

The Aggregate Implications of Regional Business Cycles*

Martin Beraja Erik Hurst Juan Ospina
University of Chicago

August 3, 2015

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) local and aggregate elasticities to the same shock can differ and (ii) the types of shocks driving the local and aggregate business cycles can differ. We develop a semi-structural methodology that combines regional and aggregate data to jointly identify the shocks determining employment, prices and wages at both the aggregate and local. Using household and scanner data, we document that consumer prices and nominal wages are quite flexible at local levels. The strong cross-region relationship between wage growth and employment growth stands in sharp contrast to the aggregate time series patterns during the Great Recession. Applying our procedure, we find that a combination of both “demand” and “supply” shocks are necessary to account for the joint dynamics of aggregate prices, wages and employment during the 2007-2012 period in the US while only “demand” shocks are necessary to explain most of the observed cross state variation. We conclude that the wage stickiness necessary to get demand shocks to be the primary cause of aggregate employment declines during the Great Recession is inconsistent with the flexibility of wages estimated from cross-region variation.

*First draft: May 2014. A previous version of this paper circulated as "The Regional Evolution of Prices and Wages During the Great Recession". We thank Mark Aguiar, Manuel Amador, Mark Bils, Adam Guren, Simon Gilchrist, Loukas Karabarbounis, Pat Kehoe, Virgiliu Midrigan, Elena Pastorino, Harald Uhlig and Joe Vavra for their very helpful comments and suggestions. We also thank Bob Hall for encouraging us to write this paper. Finally, we thank seminar participants at Berkeley, Boston University, Brown, Chicago, Chicago Federal Reserve, Columbia, Harvard, Michigan, Minneapolis Federal Reserve, MIT, Princeton, Rochester, UCLA, NBER's Summer Institute Prices Program, Yale's Cowles Conference on Macroeconomics, the Board of Governors of the Federal Reserve, and the Bank of England. Any remaining errors are our own. E-mail: martinberaja@gmail.com, erik.hurst@chicagobooth.edu, and juan.jose.ospina@gmail.com.

1 Introduction

Regional business cycles during the Great Recession in the US were strikingly different than their aggregate counterpart. This observation is the cornerstone on which this paper is built. We argue that the aggregation of these regions cannot be arbitrary. The particular regional business cycle 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.

There is a large and growing literature using regional variation to learn about the determinants of aggregate economic variables.¹ However, making inferences about the aggregate economy using only regional variation is complicated by two issues. First, the local elasticity to a given shock may differ from the aggregate elasticity to the same shock because of both factor mobility and 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. Since the aggregate effects of a given shock 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. Using regional data, we document that nominal wages and prices are fairly flexible. The extent to which nominal wages are flexible is a key restriction that places structure on the aggregate economy. Given the estimated flexibility of nominal wages, 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.. In contrast, we find that “demand” shocks explain most of the observed employment, price and wage dynamics across U.S. regions. The “supply” shocks we estimate are mostly national and as a result get differenced out when using cross region variation. Our 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 two to three times larger than the aggregate employment elasticity to a similarly sized aggregate demand shock because of factor mobility and general equilibrium effects. These results suggest that even when the aggregate and regional shocks are the same, it is hard to directly draw inferences about the aggregate economy using elasticities estimated from cross-region 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 Nielsen’s *Retail Scanner Database* to compute price indices for each U.S. state. As we

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

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. 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. Given that consumer packaged goods are relatively more tradable than the typical good in a household's consumption basket, the regional variation in prices we estimate from consumer packaged goods is likely to be a lower bound on the regional variation in prices for a composite consumption good. Using data on local price indices provided by the BLS for 27 metro areas, we show that the cross region relationship between employment growth and food price growth is nearly identical to what we get using our scanner price index during the Great Recession. Furthermore, as predicted, the cross-regional variation between employment growth and price growth for a broad basket of goods within the BLS data is larger than the employment growth-food price growth relationship.

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

The cross region patterns that we document stand in sharp contrast with the well documented aggregate time series patterns for prices and wages during the same time period. It is in this sense that regional business cycles differed from their aggregate counterpart. 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. The robust growth in nominal wages and consumer prices during the recession is viewed as a puzzle for those that believe that the lack of aggregate demand was the primary cause of the Great Recession.² Recently, a literature has emerged trying to explain the missing disinflation and the missing wage declines during this time period.³ The key point we wish to make with these new facts is that while aggregate wages appear sticky during the Great Recession using time series vari-

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

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

ation, local wages were strongly correlated with measures of local employment growth using cross-state variation. Any model that is calibrated to match the relationship between wage growth and employment growth at the aggregate level needs to confront the fact that there is a strong relationship between local employment growth and local wage growth at business cycle frequencies.

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 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 the causes of aggregate business cycles.

We start by describing a simple model of a monetary union with many islands linked by trade in intermediate goods which are used in the production of a non-tradable final consumption good and by trade in a risk-free asset. The nominal interest rate on this asset follows a rule that endogenously responds to aggregate variables and is set at the union level. Labor is the only other input in production, which is not mobile across islands. We assume that nominal wages are only partially flexible. This is the only nominal rigidity in the model. Finally, the model includes a series of shocks: a shock to the household's discount rate, shocks to non-tradable and tradable productivity/mark-up, a shock to the household's taste for leisure, and a monetary policy shock.⁴ All shocks have both local and aggregate components where by definition the weighted average of the local shocks sum to zero. We show that, under relatively few assumptions, the log-linearized economy aggregates allowing us to study the aggregate and local behavior separately.

The model is used for two purposes. First, the model specifies a wage setting equation for the local and the aggregate economies. Given our specification of the nominal rigidity, the local wage setting equation is a function of a wage stickiness parameter and parameters of the household's marginal rate of substitution. For a given set of preferences, there is a direct mapping between the parameters of the local and aggregate wage setting equations. These equations will be important to our procedure that identifies the underlying aggregate and local shocks. The model also provides a framework to quantify the differences between aggregate and local elasticities to a given shock. While it is well known that aggregate elasticities to a given shock will differ from local elasticities to the same shock because of general equilibrium forces, there have been few papers that quantitatively document the difference between these elasticities.⁵ Using a param-

⁴As discussed below, nothing in our estimation procedure allows us to tease out the difference between a productivity shock, a mark-up shock, or any other non labor marginal cost shock faced by firms. As a result, we refer to these shocks in the model as "productivity/mark-up" shocks.

⁵One notable exception is Nakamura and Steinsson (2014) that compare government spending multipliers at the local and aggregate level.

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

Turning to the estimation of the local and aggregate shocks, we consider a broader class of models than the simple model outlined above. The broader class of models, however, does nest our simple model. In particular, we show that the aggregate and local equilibria can be represented as vector autoregression (VAR) in prices, nominal wages, and employment with three shocks. We refer to the three shocks as a "demand shock" (which is a combination of the discount rate and monetary policy shock), a productivity/mark-up shock (which is a combination of the productivity/mark-up shocks in the tradable and non tradable sectors) and the taste for leisure shock. To back out the aggregate (local) shocks, we estimate the aggregate (local) VAR. Our estimation of the VAR is semi-structural in that we impose the aggregate (local) wage setting equation as an additional restriction to help identify the VAR.⁶ The aggregate (local) wage setting equation implies a series of particular linear restrictions linking the reduced form errors to the underlying structural shocks. When we impose these linear restrictions, the VAR becomes identified. In essence, our procedure uses some elements of theory to help identify the underlying economic shocks. To formalize things, our VAR nests models where the wage setting equation is equivalent to the one developed in our simple theory. We can be agnostic with respect to the other components of these models aside from the wage setting equation and still identify the underlying shocks.

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 of the wage setting equation in our base specification.⁷ 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

⁶We 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 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).

⁷In our base specification, we assume GHH preferences such that there is no income effect on labor supply. However, as a robustness exercise, we consider a broader set of preferences that allow for income effects on labor supply. The results of the two specifications are similar in terms of our shock decompositions.

the Great Recession or (2) variation in local discount rate shocks and/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/or that housing price variation during the 2007-2009 period can help us isolate movements in local labor demand.⁸ Across a variety of specifications and identification procedures, we estimate only a modest amount of wage stickiness. These results are not surprising given the reduced form relationships we document in the descriptive part of the paper. The amount of flexibility we estimate using the local variation is much greater than estimates of wage flexibility obtained using only aggregate time series variation during recent years.

With the parameterized aggregate wage setting equation, we use the VAR methodology described above to estimate the shocks driving aggregate employment, prices, and wages during the Great Recession. The results suggest during the early part of the recession (2008-2009) roughly 40 percent of the aggregate employment decline can be traced to the "demand shock" (the discount rate plus the monetary policy shocks). The leisure shock explains roughly 10 percent of the decline in aggregate employment while the productivity/mark-up shock explaining the remaining 50 percent. Over a longer period (2008-2012), the demand shock cannot explain any of the persistence in the employment decline. Instead, it is the productivity/mark-up shock explaining why employment remained low from 2010-2012. While the demand shock may have been important in the early part of the recession, it had little effect on explaining the low levels of employment in the U.S. after 2009.⁹

The regional data in our paper serve two purposes in our estimation. First, the regional data is needed to estimate the amount of wage stickiness which is a parameter of the aggregate wage setting equation. Second, the regional data is needed to estimate the local VAR. We use a similar procedure to estimate the shocks driving the local economy. Our results suggest that the discount rate shock is driving essentially all the cross-region variation in employment 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. This is why price and wage growth are so correlated with employment growth at the local level. At the aggregate level, a combination of shocks are responsible for declining aggregate employment. Some of these shocks put

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

⁹Christiano et al (2014) estimate a New Keynesian model using data from the recent recession. Although their model and identification 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. Likewise, Vavra (2014) and Berger and Vavra (2015) document that prices were very flexible during the Great Recession. They also conclude that something more than a demand shock is needed to explain aggregate employment dynamics given the missing aggregate disinflation.

downward pressure on prices and wages, while other shocks put upward pressure on prices and wages. It is our conclusion that the reason that aggregate price growth and wage growth appeared acyclical during the Great Recession was not due to extreme price and wage stickiness. Instead, we conclude that the seemingly acyclical response of prices and wages at the aggregate level was the result of the combination of shocks hitting the aggregate economy. The procedure we developed quantifies the relative magnitudes of these shocks.

Our paper contributes to many literatures. First, our work contributes to the recent surge in papers that have exploited regional variation to highlight mechanisms of importance to aggregate fluctuations. For example, Mian and Sufi (2011 and 2014), Mian, Rao, and Sufi (2013) and Midrigan and Philippon (2011) have exploited regional variation within the U.S. to explore the extent to which household leverage has contributed to the Great Recession.¹⁰ 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 trying to determine the causes of the Great Recession. In many respects, our model is more stylized than others in this literature in that we include a broad set of shocks without trying to uncover the underlying micro foundation for these shocks. However, the shocks we chose to focus on were designed to proxy for many of the popular theories about the drivers of the Great Recession. For example, our discount rate shock can be thought of as reduced form representation of tightening of household borrowing limits. For example, such shocks have been proposed by Eggertsson and Krugman (2012), Guerrieri and Lorenzoni (2011) and Mian and Sufi (2014) as an explanation of the 2008 recession. Likewise, our productivity/mark-up shock can be interpreted as anything that changes firms' demand for labor. In a reduced form sense, credit supply shocks to firms, such as those proposed by Gilchrist et al (2014), would be similar to our productivity/mark-up shock. Finally, our taste for 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)). So while our specification only allows for broad reduced form shocks, we think these shocks nest many of the prominent stories about the underlying causes of the Great Recession. As we show, the fact that prices and wages move with economic

¹⁰There has been an explosion of papers using regional data to better understand aggregate dynamics during the Great Recession. Some recent papers include: Giroud and Mueller (2015), Hagedorn et al. (2015), Mehrotra and Sergeyev (2015), and Mondragon (2015).

conditions at the local level help to discipline how aggregate prices and wages should have moved in response to different types of shocks.

Finally, there is some recent work using scanner data to explore the relationship between local economic conditions and prices. Contemporaneously, Coibion et al. (2014) use data from Symphony IRI to examine regional variation in prices during the 2000-2011 period. Kaplan and Menzio (2014) use data from Nielsen's Homescan data to examine how the variance of prices paid change with economic conditions. Stroebel and Vavra (2014) use the IRI data to document the relationship between local house price movements and local retail prices. 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, Fitzgerald and Nicolini (2014) 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 will post these indices so other researchers can use them in their research going forward.¹¹

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.¹² 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

¹¹Handbury and Weinstein (2011) use Nielsen's 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, the literature referenced above all document that local prices do move with business cycle frequencies.

¹²The 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/>.

merchandising stores.¹³

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 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. Online Appendix Table R1 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 are similar in spirit to how the BLS constructs the CPI.¹⁴ While we briefly outline our procedure in this sub-section, the full details of the procedure are discussed in the Online Data and Robustness Appendix that accompanies our paper. Our scanner price indices are built in two stages. In the first stage, we aggregate the prices of goods within the roughly 1,000 categories described above. For our base indices, a good is 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 the same UPC 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

¹³It 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.

¹⁴There is a large literature discussing the construction of price indices. See, for example, Diewert (1976). Cage et al (2003) discuss the reasons behind the introduction of the BLS's Chained Consumer Price Index. Melser (2011) discuss problems that arise with the construction of price indices with scanner data. In particular, if the quantity weights are updated too frequently the price index will exhibit "chain drift". This concern motivated us to follow the BLS procedure and keep the quantity weights fixed for a year when computing our indices rather than updating the quantities every month. Such problems are further discussed in Diewert et al. (2011).

shopping experience) to a low cost store when local economic conditions deteriorate.¹⁵ 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 period t , with a base year y , in geography k . For our analysis, geographies will either be U.S. states or the country as a whole. $p_{i,t,k}$ is the price at time t of the specific good i in geography k and $\bar{q}_{i,t-1,k}$ is the average monthly quantity sold of good i in the prior year in location k . By fixing quantities at their prior year's level, we are holding fixed household's consumption patterns as prices change. We update the basket of goods each year, and chain the resulting indices to produce one chained index for each category in each geography. Fixing quantities at a lagged level implies that the price changes we document below with changing local economic conditions is not the result of changing household consumption patterns.

The second stage of our price indices also 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:

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

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

¹⁵In 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.

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 occur. When computing our price indices, we restrict our sample to only include (1) goods that had positive sales in the prior year and (2) goods that had positive sales in every month of the current year. Online Table Appendix A1 shows the percent of sales included within the price index for each sample year.

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 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:

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.

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

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 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 use an estimate of the share of total local consumption at the state level that is imported from outside the state. 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).

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 1 and 3 and the quantitative implication of our estimation procedures are

¹⁸The level of analysis in Nakamura and Steinsson (2014) is U.S. regions. They define 10 regions - the nine Census divisions with "South Atlantic" segmented into two separate 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.

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.

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 per year between 2001-2004 and around 2 million respondents per year 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 males 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.¹⁹ These restrictions result in our sample being comprised of males with a strong attachment to the labor force.

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, 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 .²⁰

¹⁹Total 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.

²⁰These adjustments slightly mutes the relationship between employment growth and wage growth across states. In other words, if we do not adjust for the changing labor force composition, we get

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. The patterns in our data replicate the aggregate nominal wage growth patterns documented by many others in the literature.²¹ 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 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 2010.²² For the results in Figure 3, we use our price index where a good is a given UPC within a state. 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:

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

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.

slightly stronger relationships between employment growth and wage growth across states during the Great Recession.

²¹See, for example, Daly and Hobijn (2015).

²²Our 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.

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 $(-0.46 * 2)$.²³

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 has 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.²⁴ 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

²³It has been argued that food prices are more flexible than the prices of other types of goods. However, we are not identifying the flexibility of food prices in the aggregate. Any fixed amount of price flexibility for food gets differenced out in our cross-region analysis. The variation we are identifying is in local distribution costs.

²⁴The 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.

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).

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.

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 prices for the composite consumption good in the low unemployment change states relative to the high unemployment change states. The difference in prices between low and high unemployment change states is essentially the mirror image of the difference 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.

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

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 composite 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. Additionally, these patterns underlie our estimates of wage stickiness in our formal empirical work below.

The second and third columns of Table 2 show the coefficient on the change in local economic activity between 2007 and 2010 from a regression of real wage growth in a given state during that time period on the change in local economic activity. In column 2, we compute local real wages by deflating local nominal wage growth by the growth in the local scanner price index. In column 3, we compute local real wages by deflating local nominal wage growth with the growth in the prices of a composite local consumption good. As discussed above, we scale the growth in the scanner price index by a factor of 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.

3.3 Robustness

Part of the goal of this paper is to document a set of facts about the extent to which prices and wages vary with local economic conditions during the Great Recession. As we discuss throughout the paper, the extent to which the prices and wages vary with local economic conditions can inform the extent to which wages and prices are sticky helping us to infer the drivers of aggregate business cycle conditions. While we made our own local indices of prices and wages using detailed micro data, there does exist alternate measures of local prices and wages released by government agencies. As we discuss below, these alternate local indices have some limitations. Despite these limitations, the patterns we document about the relationship between changes in local prices and wages and changes in local economic conditions remain even within these alternate data sets.

The U.S. Bureau of Labor Statistics produces 27 metro areas price indices at various degrees of time aggregation.²⁵ For each metro area, the BLS publishes multiple price indices (food, nondurables, etc.). Although there are only 27 MSAs, we can still explore the relationship between unemployment growth between 2007 and 2010 and the cumulative inflation rate between 2007 and 2010 using these data. We then can compare the relationship between local unemployment growth and local scanner price inflation that we document in Section 3 with the relationship between local unemployment growth and local inflation computed using the BLS data. Given the price indices are only provided semi-annually for many MSAs, we compare the change in the unemployment rate and prices from the first half of 2007 to the latter half of 2010. When data is provided monthly or bi-monthly, we simply take the geometric average over the first and last half of the year to make the data semi-annual. As with the results above, we use data from the BLS's local area unemployment statistics to measure the percentage point change in the local unemployment rate.

Using this data, we regress the 3-year inflation rate at the MSA level on the 3 year change in the unemployment rate. We run this regression for various inflation measures (food, services, all goods less housing, etc.). The results of these simple regressions using the BLS data are strikingly similar to our results using the scanner data. For example, the results using the scanner index find that a 1 percentage point increase in the local unemployment rate is associated with a 0.38 decline in the local inflation rate (for the specification where goods are defined as UPC-store pairs). Within the BLS data,

²⁵The 27 MSAs are (in order of reporting frequency): Chicago, Los Angeles, New York, Atlanta, Boston, Cleveland, Dallas-Fort Worth, Detroit, Houston, Miami, Philadelphia, San Francisco, Seattle, Washington, Anchorage, Cincinnati, Denver, Honolulu, Kansas City, Milwaukee, Minneapolis, Phoenix, Pittsburgh, Portland, St. Louis, San Diego, and Tampa. The first three MSAs have price indices that are reported at monthly frequencies. The next 11 MSAs have price indices that are reported at bi-monthly frequencies. The last 13 price indices are reported semi-annually.

we find that a 1 percentage point increase in the local unemployment rate is associated with a 0.34 percentage point decline in the local food inflation rate (standard error = 0.22).

Also, as predicted by our simple model above, the relationship between the inflation rate for all goods and the local unemployment rate change should be higher than the relationship between food inflation and the local unemployment rate change if food is relatively more tradable than the local consumption good. We cannot test this within our scanner data. However, the BLS data allows us to test this prediction directly. Within the BLS data, the relationship between the inflation rate for all goods with the change in the unemployment rate is in fact higher at -0.47 (standard error = 0.15). The fact that the coefficient is larger in magnitude is consistent with our belief that the variation in packaged food data across regions should be a lower bound for the variation in the average consumption good given that the packaged goods in our dataset are relatively more tradable.

The patterns that we document for the 2007-2010 period with the BLS data are very consistent with the results in Fitzgerald and Nicolini (2014) that use the BLS MSA level price indices to show these relationships over a much longer time period. Fitzgerald and Nicolini (2014) find that over the period of 1976-2010, a 1 percentage point increase in the local unemployment rate is associated with a 0.3 percentage point decline in the local annual inflation rate. In summary, a limitation of our price data is that it only covers goods sold in grocery, pharmacy and mass-merchandising stores. The fact that the patterns we uncover are nearly identical for similar goods using the BLS metro area price indices is very reassuring. Moreover, as predicted, using broader measures of goods in the BLS data seems to only strengthen the cross-region variation. Our data is an advance over the BLS metro level data in that can be calculated for every state and, if one desires, a much larger set of metro areas.

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

Despite these differences, the cross state correlation between growth in our composition adjusted wage index from the ACS and the growth in earnings-per-worker from

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

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

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

4.1 Firms and Households

Producers of tradable intermediates x in island i use local labor N_i^x and face nominal wages W_i (equalized across sectors) and prices Q (equalized across islands k). Their

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

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.²⁷

Households preferences are given by

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

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

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

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

A well known issue in the international macroeconomics literature is that under market incompleteness of the type we just described there is no stationary distribution for bond holdings across islands in the log-linearized economy; and all other island variables in the model have unit roots. This is problematic for reasons both theoretical (we will like to study log-deviations from a deterministic steady state) and empirical (regional data for the US does not suggest the presence of such unit roots). We follow Schmitt-Grohe and Uribe (2003) and let ρ_{kt} be the endogenous component of the dis-

²⁷It is worth noting that all model shocks will generate endogenous variation in markups given our assumption of decreasing returns to scale. Additionally, what we call a "productivity shock" is isomorphic to any shifter of unit labor costs and, hence, labor demand schedules. Later we will refer to it as the productivity/markup shock. We do not attempt to distinguish between the different interpretations of this shock in this paper.

count factor that satisfies $\rho_{kt+1} = \rho_{kt} + \Phi(\cdot)$ for some function $\Phi(\cdot)$ of the average per capita variables in an island. As such, agents do not internalize this dependence when making their choices. This modification induces stationarity for an appropriately chosen function $\Phi(\cdot)$. They show that alternative stationary inducing modifications (a specification with internalization, a debt-elastic interest rate or convex portfolio adjustment costs) all deliver similar quantitative results in the context of a small open economy real business cycle model.

4.2 Sticky wages

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

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

Given our assumption on household preferences, $P_{kt}e^{\epsilon_{kt}}(N_{kt})^{\frac{1}{\phi}}$ is the marginal rate of substitution between labor and consumption and the parameter λ measures the degree of nominal wage stickiness. In particular, when $\lambda = 1$ wages are fully flexible and when $\lambda = 0$ they are fixed. This implies that workers will be off their labor supply curves whenever $\lambda < 1$. A similar specification has been used by Shimer (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.²⁸ In fact, this wage setting rule can be derived from the monopsonistic competition setup in the case where agents are myopic about the future; or the labor market search setup in the special case where firms make take it or leave it offers and the probability of being employed in the future is independent of the current employment status. While there is no forward looking component in the reset wage in our base specification, we consider the implications of including forward looking behavior in Section 7.4 below.

²⁸Nothing prevents the target wage from having a forward looking component. This is the case in Shimer (2009) where the target wage is determined through axiomatic Nash wage bargaining.

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 an interest rate rule $i_t = i(\cdot)e^{\mu_t}$ and given exogenous processes $\{z_{kt}^x, z_{kt}^y, \epsilon_{kt}, \delta_{kt}, \mu_t\}$, they are consistent with household utility maximization and firm profit maximization and such that the following market clearing conditions hold:

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

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. First, define $\gamma_{kt} \equiv \delta_{kt} - \delta_{kt-1} + \mu_t$. Then,

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

with $\sum_k v_{kt}^y = \sum_k v_{kt}^x = \sum_k v_{kt}^\gamma = \sum_k v_{kt}^\epsilon = 0$. By assumption, we assume the weighted average of the regional shocks sum to zero in all periods.

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, $v_{kt}^y, v_{kt}^x, v_{kt}^\gamma, v_{kt}^\epsilon$ 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 all islands are identical with respect to their underlying production parameters (α , β , and θ), their underlying utility parameters (σ and ϕ) and the degree of wage stickiness (λ).²⁹ Our second assumption is that the islands are identical in the steady state and that price and wage inflation are zero. The last assumption is that the joint distribution of island-specific shocks is such that its cross-sectional summation is zero. If K , the number of islands, is large this holds in the limit because of the law of large numbers. We log-linearize the model around this steady state and show that it aggregates up to a representative economy where all aggregate variables are independent of any cross-sectional considerations to a first order approximation.³⁰ We denote with lowercase letters a variable's log-deviation from its steady state. Also, variables without a k subscript represent aggregates. For example, $n_{kt} \equiv \log\left(\frac{N_{kt}}{N}\right)$ and $n_t \equiv \sum_k \frac{1}{K} n_{kt}$. We assume that the monetary authority announces the nominal interest rate rule in log-linearized form: $i_{t+1} = \varphi_\pi \mathbb{E}_t[\pi_{t+1}] + \varphi_y (y_t - y_t^*) + \mu_{t+1}$ where π_t is the aggregate inflation rate and $y_t - y_t^*$ is the output gap; defined as the difference between output and the flexible wage equilibrium output for the same realization of shocks. Finally, we assume that the endogenous component of the discount factor is $\Phi(\cdot) = \Phi_0 (c_{kt} - c_t)$.³¹

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

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

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.

²⁹When implementing our procedure empirically using data on US states, we discuss the plausibility of this assumption. Given that the broad industrial composition at the state level does not differ much across states, the assumption that productivity parameters and wage stickiness are roughly similar across states is not dramatically at odds with the data. As a robustness exercise, we estimate our key equations with industry fixed effects and show that our key cross section estimates are unchanged.

³⁰The 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.

³¹When $\Phi_0 > 0$ this will be enough to induce stationarity of island level variables in log-deviations from the aggregate. At the same time, since $\Phi(\cdot)$ depends only on these deviations, the aggregate equilibrium will feature a constant endogenous discount factor ρ .

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

Proof. The following equations characterize the log-linearized equilibrium

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

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 $\{\pi_t^w, w_t^r, n_t\}$ is characterized by

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

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

To prove Lemma 2, just take log-deviations from the aggregate in the original sys-

tem. This results in the system characterizing the evolution of $\{\tilde{p}_t, \tilde{w}_t, \tilde{n}_t^y, \tilde{n}_t^x\}$ for given $\{\tilde{z}_t^y, \tilde{z}_t^x, \tilde{\gamma}_t, \tilde{\epsilon}_t\}$,

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

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

The guess solution to the aggregate system is $[\pi_t^w, w_t^r, n_t]' = Q[\gamma_t, z_t, \epsilon_t]'$ and the solution to the regional system is $[\tilde{w}_t, \tilde{b}_t, \tilde{p}_t, \tilde{n}_t^y]' = \tilde{P}[\tilde{w}_{t-1}, \tilde{b}_{t-1}] + \tilde{Q}[\tilde{\gamma}_t, \tilde{z}_t^y, \tilde{z}_t^x, \epsilon_t]'$.

4.6 Aggregate vs. local shock elasticities

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

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

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

On other hand, since island level economies in deviations from the aggregate are

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

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

While it is generally understood that local and aggregate elasticities can differ, there has been little quantitative work assessing the potential size of these differences. A parameterized version of our model can allow us to directly compute the local and aggregate employment elasticities to different types of shocks. To this end, Table 4 quantifies the employment impact elasticities to each of the shocks in the full model. Table 3 discusses the parameterization of our model. Most of the parameters are standard from the literature or are chosen to match the labor share in the tradable and nontradable sectors. The Online Data and Robustness Appendix that accompanies the paper has an extended discussion of our baseline parameter choice. For our base specification, we use estimates of λ and ϕ of 2 and 0.7, respectively. These are the parameters that show up in the aggregate and local wage setting equations. The value of these parameters are the ones that we estimate using local variation in Section 6.

Column 1 of Table 4 shows our base estimates of the local and aggregate employment elasticities. In columns 2 - 6 of Table 4, we show how the elasticities change across alternate parameterization. Specifically, in column 2, we re-compute the elasticities reducing the Frisch elasticity of labor supply (ϕ) from 2 to 1. In column 3, we make wages more sticky by reducing λ from 0.7 to 0.5. In column 4, we reduce both λ and ϕ to 1 and 0.5, respectively. In column 5, we set $\beta = 0$ shutting down the possibility to substitute labor for intermediate goods in the production of final goods. In the final column, we shut down the endogenous feedback in the nominal interest rate to changes in the employment gap such that φ_y is set to zero. This last case could be thought of as a representation of the economy when we are getting closer to the zero lower bound and nominal interest rates do not respond to the employment gap.

In our base specification, we find that the regional employment elasticity to a discount rate shock is between 2 and 2.5 times larger than the aggregate employment elasticity

to a discount rate shock. This implies that a growing literature that uses cross region variation to estimate local employment elasticities to demand shocks dramatically overstate employment responses when they apply those local elasticities to the aggregate. Across the different parameterizations of the wage setting rule shown in columns 2-4, the conclusion remains unchanged. Local employment elasticities to discount rate shocks are always two to three times larger than the aggregate employment elasticities. In columns 5 and 6, we see the importance of the general equilibrium forces. As nominal interest rates do not respond to changes in the employment gap, the local employment elasticity to a discount rate shock is only 1.6 times greater than the aggregate employment elasticity. Likewise, as β goes to zero, the local employment elasticity to a discount rate shock is about 4 times the aggregate employment elasticity. Notice, the estimates show that the discount rate shock is much larger at the local level than at the aggregate level. For the productivity/mark-up shocks and the taste for leisure shocks, the local employment elasticities are smaller than their aggregate counterparts. Much of this has to do with the specification of the nominal interest rate rule.

Tables 5 and 6 summarizes the aggregate and regional impulse responses, respectively, for all variables and shocks in our benchmark calibration. We get a sense of the impulse response by showing the results upon impact (the short run elasticities) and after 5 years (the long run elasticities). These tables allows one to assess the model predictions. As seen from Table 5, an aggregate negative discount rate shock (households become less patient) lowers aggregate employment, lowers aggregate prices, and lowers (slightly) aggregate real wages. Conversely, an aggregate negative productivity shock lowers aggregate employment, raises aggregate prices, and raises aggregate real wages. We will compare these to the estimated impulse responses in the vector-autoregression model from Section 5.

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

The above model was designed to (1) link the aggregate and regional economies, (2) specify the local and aggregate wage setting equations, and (3) provide a quantitative assessment of aggregate and local elasticities to the shocks embedded in the model. However, the simple model abstracts from many features which could be important in quantifying the underlying aggregate and local shocks. In this section, we develop a procedure that allows us to estimate the aggregate and local shocks under a broader class of models. But, we can be agnostic with respect to the other components of the model aside from the wage setting equation and still identify the underlying shocks. To do so, we estimate a semi-structural VAR where the structure comes from imposing theoretical restrictions to help identify the underlying shocks. As an overview, the theoretical restrictions imposed by theory implies a series of particular linear restrictions linking the observable variables to the underlying shocks. When we impose these linear

restrictions, the VAR becomes fully identified without relying on ordering restrictions, sign restrictions or long run restrictions. In essence, our procedure uses some elements of theory to help identify the underlying economic shocks. To formalize things, we use the wage setting equation developed above to generate the additional theoretical restrictions while being agnostic with respect to the rest of the model.

5.1 Estimating the Aggregate shocks

The recursive solution to the equilibrium system of equations in Lemma 1 can be written in reduced form as a VAR(∞) in $\{p_t, w_t, n_t\}$.^{32 33}

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

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.

The benefit of estimating VARs to uncover underlying shocks is that the VARs nest a larger class of models. However, additional restrictions are needed to identify the shocks. Some identification schemes rely on ordering of the shocks. These identification schemes are often made without theoretical justification. Other identification schemes rely on long run restrictions. In terms of identification, we will use one component of the theory developed above to identify the VAR. In particular, we will use the wage setting equation. As a result, we do not need to take a stance on the rest of the underlying model (aside from our assumptions of the linear system of equation and the orthogonality of the underlying shocks). For example, their may be price stickiness in the firm's setting of prices or habits in an individual's consumption decision. As a result, our VAR will be consistent with models where the following wage setting equation in log-linearized form holds:

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

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 this one structural equation imposes several linear constraints that the reduced form errors must satisfy. At the end of Section 7, we show the robustness of our estimation procedure to other wage setting equations.

The first step in our procedure consists of estimating the reduced form VAR to obtain

³²With some abuse of notation $\{p_t, w_t, n_t\}$ are the growth rates of prices, wages and employment.

³³The 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.

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 above wage setting equation and constructing the reduced form expectational errors we obtain:

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

Similarly, constructing $\mathbb{E}_{t-1}(\cdot) - \mathbb{E}_{t-2}(\cdot)$, we 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. If the VAR includes two lags, we can construct $\mathbb{E}_{t-1}(\cdot) - \mathbb{E}_{t-2}(\cdot)$.

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

For completeness, the matrix Λ solves the system:

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

5.2 Estimating the Regional Shocks

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

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

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,

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

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 with the procedure for identifying aggregate shocks, we can identify the impulse matrix $\tilde{\Lambda}$ with these extra linear restrictions and label the shocks accordingly using the same restrictions on the co-movement on impact of the variables in the VAR.

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:

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

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

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

In order for regional data to be used to estimate λ and ϕ , one of the following must hold: (1) the labor supply shock has no regional component ($v_{kt}^\epsilon = 0$) or (2) the regional component of the labor supply shock must be uncorrelated with changes in local economic activity (i.e., $cov(v_{kt}^\epsilon, n_{kt}) = 0$ 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:

$$w_{kt} = b_0 + b_1 p_{kt} + b_2 n_{kt} + b_3 w_{kt-1} + \Psi D_t + \Gamma X_k + e_{kt}$$

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-region variation to estimate the parameters. We estimate this equation annually because we only have annual measures of wages at the state level. Our annual nominal wage measures at the state level are the composition adjusted nominal log wages computed from the American Community Survey discussed above. 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 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 on which our wage indices (price indices) are based, we can split the sample in each year and compute two measures of wage indices (price indices) for each state within each year. For example, if we have 1 million observations in the 2007 American Community Survey, we split the sample into two distinct samples with 500,000 observations each. Within each sample, we can compute a wage

measure for each state. The wage measures within each sub-sample, will be measured with error. We can use the growth rates in wages in one half of the samples as an instrument for growth rate in wages in the other half of the samples. We discuss these procedure in detail in the Online Robustness 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 provide some evidence suggesting that this assumption may be valid during the 2007-2010 period. 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)). The necessary 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

There is essentially no variation in state labor income tax rates or variation in Supplemental Nutrition Assistance Program (SNAP) benefits per recipient across U.S. states during the Great Recession. We show details of this analysis in the Online Data and Robustness Appendix. SNAP benefits per recipient are set at the federal level. The average recipient received an increase about 33 percent during the 2007 to 2010 period. Yet, there was essentially no regional variation in the increase. The increase in SNAP benefits per recipient at the federal level was one of the factors emphasized by Mulligan as deterring aggregate labor supply during this period. Another new federal program that was means tested and was argued to possibly discourage work effort was the Home Affordable Modification Program (HAMP). HAMP was designed to help homeowners who were underwater renegotiate their mortgage. The program was authorized in early 2009 but there were no significant modifications taking place until mid 2010. By the end of 2010, only about 0.5% of households had participated in the program. As we show in the Online Appendix, the extent to which households participated in the program varied slightly with the state's employment growth between 2007 and 2010. While these programs may have differentially discouraged labor supply across states during later periods, the fact that the program did not have much of an effect through 2010 means that it is unlikely it can explain the differential labor supply patterns across states through 2010. However, as a robustness specification, we restrict our analysis to pre-2010 data

³⁴In the Online Data and Robustness 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. We also discuss more fully our industry controls, our measures of migration and instrumenting procedure.

prior to the program's inception. Additionally, when extending our analysis through 2010, we exclude the high mortgage default sand states as a further robustness exercise.

One policy that has received consider attention is the differential extension of the duration of unemployment benefits across states during the Great Recession. By law in 2010, weeks of unemployment benefits were tied to the state's unemployment rate. As of 2010, most U.S. states had a duration of unemployment benefits that was close to the maximum of 99 weeks. These states with at least 86 weeks of unemployment benefits represent roughly 90 percent of the U.S. population. However some smaller states, mostly in the Plains region of the U.S., had small employment declines and only had an extension of unemployment benefits from 60-85 weeks.³⁵ While some smaller states had benefit extensions that were much lower than other states, it should be noted that by 2010 almost all states had unemployment benefit extensions that were close to the maximum. Despite this, we still perform two additional robustness exercises. First, when using our full time period, we exclude any state that had less than 85 weeks of unemployment benefit extensions. Given that these were only a handful of smaller states, such exclusion had no effect on our OLS estimates. Additionally, as noted above, we estimate our key parameters using only data prior to 2010 prior to the unemployment duration extensions taking place.

Finally, 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.³⁶ 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 dur-

³⁵States also had some discretion as to whether they opted into the program. This explains why some states did not have the maximum weeks of unemployment benefits even when their unemployment rate was higher.

³⁶For 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 1.8 percent.

ing the 2007-2011 period (in our base specification) or during the 2007-2009 period (in our restricted specification) is orthogonal to movements in local labor supply shocks. This seems like a plausible assumption for the 2007-2009 period as state policy changes did not occur prior to 2009. In the Online Robustness 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.

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

6.2 Estimates of λ and ϕ

Column 1 of Table 7 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. 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. Our base estimate of $\lambda = 0.69$ suggests only a modest amount of wage stickiness.

Columns 2-4 of Table 7 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.³⁷ 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 roughly 1.0 and 2.0. 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 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 if wages are fairly flexible. Put another way, if wages were sticky enough in the aggregate to have demand shocks be the primary driver of aggregate employment declines during the recent recession, we would not observe wages moving as much as they did in the cross section during the same time period.

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 employment fell during the Great Recession. However, exploiting variation across regions, there appears to be a significant negative relationship between nominal wages and local employment 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

³⁷We 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.

and to assess their contributions to the behavior of prices, wages and employment during this period at both the aggregate and local level.

7.1 Findings in the aggregate

We follow the procedure described in Section 5.1. We first estimate the VAR with two lags in aggregate employment, 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 $\hat{V} = \frac{UU'}{Years - \#Variables \#Lags}$.

The aggregate variables we construct are comparable to our regional measures. Given that our cross-sectional equations are estimated using annual data, we analogously define our aggregate data at annual frequencies. We use data from the CPI-U to create our measure of aggregate prices. Specifically, we take log-change in the CPI's between the second quarter of year t and $t - 1$ for our measure of p_t . For n , we use BLS data on the aggregate employment to population rate of all males 25-54. We choose this age range so as to abstract from the downward trend in employment rates due to the aging of the population over the last 30 years.³⁸ 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 time. Specifically, 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. These controls are similar to the way we adjusted our aggregate wage index using ACS data in Section 3. 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. As with the CPI, we take the log-change in this wage measure between the second quarter of year t and $t - 1$ for our measure of w_t .

Figures 7, 8, and 9 report the impulse response of aggregate employment, nominal wages and price growth to each of the shocks when we use our benchmark estimates for λ and ϕ reported in column 1 of Table 7 ($\lambda = 0.69$ and $\phi = 2.2$). Figure 7 shows their behavior after an initial discount rate/monetary (γ) shock of the same magnitude and sign as in 2008. Qualitatively, after a negative discount rate shock both prices and employment fall sharply relative to trend while real wages decline slightly relative to trend. These results are identical to the theoretical predictions shown in Table 5.

³⁸We detrend all data when estimating the VAR. Specifically, we allow for a linear trend in the employment to population ratio between 1978 and 2007. For the price inflation rate and the nominal wage inflation rate, we use an HP filter (with a smoothing parameter of 100). Given that we detrend the data, our results are essentially unchanged when we use the employment to population ratio for all individuals as opposed to using it just for prime age males.

Figure 8 shows the impulse responses to a 2008 productivity/markup (z) shock. Prices relative to trend increase on impact while employment growth falls sharply. Nominal wages, however, only decline slightly (with a lag). Again, these predictions match the predictions of our simple theoretical model shown in Table 5. While both negative γ and z shocks reduce employment, the γ shock puts downward pressure on prices while the z shock puts upward pressure on prices. Figure 9 shows the impulse response of employment, prices, and nominal wages to the leisure shock. Upon impact, the leisure shock reduces employment and prices while it increases wages. Again, these predictions match the predictions from the simple model. It is the leisure shock that is putting upward pressure on nominal wages in the aggregate.

We turn now to the cumulative response of each individual variable when we feed the VAR with the sequence of shocks between 2008 and 2012 one at a time.³⁹ 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 4 percent between 2008-2009 (relative to trend) and remained at the low level thereafter. The counterfactual exercise shows that the productivity/markup and discount rate shocks contributed about the same amount to the initial decline during the 2008-2009 period (each explaining roughly 40 percent of the aggregate employment decline). However, the discount rate/monetary does not explain any of the persistence in the employment decline post 2009. Instead, it is the productivity/mark-up shock that explains the sluggish response of employment post 2009. The leisure shock barely contributed to the observed employment decline either on impact or over the longer 2008-2012 period.

Figures 12 and 13 provides insight into the shock decomposition as well as to helping understand the "missing deflation puzzle" and "missing nominal wage decline puzzle". Figure 12 shows the counterfactual price response to each of the shocks. With respect to the data, aggregate prices fell relative to trend between 2008 and 2009. However, prices quickly stabilized relative to trend. This is the sense that there was "missing deflation". Despite the weak employment situation post 2009, prices were growing relative to trend. Both the discount rate/monetary shock and the leisure shock put downward pressure on aggregate prices. However, the productivity/markup shock put upward pressure on aggregate prices. The counterfactual analysis shows that if the economy was only hit with the productivity/mark-up shock, prices would have been rising (by upwards of 2 percent) relative to trend during the Great Recession. According to our procedure, it is this countervailing productivity/markup shock that arises as the explanation for the missing deflation puzzle - particularly post 2009.

Figure 13 shows the cumulative nominal wage response to each of the shocks. Again, the figure shows the missing nominal wage puzzle during the Great Recession. Throughout the early part of the recession and the entire recession, nominal wage growth was close to zero (relative to trend). However, our model shows - like conventional wisdom

³⁹For the interested reader, the actual realizations of the shocks we estimate can be seen in Figure 10.

- if the economy only experienced the discount rate/monetary shock, nominal wages should have fallen by roughly 2.5 percent relative to trend by 2010 and remained well below trend in 2012. The leisure and productivity/markup shocks, however, are necessary to explain why nominal wages did not fall during the Great Recession. Some of the exit in the labor force due to the leisure shock put upward pressure on nominal wages. Likewise, the productivity shock also put upward pressure on nominal wages. In summary, our methodology suggests that demand shocks cannot be solely responsible for the employment decline during the Great Recession. If demand shocks were solely responsible, price inflation would have been lower - particularly post 2009 and nominal wages would have fallen. Our method suggests that a combination of productivity shocks and leisure shocks are needed to explain the missing deflation and nominal wage declines. Moreover, our methodology allows us to quantify the relative importance of each shock in explaining the aggregate movement in wages, prices, and employment during the Great Recession.

7.2 Robustness to Estimated Parameters λ and ϕ

How do our estimated parameters affect our employment, price and wage counterfactuals? In Table 8, we report the contribution of each shock to the explanation of aggregate employment declines implied by different combinations of $\{\phi, \lambda\}$. We do this for both the initial years of the recession (2008 to 2009) and over the longer period encompassing the recovery (2008 to 2012). The table shows that the qualitative conclusions of the previous section still hold for the range of $\{\phi, \lambda\}$ estimates using alternative specifications that we report in Table 7. These go from roughly 0.5 to 0.8 for λ and from roughly 1.0 to 2.5 for ϕ . When reading Table 8, each cell shows the decomposition of how much of the employment change during the time period can be attributed to the discount rate/monetary shock (γ) and how much can be explained by the productivity mark up shock (z). The sum of all three shocks should sum to 100 percent. So, the difference between the sum of the γ and z contributions and 100 percent is attributed to the leisure (ϵ) shock.

Tables 8 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.⁴⁰ However, suppose we used a much lower elasticity $\phi = 0.5$ instead which is in line with some microeconomic estimates in the literature. In this case, the leisure shock would

⁴⁰This 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.

account for a much larger fraction of the employment decline in the Great Recession. This comes out of the VAR estimation. However, the intuition for this result is straightforward and in line with the simple theory we wrote down above. If labor supply is fairly elastic, 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.

The intuition for the decomposition between discount rate and productivity/markup shocks is more subtle but also in line with the simple theory we wrote down. We find that the degree of wage flexibility λ affects the relative importance of one vis a vis the other within the remaining unexplained part by the leisure shock. For example, suppose we increased the degree of wage flexibility to $\lambda = 0.9$ (very little wage rigidities) instead of using our estimated $\lambda = 0.69$. Then the productivity/markup shock would account for a much larger fraction of the employment decline in the Great Recession. Theoretically, it is clear that when λ is large, the discount rate shocks do not matter much for the determination of employment. To see this, consider the extreme case where wages are perfectly flexible and the demand shock is only composed of the monetary shock. Then the equilibrium in the simple theoretical model satisfies monetary neutrality. We formalized this point in Section 4.6 when we derived the model's implied elasticity of aggregate employment to a discount rate shock. Conversely, when wages are very rigid ($\lambda = 0.1$), our procedure is suggesting that demand shocks can explain essentially all of the decline in the early part of the recession and much of the persistence in employment declines during the 2008-2012 period. Given that all of our estimates of λ are around 0.7, the demand shock cannot have been the sole cause of the recession because if it did nominal wages would have fallen and prices would have declined relative to trend after 2009.

7.3 Findings at the regional level

We follow the procedure in Section 5.2. We first estimate the VAR with two lags in state-level non-tradable employment⁴¹, price and nominal wage growth via OLS equation by equation. All variables are in expressed in log-deviations from their average weighted by population in 2006. We pool all data between 2006 and 2011, and estimate common autoregressive coefficients and reduced form errors covariance matrix for all states. In our benchmark specification, we set $\{\phi, \lambda\}$ equal to our estimates from Section 6.2. We set $\theta = 0.55$ to match the labor share in the manufacturing sector in the US and $\frac{N^y}{N} = 0.85$ to match the share of total employment in the service sector plus self-employed/family workers as reported in the BLS.

Table 9 summarizes the contribution of the discount rate shock and the combined productivity/markup shocks to non-tradable employment, wages and prices. We define

⁴¹We define non-tradable employment in a state as the employment rate in the service and retail sectors combined.

shock j 's contribution to the change in variable y between 2007 and 2010 as

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

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

For our benchmark $(\lambda, \phi) = (0.7, 2)$ we find that the discount rate shock contributed 79 percent to the change in non-tradable employment between 2007 and 2010 across all states; 41 percent to local price changes and 24 percent to local wage changes. We obtain similar numbers for $(\lambda, \phi) = (0.5, 1)$. We conclude that the discount rate shock was the main driver of regional variation in non-tradable employment.

Table 10 summarizes a second set of results from the regional counterfactuals. We characterize the joint distribution of cumulative growth rates between 2007 and 2010 for each variable across states with two statistics: the variance and the correlation with each other. We compare the actual statistics from our data in Section 3 with the counterfactuals obtained when simulating one shock at a time. Again, we find that the discount rate shock alone can generate 98 percent of the cross-state variance of non-tradable employment growth; 69 percent of the price growth variance; and 47 percent of the nominal wage growth variance. Moreover, it reproduces the right sign for the cross-state correlations of price growth and non-tradable employment growth; nominal wage growth and non-tradable employment growth; and nominal wage growth and price growth. Although, quantitatively, it generates a larger correlation between prices and non-tradable employment than in the data. Both productivity/markup shocks combined can explain only 35 percent of the non-tradable employment growth variance across states. They do as good a job as the discount rate in explaining 85 percent of the variation in price growth and a worse job in explaining 26 percent of nominal wage growth variance. However, they imply negative correlations between price/wage growth and non-tradable employment growth. The opposite is observed in the data. We conclude that the discount rate shock alone does a fairly good job in reproducing the regional patterns that we documented in Section 3. To improve the estimated fit, nontradable and tradable productivity/markup shocks are needed. However, their contribution to aggregate employment changes are small.

7.4 Robustness to Alternate Identifying Assumptions

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

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

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

In Appendix A1, we derive the aggregate and local wage setting equations under a broad set of utility functions where consumption and leisure are separable. This class of utility functions allow for arbitrarily large income and substitution effects. However, as we show in the appendix, the use of local consumption data allows us to estimate the extent of wage stickiness as well as to estimate the parameters that encompass both the income and substitution effects on labor supply. In particular, we can estimate the following equation using local data:

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

This equation is identical to our estimating equation above aside from the addition of local consumption growth. As outlined in Appendix A1, the coefficients \tilde{b}_1 and \tilde{b}_3 should sum to 1 even under the broader preference specification. We impose this restriction when estimating the modified equation. For our measure of local real consumption growth, we use the change in real retail expenditures at the state level computed within the Nielsen sample. To get real expenditures, we deflate nominal expenditures by our measure of local prices.⁴² For our base specification, our estimates of \tilde{b}_1 , \tilde{b}_2 , and \tilde{b}_4 are, respectively, 0.72 (standard error = 0.12), 0.25 (standard error = 0.08), and 0.16 (standard error = 0.06). The coefficient on real consumption growth (\tilde{b}_4) is positive and significant saying that there is an estimated income effect on labor supply. Controlling for the income effect, our estimate of λ (\tilde{b}_1) is slightly higher than our base specification without allowing for an income effect.

For the aggregate wage setting equation, we can substitute out consumption growth using the model definition ($c_t = w_t + n_t - p_t$). Substituting out consumption growth, Appendix A1 shows that the aggregate wage setting equation still takes the following form:

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

where ω is a parameter that represents the strength of the income effect on labor supply (and maps directly to \tilde{b}_4 from the above local labor supply regression). Aside from the scaling of the aggregate labor supply shock, this equation is identical to the identification restriction we imposed when estimating the aggregate VAR. The only

⁴²This measure of real expenditures is (1) highly correlated with measures of local employment and (2) highly correlated with the BEA's recent state level personal expenditures measure. Our results are similar if we use the BEA's local consumption measure. However, we prefer our measure given that much of the BEA's local consumption measure is imputed (where the imputation uses local employment measures).

difference is that there is no longer a one-to-one mapping between λ and λ/ϕ in the above aggregate restriction on the VAR and the structural parameters \tilde{b}_1 and \tilde{b}_2 from the local wage setting equation. Instead, λ and λ/ϕ aggregate constraint are both functions of \tilde{b}_1 , \tilde{b}_2 , and \tilde{b}_4 . However, as shown in Appendix A1, there is still a specific mapping between the parameters we estimate from the local regression (\tilde{b}_1 , \tilde{b}_2 , and \tilde{b}_4) and the aggregate parameters we need to identify the VAR (λ and ϕ). With the correctly specified λ and ϕ , we can just use the matrix in Table 8 to read off the decomposition of shocks during the Great Recession. While λ and ϕ are no longer structural parameters (instead being combinations of structural parameters), knowing them still fully identifies the aggregate VAR. Using our estimates of \tilde{b}_1 , \tilde{b}_2 , and \tilde{b}_4 and the procedure developed in Appendix A1, we find that our new estimates of λ and ϕ which allow for income effects on labor supply to be 0.68 and 2.0, respectively. These parameters are nearly identical to our base specification without income effects. Having local consumption data allows us to control for income effects on labor supply in our local wage setting equations. The take away from this robustness exercise is that abstracting from preferences that allow for an income effect on labor supply is not biasing our decomposition of the shocks driving aggregate employment declines during the Great Recession in any meaningful way.

In Appendix B, we re-specify our wage setting equation allowing for forward looking behavior when wages are reset. In this Appendix, we show that ignoring forward looking wage setting behavior biases up our estimates of wage stickiness. That is, the amount of wage stickiness that we estimate from our local equation is too large relative to the true amount of wage stickiness. We show that the amount of the bias is dependent on two parameters: (1) the extent to which firms put weight on forward looking behavior when setting wages (which we call κ in Appendix B) and (2) the underlying persistence process of local wages (which we call $\bar{\rho}_w$). Imposing that at the aggregate level that the monetary authority wants to stabilize expected nominal wage growth, we can quantify the extent to which our estimates of λ from the local regressions are biased upwards as a function of these two parameters. Under a range of plausible parameter estimates for κ and $\bar{\rho}_w$, we show that the bias is quite small. For our base specification, where our estimated $\lambda = 0.69$, we show that $\kappa(1 - \bar{\rho}_w)$ must exceed 2.5 for the true λ to be lower than 0.4. These is an order of magnitude larger than any plausible parametrization for either κ or $\bar{\rho}_w$. Moreover, as seen in Table 8, our estimates of the role of demand shocks in explaining employment decline during the Great Recession are similar between $\lambda = 0.69$ and $\lambda = 0.4$. The results suggest that our abstraction from including expectations in our wage setting equation is not quantitatively altering the paper's conclusions.

8 Conclusion

Regional business cycles during the Great Recession in the US were strikingly different than their aggregate counterpart. This is the cornerstone observation on which we built this paper. Yet, the aggregate US economy is just a collection of these regions connected

by trade of goods and assets. We argued 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 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 em-

pirical 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.

References

- An, Sungbae and Frank Schorfheide, 2007. "Bayesian Analysis of DGSE Models", *Econometric Reviews*, 24(2-4), 113-73.
- Autor, David, David Dorn, and Gordon Hanson, 2013. "The China Syndrome: Local Labor Market Effects of Import Competitions in the United States", *American Economic Review*, 103(6), 2121-2168.
- Ball, Laurence and Sandeep Mazumder, 2011. "Inflation Dynamics and the Great Recession", *Brookings Papers on Economic Activity*, 42(1), 337-405.
- Berger, David and Joe Vavra (2015). "Dynamics of the U.S. Price Distribution", Working Paper.
- Blanchard, Olivier and Lawrence Katz, 1991. "Regional Evolutions", *Brookings Papers on Economic Activity*, 1992:1, 1-75.
- Burstein, A. T., Neves, J. C., & Rebelo, S., 2003. Distribution costs and real exchange rate dynamics during exchange-rate-based stabilizations. *Journal of Monetary Economics*, 50(6), 1189-1214.
- Cage, Robert, John Greenless and Patrick Jackman, 2003. "Introducing the Chained Consumer Price Index," U.S. Bureau of Labor Statistics.
- Carriero, Andrea and Raffaella Giacomini, 2011. "How Useful are No-Arbitrage Restrictions for Forecasting the Term Structure of Interest Rates", *Journal of Econometrics*, 164(1), 21-34.
- Charles, Kerwin, Erik Hurst, and Matthew Notowidigdo, 2014. "Housing Booms, Labor Market Outcomes, and Educational Attainment", *University of Chicago Working Paper*.
- Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber, 2011. "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins", *American Economic Review Papers and Proceedings*, 471-5.
- Christiano, Lawrence, Martin Eichenbaum, and Mathias Trabandt, 2014. "Understanding the Great Recession", working paper.
- Christiano, Lawrence, Martin Eichenbaum, and Mathias Trabandt, 2013. "Unemployment and Business Cycles", *NBER Working Paper*, 19265.
- Coibion, Olivier, Yuriy Gorodnichenko, and Gee Hee Hong, 2014. "The Cyclicity of Sales, Regular and Effective Prices: Business Cycle and Policy Implications", University of Texas Working Paper.

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

Del Negro, Marco, Marc Giannoni, and Frank Schorfheide, 2014. "Inflation in the Great Recession and New Keynesian Models", *Federal Reserve Bank of New York Staff Reports*, 618.

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

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

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

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

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

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

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

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

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

Hansbury, Jessie and David Weinstein, 2013. "Is New Geography Right? Evidence from Price Data", *NBER Working Paper #17067*.

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

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

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

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

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

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

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

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

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

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

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

Nakamura, Emi and Jon Steinsson, 2010. "Monetary Non-Neutrality in a Multi-Sector Menu Cost Model", *Quarterly Journal of Economics*, 125(3), 961-1013.

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

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

Shimer, Rob, 2010. *Labor Markets and Business Cycles*, Princeton University Press.

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

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

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

A Alternative wage setting specifications

A.1 Preferences with wealth effects in labor supply

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

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

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

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

$$\begin{aligned} w_t &= \hat{\lambda}(p_t + \epsilon_t + \left(\frac{1}{\hat{\phi}} + \omega\right)n_t + \omega c_t) + (1 - \hat{\lambda})w_{t-1} \\ \tilde{w}_{kt} &= \hat{\lambda}(\tilde{p}_{kt} + \epsilon_{kt} + \left(\frac{1}{\hat{\phi}} + \omega\right)\tilde{n}_{kt} + \omega \tilde{c}_{kt}) + (1 - \hat{\lambda})\tilde{w}_{kt-1} \end{aligned}$$

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

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

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

$$\tilde{w}_{kt} = \lambda \tilde{p}_{kt} + \frac{\lambda}{\hat{\phi}} \tilde{n}_{kt} + (1 - \lambda)\tilde{w}_{kt-1} + \frac{\lambda\omega}{1 - \omega} (\tilde{p}_{kt} + \tilde{c}_{kt} - (\tilde{w}_{kt} + \tilde{n}_{kt})) + \frac{\lambda}{1 - \omega} \tilde{\epsilon}_{kt} \quad (7)$$

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

A.2 Forward looking wages

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

setting equation.

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

$$w_t = \lambda(p_t + \epsilon_t + \frac{1}{\phi}n_t) + \lambda\kappa\mathbb{E}_t[w_{t+1} - w_t] + (1 - \lambda)w_{t-1}$$

$$\tilde{w}_{kt} = \lambda(\tilde{p}_{kt} + \tilde{\epsilon}_{kt} + \frac{1}{\phi}\tilde{n}_{kt}) + \lambda\kappa\mathbb{E}_t[\tilde{w}_{kt+1} - \tilde{w}_{kt}] + (1 - \lambda)\tilde{w}_{kt-1}$$

where κ parametrizes the importance of the forward looking term.

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

We obtain,

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

$$\tilde{w}_{kt} = \frac{\lambda}{1 + \lambda\kappa(1 - \tilde{\rho}_w)}(\tilde{p}_{kt} + \tilde{\epsilon}_{kt} + \frac{1}{\phi}\tilde{n}_{kt}) + \frac{1 - \lambda}{1 + \lambda\kappa(1 - \tilde{\rho}_w)}\tilde{w}_{kt-1}$$

Then, we can write,

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

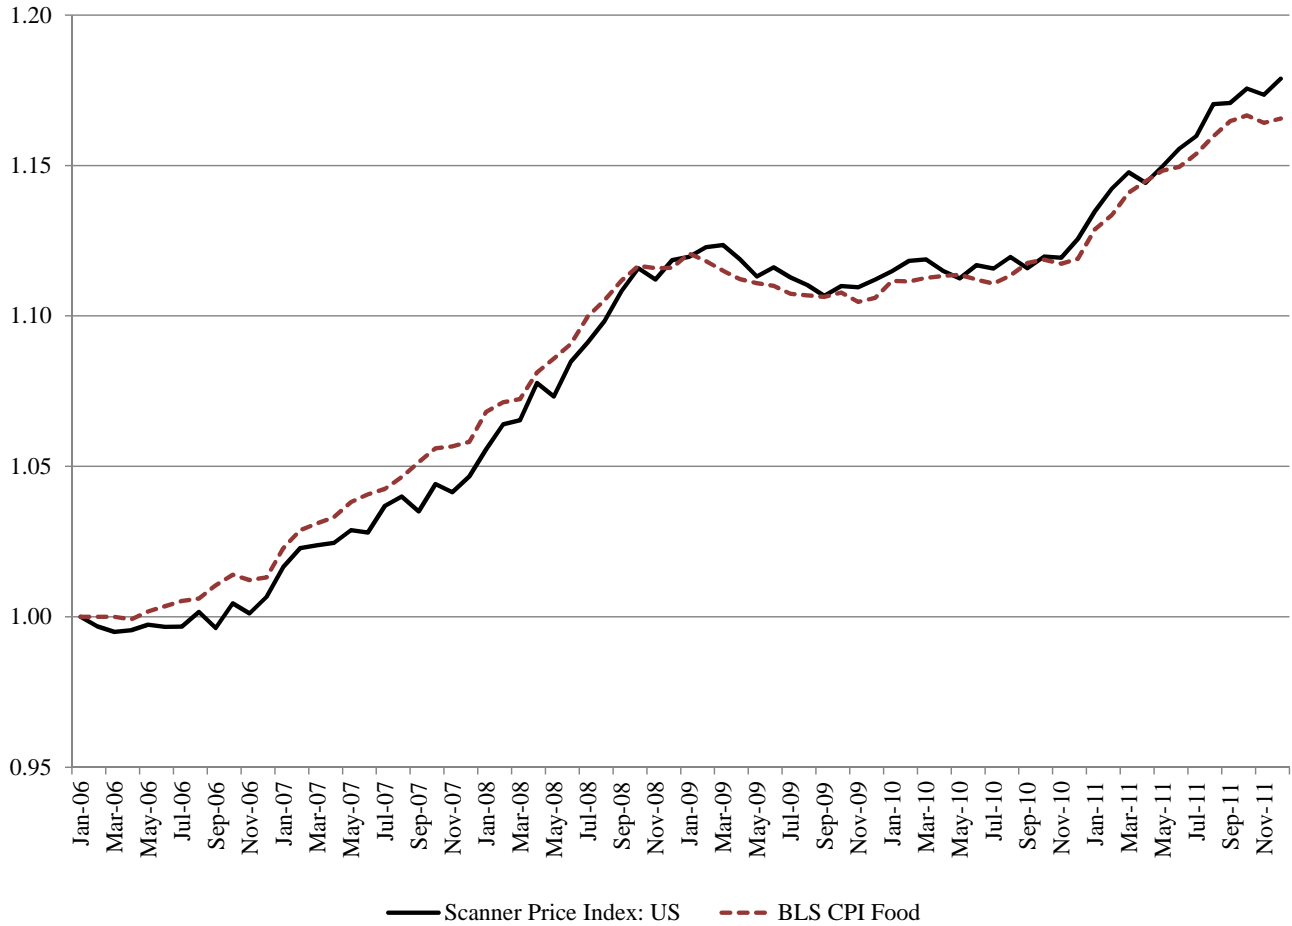
where $\beta_w \equiv \frac{1 - \lambda}{1 + \lambda\kappa(1 - \tilde{\rho}_w)}$. From this expression we see that our estimates for λ using cross-state variation are upward biased. However, we can get a notion on the magnitude of the bias by asking what would $\kappa(1 - \tilde{\rho}_w)$ have to be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

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

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

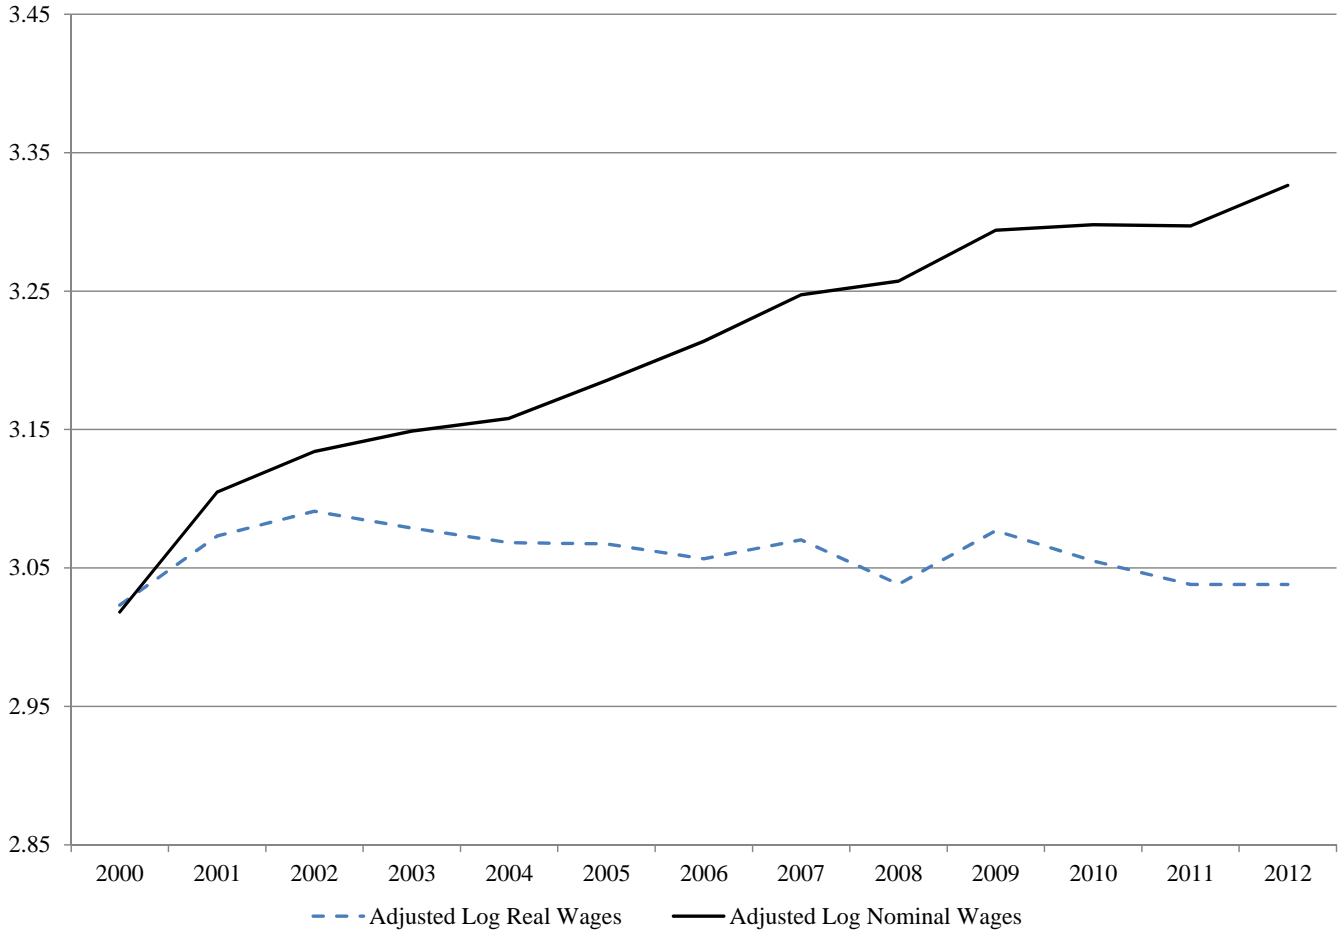
B Figures and Tables

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



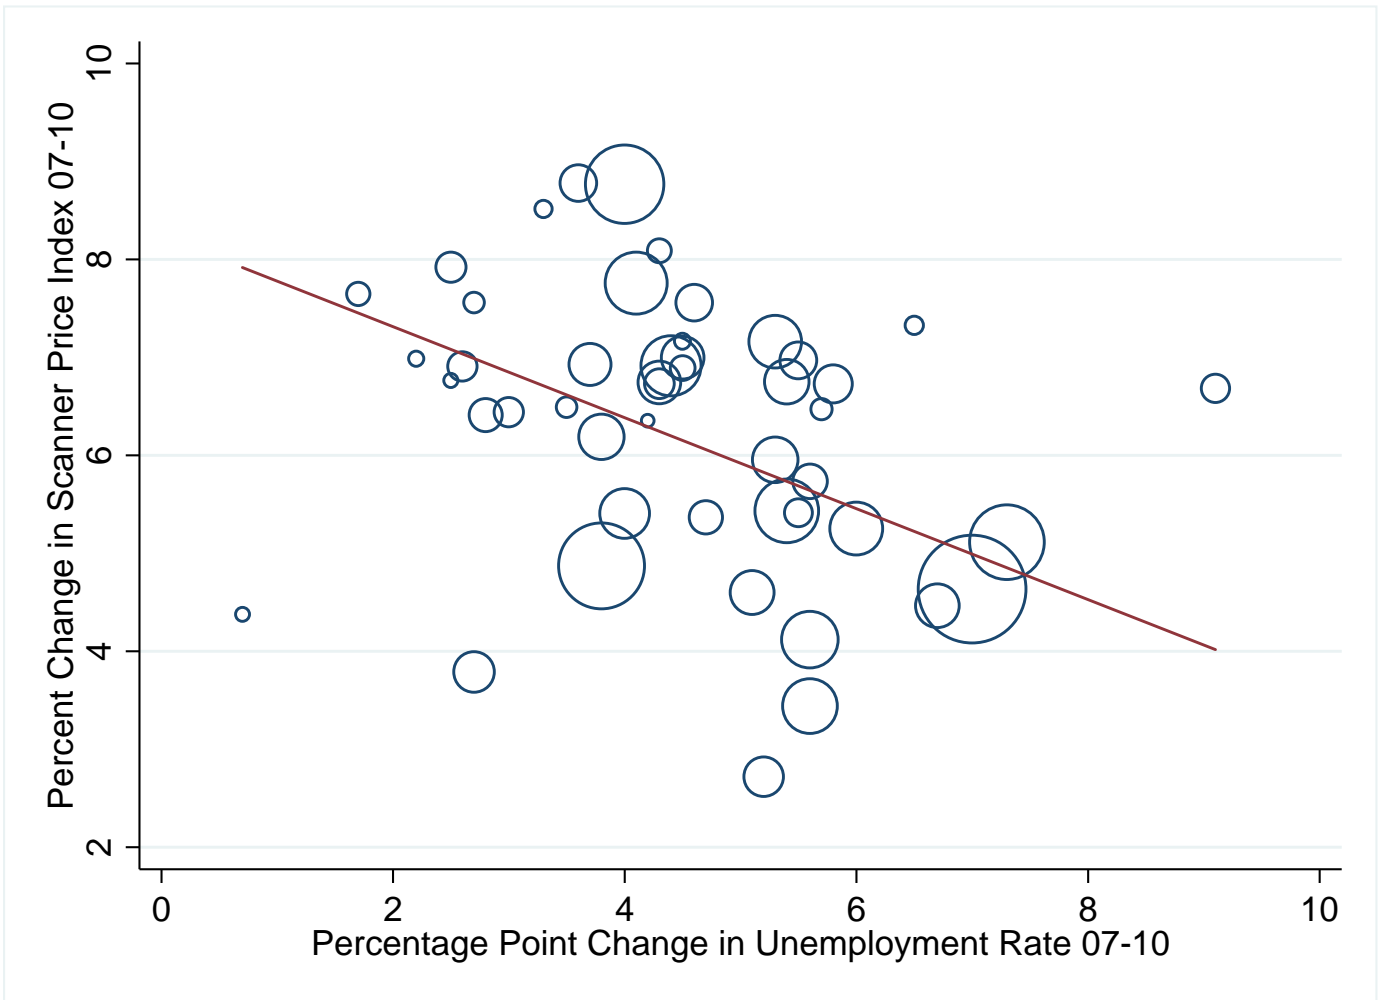
Note: In this figure, we compare our retail price index for the U.S. as whole to the CPI food price index. Given that the goods in our price index come predominantly from grocery, pharmacy, and mass merchandising stores, we thought the food CPI was an appropriate benchmark. For the Nielsen retail price 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)



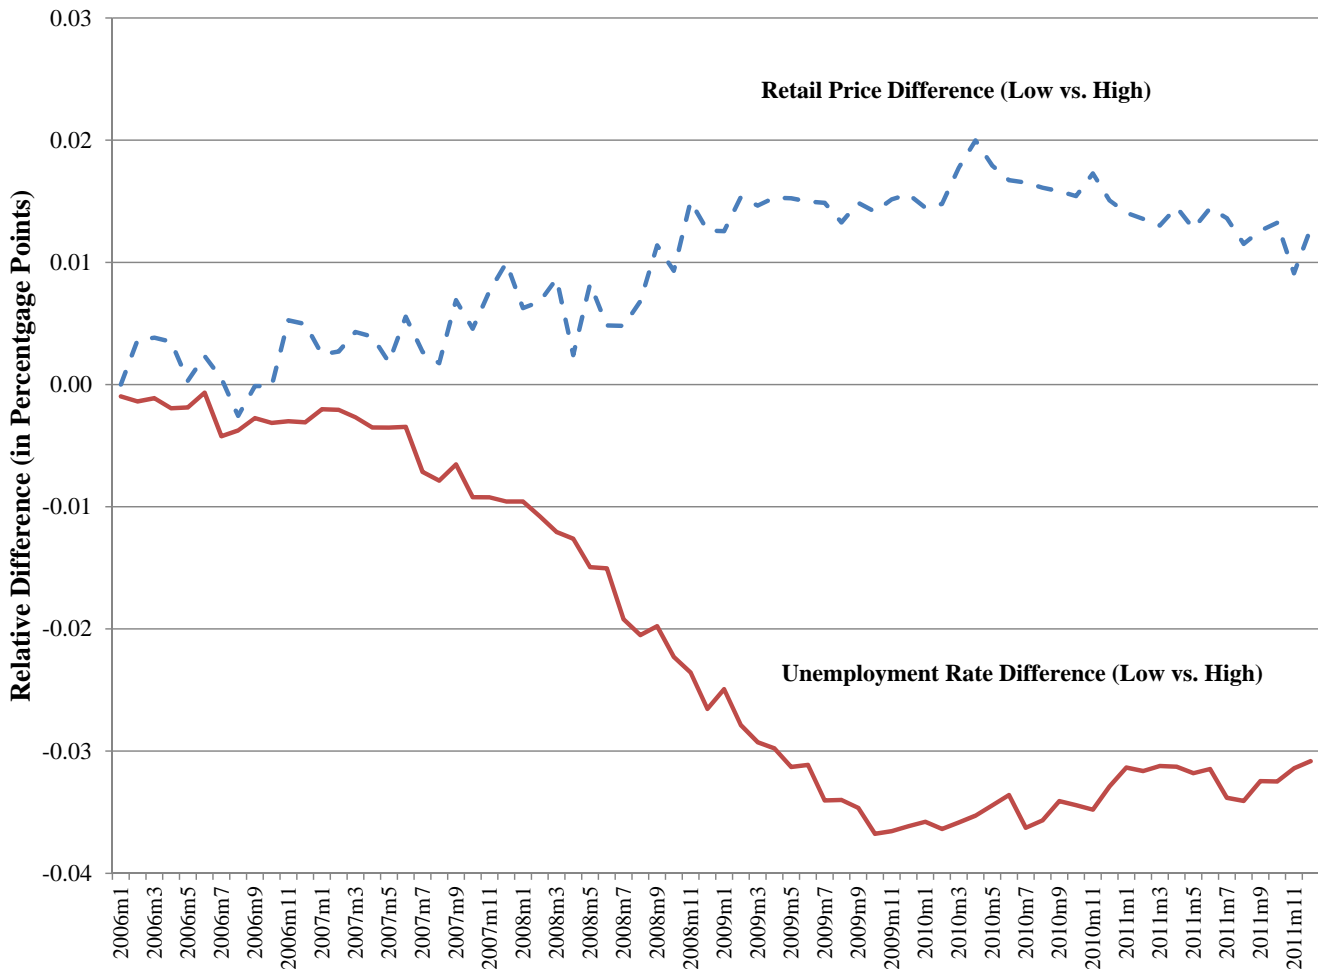
Note: Figure shows the evolution of aggregate real and nominal log wages within the U.S. between 2000 and 2012. Nominal wages were computed using data from the American Community Survey. The sample is restricted to only those individuals who are currently employed, who report usually working 30 hours 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)



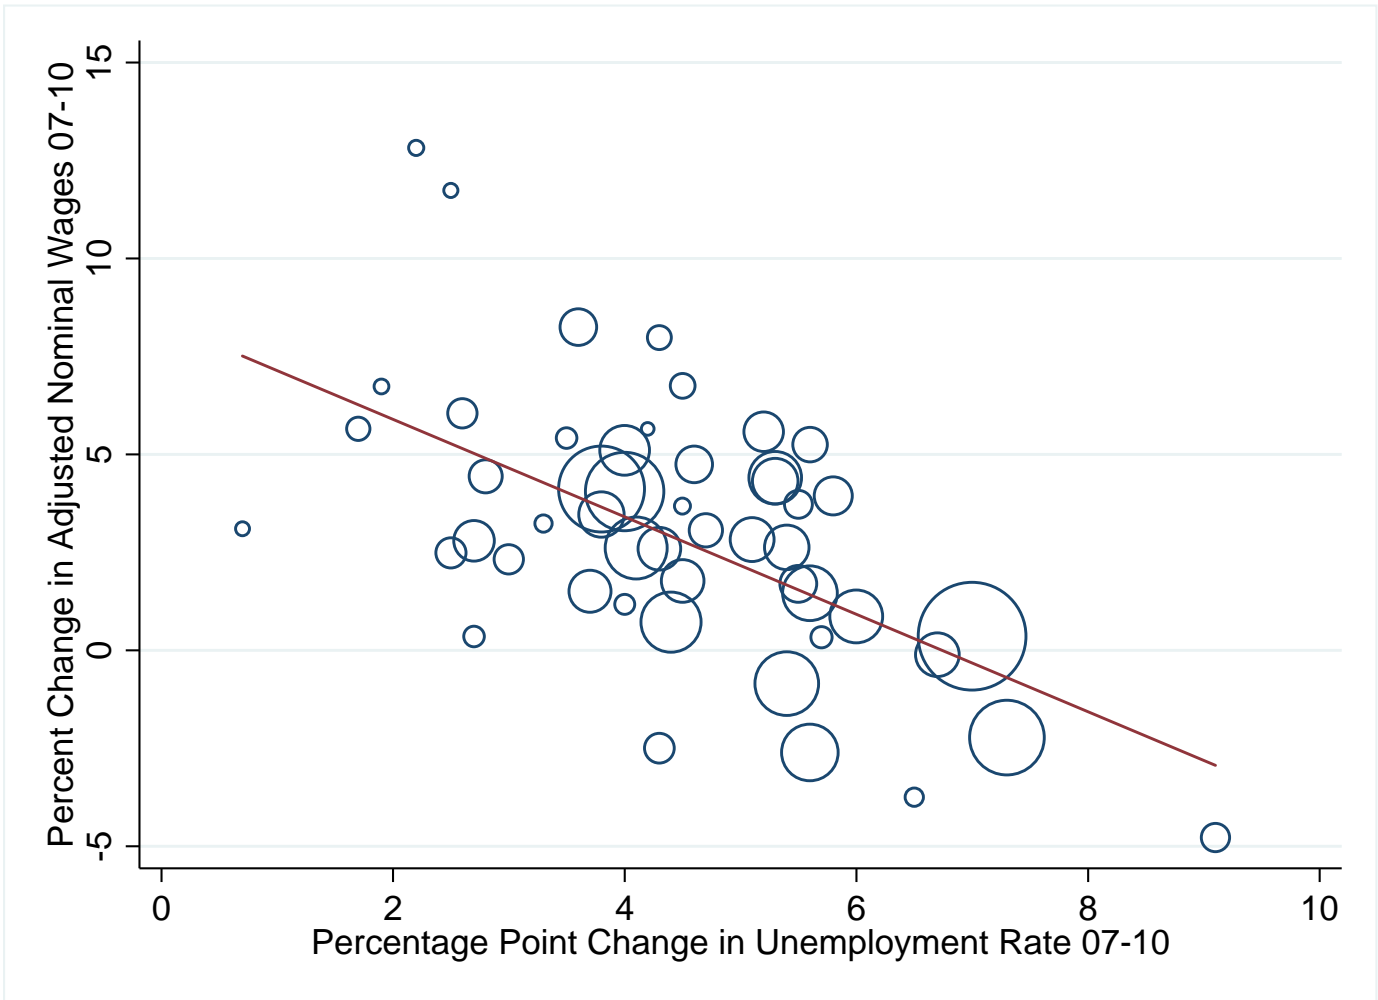
Note: Figure shows a simple scatter plot of the percentage point change in the BLS unemployment rate in the state between 2007 and 2010 against the cumulative percent change in our retail price index based on the Nielsen scanner data during the same period. The retail price index is computed where each 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 i£; 2011M12



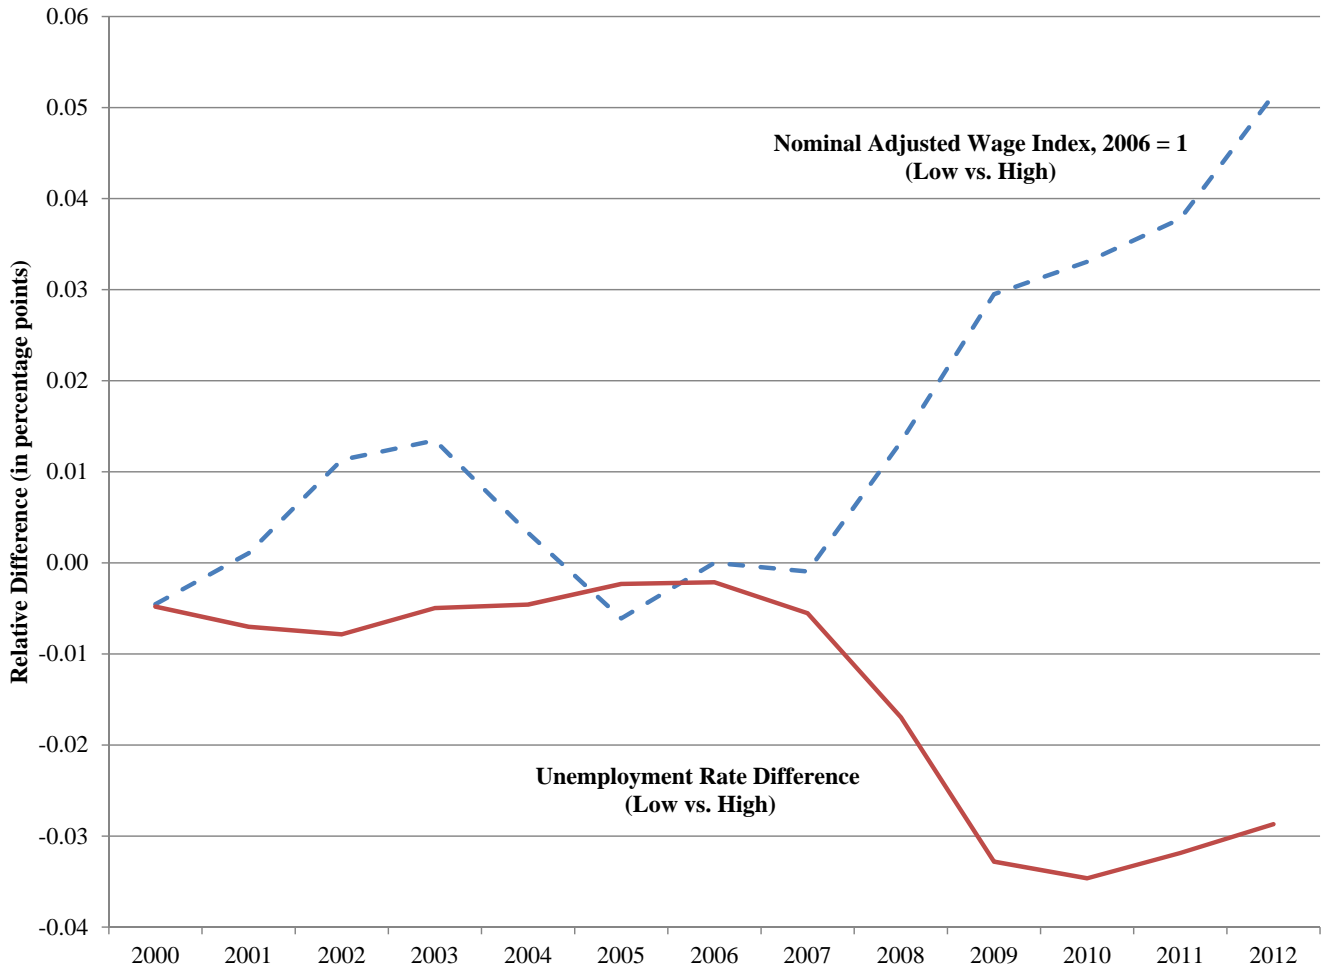
Note: Figure shows the trends in the relative monthly unemployment rate between low and high unemployment change states against the trends in the relative monthly retail price index between low and high unemployment change states. High unemployment change states are the top one-third of all states with 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)



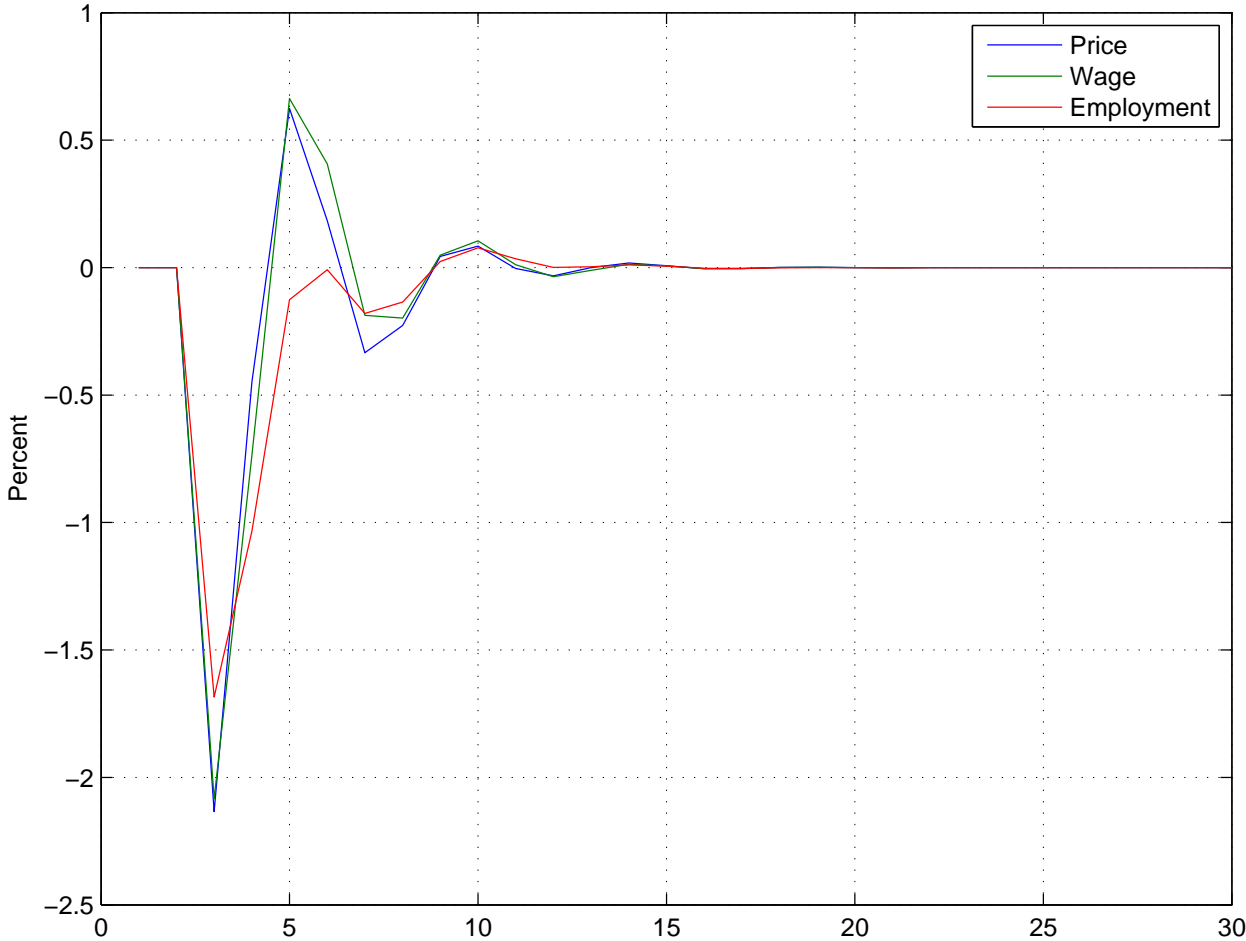
Note: Figure shows a simple scatter plot of the percentage point change in the BLS unemployment rate in the state between 2007 and 2010 against nominal wage growth during the same period. The retail price index is computed where each good is based on a UPC within the state (as opposed to a UPC-store). 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



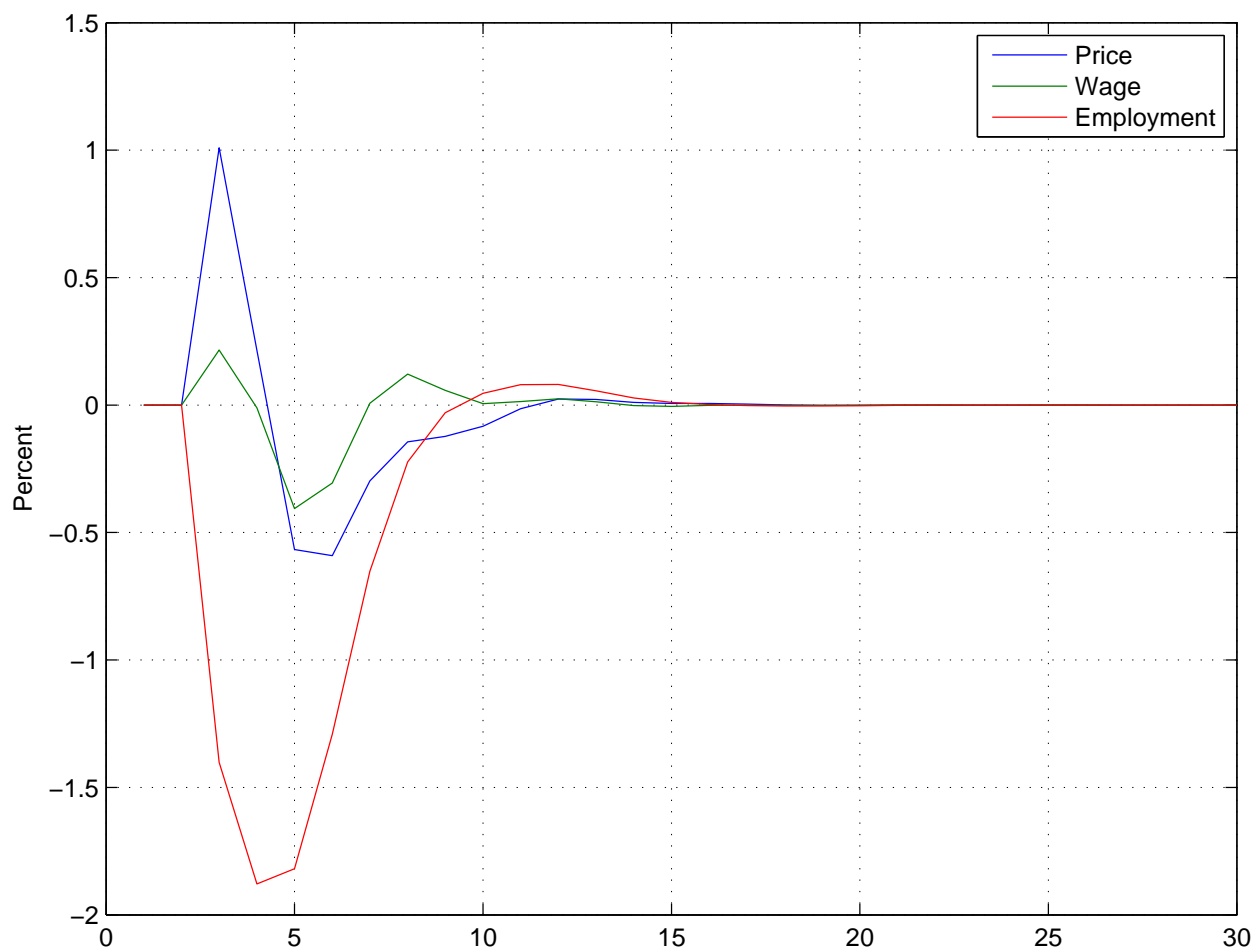
Note: Figure shows the trends in the relative annual unemployment rate between low and high unemployment change states against the trends in the adjusted nominal wage index between low and high unemployment change states. High unemployment change states are the top one-third of all states with respect 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



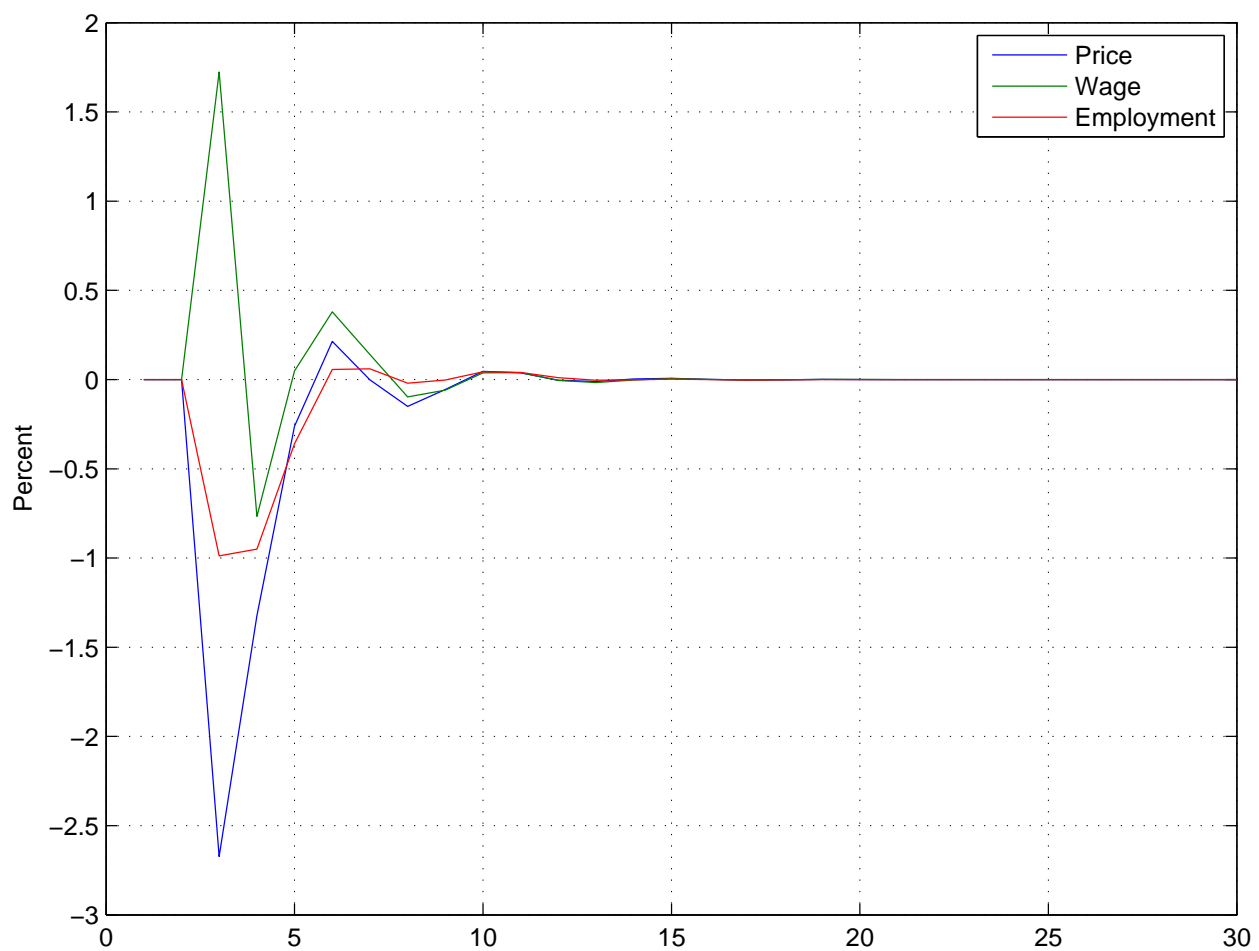
Note: Figure shows the impulse response to a Discount rate shock of the same sign and magnitude as the one we estimate for 2008. The horizontal axis are years after the shock.

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

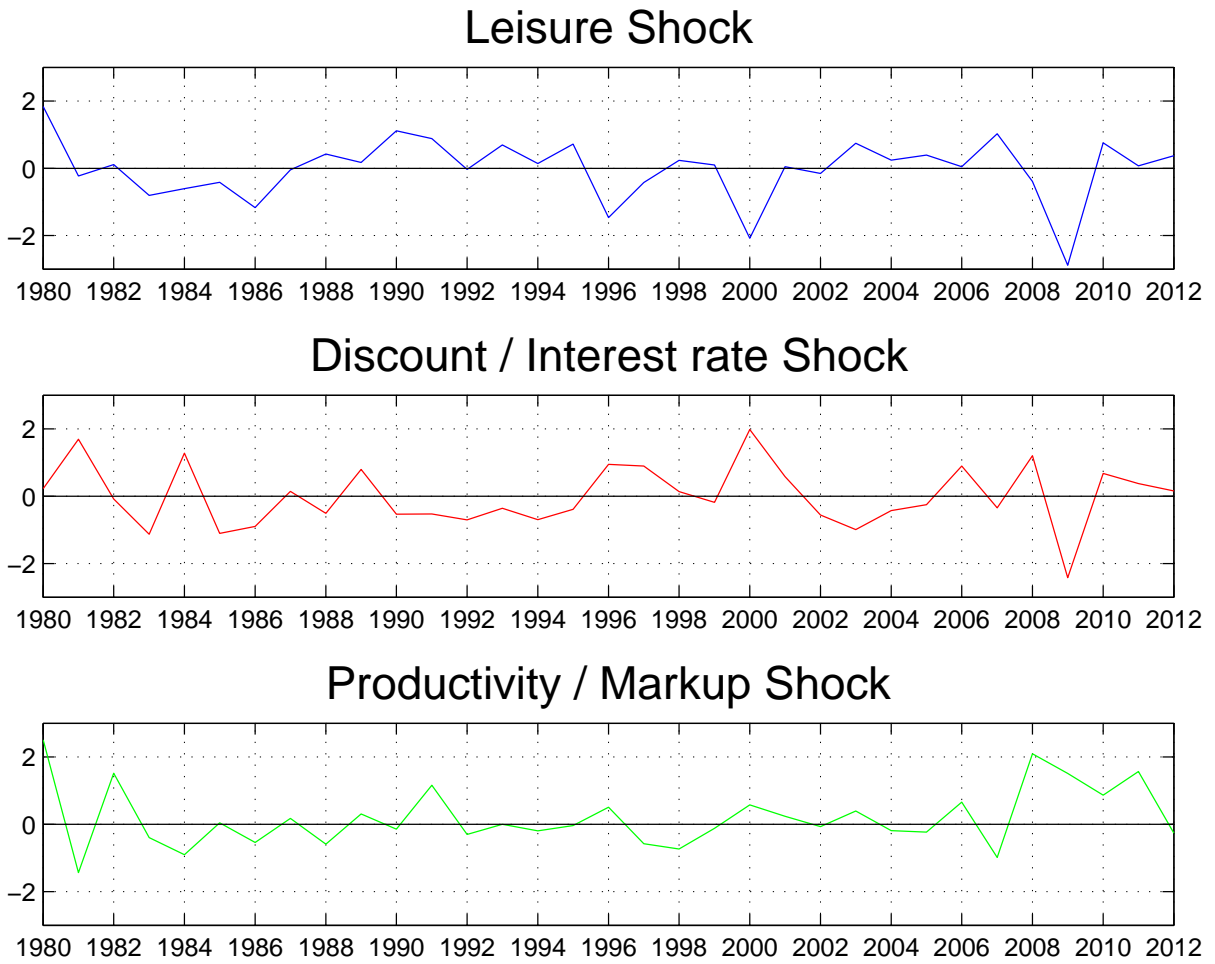
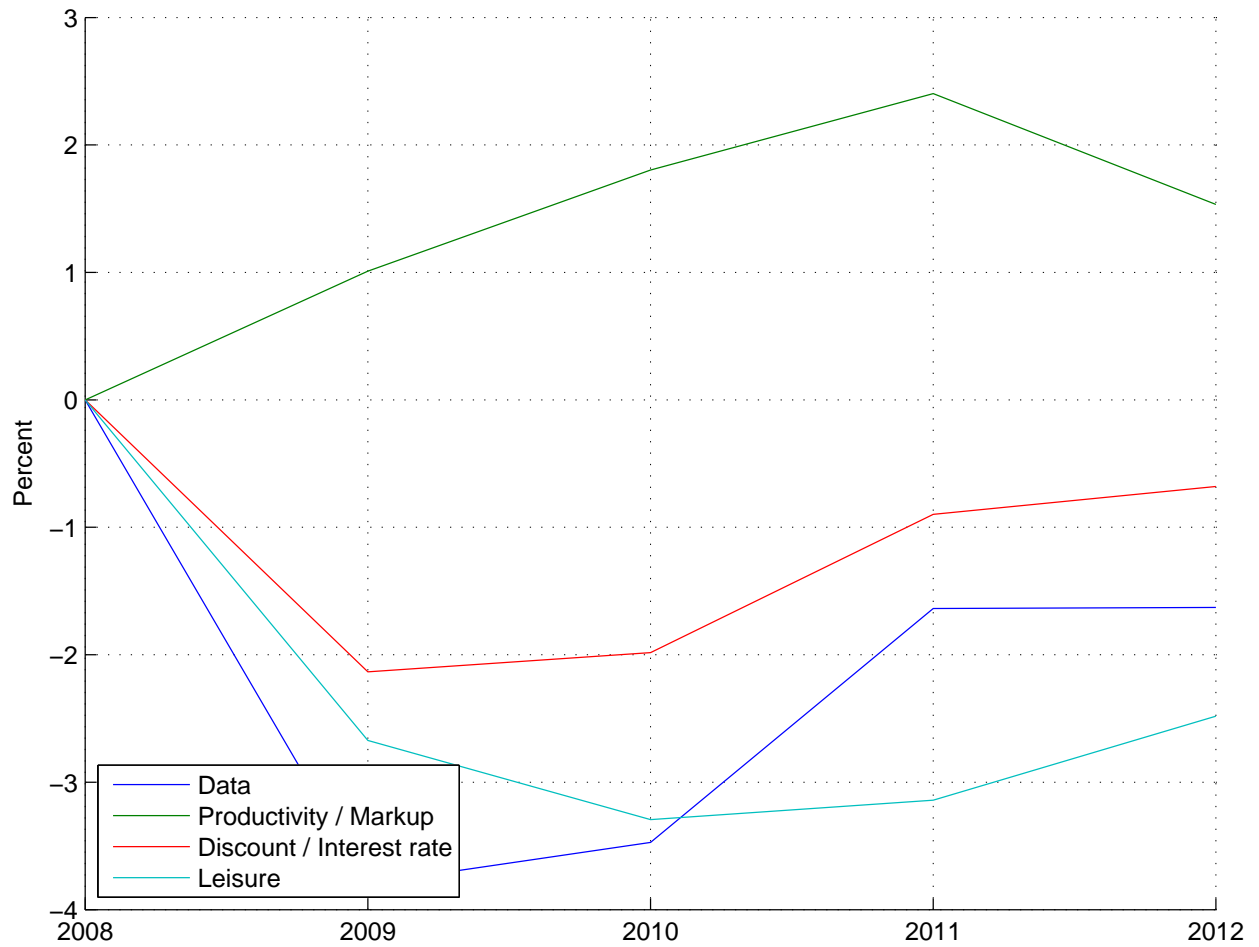


Figure 11: Counterfactual Employment Response



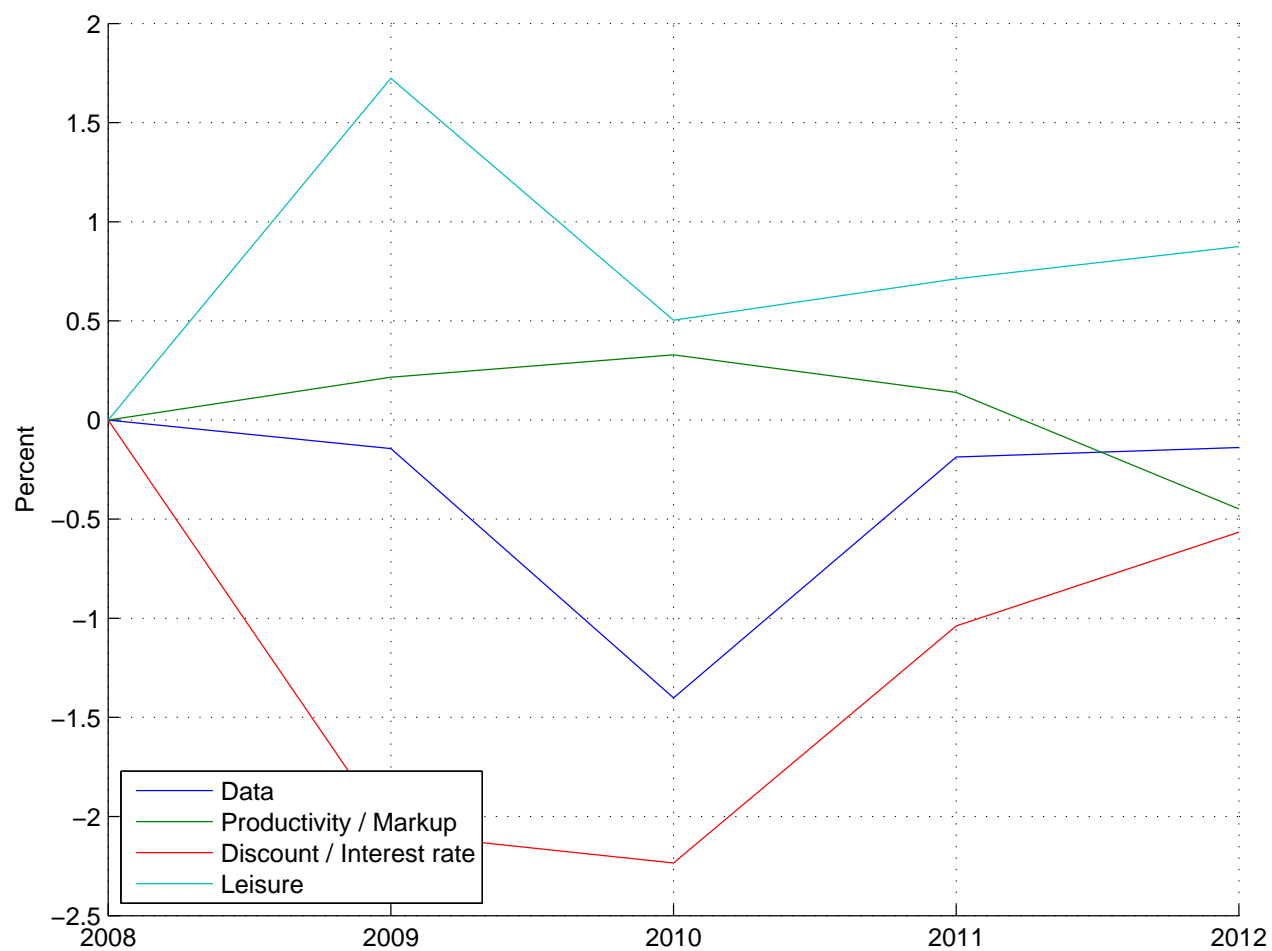
Note: Figure shows the cumulative response of Employment when we feed the VAR with the sequence of shocks between 2008 and 2012; one at a time.

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: Calibration

	Value	Target
α	0.62	Aggregate labor share in the non-tradable sector in 2006
θ	0.55	Aggregate labor share in the tradable sector in 2006
β	0.16	Aggregate labor share in 2006, $\alpha + \theta\beta=0.61$
σ	2	Intertemporal elasticity of substitution
ϕ	2	Frisch elasticity. Own estimate.
ϑ_p	1.5	Taylor Rule from Galí (2011)
ϑ_y	0.5	Taylor Rule from Galí (2011)
Φ_0	0.01	Trade balance-output ratio volatility = 2 from Mendoza (1991).
R	0.03	Real interest rate
X	0.17	Median tradable output share in the US 2006.
B	2.25	Median net worth to output ratio in the Us 2006.
ρ_γ	0.9	Persistence of discount rate shock
ρ_z	0.9	Persistence of productivity shock.
ρ_ϵ	0.9	Persistence of leisure shock

Table 4: Aggregate v. Regional Employment Impact Elasticities

β		0.16				0	0.16
ϑ_y		0.5				0.5	0
(λ, ϕ)		(0.7, 2)	(0.7, 1)	(0.5, 2)	(0.5, 1)	(0.7, 2)	(0.7, 2)
Aggregate	γ	0.77	0.53	1.29	0.98	0.74	1.1
	z	0.91	0.57	0.67	0.4	0.83	0.76
	ϵ	-1.13	-0.72	-1.03	-0.67	-1.03	-1.07
Regional	γ	1.76	1.3	2.95	2.33	2.86	1.76
	z^y	0.54	0.41	-0.01	0.09	0.23	0.54
	z^x	0.13	0	0.31	0.14	0.1	0.13
	ϵ	-0.71	-0.54	-0.28	-0.32	-0.5	-0.71

Note: The table summarizes the response of employment on impact to each of the shocks. The units are percentage deviations from the steady state, in the case of aggregate employment, and percentage deviations from the aggregate in the case of regional employment.

Table 5: Aggregate Elasticities ($\lambda = 0.7, \phi = 2$)

		π^p	π^w	w^r	n
Short	γ	164	142	-22	77
	z	-138	-65	74	92
	ϵ	-7	26	33	-113
Long	γ	101	103	-16	56
	z	-41	-47	54	67
	ϵ	21	19	24	-82

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

Table 6: Regional Elasticities ($\lambda = 0.7, \phi = 2$)

		π^p	π^w	w^r	n
Short	γ	254	240	-14	176
	z^y	-186	-111	75	54
	z^x	36	38	1	13
	ϵ	61	87	27	-71
Long	γ	-28	-30	19	12
	z^y	17	13	40	92
	z^x	-2	-2	6	-7
	ϵ	-6	-10	31	-93

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

Table 7: Estimates of λ and $\frac{\lambda}{\phi}$ using Cross- Region Data

	Specification							
	OLS				IV			
	2007-2011				2007-2009		2007-2011	2007-2009
(1)	(2)	(3)	(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 8: Discount/Interest rate (γ) and Productivity/Markup (z) shocks' contribution to aggregate employment change

λ	2008 to 2009						2008 to 2012				
	ϕ						ϕ				
		0.5	1	2	3	4	0.5	1	2	3	4
0.1	γ	104	69	110	104	105	73	-18	64	88	86
	z	-5	33	-8	-3	-4	16	112	33	9	12
0.3	γ	59	71	42	106	78	-4	-10	-21	70	-10
	z	-1	16	57	-5	23	73	101	120	30	110
0.5	γ	19	51	35	13	101	-7	-5	-16	-5	61
	z	7	13	52	81	-4	71	98	119	110	44
0.7	γ	0	13	41	33	11	20	-4	-9	-12	-2
	z	9	32	34	53	79	44	99	118	123	112
0.9	γ	*	1	1	1	4	*	16	18	17	8
	z	*	34	66	78	80	*	82	95	97	107

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

Table 9: Discount rate (γ) and Productivity/Markup (z) shocks' contribution to change in regional variables

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

Note: The table shows the contribution of each regional shock to each regional wage, price and non-tradable employment change between 2007-2010. As defined in Section 7.3, for each variable

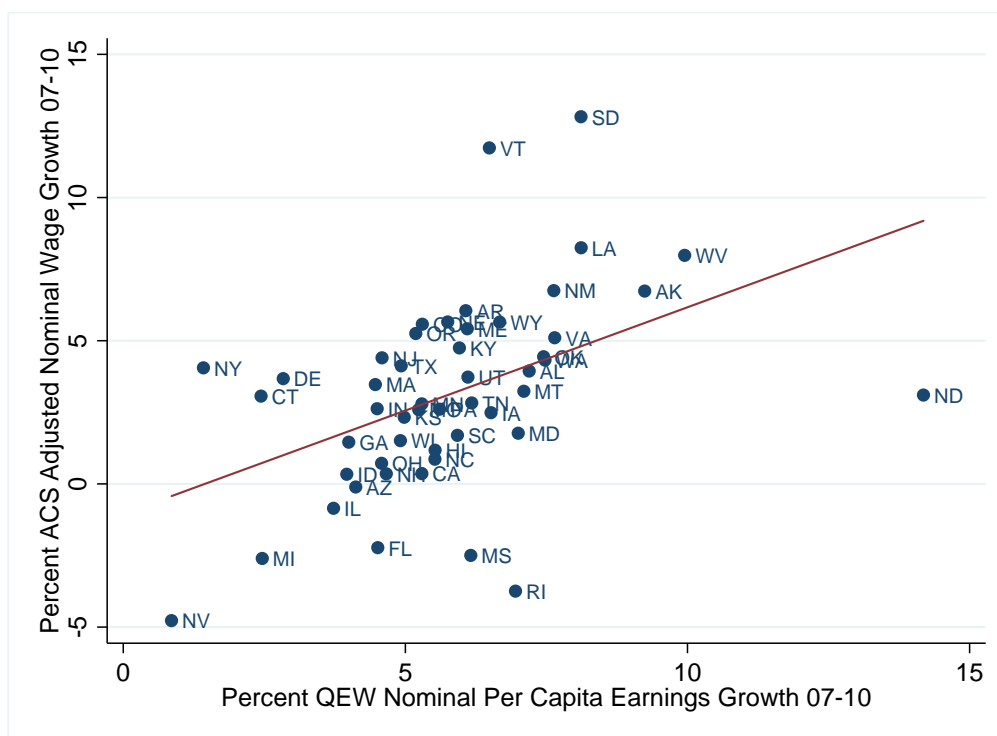
$$y \text{ and shock } j \text{ we compute } \tilde{\zeta}_{y^j}^j = \frac{1}{\sum_k \omega_k (\Delta^j \tilde{y}_k - \Delta \tilde{y}_k)^2}.$$

Table 10: Regional counterfactual statistics

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

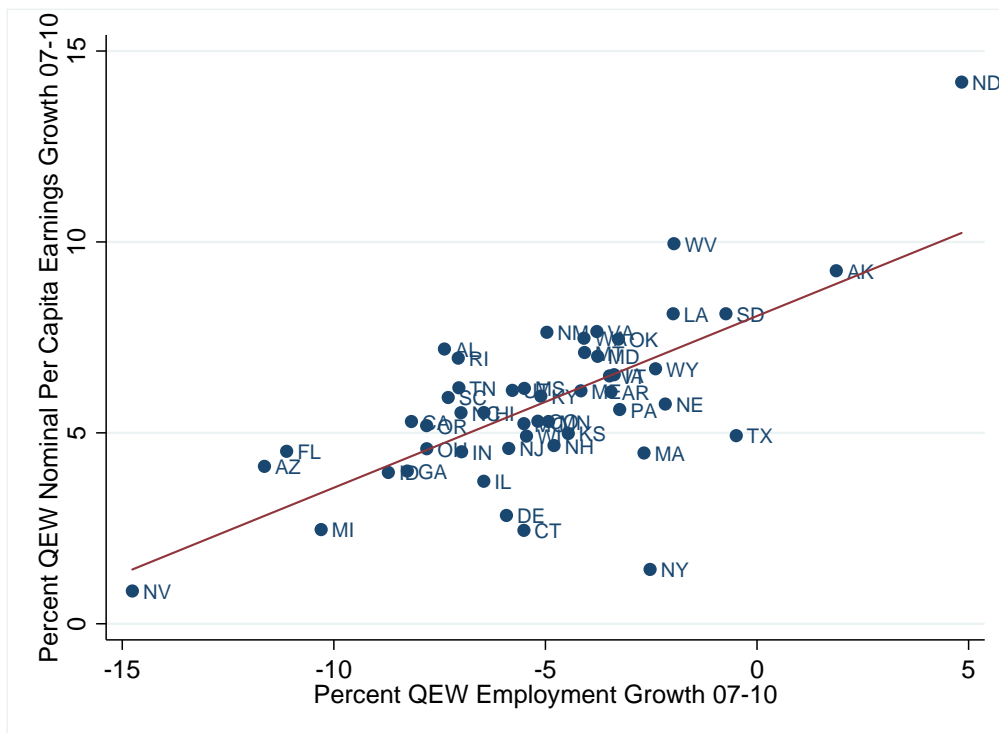
Note: The first three lines in the table show the counterfactual variance across states relative to the actual variance of the total percent change in each variable between 2007-2010. The last three lines show the population weighted OLS coefficient corresponding to each variable pair. For example, β_{p,n^y}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

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



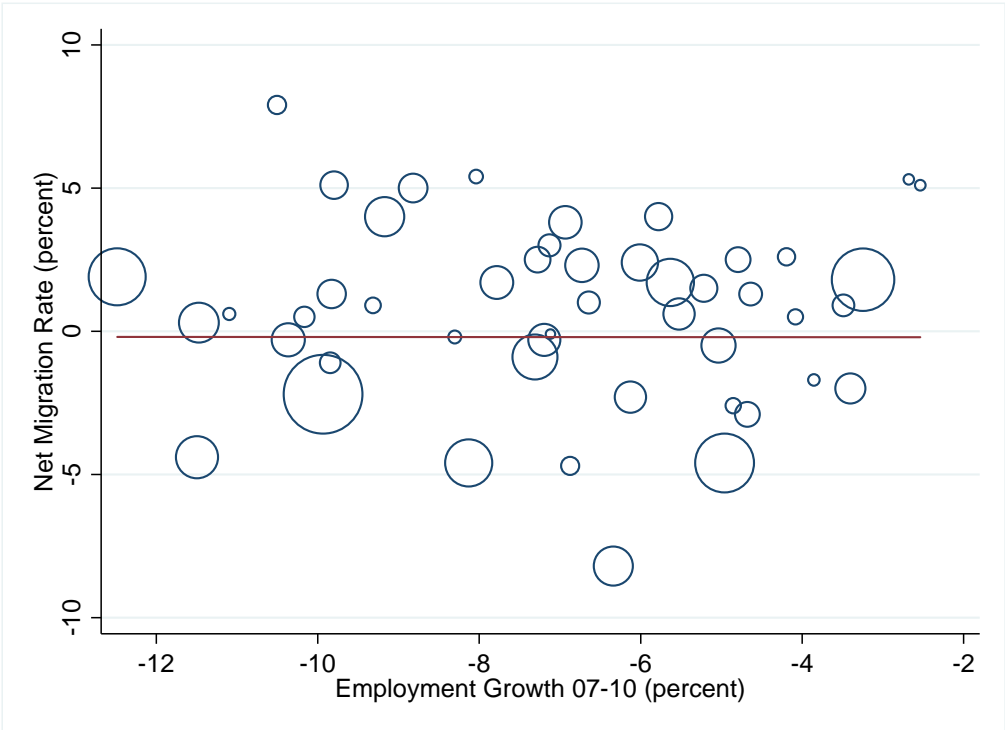
Note: Figure shows the simple scatter plot of the growth in our ACS adjusted nominal wage index from 2007-2010 against the nominal per worker earnings growth from 2007-2010 from the BLS's QEW. Each observation is a U.S. state. The slope of the regression line is 0.72 with a standard error of 0.20.

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



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

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



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