House Prices, Local Demand, and Retail Prices

Johannes Stroebel

New York University, National Bureau of Economic Research, and Centre for Economic Policy Research

Joseph Vavra

University of Chicago and National Bureau of Economic Research

We document a causal response of local retail prices to changes in local house prices, with elasticities of 15–20 percent across housing cycles. These price responses are largest in zip codes with many homeowners and are driven by changes in markups rather than local costs. We argue that markups rise with house prices because greater housing wealth reduces homeowners' demand elasticity, and firms raise markups in response. Shopping data confirm that house price changes affect the price sensitivity of homeowners, but not that of renters. Our evidence suggests a new source of business cycle markup variation.

How do prices and markups respond to demand shocks? This question is of central importance for business cycle modeling, and a large empirical literature has tried to provide answers using aggregate time-series data.

However, this approach requires strong assumptions, both to identify aggregate demand shocks and to measure markups; consequently, the literature has arrived at conflicting conclusions regarding the cyclicality of markups (see Nekarda and Ramey [2013] for a review). Analyzing only time-series data also makes it hard to isolate the channels that explain any observed relationship.

In this paper, we instead turn to microdata to provide direct causal evidence on the response of retail price-setting and household shopping behavior to changes in wealth and demand, and in doing so we propose a new channel for business cycle variation of markups. In a series of papers, Mian and Sufi (2011, 2014a, 2014b) and Mian, Rao, and Sufi (2013) document that local house price movements have strong effects on local demand. In this paper, we link retailer scanner price data and household shopping data to zip code–level house prices to identify the response of price-setting and shopping behavior to these house price–induced local demand shocks. We provide evidence that households' elasticity of demand, and thus firms' markups, vary substantially in response to these shocks.

Our first result is that retail prices increase much more in regions with higher house price growth. We argue for a causal relationship from house prices to retail prices using two alternative and complementary identification strategies. First, we follow the identification strategy in Mian and Sufi (2011) and use measures of local housing supply elasticity constructed by Gyourko, Saiz, and Summers (2008) and Saiz (2010) as instruments for local house price movements. Across a variety of empirical specifications, we estimate an elasticity of local retail prices to house price movements of 15–20 percent. This elasticity is highly significant, and its magnitude implies that house price–induced local demand shocks account for roughly two-thirds of the substantial inflation differences across regions in our sample.

Our second identification strategy exploits variation in home ownership rates across zip codes. The same change in house prices will induce different real wealth and demand effects for homeowners and renters, since they differ in their net asset position in housing.[1] Consistent with

    [1] House price increases relax borrowing constraints and raise wealth for homeowners. Neither effect should be present for renters. Higher rents (either explicit or implicit rents when living in owner-occupied housing) affect both renters and homeowners the same way. Therefore, higher house prices raise wealth and credit access of homeowners relative to renters.

these differential demand effects, we show that there is a strong interaction between home ownership rates and the relationship between house prices and retail prices. In zip codes with a high home ownership rate, house price increases lead to the largest increases in retail prices, while in zip codes with the lowest home ownership rates, house price increases lead to (statistically insignificant) declines in retail prices.

Taken together, we believe that our two identification strategies provide compelling evidence for a causal effect of house price–induced local demand shocks on local retail prices, since it is difficult to jointly explain both results via confounding explanations such as local supply shocks. Our first empirical results instrument for changes in house prices using measures of the housing supply elasticity, and it is unclear why possible confounding supply-side shocks would be particularly strong in regions with lower housing supply elasticity. More importantly, any shock that might violate the exclusion restriction in our instrumental variables strategy would need to also vary with local home ownership rates, which dramatically narrows the list of potential concerns.

To provide further evidence for our causal interpretation of the observed relationship, we show that the relationship between house prices and retail prices survives an extensive set of robustness checks. In particular, we document that our results are not driven by changes in store or product quality, changes in income or gentrification patterns, differences in the employment mix across locations, or store entry and exit. We also show that our results hold with coast and region fixed effects, so that our instrumental variables results are not driven by a spurious correlation between supply elasticity and region-specific shocks. Finally, our results hold when dropping the "sand states" that saw the largest housing bubbles as well as other outliers and so are not driven by unusual observations.

After arguing for a causal relationship between house prices and retail prices, we next consider why increases in house prices lead to higher retail prices. By definition, an increase in retail prices must be driven by either an increase in markups or an increase in marginal costs. While we believe that identifying either channel would be interesting, we provide several pieces of evidence that support markup variation as the primary explanation for our empirical patterns.

First, our retail price data include only tradable goods in grocery stores and drugstores. These goods are not produced locally, so their wholesale cost should be independent of any local shocks. Indeed, we exploit a novel data set containing Universal Product Code (UPC) level wholesale costs from around the United States to show that geographic variation in wholesale costs is very small. Since wholesale costs represent nearly three-quarters of total costs and an even larger fraction of marginal costs in our stores, it is thus unlikely that geographic variation in marginal costs drives our retail price patterns. To provide additional support for this argument, we

supplement our primary analysis using data from a large national retailer, which include measures of both marginal costs and markups. We use these high-quality internal cost measures to directly show that this retail chain raises markups in locations with increasing house prices. Again, this house price effect on markups is strongest where home ownership rates are high.

While wholesale costs are the primary component of our retailers' marginal costs and do not drive our retail price patterns, we next directly consider two additional cost channels that might affect retail prices: local labor costs might rise in response to increased local demand, or local retail rents may rise.

Since labor costs are a small fraction of overall marginal cost for the stores in our data, explaining retail price movements through this channel would require extremely large responses of local labor costs to local demand. Consistent with this channel not being important, we find that controlling for local wages and a variety of other labor market conditions does not change our estimates.

Next, we provide evidence that our retail price results do not reflect a pass-through of local retail rents or land prices. First, and most importantly, pass-through of these costs cannot explain the fact that retail prices rise much more quickly with house prices in locations with high home ownership rates. If the relationship between retail prices and house prices was driven by pass-through of local land prices or rents, then local home ownership rates should instead be irrelevant. Second, we match our data with information on local retail rents and find that they do not affect our estimates. Finally, we exclude stores in high-rent locations from our analysis (since they should have the highest fraction of rent in total costs) and obtain near-identical estimates.

Together, wholesale inventory costs, labor costs, and rent overhead represent essentially 100 percent of marginal costs for our retailers. Thus, if the variation in retail prices is not driven by variation in these costs, it must be driven by variation in markups.

Why would firms raise markups in response to positive house price shocks, and more so in regions with more homeowners? In the final empirical section of our paper, we document that positive housing wealth effects lead households to become less price sensitive, prompting firms to increase their "natural markups" as the elasticity of demand falls.[2] We use data on individual household shopping behavior from Nielsen Homescan to show that when house prices rise, homeowners increase their nominal spending but purchase fewer goods with a coupon and reduce the fraction of spending on generics and on items that are on sale. In contrast,

---

[2] Following the New Keynesian literature, we refer to the markup under fully flexible prices as the natural markup.

the behavior of renters is essentially unchanged; if anything, renters spend less and become more price sensitive as house prices increase.

Why would homeowners become less price sensitive as house prices increase? If homeowners face the possibility of moving in the future, increases in local house prices directly lead to increases in expected wealth. In many models in which the value of leisure rises with wealth, as households become wealthier, they will allocate less time to shopping for lower prices and thus become less price sensitive (see the discussions in Alessandria [2009], Aguiar, Hurst, and Karabarbounis [2013], Huo and Ríos-Rull [2015], and Kaplan and Menzio [2016]). Similarly, in the presence of home equity extraction, higher house prices can alleviate credit constraints, relaxing homeowners' budget constraint and similarly reducing their price sensitivity (see Berger et al. 2018). Since house price changes have different wealth and collateral effects for homeowners and renters, this naturally explains the observed difference in shopping responses and rationalizes our earlier findings that retail price responses to house price changes depend on local home ownership rates.

Taken together, our empirical results provide evidence of an important link between changes in household wealth, shopping behavior, and firm price setting. Positive shocks to wealth cause households to become less price sensitive, and firms respond by raising markups and prices.

Implications.—Our results have direct implications for understanding the consequences of the recent housing boom and bust, which was central to the Great Recession. We show that in addition to the well-documented effects of house prices on spending (e.g., Mian and Sufi 2014a), there were important effects on local prices. Indeed, our evidence implies that part of the variation in local spending captures price variation rather than variation in real spending.

Even though our results are most directly informative about business cycle movements related to housing, we believe they provide useful insights for understanding business cycles in general. While recessions can be caused by many factors, as long as they lead to more price-sensitive shopping behavior, then the mechanisms we identify will apply.[3] There

---

[3] Our empirical evidence of a strong interaction between household shopping behavior and firm price setting supports the theoretical models of Huo and Ríos-Rull (2013) and Kaplan and Menzio (2016), who argue that such interactions can give rise to recessions. Our time-series analysis also complements cross-sectional analyses of relationships between household characteristics and firm price setting. For example, Handbury (2012) estimates nonhomothetic price indices that vary with household wealth in the cross section, and Manova and Zhang (2012) show that exporters set higher prices in wealthier product markets. However, the forces driving these long-run, static relationships between wealth and prices might be irrelevant for understanding the effects of business cycle fluctuations. For example, permanent differences in tastes could explain the static relationships but would not generate the changes across time in individual household behavior that we document.

is indeed growing empirical evidence for this type of shopping response during recessions: Aguiar et al. (2013) show that time spent on shopping increased during recessions, and Krueger and Mueller (2010) and Nevo and Wong (2019) show that many measures of shopping intensity rose during the Great Recession. Dube, Hitsch, and Rossi (2018) find that spending shares on generic goods fall with household income and wealth in the Great Recession, although these particular effects are small. They also find that there are significant price responses of these goods to local economic conditions. The fact that our house price–induced demand shocks are large and unanticipated and generate the kinds of changes in household shopping behavior and demand elasticity observed in recessions contrasts our approach with existing micro studies of demand shocks such as Warner and Barsky (1995), Chevalier, Kashyap, and Rossi (2003), Gicheva, Hastings, and Villas-Boas (2010), and Gagnon and Lopez-Salido (2014). These papers study responses to predictable seasonal holidays, changes in gasoline prices, store strikes, and temporary weather events. While these demand shocks are interesting in their own right, it is less clear whether they are informative for understanding business cycles.

To our knowledge, Coibion, Gorodnichenko, and Hong (2015) are the first to analyze geographic variation in price setting to inform aggregate business cycles. They use the same scanner data as we do to find that prices respond little to local unemployment rates. Beraja, Hurst, and Ospina (2019) use a broader set of scanner data that are available only beginning in 2006 and find a larger response of prices to unemployment rates. Our focus on exogenous changes in house prices allows us to isolate demand shocks, while local unemployment rates reflect a combination of local supply and demand factors.[4] We also jointly analyze household shopping behavior and firm price setting to argue that the relationship between house prices and retail prices reflects markup variation driven by changes in households' price sensitivity.

This type of markup variation has significant implications for business cycle modeling. In New Keynesian models, increases in demand drive up nominal marginal costs, and sticky prices mean that average markups fall and real economic activity rises. In the simplest versions of these models, flexible price natural markups are constant so that if pricing frictions are removed, then actual, realized markups are also constant. This means that all variation in realized markups in these models is driven by sticky prices, which induce a "gap" between the natural and the realized markup. Our

---

[4] The conflicting findings of Coibion et al. (2015) and Beraja et al. (2019) could reflect time-varying confounding shocks, since supply and demand shocks imply opposite correlations between prices and unemployment. In addition, even large increases in unemployment affect only a small part of the population directly, which reduces their econometric power for identifying demand effects. In contrast, house price changes affect many more households.

results suggest that even with no pricing frictions, realized markups can change for a second and complementary reason: countercyclical household shopping intensity leads to higher natural markups in booms.

It is important to note that a procyclical natural markup need not imply a procyclical realized markup if markup gaps are countercyclical, but it does suggest that modeling the endogenous interaction between household shopping intensity and firm pricing behavior might improve our understanding of the monetary transmission mechanism. Indeed, medium-scale dynamic stochastic general equilibrium models such as Smets and Wouters (2007), Christiano, Motto, and Rostagno (2010), and Justiniano, Primiceri, and Tambalotti (2011) introduce natural markup ("cost-push") shocks to match aggregate time-series data. One can interpret our results as a potential microfoundation for such shocks. However, in these models, changes in natural markups are treated as exogenous policy-invariant shocks. In contrast, our evidence suggests that natural markups will respond endogenously to changes in monetary policy.

The conclusion that realized markups vary for reasons besides sticky prices also complicates the interpretation of the large literature using aggregate time-series data to measure markup cyclicality (e.g., Domowitz, Hubbard, and Petersen 1986; Bils 1987; Rotemberg and Woodford 1999; Gali, Gertler, and Lopez-Salido 2007; Nekarda and Ramey 2013).[5] These papers measure movements in realized markups and often interpret their results as tests of New Keynesian mechanisms. However, if natural markups are procyclical while sticky price–induced markup gaps are countercyclical, then the realized markup measured in the data will depend on the relative strength of these two forces, and tests assuming a constant natural markup will likely be misspecified. Moreover, if the relative strength of the two forces varies over time (see Vavra 2014), then this can potentially reconcile conflicting conclusions about the importance of price stickiness in explaining markup variation in the literature.

In addition to contributing a new source of identification to a long-running empirical debate in macroeconomics, which has typically relied on vector autoregression analyses of aggregate time-series relationships, our results also have a wide range of implications that stretch beyond macroeconomics. For example, the response of local prices to local house price movements is central to many models in urban economics and for understanding the incidence of local labor market shocks. Our paper directly informs this important and previously unobserved parameter. We also provide novel evidence on the industrial organization of wholesale retailer relationships and their interactions with local economic conditions. Many recent papers have argued that pricing decisions are largely

---

[5] Analogous concerns also apply to the voluminous literature focusing on testing New Keynesian Phillips curves.

made at the wholesale rather than the retail level (e.g., Nakamura and Zerom 2010). However, this literature has focused on the pass-through of *cost* shocks. Our empirical evidence shows that the exact opposite conclusion arises when studying the price response to *demand* shocks. This implies that fully understanding the behavior of prices requires a comprehensive analysis of the interaction of retailers and their suppliers. Retailers are crucial in determining how prices and markups respond to changes in customer demand, while wholesalers play a crucial role in transmitting upstream cost shocks into final prices. These and other implications are developed in more detail in online appendix D, which also provides a more extensive discussion of the primary implications of our paper for business cycle modeling.

I. Data Description

To conduct the empirical analysis we combine a number of data sets. We begin by describing the construction of our key dependent variables and then detail the sources for our other data.

A. Retail Price Data: IRI Data

Our primary retail price data set covers various grocery stores and drugstores from 2001 to 2011 and is provided by IRI Worldwide.[6] The data set includes store-week-UPC sales and quantity information for products in 31 categories, representing roughly 15 percent of household spending in the Consumer Expenditure Survey.[7] We also obtained the zip code location of each store in the data from IRI Worldwide. These zip code identifiers are not part of the standard academic data release, and we believe we are the first to exploit them.[8] In total, the data cover around 7,200 stores in over 2,400 zip codes. There are many retailers in each metropolitan area. For example, the Chicago market contains observations from 131 unique retailers. Appendix figure A1 shows the geographic distribution of the stores in our data.

---

[6] See https://www.iriworldwide.com/en-US/solutions/Academic-Data-Set for information on accessing the data and Bronnenberg, Kruger, and Mela (2008) for additional description.

[7] These categories cover mostly processed food and beverages, cleaning products, and personal hygiene products and so are most similar to the Bureau of Labor Statistics (BLS) food-at-home price index. UPCs consist of 12 numeric digits that are uniquely assigned to each trade item; different packaging and sizes of the same item will be assigned unique UPCs.

[8] The standard academic data release includes geographic indicators for only 47 broad geographic markets, often covering a major metropolitan area (e.g., Chicago), but sometimes covering regions with numerous metropolitan statistical areas (MSAs; e.g., New England).

While the raw data are sampled weekly, we construct quarterly price indices, since this reduces high-frequency noise and makes the time unit comparable to that of our house price measures. Let t index the quarter of observation, l a geographic location (MSA or zip code), c a product category, and i an individual UPC-store pair (henceforth item).[9] We construct the price of an item by dividing its total dollar value of sales (TS) by the total quantity of units sold (TQ). That is, $P_{i,l,c,t} = TS_{i,l,c,t}/TQ_{i,l,c,t}$.

Total sales are inclusive of retailer discounts and promotions but exclude manufacturer coupons. In our benchmark specification, we include all observed prices when constructing our price indices, since we are interested in how the broadest price aggregate responds to local demand. We later show the robustness of our results to using price indices constructed when excluding "sales" prices.

We next describe the construction of our location-specific price indices from these individual price observations. This construction necessarily entails various measurement choices. In the body of the paper we concentrate on a single benchmark price index, but in appendix C we show that our empirical results continue to hold for price indices constructed under various alternative assumptions.

Since we are interested in constructing price indices across time, we include an item only if it has positive sales in consecutive quarters. After constructing item-level prices, we create location-specific price indices using a procedure similar to the construction of the Consumer Price Index (CPI) by the BLS. In particular, we construct a geometric weight price index with a consumption basket that is chained annually.[10] Let

$$\omega_{i,l,c,y(t)} = \frac{TS_{i,l,c,y(t)}}{\sum_{i \in c} TS_{i,l,c,y(t)}}$$

be an item's share in a category's annual revenue, where $y(t)$ indexes the year in which quarter t is observed. In our benchmark results, we construct these revenue weights separately for each location to allow for spatial variation in item importance. That is, ω is indexed by l. In appendix C, we also redo our analysis using national revenue weights, so that ω is no longer indexed by l, and using constant geographic weights, so that ω is no longer indexed by t. Under these alternative constructions, location-specific changes in household purchases, in product composition, or in product quality do not affect location-specific price indices. Our findings are robust to these alternative weights, which implies that

[9] We track the price of identical items (UPC-store pairs) across time, so that changes in quality or issues with comparing nonidentical products are not relevant for our results (quality changes across time will typically be associated with new UPCs). In particular, our price index is not affected by changes in the composition of goods or stores over time.

[10] We chain our results annually rather than at higher frequencies to avoid "chain drift."

the retail price responses we document require actual changes in price posting behavior and cannot be explained by shifting weights.

We construct our price index in two steps. We first construct a category-level price index:

$$\frac{P_{l,c,t+1}}{P_{l,c,t}} = \prod_i \left(\frac{P_{i,l,c,t+1}}{P_{i,l,c,t}}\right)^{\omega_{i,l,c,y(t)}}.$$

We then construct an overall location-specific price index by weighting these category price indices by the revenue share of a particular category,

$$\omega_{l,c,y(t)} = \frac{\sum_{i \in c} TS_{i,l,c,y(t)}}{\sum_i TS_{i,l,y(t)}}:$$

$$\frac{P_{l,t+1}}{P_{l,t}} = \prod_c \left(\frac{P_{l,c,t+1}}{P_{l,c,t}}\right)^{\omega_{l,c,y(t)}}.$$

Panel A of figure 1 shows that a nationally aggregated version of our price index reproduces the behavior of the BLS food-at-home CPI.[11] While they do not match precisely, this is not surprising, since the categories and products sampled are not identical. The BLS also produces food-at-home CPIs for 27 metro areas, of which 19 overlap with locations in the IRI data. Panel B of figure 1 compares changes in our MSA-level price indices to changes in these BLS indices. Again, there is a strong correlation between changes in our MSA price indices and those published by the BLS. The relationship is not perfect, but this is even less surprising for these disaggregated indices.[12] This figure also shows large variation across MSAs in retail price movements.

Finally, panel C of figure 1 shows that the cross-sectional variation in the food-at-home CPI produced by the BLS is very similar to the cross-sectional variation in the broader CPI including all products. This sug-

[11] We normalize indices to 1 in period $t = 0$, so our computation requires information on how prices change across time but not information on how products are priced across locations at a point in time. The reason is that we are interested in responses to demand shocks at business cycle frequencies, not in permanent price differences across locations. This specification in changes also allows us to avoid the biases discussed in Handbury and Weinstein (2015). Nevertheless, it is worth noting that methodological differences make our time-series results perfectly consistent with their conclusion that retail price levels do not vary across locations. Their paper uses household-level data with controls for income as well as retailer fixed effects. This specification absorbs most of our variation of interest since any correlation between retail prices and house price levels that is related to household demographics or retailer location will be absorbed by their controls.

[12] For most MSAs, the increase in the CPI is modestly larger than the increase in the IRI index since we use a chained index while the BLS uses a fixed basket. Sampling error is also less of a concern in IRI data since they have 20 times more observations per MSA-quarter than the BLS and cover substantially more markets.
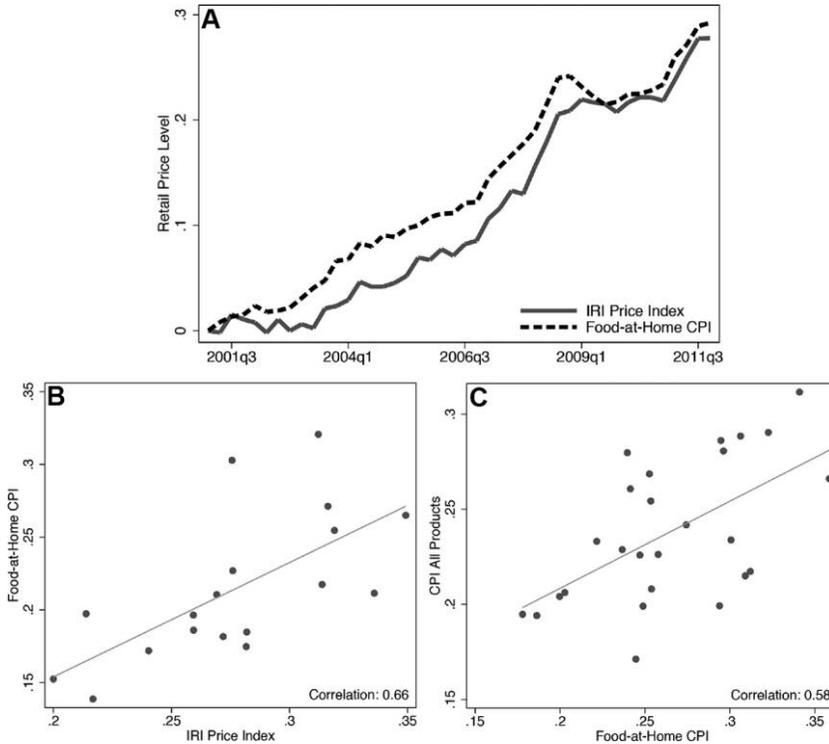
FIG. 1.—Price index versus BLS. *A*, IRI data versus food-at-home CPI. *B*, Metro-level comparison of IRI versus food-at-home CPI: $\log(P_{2011}) - \log(P_{2001})$. *C*, Metro-level comparison of food-at-home CPI versus CPI all products: $\log(P_{2011}) - \log(P_{2001})$. The figure shows comparisons of our price indices produced with IRI data to price indices provided by the BLS. Panel A compares our aggregate price index to the food-at-home CPI. Panel B compares the change in prices between 2001 and 2011 using our local price indices to the change in the metro area food-at-home price indices provided by the BLS for the set of MSAs for which we have overlapping data. Panel C compares the change in prices between 2001 and 2011 of metro area food-at-home prices to the change in "all product" prices from the BLS.

gests that the retail price responses to house prices that we document are likely to generalize to a broader set of goods than that covered by our IRI data.

B.   *Wholesale Cost Data*

We use the IRI data set as our primary measure of retail prices, since it covers many chains and has large geographic coverage. Unfortunately, IRI does not collect data on wholesale costs. Since the second half of our paper focuses on decomposing price changes into changes in markups and marginal cost, we supplement our primary IRI data with two addi-

tional data sets that do contain measures of wholesale costs. First, we use Nielsen PromoData, which collects information from one confidential grocery wholesaler in each of 32 markets for the period 2006–12.[13] This data set contains UPC-level wholesale prices for each date in each market. Overall there are more than 34,000 UPCs in PromoData.

While IRI data contain retail prices and not wholesale costs, PromoData contain wholesale costs and not retail prices. Since we are ultimately interested in measuring markups, we supplement our analysis with one additional data set from a large retail chain, which contains reliable measures for both prices and costs. This retailer reports UPC-store-level information for more than 125,000 UPCs from 250 stores in 39 MSAs, from January 2004 to June 2007. For each item, there is information on wholesale costs, adjusted gross profits, gross prices (i.e., list prices), and net prices (i.e., list prices net of rebates, promotions, and coupons). We follow Gopinath et al. (2011) and measure the marginal cost for each item as the difference between net price and adjusted gross profit. This represents the retailer's cost, net of discounts and inclusive of shipping costs, and is the cost measure used in their pricing decisions (see Eichenbaum, Jaimovich, and Rebelo 2011). Using these data, we construct a price, marginal cost, and markup index for each location, using the approach described in Section I.A.[14]

C. Shopping Data

We use Homescan data from AC Nielsen to measure household-level shopping behavior.[15] The data set contains a weekly household-level panel for the period 2004–11. The panel has large coverage, with 125,000 households in over 20,000 zip codes recording prices for 400 million unique transactions. The product coverage is somewhat broader than that in the IRI data and essentially captures broad nonservice retail spending. Roughly half of expenditures are in grocery stores, a third of expenditures are in discount/warehouse club stores, and the remaining expenditures are split among smaller categories such as pet stores, liquor stores, and electronics stores.

---

[13] There is a single wholesaler for each market but potentially different wholesalers across markets. Because of the anonymous nature of the wholesaler information, there is no way to identify which wholesalers are related across markets. While the raw data set contains information on more than 32 markets, many of these have very little data, so we concentrate on the largest 32 markets. Results are not sensitive to this restriction.

[14] This cost measure does not include wages or other fixed overhead. We address this explicitly in our empirical analysis. We concentrate on prices inclusive of discounts, but results are similar using gross prices.

[15] These data are available for academic research through a partnership with the Kilts Center at the University of Chicago, Booth School of Business. See http://research .chicagobooth.edu/nielsen for more details on the data.

Households report detailed information about their shopping trips using a bar code scanning device provided by Nielsen. After a shopping trip, households enter information including the date and store location. They then scan the bar codes of all purchased items. The price is collected in one of two ways: for trips to stores that partner with Nielsen, the average price of the UPC for that store-week is automatically recorded. For trips to stores not partnered with Nielsen, households hand-enter the price paid from their receipt. In addition to the price, households also record whether a product was purchased while "on sale" or using a coupon. In addition, since we know the UPC of each item, information is available on whether a product is generic or name brand. We use this information to construct quarterly expenditure shares for goods purchased in these categories for each household.

While panelists are not paid, Nielsen provides incentives such as sweepstakes to elicit accurate reporting and reduce panel attrition. Projection weights are provided to make the sample representative of the overall US population.[16] A broad set of demographic information is collected, including age, education, employment, marital status, and type of residence. Nielsen maintains a purchasing threshold that must be met over a 12-month period in order to eliminate households that report only a small fraction of their expenditures. The annual attrition rate of panelists is roughly 20 percent, and new households are regularly added to the sample to replace exiting households.

D. Other Data

We also use a number of other data sets in our analysis. We obtain house price indices at both the zip code level and the MSA level from CoreLogic, which computes repeat sales price indices from individual transactions data. We also use information on average effective retail rents from 2000–2014 for 45 MSAs. These rent data are compiled by the REIS Corporation from surveys of property managers and leasing agents and include quarterly information on the average rent paid per square foot of retail space. Home ownership rates by zip code come from the 2000 census. Data on education levels, age, and population density come from the respective waves of the American Community Survey (ACS). We obtain wage data from the Quarterly Census of Employment and Wages conducted by the BLS. Employment shares and information on the number of retail establishments come from the County Business Patterns produced by the US Census Bureau, and we classify North American Industry Classification

---

[16] We use these projection weights in all reported results, but our results are similar when weighting households equally.

System sectors into tradable and construction using the definitions in Mian and Sufi (2014b).

II.   Empirical Analysis

We next provide an overview of our empirical strategy for identifying the impact of house price changes on retail prices. We use two complementary identification strategies to show that our relationship is causal and that house price–induced demand shocks drive changes in retail prices.

Our first approach uses across-MSA variation in housing supply elasticity as an instrument for house price changes. This approach isolates differences in house price growth that are plausibly orthogonal to other factors that might directly influence retail prices.

Our second approach exploits a unique feature of house price changes to provide additional evidence that they causally influence retail price. In particular, house price changes induce differential wealth effects for homeowners and renters due to these households' different net housing positions. Consistent with this, we show that the relationship between house prices and retail prices depends strongly on local home ownership rates.

The use of these two complementary identification strategies substantially reduces the set of confounding explanations for our results, since geographic variation in home ownership rates is quite distinct from geographic variation in housing supply elasticity. In particular, alternative stories must explain not just why housing supply elasticity would not satisfy the instrumental variables' exclusion restriction, but also why such violations would then interact with local home ownership rates.

In addition to documenting a causal link from house prices to retail prices, we provide evidence on the economic mechanism driving this relationship. In general, an increase in retail prices must reflect an increase in marginal costs or an increase in markups. We argue that our results primarily reflect markup movements by first showing that our patterns are not driven by changes in observable costs. We then present direct evidence that households become less price sensitive after their housing wealth rises; this increases firms' optimal markups. Just as suggested by our retail price results, we show that this change in household price sensitivity differs strongly by home ownership status.

A.   Price-Setting Behavior: MSA Level

We first analyze the relationship between house prices and retail prices. We split the sample into the periods 2001–6, when house prices in the United States were generally rising, and 2007–11, when house prices were generally falling. This allows for an asymmetric impact of house price increases and decreases on retail prices. We begin by sorting MSAs

into quintiles by their house price growth over the housing boom and housing bust. The top row of figure 2 shows how retail prices evolve for MSAs in the top and bottom quintiles of house price growth over each period. Clearly, retail price growth was significantly stronger in those MSAs that experienced higher house price growth.[17]

The middle row of figure 2 shows the more disaggregated correlation between MSA-level house price growth and retail price growth over the periods 2001–6 (panel C) and 2007–11 (panel D). In both periods there is a strong positive correlation between house price growth and retail price growth. This positive relationship is confirmed by ordinary least squares (OLS) regressions of retail price changes on house price changes, shown in column 1 of table 1. Appendix table A1 provides summary statistics on the dependent variable and controls. The estimated coefficient suggests an elasticity of retail prices to house prices of about 6–8 percent.[18] In column 2, we also include controls for changes in economic conditions such as the unemployment rate, wages, and employment shares in the grocery, retail, construction, and nontradable sectors. The estimated elasticity of retail prices to house prices is unaffected.[19]

However, even after the inclusion of control variables, these estimates do not establish causality, since there might be an unobserved third factor, such as time-varying productivity, that could simultaneously move both house prices and retail prices. If we cannot directly control for this third factor in the OLS regression, we will obtain a biased estimate of the elasticity of retail prices to house prices.

B. *Price-Setting Behavior: Instrumental Variables Identification Strategy*

Our first approach to dealing with possible omitted variables bias is to exploit instrumental variables that are correlated with house price changes over our periods of interest but that do not directly affect retail price

---

[17] While the difference in retail prices between high and low house price–growth MSAs during the bust is smaller than during the boom, the elasticity is higher because the difference in house price changes is smaller in the bust. In addition, sorting over 2001–11 house price growth rather than separately over the boom and the bust produces similar patterns.

[18] Figure 2 shows that there are some outliers in house price changes that might affect our estimates. Repeating regressions excluding the top and bottom 5 percent of MSAs by house price growth raises elasticities to 8–9 percent, and equality with the instrumental variable (IV) coefficients shown below cannot be rejected.

[19] The positive coefficient on unemployment rate changes might appear to conflict with our argument that house prices move as a result of a markup response to demand shocks. However, this coefficient is estimated after controlling for house price changes, which were the primary driver of cross-regional differences in changes in the unemployment rate in our sample. Indeed, when we regress retail price changes on changes in the unemployment rate without also controlling for house price changes, we recover a negative relationship. See also n. 4.
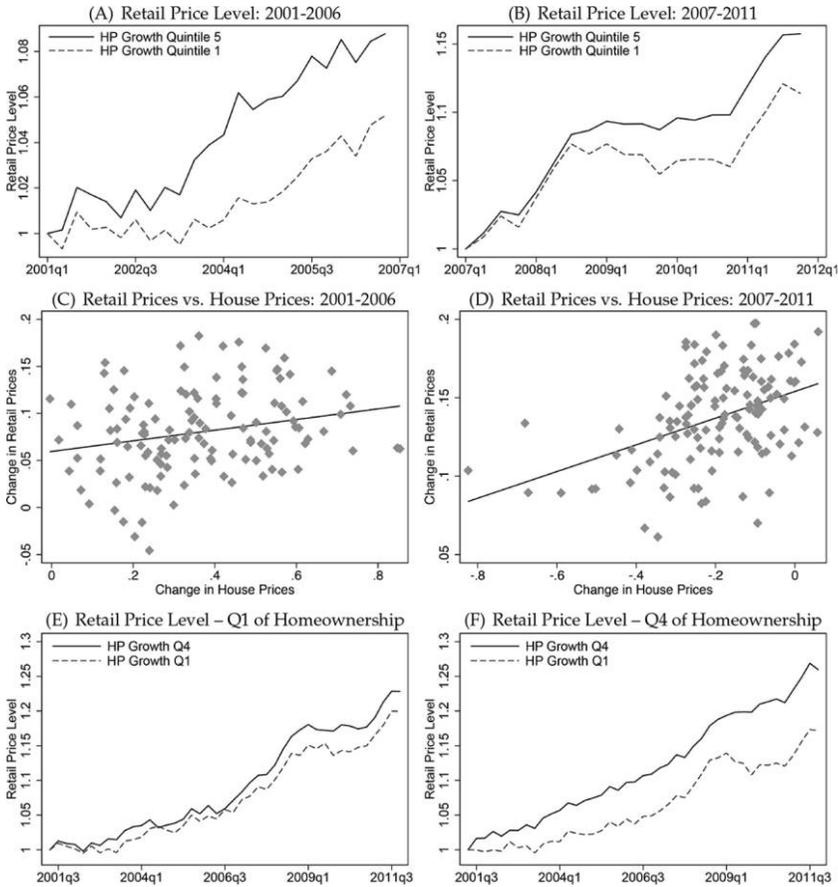
FIG. 2.—Retail prices versus house prices. The top row shows the average retail price level over time for MSAs in the top quintile (solid line) and bottom quintile (dashed line) of house price appreciation for the period 2001–6 (panel A) and the period 2007–11 (panel B). The middle row shows the MSA-level correlation between changes in house prices and changes in retail prices for the period 2001–6 (panel C) and the period 2007–11 (panel D), as well as the line of best fit. The bottom row shows the average retail price level over time for zip codes in the top quartile (solid line) and bottom quartile (dashed line) of house price appreciation between 2001 and 2011. Panel E shows results of zip codes in the bottom quartile of the home ownership rate distribution. Panel F shows results of zip codes in the top quartile of the home ownership rate distribution.

growth. In particular, we follow an extensive literature that exploits across-MSA variation in housing supply elasticity to instrument for house price changes (see, e.g., Mian and Sufi 2011, 2014a; Adelino, Schoar, and Severino 2015; Giroud and Mueller 2017). The intuition for this instrument is that for a fixed housing demand shock during the housing boom, house prices should rise more in areas where housing supply is less

TABLE 1
RETAIL PRICES VERSUS HOUSE PRICES: MSA-LEVEL ANALYSIS

	DEPENDENT VARIABLE: Δ RETAIL PRICES					
	OLS		IV Saiz		IV Wharton	
	(1)	(2)	(3)	(4)	(5)	(6)
A. Time Period: 2001–6						
Δ house prices	.057***	.068***	.129***	.153***	.224***	.230***
	(.020)	(.023)	(.042)	(.058)	(.048)	(.048)
Δ share grocery retail employment		−.068		.132		.219
		(.360)		(.376)		(.391)
Δ share nontradable employment		.073		−.082		−.130
		(.182)		(.175)		(.174)
Δ share construction employment		−.060		.012		.039
		(.098)		(.114)		(.130)
Δ unemployment		.039**		.070**		.095***
		(.018)		(.029)		(.026)
Δ wage		.039		.038		−.005
		(.055)		(.060)		(.061)
Observations	125	125	112	112	112	112
B. Time Period: 2007–11						
Δ house prices	.085***	.086***	.124***	.146***	.147***	.157***
	(.015)	(.018)	(.041)	(.049)	(.048)	(.043)
Δ share grocery retail employment		−.090		.008		−.000
		(.264)		(.275)		(.282)
Δ share nontradable employment		.086		−.000		−.003
		(.139)		(.169)		(.172)
Δ share construction employment		.050		−.024		−.040
		(.127)		(.135)		(.147)
Δ unemployment		.000		.017		.019
		(.011)		(.015)		(.014)
Δ wage		−.030		−.060		−.063
		(.044)		(.048)		(.049)
Observations	126	126	112	112	112	112

NOTE.—The table shows results from the following OLS regression: $\Delta \log (\text{RetailPrice})_m = \beta \Delta \log (\text{HousePrice})_m + \gamma X_m + \varepsilon_m$ in cols. 1 and 2 and from IV regression (2) in the other columns. All specifications include a constant that is not reported. The unit of observation is an MSA, and the dependent variable is the change in retail prices in 2001–6 (panel A) and 2007–11 (panel B). We instrument for the change in house prices using measures of the housing supply elasticity provided by Saiz (2010) in cols. 3 and 4 and the Wharton Regulation Index described in Gyourko et al. (2008) in cols. 5 and 6. Robust standard errors are in parentheses.
* Significant at $p < .10$.
** Significant at $p < .05$.
*** Significant at $p < .01$.

elastic.[20] During the housing bust, it is then precisely those areas where house prices rose most that see the largest declines in house prices (Glaeser 2013).

We use two measures of housing supply elasticity as instruments for house price changes: the primarily geography-based measure of Saiz (2010) and the Wharton Regulation Index from Gyourko et al. (2008). Saiz (2010) uses information on the geography of a metropolitan area to measure the ease with which new housing can be constructed. The index assigns a high elasticity to areas with a flat topology without many water bodies, such as lakes and oceans. Gyourko et al. (2008) conduct a nationwide survey to construct a measure of local regulatory environments pertaining to land use or housing. Their index aggregates information on who can approve or veto zoning requests, and particulars of local land use regulation, such as the review time for project changes. In areas with a tighter regulatory environment, the housing supply can be expanded less easily in response to a demand shock, and house prices should therefore increase more. Appendix table A2 presents results from the first-stage regression (1). Both instruments are highly predictive of house price changes over both periods, with low-elasticity MSAs experiencing larger house price gains during the housing boom and larger house price drops during the housing bust.[21]

The exclusion restriction requires that housing supply elasticity affects retail prices only through its impact on house prices (see app. A for a formal statement of the exclusion restriction). To provide some evidence for the validity of the Saiz (2010) instrument, Mian and Sufi (2011, 2014a) show that wage growth did not accelerate differentially in elastic and inelastic core-based statistical areas between 2002 and 2006. The authors also show that the instrument is uncorrelated with the 2006 employment share in construction, construction employment growth in the period 2002–5, and population growth over the same period. Consistent with this, we find no relationship between housing supply elasticity and income growth in our sample: during the housing boom, income growth has a correlation of .040 with the Saiz (2010) instrument and −.007 with

---

[20] Importantly, this instrument generates differential regional house price movements through a differential propagation of a national housing demand shock, not through inducing a shock of its own. That is, it is the interaction of supply elasticity with the national shock that allows a time-invariant instrument to predict time-varying effects. This national demand shock could, e.g., result from the relaxation in down payment requirements or a decline in interest rates. See app. A for additional discussion and app. B for some alternative, related empirical specifications.

[21] Unsurprisingly, the power of the instrument is significantly stronger during the housing boom than during the housing bust. The first-stage F-statistics of the Saiz (2010) instrument are 42.4 for 2001–6 and 19.9 for 2007–11. They are 52.5 and 17.2, respectively, for the Gyourko et al. (2008) instrument. Supply elasticity has predictive power during the bust because it reflects ex post unraveling of the differential house price bubble.

the Wharton Regulation Index. These correlations are $-.224$ and $.054$, respectively, for the housing bust and are never statistically significant (see also Davidoff [2013] for a discussion of the exclusion restriction). It is also important to recall that the main objective in our paper is to document the response of retail prices and markups to shocks that affect demand elasticity. While we believe that the IV approach in this section and the home ownership interaction in the next section strongly point to a causal link from house price–induced local demand shocks to retail prices, it is worth noting that many potential violations of the exclusion restriction in our IV approach involve a correlation between housing supply elasticity and local demand factors. While we find no evidence for such a correlation, its presence would only mildly change our interpretation. In particular, while not all of the markup response would then represent a response to changes in house prices, it would still represent a markup response to demand shocks that shift demand elasticity. This would have the same implication for business cycle models as our preferred causal interpretation.

The exclusion restriction could be potentially violated if changes in the degree of local retail competition were correlated with the housing supply elasticity. This might occur if the regulatory or geographic environment hindered the entry of new retail stores. In Section II.D, we directly address this concern and show that differential changes in competition do not explain our results.

The first and second stages of the IV regression are given by equations (1) and (2), respectively:

$$\Delta \log(\text{HousePrice})_m = \rho \text{SupplyElasticity}_m + \delta X_m + \epsilon_m, \tag{1}$$

$$\Delta \log(\text{RetailPrice})_m = \beta \Delta \log (\widehat{\text{HousePrice}})_m + \gamma X_m + \varepsilon_m. \tag{2}$$

The unit of observation is an MSA, denoted by m. We estimate these regressions separately for the housing boom (2001–6) and bust (2007–11). The dependent variable in the second-stage regression is the change in retail prices over the period of interest. The coefficient of interest is β, which captures the causal effect of house price growth on retail price growth; X_m is a vector of controls.

We first present results using the housing supply elasticity from Saiz (2010) as an instrument for house price changes. Column 3 of table 1 presents estimates from the second-stage regression. The elasticity of retail prices to house prices is about 12–13 percent during both the housing boom and the housing bust. These elasticities are about two times as large as the estimates from the OLS regressions presented in columns 1 and 2. This is consistent with the presence of local productivity shocks, which

would lower retail prices but raise house prices. In addition, any measurement error in house price growth will also bias down the OLS estimates.[22] Column 4 includes control variables for local economic conditions. The robustness of the estimated coefficients to the addition of these controls, as well as additional controls for changes in income and demographics that will be added in Section II.D, helps to alleviate concerns about whether our instruments satisfy the exclusion restriction.

Column 5 and 6 of table 1 show the IV estimates using the Wharton Regulation Index as an alternative measure of housing supply elasticity to instrument for house price changes. The estimated elasticity of retail prices to house prices is slightly stronger, with estimates between 15 percent and 22 percent depending on the exact specification.

We next provide additional robustness checks to the results presented above. We first address the geographic clustering of our measures of housing supply elasticity, which raises concerns that they might be correlated with unobserved regional shocks. To show that such unobserved shocks do not explain our results, we add geographic controls to the IV regression. In columns 1–3 of appendix table A5, we add a coastal indicator, four census region fixed effects, and nine census division fixed effects, respectively. The estimated elasticity of retail prices to house prices is unchanged, suggesting that it is not explained by regional shocks. Next, while we believe that using the broadest price index available is the appropriate benchmark, a large literature has explored the implications of sales for monetary policy. Column 4 of table A5 shows that our results are robust to excluding temporary "sales" prices from the price index. We also want to ensure that our results are not driven by extreme outliers. Column 5 of table A5 excludes the MSAs with the largest and smallest 5 percent house price growth; column 6 drops observations from states that experienced some of the largest swings in house prices: California, Arizona, and Florida. Our results are robust across these specifications. Finally, one might worry that the relationship between our instruments and house price changes could be nonlinear. Column 7 includes the squared and cubed measures of supply elasticity as additional regressors in the first stage; the results are essentially unchanged.

---

[22] Whether one would expect the IV estimates to be larger or smaller than the OLS estimates depends on whether omitted variables or shocks would primarily induce a positive or a negative correlation between house prices and retail prices. As discussed, local productivity shocks might lower retail prices and increase house prices. On the other hand, local demand shocks might lead to higher house prices and higher retail prices. The IV estimates remove the effects of both types of potential omitted shocks, as well as the attenuation bias associated with measurement error. Indeed, if we exclude the largest outliers in house price growth, which are more likely to suffer from measurement error, the OLS coefficients increase substantially and we cannot reject equality with the IV regressions. See n. 18.

C.   Price-Setting Behavior: Zip Code–Level Identification Strategy

In the previous section we measured both house prices and retail prices at the MSA level. There are some advantages of these MSA-level estimates relative to estimates using house price and retail price measures at more disaggregated levels such as zip codes. First, nearly all grocery spending for a household should occur within MSAs, but this may not hold for zip codes. Second, both house price changes and retail price changes are measured more precisely for MSAs than for zip codes.[23] Third, our housing supply elasticity instruments do not vary at the zip code level. Therefore, we think the elasticities at the MSA level are the most reasonable to take away from our analysis.

Nevertheless, we now extend our analysis to the zip code level, because the large variation in home ownership rates across zip codes allows us to explore a separate, complementary identification strategy. In particular, the same change in house prices will induce different demand effects for homeowners and renters, since these households differ in their net asset position in housing. While house price increases can raise wealth or relax borrowing constraints for homeowners, they have no such effects on renters. If house prices are capitalized into apartment rents or renters plan to purchase in the future, then higher house prices might even represent negative wealth shocks for renters.[24] Thus, if the positive relationship between retail prices and house prices is truly driven by house price–induced demand effects, we would expect a stronger relationship in zip codes with high home ownership rates.

To explore this prediction, the bottom row of figure 2 shows the average retail price level for zip codes in the top and bottom quartiles of house price growth between 2001 and 2011. Panel E focuses on zip codes in the bottom quarter of the home ownership rate distribution (average of 46 percent) and panel F on zip codes in the top quarter of the home ownership rate distribution (average of 86 percent). Those zip codes with larger house price increases have higher retail price growth. However, as one would expect if house price effects work through a wealth channel, the differential price growth is much larger in zip codes with higher home ownership rates than it is in zip codes with lower home ownership rates.

---

[23] The repeat sales house price index at the zip code level is often based on a small number of sales: an analysis of transaction-level data for Los Angeles since 1994, e.g., reveals that in the median zip code–quarter there were 62 arm's-length housing transactions and many fewer transactions that were part of a repeat-sales pair.

[24] We analyze apartment rent data from REIS and find that the elasticity of cumulative apartment rent growth to cumulative house price growth over the housing boom was 0.34. The same elasticity was 0.10 over the housing bust (see fig. A2). The less-than-full pass-through of house price movements to rents, even over such relatively long horizons, is consistent with swings in the price-rent ratio over this period (see, e.g., Sinai 2013).

Regression (3) formalizes this insight. As before, we estimate this specification separately for the housing boom period and the housing bust period. Since we do not have housing supply measures at the zip code level, we focus on OLS estimates:

$$\Delta \log(\text{RetailPrice})_z = \beta \Delta \log(\text{HousePrice})_z + \gamma \text{HomeOwnershipRate}_z$$
$$+ \delta \Delta \log(\text{HousePrice})_z \times \text{HomeOwnershipRate}_z$$
$$+ \psi X_z + \varepsilon_z. \tag{3}$$

The results of this regression are presented in table 2. Columns 1 and 5 show the elasticity of retail prices to house prices without controlling for other covariates for the periods 2001–6 and 2007–11, respectively. The estimated elasticities are approximately 50 percent of the size of the MSA-level OLS estimates presented in table 1. As discussed above, this likely reflects attenuation bias relative to the MSA specifications, due to greater measurement error of zip code–level house prices, plus the fact that some fraction of household spending will occur outside of a household's zip code of residence. The addition of control variables in columns 2 and 6 has little effect on the estimated elasticities.

Importantly, columns 3 and 7 of table 2 interact house price changes with the home ownership rate in the zip code. The results show that house price increases are associated with particularly large increases in retail prices in zip codes with high home ownership rates.[25] For zip codes with low home ownership rates, the effect of higher house prices on retail prices is, if anything, negative, although this point estimate is not statistically significant.[26]

These results significantly strengthen the argument for a causal effect of house prices on retail prices. In particular, any omitted variables that might be correlated with our housing supply elasticity instruments in Section II.A, and that would thus violate the exclusion restriction, would also have to have a differential impact on homeowners and renters in order to explain our results.

One concern with the interpretation of the home ownership rate interaction is that home ownership rates might proxy for effects of other neighborhood characteristics. For example, high–home ownership zip codes have lower population density and so might have inhabitants who do

---

[25] The implicit identification assumption is that home ownership rates (conditional on controls) are orthogonal to shocks that simultaneously move house prices and retail prices. See below for further discussion of additional controls.

[26] One might worry that this relationship is driven by larger measurement error in house prices in areas with more renters, which could lead to larger attenuation bias. However, there is no strong relationship between home ownership rates and turnover: in zip codes in the bottom quartile of the home ownership distribution, about 2.1 percent of the housing stock turned over every year between June 2008 and March 2015; in zip codes in the top quartile of the home ownership rate distribution, this share was 2.2 percent. In addition, all results persist if we measure the change in house prices in regression (3) at the MSA level.

TABLE 2
RETAIL PRICES VERSUS HOUSE PRICES: ZIP CODE-LEVEL ANALYSIS

	Dependent Variable: Δ Retail Prices							
	Period: 2001–6				Period: 2007–11			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ house prices	.023***	.048***	−.045	−.170*	.040***	.030***	−.035	−.072
	(.007)	(.009)	(.032)	(.095)	(.007)	(.008)	(.033)	(.084)
Δ unemployment		.056***	.059***	.057***		−.015**	−.016**	−.014*
		(.012)	(.012)	(.013)		(.008)	(.008)	(.008)
Δ wage		.058**	.048	.047		.011	.006	.007
		(.029)	(.029)	(.029)		(.024)	(.025)	(.025)
Δ share grocery retail employment		−.230	−.252	−.241		.131	.119	.078
		(.300)	(.297)	(.295)		(.225)	(.222)	(.216)
Δ share nontradable employment		.080	.095	.085		.018	.026	.037
		(.123)	(.123)	(.122)		(.101)	(.101)	(.101)
Δ share construction employment		−.152**	−.177***	−.184***		.072	.078	.114
		(.065)	(.065)	(.071)		(.071)	(.071)	(.076)
Home ownership rate			−.063**	−.120***			.030	.021
			(.027)	(.046)			(.019)	(.030)
Δ house prices × home ownership rate			.142***	.222***			.095**	.123*
			(.047)	(.081)			(.047)	(.071)
Population density				.001				−.000
				(.002)				(.001)
Δ house prices × population density				−.002				.002
				(.003)				(.003)
Share below 35 years				−.002**				−.000
				(.001)				(.001)
Δ house prices × share below 35 years				.003*				.000
				(.002)				(.002)
Observations	708	708	708	708	846	846	846	846

NOTE.—The table shows results from regression (3). All specifications include a constant that is not reported. The unit of observation is a zip code, and the dependent variable is the change in retail prices in 2001–6 in cols. 1–4 and the change in retail prices in 2007–11 in cols. 5–8. Population density is measured in 1,000 inhabitants per square mile. Robust standard errors are in parentheses.

  * Significant at $p < .10$.
 ** Significant at $p < .05$.
*** Significant at $p < .01$.

more grocery shopping within the zip code. This could explain the larger measured response of local retail prices to local house prices in those areas, even with no differential wealth effects. Similarly, low–home ownership zip codes also have younger inhabitants, who might be less responsive to house price changes for reasons unrelated to home ownership status. To see whether these factors can explain our findings, columns 4 and 8 of table 2 include controls for the population density and the share of inhabitants under age 35, as well as their interaction with the change in house prices. Furthermore, table A3 additionally controls for zip code–level income, racial composition, and education levels and their interaction with house price changes. Reassuringly, the estimated coefficient on the interaction of house price changes and home ownership rates is, if anything, slightly larger in these specifications.

D.   Changes in Markups or Pass-Through of Changes in Marginal Cost?

The previous sections provide evidence of a strong response of retail prices to house price–induced demand shocks. By definition, a change in retail prices can be decomposed into a change in marginal costs and a change in markups. While either channel would be interesting, we next provide evidence that the relationship between house prices and retail prices is driven largely by markup variation.

We begin by decomposing retailers' marginal cost into its various components, before analyzing these components separately. For the typical grocery store, the wholesale cost of goods sold makes up approximately 75 percent of total costs.[27] It is more difficult to decompose the remaining 25 percent, but the majority of those costs represent fixed overheads (e.g., store rental costs, utilities, and corporate salaries) rather than costs that directly vary with sales. Thus, wholesale costs make up substantially more than 75 percent of all marginal costs. Since our data include only tradable goods, which are generally not produced locally, these wholesale costs are also likely to be insensitive to local demand shocks.[28]

---

[27] For example, Safeway's 2013 10-K reports a cost of goods sold of $26.6 billion, compared to operating and administrative expenses (which include store occupancy costs, wages, employee benefits, rent, depreciation, and utilities) of $8.9 billion. Walmart reported "cost of sales" of $385 billion, compared to "operating, selling, and administrative expenses" of $91.3 billion.

[28] In commodity flows survey data, only 24 percent of food and beverage shipments by gross value added are shipped less than 50 miles. However, net rather than gross value added is the relevant object for determining local marginal cost shares, and local distribution inflates gross relative to net value added in local production. The Bureau of Economic Analysis reports that trucking/warehousing has a 12.4 percent intermediate share in food and beverage, which, together with the 24 percent local share in gross value added, implies that less than 3 percent of inventory input costs for food and beverage stores are determined within an MSA.

Indeed, in the next section we use Nielsen PromoData to provide direct evidence that the geographic variation in wholesale costs is very small, consistent with wholesale costs not being determined locally. We also use data from a large national retailer that allow us to directly measure both marginal costs and markups in the same setting and show that the estimated price response captures movements in markups. In addition, our results are robust to excluding products from our price index construction that likely have a larger local component of wholesale costs (e.g., milk). We also directly control for changes in retail rents and labor costs that could affect retailers' (small) nonwholesale cost component of marginal cost. Our results are unaffected by the addition of these control variables.

We conclude that changes in local demand are largely uncorrelated with retailers' marginal cost, so that the increase in retail prices we observe mostly reflects higher markups. This interpretation arises naturally from the changes in shopping behavior we document in Section II.E and can also rationalize the homeownership interactions in Section II.C.

1. Direct Evidence on Wholesale Costs and Markups

We next provide three separate pieces of evidence that our results are not driven by variation in wholesale costs in response to local demand shocks. Since these wholesale costs make up the vast majority of all marginal costs, this provides strong support for an interpretation of our pricing results as driven by markup variation in response to these local demand shocks.

Nielsen PromoData.—We begin by considering the Nielsen PromoData described in Section I.B. We first examine the unconditional distribution of wholesale costs across markets. For each UPC observed in at least two markets, we compute

$$\widehat{w_{imt}} = \frac{w_{imt}}{\sum_{m \in M}(w_{imt}/|M|)},$$

where w_{imt} is the average wholesale price of item i in market m in quarter t, and M indexes the set of markets in which this item is sold in quarter t. That is, $\widehat{w_{imt}}$ tells us how far a particular item's wholesale cost in a given market is from that item's average cost across all markets. Finally, to measure overall variation in wholesale costs across markets, we compute the average of $\widehat{w_{imt}}$ across all items sold in market m in quarter t, I_{mt}, and then average this over all quarters in which we observe that market:[29]

---

[29] Computing these numbers quarter by quarter produces similar results. Note, also, that since PromoData contain only prices and not quantities, we weight all items equally.

$$\overline{w_m} = \sum_{t \in T} \frac{1}{|T|} \sum_{i \in I_m} \frac{\widehat{w_{imt}}}{|I_{mt}|}.$$

Overall, we find that 26 out of 32 markets have average wholesale costs within 1 percent of the national average: $\overline{w_m} < 0.01$. The market with the most expensive wholesale costs, Indianapolis, is 2.9 percent above the national average, while the market with the least expensive wholesale costs, Seattle, is 2.3 percent below the national average. In addition, we also calculate the fraction of individual items that are sold at the cross-market modal price in a given quarter. Overall, wholesale costs for 78 percent of items are exactly equal to the modal price, and more than 90 percent of items are within 5 percent of the modal price.

Why is there so little wholesale cost variation relative to retail price variation in the data? A full exploration of this point is beyond the scope of this paper and is not essential for our conclusions, but one potential explanation is that retailers have much better information on the elasticity of customer demand in their stores than wholesalers do. In a full-information environment, wholesalers could potentially charge higher prices to retailers with less elastic demand and higher markups, but retailers have no incentive to disclose this information to wholesalers, making such price discrimination difficult in practice. It is also important to note that this empirical finding of little wholesale cost variation is highly consistent with a legal restriction against such variation. In particular, the Robinson-Patman Anti–Price Discrimination Act explicitly limits the extent to which wholesalers can charge different prices to different retail customers.

These findings imply that, while wholesale costs are not perfectly equal across space, their spatial variation is very small. This is especially true once individual items are aggregated to the market averages relevant for our analysis. It is also important to note that even if wholesale costs did vary substantially across space, this would not necessarily alter our interpretation of whether house prices affect retail prices through moving markups or marginal costs. Indeed, our interpretation requires only the much weaker condition that wholesale costs not rise with house prices. However, PromoData has a limited time dimension with which to test this condition, and many markets are in the sample for only 1 or 2 years. In addition, we cannot compare any variation in observed wholesale costs to variation in retail prices or markups, since we do not observe this information in PromoData.

Large retailer data.—Thus, to provide further evidence that we are capturing changes in markups rather than the pass-through of marginal costs, we turn to the data from a large US retail chain described in Section I.B. These data include a complete measure of wholesale costs (inclusive of various vendor rebates and discounts), which the retailer uses when determining prices and thus markups.

We begin by repeating the above procedure to compute the across-MSA variation in wholesale costs in this data set. The range of MSA-level deviations of wholesale cost from the national average is between -1.4 percent and $+2.6$ percent, with the majority of markets clustered extremely close to the national average. This is in line with the findings from PromoData and is substantially less variation than we observe for markups. Indeed, when we compute MSA-level deviations of the markup from the national average, we find that the range of markup deviations is six times larger than that of wholesale cost deviations, with an across-MSA variance that is nearly 30 times larger.[30]

We next turn to a direct test of our question of interest: do changes in house prices lead to changes in marginal costs or changes in markups? As discussed, a law of one price for wholesale costs is sufficient but not necessary to interpret our retail price effects as markup changes. We instead require only the much weaker condition that marginal cost does not rise with house prices. To explore whether this condition holds, we use the data from our large retailer to construct zip code–level price, marginal cost, and markup indices from January 2004 to June 2007.[31] We then run regression (3) using changes in the price index, the markup index, and the marginal cost index as the dependent variables.

Table 3 presents the results. Column 1 shows that, on average, zip codes with higher house price growth see an increase in retail prices. Interestingly, despite the fact that these data cover only one retailer and a different time period, the estimated elasticity is similar to that in table 2, which was estimated using the broader IRI data. Importantly, columns 2 and 3 show that this increase in retail prices represents an increase in markups rather than a pass-through of marginal cost.[32]

Columns 4–6 of table 3 interact house price changes with the home ownership rate in the zip code. Columns 7–9 also control for the interaction of house price changes with demographic variables such as

---

[30] The retailer agreement does not allow us to report the actual levels of markups for specific stores. While we compute a statistic slightly different from that reported by Gopinath et al. (2011) and use all products instead of those matched between the United States and Canada, our results are consistent with their table 2. This shows that prices vary significantly more across stores than costs. We also reproduce their result that prices respond significantly to nonlocal costs after controlling for local costs. However, the economic magnitude of this effect is small. Adding nonlocal costs increases the regression R^2 from only .42 to .43, which is consistent with local costs capturing most information about non-local costs.

[31] Since our data contain information from only 39 MSAs with house price information, many with a single store, we do not repeat the MSA-level analysis.

[32] The negative house price coefficient on marginal cost likely reflects volume discounting: when house prices increase, the resulting increase in demand allows the retailer to take advantage of discounts from wholesalers. Since marginal cost in our primary specification includes vendor discounts and rebates, this can generate the observed negative relationship. If we instead use costs excluding discounts and promotions as the dependent variable, this negative coefficient disappears.

TABLE 3
ZIP CODE PRICING RESULTS: LARGE RETAILER

	Δ RP (1)	Δ Markups (2)	Δ MC (3)	Δ RP (4)	Δ Markups (5)	Δ MC (6)	Δ RP (7)	Δ Markups (8)	Δ MC (9)
Δ house prices	.018**	.039***	−.022***	−.065*	−.032	−.035	−.093	−.132*	.041
	(.008)	(.007)	(.007)	(.036)	(.031)	(.033)	(.083)	(.077)	(.080)
Δ unemployment	.029**	.044***	−.014*	.028*	.043**	−.014*	.026*	.045***	−.018**
	(.014)	(.017)	(.007)	(.014)	(.017)	(.008)	(.014)	(.017)	(.008)
Δ wage	.049**	.058***	−.006	.050**	.058***	−.005	.053**	.051**	.006
	(.019)	(.019)	(.015)	(.019)	(.020)	(.015)	(.021)	(.023)	(.016)
Δ share retail employment	−.112	.315	−.415**	−.119	.309	−.418**	−.099	.278	−.365**
	(.156)	(.226)	(.163)	(.152)	(.225)	(.164)	(.152)	(.214)	(.161)
Δ share nontradable employment	.193**	.181*	.022	.182**	.174	.018	.176**	.182*	.004
	(.091)	(.108)	(.064)	(.088)	(.106)	(.064)	(.088)	(.101)	(.063)
Δ share construction employment	.094	−.119	.218***	.089	−.123	.217***	.074	−.118	.195***
	(.077)	(.077)	(.053)	(.075)	(.077)	(.053)	(.077)	(.078)	(.054)
Home ownership rate				−.016	−.017	.000	−.038**	−.041**	.003
				(.011)	(.011)	(.012)	(.019)	(.019)	(.020)
Δ house prices × home ownership rate				.115**	.098**	.019	.146**	.172***	−.026
				(.047)	(.041)	(.045)	(.072)	(.065)	(.068)
Population density							−.001	.001	−.002*
							(.001)	(.001)	(.001)
Δ house prices × population density							.005	−.000	.005
							(.004)	(.004)	(.004)
Share below 35 years							−.000	−.001*	.001
							(.000)	(.000)	(.001)
Δ house prices × share below 35 years							−.001	.002	−.003
							(.002)	(.002)	(.002)
R^2	.071	.277	.156	.092	.284	.150	.088	.284	.160
Observations	192	192	192	192	192	192	192	192	192

NOTE.—The table shows results from regression (3). All specifications include a constant that is not reported. The unit of observation is a zip code; the dependent variable is the change in retail prices (cols. 1, 4, and 7), the change in retail markups (cols. 2, 5, and 8), and the change in marginal costs (cols. 3, 6, and 9) for a large national retailer between January 2004 and June 2007. See Sec. II.D.1 for a discussion of what is included in this measure of marginal cost. Robust standard errors are in parentheses.

  * Significant at $p < .10$.
  ** Significant at $p < .05$.
  *** Significant at $p < .01$.

population density and age composition. Consistent with results in Section II.C, the response of markups and prices to changes in house prices is increasing in the home ownership rate of the zip code.[33] In zip codes with the lowest home ownership rates, increases in house prices actually cause retail prices and markups to fall.

Products with a large component of local marginal cost.—While the previous results document that changes in local wholesale costs are not quantitatively important in general and are uncorrelated with house price changes, certain products do have a larger local cost component. If changes in local marginal costs were important, we would expect that those goods would contribute significantly to our estimated elasticity. In column 1 of table 4, we repeat the empirical analysis from table 1 using a retail price index that excludes product categories classified as "perishable" or as "liquid" by Bronnenberg et al. (2008). Perishable products, such as milk, are more likely to be sourced locally while liquid products, such as carbonated beverages, are frequently bottled locally, which might make their costs more sensitive to local shocks. However, we obtain very similar elasticities when excluding these potentially problematic product categories from our local retail price indices, confirming that a pass-through of local marginal costs is unlikely to explain our findings.

2.   Labor Costs

The previous sections documented that wholesale costs, which constitute the vast majority of retailers' marginal costs, vary little across geographies and do not rise with local house price shocks. We next address two other cost pass-through channels: labor costs and retail rents. If there was an increase in the shadow cost of labor, for example, because of higher wages, retail prices might increase as retailers pass through this (small) component of marginal cost.

However, the controls for changes in the unemployment rate, average weekly wages, and employment shares in our baseline regressions in table 1 already suggest that our findings are unlikely to be explained by pass-through of labor costs. In addition, columns 3 and 4 of table 4 show the robustness of the results in table 1 to alternative labor market controls.

---

[33] Our results might seem surprising in light of Eichenbaum et al.'s (2011) finding that this "retailer resets its reference prices so that variations in the realized markup fall within a reasonably small interval" (250). However, their finding refers to variation in markups within a given reference price spell, while our statistic computes markup variation both within and across spells. Small markup variation within a given reference price spell is consistent with substantial markup variation when reference prices change. Furthermore, even if there was no across-spell variation, one can still generate relatively large increases in markups of 22 percent, while remaining within the $+/-10$ percent within-spell markup interval they document.

TABLE 4
MARKUP OR MARGINAL COST?

	DEPENDENT VARIABLE: Δ RETAIL PRICES								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
A. Instrument with Saiz Supply Elasticity, 2001–6									
Δ house prices	.145***	.115**	.111**	.168**	.159***	.148**	.163***	.157***	.151***
	(.065)	(.046)	(.047)	(.048)	(.059)	(.058)	(.056)	(.058)	(.058)
B. Instrument with Saiz Supply Elasticity, 2007–11									
Δ house prices	.121***	.131***	.137***	.109	.127**	.129***	.140***	.152***	.134***
	(.059)	(.039)	(.044)	(.075)	(.049)	(.043)	(.051)	(.046)	(.045)
C. Instrument with Wharton Regulation Index, 2001–6									
Δ house prices	.194***	.196***	.195***	.265***	.224***	.259***	.272***	.262***	.223***
	(.052)	(.045)	(.042)	(.087)	(.048)	(.051)	(.055)	(.053)	(.047)
D. Instrument with Wharton Regulation Index, 2007–11									
Δ house prices	.246***	.155***	.165***	.145***	.159***	.159***	.149***	.140***	.154***
	(.054)	(.042)	(.044)	(.054)	(.045)	(.042)	(.043)	(.043)	(.042)

1420

Controls Robustness	Exclude liquids & perishable goods	Average unemployment rate	Control for employment to population	Control for low education wage & unemployment	Drop high retail rent cities	Control for changes in income	Control for changes in education	Control for population growth	Control for entry & exit
	✓	✓	✓	✓	✓	✓	✓	✓	✓

Note.—The table shows results from regression (2). All specifications include a constant that is not reported. The unit of observation is an MSA, and the dependent variable is the change in retail prices in 2001–6 (panels A and C) and 2007–11 (panels B and D). In panels A and B we instrument for the change in house prices using the housing supply elasticity measure provided by Saiz (2010); in panels C and D we instrument for house price changes using the Wharton Regulation Index described in Gyourko et al. (2008). All specifications control for changes in the unemployment rate, changes in wages, and changes in the employment share in the construction, nontradable, and grocery retail sectors. Col. 1 drops all product categories classified as "perishable" in Bronnenberg et al. (2008), as well as all liquids from our construction of the local price index. Col. 2 controls for the average unemployment rate over the sample rather than for changes in the unemployment rate. Col. 3 controls for changes in the employment-to-population ratio rather than changes in the unemployment rate. Col. 4 controls for changes in the wage and unemployment of lower-educated people in the ACS, defined as those with at most a high school diploma. Col. 5 drops the six cities with the highest level of retail rents (Boston, Chicago, New York, Los Angeles, San Francisco, and Washington, DC). Col. 6 controls for changes in income using data from the Internal Revenue Service. Col. 7 controls for changes in the share of people who have completed high school and changes in the share of people who have completed a bachelor's degree. Col. 8 controls for population growth using data from the annual population estimates for MSAs produced by the US Census Bureau. Col. 9 controls for changes in the number of grocery retail establishments per 1,000 citizens. Robust standard errors are in parentheses.

* Significant at $p < .10$.

** Significant at $p < .05$.

*** Significant at $p < .01$.

First, local unemployment measures can be sensitive to measured local labor force participation, but using changes in the employment-to-population ratio leaves our results unchanged. Second, grocery stores tend to hire labor that is less educated than the average population, but controlling for changes in wages or unemployment among those with at most a high school diploma in the ACS yields nearly identical results.

3.   Retail Rents

We next explore whether a pass-through of higher commercial rents can explain the retail price response to house prices.[34] The most important piece of evidence against this channel is that the response of house prices to retail prices grows with local home ownership rates and is essentially zero in areas with mainly renters (see Sec. II.C). An increase in local rents should affect a firm's costs in the same way whether the firm is located in an area with many or few homeowners. Thus, an explanation for our price patterns that relies on pass-through of local land prices into commercial rents and retail prices will struggle to explain the observed home ownership interaction.

Nevertheless, we next directly control for changes in retail rents in our empirical specifications, using data on annual effective retail rents that we obtained from REIS for 45 MSAs. Appendix figure A2 shows the relationship between changes in house prices and changes in retail rents over our sample period. The elasticity of cumulative changes in retail rents to cumulative changes in house prices over the housing boom is 0.2; it is 0.08 in the housing bust. This relatively low pass-through of house prices to retail rents is consistent with the long duration of retail lease contracts. As a first back-of-the-envelope calculation, even if retail rent made up an unrealistic 20 percent of marginal costs, these estimates suggest that rent pass-through could explain at most one-fourth of our retail price movements.[35]

To assess more formally the extent to which the (small) changes in retail rents can explain our results, table 5 includes the average retail rent as a control variable in regression (2). While the statistical significance of

---

[34] Whether rent is a fixed or marginal cost depends on the time horizon. In the short run, rents are largely a fixed cost (Gopinath et al. 2011). With entry and exit at longer horizons, an increase in fixed overhead costs would lead to a decline in the number of stores, and the resulting reduction in competition should lead to an increase in markups. As long as other costs remained constant, this pass-through channel would still represent an increase in markups.

[35] If 20 percent of marginal costs are retail rents and the rent elasticity is 20 percent, then a doubling of house prices will increase marginal costs by $20\% \times 20\% = 4$ percent. This is only one-fourth of the total increase in retail prices observed in our regressions.

TABLE 5
CONTROLLING FOR RETAIL RENT

	OLS			IV Saiz			IV Wharton		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	A. Time Period, 2001–6								
Δ house prices	.063**	.084**	.074**	.088*	.138	.122	.188***	.459*	.470*
	(.028)	(.034)	(.036)	(.052)	(.131)	(.139)	(.059)	(.238)	(.269)
Δ retail rent		−.101	−.092		−.221	−.194		−1.192	−1.216
		(.122)	(.122)		(.413)	(.421)		(.771)	(.846)
Δ wage			.115			.097			−.036
			(.173)			(.170)			(.160)
Observations	45	45	45	42	42	42	42	42	42
	B. Time Period, 2007–11								
Δ house prices	.104***	.114***	.109***	.105***	.114***	.105**	.132**	.129**	.119**
	(.022)	(.023)	(.023)	(.041)	(.044)	(.041)	(.052)	(.053)	(.051)
Δ retail rent		−.121	−.120		−.233	−.232		−.275	−.270
		(.123)	(.121)		(.180)	(.163)		(.209)	(.193)
Δ wage			.133*			.125*			.120
			(.077)			(.069)			(.074)
Observations	45	45	45	42	42	42	42	42	42

NOTE.—The table shows results from regression (2). All specifications include a constant that is not reported. The unit of observation is an MSA, and the dependent variable is the change in retail prices in 2001–6 (panel A) and 2007–11 (panel B). We show results from an OLS specification (cols. 1–3), as well as IV specifications that instrument for the change in house prices using the Saiz (2010) measure of housing supply elasticity (cols. 4–6) and the Wharton Regulation Index described in Gyourko et al. (2008) (cols. 7–9). The sample is restricted to MSAs for which we observe retail rents in the REIS data. Robust standard errors are in parentheses.

  * Significant at $p < .10$.
  ** Significant at $p < .05$.
  *** Significant at $p < .01$.

the elasticity estimates declines because of the smaller sample size, our results suggest that the increase in retail prices in response to higher house prices is not driven by the pass-through of higher retail rents. If anything, controlling for changes in retail rents increases the estimated response of retail prices to changes in house prices. As further evidence, in column 5 of table 4 we exclude the six MSAs in our data with the largest level of retail rents, as identified in the 2012 Retail Research Report provided by Colliers International. In these markets, retail rents are likely to make up a larger fraction of total costs; therefore, if the pass-through of higher retail rents were a significant factor in explaining our results, we would expect the estimated elasticity to be smaller when excluding cities with high retail rents. Contrary to this, the estimated elasticity is unchanged.

4. Demographic Changes and Gentrification

We next explore whether our results are driven by migration and changing demographics rather than by changes in the behavior of individuals already living in a location. If richer, less price-sensitive households moved into a location when house prices increase or if retailers responded to an overall increase in demand due to more people living in an MSA, then this could change the interpretation of our results. In column 6 of table 4 we control for changes in income and in column 7 for changes in the fraction of population that has completed at least high school or at least a bachelor's degree. In column 8 we control for population growth over our sample period. Our estimates are unaffected by the addition of these control variables. Consistent with this, Section II.E shows that the shopping behavior of the same household does indeed change in response to house price movements.

5. Grocery Retail Entry

As discussed in Section II.A, the most pertinent potential challenge to using housing supply elasticity as an instrument for house price changes is that changes in the competitiveness of the retail sector might be correlated with both the housing supply elasticity and changes in retail prices. In particular, in areas where it is difficult to build new houses, it may also be difficult for new retail establishments to enter. In that case, areas with a low housing supply elasticity might see greater retail price growth because of greater house price growth (our effect of interest), but also because of less retail competition.

In appendix A we formalize this potential concern but show that these confounding effects can be explicitly dealt with by using data on the number of local grocery retail establishments as an additional control

in our regressions. Column 9 of table 4 directly controls for the change in the number of retail establishments per inhabitant (in addition to the share of grocery retail employment that is controlled for in all regressions reported in table 4).[36] If anything, the estimated elasticity is slightly larger. Thus, the response of retail prices to house price movements in our IV regressions is not explained by the simple effect of housing supply elasticity on local retail competition. This does not imply that supply elasticity does not affect entry, but it does imply that any such effects do not drive our house price elasticities. In addition, in order for entry effects to explain the interaction of house price changes with local home ownership rates in Section II.C, there would need to be a strong relationship between entry and owner occupancy rates. However, when we include an additional interaction between entry and house price changes in regression (3), this does not affect our previous results. Finally, table A6 shows that much of the response of retail prices to house prices occurs at relatively high frequencies, where entry is unlikely to be of quantitative significance.

While changes in grocery retail entry therefore cannot explain our findings, they are interesting in their own right and provide additional support for markup variation. In table A4 we show the effects of house prices and supply elasticity on entry. Overall, we find that, if anything, there is more entry during the housing boom in the less elastic regions that experienced larger increases in house prices. This provides further evidence that those regions experienced increases in retail markups: without an increase in profitability, it is hard to explain why there would be more entry in less elastic regions where entry is more restricted. Of course, this need not imply that lower supply elasticity increases entry holding all else equal. Our previous regressions show that areas with lower housing supply elasticity had higher house price–induced retail price growth during the housing boom, which makes entry more attractive. Columns 4 and 5 show that after controlling for house price growth, greater housing supply elasticity has an insignificant effect on entry.

6. Product or Store Quality

It is worth reemphasizing, at this point, that none of our results are explained by either changes in the composition of stores or changes in the composition of products within a store. The reason is that our price indices measure price changes for the same UPC in the same store. If a low-quality product is replaced with a higher-quality product with a higher

---

[36] We obtain similar results when using changes in the absolute number of establishments and when controlling for establishment levels rather than changes. Similar controls in our OLS specification also slightly increase our elasticity estimates.

price, this product substitution itself does not affect our price index. Similarly, if higher house prices lead to the entry of higher-quality stores that charge higher prices, this also does not affect our price index.

A second concern might be that changes within a particular store could affect the quality of the same UPC over time. However, while it might be conceivable that, as house prices go up, stores increase the "freshness" of their produce, this is much less likely for the type of processed foods and toiletries that we observe in our data. Furthermore, this would result in higher shipping and inventory costs, yet we observe no change in marginal cost in our large retailer data.

Finally, one might be concerned about time-varying changes in the "shopping experience" of buying identical goods. Even if these were important, many changes to the shopping experience, such as upgrading the store, are fixed costs, so a pass-through of these costs would reflect an increase in markups. Conversely, most changes to the shopping experience that increase the marginal cost of selling a particular product, such as hiring more staff to reduce checkout lines, should be picked up by the controls for labor shares and other measures of marginal cost. Finally, while grocery stores might be renovated during housing boom periods, it is unlikely that store owners actively degrade stores during the housing bust, and differential depreciation during the housing bust would operate on a longer time scale and thus cannot explain the results at a quarterly frequency described in appendix B.

E.   Shopping Behavior

In the previous sections we documented a positive, causal relationship between house prices and retail prices. We argued that this relationship is not driven by an increase in retailers' marginal costs and is therefore best explained by an increase in retail markups. In this section, we provide further evidence on why retailers raise markups following house price increases, arguing that this is the optimal response to a decrease in the overall price elasticity. In particular, we show that higher house prices lead homeowners to increase their nominal spending and to become less price sensitive, while renters, if anything, exhibit the opposite behavior. Such a demand elasticity response is a natural feature of many models in which the value of leisure rises with wealth, so that wealthier households allocate less time to shopping for cheaper prices and thus become less price sensitive (see Alessandria 2009; Aguiar et al. 2013; Huo and Ríos-Rull 2015; Kaplan and Menzio 2016). This decrease in the demand elasticity faced by firms increases their optimal markup.

We use household-level information on purchasing behavior from Nielsen Homescan to analyze how changes in house prices affect house-

hold shopping behavior. Motivated by the differential response of retail prices to house prices in zip codes with different home ownership rates, we allow homeowners and renters to respond differently to house price changes.[37] The dependent variable in regression (4) captures the shopping behavior of household i, in zip code z, in quarter q:[38]

$$\text{ShoppingOutcome}_{i,z,q} = \psi_{i,z} + \xi_q + \beta_1 \log(\text{HousePrices})_{z,q}$$
$$+ \beta_2 \text{Homeowner}_{i,q} + \beta_3 \log(\text{HousePrices})_{z,q}$$
$$\times \text{Homeowner}_{i,q} + \gamma X_z + \epsilon_{i,z,q}. \tag{4}$$

In our baseline estimates, we measure local house prices at the quarter × zip code level, but our results are robust to measuring house prices at the quarter × MSA level. We include quarter fixed effects, ξ_q, to control for any aggregate time trends. Importantly, we also control for household × zip code fixed effects, $\psi_{i,z}$. This keeps constant any household-specific determinants of shopping behavior, such as the disutility from comparing prices or the baseline preference for generic goods. In addition, it ensures that the effects are not driven by changes in households' shopping behavior as they move to zip codes with different levels of house prices. Instead, our estimates come from observing how the behavior of the same household living in the same zip code changes as house prices vary. The parameter β_1 is informative for how changes in house prices affect the shopping behavior of renters, while β_3 captures the differential effect of house prices on homeowners.

Columns 1 and 2 of table 6 show that increases in house prices lead to more retail spending by homeowners but to somewhat reduced spending by renters (though that effect is not statistically significant). Owing to household fixed effects in regression (4), the coefficient on homeowner, β_2, is identified only off households that change tenure, so the negative coefficient suggests that households spend less after a house purchase, perhaps because of down payment constraints. Our results are insensitive to dropping these households, and if we remove household fixed effects, homeowners have higher expenditures and are less price sensitive than renters. Overall, the estimate of β_3 is consistent with homeowners

---

[37] We identify households in one-family noncondo buildings as homeowners and those in three-plus-family noncondo buildings as renters. Replacing household-level home ownership with zip code–level home ownership rates delivers similar results.

[38] A quarterly specification rather than the previous long-difference specification is necessary in Homescan data because panel turnover means that only a small and selected set of households are observed over the full sample. To facilitate comparability, app. B presents quarterly specifications of our retail price results.

TABLE 6
Effect of House Prices on Shopping Behavior

DEPENDENT VARIABLE

	Log(Expenditure)			Share "Deal"			Share Generic			Share Coupon		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Log(house price)	−.024	−.012	.408***	−.004	.002	.026	.002	.002	−.017	.001	−.000	.021
	(.016)	(.016)	(.077)	(.005)	(.006)	(.024)	(.003)	(.003)	(.014)	(.003)	(.003)	(.014)
$1_{homeowner}$	−.209***	−.203***	−.253***	.085***	.084***	.073**	.036**	.036**	.033**	.049***	.048***	.028**
	(.077)	(.077)	(.080)	(.028)	(.028)	(.028)	(.014)	(.014)	(.015)	(.013)	(.013)	(.014)
Log(house price) × $1_{homeowner}$	.048***	.046***	.055***	−.017***	−.017***	−.014**	−.007***	−.007***	−.007**	−.009***	−.009***	−.005*
	(.015)	(.015)	(.016)	(.005)	(.005)	(.006)	(.003)	(.003)	(.003)	(.003)	(.003)	(.003)
Unemployment rate		.384***	.420***		.097***	.091***		−.015	−.019		−.045***	−.046***
		(.083)	(.083)		(.026)	(.025)		(.015)	(.015)		(.015)	(.015)
Average weekly wage		.013	.016		−.017***	−.017***		−.004	−.004		−.005*	−.005*
		(.016)	(.016)		(.005)	(.005)		(.003)	(.003)		(.003)	(.003)
Share grocery retail Employment		−.008	.020		−.014	−.019		.005	.003		.015	.013
		(.082)	(.082)		(.026)	(.026)		(.016)	(.016)		(.015)	(.015)
Share nontradable Employment		.120**	.108**		.027*	.028*		.009	.009		.003	.003
		(.048)	(.048)		(.015)	(.015)		(.009)	(.009)		(.008)	(.008)

1428

Share construction	✓			✓			✓			✓		
Employment		.171*** (.049)	.167*** (.049)		−.013 (.015)	−.016 (.015)		.005 (.009)	.007 (.009)		−.015* (.008)	−.015* (.008)
Quarter fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Household × zip code fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Household characteristics × log(house price)			✓			✓			✓			✓
R^2	.721	.721	.722	.872	.872	.872	.740	.740	.741	.770	.770	.771
$\bar{y}$	6.698	6.698	6.698	.277	.277	.277	.176	.176	.176	.077	.077	.077
Observations	839,176	839,176	839,176	839,176	839,176	839,176	839,176	839,176	839,176	839,176	839,176	839,176

NOTE.—The table shows results from regression (4). All specifications include a constant that is not reported. The unit of observation is a household-quarter; the sample is 2004–11. The dependent variables are the log of total household expenditure (cols. 1–3), the expenditure share of products that are on sale (cols. 4–6), the expenditure share of generic products (cols. 7–9), and the expenditure share of products purchased with a coupon (cols. 10–12). House prices are measured at the zip code level. All specifications include household × zip code fixed effects and quarter fixed effects. In cols. 2, 3, 5, 6, 8, 9, 11, and 12, we also include additional control variables. In cols. 3, 6, 9, and 12, we additionally control for the following household characteristics interacted with log house prices: household income, age of head of household, race of head of household, marital status of head of household, and education level of head of household. Each observation is weighted by the household sampling weight. Standard errors are clustered at the zip code × quarter level.

  * Significant at $p < .10$.

  ** Significant at $p < .05$.

  *** Significant at $p < .01$.

1429

consuming out of their increased housing wealth.[39] These estimates highlight the first of two channels through which higher house prices reduce the average price elasticity faced by retailers. As house prices increase, the share of total expenditure coming from lower-elasticity homeowners increases relative to the share of expenditures from higher-elasticity renters. In addition to this "extensive margin" effect, we document below an "intensive margin" effect, whereby each individual homeowner becomes even less price elastic as house prices increase.

One potential concern with the specifications in columns 1 and 2 is that the differential behavior by homeowners might also be picking up differential behavior along other dimensions that are correlated with homeownership, such as income and age. To test whether this is the case, in column 3 we include additional controls for the interaction of log house prices with household income, age of the household head, race of the household head, marital status of the household head, and education level of the household head. Reassuringly, the coefficient on the interaction of home ownership status and log house prices remains unaffected.[40] This supports the interpretation that our estimates are picking up differential wealth effects of house price changes for homeowners and renters.[41]

---

[39] Given our previous findings that higher house prices also cause higher retail prices, it is interesting to ask what fraction of this increase in nominal expenditures is driven by higher prices. We find that across those individuals living in areas for which we can construct an IRI price index, the house price effect on real expenditures is about 30 percent smaller than the effect on nominal expenditures. However, we observe an IRI price index for the locations of only about 70 percent of respondents in the Homescan data. When we construct an alternative price index based on Homescan transaction prices, the real share of the nominal increase falls slightly. See Kaplan, Mitman, and Violante (2016) for recent related estimates.

[40] We also control for uninteracted levels of these characteristics. Because of household fixed effects, these are identified only in the small number of cases in which there is a within-household change in these variables. In the interest of space, we do not report coefficients on the additional controls or their interactions with house prices. Since we include interactions of house prices with these additional controls, the single coefficient, β_1, loses its interpretation as the effect of house prices on renters.

[41] It is interesting to consider other dimensions of heterogeneity that might predict different responses to house price changes. Unfortunately, theoretical predictions for other characteristics in our data are ambiguous. To see this, decompose the total expenditure response to house prices into two components, the expenditure response to wealth (or borrowing constraint) changes and the response of wealth (or borrowing constraints) to house price changes:

$$\frac{d\text{Exp}}{d\text{HP}} = \frac{d\text{Exp}}{d\text{Wealth} + \text{Borrow}} \times \frac{d\text{Wealth} + \text{Borrow}}{d\text{HP}}.$$

Since $(d\text{Wealth} + \text{Borrow})/d\text{HP}$ has (weakly) opposite signs for homeowners and renters and $d\text{Exp}/(d\text{Wealth} + \text{Borrow})$ is (weakly) positive for everyone, the difference in $d\text{Exp}/d\text{HP}$ for homeowners and renters follows naturally. However, many other sources of within-homeowner heterogeneity give rise to ambiguous predictions. For example, consider young vs. old homeowners. Since both parts of the decomposition are positive for homeowners, any differential age effects depend on how magnitudes rather than signs vary with age. In

We now turn to "intensive margin" effects, showing that individual home-owners become less price sensitive as house prices rise, consistent with our previous findings that retail prices increase more with house prices in areas with more homeowners. In columns 4–6 of table 6, the dependent variable is the expenditure share on goods that are on sale. We find that, as house prices increase, homeowners are less likely to purchase goods that are on sale. In columns 7–9, we use the share of purchases of cheaper generic goods as the dependent variable. A higher share of generic purchases again suggests higher price sensitivity. In columns 10–12, the dependent variable is the share of purchases made with a coupon, another measure of how intensely consumers substitute time for lower prices. As before, these three measures of price sensitivity decrease with house prices for homeowners but do not move significantly for renters.

One might be concerned that changing expenditure shares could reflect changes in the composition of goods purchased by households as their housing wealth increases rather than changes in households' shopping intensity and price sensitivity for a given good. For example, a decline in the expenditure share on sale items could reflect either a reduction in the shopping intensity devoted to the same goods or a change in the composition of purchases toward goods that are on sale less often. In the latter case we would see changes in expenditure shares, but this would not necessarily indicate a decline in price sensitivity. To show this is not the case, table A7 presents results from a version of regression (4) in which the unit of observation is a shopping outcome for each household × quarter × product category. As before, we include household × zip code fixed effects but also add product category × quarter fixed effects.

Columns 1–3 show that, for homeowners, higher house prices lead to higher total expenditures within each product category; higher house prices lead to lower expenditures for renters, though the effect is again not statistically significant. The magnitude is similar to that in table 6. Columns 4–6 show that the share of products bought on sale within each product category varies with house prices in the same way as when we pool across product categories. Similar results are obtained when looking at the share of goods purchased with a coupon and the share of generic goods purchased. This suggests that the observed changes in ex-

---

the Survey of Consumer Finances, housing wealth is a relatively constant share of total wealth among owners over the life cycle except at the very end of life. Therefore, we would expect the life cycle profile of $d\text{Exp}/d\text{HP}$ to be driven by the life cycle profile of $d\text{Exp}/(d\text{Wealth} + \text{Borrow})$. However, there is no consensus on this profile, either in the data or in theory (see Sahm, Shapiro, and Slemrod [2010], Kaplan and Violante [2014], Berger et al. [2018], and papers cited therein). Other household characteristics in our data (e.g., income) involve similar ambiguity, so we restrict attention to the difference between renters and homeowners. However, studying the empirical relationship between housing equity, liquidity, and shopping behavior is an interesting avenue for future research that could potentially be explored by linking credit bureau information with the Nielsen Homescan data.

penditure shares are truly driven by changing household price sensitivity, and not by compositional changes in the types of products purchased.

Finally, one might be interested in analyzing the extent to which our findings are driven by changes in the share of goods that are on sale (or changes in the local availability of generics or coupons) rather than by changes in households' effort in searching for these sales. That is, we want to isolate changes in purchases that are driven by changes in household behavior from those driven by changes in firm behavior. To do this, we would ideally like to include zip code × quarter fixed effects to capture time variation in the propensity of goods in a zip code to be on sale. However, this removes almost all of the variation, since we often observe only one household per zip code. In table A8 we thus repeat regression (4) including MSA × quarter fixed effects, both with and without including other household characteristics interacted with the log of house prices. This controls for MSA-level changes in the share of goods offered on sale in response to changes in house prices. The estimated interaction between house prices and home ownership status remains economically and statistically significant and is, if anything, slightly larger in magnitude than that in table 6.

The evidence in this section shows that wealth effects from higher house prices make homeowners less price elastic and renters more price elastic. Therefore, as house prices increase, retailers can increase their markups, in particular in areas with many homeowners.

F.   Interpretation and Discussion of Magnitude

In the previous sections, we concentrated on establishing a causal relationship between changes in house prices and changes in markups and on documenting changes in the price elasticities of homeowners in response to house price increases. We next interpret the magnitude of our effects and their implications for price variation across locations. An important first step is to discuss the external validity of our findings, since the scanner data that allow us to precisely measure local prices and establish causality are limited to mostly food items, and other prices could potentially respond differently to house price changes.

There are several reasons why we believe our results extend to prices in general. Panel C of figure 1 shows that, across cities, the BLS food-at-home CPI moves closely with the broader CPI, which suggests that the price movements we identify generalize to a large set of goods. We also run a simple OLS regression of changes in these BLS price indices on changes in local house prices. Despite having only 26 observations in this regression, we find significant responses of both the broad CPI and the food-at-home CPI to house price movements. The implied elasticities are

nearly identical both to each other and to our regressions using IRI data.[42] Finally, in addition to the pooled results discussed in Section II.E, we found that our measures of household shopping intensity respond similarly to house price movements across a large variety of product categories, so there is reason to believe that markups should also respond similarly.[43] Thus, we believe that our estimated elasticities provide a reasonable baseline for the response of more general price levels to house price movements.

How much of the across-MSA variation in inflation can be explained by differential house price movements? Over the housing boom, the 90–10 percentile difference in house price growth across MSAs was 45 percent. Multiplying this difference by our estimated boom elasticities of 15–23 percent implies that moving from the 10th percentile of MSA house price growth to the 90th percentile of MSA house price growth generates an increase in relative retail prices of 7–10 percent.[44] This compares to an overall 90–10 difference in retail price changes of 12 percent. The same calculation in the housing bust implies that house price differences generate a 5–6 percent 90–10 retail price movement, compared with an actual 90–10 difference of 8.4 percent. Given that the differential housing boom-bust across locations was one of the most important regional factors during this time period, we think it is indeed plausible that a substantial part of the regional variation in retail price changes can be explained through this channel.

Is the magnitude of markup variation implied by our observed retail price movements plausible? If markup variation explains all of the ob-

---

[42] Specifically, when we use the overall CPI, we find elasticities of 0.106 (standard error [SE] = 0.040) for the full period (2001–11), as well as elasticities of 0.079 (SE = 0.020) for the boom (2001–6) and elasticities of 0.051 (SE = 0.026) for the bust (2007–11). When we analyze the food-at-home CPI, elasticities are 0.062 (SE = 0.018) for the boom, 0.042 (SE = 0.054) for the bust, and 0.085 (SE = 0.034) for the full period. The estimated elasticities for the bust period in this sample are significantly downward biased by the outlier of Miami, which has by far the largest house price drop but has about average retail price increases. When dropping this observation, elasticities during the bust increase to 0.095 (SE = 0.054) for the food-at-home CPI and to 0.083 (SE = 0.023) for the overall CPI. Dropping Miami leaves the estimates for the other elasticities relatively unaffected.

[43] One might be concerned that since food spending rises less with wealth than other consumption, implied price responses in other categories might be implausibly large. However, even if the level of demand for different goods responds differently to shocks, this need not imply that the elasticity of demand also does so. We believe it is plausible that the elasticity of demand for food falls substantially with wealth shocks, even if most of the additional consumption occurs in nonfood items.

[44] Our conclusion that, across space, stores are varying markups substantially might seem at odds with the conclusion in Kaplan and Menzio (2016) that store-specific pricing effects explain only 10 percent of overall price variation. However, overall price variation across households is huge in their data, so the 10 percent, which is due to the store-specific component, is still large in an absolute sense. More importantly, their analysis computes price dispersion within a market rather than across markets, and so most of our effects of interest would not be captured in their price measures.

served elasticities of retail prices to house prices, then a 10 percent increase in house prices implies markup changes of 1.5–2 percentage points. We directly observe an average markup of roughly 45 percent for our large anonymous retailer, so a 2 percentage point movement does not seem unreasonable.[45] Assuming a constant elasticity of substitution, this implies a reduction from 3.22 to 3.13 in this elasticity.

Finally, it is important to note that all of our empirical estimates measure the response of prices and markups to house prices at medium-run business cycle frequencies. Section II.D.5 shows that for the time horizons in our sample, store entry plays only a small role. However, over longer time horizons, entry should diminish these initial markup responses. That is, our conclusion that natural markups are procyclical does not imply that there will be trend growth in markups in the long run, even if average household wealth increases over time.

III.   Conclusion

We link detailed geographically disaggregated data on local house prices, retail prices, and household shopping behavior to provide new evidence on how the economy responds to changes in demand. We argue that increases in house prices lead to changes in demand for homeowners who become less price sensitive and that firms respond by raising markups. Consistent with this interpretation, we find much stronger retail price responses to changes in house prices when home ownership rates are high. We also find evidence of differential shopping responses to house price changes for homeowners and renters. The economic magnitude of our price effect is large but not implausible: we estimate elasticities of retail prices to house prices of 15–20 percent and show that this channel can explain a large fraction of geographic variation in retail price changes. As discussed in the introduction (and in more detail in app. D), the conclusion that household price sensitivity declines with housing wealth and that this leads firms to raise markups has wide-ranging implications for business cycle modeling, industrial organization, and urban economics.

References

Adelino, Manuel, Antoinette Schoar, and Felipe Severino. 2015. "House Prices, Collateral and Self-Employment." *J. Financial Econ.* 117 (2): 288–306.
Aguiar, Mark, Erik Hurst, and Loukas Karabarbounis. 2013. "Time Use during the Great Recession." *A.E.R.* 103 (5): 1664–96.

----

[45] In the Census of Retail Trade, the average retail markup is only modestly lower at 39 percent.

Alessandria, George. 2009. "Consumer Search, Price Dispersion, and International Relative Price Fluctuations." *Internat. Econ. Rev.* 50 (3): 803–29.

Beraja, Martin, Erik Hurst, and Juan Ospina. 2019. "The Aggregate Implications of Regional Business Cycles." Working Paper no. 21956, NBER, Cambridge, MA.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra. 2018. "House Prices and Consumer Spending." *Rev. Econ. Studies* 85 (3): 1502–42.

Bils, Mark. 1987. "The Cyclical Behavior of Marginal Cost and Price." *A.E.R.* 77:838–55.

Bronnenberg, Bart J., Michael W. Kruger, and Carl F. Mela. 2008. "Database Paper—the IRI Marketing Data Set." *Marketing Sci.* 27 (4): 745–48.

Chevalier, Judith A., Anil K. Kashyap, and Peter E. Rossi. 2003. "Why Don't Prices Rise during Periods of Peak Demand? Evidence from Scanner Data." *A.E.R.* 93:15–37.

Christiano, Lawrence J., Roberto Motto, and Massimo Rostagno. 2010. "Financial Factors in Economic Fluctuations." Working Paper no. 1192, European Central Bank, Frankfurt.

Coibion, Olivier, Yuriy Gorodnichenko, and Gee Hee Hong. 2015. "The Cyclicality of Sales, Regular and Effective Prices: Business Cycle and Policy Implications." *A.E.R.* 105 (3): 993–1029.

Davidoff, Thomas. 2013. "Supply Elasticity and the Housing Cycle of the 2000s." *Real Estate Econ.* 41 (4): 793–813.

Domowitz, Ian, R. Glenn Hubbard, and Bruce C. Petersen. 1986. "Business Cycles and the Relationship between Concentration and Price-Cost Margins." *Rand J. Econ.* 17:1–17.

Dube, Jean-Pierre, Gunter Hitsch, and Peter Rossi. 2018. "Income and Wealth Effects on Private-Label Demand: Evidence from the Great Recession." *Marketing Sci.* 37 (1): 22–53.

Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo. 2011. "Reference Prices, Costs, and Nominal Rigidities." *A.E.R.* 101:234–62.

Gagnon, Etienne, and David Lopez-Salido. 2014. "Small Price Responses to Large Demand Shocks." Technical report, Board Governors, Fed. Reserve System, Washington, DC.

Gali, Jordi, Mark Gertler, and J. David Lopez-Salido. 2007. "Markups, Gaps, and the Welfare Costs of Business Fluctuations." *Rev. Econ. and Statis.* 89 (1): 44–59.

Gicheva, Dora, Justine Hastings, and Sofia Villas-Boas. 2010. "Investigating Income Effects in Scanner Data: Do Gasoline Prices Affect Grocery Purchases?" *A.E.R.* 100:480–84.

Giroud, Xavier, and Holger M. Mueller. 2017. "Firm Leverage, Consumer Demand, and Employment Losses during the Great Recession." *Q.J.E.* 132:271–316.

Glaeser, Edward L. 2013. "A Nation of Gamblers: Real Estate Speculation and American History." Technical report, NBER, Cambridge, MA.

Gopinath, Gita, Pierre-Olivier Gourinchas, Chang-Tai Hsieh, and Nicholas Li. 2011. "International Prices, Costs, and Markup Differences." *A.E.R.* 101 (6): 2450–86.

Gyourko, Joseph, Albert Saiz, and Anita Summers. 2008. "A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index." *Urban Studies* 45 (3): 693–729.

Handbury, Jessie. 2012. "Are Poor Cities Cheap for Everyone? Non-homotheticity and the Cost of Living across US Cities." Research Paper no. 71, Wharton School, Univ. Pennsylvania.

Handbury, Jessie, and David E. Weinstein. 2015. "Goods Prices and Availability in Cities." *Rev. Econ. Studies* 82 (1): 258–96.

Huo, Zhen, and José-Víctor Ríos-Rull. 2013. "Paradox of Thrift Recessions." Technical report, NBER, Cambridge, MA.

———. 2015. "Tightening Financial Frictions on Households, Recessions, and Price Reallocations." *Rev. Econ. Dynamics* 18 (1): 118–39.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti. 2011. "Investment Shocks and the Relative Price of Investment." *Rev. Econ. Dynamics* 14 (1): 102–21.

Kaplan, Greg, and Guido Menzio. 2016. "Shopping Externalities and Self-Fulfilling Unemployment Fluctuations." *J.P.E.* 124 (3): 771–825.

Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante. 2016. "Non-durable Consumption and Housing Net Worth in the Great Recession: Evidence from Easily Accessible Data." Technical report, NBER, Cambridge, MA.

Kaplan, Greg, and Giovanni L. Violante. 2014. "A Model of the Consumption Response to Fiscal Stimulus Payments." *Econometrica* 82 (4): 1199–1239.

Krueger, Alan B., and Andreas Mueller. 2010. "Job Search and Unemployment Insurance: New Evidence from Time Use Data." *J. Public Econ.* 94 (3): 298–307.

Manova, Kalina, and Zhiwei Zhang. 2012. "Export Prices across Firms and Destinations." *Q.J.E.* 127:379–436.

Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump." *Q.J.E.* 128 (4): 1687–1726.

Mian, Atif, and Amir Sufi. 2011. "House Prices, Home Equity–Based Borrowing, and the US Household Leverage Crisis." *A.E.R.* 101 (5): 2132–56.

———. 2014a. "House Price Gains and US Household Spending from 2002 to 2006." Technical report, NBER, Cambridge, MA.

———. 2014b. "What Explains the 2007–2009 Drop in Employment?" *Econometrica* 82 (6): 2197–2223.

Nakamura, Emi, and Dawit Zerom. 2010. "Accounting for Incomplete Pass-Through." *Rev. Econ. Studies* 77 (3): 1192–1230.

Nekarda, Christopher J., and Valerie A. Ramey. 2013. "The Cyclical Behavior of the Price-Cost Markup." Technical report, NBER, Cambridge, MA.

Nevo, Aviv, and Arlene Wong. 2019. "The Elasticity of Substitution between Time and Market Goods: Evidence from the Great Recession." *Internat. Econ. Rev.* 60:25–51.

Rotemberg, Julio J., and Michael Woodford. 1999. "The Cyclical Behavior of Prices and Costs." In *Handbook of Macroeconomics*, vol. 1B, edited by John B. Taylor and Michael Woodford, 1051–1135. Amsterdam: Elsevier.

Sahm, Claudia R., Matthew D. Shapiro, and Joel Slemrod. 2010. "Household Response to the 2008 Tax Rebate: Survey Evidence and Aggregate Implications." In *Tax Policy and the Economy*, vol. 24, edited by Jeffrey R. Brown, 69–110. Chicago: Univ. Chicago Press (for NBER).

Saiz, Albert. 2010. "The Geographic Determinants of Housing Supply." *Q.J.E.* 125 (3): 1253–96.

Sinai, Todd. 2013. "House Price Moments in Boom-Bust Cycles." In *Housing and the Financial Crisis*, edited by Edward L. Glaeser and Todd Sinai. Chicago: Univ. Chicago Press (for NBER).

Smets, Frank, and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *A.E.R.* 97 (3): 586–606.

Vavra, Joseph. 2014. "Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation." *Q.J.E.* 129 (1): 215–58.

Warner, Elizabeth J., and Robert B. Barsky. 1995. "The Timing and Magnitude of Retail Store Markdowns: Evidence from Weekends and Holidays." *Q.J.E.* 110:321–52.