}	0.57	2.12	0.19
$\beta_{p,w}$	0.32	0.15	0.33

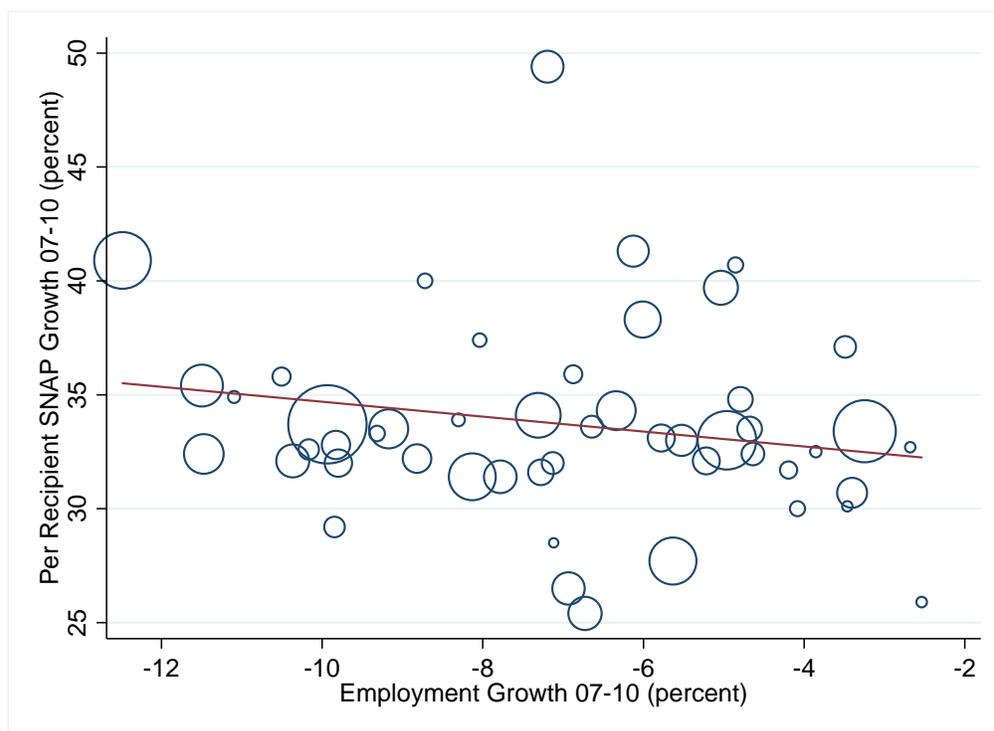
Note: The first three lines in the table show the counterfactual variance across states relative to the actual variance of the total percent change in each variable between 2007-2010. The last three lines show the population weighted OLS coefficient corresponding to each variable pair. For example, β_{p,n^y} 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 both z^x, z^y shocks and no γ shock. The third column corresponds to the counterfactual with the γ shock alone.

Figure A1: Change in State Tax Rate vs. State Employment Growth: 2007-2010



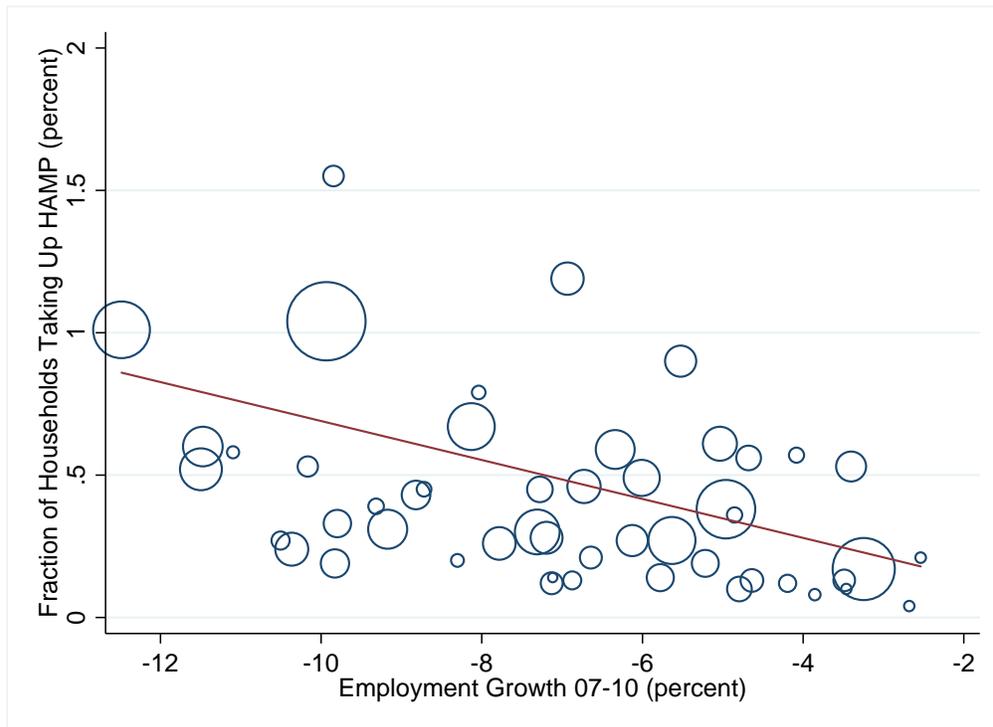
Note: Figure shows the change in the average marginal tax rate in a state between 2007 and 2010 against employment growth in the state during the same period. Employment growth comes from the BLS and is defined in the text. To compute the average marginal tax rate in the state we use data from the American Community Survey and state tax rate formulas from taxfoundation.org. Using the American Community Survey, we compute the fraction of state residents in 15 labor income bins as well as the mean income within each bin. We then compute the marginal tax rate in that bin. Averaging over the bins, we get the state's average marginal tax rate. Our procedure does not account for any state level deductions or exemptions. Additionally, it assumes no one files jointly. It is meant to give a summary statistic for the state's average marginal tax rate.

Figure A2: State SNAP Growth vs. State Employment Growth: 2007-2010



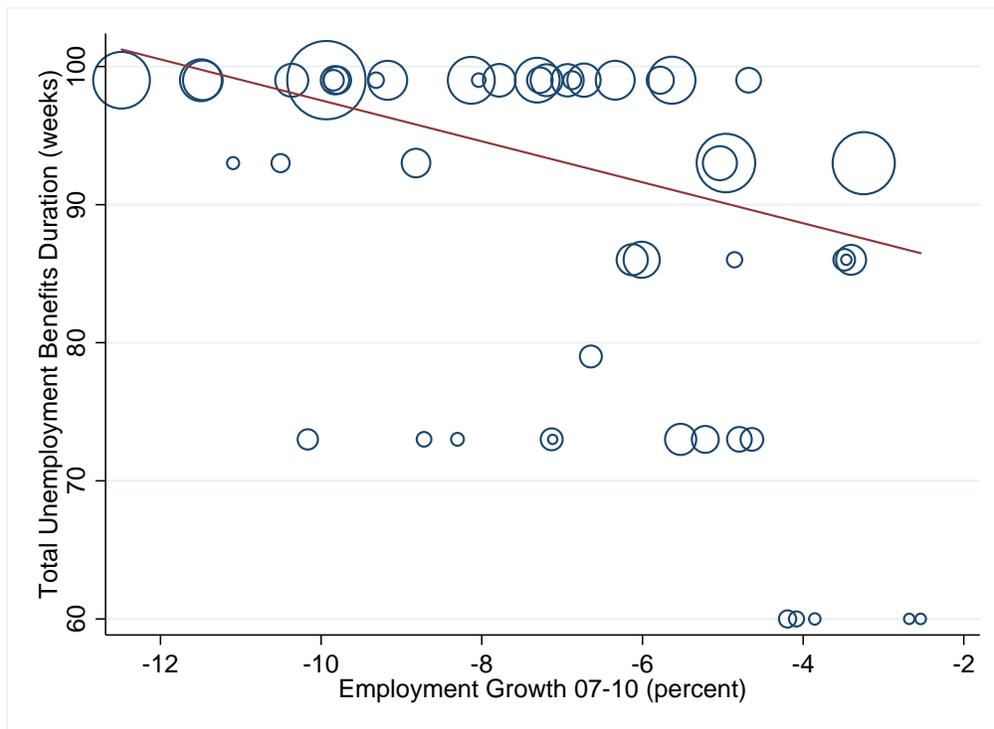
Note: Figure shows the change in SNAP payment growth per recipient at the state level between 2007 and 2010 against employment growth in the state during the same period. Employment growth comes from the BLS and is defined in the text. SNAP growth per recipient was collected from <http://www.fns.usda.gov>

Figure A3: State HAMP Take-Up vs. State Employment Growth 2007-2010



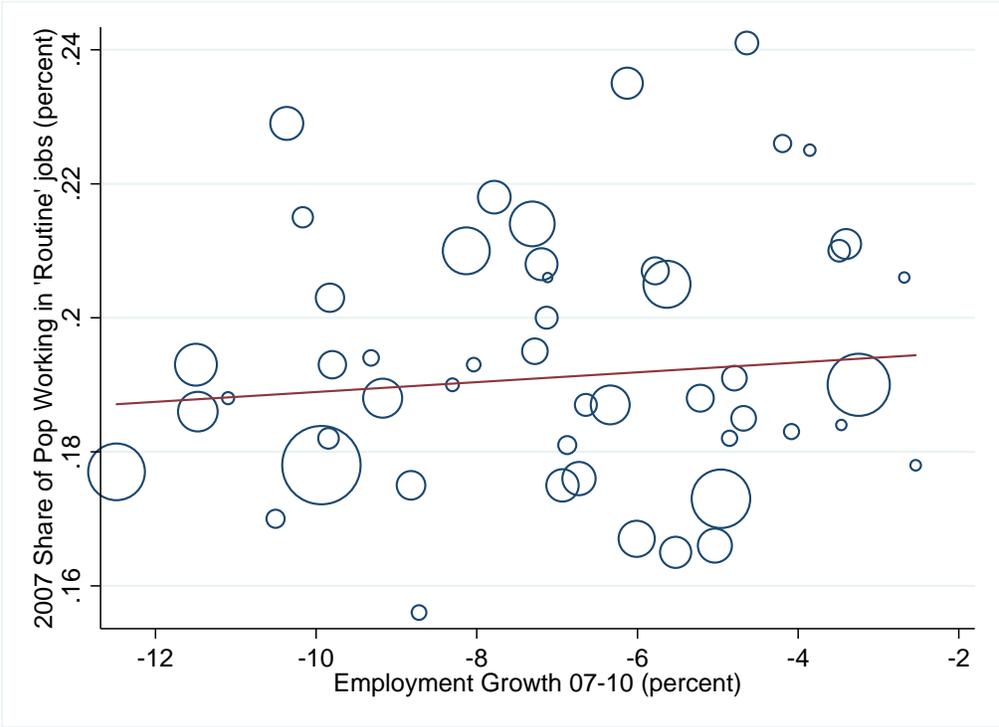
Note: Figure shows number of households participating in HAMP programs in 2010 against employment growth in the state during 2007-2010. Employment growth comes from the BLS and is defined in the text. HAMP participation comes from <http://www.treasury.gov>

Figure A4: Max Unemployment Benefit Receipt in 2010 vs. State Employment Growth 2007-2010



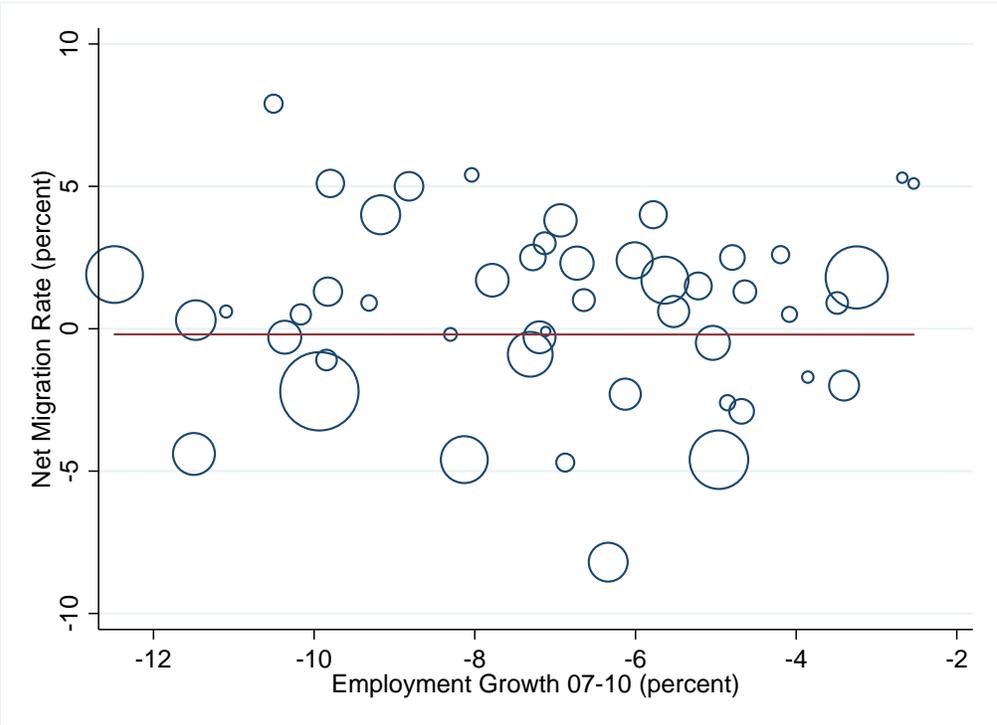
Note: Figure shows the maximum number of unemployment benefits allowed in state in 2010 against employment growth in the state during 2007-2010. Employment growth comes from the BLS and is defined in the text.

Figure A5: Routine Share of Employment in 2007 vs. State Employment Growth 2007-2010



Note: Figure shows the routine share of employment in the state in year 2007 against employment growth in the state during 2007-2010. Routine employment is defined as anyone working in a manufacturing or administrative job. Routine employment shares are computed from the 2007 American Community Survey

Figure A6: 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.