

The Aggregate Implications of Regional Business Cycles

Martin Beraja Erik Hurst Juan Ospina
University of Chicago

May 15, 2015

Preliminary

Abstract

Inferences about the determinants of aggregate business cycles from cross-region variation is possible, but should be conducted with caution. In a model of a monetary union we make the case that regional economies differ from their aggregate counterparts in two important respects: (i) the elasticities to certain types of shocks and (ii) the magnitudes of the shocks themselves. We develop a semi-structural methodology that combines regional and aggregate data to identify the shocks determining employment, prices and wages at both the aggregate and local level as well as recovering the local and aggregate elasticities to a given shock. Moreover, we formalize conditions under which regional variation may be used in the context of our methodology to inform about aggregate business cycles. We document that US states that experienced smaller employment declines between 2007 and 2010 had larger consumer price increases, nominal wage increases and real wage increases. These cross-region patterns stand in sharp contrast with the corresponding aggregate time series patterns; a reflection of (i) and/or (ii). 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 the US. On the other hand, we find that “demand” shocks explain most of the observed dynamics across states. Finally, we estimate that the local elasticity of employment to a “demand” shock is larger than the aggregate elasticity to the same shock. These results cast doubts on a large and growing literature only using cross-region variation to explain aggregate fluctuations.

1 Introduction

It is tempting to think that regional economies are smaller counterparts of the aggregate economy to which they belong. In fact, there is a large and growing literature that uses regional variation to learn about the determinants of aggregate economic variables.¹ We argue that this inference is less than straightforward because regional economies differ from their aggregate counterparts in two important respects. First, the local elasticity to a given shock may differ from the aggregate elasticity to the same shock because of general equilibrium effects.² For example, if either monetary or fiscal policy endogenously respond to aggregate variables, the aggregate employment, price and wage response to a given shock may be much smaller than the local response. Likewise, factor and product mobility may also result in aggregate elasticities to a given shock being smaller than local elasticities. Second, the type of shocks driving most of the regional variation may be different than the shocks driving most of the aggregate variation. For instance, some shocks may cause a positive covariance between employment and prices while other shocks may cause a negative covariance between employment and prices. If the aggregate covariance between employment and prices differs from the regional covariance it may just be that the combination of the shocks driving the aggregate time series data are different than the shocks causing the cross-region variation. Since the aggregate effects of a given shock almost always get differenced out when making inferences using cross-region variation, it is difficult to use regional variation to uncover the forces that are important in shaping the evolution of aggregate economic variables.

In this paper, we develop a methodology that combines regional and aggregate data to identify the shocks determining employment, prices and wages at both the aggregate and local level as well as recovering the local and aggregate elasticities to a given shock. 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 within the U.S.. The reason that we conclude that demand shocks cannot explain the bulk of the employment decline during the Great Recession is that we estimate wages are fairly flexible using cross region data. In contrast with the aggregate results, we find that “demand” shocks explain most of the observed employment, price and wage dynamics across states. These results suggest that only using cross-region variation to explain aggregate fluctuations is insufficient when some shocks do not have a substantive regional component. Lastly, we quantify that the local employment elasticity to a local demand shock is larger than the aggregate employment elasticity to a similarly sized aggregate demand shock. These results suggest that even when the aggregate and regional shocks are the same, it is hard to draw inferences about the aggregate economy using regional variation.

We begin the paper by documenting a series of new facts about the variation in prices and wages across U.S. states during the Great Recession. To do this, we use data from

¹See, for example, Autor et al (2013), Charles et al (2015), Hagedorn et al (2015), Mehrotra and Sergeyev (2015), Mian and Sufi (2014) and Mondragon (2015).

²We use the term elasticity to refer to the cumulative response to a shock over a fixed time interval.

Nielsen's *Retail Scanner Database* to compute price indices for each U.S. state. As we discuss in detail below, the *Retail Scanner Database* (RSB) includes prices and quantities for given UPC codes at over 40,000 stores at a weekly frequency from 2006 through 2011. In 2011, the RSB includes \$236 billion in sales. Most of the data come from grocery, pharmacy and mass merchandising stores. We show that an aggregate price index created with this data matches the BLS's Food CPI nearly identically. While the price indices we create from this data are based mostly on consumer packaged goods, we show how under certain assumptions the indices can be scaled to be representative of a composite local consumption good. Likewise, we use data from the U.S. Census's *American Community Survey* (ACS) to make composition adjusted nominal wage indices for each U.S. state during the 2006 to 2011 period. Using these indices, we show that states that experienced smaller unemployment increases (employment declines) between 2007 and 2010 had much larger consumer price increases, much larger nominal wage increases and larger real wage increases.³

The cross region patterns that we document stand in sharp contrast with the aggregate time series patterns for prices and wages during the same time period. As both aggregate output and employment contracted sharply within the U.S. during the 2007-2012 period, aggregate consumer price growth and aggregate nominal wage growth remained robust.⁴ 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.⁵ Recently, a literature has emerged trying to explain the missing disinflation and the missing wage declines during this time period.⁶ The key points we wish to make with these new facts is that while aggregate wages (composition adjusted) appear "sticky" during the Great Recession using time series variation, local wages (composition adjusted) were strongly correlated with measures of local economic activity using cross-state variation. Likewise, while aggregate price growth was unrelated to aggregate employment growth during the Great Recession, local price growth and local measures of employment were strongly correlated.

Having documented the contrasting behavior of aggregate and regional variation in prices, wages and employment within the U.S. during the Great Recession, we ask two questions. Were the aggregate and regional patterns different because the underlying shocks were the same but the elasticities differed because of general equilibrium effects? Or, were they different because the shocks that drove the cross-region variation were just different than the shocks that drove the aggregate time series? We propose a

³Although these local price and wage indices will be necessary for our procedure to identify both the aggregate and local shocks, we view the creation of these indices as an innovation in their own right which could be useful to researchers in variety of applications. Our wage and price indices will be posted on our webpage for use by other researchers.

⁴The one exception was during 2008Q4 and 2009Q1 where the aggregate CPI fell sharply. This decline was concentrated in energy prices.

⁵See, 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).

⁶See, for example, Del Negro et al. (2014).

semi-structural methodology that allow us to both answer these questions and describe conditions under which certain aspects of the observed regional variation may be used to inform about aggregate business cycles.

We start by describing a simple model of a monetary union with many islands linked by trade in intermediate goods, used in the production of non-tradable final consumption goods, and a risk-free asset. The nominal interest rate on this asset follows a rule that endogenously respond to aggregate variables and is set at the union level. Labor is the only other input in production, which is not mobile across islands. We allow for multiple shocks potentially having both an aggregate and local component. Furthermore, we assume that nominal wages are only partially flexible. This is the only nominal rigidity in the model.⁷ We show that, under relatively few assumptions, the log-linearized economy aggregates allowing us to study the aggregate and local behavior separately. Moreover, we show that the aggregate and local equilibria can be represented as a finite vector autoregression (VAR). Then, we formalize our initial intuition: the local elasticity of prices, wages, and employment to the local component of a given shock differs from the aggregate elasticity of prices, wages and employment to the aggregate component of that same shock. Within the model, there are two forces that make the local elasticities different from the aggregate elasticities: the endogenous part of the nominal interest rate rule and the possibility to substitute labor for intermediate goods and/or shift labor across sectors at the local level. The former gets differenced out from the cross-region variation but shows up in the aggregate elasticities, while the opposite is true for the later.

Our next goal is to estimate the magnitudes of both the aggregate and local shocks and to quantify the local and aggregate elasticities of employment, wages, and prices to a given shock. To do this, we estimate both aggregate and local VARs. Our estimation procedure allows for a more general class of models than our simple illustrative model but the endogenous variables and shocks are similar. However, to identify the aggregate VAR, we take the simple model's aggregate wage setting equation seriously. Given our model assumptions, there are only two parameters in the aggregate wage setting equation: the Frisch elasticity of labor supply and a wage stickiness parameter. Assuming the aggregate wage setting equation holds and the Frisch elasticity of labor supply and wage stickiness parameters are known, we show that the aggregate VAR is completely identified with no additional assumptions. The reason is that by specifying a structural equation we can decompose certain linear combinations of the reduced form shocks into known linear combinations of the structural shocks. This, together with usual orthogonalization conditions, allow us to identify the impulse response matrix and shock realizations. Likewise, we show that specifying this structural equation also helps us identify the local VAR. We view this identification scheme as an additional contribution of our paper and as part of a growing literature developing "hybrid" methods that, for instance, constructs optimal combinations of econometric and theoretical

⁷As we show below, local prices respond quickly to changes in local employment while local wages respond sluggishly. For this reason, we model wages as being rigid as opposed to prices.

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).

The broad set of shocks we aim to identify in the VAR (and are included in our simple theoretical model as well) are ones that have been emphasized in the literature as being potentially important in explaining 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 or a decline in household wealth. For example, such shocks have been proposed by Eggertsson and Krugman (2012), Guerrieri and Lorenzoni (2011) and Mian and Sufi (2014) as an explanation of the 2008 recession. The model also includes a shock to the nominal interest rate rule. This is a separate "demand" shock within the model. Given our VAR, we will only be able to identify the net effect of the discount rate shock and the shock to the nominal interest rate rule. Additionally, we allow for shocks to firm's productivity. We refer to such shocks as "supply" shocks. These shocks are modeled as pure productivity shocks. However, they could also be interpreted as anything that changes firms' demand for labor or as shocks to the firm's mark-up. As will be seen in our model set up, the marginal cost shock or mark-up shock interpretation is isomorphic to our productivity shock interpretation. Without bring in additional data, we cannot distinguish between the different interpretations in our VAR. For example, credit supply shocks to firms, such as those proposed by Gilchrist et al (2014), would be similar to our productivity shock. Finally, we include 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 (e.g., Mulligan (2012) or as a reduced form representation of a skill mismatch story within the labor market (e.g., Charles et al. (2013)).

As mentioned, the shock identification procedure requires parametrizing the structural wage setting equation. We argue that the regional data on prices, wages and employment during the 2006-2011 period can be used to estimate the Frisch elasticity of labor supply and the amount of wage stickiness (which are the only parameters in this equation). In order for regional data to be used to parameterize the local and aggregate wage setting equations we need one of the following two assumptions to hold: (1) the local component of the taste for leisure shock is zero for all regions during the Great Recession or (2) variation in either local discount rate shocks or local productivity shocks can be identified. The local wage setting equation is akin to a local labor supply curve with sticky wages. These assumptions basically state that the parameters of the local labor supply curve can be identified if there are no local shocks to labor supply or if shocks to local labor demand can be isolated. Clearly, over different periods of time, these assumptions may not hold. However, we provide evidence that the local component of the taste for leisure shock may be small during the Great Recession and that housing price variation during the 2007-2009 period can help us isolate movements

in local labor demand.⁸ It is worth noting that if the above assumptions hold and if the parameters in the wage setting equation are the same in the aggregate and at the local level, cross-region variation may be used to discipline the aggregate economy's response to certain type shocks via our VAR identification procedure. Assuming these conditions hold and using state level data during the 2006-2011 period, we estimate a range of Frisch elasticities of labor supply from 1 to 2 across our various specifications. Additionally, we estimate only a modest amount of wage stickiness.

With the parameterized aggregate and local wage setting equations, we use the VAR methodology described above to estimate the shocks driving aggregate and local employment, prices, and wages during the Great Recession. Using the impulse responses, we also estimate the relevant aggregate and local employment, price and wage elasticities to each individual shock. The results suggest that something akin to a discount rate shock is driving essentially all the cross-region variation in employment, wages and prices during the Great Recession. For those papers that view the world through cross-region variation it looks like demand (discount rate) shocks are very important. However, at the aggregate level, the discount rate shock only explained roughly 30 percent of the employment decline during the 2008-2010 period and essentially none of the decline during the 2008-2012 period. Instead, we estimate that a combination of discount rate shocks and productivity/mark-up shocks are important for explaining the dynamics of aggregate employment during the Great Recession. We estimate that if not for the productivity/mark-up shocks, both aggregate prices would have declined more. The productivity/mark-up shock was putting downward pressure on employment and upward pressure on employment. The discount rate shock was putting downward pressure on both employment and prices. The fact that prices and employment did not move together in the aggregate time series data during the Great Recession was a result of two shocks that had relatively offsetting effects on prices. The taste for leisure shock in the aggregate is important for explaining why aggregate wages did not fall during the 2008-2012 period.

To provide intuition, we highlight that our estimated amount of wage stickiness from the cross-region data is key to our empirical results. To get large and persistent effects of aggregate demand shocks in our model, wages need to be very sticky. Using the cross-region variation, we document that wages are rather flexible. Again, while we do estimate that wages are not perfectly flexible from the cross-region data during the Great Recession at an annual frequency, we show that the amount of stickiness we estimate is not large enough to make aggregate demand shocks have large and persistent effects on employment. We document that we would need a substantially higher amount of wage stickiness in order for aggregate demand shocks to be the primary explanation for the weak employment situation within the U.S. during the Great Recession. Such wage stickiness, however, is at odds with the amount of wage adjustments seen using cross-region variation.

In summary, our results explain that the co-movement of prices and employment on

⁸This is a similar assumption to Mian and Sufi (2014) or Mehrotra and Sergeyev (2015).

the one hand and wages and employment on the other hand differ between the aggregate time series and the cross-region variation both because the shocks differ and because the elasticities to a given shock differ. In particular, we show that the local elasticity of employment to discount rate shocks is much larger than the aggregate elasticity to the same shock. The reason is that the nominal interest rate rule mitigates the aggregate effects but gets differenced out from the cross-region analysis. We also find that the productivity/mark-up shock is primarily an aggregate shock. Again, because this shock is aggregate, it gets differenced out in the cross-region analysis. The combination of these two factors implies that identifying shocks by looking at the covariance of prices, employment and wages at the local level can be problematic when trying to infer the drivers of the aggregate economy.

It is worth highlighting the value of the contribution of creating local price and wage indices with respect to this project. The local price and wage series serve two purposes. First, the local price and wage data are necessary to estimate the parameters of the aggregate and local wage setting equations. These parameters are key to the restrictions imposed on our aggregate and local VARs. While there are many existing estimates of the Frisch elasticity of labor supply, there are very few estimates of wage stickiness. As noted above, the fact that we are estimating wages to be only moderately sticky at the local level disciplines the potential role for demand shocks in explaining the persistent decline in employment during the 2008-2011 period. Second, the local price and wage data are necessary to infer the local shocks and the local elasticities to a given shock. Without the local price data, for example, it is impossible to tell whether the cross sectional differences in employment changes during the Great Recession are due to cross sectional differences in "demand" (as in Mian and Sufi 2014) or due to cross sectional differences in "supply" (as in Mehrotra and Sergeyev (2015)).

Finally, we want to stress that there are two limitations of our analysis. First, given our methodology, we do not point to what specific "aggregate demand" shock or "aggregate supply" shock drove the Great Recession. For example, we cannot distinguish between a tightening of household borrowing constraints vs. households wanting to deleverage. Likewise, we do not distinguish whether the supply shock is due to a decline in productivity or an increase in mark-ups. Despite that, we think our conclusions are important in the extent that we quantify the relative importance of broad types of shocks. This finding will hopefully guide researchers to focus on exploring the origins of these broad shocks in future research. Second, the VAR identification procedure is inherently linear. Any non-linearity, perhaps due to downward sticky nominal wages or the zero lower bound, are not easily accommodated.⁹

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

⁹We return to a more detailed discussion of the zero lower bound in the conclusion of the paper.

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. With respect to the latter innovation, our paper is similar in spirit to Nakamura and Steinsson (2014).

Second, our paper contributes to the recent literature 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.¹⁰ 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 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,

¹⁰There 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.

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 And Wage Indices

2.1 Local Price Indices

2.1.1 Price Data

To construct state level price indices we use the *Retail Scanner Database* collected by AC Nielsen and made available at The University of Chicago Booth School of Business.¹¹ 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.¹²

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

¹¹The 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/>.

¹²It should be noted that Walmart only recently started sharing their retail data with Nielsen. As a result, the data through 2011 does not include any Walmart stores.

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. Online Appendix Table A1 shows summary statistics for the scanner data for each year between 2006 and 2011 and for the sample as a whole.

2.1.2 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.¹³ While we briefly outline our procedure in this sub-section, the full details of the procedure are discussed in the Online Appendix that accompanies our paper. Our scanner price indices are built in two stages. In the first stage, we aggregate the prices of goods within the roughly 1,000 categories described above. For our base index, a good is either a given UPC or a given store-UPC pair. In the latter case, a UPC 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.¹⁴ 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.

Specifically, for each category, we compute:

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

where $P_{j,t,y,k}$ is category level price index for category j , in year t , with base year y , in geography k . For our analysis, geographies will either be U.S. states or the country as a

¹³There is a large literature discussing the construction of price indices. See, for example, Diewert (1976). Cage et al (2003) discuss the reasons behind the introduction of the BLS's Chained Consumer Price Index. Melser (2011) discuss problems that arise with the construction of price indices with scanner data. In particular, if the quantity weights are updated too frequently the price index will exhibit "chain drift". 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).

¹⁴In 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.

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.¹⁵ Fixing quantities at a lagged level implies that the price changes we document below with changing local economic conditions is not the result of changing household consumption patterns.

The second stage of our price indices also 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. 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.¹⁶

To benchmark our scanner price index, we compare our scanner price index for the aggregate U.S. (where each good is treated as a UPC-store pair) to the BLS's CPI for food. We chose the BLS Food CPI as a benchmark given that most of the goods in our database are food data.¹⁷ Figure 1 shows that our scanner price index matches nearly identically the BLS's Food CPI. 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

¹⁵For example, 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 procedure of fixing quantities at a lagged level is similar in spirit to the way the BLS builds category-level first stage for their price indices.

¹⁶One issue discussed in greater depth within the Online Appendix is how we deal with missing data when computing the price indices. Seasonal goods, the introduction of new goods, and the phasing out of existing goods means that missing data on month to month price changes occurs. When computing our price indices, we restrict our sample to only include (1) goods that had positive sales in the prior year and (2) goods that had positive sales in every month of the current year. Online Table Appendix A1 shows the percent of sales included within the price index for each sample year.

¹⁷Not 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.

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.

2.1.3 Computing Regional Inflation Rates Using Retail Data

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

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.

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:

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

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

Given these assumptions, we can transform the variation in the grocery sector prices

that we identify into variation in the broader consumption basket across regions. Taking logs and differencing across regions we get that the variation in log-prices of the composite good between two regions 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 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.¹⁸ Putting the two estimates together, we adjust the variation in the regional inflation rates computed using the goods in our database by a factor of 2 ($0.6/0.3$).¹⁹

We want to stress that the adjustment factor plays a minimal role in our formal quantitative work below estimating local and aggregate shocks and local and aggregate elasticities to a given shock. Our base assumption is that the grocery sector has a larger tradable share than the average good. We have explored a variety of adjustment factors between 0 and 3 and the quantitative implication of our estimation procedures are relatively robust. The main importance of the scaling factor is in the descriptive patterns of how real wages move with local economic conditions that we document in the next section. The higher the adjustment factor, the more muted are real wage movements with local measures of economic activity during the Great Recession.

¹⁸The 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 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.

¹⁹The 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.

2.2 Local Wage Indices

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.²⁰

The composition of 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 2 shows aggregate nominal and real composition adjusted log wages during the 2000-2012 period using the above method for the country as a whole. To get real wages, we deflate nominal wages by the aggregate June CPI-U with 2000 as the base year. 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. As seen from Figure 2, nominal wages increased slightly and real wage growth did not

²⁰Total 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 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.

seem to break trend during the Great Recession. The puzzle has been why wages did not decline despite the very weak aggregate labor market.

3 Regional Variation in Prices and Wages During the 2000s

3.1 Regional Variation in Prices During the 2000s

Figure 3 and Table 1 explore the extent to which our regional scanner price index is correlated with measures of local economic activity. Specifically, Figure 3 plots the percentage point change in the state’s average unemployment rate between 2007 and 2010 against the percent change in the state’s scanner price index between 2007 and 2010.²¹ For the results in Figure 3, we use our price index where a good is a given UPC within a state (as opposed to a UPC-store pair). Additionally, Figure 3 shows the variation in P^r . In other words, the results in this Figure are not adjusted for the fact that the tradable share of the goods in our sample differs from the tradable share in the composite consumption good. The unemployment rate data come from the BLS’s Local Area Unemployment Statistics. Each observation 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 3, $\Delta X_{k,07-10}$ equals the percentage point change in the state unemployment rate between 2007 and 2010.

Figure 3 shows that there is a negative relationship between the change in the state’s unemployment rate between 2007 and 2010 and the change in the state’s price level between 2007 and 2010. The estimate of β_1 for this specification is -0.46 (standard error = 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 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 of 2, 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).

²¹Our 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.

Table 1 shows different estimates of β_1 from the above regression with different measures of changing local economic activity ($\Delta X_{k,07-10}$). 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 1 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 3). Other local economic measures in the subsequent rows include the percent change in state per-capita nominal GDP, the percent change in state per-capita total hours worked, the percent change in state housing prices, and the percent change in the state employment rate.²² Additionally, in some of the empirical work below, we isolate movements in local employment that were correlated with local housing price changes. The last two rows of Table 1 isolate the relationship between local price growth and local unemployment changes (row 6) and local employment changes (row 7) that are correlated with changes in local house price growth.

As seen from the results in Table 1, all measures of the change in economic activity are correlated with the change in local prices. As local economic conditions deteriorated during the Great Recession (higher change in the unemployment, lower growth rate in the employment rate, lower house price growth, lower change in hours and GDP per capita), the lower the price inflation during Great Recession. 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 4 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 3 compared long differences in both the unemployment rate and the inflation rate. In Figure 4, 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.

²²The 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.

In Figure 4 , 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 4 , 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 4 is the difference between the average retail price level in the low unemployment change states relative to the high unemployment change states.²³ 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. The simple correlation between the two series in Figure 4 is -0.93. Moreover, there does not seem to be a delay between when the unemployment rate changed and when the prices changed. For example, 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.²⁴

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

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

Figure 5 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 also 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 2 (column 1) we show that the growth in local nominal wages was highly correlated

²³To 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.

²⁴In the Online Appendix that accompanies the paper, we discuss how our scanner price indices compare to the BLS's local price indices. The BLS releases price indices for 27 MSAs. We show that even within the BLS price indices during the 2007-2010 period, a 1 percentage point increase in the local unemployment rate was associated with a 0.34 percentage point decline in the food inflation rate (over the three year period). This is nearly identical to the relationship in Figure 3. The relationship between the local unemployment rate and the total inflation rate was even stronger. These relationships give us confidence in our scanner index and allows us to have local price indices for all U.S. states (not just a small subset of MSAs).

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 6 shows the relative difference in state unemployment rates and the adjusted nominal wage index during the 2000s for high and low unemployment change states during the Great Recession. The high and low unemployment change states are defined similarly as in Figure 4. The figure is similar in spirit to Figure 4 except for three things: (1) we are plotting the adjusted nominal wage index instead of the scanner price index, (2) the frequency of the wage data is annual as opposed to monthly, and (3) we show the patterns from 2000-2012 instead of 2006-2011. We normalize the adjusted wage index so that it is 1 in 2006 for all states. The high and low unemployment change states during the Great Recession had roughly similar wages and unemployment rates during the 2000-2006 period. However, after 2007, the adjusted nominal wage index and the unemployment rates started to diverge. For example, in 2010, adjusted nominal wages were 3.5 percent higher in the low unemployment states relative to the high unemployment change states compared to their 2006 levels. One main insights from Figure 6 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 is different than the scanner price indices highlighted in Figure 4 in that the scanner price differences stabilized once the differences in real activity stabilized. The fact that nominal wages continue to diverge motivates our modeling assumption of sticky nominal wages as opposed to sticky prices.

The second and third columns of Table 2 show the coefficient on the change in local economic activity between 2007 and 2010 from a regression of real wage growth in a given state during that time period on the change in local economic activity. In column 2, we compute local real wages by deflating local nominal wage growth by the growth in the local scanner price index. In column 3, we compute local real wages by deflating local nominal wage growth the growth in the prices of a composite local consumption good. As discussed above, we scale the growth in the scanner price index by a factor of two to account for the fact that grocery/mass merchandising goods have a higher tradable share than the composite consumption good. Not surprising, the coefficients in column 2 of Table 2 are roughly equal to the coefficients from column 1 of Table 2 less the coefficients from columns 2 and 4 of Table 1. In all specifications, real wages fell as measures of local economic conditions worsened. For example, a 1 percentage point increase in the unemployment rate was associated with a 0.5 and 0.9 percentage point decline in real wage growth during the 2007 to 2010 period (depending on the scaling factor).

In summary, despite aggregate wages appearing sticky during the Great Recession, there is a strong relationship between local prices and local real activity across U.S. states during this time period.

4 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.²⁵ Finally, we assume labor is mobile across sectors but not across islands.²⁶ Throughout we assume that parameters governing preferences and production are identical across islands and these only differ, potentially, in the shocks that hit them.

4.1 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

²⁵We 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.

²⁶We 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...

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

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

4.2 Sticky wages

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

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

Given our assumption on household preferences, $P_{kt}e^{\epsilon_{kt}}(N_{kt})^{\frac{1}{\phi}}$ is the marginal rate of substitution between labor and consumption and the parameter λ measures the degree of nominal wage stickiness. In particular, when $\lambda = 1$ wages are fully flexible and when $\lambda = 0$ they are fixed. This implies that workers will be off their labor supply curves whenever $\lambda < 1$. A similar specification has been used by Shimer (2009) and, more recently, by Midrigan and Philippon (2011). Shimer (2009) argues that in labor market search models there is typically an interval of wages that both the workers are willing to accept and firms willing to pay. To resolve this wage indeterminacy he considers a wage setting rule that is a weighted average of a target wage and the past wage. The target wage in our case is the value of the marginal rate of substitution.

Popular alternatives in the literature include the wage bargaining model in the spirit of Hall and Milgrom (2008) as in Christiano, Eichenbaum and Trabandt (2013); and 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 that is absent in the former.²⁷ In fact, this wage setting rule can be derived from the monopsonistic competition setup in the case where agents are myopic about the future; or the labor market search setup in the special case where firms make take it or leave it offers and the probability of being employed in the future is independent of the current employment status.

4.3 Equilibrium

An equilibrium is a collection of prices $\{P_{kt}, W_{kt}, Q_t\}$ and quantities $\{C_{kt}, N_{kt}, B_{kt}, N_{kt}^x, N_{kt}^y, X_{kt}\}$ 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

²⁷The same would be true if the target wage would have a forward looking component. This is the case in Shimer (2009) where the target wage is determined through axiomatic Nash wage bargaining.

profit maximization and such that the following market clearing conditions hold:

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

4.4 Shocks

We assume the exogenous shocks follow an AR(1) process, with an identical autoregressive coefficient across islands (and sectors in the case of productivity), and that the innovations are iid, mean zero, random variables with an aggregate and island specific component. Formally,

$$\begin{aligned}
z_{kt}^y &= \rho_z z_{kt-1}^y + \sigma_z u_t^y + \tilde{\sigma}_y v_{kt}^y \\
z_{kt}^x &= \rho_z z_{kt-1}^x + \sigma_z u_t^x + \tilde{\sigma}_x v_{kt}^x \\
\gamma_{kt} &= \rho_\gamma \gamma_{kt-1} + \sigma_\gamma u_t^\gamma + \tilde{\sigma}_\gamma v_{kt}^\gamma \\
\epsilon_{kt} &= \rho_\epsilon \epsilon_{kt-1} + \sigma_\epsilon u_t^\epsilon + \tilde{\sigma}_\epsilon v_{kt}^\epsilon
\end{aligned}$$

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

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 *Productivity/Markup*, *Discount rate and Leisure shocks* respectively. These are the innovations that the econometric procedure aims to identify. Analogously, v_{kt}^y , v_{kt}^x , v_{kt}^γ , v_{kt}^ϵ are the *Regional shocks*. The interpretation of the Leisure and Productivity/Markup shocks is relatively straightforward given our model environment. They are shifters of households and firms' labor supply and demand schedules respectively. On the other hand, what we identify as a Discount rate 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.

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 assump-

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

tion is that the islands are identical in the steady state and that price and wage inflation are zero. The last assumption is that the joint distribution of island-specific shocks is such that its cross-sectional summation is zero. If K , the number of islands, is large this holds in the limit because of the law of large numbers. We log-linearize the model around this steady state and show that it aggregates up to a representative economy where all aggregate variables are independent of any cross-sectional considerations to a first order approximation.²⁹ We denote with lowercase letters the log-growth rates of variables. Also, variables without a 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. For simplicity, we assume it only depends on nominal wage inflation. In log-linearized form it can be written as: $i_{t+1} = \varphi_w w_t + \mu_{t+1}$ 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.

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 $q_t = i_t = 0 \forall t$.*

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, 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.

Proof. For a given sequence of aggregate variables $\{p_t, n_t, q_t\}$, the following equations characterize the log-linearized equilibrium

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

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

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

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

$$\{\tilde{z}_t^y, \tilde{z}_t^x, \tilde{\gamma}_t, \tilde{\epsilon}_t\},$$

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

$$\tilde{w}_t = \tilde{z}_t^x - \tilde{z}_{t-1}^x - (1 - \theta)\tilde{n}_t - (1 - \theta)\frac{N - N^x}{N}(\tilde{n}_t^x - \tilde{n}_t^y)$$

$$0 = \mathbb{E}_t(\tilde{m}u_{t+1} - \tilde{p}_{t+1} + \tilde{\gamma}_{t+1})$$

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

$$\tilde{w}_t = \lambda(\tilde{p}_t + \tilde{\epsilon}_t - \tilde{\epsilon}_{t-1} + \frac{1}{\phi}\tilde{n}_t) + (1 - \lambda)\tilde{w}_{t-1}$$

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 Aggregate v. local shock elasticities

This section compares the island level and aggregate response of employment to an unanticipated discount rate shock. We consider the special case where there is an endowment of the tradable good and no labor is used in its production, i.e. $\theta = 0$. Focusing on a discount rate shock in this special case makes the comparison very transparent. We let $\zeta_0^{agg} \equiv \frac{dn_0}{d\gamma_0}$ and $\zeta_0^{reg} \equiv \frac{d\tilde{n}_0}{d\tilde{\gamma}_0}$ be the impact elasticities to the shock. By solving for the recursive laws of motion in equilibrium we obtain,

$$\zeta_0^{agg} = - \frac{(1 - \lambda)\frac{\rho\gamma}{1 - \rho\gamma}}{\hat{\sigma} + \frac{\lambda(1 - \alpha + \frac{1}{\phi})}{1 - \rho\gamma}(\varphi_w - 1)}$$

$$\zeta_0^{reg} = - \frac{(1 - \lambda)\frac{\rho\gamma}{1 - \rho\gamma}}{\hat{\sigma} - \left(1 - \frac{\sigma(1 - \lambda)}{1 - \alpha}\right)\beta}$$

where $\hat{\sigma} = \frac{\sigma(1 - \lambda)}{(1 - \alpha)} \left(\alpha - \frac{1 + \phi}{\phi \frac{Y}{N}} \right) + (1 - \alpha) + \frac{\lambda}{\phi}$.

The nominal interest rate rule endogenous response to aggregate nominal wage growth φ_w reduces the aggregate employment impact elasticity (in absolute value) to an unanticipated discount rate shock. Additionally, if $\frac{\sigma(1 - \lambda)}{1 - \alpha} < 1$ the possibility to substitute labor for intermediate goods in the production of final consumption goods increases the regional employment elasticity (in absolute value) to the shock. Both forces make the aggregate elasticity smaller than the regional response.

It is also interesting to compare how the elasticities change with the degree of nominal wage stickiness. Our identification procedure, will allow us to do this exercise when we

estimate the impulse response to a discount rate shock. When $\frac{1+\phi}{\phi} - (\alpha + \beta) > 0$ and $\phi_w - \rho_\gamma > 0$ both elasticities are decreasing (in absolute value) in λ . In particular, employment does not respond to discount rate shocks at all in the limit when wages are perfectly flexible ($\lambda \rightarrow 1$).

5 Elasticities and Shocks: A Semi-structural Approach

5.1 Estimating aggregate shocks and elasticities

The recursive solution to the equilibrium system of equations in Lemma 1 can be written in reduced form as a VAR(∞) in $\{p_t, w_t, n_t\}$.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 shocks. Hence, identification of the shocks is identification of these matrices.

This reduced form vector autoregression representation is consistent with a much more general class of models than the one we have introduced in the previous section. For example, the labor demand equation in Lemma 1 could have a forward looking component if there were adjustment costs to employment. From here on we will be working within a subset of these more general class of models. A subset of models where the wage setting equation in log-linearized form holds,

$$w_t = \lambda(p_t + \epsilon_t - \epsilon_{t-1} + \frac{1}{\phi}n_t) + (1 - \lambda)w_{t-1}$$

This will be our key identification assumption. We will show that if this one equation is specified, we can identify the aggregate shocks and elasticities while being agnostic about the remaining structural equations describing the economy. This is because the structural equation imposes several linear constraints that the reduced form errors must satisfy.

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

We now derive the identification restrictions that will allow us to estimate Λ and the shocks. Applying the conditional expectation operator $\mathbb{E}_{t-1}(\cdot)$ on both sides of the wage

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(1,2) and hence as a VAR(∞) if the moving average part of the process is invertible.

setting equation and constructing the reduced form expectational errors we obtain,

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

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

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

where ρ_1 is the matrix collecting the first order autoregressive coefficients in the reduced form VAR.

Equations (5) and (6) have to hold for all realizations of the shocks. In particular, equation (5) gives us two linear restrictions in the elements of Λ for given parameters in the wage setting equation when there are either contemporaneous discount rate or productivity/markup shocks. Moreover, from equation (6) we can obtain one extra linear restriction that holds when there is a lagged discount rate shock or a lagged productivity/markup shock. These three restrictions, together with six restrictions coming from the orthogonalization of the shocks are sufficient to exactly identify all nine elements in the Λ matrix.³¹ Intuitively, (5) allows us to “separate” the leisure shock from the discount rate and productivity/markup shocks combined; and (6) “separates” the discount rate from the productivity/markup shock. For completeness, the matrix Λ solves the system:

$$\begin{aligned} \begin{bmatrix} \lambda & -1 & \frac{\lambda}{\phi} \end{bmatrix} \Lambda \begin{bmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} &= [0 \ 0] \\ \left(\begin{bmatrix} \lambda & -1 & \frac{\lambda}{\phi} \end{bmatrix} \rho_1 + \begin{bmatrix} 0 & 1 - \lambda & 0 \end{bmatrix} \right) \Lambda \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} &= 0 \\ \Lambda \Lambda' &= V \end{aligned}$$

It is worth noting that there is a sense in which the shocks are not identified. The issue arises because the above procedure does not allow us to label which of the shocks that do not appear in the wage setting equation corresponds to u_t^z or u_t^γ . Basically, the linear restrictions from equation (5) and (6) are identical for both discount rate and productivity/markup shocks. A solution to this labeling problem is to use the theoretical co-movement on impact of employment, wages and prices after a u_t^γ and u_t^y shock,

³¹Actually, equation (6) gives two independent restrictions for each shock. In this sense, the system is over identified and the extra independent condition can be used ex-post to test the model via this over identifying restriction.

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 discount rate shock moves prices and employment in the same direction on impact.

5.2 Estimating the Regional Shocks and Elasticities

The procedure for estimating regional shocks and elasticities follows a similar logic than the one for aggregate shocks.

The recursive solution to the equilibrium system of equations in Lemma 2 can be written in reduced form as a VAR(∞) in $\{\tilde{p}_t, \tilde{w}_t, \tilde{n}_t^y\}$ when $\tilde{v}_t^\epsilon = 0$.³²

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

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

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

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

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

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

$$\left[1 \quad - \left(\frac{1}{\lambda} + \frac{N - N^y}{N\phi(1-\theta)} \right) \quad \frac{N^y}{N\phi} \right] \Lambda \begin{bmatrix} \tilde{v}_t^y \\ \tilde{v}_t^x \\ \tilde{v}_t^\gamma \end{bmatrix} + \frac{N - N^y}{N\phi(1-\theta)} \sigma^x \tilde{v}_t^x = 0 \quad (5)$$

³²In the appendix we provide supporting evidence for this assumption during the Great Recession.

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

$$\left(\begin{bmatrix} 0 & \frac{1-\lambda}{\lambda} & 0 \end{bmatrix} + \begin{bmatrix} 1 & -\left(\frac{1}{\lambda} + \frac{N-N^y}{N\phi(1-\theta)}\right) & \frac{N^y}{N\phi} \end{bmatrix} \tilde{\rho}_1 \right) \Lambda \begin{bmatrix} \tilde{v}_{t-1}^y \\ \tilde{v}_{t-1}^x \\ \tilde{v}_{t-1}^\gamma \end{bmatrix} - \frac{1}{\phi} \frac{N - N^y}{N(1-\theta)} \rho^x \sigma^x \tilde{v}_{t-1}^x = 0 \quad (6)$$

where $\tilde{\rho}_1$ is the matrix collecting the first order autoregressive coefficients in the reduced form VAR. As before, we can identify the impulse matrix $\tilde{\Lambda}$ with these extra linear restrictions.

5.3 Using Cross-region Variation to Make Inferences about the Aggregate

Under certain conditions, the regional data can be used to estimate the structural parameters of the aggregate wage setting equation. Hence, the observed cross-region variation will have aggregate implications when combined with the identification procedure described in Section 5.1. For example, different estimates of λ from the cross-section of regions will imply different impulse response matrices Λ in the aggregate and different shock decompositions.

We now present this conditions. Under the maintained assumption that both the local and aggregate labor supply curves share similar parameters we have that,

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

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 leisure 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 leisure preference. However, in order to for the estimates of ϕ and λ to be unbiased, the regional leisure shocks or the measurement error in the variables have to be uncorrelated with local price and employment growth.³³ The intuition is straight-forward. If there is no systematic variation in the wage setting equation 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.

As a rule, it is probably unrealistic to assume that there is either no regional component to leisure shocks. However, during certain periods, this 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.

³³Both employment and price growth are endogenous to the local leisure 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.

Below, we argue that the assumption that the regional component of leisure shocks was both small and uncorrelated with local price and employment growth during the Great Recession.

6 Estimating Parameters of the Wage Setting Equation Using Regional Data

Given the above assumptions, the aggregate and local wage setting equations can be expressed as:

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

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

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

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

6.1 Estimating Equation and Identification Assumptions

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

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

where $b_1 = \lambda$, $b_2 = \lambda/\phi$, $b_3 = (1 - \lambda)$, and $b_t = \lambda(u_t^\epsilon - (1 - \rho_\epsilon)\epsilon_{t-1})$. Any aggregate labor supply shocks are embedded the constant term. The local error term includes λv_{kt}^ϵ as well as measurement error for the local economic variables. We estimate this

equation pooling together all annual employment, price and wage data for years between 2007 and 2011. When estimating the above regression, we include year fixed effects (D_t). This ensures that we are only using the cross-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 above. w_{kt} , therefore, is just the log-growth rate in adjusted nominal wages within the state between year t and $t + 1$. Our measure of employment growth at the state level is calculated using data from the U.S. Bureau of Labor Statistics. The BLS reports annual employment counts and population numbers for each state in each year. We divide the employment counts by population to make an annual employment rate measure for each state. n_{kt} is the log-change in the employment rate between year t and $t + 1$. 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. Finally, in some specifications we include controls for the state's industry mix in 2007. This allows for the potential that local labor supply shocks, to the extent that they exist, may be correlated with the state's industry structure. Given that we have observations on 48 states for 4 years of growth rate data, our estimating equation includes 192 observations in our base specification. We also show results using data from 2007-2009 before the large changes in unemployment benefits extension starting in 2010.

Two additional comments are needed about our estimating equation. First, the theory developed above implies that $b_1 + b_3 = 1$. We impose this condition when estimating the cross region regression. Second, we believe our measures of local wage growth and price growth are measured with error. The measurement error, if classical, will attenuate our estimates of b_1 and b_3 . Additionally, because we are regressing wage growth on lagged wage growth, any classical measurement error in wages in year t will cause a negative relationship between wage growth today and lagged wage growth. We take these measurement error concerns seriously when estimating the above regression. Specifically, given the large sample sizes 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 Online Data Appendix that accompanies the paper. As we show in that appendix, the procedure dramatically corrects the attenuation bias from measurement error in our estimates.

In order to recover unbiased estimates of λ and ϕ via OLS, $v_{kt}^e = 0$. The assumption that there are no local labor supply shocks cannot generically be true. However, we pro-

vide some evidence suggesting that this assumption may be valid during the 2007-2011. We briefly summarize some of our findings here while a more detailed discussion can be found in the Online Data Appendix. Broadly, the labor supply shock can be thought of proxying for actual changes in household preferences, changes in government policies that discourage work (e.g., Mulligan 2012), or the skill mismatch story where workers in declining sectors face frictions in transitioning to growing sectors (e.g., Charles et al (2014) or Jaimovich and Siu (2014)). One main identification assumption when we estimate the OLS specification is that household preferences for leisure did not differentially change across states during the 2007-2011 period. In terms of policy changes, we explore the extent to which policies that discourage work changed differentially across U.S. states during the Great Recession. We examined four such policies: state income taxes, state expansion of food assistance programs, state expansion of weeks of unemployment benefits, and state policies to help underwater homeowners.³⁴ For each policy, we explore the standard deviation of changes in the policy across states and the extent to which any cross-state variation in the policy is correlated with either local employment or price growth.

As seen from Online Appendix Figures A1 and A2, there was essentially no variation in state income tax rates across states during the Great Recession nor was there any variation in food assistance benefits per recipient.³⁵ For example, while food assistance benefits per recipient increased by roughly 35 percent nationally during the 2007-2010 period, there was essentially no variation across U.S. states. Given that food assistance programs are means tested, Mulligan (2012) has argued that the expansion of food assistance programs may have discouraged labor supply during the Great Recession. Online Appendix Figures A3 and A4 show the changes in unemployment benefit duration across states and the change in Federal mortgage assistance programs (HAMP) across states. Mortgage assistance programs were also means tested potentially discouraging labor supply. As seen from Figures A3 and A4, there is some variation in these programs across states and that variation is correlated with local employment and price movements during the time period. But we wish to stress a few additional facts. First, as seen in these appendix figures, the variation that does exist was concentrated in only a handful of states. By law in 2010, weeks of unemployment benefits were tied to the state's unemployment rate. 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 Likewise, mortgage assistance programs were initially only con-

³⁴In the Online Data Appendix, we discuss in detail the specific policies we examined and the specific data sources we used to get the spatial variation in these policies.

³⁵ For our tax measure, we create an average marginal tax rate for each state during each year. For our food assistance measure, we measure SNAP benefits per recipient. SNAP is the recent version of the federal Food Stamps program. In total, SNAP benefits increased in states with higher unemployment rates. However, the increase is solely driven by the increase in the number of recipients not the increase in benefits per recipient.

³⁶States also had some discretion as to whether they opted into the program. This explains why some

concentrated in a handful of states (CA, AZ, NV, and FL). Second, and most importantly, both the extension of local unemployment benefits and federal home owner assistance programs did not occur until 2010. As a result, prior to 2009, there was no regional variation in these policies.

To account for the potential that certain government policies changed differentially across states during the Great Recession, we perform a series of robustness specifications to our base OLS specification. First, we re-estimate our key equation using only data up through 2009. Prior to 2010, there was essentially no changing policy variation across U.S. states. As we show, our estimates using data from the full 2007-2011 period are nearly identical to our estimates using data only prior to 2010. Second, we re-estimate our base specification excluding states that had less than 85 weeks of unemployment benefit extensions and then separately excluding states that had substantial Federal loan modifications during the 2010 and 2011 period. Again, excluding states that had some modest policy change post 2009 did not alter our key estimates of λ and ϕ in any substantive way.

Online Appendix Figure A5 examines the extent to which industry mix varies across U.S. states. In recent working papers, Jaimovich and Siu (2014) and Charles et al. (2015) discuss how skill mismatch stories from declining routine jobs and the manufacturing sector, respectively, could have contributed to declines in aggregate labor supply. If workers from these declining sectors do not have the skills to enter growing sectors at average wages, they may leave the labor force. This could manifest itself as a something akin to our reduced form labor supply shock. While the manufacturing and routine share of workers in 2007 was fairly similar across states, there was some potential for variation in exposure across states to declines in these industries.³⁷ To account for this possibility, we include a vector of industry controls in our OLS regression to isolate variation in employment, prices, and wages across states that are orthogonal to the state's industry mix.

While we try to defend that OLS estimation of the above equation yields unbiased estimates of λ and ϕ using cross state variation during the Great Recession, it is impossible to completely rule out that labor supply shocks are causing some of the variation in state business cycles during this period. To further explore the robustness of our results, we also estimate an IV specification of the above equation. Following the work of many recent papers including Mian and Sufi (2014), we use contemporaneous and lagged variation in local house prices as our instruments for local employment and price growth and lagged wage growth. The argument is that local house price variation during the 2007-2011 period (in our base specification) or during the 2007-2009 period (in our restricted specification) is orthogonal to movements in local labor supply shocks. This seems like a plausible assumption for the 2007-2009 period as state policy changes

states did not have the maximum weeks of unemployment benefits even when their unemployment rate was higher.

³⁷For example, the routine share (as measured by employment in manufacturing and administrative occupations) averaged 18.8 percent across states in 2007 (population weighted). The standard deviation in the routine share across states (population weighted) was only 1.8 percent.

did not occur prior to 2009. In the Online Appendix that accompanies the paper, we discuss the IV procedure in detail. We also show that contemporaneous housing price growth strongly predicts contemporaneous employment growth and lagged measures of housing growth predicts lagged wage growth.

Finally, it is worth discussing no migration assumption that we have imposed throughout. If individuals were more likely to migrate out of poor performing states and into better performing states, our estimated labor supply elasticities from the state regression may be larger than the aggregate labor supply elasticity. While theoretically interstate migration could be problematic for our results, empirically it is just not the case. Using data from the 2010 American Community Survey, we can compute both the in-migrants and the out-migrants to and from each state. Given this data, we can compute a net-migration rate for each state. As documented by others, we find that the net migration rate was very low during the Great Recession (Yagan 2014). This can be seen from Appendix Figure A6. Both the low level of interstate migration and the fact that it is uncorrelated with employment growth during this period makes us confident that our estimated parameters of our local wage setting curve can be applied to the aggregate.

6.2 Estimates of λ and ϕ

Column 1 of Table 3 shows the estimates of our base OLS specification where we use all data from 2007-2011 and do not include any additional controls. Our base estimates are $b_1 = 0.69$ (standard error = 0.13) and $b_2 = 0.31$ (standard error = 0.08). As noted above, b_1 is λ and b_2 is λ/ϕ . Given our base estimates, the cross sectional variation in prices and wages implies a labor supply elasticity of 2.2. Standard macro models imply a labor supply elasticity of 2 to 4 based on time series variation. The estimates from the cross-section of states are in-line with these macro time series estimates. 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 6 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.

Columns 2-4 show a variety of robustness checks for our base estimates. In column 2 we include industry controls. In particular, we include the share of workers in 2007 working in manufacturing occupations or in routine occupations. In column 3, we use the actual change in retail prices as opposed to scaling the local retail price difference for the fact that retail grocery sector is more tradable than the composite local consumption good. Both including controls for local industry mix and changing the scaling on local retail price variation does not effect our estimates of λ and ϕ in any meaningful way. In column 4, we show that our estimates of λ remain relatively high even if we constrain $\phi = 1$. In columns 5 and 6, we re-estimate our base specification with and without industry controls using only data from 2007-2009 prior to the changes in national policy

extending unemployment benefit duration and modifying mortgages for underwater homeowners. Again, our estimate λ and ϕ remain 0.73 and 1.9, respectively.³⁸ Finally, in columns (7) and (8) we show our IV estimates for the 2007-2011 period and the 2007-2009 period where we instrument local employment growth and local price growth with contemporaneous and one lag of local house price growth. Our estimates of λ and λ/ϕ are 0.77 (standard error = 0.13) and 0.76 (standard error = 0.17) implying an estimated Frisch elasticity of 1.0.

Regardless of our specification we estimate labor supply elasticities of between 1.0 and 2.5. More importantly, all of our estimates imply a fair degree of wage flexibility. Without constraining the labor supply elasticity, all of our estimates of λ range from 0.69 to 0.79. This is consistent with the patterns shown in Figure 5 where local wages moved quite a bit with local economic conditions during the Great Recession. Even when we constrained the labor supply elasticity to be 1.0 in our base specification, wages were found to be fairly flexible. The estimation that wages are fairly flexible is a key insight driving that both is important for the interpretation of our results and has broader implications for the literature. Linking our estimates back to the prior section, it is hard to get aggregate demand shocks to be the primary shock driving economic conditions during the Great Recession because wages are fairly flexible. Put another way, if wages were sticky enough in the aggregate to have demand shocks be the primary driver of aggregate employment declines during the recent recession, we would not observe wages moving as much as they did in the cross section during the same time period.

7 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

³⁸We also estimated our base specification excluding any state that had an unemployment benefit extension less than 85 weeks. Additionally, we estimated our base specification excluding CA, NV, AZ, and FL. In both cases, our estimates were nearly identical to our base specification in column 1 of Table 3.

this period.

7.1 Findings in the aggregate

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

The aggregate variables we construct are comparable to our regional measures. Given that our cross-sectional equations are estimated using annual data, we analogously define our aggregate data at annual frequencies. We use data from the CPI-U to create our measure of aggregate prices. Specifically, we take log-change in the simple average of the monthly CPI's during year t for our measure of p_t .³⁹ For n , we use BLS data on the aggregate employment rate of all individuals of workers 25-54.⁴⁰ We choose this age range so as to abstract from the downward trend in employment rates due to the aging of the population.⁴¹ Finally, we use data from the Current Population Survey (CPS) to construct our aggregate wage measure w_t . 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 create 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.⁴² 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.⁴³

We start by reporting the impulse response of aggregate employment, nominal wages and price growth to each of the shocks when we use our benchmark estimates for λ, ϕ reported in Section 6.2. Figure 7 shows their behavior after an initial 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 discount rate shock

³⁹All results in the paper are nearly identical if we use annual measures of the PCE as our aggregate price measure.

⁴⁰We downloaded the monthly data on prime age employment rates directly from the Federal Reserve Economic Data site maintained by the St. Louis Federal Reserve.

⁴¹Given 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.

⁴²To the extent possible, the controls in this regression were created to match the controls used to adjust our regional wage data from the ACS.

⁴³In the data appendix we also discuss how we deal with the change in income measurement that occurred in the CPS during 1994.

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 8 shows the impulse responses to a 2008 productivity/markup 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 leisure shock as shown in Figure 9. The employment growth response is rather small, while price growth declines 0.5 percent on impact 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.⁴⁴ 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 11 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 productivity/markup and discount rate shocks contributed about the same to the initial decline during the 2007-2010 period. Nonetheless, absent the sequence of productivity/markup shocks, aggregate employment in the US would have recovered much faster. This is in line with the difference in persistence in the response to each shock that we noted earlier on. The leisure shock, however, barely contributed to the observed employment decline.

Figure 12 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 discount rate shocks alone shows that prices would have decreased by more than 1 percent in the absence of a leisure and productivity/markup 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 productivity/markup shocks alone shows that prices would have increased even more; approximately 5 percent. We note that, according to our procedure, it is this countervailing productivity/markup shock that arises as an explanation for the missing deflation puzzle.

Finally, Figure 13 shows the cumulative nominal wage response. Similar to the price response, if the US economy would have been hit by the sequence of discount rate shocks alone, nominal wages would have declined by approximately 1.5 percent between 2007 and 2012. The combination of the leisure and productivity/markup shocks was behind

⁴⁴For the interested reader, the actual realizations of the shocks we estimate can be seen in Figure 10.

the actual 3 percent increase during this period.

7.1.1 Robustness and Discussion

What patterns across regions in the US are behind our particular decomposition? In Tables 4-6 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). The tables show that the qualitative conclusions of the previous section still hold for the range of $\{\phi, \lambda\}$ estimates using alternative specifications that we report in Table 3. These go from 0.46 to 0.8 for λ and 0.8 to 2.2 for ϕ .

Tables 4-6 offer several further results worth discussing. First, we observe that the relative importance of the leisure shock vis a vis the discount rate and productivity/markup shocks combined is governed by the Frisch labor supply elasticity ϕ . We estimate a relatively large elasticity; in the range of that used to calibrate standard macro models.⁴⁵ However, suppose we used 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 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.

The intuition for the decomposition between discount rate and productivity/markup shocks is more subtle. We find that the degree of wage flexibility λ affects the relative importance of one vis a vis the other within the remaining unexplained part by the leisure shock. For example, suppose we increased the degree of wage flexibility to $\lambda = 1$ (no wage rigidities) instead of using our estimated $\lambda = 0.69$. Then the productivity/markup 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 somewhat persistent, opens the door for the discount/rate shock to play a larger role in the aggregate. Theoretically, it is clear that when λ is large discount rate shocks do not matter much for the determination of employment. To see this, consider the extreme case where wages are perfectly flexible and the demand shock is only composed of the monetary shock. Then the equilibrium in the model satisfies monetary neutrality. We formalized this point in Section 4.6 when we derived the model's implied elasticity of aggregate employment to a discount rate shock.

⁴⁵This 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.

7.2 Findings at the regional level

We follow the procedure in Section 5.2. We first estimate the VAR with two lags in state-level non-tradable employment⁴⁶, price and nominal wage growth via OLS equation by equation. We pool all data between 2006 and 2011, and estimate common autoregressive coefficients and reduced form errors covariance matrix for all states. In our benchmark specification, we set $\{\phi, \lambda\}$ equal to our estimates from Section 6.2. We set $\theta = 0.6$ to match the labor share in the manufacturing sector in the US and $\frac{N^y}{N} = 0.85$ to match the share of total employment in the service sector plus self-employed/family workers as reported in the BLS.

Table 7 summarizes the main results from the regional counterfactuals. We characterize the joint distribution of cumulative growth rates between 2007 and 2010 for each variable across states with two statistics: the variance and the correlation with each other. We find that the discount rate shock alone can explain 59 percent of the cross-state variance of non-tradable employment growth; 86 percent of the price growth variance; and 56 percent of the nominal wage growth variance. Moreover, it reproduces the right sign for the cross-state correlations of price growth and non-tradable employment growth; nominal wage growth and non-tradable employment growth; and nominal wage growth and price growth. Although, quantitatively, it generates larger correlations than in the data. Both productivity/markup shocks combined can explain only 33 percent of the price growth variance across states. They do as good a job as the discount rate in explaining 63 percent of the variation in non-tradable employment growth and a better job in explaining 90 percent of nominal wage growth variation. However, they imply negative correlations between price growth and non-tradable employment growth, and between price growth and nominal wage growth. The opposite is observed in the data. We conclude that the discount rate shock alone does a fairly good job in reproducing the regional patterns that we documented in Section 3.

7.3 Aggregate v. regional elasticities

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.

In this paper, we show that making inferences about the aggregate economy using

⁴⁶We define non-tradable employment in a state as the employment rate in the service and retail sectors combined

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

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

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

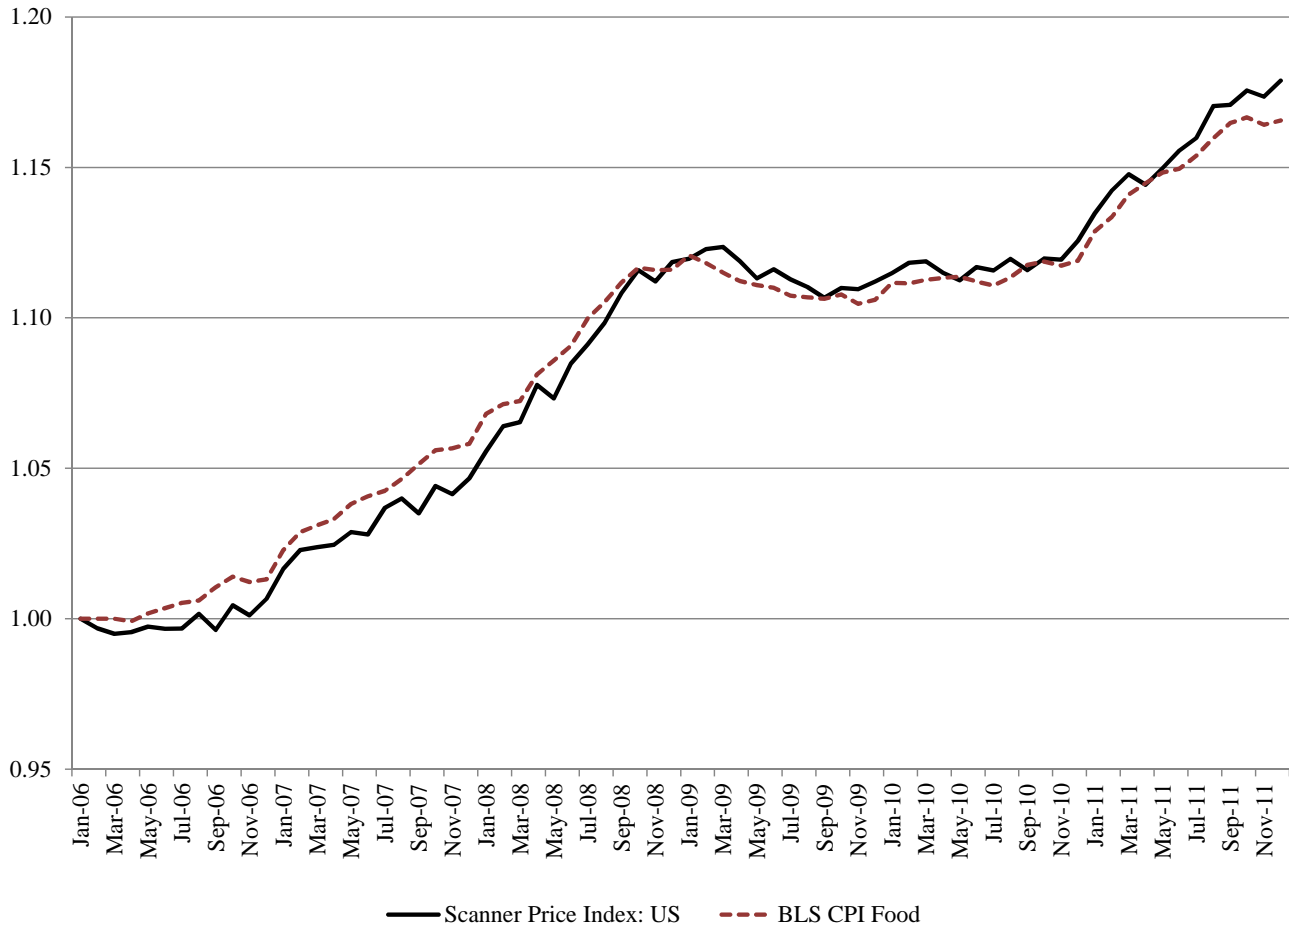
The third take away of our paper is the development of an econometric procedure that allows us to estimate both the aggregate and local shocks as well as aggregate and local elasticities to a given shock using hybrid method that merges restrictions imposed by a theoretical model when estimating a VAR.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

Our fourth take is most important for the goals of the paper. Using our various empirical components, we show that a combination of both "demand" and "supply" shocks are necessary to account for the joint dynamics of aggregate prices, wages and employment during the 2007-2012 period within the U.S.. In contrast with the aggregate results, we find that "demand" shocks explain most of the observed employment, price and wage dynamics across states. These results suggest that only using cross-region variation to

explain aggregate fluctuations is insufficient when some shocks do not have a substantive regional component. The reason that aggregate prices and wages are not falling is not because wages and prices are completely sticky. The reason aggregate prices and wages are not falling is that the series of shocks experienced by the aggregate economy are such that some shocks are putting downward pressure on prices and wages while other shocks are putting upward pressure on prices and wages. In the cross-section, however, the demand shocks are causing prices, wages and employment to covary positively. Lastly, we quantify that the local employment elasticity to a local demand shock is substantially larger than the aggregate employment elasticity to a similarly sized aggregate demand shock. These results suggest that even when the aggregate and regional shocks are the same, it is hard to draw inferences about the aggregate economy using regional variation. Collectively these results suggest that researchers should be cautious when extrapolating variation across regions to make statements about aggregate dynamics.

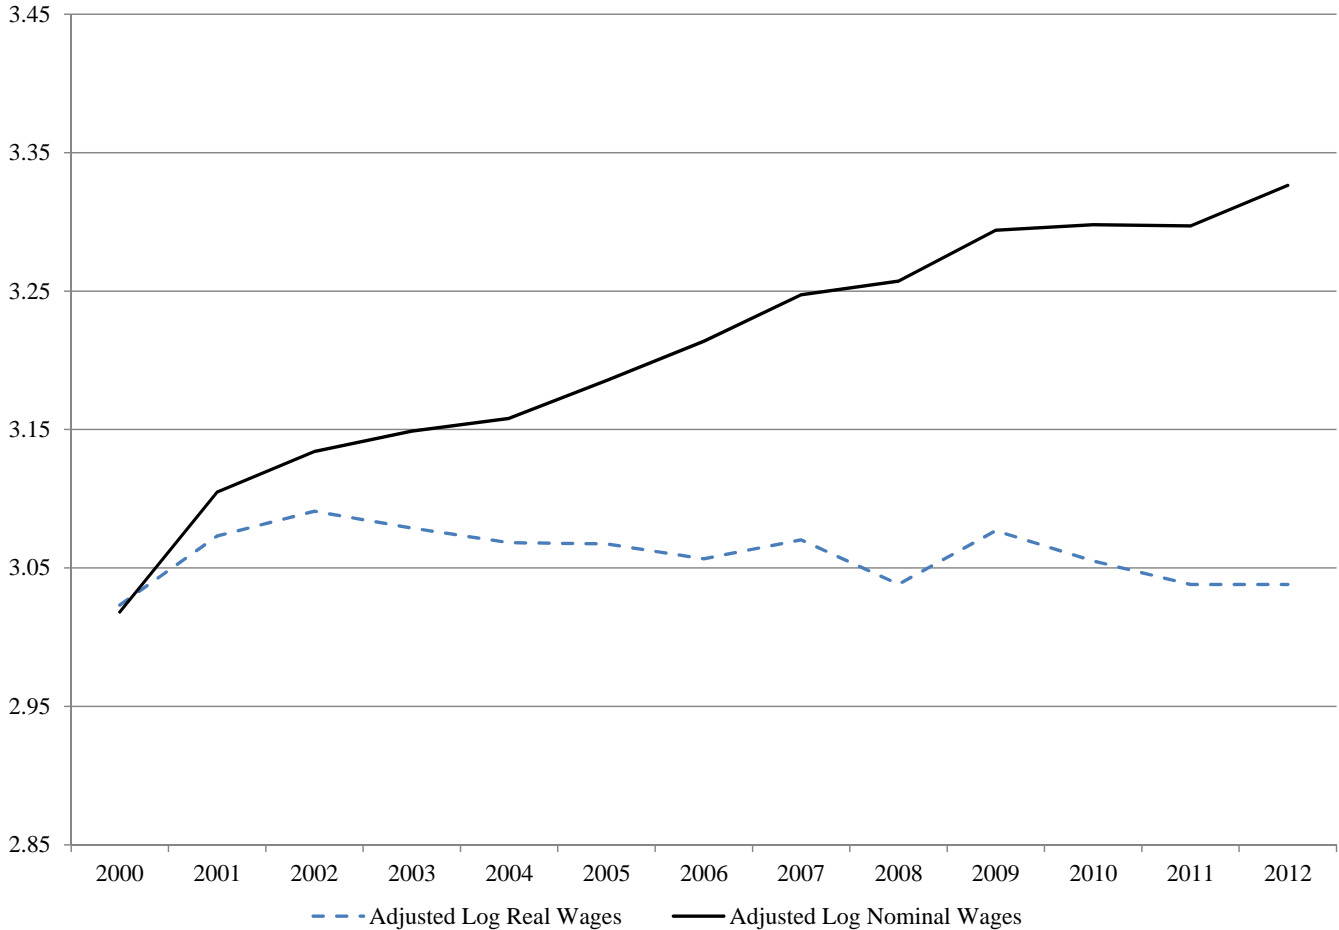
A Figures and Tables

Figure 1: Nielsen Retail Price Index vs. CPI Food Price Index, 2006M1 to 2011M12



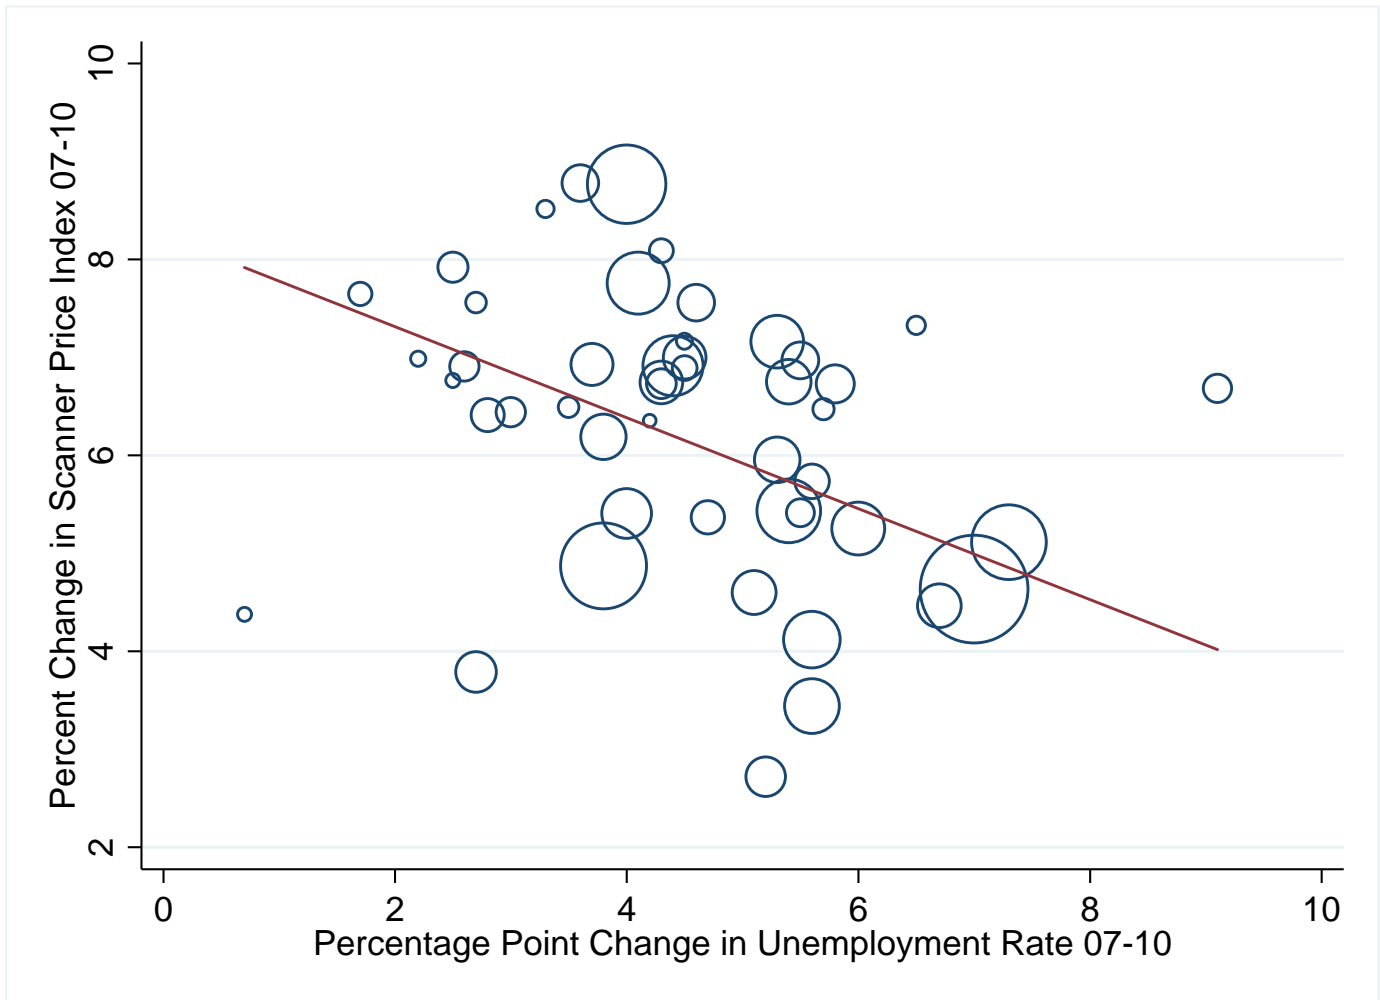
Note: In this figure, we compare our retail price index for the U.S. as whole to the CPI food price index. Given that the goods in our price index come predominantly from grocery, pharmacy, and mass merchandising stores, we thought the food CPI was an appropriate benchmark. For the Nielsen retail price 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 2: The Evolution of Aggregate Real and Nominal Log Wages, Annual Data (2000-2012)



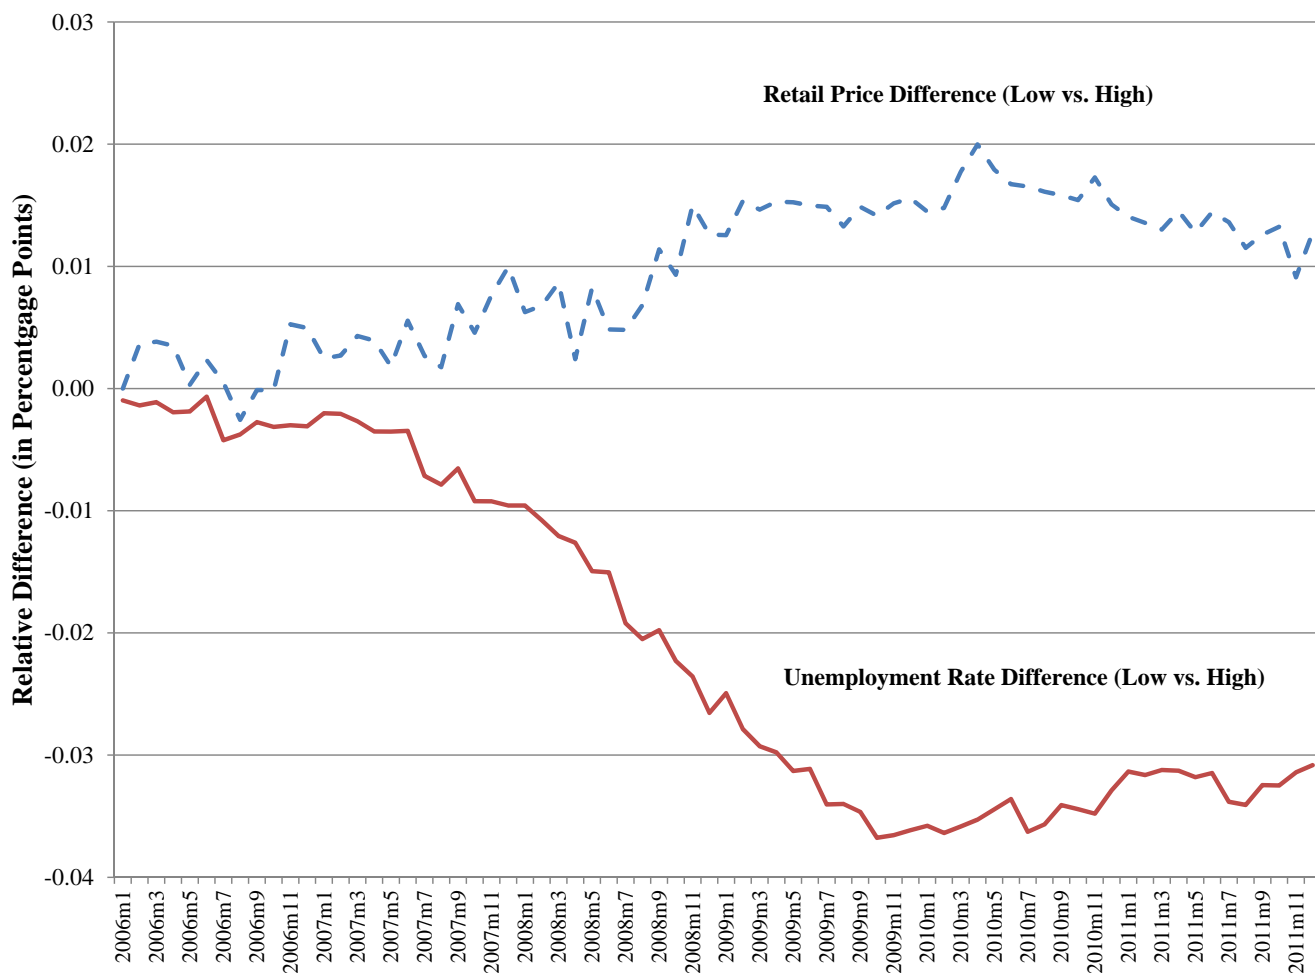
Note: Figure shows the evolution of aggregate real and nominal log wages within the U.S. between 2000 and 2012. Nominal wages were computed using data from the American Community Survey. The sample is restricted to only those individuals who are currently employed, who report usually working 30 hours 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: Change in State Unemployment Rate vs. Cumulative State Retail Price Inflation (2007-2010)



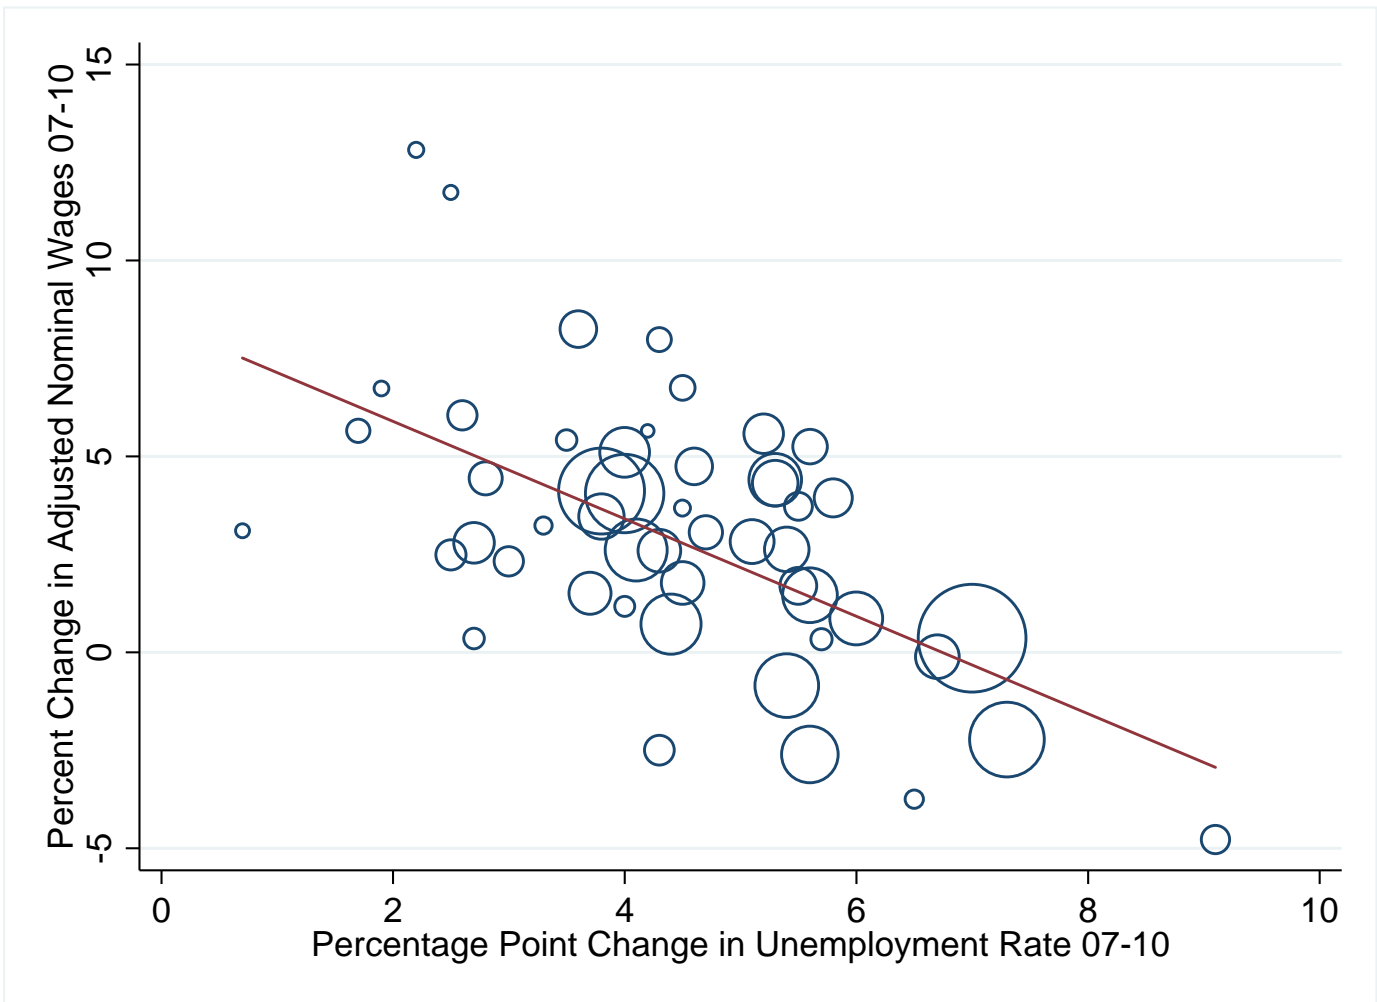
Note: Figure shows a simple scatter plot of the percentage point change in the BLS unemployment rate in the state between 2007 and 2010 against the cumulative percent change in our retail price index based on the Nielsen scanner data during the same period. The retail price index is computed where each 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 4: Differential Retail Prices and Unemployment between Low and High Unemployment Change States, 2006M1 to 2011M12



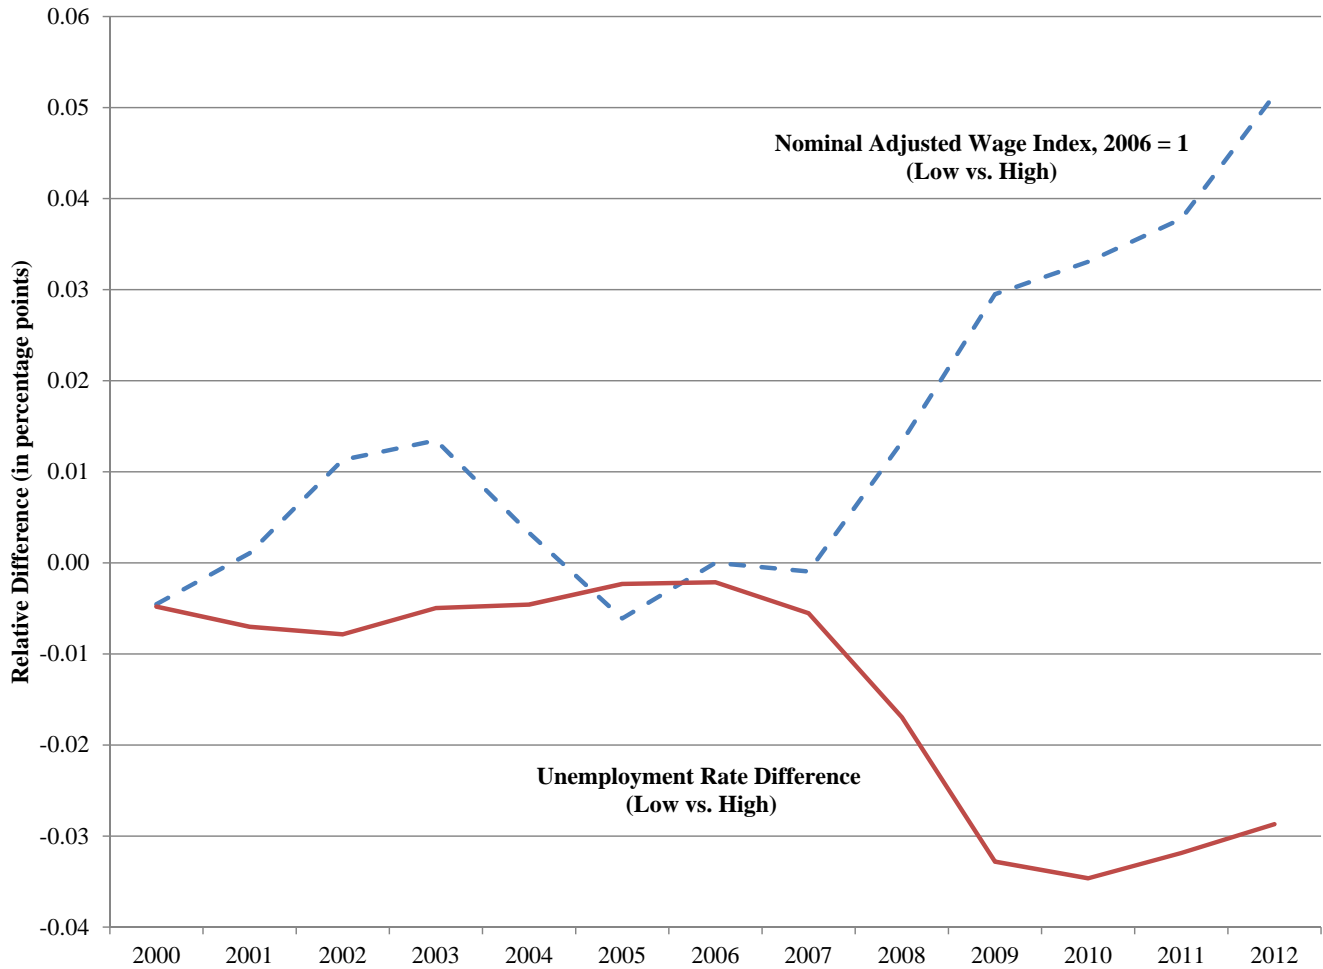
Note: Figure shows the trends in the relative monthly unemployment rate between low and high unemployment change states against the trends in the relative monthly retail price index between low and high unemployment change states. High unemployment change states are the top one-third of all states with 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 5: Change in State Unemployment Rate vs. State Nominal Wage Growth (2007-2010)



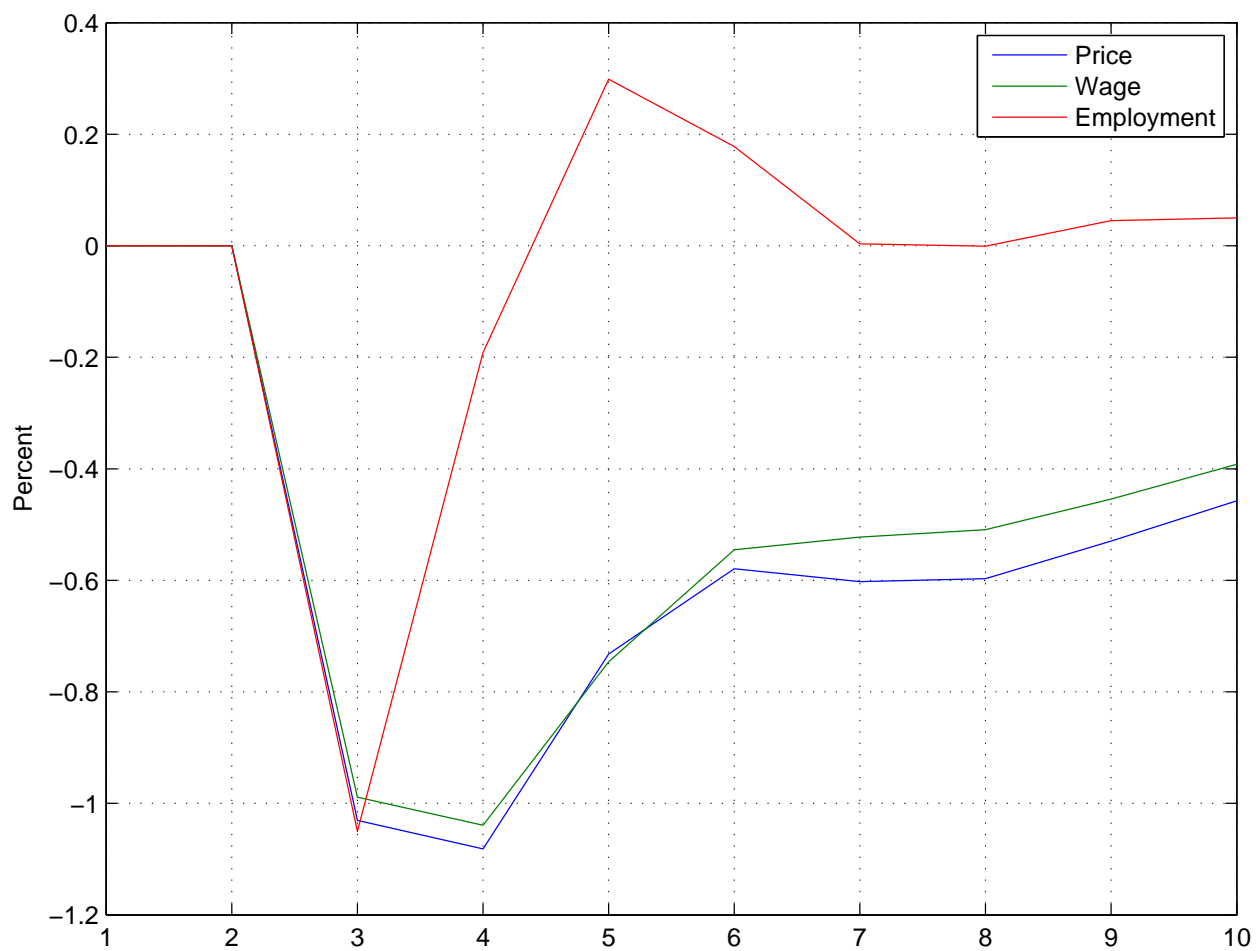
Note: Figure shows a simple scatter plot of the percentage point change in the BLS unemployment rate in the state between 2007 and 2010 against nominal wage growth during the same period. The retail price index is computed where each good is based on a UPC within the state (as opposed to a UPC-store). 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 6: Differential Adjusted Nominal Wage Index and Unemployment between Low and High Unemployment Change States, 2000-2012



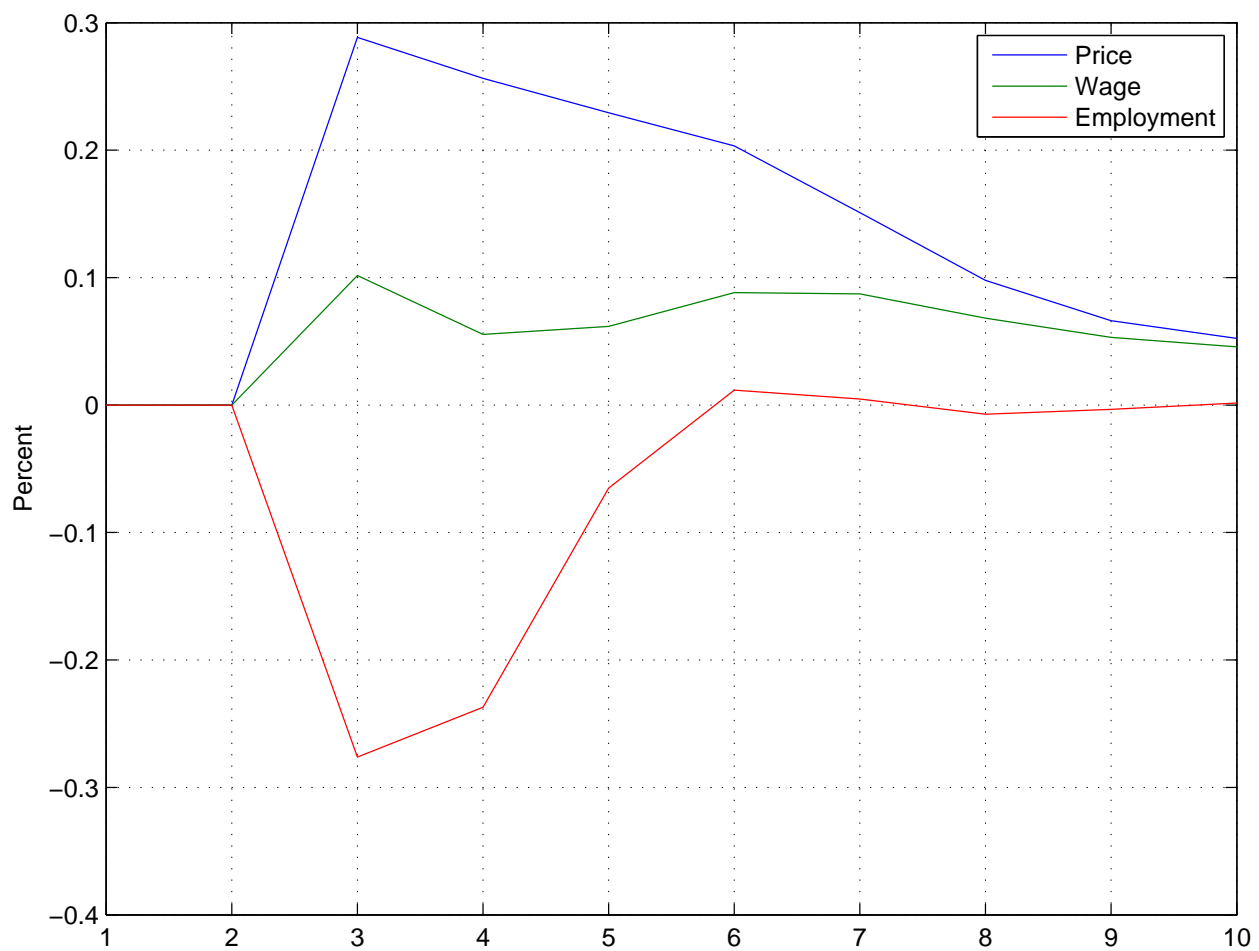
Note: Figure shows the trends in the relative annual unemployment rate between low and high unemployment change states against the trends in the adjusted nominal wage index between low and high unemployment change states. High unemployment change states are the top one-third of all states with respect 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 year, we average the unemployment rate and price index across states weighting each state by their population. The adjusted nominal wage index is adjusted for demographic composition as discussed in the text. We normalize the index to 1 for all states in 2006.

Figure 7: Impulse Response to 2008 Discount rate Shock



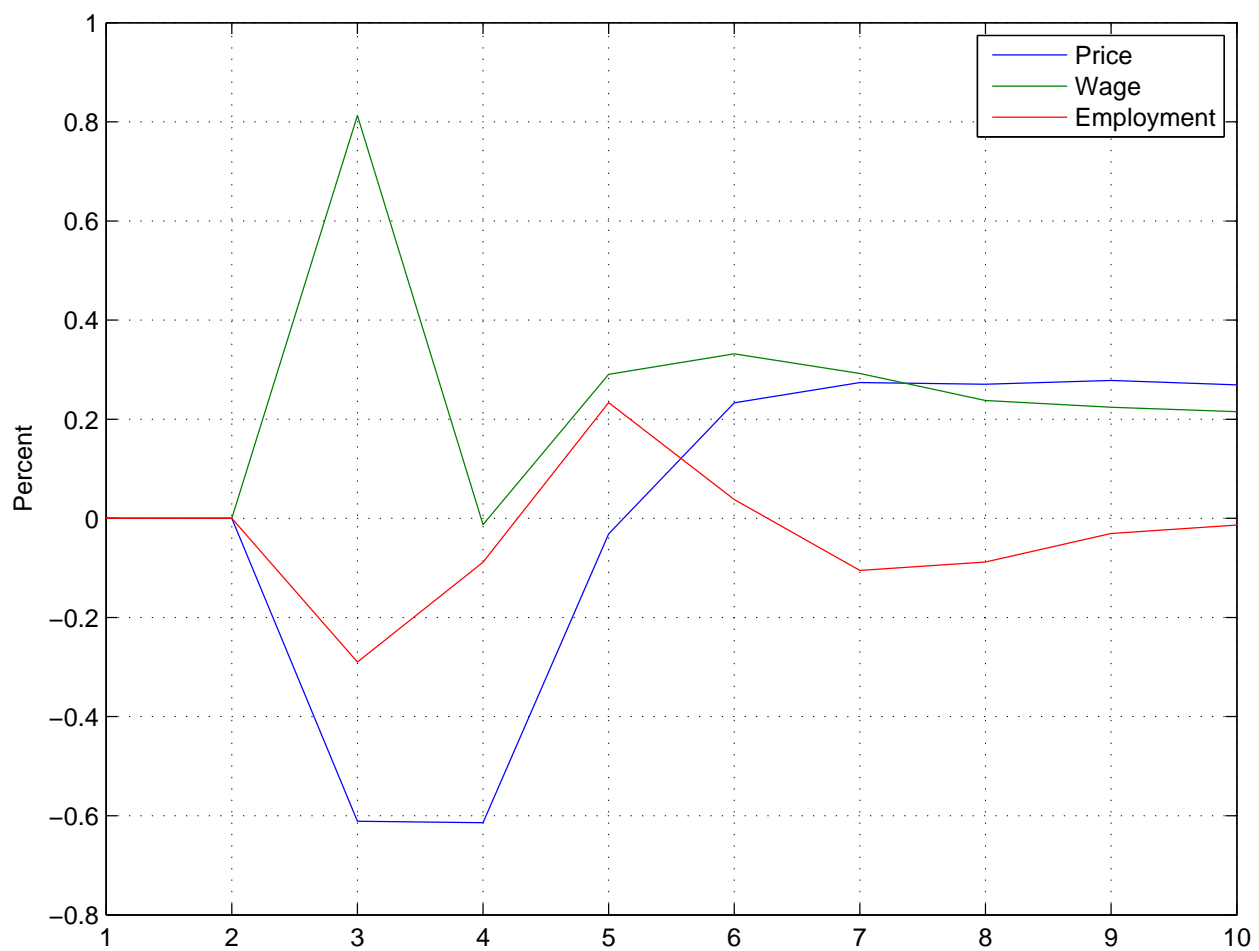
Note: Figure shows the impulse response to a Discount rate shock of the same sign and magnitude as the one we estimate for 2008. The horizontal axis are years after the shock.

Figure 8: Impulse Response to 2008 Productivity / Markup Shock



Note: Figure shows the impulse response to a Productivity/Markup 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 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 10: Shock Time Series

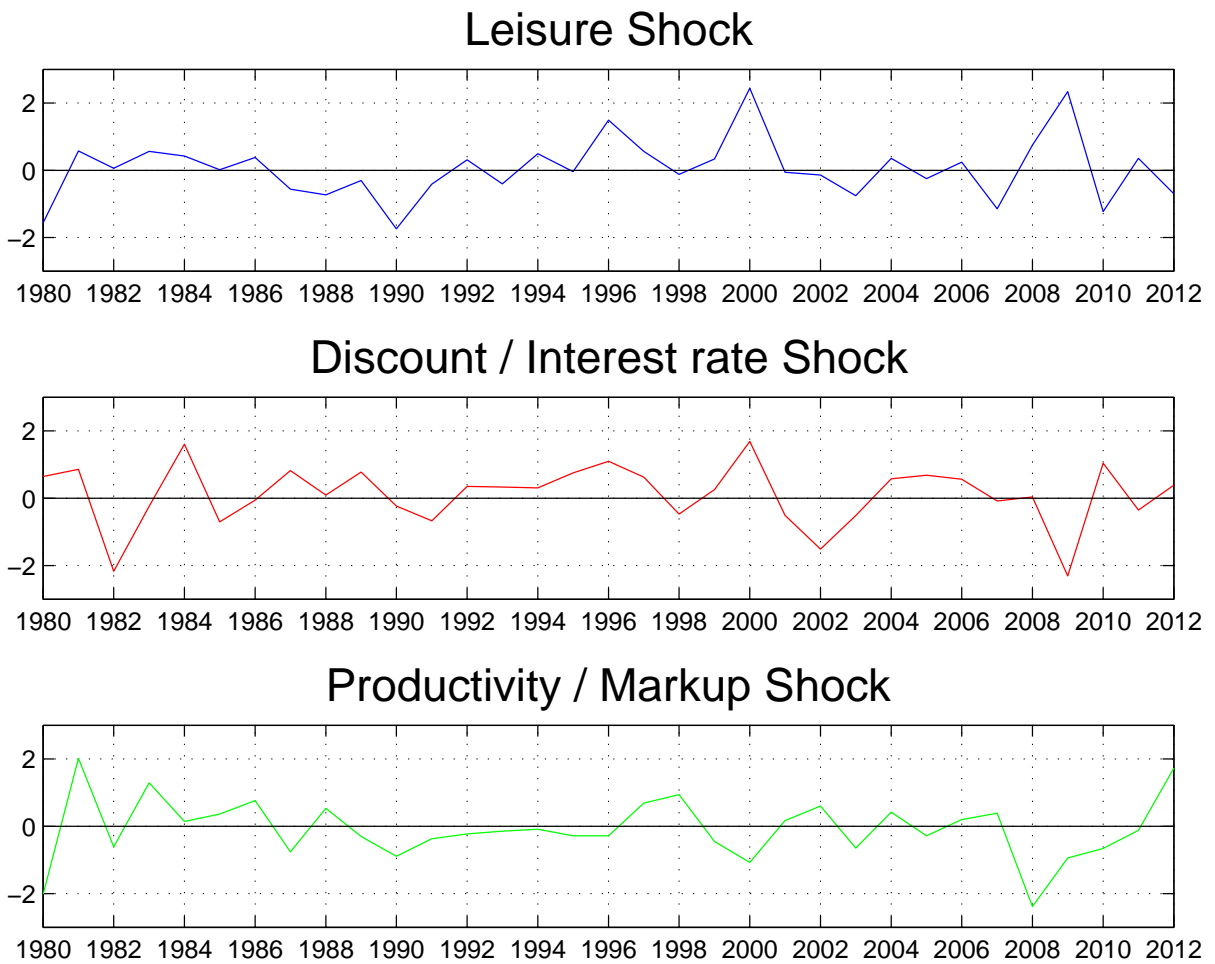
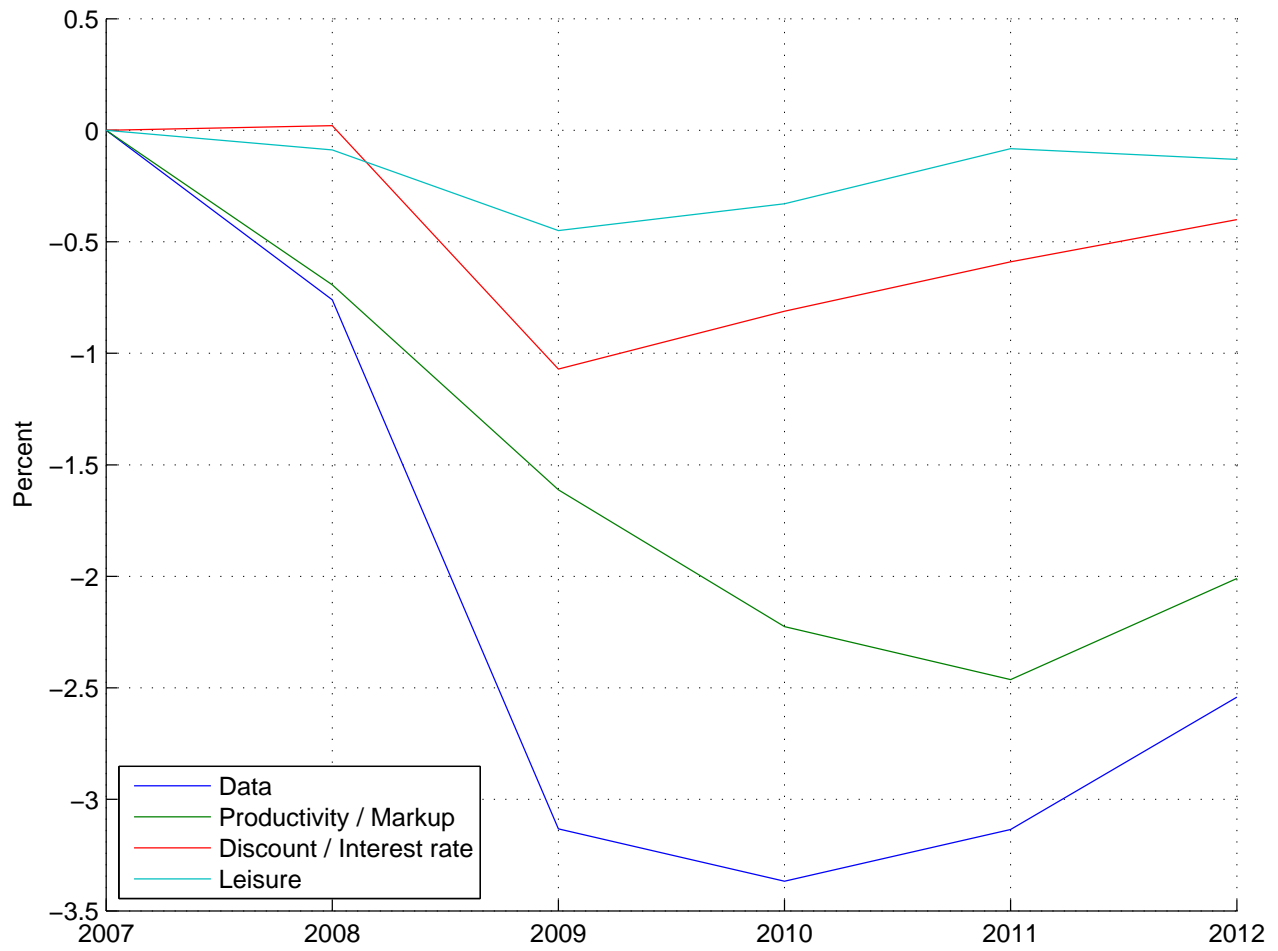
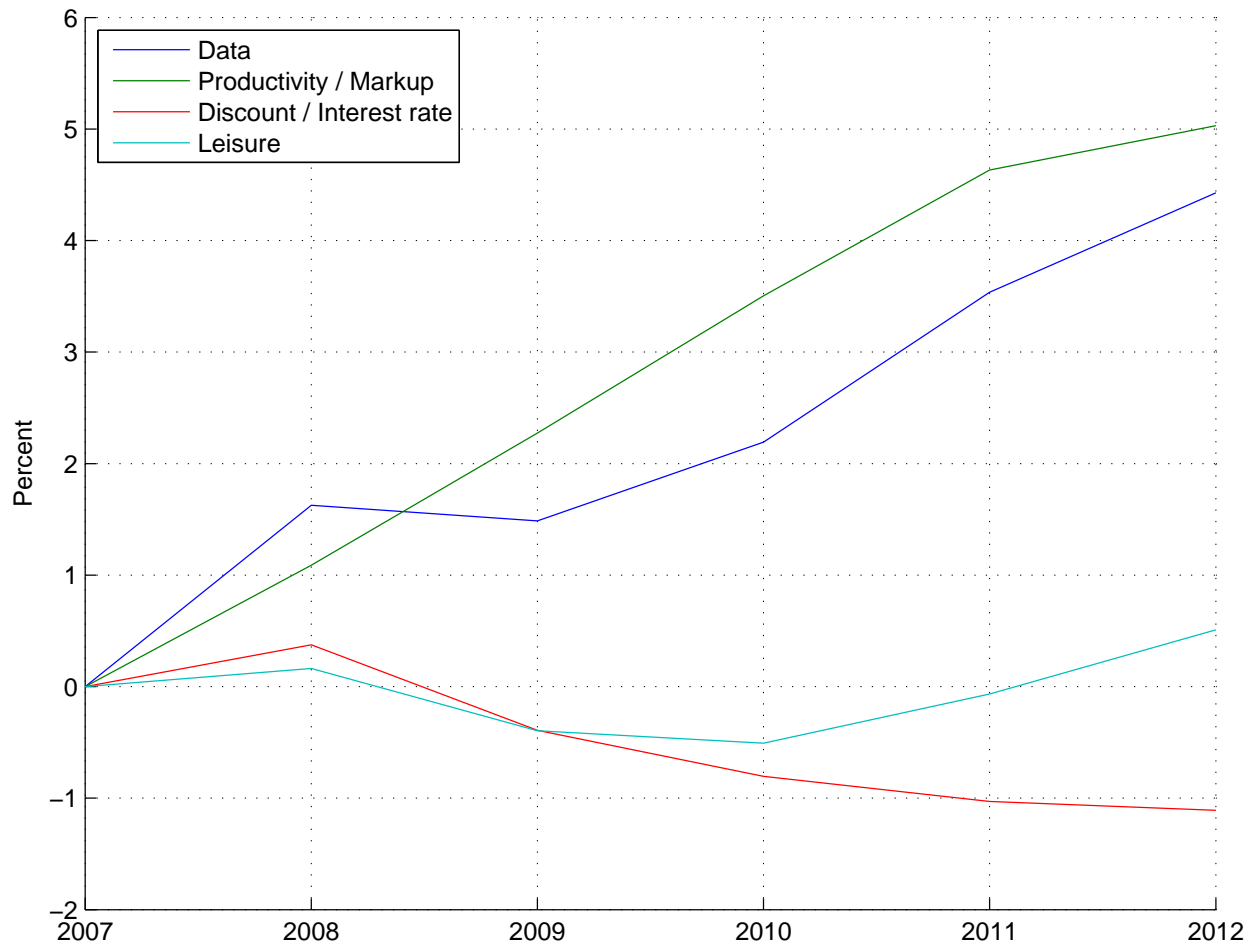


Figure 11: Counterfactual Employment Response



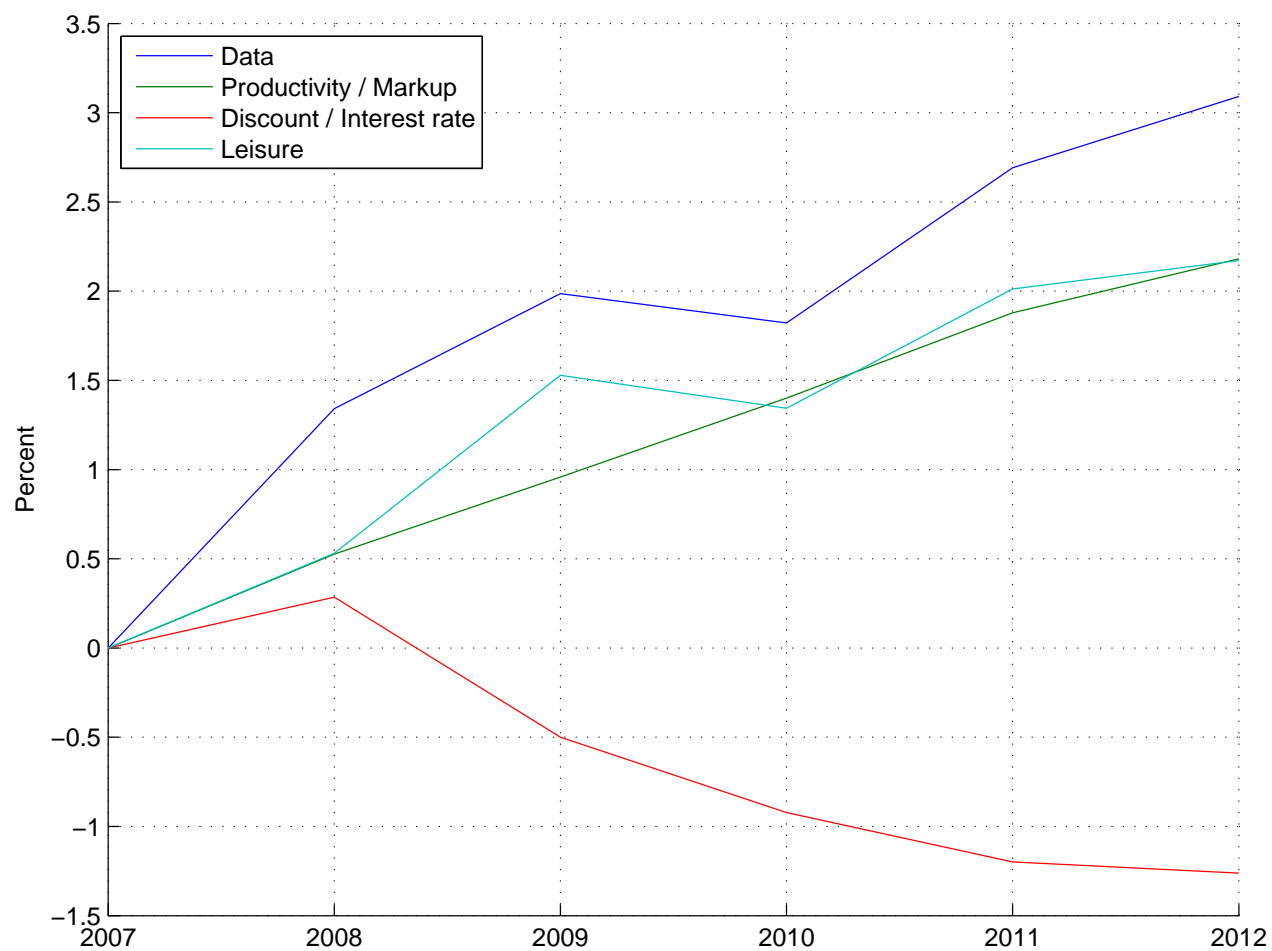
Note: Figure shows the cumulative response of Employment when we feed the VAR with the sequence of shocks between 2008 and 2012; one at a time.

Figure 12: 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 13: Counterfactual Wage Response



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

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

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

	Dependent Variable		
	Nominal Wage Growth	Real Wage Growth (No Scaling of Price Index)	Real Wage Growth (Scaled Price Index)
Change in Unemployment Rate (Percentage Point)	-1.244 (0.205)	-0.871 (0.263)	-0.495 (0.354)
Per-Capita GDP Growth (Percent)	0.487 (0.060)	0.401 (0.079)	0.316 (0.113)
Per-Capita Hours Growth (Percent)	0.653 (0.120)	0.508 (0.146)	0.357 (0.200)
House Price Growth (Percent)	0.113 (0.019)	0.081 (0.024)	0.050 (0.032)
Employment Rate Growth (Percent)	0.618 (0.109)	0.536 (0.128)	0.458 (0.173)
IV: Change in Unemployment Rate (Percentage Point)	-1.500 (0.253)	-1.082 (0.321)	-0.669 (0.430)
IV: Employment Growth	1.011 (0.195)	0.732 (0.209)	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 3: Estimates of λ and $\frac{\lambda}{\phi}$ using Cross- Region Data

	Specification							
	OLS				IV			
	2007-2011		2007-2011		2007-2009		2007-2009	
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	
λ	0.69 (0.13)	0.69 (0.13)	0.75 (0.13)	0.46 (0.07)	0.73 (0.17)	0.73 (0.17)	0.77 (0.13)	0.79 (0.18)
$\lambda\phi$	0.31 (0.08)	0.32 (0.08)	0.31 (0.07)	0.46 (0.07)	0.39 (0.09)	0.39 (0.10)	0.76 (0.17)	0.99 (0.25)
Implied ϕ	2.2	2.2	2.4	1.0	1.9	1.9	1.0	0.8
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry Controls	No	Yes	Yes	Yes	No	Yes	Yes	Yes
No Scaled Prices	No	No	Yes	No	No	No	No	No
Set $\phi = 1$	No	No	No	Yes	No	No	No	No

Note: Table shows the estimates of λ and $\frac{\lambda}{\phi}$ from our base wage setting specification using the regional data. Each observation in the regression is state-year pair. Each column shows the results from different regressions. The regressions differ in the years covered and additional control variables added. The first four columns show the OLS results using all local data between 2007 and 2011. Columns 5 and 6 show OLS results using only data from 2007 through 2009. The final two columns show IV results for the different time periods. In the IV specifications, we instrument contemporaneous employment and price growth with contemporaneous and lagged house price growth. We adjust for measurement error in wage growth, lagged wage growth, and price growth using the split sample methodology discussed in the Online Data Appendix. All regressions included year fixed effects. All standard errors are clustered at the state level.

Table 4: 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 5: 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 6: 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

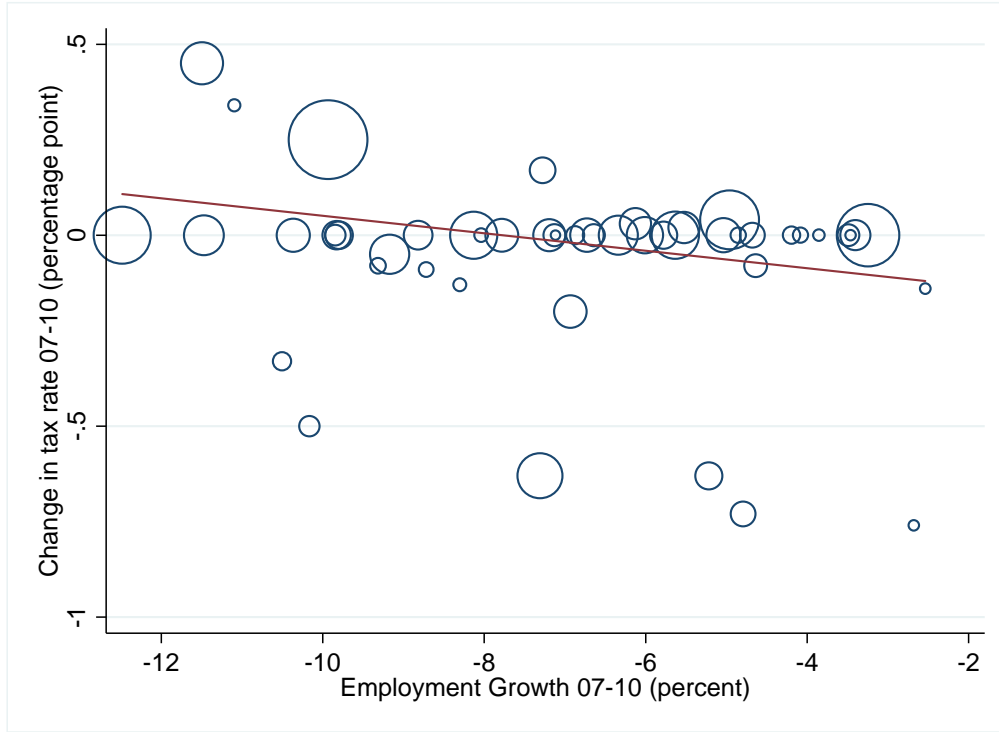
Note: Table shows the percent contribution of the demand and supply shocks to the aggregate wage level change implied by our procedure for different combinations of the parameters. For a given pair $\{\phi, \lambda\}$, the ' γ ' entry corresponds to the demand shock. The ' z ' entry to the supply shock. The percent contribution of the leisure shock can be calculated by subtracting the sum of both entries from 100.

Table 7: Regional counterfactuals

	Data	Discount rate	Productivity/Markup
$\sigma_n^2/(\sigma_n^{data})^2$	1	0.59	0.63
$\sigma_p^2/(\sigma_p^{data})^2$	1	0.86	0.33
$\sigma_w^2/(\sigma_w^{data})^2$	1	0.56	0.9
β_{p,n^y}	0.46	1.63	-0.52
β_{w,n^y}	0.57	1.09	0.46
$\beta_{p,w}$	0.32	0.67	-1.78

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

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



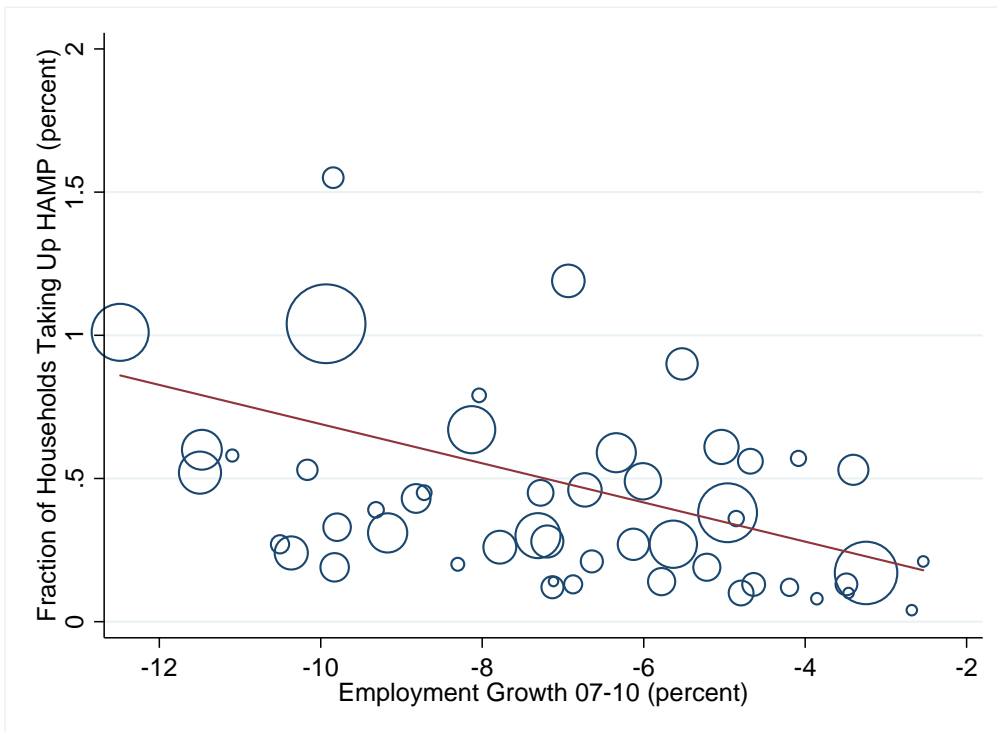
Note: Figure shows the change in the average marginal tax rate in a state between 2007 and 2010 against employment growth in the state during the same period. Employment growth comes from the BLS and is defined in the text. To compute the average marginal tax rate in the state we use data from the American 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



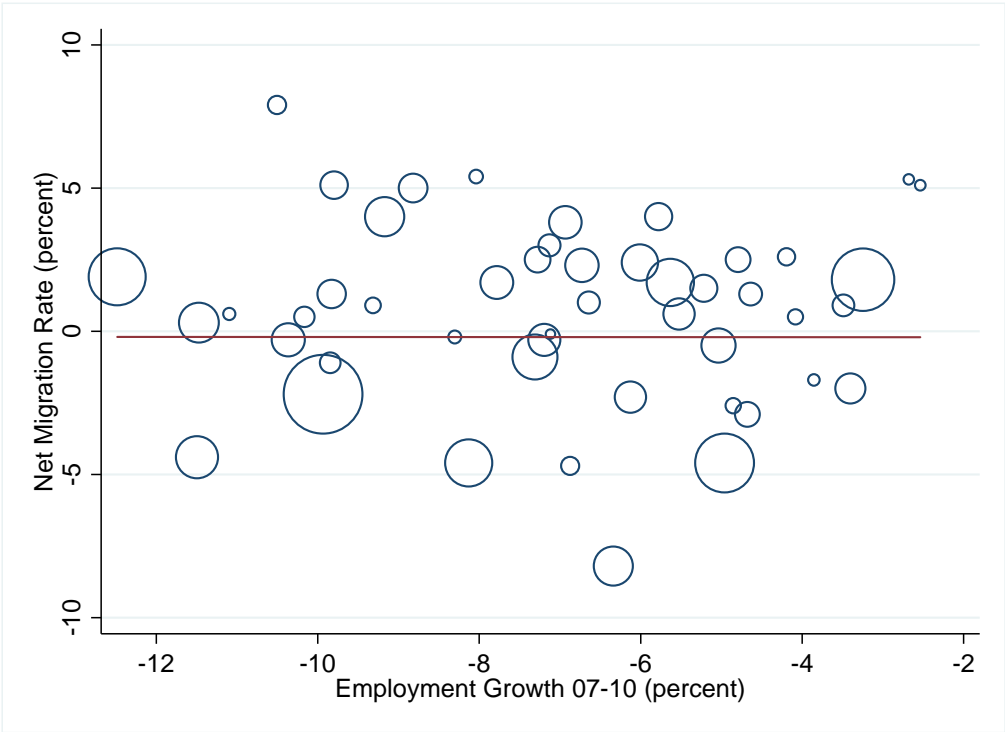
Note: Figure shows the maximum number of unemployment benefits allowed in state in 2010 against employment growth in the state during 2007-2010. Employment growth comes from the BLS and is defined in the text.

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.