Demand Stimulus, Inflation and Marginal Costs: Empirical Evidence*

Marcus Hagedorn†
University of Oslo and CEPR

Jessie Handbury‡
University of Pennsylvania

Iourii Manovskii§
University of Pennsylvania and NBER

Abstract

The relationship between government spending, marginal costs and inflation is at the core of New Keynesian theories of price setting. In particular, the impact of a demand stimulus on marginal costs and inflation determines the size of the fiscal multiplier. Yet, despite the fundamental importance of this relationship to theory, documenting it in the data has proved elusive.

In this paper we estimate the effect of a fiscal expansion on inflation using the Nielsen Retail Scanner Data. To identify this response of inflation to a demand stimulus (an increase in the generosity of unemployment insurance) we exploit a policy discontinuity at U.S. state borders. We find that inflation (and also spending) responds with a delay in a hump-shaped manner.

To assess the impact on marginal costs we explore whether this response is consistent with two leading models of price setting, the sticky price model and the sticky information model, which are known to have different implications for the size of the fiscal multiplier. We find that the sticky information model is consistent with our empirical findings whereas the sticky price model is not.

Keywords: Aggregate Demand, Inflation, Marginal Costs, Sticky Prices, Fiscal Stimulus, Unemployment Insurance, Information Frictions

*This Draft: March 30, 2015. We thank Arash Nikoei, Eric Sims and Nicolas Trachter who formally discussed the paper as well as seminar participants at the Federal Reserve Board, IIES (Stockholm), NYU, UCLA, the St. Louis Fed Annual Fall Conference, the "Aggregate Demand, the Labor Market and Macroeconomic Policy" conference at Cambridge University, the Riksbank-EABCN conference on "Inequality and Macroeconomics", and the S&M workshop at the Philadelphia Fed for helpful comments. Support from the National Science Foundation Grant No. SES-0922406 is gratefully acknowledged.

†University of Oslo, Department of Economics, Box 1095 Blindern, 0317 Oslo, Norway. Email: marcus.hagedorn07@gmail.com

‡The Wharton School, University of Pennsylvania, Real Estate Department, 1463 Steinberg Hall-Dietrich Hall, 3620 Locust Walk, Philadelphia, PA 19104-6302, Email: handbury@wharton.upenn.edu

§University of Pennsylvania, Department of Economics, 160 McNeil Building, 3718 Locust Walk, Philadelphia PA 19104. Email: manovski@econ.upenn.edu.

1 Introduction

The objective of this paper is to document empirically the relationship between a fiscal stimulus, marginal costs, and inflation. This relationship is at the core of the New Keynesian theory which is widely used to assess various macroeconomic stabilization policies. In particular, this relationship is crucial in determining the size of the fiscal multiplier implied by this model. For example, using the standard sticky-price version of the theory, Christiano, Eichenbaum, and Rebelo (2011) and Farhi and Werning (2013) find the multiplier to be large whenever nominal interest rates are not responsive to inflation. The model mechanism for this result is simple. Government spending increases marginal costs which leads firms to increase prices, resulting in higher inflation. With fixed nominal interest rates - with the leading example a liquidity trap with a nominal interest rate of zero - the increase in inflation leads to a one-for-one reduction in the real interest rate, boosting current private spending. This increase in demand leads to further increases in inflation, and so on, explaining the quantitatively important feedback loop. In contrast, the dynamic response of marginal costs and inflation to a fiscal stimulus in the version of the theory based on sticky information developed by Mankiw and Reis (2002) implies that the fiscal multiplier is below one.

While it has been long recognized that a short run link between real activity - including government spending, marginal costs, and inflation is at the heart of the theory, e.g., Gali and Gertler (1999), documenting this relationship in the data has proved elusive. The empirical literature focused primarily on describing the relationship between inflation and measures of economic slack such as the unemployment rate. This evidence is, however, difficult to interpret because while the unemployment rate increases and inflation decreases in response to a negative demand shock, the unemployment rate and the inflation rate both increase in response to a negative technology shock. In addition, measurement of economic slack or marginal costs in the data is highly model dependent. Another voluminous and contentious branch of the literature that attempted to measure fiscal multipliers using Structural VARs or using the regional variation in government spending (see Ramey (2011) for a survey), has not assessed the link between government spending and inflation. This seems surprising as the effect on inflation is not only the key channel but the path of consumption is independent of the path of government spending given the inflation path (and a path for the nominal interest rate), as pointed out by Farhi and Werning (2013).

In this paper we attempt to fill this gap. Given that it appears exceedingly challenging to identify the impact of fiscal spending on marginal costs and inflation using aggregate

data, we propose to exploit a novel source of substantial exogenous time-series and cross-sectional variation in government spending across U.S. states and counties. Exploiting the cross-sectional variation is attractive for a number of reasons. The vast amount of exogenous variation leads to fairly large sample sizes which allows us to obtain precise estimates. As these locations constitute a monetary union, it is easy to control for the effects of the common monetary policy. Finally, the variation in government spending that we consider represents pure transfers from the rest of the union for which, as we show below, the theory provides unambiguous predictions regarding the response of local inflation and marginal costs.

To exploit this variation we require data on the evolution of prices at the same level of geographic aggregation. To construct local price indexes we use data from Nielsen's Retail Scanner Database. This is an enormous dataset that provides over the 2006-2012 period weekly records of quantities and prices of nearly all goods sold at approximately 40 thousand grocery, pharmacy, and mass merchandising stores spread out across the U.S. All goods are precisely identified using their UPC codes. Beraja, Hurst, and Ospina (2015) also construct local indexes using these data and show that aggregating these price data tracks the CPI indexes provided by the BLS for the corresponding product categories nearly identically.

To identify exogenous variation in government spending across locations we consider the transfers from the federal government through the unemployment insurance system. The period for which we have price and quantity data from the Nielsen's Retail Scanner Database largely coincides with the Great Recession and its aftermath. During this period there was an unprecedented expansion of the unemployment insurance payments with federally financed unemployment benefit duration rising by up to 73 weeks, but with substantial variation across U.S. states. We document below that when a state increases available benefit duration, the amount of transfers it receives from the federal government to pay the newly-eligible unemployed jumps on impact and remains stable at the new higher level for a long time. The amount of transfers through this policy is also large. For example, in 2011, these transfers accounted for 1.5% of average personal income with a wide heterogeneity across states.

The transfers are obviously endogenous to the economic conditions in the state. To obtain the necessary exogenous variation, our empirical strategy exploits a policy discontinuity at state borders. A key feature of the U.S. unemployment insurance system is that policies are determined at the state level and apply to all locations within a state. This allows us to identify the effects of exogenous variation in transfers by comparing the evolution of sales and

inflation in counties that border each other but belong to different states.[1] Locations separated by a state border share a common monetary policy and are expected to have similar labor and product markets due to the same geography, climate, access to transportation, agglomeration benefits, access to specialized labor and supplies, etc. Indeed, we provide direct evidence that economic shocks do not stop at the state border but evolve smoothly across borders. The key feature that sets these locations apart is the difference in policies on the two sides of the border. This policy discontinuity allows to identify its pricing implications. A fundamentally similar identification strategy was used, among others, by Holmes (1998) to identify the impact of right-to-work laws on location of manufacturing industry and by Dube, Lester, and Reich (2010) to identify the effect of minimum wage laws on earnings and employment of low-wage workers. We explicitly control for the effects of other policy changes at the state level (that could be correlated with the expansion of unemployment benefit durations) to ensure that our estimates isolate the effects of unemployment benefit extensions.

Armed with the required data and methodology, we are able to document empirically the effect of exogenous variation in fiscal transfers on retail sales and prices. We find that these transfers represent a substantial demand stimulus. In particular, we show that, controlling for state unemployment, sales increase significantly in a hump-shaped fashion with a peak response of between 7 and 8 percent occurring one to two years after the increase in benefit duration.[2] The impact of benefit extensions on inflation is also hump-shaped and gradual with a peak response of 1.2% reached 7 quarters after the change in benefit duration.

This descriptive evidence appears consistent with an effective expansionary fiscal policy at the local level. However, to understand the implication of these findings for the effectiveness of a demand stimulus a general equilibrium model is required. We consider the benchmark New Keynesian model with sticky prices as well as the Sticky Information Model. For the model to be useful for a policy analysis at the aggregate level, it has to be consistent with the evidence at the disaggregated level.

The theoretical foundations underlying our analysis of the effects of a fiscal stimulus in a currency union of small open economies (counties in our analysis) using a sticky-price New Keynesian model were developed in Farhi and Werning (2013). They show that in this model an expansionary fiscal policy (in particular, a fiscal transfer to a member of the currency union) raises firms' marginal costs which induces firms to increase prices. While the

[1] A Map of U.S. state and county borders can be found in Appendix Figure A-1.

[2] Controlling for state unemployment eliminates the potential positive effect of unemployment benefit extensions on unemployment, the magnitude of which remains a subject of a lively debate in the literature.

relevant deviations of marginal costs from the frictionless benchmark in response to fiscal stimulus are difficult to measure directly, the theory of price setting underlying the New Keynesian model with sticky prices allows us to do so using the descriptive estimates of the dynamic response of inflation. Specifically, the New Keynesian Phillips Curve implies that the deviations of marginal costs are equal to the quasi-differenced inflation (inflation in period t minus discounted inflation in period $t+1$). Comparing the response of quasi-differenced inflation across border counties to the extensions of UI benefits, we find a zero impact both economically and statistically at all leads and lags. Thus, viewed through the lens of the sticky-price New Keynesian model, the dynamics of price response that we document implies that fiscal stimulus has no effect on marginal costs at all leads and lags. As the inflation predicted by this model is equal to the present discounted value of future deviations of marginal costs, the finding of no effect of fiscal stimulus on marginal costs implies that there is also no impact on inflation in the model. Thus, the model fails to replicate the dynamic path of inflation that we measure in the data. Conducting a robustness analysis we show that these findings are insensitive to allowing for price indexation and trend inflation as well as to different discounting of future marginal costs in firms' pricing decisions. Moreover, the results apply to goods with any observed degree of price stickiness and are not affected by controlling for store chain-wide pricing policies. As our measurement approach exploited only the optimality of firms' decisions given the theory of pricing frictions, the evidence calls this theory into question.

An alternative theory of price setting based on the assumption that information disseminates slowly among the price-setting firms, was developed by Mankiw and Reis (2002). Contrasting this theory with the descriptive evidence that we documented, we find that it can generate the dynamic path of marginal costs (and inflation) consistent with the data. This is possible, however, only if the information frictions are relatively small. Moreover, regardless of the size of these frictions, the fiscal multipliers generated by this model are strictly below one.

The remainder of the paper is organized as follows. In Section 2 we describe the data and variable construction. In Section 3 we show that unemployment benefit extensions represent significant demand stimulus and document the effects of this stimulus on sales and inflation. In Section 4 we measure the effects of unemployment benefit extensions on the marginal costs and inflation rate using the New Keynesian model with sticky prices and find them to be virtually zero, contradicting our direct observations in the data. The corresponding analysis of the Sticky Information model is in Section 5. Section 6 concludes.

2 Data

In this section we describe the data sources we use to in our empirical analysis. In particular we use the Nielsen Retail Scanner Data to construct an inflation index which matches the definitions in the theoretical models of price setting. In these models firms face some frictions when changing prices, which reflect changes in firms' marginal costs. This implies for our analysis that the inflation index should not be contaminated by changing consumer tastes or shop-switching behavior of consumers. Instead we have to measure how a store prices a specific good and how this price changes over time due to changes in costs or demand. We therefore consider how each store in the Nielsen Retail Scanner Data changes the price of each good sold. Proceeding this way ensures that our inflation index does not change for example if consumers switch to cheaper stores in a recession, what it would if we computed how the price of the average consumer's basket changes.

2.1 Price Data

Inflation is calculated using the Nielsen Retail Scanner Data, which we obtained through the Kilts Center for Marketing at the University of Chicago Booth School of Business.[3] The database includes weekly sales records for over 35,000 participating grocery, drug, mass merchandiser, and other stores in 52 major markets across the US, as well as information on the county where each store is located. Stores report the value and quantity sold of each product sold at a weekly frequency. This high level of detail enables us to measure inflation at the county level by comparing the prices charged for identical products sold in the same store in different time periods, whereby controlling for factors such as product quality and store amenities.

The Kilts-Nielsen data uniquely identifies each product using a Universal Product Code (UPC) as well as a UPC version control number that reflects any changes to the observable characteristics of the product, such as package size or brand name. We will follow the same convention, considering two products to be identical only if they have the same UPC and UPC version control number. We will henceforth refer to this pair of identifiers as a product or UPC and denote it using u.

Inflation Index Definitions We measure two different types of inflation. One is intended to represent inflation in posted prices, the other one is intended to measure consumer inflation.

[3] http://research.chicagobooth.edu/nielsen/datasets/scanner-data.aspx

All of these inflation indexes are geometric averages of price changes and each is differentiated by the weights that are applied to each price change in the geometric average.

We study these indexes at a monthly and quarterly frequency, so we first aggregate the Nielsen sales data across weeks, defining the sales-weighted average price charged by store s for UPC (or product) u in time period t as

$$p_{u,s,t} = \frac{v_{u,s,t}}{q_{u,s,t}}, \tag{1}$$

where $v_{u,s,t}$ and $q_{u,s,t}$ respectively denote the total value and quantity purchased of UPC u in store s in time period t.

The index for a county c in month t is a weighted geometric average of these price changes across all of the products and county-c stores for which we observe positive sales in both period t and period $t - dt$:

$$1 + \pi^k_{c,t,dt} = \prod_{u \in U_{s,[t-dt,t]}} \prod_{s \in S_{c,[t-dt,t]}} \left(\frac{p_{u,s,t}}{p_{u,s,t-dt}} \right)^{w^k_{u,s,t,dt}}, \tag{2}$$

where $U_{s,[t-dt,t]}$ is the set of UPCs sold in store s in months $t - dt$ and t; $S_{c,[t-dt,t]}$ is the set of county-c stores whose sales are observed in the sample in months $t - dt$ and t; and $w^k_{u,s,t}$ is a weight unique to the index being calculated. The posted price inflation indexes are calculated using equal weights, which are formally defined as follows:

$$w^{Equal}_{u,c,t,dt} = \frac{1}{\sum_{u' \in U_{s,[t-dt,t]}} \sum_{s' \in S_{c,[t-dt,t]}} \mathbb{I}[v_{u',s',t-dt} > 0] \mathbb{I}[v_{u',s',t} > 0]}. \tag{3}$$

We measure consumer inflation using a Tornqvist index:

$$w^T_{u,c,t,dt} = \frac{1}{2} \left(\frac{v_{u,s,t-dt}}{\sum_{u' \in U_{s,[t-dt,t]}} \sum_{s' \in S_{c,[t-dt,t]}} v_{u',s',t-dt}} + \frac{v_{u,s,t}}{\sum_{u' \in U_{s,[t-dt,t]}} \sum_{s' \in S_{c,[t-dt,t]}} v_{u',s',t}} \right). \tag{4}$$

In our base analysis, we calculate both indexes at a monthly frequency, where t denotes a month, with a "base-gap," dt, of 1 and 12 months.

Implementing the Inflation Index Calculation The size of the dataset implies that we must construct each of the indexes above by first calculating "sub-indexes" for each of 800 or so product categories. We will follow Nielsen's naming convention and refer to these categories as modules. Let U_m denote the set of UPCs categorized in module m. We can define a module-specific inflation index for a county c in period t relative to period $t - dt$ as follows:

$$1 + \pi^k_{m,c,t,dt} = \prod_{u \in U_{m,s,[t-dt,t]}} \prod_{s \in S_{c,[t-dt,t]}} \left(\frac{p_{u,s,t}}{p_{u,s,t-dt}} \right)^{w^k_{u,m,s,t,dt}}, \tag{5}$$

where $U_{m,s,[t-dt,t]} = U_{s,[t-dt,t]} \bigcap U_m$ is the set of UPCs in module m that are sold in store s in period $t-dt$ and t; $S_{c,[t-dt,t]}$ is the set of county-c stores whose sales are observed in the sample in periods $t-dt$ and t; and $w_{u,m,s,t,dt}$ is a module-specific and time-specific weight for a UPC relative to the other UPCs in that module given by equations 3 through 4, where the denominator is summed only over module m UPCs rather than all UPCs (i.e., replace $U_{s,[t-dt,t]}$ with $U_{m,s,[t-dt,t]}$). For example, the within-module Tornqvist weight becomes:

$$w^T_{u,m,s,t,dt} = \frac{1}{2}\left(\frac{v_{u,s,t-dt}}{\sum_{u' \in U_{m,s,[t-dt,t]}} \sum_{s' \in S_{c,[t-dt,t]}} v_{u',s',t-dt}} + \frac{v_{u,s,t}}{\sum_{u' \in U_{m,s,[t-dt,t]}} \sum_{s' \in S_{c,[t-dt,t]}} v_{u',s',t}}\right) \tag{6}$$

We aggregate across modules $m = 1, ..., M$, using a weighted geometric average:

$$1 + \tilde{\pi}^k_{c,t,dt} = \prod_{m=1}^{M} \left(1 + \pi^k_{m,c,t,dt}\right)^{w^k_{m,c,t}}. \tag{7}$$

with module weights, $w^k_{m,c,t,dt}$ given by the module-level equivalents for the weights given by equation 3. For example, the module-level Tornqvist weight is:

$$w^k_{m,c,t,dt} = \frac{1}{2}\left(\frac{\sum_{u' \in U_{m,s,[t-dt,t]}} \sum_{s' \in S_{c,[t-dt,t]}} v_{u',s',t-dt}}{\sum_{m \in M} \sum_{u' \in U_{m,s,[t-dt,t]}} \sum_{s' \in S_{c,[t-dt,t]}} v_{u',s',t-dt}} + \frac{\sum_{u' \in U_{m,s,[t-dt,t]}} \sum_{s' \in S_{c,[t-dt,t]}} v_{u',s',t}}{\sum_{m \in M} \sum_{u' \in U_{m,s,[t-dt,t]}} \sum_{s' \in S_{c,[t-dt,t]}} v_{u',s',t}}\right). \tag{8}$$

The two-tiered Tornqvist index produces inflation rates that are almost identical to those of the standard Tornqvist defined in equations 2 and 4 above.

County Price Levels County-level indexes are calculated using spatial analogs of the inflation indexes outlined above. Where the inflation indexes compare the prices charged locally in the current period, t, relative to the prices charged locally in a past period $t-dt$, the county-level price level indexes compare the prices charged locally in the current period to those charged for the same products nationally. The price-level index for a county c in month t, for example, is a weighted geometric average of prices charged across all of the products sold in county-c stores relative to the sales-weighted average price charged for those same products nationally in the same month:

$$P^k_{c,t,dt} = \prod_{u \in U_{s,[t-dt,t]}} \prod_{s \in S_{c,[t-dt,t]}} \left(\frac{p_{u,s,t}}{p_{u,t}}\right)^{w^k_{u,s,t}}, \tag{9}$$

CPI Category:	Food	All items
New York	.974	.96
Boston	.992	.923
Chicago	.983	.914
Washington	.988	.97

Table 1: Correlations between Kilts-Nielsen CPI for Border Counties in MSAs and Respective MSA-Level BLS Indexes

where $p_{u,t} = \frac{\sum_{c \in C} \sum_{s' \in S_{c,t}} v_{u,s',t} p_{u,s',t}}{\sum_{c \in C} \sum_{s' \in S_{c,t}} v_{u,s',t}}$ is the sales-weighted national average price of UPC u in period t and the weights $w^k_{u,s,t}$ are defined as above with the national market share taking the place of the local lagged market share in the Tornqvist index:

$$w^T_{u,c,t} = \frac{1}{2} \left(\frac{\sum_{c \in C} \sum_{s' \in S_{c,t}} v_{u,s,t}}{\sum_{u' \in U_{s,t}} \sum_{c \in C} \sum_{s' \in S_{c,t}} v_{u',s',t}} + \frac{v_{u,s,t}}{\sum_{u' \in U_{s,t}} \sum_{s' \in S_{c,t}} v_{u',s',t}} \right). \tag{10}$$

Measuring and using retail prices as a combination of local marginal costs and goods imported by the county is based on results by Burstein, Eichenbaum, and Rebelo (2005). These authors find that "retail prices of tradable goods comprise two important nontradable components: distribution costs and local goods. We define distribution costs as wholesale and retail services, marketing and advertising, and local transportation services. We define local goods as goods that are produced solely for the domestic market. Thus the retail prices of tradable goods do not accurately reflect the prices of pure traded goods at the dock, that is, the prices of goods that are actually traded, exclusive of distribution costs."

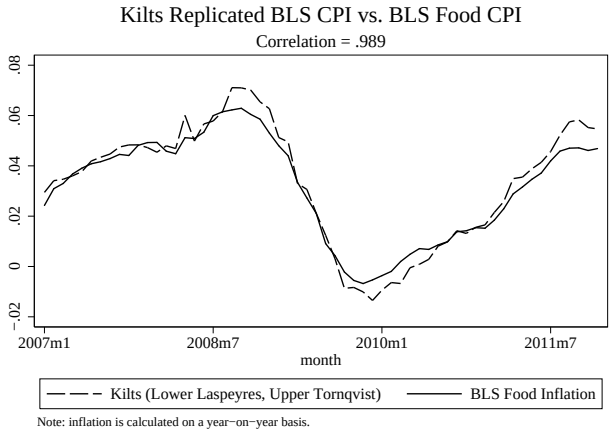


Figure 1: Comparison Kilts Data and BLS

Rough Comparison with BLS Data

2.2 Unemployment Benefit Duration Data

For our empirical analysis we have data available from 2006-2012. To identify the role of unemployment benefit extensions on labor market outcomes, we focus our analysis on a sample of county pairs that are in different states and share a border.[4] There are 1,107 such pairs for which we have complete data.

Data on unemployment benefit durations in each state is based on trigger reports provided by the Department of Labor. These reports contain detailed information for each of the states regarding the eligibility and adoption of the two unemployment insurance programs over our primary sample period: Extended Benefits program (EB) and Emergency Unemployment Compensation (EUC08).[5]

The EB program allows for 13 or 20 weeks of extra benefits in states with elevated unemployment rates. The EB program is a joint state and federal program. The federal government pays for half of the cost, and determines a set of "triggers" related to the insured and total unemployment state rates that the states can adopt to qualify for extended benefits. At the onset of the recession, many states chose to opt out of the program or only adopt high triggers. The American Recovery and Reinvestment Act of 2009 turned this into a federally funded program (with 100% Federal funding currently scheduled to expire on December 31, 2013). Following this, many states joined the program and several states adopted lower triggers to qualify for the program.

The EUC08 program enacted in June 2008, on the other hand, has been a federal program since its onset. The program started by allowing for an extra 13 weeks of benefits to all states and was gradually expanded to have 4 tiers, providing potentially 53 weeks of federally financed additional benefits. The availability of each tier is dependent on state unemployment rates. The trigger reports contain the specifics of when each state was eligible and activated the EB program and different tiers of the EUC08 program. We have constructed the data through December 2012.

There is a substantial heterogeneity in the actual unemployment benefit durations across time and across the U.S. states. Appendix Figure A-2 presents some snapshots that illustrate the extent of this variation. Among 1,107 border county pairs used in our analysis, 1,079 have different benefits for at least one quarter. The median county pair has different benefit durations for 11 quarters during 2008-2012. The difference in available benefit duration

[4]Data on county pairs are provided by Arindrajit Dube and were used in Dube, Lester, and Reich (2010).

[5]See http://ows.doleta.gov/unemploy/trigger/ for trigger reports on the EB program and http://ows.doleta.gov/unemploy/euc_trigger/ for reports on the EUC08 program.

within a county-pair ranges from 0 to 17 quarters.

Some of the data series used in the analysis are available at a monthly frequency while others are quarterly. Therefore, we aggregate all monthly data to obtain quarterly frequency. Logs are taken after aggregation.

3 UI Benefit Durations, Demand, and Inflation

In this Section we document how the inflation rate and consumer spending respond to fiscal policy. The fiscal policy measure, we use, is an increase in the duration of UI benefits. Since UI benefits are a state policy, we have substantial variation of UI benefit duration not only over time within a state but also across states. A key feature of the U.S. unemployment insurance system is that unemployment insurance policies are determined at the state level and apply to all locations within a state. Our empirical strategy exploits this policy discontinuity at state borders to measure the effects of an increase in UI benefit duration on inflation and spending. One cannot, however, infer the effects of benefit extensions by simply relating benefit duration to inflation or spending in a panel of states because of a potential policy endogeneity: it might be the states that have a large increase in unemployment, then not only experience an increase in benefit eligibility but also a decrease in spending or inflation. Thus it might be that some negative shock drives unemployment up, benefit duration up, and inflation and spending down leading to a biased estimate of duration on inflation and spending. In Hagedorn, Karahan, Mitman, and Manovskii (2013) we show, building on work of Holmes (1998) and Dube, Lester, and Reich (2010), that the endogeneity problem can be overcome by comparing the evolution of unemployment, inflation or spending in counties that border each other but belong to different states. We apply this identification strategy here as well and measure the differential response of inflation and spending of neighboring counties which belong to different states to changes in benefit duration.

Using benefit duration as a measure of fiscal policy to measure the effects on inflation and spending makes it necessary to document several facts. First, we show that increasing the duration of benefits leads to an increase in transfers from the federal government to the state, where benefits increase. Since these transfers are paid by the federal government and therefore state taxes do not change, an increase in benefits means that the unemployed who receive benefit payments from this state have more resources to spend.

For how long an unemployed receives benefits depends on the state where he worked last, not where he lives. We therefore document next that the majority of workers receive

benefits from the state where they live, that is their last job was in the same state as where they resided (work home bias). Together with the first fact, this implies that an increase in benefit duration increases the purchasing power of the residents of the same state.

Finally, we document that spending is concentrated in the state/county of residence (home bias in spending), implying that the increase in benefit duration leads to an increase in demand in the counties of this state.[6]

3.1 Benefits and Transfers

All payments outside the regular UI system were financed by the federal government and thus lead to transfers T_t^s to state s in every month t. The following regression summarizes the dynamic response of total (extended) UI Transfers at different lags.

$$\Delta \log(T_{t+k}^s) = \varphi_k \Delta \log(b_t^s) + \epsilon_{s,t},$$

where φ_k is the effect of a change in benefits between months t and $t-1$ on the change in transfers between months $t+k$ and $t+k-1$ for $k = 0, 1, \ldots, 15$, that is $\Delta \log(T_{t+k}^s) = \log(T_{t+k}^s) - \log(T_{t+k-1}^s)$ and $\Delta \log(b_t^s) = b_t^s - b_{t-1}^s$.

The result in Figure 2 show that in the month in which the duration is extended, transfers change significantly. After that period changes are insignificant. This finding reflects the fact that benefits are paid retroactively, meaning that if benefits are extended for example from 39 to 53 weeks then those who have been unemployed for longer than 39 weeks at the time of extension all receive their UI checks. We obtain the same results when we use transfer per unemployed instead of transfers as shown in Figure 3.

We also explore the consequences of extending benefits on the accumulated change in transfers, $\Delta_k \log(T_t^s) = \log(T_{t+k}^s) - \log(T_{t-1}^s)$, that is we implement the regression

$$\Delta_k \log(T_t^s) = \varphi_k \Delta \log(b_t^s) + \epsilon_{s,t},$$

where Figure 4 shows φ_k for $k = 1, \ldots, 12$.

We again find the immediate spike in transfers which flattens out after about nine months. From this we conclude that extending the duration of benefits is an increase in transfers from the federal level to the state which is concentrated in the period when the duration increases.

[6] To see that this is important consider the counterfactual case where households located in some county i spend half of their income in county i and the other half in the neighboring county j of a different state. As a result if benefit duration increases for county i but not for county j, we would see an increase in demand for both counties of the same magnitude, that is the difference in demand would not change and we would not expect to see a different response of inflation or spending in the two counties. However, this is not the case. The home bias in spending means that demand increases more in county i than in county j and we can use this to measure the impact on inflation.

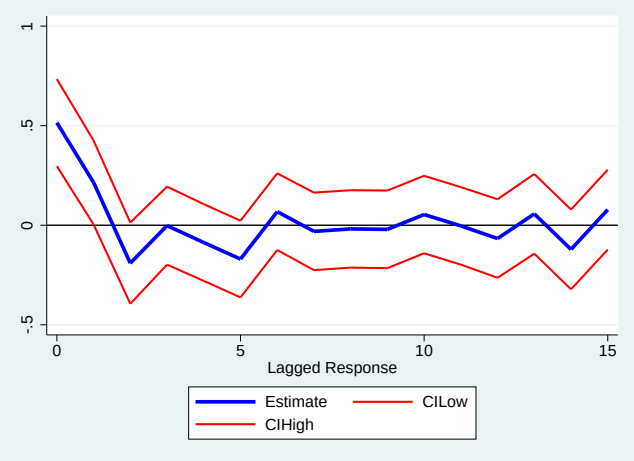


Figure 2: Benefit extensions and Change of Transfers

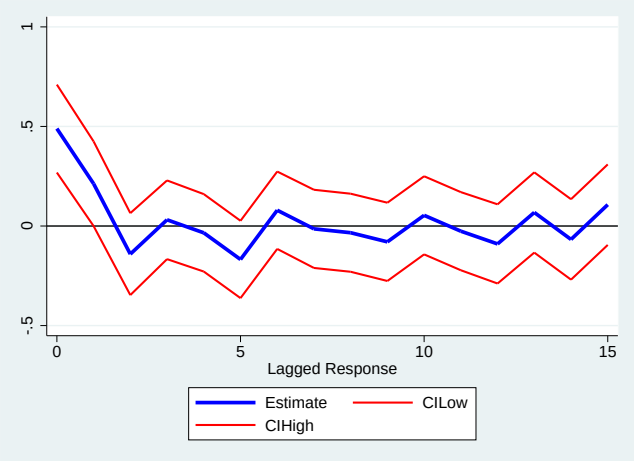


Figure 3: Benefit extensions and Change of Transfers per Unemployed

As explained above, to be able to measure how this increase in transfers transforms into spending and how it affects inflation we have verify that two home biases, the work home bias and home bias in spending. This is indeed what we find in the data. Detailed calculations will be provided in the next draft, but overview of the results is as follows.

First the amount of benefits paid depend on the location of the (last) job and not on where the unemployed lives. Because of this institutional feature we measure how many of those located in county i also work in the same state. We find that this applies to vast majority of all employed. As a result, transfers paid according to the rules of state s accrue predominantly to unemployed located in state i and not to counties in the neighboring states. We have therefore have a work home bias.

We also find a home bias in spending. Using Nielsen Consumer Panel Data we find that retail spending of residents of county i is very heavily concentrated on the stores located in

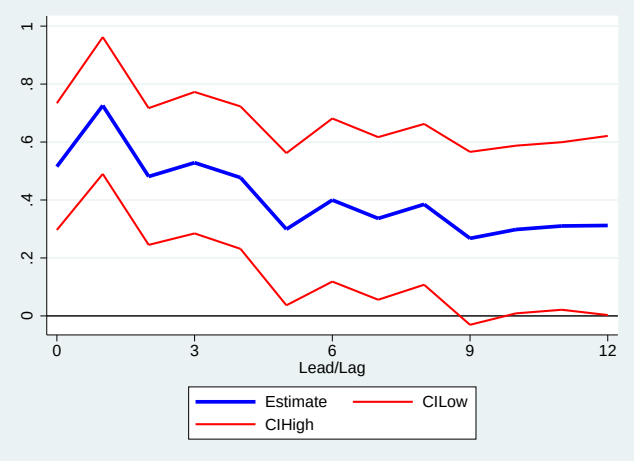


Figure 4: Benefit extensions and Change of Transfers Accumulated

county i with very little retail spending occurring in counties in neighboring states. We can therefore find that increasing benefits in state s not only increases transfers to those living in counties of state s but the transfer recipients also concentrate their spending on their home county. We thus conclude that increasing benefits are a demand stimulus.

Using the same data we also find that households do not systematically change their shopping behavior in response to changes in unemployment benefits. For example, households do not switch to states with less generous benefits to take advantage of lower prices. Thus, we confirm that shopping search behavior does not vary systematically with changes in benefits, validating our use of a simple and transparent specification that ignores shopping decisions.

3.2 Demand Stimulus

In the previous section we documented that an increase in the duration of benefits leads to an increase in transfers from the federal level to the state. In this section we show that this increase also stimulates demand. A simple back of the envelope calculation already shows that this stimulus is quantitatively significant.

In 2011 extended benefits increased average personal income by 1.5%.

On average an US individual spends about 6% of income on food at home (retail), but the average unemployed spends about $20 - 30\%$ of income on food at home (retail).

Therefore the total percentage increase in spending on food due to transfer T:

$$\begin{aligned}\frac{0.3T}{RetailSpending} &= 0.3\frac{T}{USPersonalIncome}\frac{USPers.Income}{RetailSpending} \\ &= 0.3 * 0.015 * \frac{1}{0.06} \\ &= 0.075\end{aligned}$$

To investigate the effect of benefit extension on sales in county i at time t, Q_t^i, we use the following regression:

$$\log(Q_{t+k}^i) = \varphi_k \log(b_t^i) + \chi_k \log(u_{t+k}^i) + \epsilon_{t+k}^i, \tag{11}$$

To control for aggregate demand we add the county unemployment rate, $\log(u_t^i)$, to the regression.

We difference between border counties within a pair p when implementing this regression in the data so that we regress the difference in sales on the difference in benefits between the two counties:

$$\Delta_p \log(Q_{t+k}) = \varphi_k \Delta_p \log(b_t) + \chi_k \Delta_p \log(u_{t+k}) + \psi_p + \delta_t + \Delta_p \epsilon_{t-k}, \tag{12}$$

where Δ_p is the difference between counties i, j, within a pair p, for example $\Delta_p \log(Q_{t+k}) = \log(Q_{t+k}^i) - \log(Q_{t+k}^j)$ is the difference in (log) sales in period $t+k$ between counties i and j. The regression includes county-pair fixed effects ψ_p (for each pair p) and time effects δ_t.

The regression yields coefficients φ_k for $k = 0, 1, \ldots, 12$. These coefficients give the change of sales k quarters ahead in response to an increase in today's benefit duration.

To measure potential anticipatory effects, we also estimate a mirrored response function, where we regress the change in sales between period t and $t-k-1$ on a change in benefits in period t. For country i we implement:

$$\Delta_p \log(Q_{t-k}) = \varphi_{-k} \Delta_p \log(b_t) + \chi_k \Delta_p \log(u_{t-k}) + \psi_p + \delta_t + \Delta_p \epsilon_{t-k}, \tag{13}$$

where $\Delta_p \log(Q_{t-k}) = \log(Q_{t-k}^i) - \log(Q_{t-k}^j)$ is the difference in (log) sales in period $t-k$ between counties i and j. This yields coefficients φ_{-k} for $k = 1, \ldots, 8$. Figure 5 shows the coefficients φ_k for $k = -8, \ldots, 0, \ldots, 12$ we obtain for both the actual and the mirrored response function of sales.

We find no response of sales before benefit duration increased. After the increase has realized, sales respond in hump-shaped way with a peak response between 7 and 8 percent. An estimate of 7 percent yields for two counties with a duration of UI benefits of 50 and 70

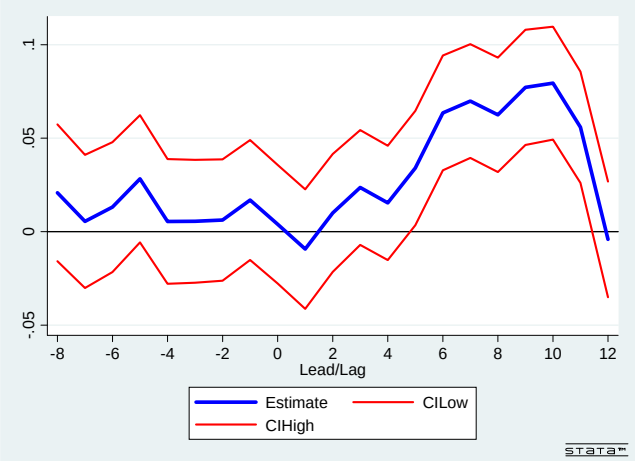


Figure 5: Lagged and Lead Response of Sales

weeks respectively a difference in sales of 2.4%, keeping the unemployment rate unchanged. As the unemployment rate increases in response to an increase in UI duration, the total observed effect - the sum of the direct positive effect of the increase in transfers and the negative effect of a higher unemployment rate - would be smaller than 2.4%. How much smaller depends on the increase in the unemployment rate, which is investigated in a large and active literature (see for example Hagedorn, Karahan, Mitman, and Manovskii (2013)) but not in this paper.

3.3 Inflation

Having established that increases in benefit duration means an increase in a state's purchasing power and that this leads to increase in demand, we now measure the dynamic effect on inflation. To this aim, define the change in the price between the current period (period 0) and the future period k, as

$$\Pi_k^i = \sum_{t=0}^{k} \pi_t^i, \tag{14}$$

where π_t^i is period t inflation rate in county i, that is the log price change between period $t-1$ and t in this county.

Then, differencing between border counties within a pair p and adding pair-fixed effects ψ_p and time effects δ_t yields:

$$\Delta_p \Pi_k = \lambda_k \Delta \log(b_{p,t}) + \psi_p + \delta_t + \Delta\epsilon_{p,t}, \tag{15}$$

where Δ_p is again the difference operator over counties in the same pair. More specifically, if counties i and j are in the same border-county pair p, then $\Delta_p \Pi_k = \Pi_k^i - \Pi_k^j$.

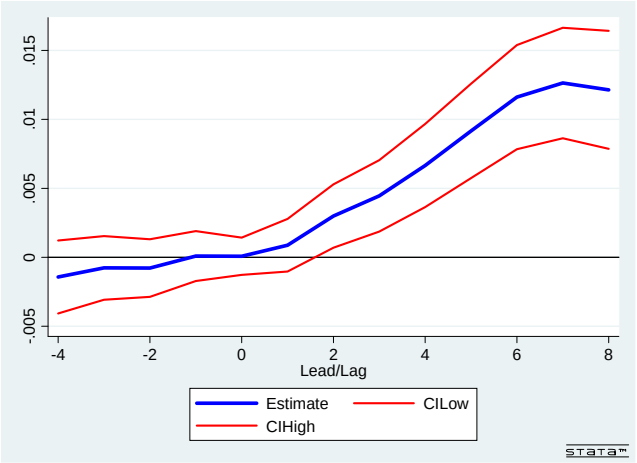


Figure 6: Benefit Effects on Inflation, Leads and Lags

Again, to measure potential anticipatory effects, we also estimate a mirrored response function, where we regress the difference lagged prices on current benefits after differencing between border counties within a pair p:

$$\Delta_p \Pi_{-k} = \lambda_{-k} \Delta \log(b_{p,t}) + \psi_p + \delta_t + \Delta \epsilon_{p,t}, \tag{16}$$

where $\Delta_p \Pi_{-k} = \Pi^i_{-k} - \Pi^j_{-k}$. This yields coefficients λ_{-k} for $k = 1, \ldots, 4$. Figure 5 shows the coefficients φ_k for $k = -4, \ldots, 0, \ldots, 8$ we obtain for both the actual and the mirrored response function of inflation. We find no response of prices before benefit duration increased and a hump-shaped response afterwards with a peak response of about 1.2%.

This finding implies that an increase in the duration of benefits from 50 to 70 weeks leads to a 0.4% peak price difference, an increase from 26 to 99 weeks would lead to a 1.7% price increase.

Consistent with previous results (Hagedorn, Karahan, Mitman, and Manovskii (2013)) we also find here that using the variation in the difference in benefit duration across neighboring counties (which belong to different states) is exogenous. This addresses the potential concern that a shock to the state unemployment rate triggers both a change in benefits and in county inflation leading to a biased estimate of the effects of benefits on county inflation. To see that this concern is not warranted we now just add the difference in state unemployment u^s to the previous regressions

$$\Delta_p \Pi_k = \lambda_k \Delta \log(b_{p,t}) + \mu \Delta u^s_t + \psi_p + \delta_t + \Delta \epsilon_{p,t}, \tag{17}$$

so that the above mentioned potential endogeneity cannot arise. Figure 7 shows the estimated coefficients λ_k for $k = -4, -3, \ldots, 0, 1, \ldots, 8$, which establishes that the coefficients virtually

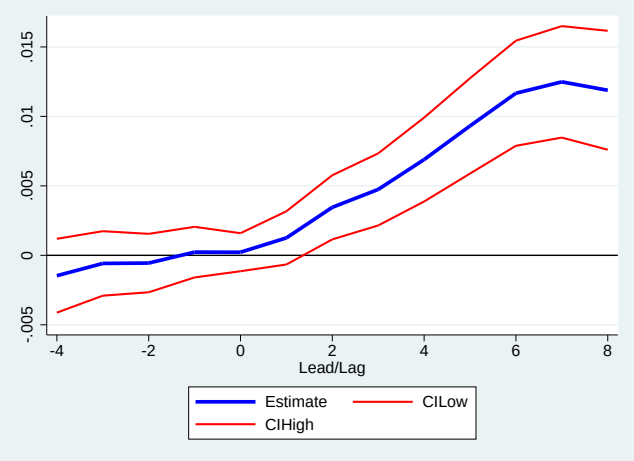


Figure 7: Endogeneity Test: Controlling for state unemployment

do not change in comparison to the case where we do not control for the state unemployment rate (shown in Figure 6). Since the coefficients would change if an endogeneity problem was affecting our results, we conclude that our regressor is exogenous and our estimates are therefore unbiased.

Our findings for the response of local prices also imply that the shopping behavior does not adjust to the change in benefits as one would expect if the considered pair of counties was a frictionless perfectly integrated market. Indeed, we find that a change in benefits does not lead to a change in shopping behavior to take advantage of relative price differentials until these are equalized as the Law of one Price would suggest in a world without mobility costs. Instead we find that the relative price increases. The message that prices respond to local labor market conditions is also broadly consistent with the findings based on the same dataset in Beraja, Hurst, and Ospina (2015).

4 The New Keynesian Model

In the previous Section we have documented that prices increase in response to an expansionary fiscal policy. This reduced form evidence is consistent with the channel which renders fiscal policy effective. Inflation rises if government spending or transfers are increased and as a result real interest rate fall stimulating private consumption expenditures. The argument is however not complete as it only shows the positive price response at the local level whereas one has to know what happens at the aggregate level. To translate the disaggregated results to the aggregate level a theoretical model is needed. In this Section we consider the New Keynesian Model with Calvo pricing and in the next section we consider the Sticky Infor-

mation Model. In these models an expansionary fiscal policy raises firms' marginal costs which induces firms to increase prices. For the model to be useful for a policy analysis at the aggregate level, it has to be consistent with the evidence at the disaggregated level.

We therefore view the data in this section through the lens of the New Keynesian model which is the standard framework to study monetary and fiscal policy. We start with a description of the Calvo pricing and how we use this theoretical restriction when we derive an empirical specification.

4.1 Theory

The theoretical framework we use considers each county as a small open economy in the currency union (U.S.). As a result all counties share a common U.S.-wide monetary policy. We assume a home bias in spending at the county level which we have already established in the data in the previous section. The final consumption good combines local goods/services and imported goods from the rest of the US. No further assumptions on trade or financial markets are made. In terms of firms' price setting we consider Calvo Pricing here and sticky information pricing in the next section.

This is the theoretical framework used in Farhi and Werning (2013) to show that the New Keynesian theory implies that an increase in government transfers to one county (here an increase in b paid for by the federal government) increases marginal costs, which then leads to an increase in the inflation rate in this county.[7]

We now first derive the implications of Calvo pricing and then establish how transfers affect inflation and marginal costs within the model.

4.1.1 Theory: New Keynesian Calvo Pricing

To estimate the macroeconomic effects of unemployment insurance on inflation, we look at firms' pricing decisions in a standard New Keynesian Model with Calvo pricing. In such a model a firm in a county which we call "home" sets (with probability $(1-\theta)$) the price to maximize expected discounted profits. Linearizing this dynamic forward-looking problem

[7] In an open economy with fixed exchange rates (such as the U.S. being the union of all its counties) this increase is due to two effects, a standard Keyensian expansion effect and a neoclassical wealth effect. As in a closed economy, transfers stimulate nominal demand for home goods and sticky prices imply then that output and marginal costs increase. In an open economy with fixed exchange rates, a second effect kicks in as transfers lead to a gradual appreciation of the terms of trade, implying a positive neoclassical wealth effect, which raises wages and thus marginal costs. These two effects have however opposite effects on consumption of home goods as the Keynesian demand expansion stimulates demand and thus output whereas the neoclassical wealth effect decreases labor supply and depresses output. The Keynesian effect dominates in the short-run whereas the neoclassical effect dominates in the long-run.

around the steady state (and some simple algebra) yields that Home inflation π_t^H can be expressed as the discounted sum of deviations of marginal costs from it steady state value:[8]

$$\pi_t^H = \lambda \sum_{k=0}^{\infty} \beta^k E_t(\hat{mc}_{t+k}^H), \tag{18}$$

where the period t deviation of marginal costs (expressed in terms of Home prices) from its steady value is denoted $\hat{mc}_t^H$, E_t is the expectation operator using information available at time t, β is the discount factor and λ is a parameter which depends on fundamentals such as θ, the markup, the degree of decreasing returns to labor and β (see for example Woodford (2003)). This equation is the result of imperfect competition which by itself would lead prices to be set at a markup over marginal cost in a world with flexible prices. With sticky prices à la Calvo, firms can adjust the price every period only with probability $1-\theta < 1$. As a result they cannot set the price to its optimal level every period but instead the price is set at an average markup where the average is computed until the expected next time when the firm is able to change its price.

Using the fact that next period's inflation can be written as

$$\pi_{t+1}^H = \lambda \sum_{k=0}^{\infty} \beta^k E_{t+1}(\hat{mc}_{t+1+k}^H) \tag{19}$$

and the law of iterated expectations yields a Bellman equation for inflation

$$\pi_t^H = \lambda \hat{mc}_t^H + \beta E_t \pi_{t+1}^H, \tag{20}$$

or equivalently for an expectational error η_t

$$\pi_t^H = \lambda \hat{mc}_t^H + \beta \pi_{t+1}^H + \eta_t. \tag{21}$$

This allows us to relate inflation to marginal costs. More precisely, the quasi-differenced inflation rate is proportional to the deviation of marginal costs from steady state:

$$\tilde{\pi}_t^H = \pi_t^H - \beta \pi_{t+1}^H = \lambda \hat{mc}_t^H + \eta_t. \tag{22}$$

This relationship between a firm's marginal cost and the quasi-differenced inflation rate is the same as in a standard Calvo pricing model in a closed economy. The only difference is that the real marginal costs in the Home county in an open economy are expressed in terms of domestic prices, so that the closed economy marginal costs and the open economy local

[8] A derivation can be found in Appendix I.1.

marginal costs differ by $p_t^{CPI} - p_t^H$, the difference of the CPI p_t^{CPI} and the PPI p_t^H.[9]
If markets are complete - the maintained assumption in Gali and Monacelli (2005) - then the marginal cost gap is equal to the output gap as it is in the closed economy case. The reason for this is the Backus Smith condition for complete insurance which tightly links home consumption/output, world consumption/output and real exchange rates such that movements in the terms of trade lead to insurance from the rest of the world.

But equation (22) is valid independent from the financial structure since it just describes firms' optimal price setting, that is, it holds both for complete and incomplete markets (Farhi and Werning (2013)). However, the marginal cost gap is not equal to the output gap if markets are incomplete as there are no insurance transfers in response to changes in the terms of trade. As a result changes in the terms of trade have different effects on Home marginal costs $\hat{mc}^H$ and the output gap.

For our analysis however the question how to relate marginal costs to observables (output gap, labor share, ..) does not arise. Instead we use the theory which implies that the quasi-differenced inflation rate is equal to marginal costs. This not only avoids the notoriously difficult task to measure the output gap in the data which would be proportional to the marginal cost gap if markets were complete. In also avoids the additional difficulties, due to the incompleteness of markets, to measure the effects of movements in the terms of trade on marginal costs. Instead our measure of the marginal cost gap is conditional on the New Keynesian theory and therefore equals $(1/\lambda)\tilde{\pi}_t^H$. As a result all our conclusions are conditional on this theory as long as we impose this model implication to measure marginal costs.

This implication however extends beyond the conventional time-dependent Calvo formulation. The same relationship between quasi-differenced inflation and marginal costs - the Phillips curve - also holds in the Gertler and Leahy (2008) model with state-dependent pricing.[10] The interpretation of the coefficient multiplying marginal costs is different, however. In particular, it does not depend on the probability to change the price as with Calvo pricing (as there is no such thing with state-dependent pricing) but instead it depends on the arrival

[9]Real marginal costs mc^H in the home county are the difference of nominal marginal costs MC^H and the price level p^H in the home county, $mc^H = MC^H - p^H$. Marginal costs can also be expressed in terms of p^{CPI}, so that they can be rewritten as CPI real marginal costs and "terms of trade", $mc^H = (W - p^{CPI}) + (p^{CPI} - p^H)$. What matters for firms' pricing decision is nominal marginal costs though, $MC^H = mc^H + p^H$, which is what enters the pricing equation.

[10]Gertler and Leahy (2008) develop an analytically tractable Phillips curve based on state-dependent pricing. They consider a local approximation around a zero inflation steady state and introduce infrequent idiosyncratic shocks. The resulting Phillips curve is a simple variant of the conventional time-dependent Calvo formulation.

rate of idiosyncratic shocks. This is however not relevant for our analysis so that our findings are not restricted to the popular time-dependent pricing but also apply to state-dependent pricing. These arguments, however do not extend to the sticky information model of Mankiw and Reis (2002). We get back to this model and its empirical implications in Section 5.

Now, denote by p the border-county pair. Then, differencing between border counties within a pair yields:

$$\Delta\tilde{\pi}^H_{p,t} = \lambda\Delta\hat{mc}^H_t + \Delta\eta_{p,t}, \tag{23}$$

where Δ the difference operator over counties in the same pair. More specifically, if counties i and j are in the same border-county pair p, then $\Delta\tilde{\pi}^H_{p,t} = \tilde{\pi}^H_{p,i,t} - \tilde{\pi}^H_{p,j,t}$, with a slight abuse of notation, $\Delta b_{p,t} = \log(b_{p,i,t}) - \log(b_{p,j,t})$ and $\Delta\eta_{p,t}$ is the difference in shocks η across counties.

The above equation relates marginal costs of home producers to PPI inflation π^H whereas in the date we measure CPI inflation π^{CPI}. CPI inflation is a combination of local marginal costs and goods imported by the county (Burstein, Eichenbaum, and Rebelo (2005)). Local marginal costs include distribution costs such as wholesale and retail services, marketing and advertising, and local transportation services. Local goods are goods which are produced solely for the domestic market.

To build on these insights we denote as in Gali and Monacelli (2005) by α the weight on foreign goods so that $1-\alpha$ is the weight on domestic goods. CPI inflation then equals

$$\pi^{CPI}_t = (1-\alpha)\pi^H_t + \alpha\pi^F_t, \tag{24}$$

where π^F_t is the inflation rate of Foreign imported goods. Taking differences over counties in the same pair eliminates the inflation for imported goods so that we obtain the identical equation for CPI inflation that we obtained for PPI inflation:

$$\Delta\tilde{\pi}^{CPI}_{p,t} = (1-\alpha)(\lambda\Delta\hat{mc}^H_t + \Delta\eta_{p,t}) = (1-\alpha)\lambda\Delta\hat{mc}^H_t + \Delta\tilde{\eta}_{p,t}. \tag{25}$$

The difference in quasi-differenced inflation across counties, $\Delta\tilde{\pi}^{CPI}_{p,t}$ is our measure of the difference of marginal costs across counties, which allows us to regress marginal costs on UI benefits. This regression informs us whether the sticky price model is consistent with the empirical relationship between UI benefits and inflation documented in the data since in the model an increase in benefits increases inflation if it increases marginal costs. We therefore estimate

$$\Delta\tilde{\pi}^{CPI}_{p,t} = \chi\Delta b_{p,t} + \Delta\epsilon_{p,t}, \tag{26}$$

where we do not restrict the sign of the coefficient χ, although it is positive in the sticky price model based on the results in Farhi and Werning (2013) independently from the financial

structure. When constructing the quasi-differences at the quarterly frequency, we set $\beta = 0.99$.

Equation (26), which will form the basis of our empirical strategy, is just the New Keynesian Phillips curve which relates the quasi-differenced inflation rate to marginal costs. Because of the forward-looking nature of price setting, a drop in inflation next period is due to a change in future marginal costs and not due to a change in today's marginal costs. Of course the change in say next quarter's marginal costs affects both this period's and next period's inflation but leaves the quasi-differenced inflation rate unaffected if today's marginal costs do not change. The quasi-difference changes only if today's marginal costs change due to a change in demand today. Changes in demand next period only affect next period's marginal costs but not today's marginal costs and therefore the quasi-differenced inflation rate is unchanged. Not taking into account (potentially different) expectations in the two counties say by just considering the difference in inflation, leads to an estimate which is a convolution of our estimate and expectations. And is impossible to interpret within the New Keynesian theory.

4.2 New Keynesian Model: Baseline Empirical Results

Table 2 contains the results of the estimation of the effect of unemployment benefit duration on marginal costs (quasi-differenced inflation) using the baseline specification in Equation (23). We find that changes in unemployment benefit duration have a very small and statistically highly insignificant short-run effect on inflation: a 1% rise in benefit duration for one quarter increases the inflation rate by 0.0004 percentage points.

Estimated equation:

$$\Delta\tilde{\pi}_{p,t} = \gamma\Delta b_{p,t} + \psi_p + \xi_t + \Delta\nu_{p,t}.$$

Table 2: Unemployment Benefit Extensions and Inflation

VARIABLES	Marginal Costs (QD-Inflation)
Weeks of Benefits	0.0004
	(0.0010)
Observations	16790
R-squared	0.0032

Note - Standard errors in parentheses

It is important to keep in mind that this finding is based on the difference across pairs

of border counties. Thus, the effects of various other shocks or policies that affect these counties symmetrically are differenced out. In particular monetary policy is the same for all counties and thus cancels out.

The point estimate of 0.0004 is statistically very insignificant. It is also economically virtually zero. The point estimate implies that increasing benefits from 26 weeks to 99 weeks increases inflation by

$$0.0004(log(99) - log(26)) = 0.0005, \tag{27}$$

that is by 0.05 percentage points. An increase from 50 weeks to 70 weeks would increase inflation by $0.0004(log(70) - log(50)) = 0.0001$, that is by 0.01 percentage points.[11] Looking at two standard deviations does not change this conclusion. Using an estimate of $2 * 0.0010 = 0.002$ yields

$$0.002(log(99) - log(26)) = 0.0027, \tag{28}$$

that is an increase by 0.27 percentage points.

Furthermore the effect of an increase in benefits today on future marginal costs is even negative (insignificantly) at most leads. Applying the same simple calculation would yield a decrease of marginal costs of the same small magnitude.

Lagged Response

Estimate: $\Delta\tilde{\pi}_{p,t+\mathbf{k}} = \gamma\Delta b_{p,t} + \psi_p + \xi_t + \Delta\nu_{p,t}$

VARIABLES	Marginal Costs (QD-Inflation)
Weeks of Benefits	0.0004
	(0.0013)
Weeks of Benefits (Lag 1)	-0.0015
	(0.0013)
Weeks of Benefits (Lag 2)	-.0004
	(0.0014)
Weeks of Benefits (Lag 3)	-0.0016
	(0.0014)
Weeks of Benefits (Lag 4)	-0.0003
	(0.0015)

Standard errors in parentheses.

[11] The conclusion does not change if one considers the 2-quarter ahead forecast, 0.0014. In this case inflation increases by $0.0004(log(99) - log(26)) = 0.0019$, that is 0.19 percentage point increase.

We conclude that the effect on inflation is zero conditional on using the New Keynesian model. This finding is also not the result of a lack of power as we have established in Section 3.3 that raising benefits has significant effects on inflation.

4.3 New Keynesian Model: Robustness

In section we analyze the effects of benefits on marginal costs and the robustness of our results to price indexation and different discounting of future marginal costs.

4.3.1 Price indexation and Trend Inflation

The textbook forward-looking new Keynesian model used in Section 4 fails to account for inflation persistence. As a response the literature has allowed for some form of price indexation, which adds a backward looking element to the determination of current inflation (Woodford (2003), Christiano, Eichenbaum, and Evans (2005)). In this Section we show that our results are unchanged add price indexation to the model or allow for trend inflation.

With price indexation, firms that cannot re-optimize price in some period choose to automatically index their price to some overall price index. The motivation is that a rule, if simple enough, saves managerial costs, so that such a rule of thumb is used between times when the prices are re-optimized. With (partial) indexation of degree γ to a reference inflation rate π^R the optimal pricing results in the following equation:

$$\pi_t - \gamma\pi_{t-1}^R = \lambda\hat{mc}_t + \beta(\pi_{t+1} - \gamma\pi_t^R) + \eta_t, \tag{29}$$

so that we obtain the quasi-differenced equation

$$\tilde{\pi}_t^\gamma = (\pi_t - \gamma\pi_{t-1}^R) - \beta(\pi_{t+1} - \gamma\pi_t^R) = \lambda\hat{mc}_t + \eta_t. \tag{30}$$

It would be in the spirit of the underlying assumption that monitoring of current and local conditions is too costly to be worthwhile that firms decide to index to some readily available (lagged) inflation rate, such as the US CPI. Whenever costs induce firms do so or at least if firms from counties in a pair index to the same price index, this indexation drops out when taking differences between counties.

Only in the (unlikely) case that firms index to a local price index, taking differences does not cancel the indexation terms since the PPI is potentially different in the two counties. Again, we say unlikely because collecting the information to index to the local price index seems similar to the information necessary to set an optimal price which is assumed to be too costly.

Table 3: Unemployment Benefits on Inflation: Price Indexation

Indexation	.1	.2	.3	.4	.5
Weeks of Benefits	-0.000599 (0.00197)	-0.000672 (0.00208)	-0.000744 (0.00220)	-0.000817 (0.00233)	-0.000889 (0.00247)
Observations	15,247	15,247	15,247	15,247	15,247
Number of Pairs	729	729	729	729	729
Indexation	.6	.7	.8	.9	1
Weeks of Benefits	-0.000962 (0.00261)	-0.00103 (0.00276)	-0.00111 (0.00292)	-0.00118 (0.00307)	-0.00125 (0.00323)
Observations	15,247	15,247	15,247	15,247	15,247
Number of Pairs	729	729	729	729	729

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Note - Results from the regression of $(\pi_t - \gamma\pi_{t-1}) - \beta(\pi_{t+1} - \gamma\pi_t) = \chi\Delta b_{p,t} + \ldots$, where the degree of price indexation γ takes values $\{0.1, 0.2, \ldots, 0.9, 1\}$.

If firms index to the local inflation rate, the quasi-differenced inflation rate is defined as

$$\tilde{\pi}_t^{\gamma} = (\pi_t - \gamma\pi_{t-1}) - \beta(\pi_{t+1} - \gamma\pi_t), \tag{31}$$

where the lagged county specific inflation terms do not drop out. Since we have no good information about γ and estimated DSGE models give numbers all over the place, we redo the benchmark for various values of γ. The degree of indexation γ used to define the quasi-differenced inflation rate takes values in $\{0.1, \ldots, 1\}$.Table 3 shows the results. We find that the effect of benefits on inflation is virtually zero for all degrees of indexation.

This conclusion about inflation also does not change if trend inflation $\bar{\pi}_t$ is not zero, but now applies to the inflation gap $\hat{\pi}_t = \pi_t - \bar{\pi}_t$ instead of to the inflation rate. Indeed, as in Cogley and Sbordone (2008) and in Mavroeidis, Plagborg-Moller, and Stock (2014), if nonresetting firms' prices are indexed to a mixture of past inflation with weight ρ and current trend inflation with weight $1 - \rho$, then

$$\hat{\pi}_t = \gamma_f E_t\hat{\pi}_{t+1} + \gamma_b\hat{\pi}_{t-1} + \lambda\hat{mc}_t + \gamma_b(\beta E_t\Delta\bar{\pi}_{t+1} - \Delta\bar{\pi}_t), \tag{32}$$

where $\gamma_f = \beta/(1+\beta\rho)$ and $\gamma_b = \rho/(1+\beta\rho)$. The last term allows for changes in the inflation

target, that means, since the inflation target is determined by the long run target of the central bank, changes in the target inflation rate. Collecting terms in the above equation yields an equation that differs from the standard Phillips curve (with indexation) by a term which depends on trend inflation only:

$$\pi_t = \gamma_f E_t \pi_{t+1} + \gamma_b \pi_{t-1} + \lambda \hat{mc}_t + (\bar{\pi}_t - \gamma_f E_t \bar{\pi}_{t+1} - \gamma_b \bar{\pi}_{t-1} + \gamma_b (\beta E_t \Delta \bar{\pi}_{t+1} - \Delta \bar{\pi}_t), \quad (33)$$

and after some simple manipulations is equivalent to

$$\underbrace{\pi_t - \rho \pi_{t-1} = \lambda \hat{mc}_t + \beta E_t(\pi_{t+1} - \rho \pi_t)}_{\text{Phillips curve with indexation } \rho} + \underbrace{(1+\beta\rho)\{(\bar{\pi}_t - \gamma_f E_t \bar{\pi}_{t+1} - \gamma_b \bar{\pi}_{t-1}) + \gamma_b(\beta E_t \Delta \bar{\pi}_{t+1} - \Delta \bar{\pi}_t)\}}_{\text{Due to Trend Inflation}}, \quad (34)$$

Since our empirical methodology involves taking differences across counties within a pair, the latter term drops out and we end up with the same specification for trend inflation as for the standard model with price indexation of degree $\rho \in [0, 1]$. As the result of no inflationary effects holds for all degrees of price indexation ranging from 0 to 1, the same conclusion applies to the model with trend inflation: Government transfers paid by foreigners have no effect on marginal costs and inflation using the New Keynesian model.

4.3.2 Discounting of Marginal Costs

In the textbook New Keynesian model future marginal costs and thus future inflation are discounted at rate β, which we also assume in our benchmark specification. In this section we explore the sensitivity of our results to this assumption. Specifically, we redo the benchmark analysis with quarterly discount rate ranging between 0.9 and 1, that is the quasi-differenced inflation rate is defined as

$$\tilde{\pi}_t = \pi_t - \hat{\beta} \pi_{t+1}, \quad (35)$$

where the discount factor $\hat{\beta}$ takes values in $\{0.9, 0.91, \ldots, 0.99, 1\}$ including the benchmark value 0.99. Table 4 shows the results.

We find that the effect of benefits on inflation is virtually zero for all discount factors. This establishes that our findings are not only robust to mismeasurement of the discount factor but also to model variation which may lead to different effective discounting by firms. For example, with endogenous investment, Woodford (2003) shows that quasi-differenced inflation is still proportional to marginal costs (correct: deviation from its steady state value). The discount factor is not necessarily exactly equal to β but close to it at shorter frequencies. The reason that future marginal costs are discounted heavier in the presence of endogenous investment is that firms do not have to increase prices now if marginal costs increase in say five

Table 4: Unemployment Benefits on Inflation: Price Indexation

	(1)	(2)	(3)	(4)	(5)	(6)
DISCOUNT FACTOR	.9	.91	.92	.93	.94	.95
Weeks of Benefits	-0.000174	-0.000156	-0.000139	-0.000122	-0.000104	-8.69e-05
	(0.00175)	(0.00176)	(0.00177)	(0.00178)	(0.00179)	(0.00180)
Observations	15,982	15,982	15,982	15,982	15,982	15,982
Number of Pairs	729	729	729	729	729	729
	(7)	(8)	(9)	(10)	(11)	
DISCOUNT FACTOR	.96	.97	.98	.99	1	
Weeks of Benefits	-6.96e-05	-5.22e-05	-3.48e-05	-1.75e-05	-1.05e-07	
	(0.00181)	(0.00182)	(0.00183)	(0.00183)	(0.00184)	
Observations	15,982	15,982	15,982	15,982	15,982	
Number of Pairs	729	729	729	729	729	

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Note - Results from the regression of $\pi_t - \hat{\beta}\pi_{t+1} = \chi \Delta b_{p,t} + \ldots$, where the discount factor $\hat{\beta}$ takes values $\{0.9, 0.91, \ldots, 0.99, 1\}$.

years but instead can invest into physical capital in the meantime. In Woodford's calibration increases in marginal costs within the next 3 years increase inflation today. This of course depends on his choice of capital adjustment costs. Obviously, with infinite capital adjustment costs the model is equivalent to a model with exogenous investment and discounting at the time preference rate. For our purposes the conclusion is that a monthly or quarterly frequency discounting at rate of time preference is reasonable as firms can hardly adjust their capital stock. But even if in a parametrization substantially different from Woodford's one this is not the case, Table 4 shows the robustness of our results to a wide range of values.

5 The Mankiw and Reis (2002) Sticky Information Model

An alternative theory to the standard New Keynesian Phillips Curve was proposed by Mankiw and Reis (2002). In this theory, prices adjust slowly since information disseminates slowly and processing information takes time. In this Section we use this theory to

measure the effect of government transfers on inflation.

In each period, a fraction θ of firms obtains new information. Whereas firms in the New Keynesian model are able to adjust the price only with some probability, in this model firms can adjust prices any time but based on potentially old information. Mankiw and Reis (2002) show that the inflation rate π_t^H satisfies:

$$\pi_t^H = \frac{\lambda\theta}{1-\theta}\hat{mc}_t^H + \theta\sum_{j=0}^{\infty}(1-\theta)^j E_{t-j}(\pi_t^H + \lambda\Delta\hat{mc}_t^H), \tag{36}$$

where E_{t-j} denotes expectations based on period $t-j$ information and $\Delta\hat{mc}_t^H = \hat{mc}_t^H - \hat{mc}_{t-1}^H$ is the change in marginal costs. Although firms are forward-looking, the inflation rate in this model has a backward-looking element built in as past expectations of current conditions matter.

Consider now an unexpected increase in transfer in period 0 and firms start obtaining this information at rate θ. As a result all previous expectations for period $-1, -2, \ldots$ are zero, $E_{-j}(\pi_k^H + \lambda\Delta\hat{mc}_k^H) = 0$ for all $k = 0, 1, \ldots$ and $j = 1, 2, \ldots$ and information obtained after period 0 is correct, $E_s(\pi_t^H + \lambda\Delta\hat{mc}_t^H) = \pi_t^H + \lambda\Delta\hat{mc}_t^H$ for $0 \leq s \leq t$. We can therefore solve for marginal costs in terms of inflation:

$$\hat{mc}_t^H = \frac{1}{\lambda}\frac{(1-\theta)^{t+1}}{1-(1-\theta)^{t+1}}\sum_{k=0}^{t}\pi_k^H. \tag{37}$$

It is easy to verify this result by plugging it into equation (36).

Obviously, marginal costs depends on the parameter θ, the speed at which information disseminates. Before we try to explore which values of θ are suggested by the data, we compute the response of $\sum_{k=0}^{t}\pi_k^H$ to an increase in the duration of benefits b. Figure 8 shows the response, which is positive for all lags $k = 0, \ldots 6$ and becomes significant for $k \geq 2$.

The next step is to compute the implications for marginal costs, assuming different values of $\theta = 0.25, 0.5, 0.6, 0.75$ and $\lambda = 0.02$, which are shown in Figure 9.

Since the inflation rate increases at all lags, the real interest rate drops, implying a falling profile for marginal costs. As cab be seen in Figure 9 marginal costs are not falling for any value of θ, as there is always an increase in marginal costs in the first period. It is however easy to overcome this small problem which is due to the empirical finding that inflation does not respond much in period 0. Mankiw and Reis (2002)) assume that firms in period 0 can incorporate their newly obtained information already in period zero prices. An alternative assumption is that prices are set before this information is obtained, so that

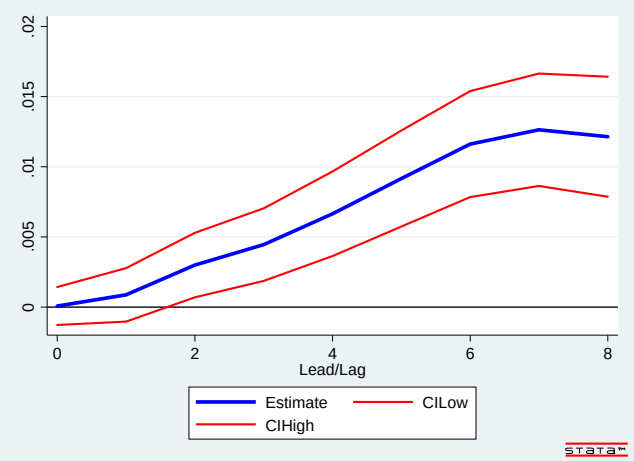


Figure 8: Response of $\sum_{k=0}^{t} \pi_k^H$ for $t = 0, \ldots, 8$

firms which obtain the information during period 0 use this information in their price setting decision in period 1, firms obtaining information in period 1 use it their prices in period 2 and so on. As a result, $E_{-j}(\pi_k^H + \lambda \Delta \hat{mc}_k^H) = 0$ for all $k = 0, 1, \ldots$ and $j = 0, 1, \ldots$ and $E_s(\pi_t^H + \lambda \Delta \hat{mc}_t^H) = \pi_t^H + \lambda \Delta \hat{mc}_t^H$ for $1 \leq s \leq t$. In particular inflation does not respond in period 0 and the dynamic response is shifted by one period, so that marginal costs now equal, for all $t \geq 1$,

$$\hat{mc}_t^H = \frac{1}{\lambda} \frac{(1-\theta)^{t+1}}{1-(1-\theta)^{t+1}} \sum_{k=1}^{t} \pi_k^H . \tag{38}$$

Given this timing assumption informs us about the value of θ, as different values imply different slopes for marginal costs.

As we can see from Figure 10, for $\theta = 0.25$ (the preferred value of Mankiw and Reis (2002)) as well as for $\theta = 0.4, 0.5$ this is not the case. Marginal costs are not monotonically falling if $\theta < 0.59$ and are monotonically decreasing only for $\theta \geq 0.59$.

5.1 Discussion

So far we have established that stimulus policy has no effect on the quasi-differenced inflation rate, which in the New Keynesian model equals marginal costs (more precisely the deviation of marginal costs from its steady state value). Since the inflation rate in this model is equal to the present discounted value of marginal costs, stimulus policy has no effect on inflation either. As we have established that the inflation rate responds to stimulus policy in the empirical analysis in Section 3.3, we conclude that the the sticky price model is inconsistent with the data.

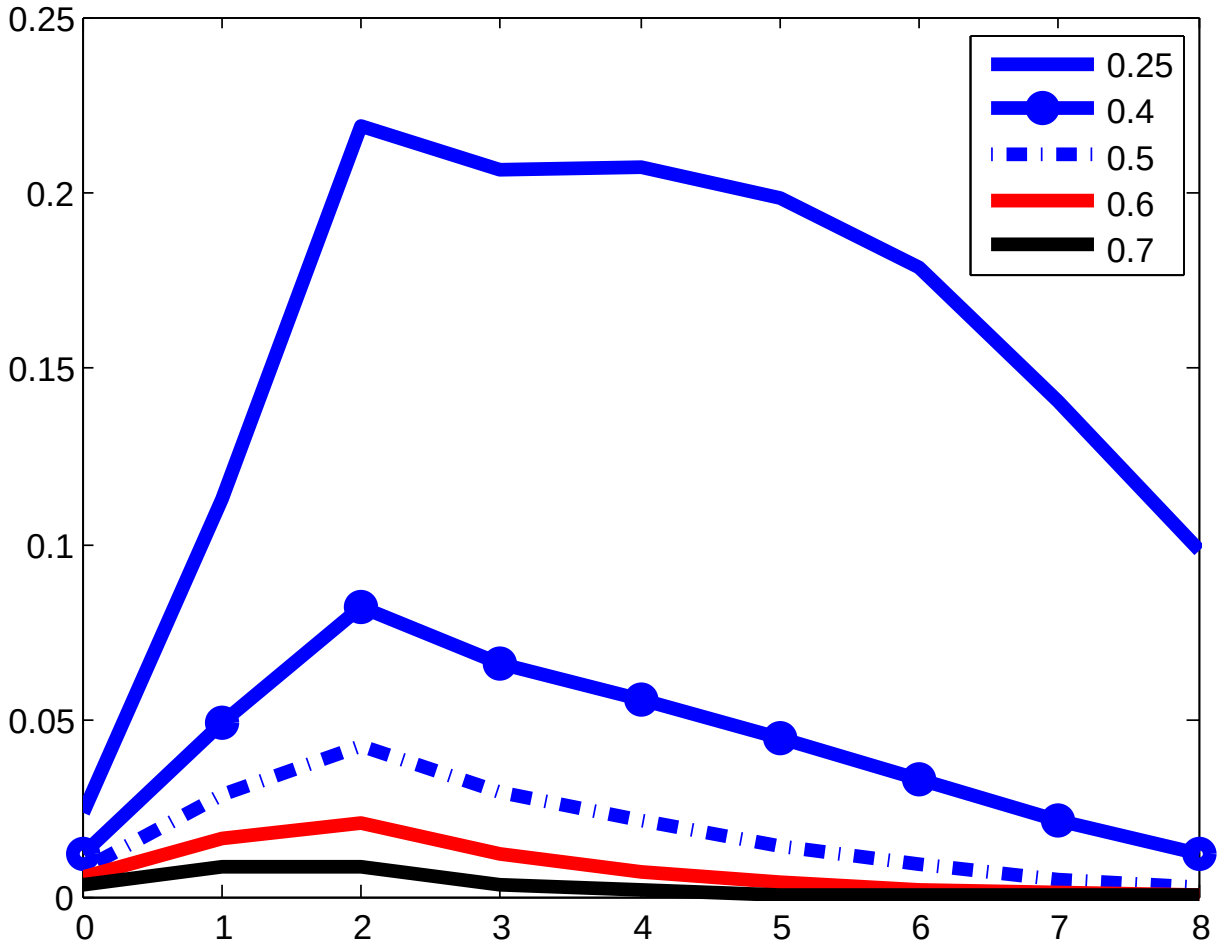


Figure 9: Dynamic Response of Marginal costs, $100\hat{mc}_k^H$ for $\theta = 0.25, 0.5, 0.6, 0.75$.

The sticky information model on other hand is consistent with the data, which begs the question why this is.

To understand this we use the sequence of marginal costs as implied by the sticky information model (see equation 38), and compute the inflation rate that the New Keynesian model would imply given this sequence of marginal costs. Figure 11 shows this "sticky price" inflation rate, marginal costs based on the sticky information model and the inflation rate in the data. Two observations emerge from this figure. First, the "sticky price" inflation rate and the inflation rate in the data are divergent. The observed inflation rate builds up over time whereas the "sticky price" inflation rate is fading out. This is due to the second observation that the inflation rate tracks marginal costs quite closely which is since the inflation rate reflects future marginal costs. These two observations together explain why the sticky price model is inconsistent with the data. With Calvo pricing, firms engage in preemptive price setting. Whenever a firm has the possibility to change a price, it takes not only current

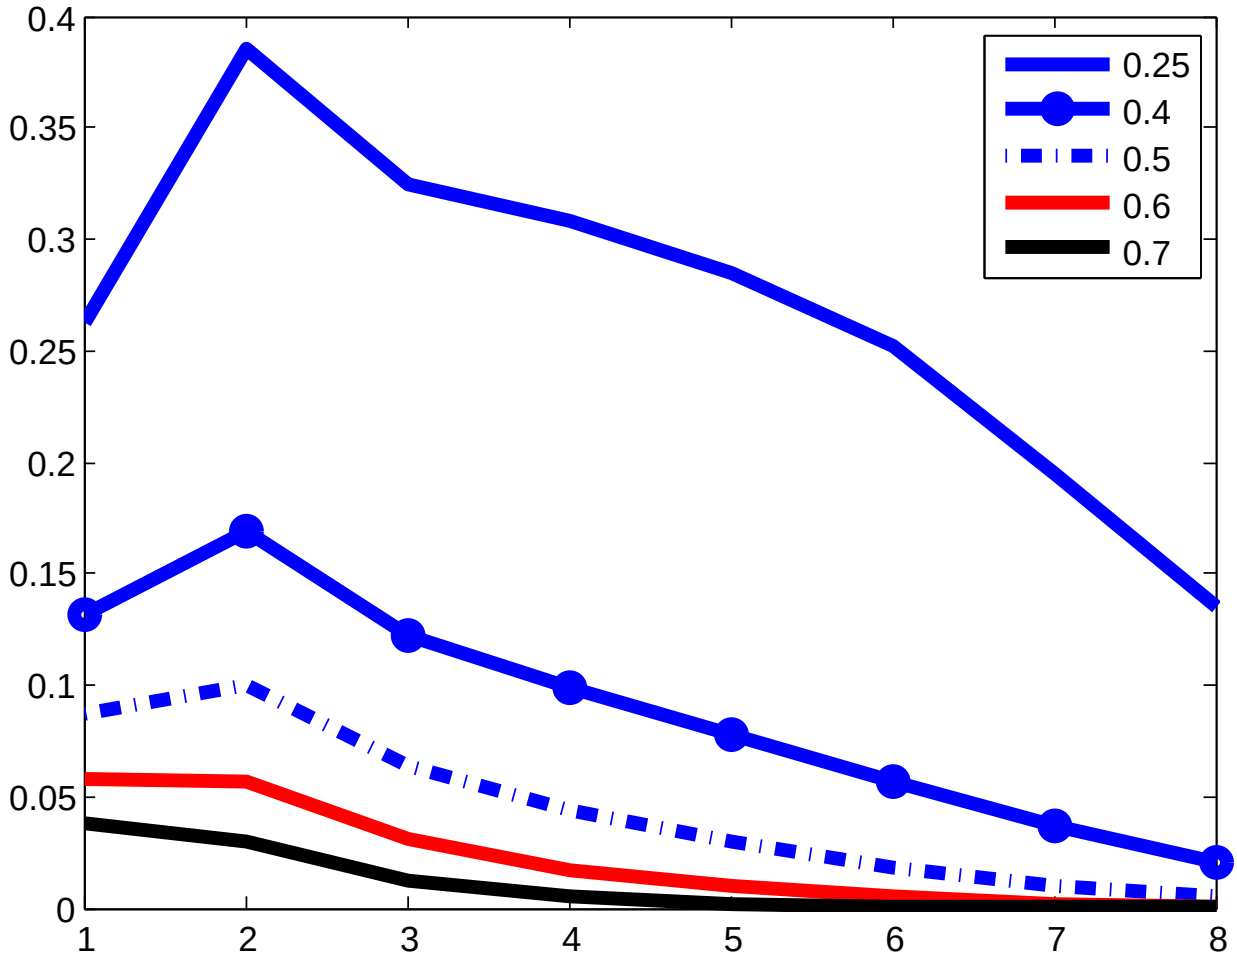


Figure 10: Dynamic Response of Marginal costs (Predetermined Prices), $100\hat{mc}_k^H$ for $\theta = 0.25, 0.5, 0.6, 0.75$.

marginal costs into account but also future marginal costs and future price changes of other firms. This is because, when other firms change prices in the future, the firm may not be able to change its price as well. As a result it has to form expectations about future marginal costs since those imply firms' future price adjustments and has to incorporate these expectations into its current price setting. This leads to preemptive price setting since future increases in marginal costs make firms adjust their prices already preemptively today. The empirical implication of this reasoning is that the inflation rate jumps up and then decays over time as shown in Figure 11, that is

$$\pi_{t+1} - \pi_t = (1-\beta)\pi_{t+1} - \lambda \hat{mc}_t < 0. \tag{39}$$

The empirical response of inflation on the other hand is not decaying but instead building up over time. That is why the sticky price model is inconsistent with the data. The intuition

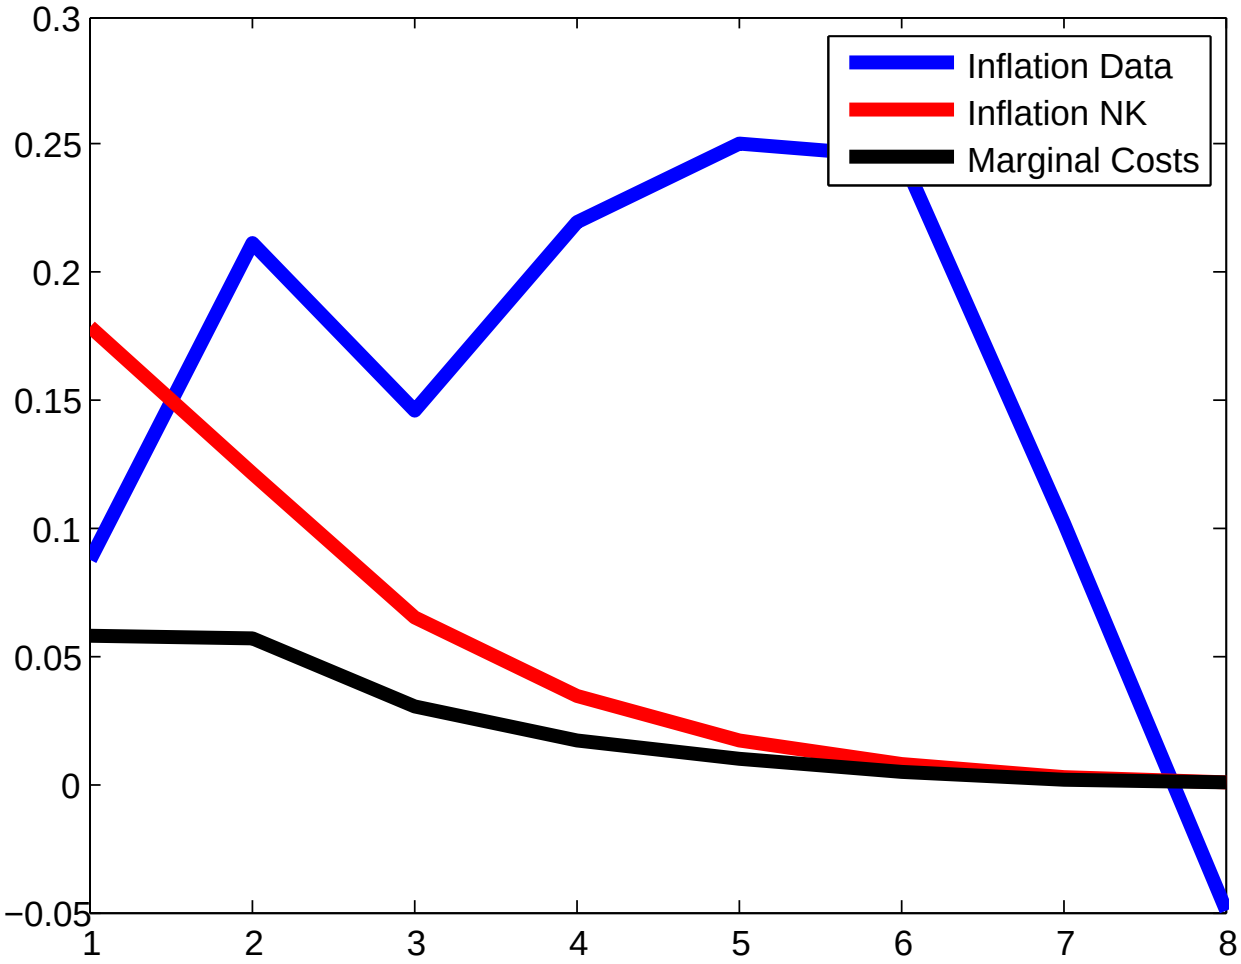


Figure 11: Sticky Prices vs. Sticky Information: Marginal Cost and Inflation

for this inconsistency provided above used the simple sticky price model but we have shown in Sections 4.3 that various modifications do not overturn this intuition and do not make the sticky price model consistent with the data.

Pricing setting in the sticky information model incorporates expectations quite differently with the implication that it is consistent with the data. A firm which obtains full information does not engage in preemptive price setting. Instead it sets prices to its current optimal level - a markup over current marginal costs. When other firms get informed in the future, the firm will readjust its price again, sot that both the newly informed firms and the previously informed firms adjust the price. The empirical implication is that - as we observe in the data - that the price increase is smaller in the beginning and increases over time.

Since our finding that inflation build up over time reflects (informed) firms optimal forward looking pricing decisions, we can also reject the the potential concern that sticky-

information expectations can be confused for adaptive expectations.

6 Conclusion

The relationship between government spending, marginal costs and inflation is at the core of New Keynesian theories of price setting. In particular, the impact of a demand stimulus on marginal costs and inflation determines the size of the fiscal multiplier. Yet, despite the fundamental importance of this relationship to theory, documenting it in the data has proved elusive.

This paper is the first to provide a direct measurement of this relationship. We estimate the effect of a fiscal expansion on inflation using the Nielsen Retail Scanner Data. To identify this response of inflation to a demand stimulus (an increase in the generosity of unemployment insurance) we exploit a policy discontinuity at U.S. state borders. We find that inflation (and also spending) responds with a delay in a hump-shaped manner.

To assess the impact on marginal costs we explore whether this response is consistent with two leading models of price setting, the sticky price model and the sticky information model, which are known to have different implications for the size of the fiscal multiplier. We find that the sticky information model is consistent with our empirical findings whereas the sticky price model is not.

Bibliography

Beraja, M., E. Hurst, and J. Ospina (2015): "The Aggregate Implications of Regional Business Cycles," Working paper, University of Chicago.

Burstein, A., M. Eichenbaum, and S. Rebelo (2005): "Large Devaluations and the Real Exchange Rate," *Journal of Political Economy*, 113(4), pp. 742–784.

Christiano, L., M. Eichenbaum, and S. Rebelo (2011): "When Is the Government Spending Multiplier Large?," *Journal of Political Economy*, 119(1), pp. 78–121.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113(1), pp. 1–45.

Cogley, T., and A. M. Sbordone (2008): "Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve," *American Economic Review*, 98(5), 2101–26.

Dube, A., W. Lester, and M. Reich (2010): "Minimum Wage Effects Across State Borders: Estimates Using Contiguous Counties," *Review of Economics and Statistics*, 92(4), 945 – 964.

Farhi, E., and I. Werning (2013): "Fiscal Multipliers: Liquidity Traps and Currency Unions," mimeo.

Gali, J. (2008): *Monetary Policy, Inflation, and the Business Cycle*. Princeton University Press.

Gali, J., and M. Gertler (1999): "Inflation Dynamics: A Structural Econometric Analysis," *Journal of Monetary Economics*, 44(2), 195 – 222.

Gali, J., and T. Monacelli (2005): "Monetary Policy and Exchange Rate Volatility in a Small Open Economy," *Review of Economic Studies*, 72, 707 – 734.

Gertler, M., and J. Leahy (2008): "A Phillips Curve with an Ss Foundation," *Journal of Political Economy*, 116(3), pp. 533–572.

Hagedorn, M., F. Karahan, K. Mitman, and I. Manovskii (2013): "Unemployment Benefits and Unemployment in the Great Recession: The Role of Macro Effects," NBER Working Papers 19499, National Bureau of Economic Research, Inc.

Holmes, T. J. (1998): "The Effect of State Policies on the Location of Manufacturing:

Evidence from State Borders," *Journal of Political Economy*, 106(4), 667–705.

Mankiw, N. G., and R. Reis (2002): "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *The Quarterly Journal of Economics*, 117(4), 1295–1328.

Mavroeidis, S., M. Plagborg-Moller, and J. Stock (2014): "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve," *Journal of Economics Literature*, 52, 124–188.

Woodford, M. (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.

APPENDICES

I Some Derivations

I.1 New Keynesian Pricing: Details

We derive equation (18) in the main text:

The optimal price for Home profit maximizing producer equals

$$p^*_{H,t} = (1-\beta\theta)\sum_{k=0}^{\infty}(\beta\theta)^k E_t(mc^H_{t+k} - (-\mu) + p^H_{t+k}),$$

reflecting the intuition provided in the main text that the firm would like to set its nominal price as markup markup over marginal costs but due to the Calvo assumption cannot do so. Instead it sets the price such that in expectation the price is close to its optimal value.

Subtracting p^H_{t-1} and using (π^H denotes the inflation rate, in steady state $mc = -\mu$):

$$p^*_{H,t} - p^H_{t-1} = (1-\beta\theta)(\sum_{k=0}^{\infty}(\beta\theta)^k E_t(\hat{mc}^H_{t+k}) + \sum_{k=0}^{\infty}(\beta\theta)^k E_t\pi^H_{t+k}),$$

which can be written recursively as

$$p^*_{H,t} - p^H_{t-1} = \beta\theta E_t(p^*_{H,t+1} - p^H_t) + (1-\beta\theta)\hat{mc}^H_t + \pi^H_t.$$

This yields the inflation equation ($\pi^H_t = (1-\theta)(p^*_{H,t} - p^H_{t-1})$):

$$\pi^H_t = \lambda\sum_{k=0}^{\infty}\beta^k E_t(\hat{mc}^H_{t+k}),$$

for $\lambda = \frac{(1-\beta\theta)(1-\theta)}{\theta}$. It is straightforward to extend this analysis to the case where firms face decreasing returns to scale in production and have different marginal costs due to different levels of production (see for example Gali (2008)).

II Appendix Figures

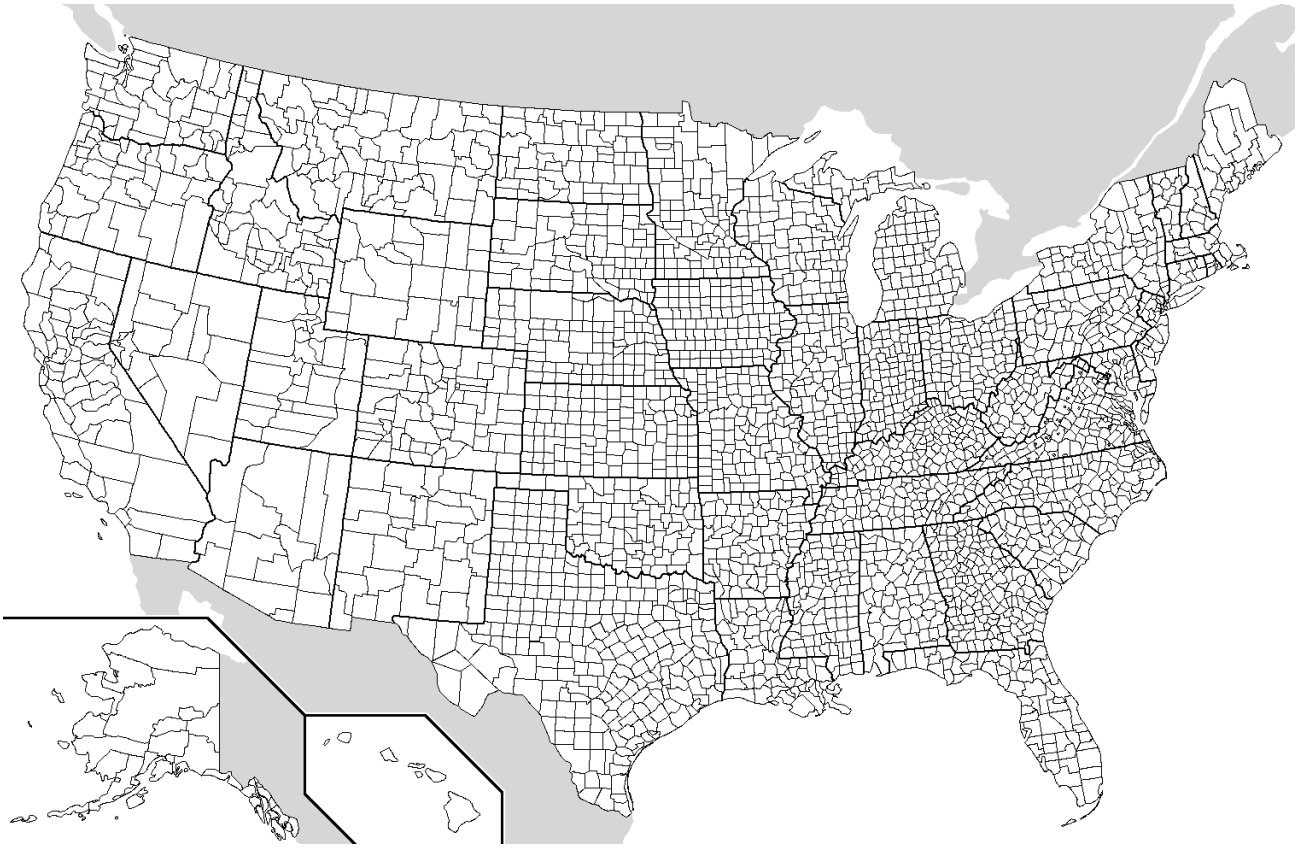
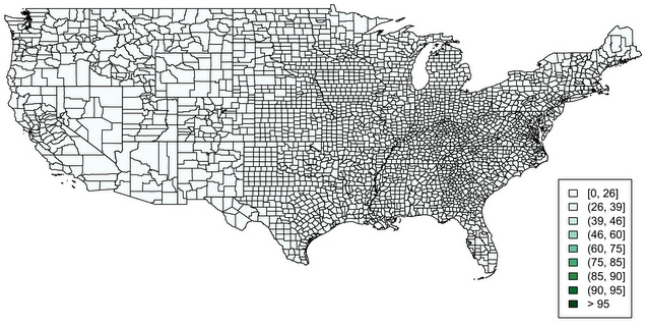
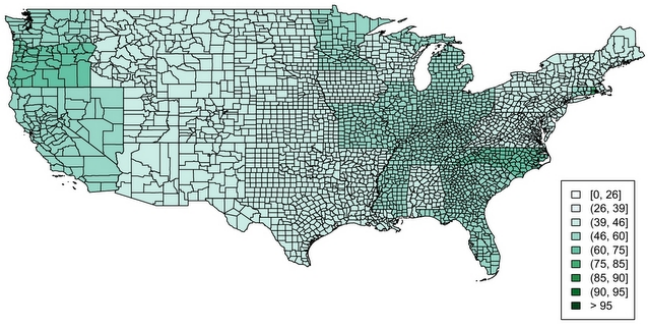


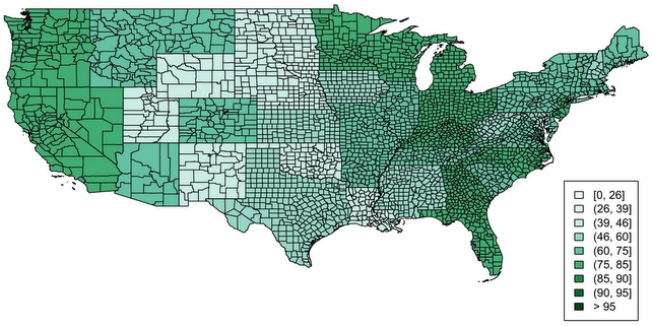
Figure A-1: Map of U.S.A. with state and county outlines.



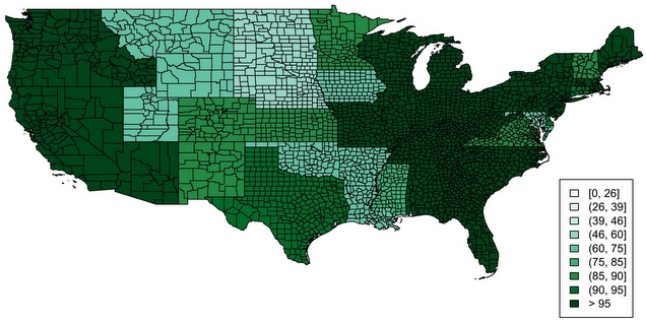
(a) June 2008.



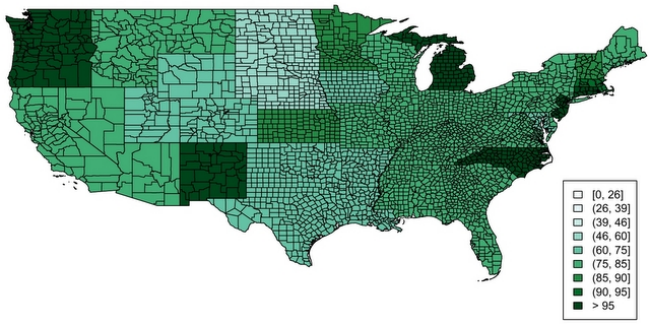
(b) December 2008.



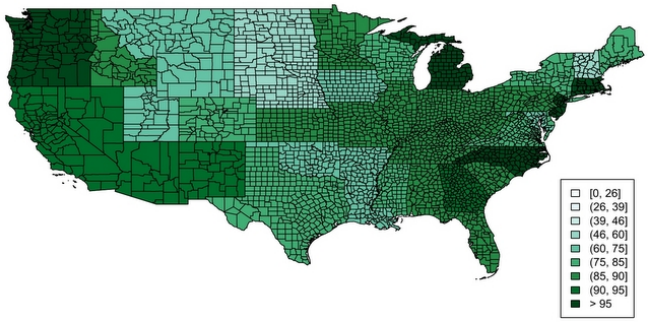
(c) June 2009.



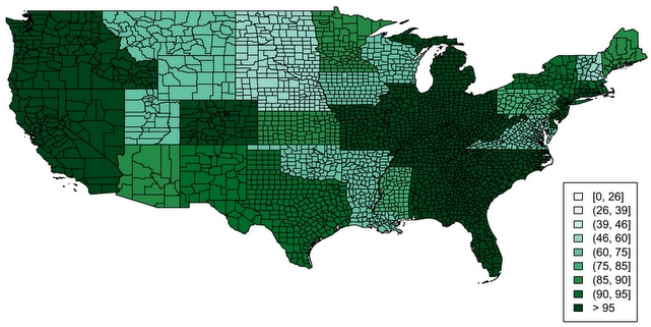
(d) December 2009.



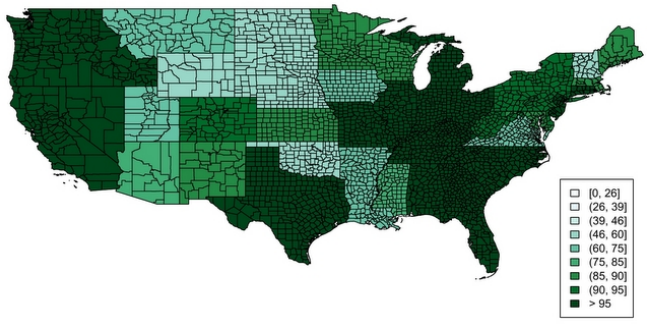
(e) June 2010.



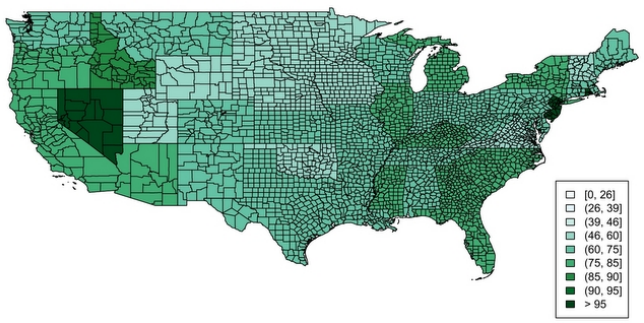
(f) December 2010.



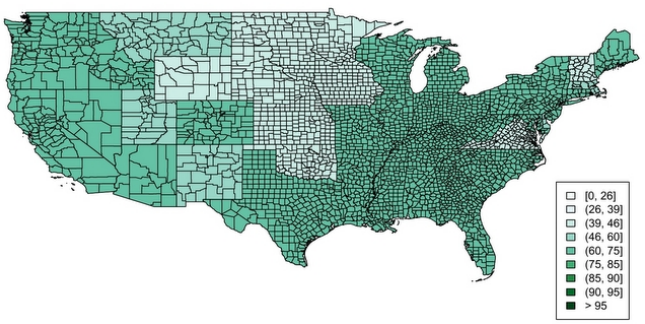
(g) June 2011.



(h) December 2011.



(i) June 2012.



(j) December 2012.

Figure A-2: Unemployment benefit duration across U.S. states during the Great Recession. Selected months.