Consumption responses to house price heterogeneity *
PRELIMINARY AND INCOMPLETE

James Graham †

November 18, 2017

Abstract

Movements in house prices may affect household consumption through wealth, collateral, income, or substitution effects. However, individual and aggregate consumption responses depend on whether house prices are moving due to aggregate, regional, local, neighborhood, or idiosyncratic shocks. I first show that there is significant city-, neighborhood-, and idiosyncratic-level variation in house prices.

Second, I show that the different components of house price movements are associated with different consumption movements. Using a large panel of consumers over the period 2004-2015, I find that aggregate price movements are associated with the largest consumption movements, however neighborhood-level price movements have a stronger effect than city-level price movements.

There are theoretical reasons to think that different components of house price movements should have differential effects on consumption. Previous work has shown that older homeowners have larger consumption responses to house price rises than younger households since older homeowners are more likely to downsize or sell their housing stock, implying future housing costs for them are lower, which generates a net positive wealth effect (Campbell and Cocco (2007)).

The same logic applies to movements across locations. Households are more likely to move across counties or neighborhoods than they are to move across cities, states, or regions. Thus, a house price increase in a particular neighborhood generates a wealth effect for households likely to move to other neighborhoods. A house price increase in a city does not generate a similar consumption response if a household never intends to leave the city.

To investigate this mechanism I then build a partial equilibrium, life-cycle model with heterogeneous agents to explore the effect on consumption of different levels of house price shocks.

*This draft: November 2017.

†New York University. Email: james.graham@nyu.edu

1 Introduction

House price cycles have long been known to have significant effects on macroeconomic dynamics. One important tranmission mechanism from house prices to the macroeconomy is through consumption. Household consumption may respond strongly to house price movements via, for example, substitution, household wealth, or collateral effects. Previous work has shown substantial evidence of house price effects on household consupmtion (see Campbell and Cocco (2007), Mian et al. (2013), Kaplan et al. (2016b), and Aladangady (2017)). However, the literature is mixed on whether the effect of house prices comes through collateral or wealth effects (see Berger et al. (2015) for a discussion).

It is important to recognize that the effect of house price movements on consumption depends on the characteristics of the households facing these movements. Buiter (2008) argues that the effect of house prices is heterogeneous across the economy. For example, young households face a long life-time of housing costs ahead of them. Any increase in prices that indicates an increase in future housing costs reflects a decrease in presented discounted wealth for these households. On the other hand, older households have a low present discounted value of future housing costs since they have a short horizon. Additionally, for life-cycle reasons, these households may want to downsize their housing stock. Thus, house price increases appear to increase the presented discounted wealth of older households, leading to higher consumption response to house price movements.

Another issue in identifying the effect of house price movements on consumption is that there is not one single house price in the economy. Rather, recent work has shown that there is significant heterogeneity in house price movements across the US (Ferreira and Gyourko (2011), Landvoigt et al. (2015), and Giacoletti (2016)). More broadly, the study of the sources and effects of idiosyncratic risk facing households is becoming more common in macroeconomics (see Heathcote et al. (2009) for an overview), yet so far little attention has been paid to heterogeneous house price movements as source of idiosyncratic risk.

How might house price heterogeneity be important for consumption responses to house price movements? I propose that house price movements matter to the extent that households are likely to experience house prices uncorrelated to the prices they currently face. Suppose a homeowner currently living in San Francisco, facing the high, and rising, house prices in that market, expects to move to Detroit in the near future. As long as San Francisco prices continue to diverge from prices in Detroit, the household enjoys an increase in life-time wealth.

Consider Figure 1. It shows that conditional on moving, households are more than twice as likely to move within their own county than to move across counties or states. Given the argument above, this suggests that all else equal, households should be more responsive to within-county house price movements than to across-county house price movements. In this paper, I investigate this propsition both empirically and within a structural model.

In this paper, I investgiate the effect of movements in different components of house prices on consumption. To do this, I first make use of a new data set on individual house transactions from across the US, made available by Zillow Research. The data I use includes 13 million individual housing transactions between 1994 and 2016. Using this data I can decompose house price movements into city (CBSA), neighborhood (zip-code), and idiosyncratic components. I show that all three components contribute to house price movements. City-level effects tend to dominate the volatility of house prices, although the idiosyncratic component plays a large role, and the volatility of neighborhood-level prices increased sharply during the housing bust. Since the idiosyncratic and neighborhood components of house prices are non-neglible, this

Figure 1: Household migration within the US

Each figure shows the proportion of households that moved house in the past year. Households may move within county, within state but across counties, across states, or move to the USA from another country. Migration information is provided at the person-evel, hence Figure 1 uses CPS-provided person-level weights. Homeowership status is provided at the household-level, so Figures 2 and 3 use CPS-provided household-level weights. Source: IPUMS-CPS.

opens the door to signifcant consumption responses to house price movements.

Unfortunately, I cannot match individual house sales in the Zillow data to household level consumption. This means that it is not possible to directly test the effect of idiosyncratic house price risk on consumption. Instead, I match the neighborhood and city components of house price movements to households by zipcode. To do this, I use the Kilts Consumer Panel data, which surveys 40,000 to 60,000 households between 2004 and 2015. I match households to corresponding national, zip-code, and CBSA house price components. In the results I have developed so far, I find that consumption is most strongly associated with the aggregate comoponent of house prices, however consumption is also significantly associated with the neighborhood level component of house prices. In contrast, the city-level component of house prices is only weakly associated with consumption, and the effect is weaker than for the neighborhood level component.

These findings are consistent with the view that households respond more strongly to the component of house prices that affects their future wealth. Since households are less likely to move across cities than they are to move witin cities, the within-city component of house prices is more relevant for household wealth and consumption than the city-wide component of house prices.

In order to explore this intuition further, I build a strucutral life-cycle model in which agents face realistic housing decisions. In particular, households can own or rent; they can choose their housing size; they can take out long-term mortgages against the value of their house; they can refinance their mortgage; and they face both city-level and neighborhood level house price risk. This part of the paper is also a work in progress, but so far I have found that the average elasticity of consumption with respect to unexpected, permanent, city-level house price shocks is around 0.2. However, there is also sigificant heterogeneity across the household's state-space (e.g. age, loan-to-housing value ratio).

This paper represents a first draft of this work, and as such there are many more improvements to both the empirical work and the model to be made. A non-exhaustive list of 'to-dos' can be found in Section 6.

2 Related literature

House price movements and their effect on household wealth are not entirely an aggregate story. Landvoigt et al. (2015) show that there is significant heterogeneity in house price movements during the recent boom and bust in the San Diego metropolitan area, with idiosyncratic returns volatility of between 8% and 14%. Giacoletti (2016) shows that for metropolitan areas in California, idiosyncratic house price risk explains between 20% and 60% of housing capital gains for long and short holding periods, respectively.

House prices are not entirely idiosyncratic, either. Ferreira and Gyourko (2011) study regional heterogeneity in the timing of house price booms. They find that MSAs and neighborhoods experienced different house price paths since the 1990s, with structural breaks beginning as early as 1997 (e.g. San Francisco) and as late as 2006. They also find one of the only economically and statistically significant explanatory variables for the beginning of these booms is local income growth. This can account for up to half of the initial jump in house prices, with the remainder unexplained.

Cross-sectional variation in house price movements has been used to identify the effect of house prices on consumption. Mian et al. (2013) study how shocks to household wealth via house price movements pass through to consumption. They make use of the Saiz (2010) instrument for supply elasticity. This helps deal with potential endogeneity between regional house price movements and consumption. Mian et al. (2013) show that similar sized shocks affect households in different parts of the wealth distribution differently. The consumption of low wealth households responds more strongly to the wealth shock than high wealth households, suggesting that the size of initial wealth holdings helps households to insure against such shocks.

In earlier work, Campbell and Cocco (2007) constructed a psuedo panel of households in the UK to investigate the effect of house prices on consumption. They show that both regional and national level house prices may affect consumption, although national house prices have a stronger effect.

Several hetereogeneous agent models have now been developed with the effect of house prices on consumption in mind. Gorea and Midrigan (2017) build a partial equilibrium lifecycle model with housing choice and long-term mortgages. They find that among homeowners that value liquidity, not all of them are borrowing constrained, and thus not all of them would increase consumption in response to a liquidity injection. Many households value additional liquidity for precautionary reasons: households who are close to paying the cost of extracting housing equity value additional cash on hand to the extent that they can avoid paying this extraction cost. These households have a large option value of waiting to extract equity. Thus, to the extent that fluctuations in house prices represent an unexpected increase in available liquidity, it is not immediately clear that households will necessarily extract and consume immediately. Rather, they may extract in order to hold a precautionary liquid buffer. Chen et al. (2013) present a similar model in this vein, exploring the effect of house price movements on borrowing via home equity extraction.

Favilukis et al. (2017) build a fully general equilibrium model and observe that when there is an aggregate house price, and house prices co-vary with business cycle shocks, these shocks make liquidity out of housing pro-cyclical, which generates counter-cyclical insurance opportunities, which generates a risk premium on housing wealth. Changes in this risk premium account for movements in equilibrium house prices. Kaplan et al. (2017) also build a general equilibrium model, in order to investigate the role of house prices in the most recent housing cycle. They find that shocks that generate the house price boom lead to large effects on consumption.

Beraja et al. (2017) build a model to explore the effect of regional house price cycles on the macroeconomy. They find that the distribution of housing equity across the economy, which is dispersed due to imperfectly correlated house price cycles, can generate large consumpton responses to other shocks.

Finally, Berger et al. (2015) investigate the theoretical effect of house prices on consumption in a partial equilibrium, incomplete markets model. For a plausibly calibrated model, they find that the effect of house prices on consumption is almost entirely due to wealth effects. However, the aggregate effects then depend on the underlying distributions of income, housing, mortgages, and so on.

3 Data

In this section I describe the house price and consumption data used in the empirical analysis. The house price data are taken from a newly available data set provided by Zillow Research. The consumption data are taken from the Kilts Consumer Panel Data. I use the house price data to investigate house price hetereogeneity, and the combine this with the consumption data in order to explore the effects of house prices on consumption.

3.1 House prices

To explore different components of house price movements, I make use of detailed micro data on individual house transactions. The data is the Zillow Transaction and Assessment Dataset (ZTRAX), made available by Zillow Research. ZTRAX contains more than 370 million public records containing information on deed transfers, mortgages, foreclosures, auctions, property characteristics, geographic information, and assessor valuations for residential and commercial properties. The data covers over 2750 US counties, and is available for up to twenty years for many of these counties.

I restrict the data to housing transactions (i.e. not mortgages or refinancing transactions) that are arm's-length and non-foreclosed sales, for properties that are non-commercial, single family residences. This restricts focus to household transactions.

After cleaning and filtering, there are 83 million transaction-level observations. Three states – Rhode Island, Tennessee, and Vermont – have various missing data in the ZTRAX database, and so are not included here. Although all states report the deeds records that the ZTRAX database is constructed from, several states either prohibit or do not require the disclosure of transactions prices.[1] For those states, a very large proportion of transactions report prices as zero. Table 4 reports the number of observations per state as well as the proportion of observations with non-zero prices. In the remainder of the analysis I drop the following states entirely due to missing price data: Alaska, Idaho, Indiana, Kansas, Maine, Missouri, New Mexico, Utah, and Wyoming.

[1] See http://www.zillowgroup.com/news/chronicles-of-data-collection-ii-non-disclosure-states/ for more details.

In Appendix B, Figures 8 and 9 compare, for each state, the log-median sale price (for non-zero price transactions) in ZTRAX to the log of the all transactions house price index from FRED. The states with many non-zero prices display a very poor match between the two series. For states with a lot of available data, the median ZTRAX price is very close to the all-transactions index. The following states provide a good fit to the FRED data: Arizona, California, Colorado, Connecticut, Washington DC, Delaware, Florida, Georgia, Hawaii, Iowa, Illinois, Kentucky, Massachusetts, Maryland, Minnesota, North Carolina, New Hampshire, New Jersey, Nevada, New York, Oregon, Pennsylvania, South Carolina, Virginia, Washington, Wisconsin. I focus on these states in the data analysis.

3.2 Consumption panel data

In order to investigate the effect of house price movements on consumption, I make use of Neilsen Consumer Panel data. The data comprise a panel of between 40,000 and 60,000 households, covering the years 2004 to 2015. Households report, via an in-home scanning device, the details of all purchases made during the survey period. Panelists in the sample are geographically dispersed throughout the country, and the survey is designed to be demographically balance. In particular, surveyed households are balanced across: age of household head(s); education of household head(s); occupation of household head(s); household income; household size; presence of children; race; whether Hispanic.

House prices are likely to differentially affect the consumption of homeowners and non-homeowners. Unfortunately, however, the panel data do not provide information on home ownership-status. Stroebel and Vavra (2014) explore the relationship between house prices and shopping behavior using the Consumer Panel Data, and they infer ownership from households' reported 'type of residence'. This variable reports whether a household lives in a one-, two-, or three-family house, and also whether the house is a condo or co-op. Homeowners are assumed to be those living in single-family, non-condo/co-op homes. Other households are assumed to be renters. The weighted-proportion of households living in single-family homes ranges from 0.71 to 0.74. From 2004 to 2015, the homeownership rate for the US as a whole fell from 69% to 63.7%.[2]

Households in the panel report purchases from every shopping trip conducted during their time in the survey. I aggregate each household's total consumption for the year. I drop an observation for a household in a given year if the household did not make one or more purchases in at least 10 months of that year. Although households are required to report every purchase that they make on all shopping trips, households may purchase goods that are not coded by Neilsen, and which do not make it into their reported expenditures for the year. The Kilts Center reports that the consumption goods featured in the Consumer Panel Data account for approximately 30 percent of all household consumption categories (for Marketing (2016)).

Usefully, geographic information for state, county, and zip-code are reported for each household. Using this information, I can match each household in the panel to the zip- and CBSA-level house price components. I drop households that cannot be matched to a CBSA.[3] In Appendix C, Table **??** reports the total number of households in the panel prior to and after filtering on total consumption and geographic area.

[2]Homeownership rate for the United States ('USHOWN'), from FRED.

[3]CBSAs cover both of the older Metropolitan Statiscal Area and Micropolitan Statisical Area designitions. Households that cannot be matched to one of these areas likely live in regions in which there is little, if any, house price information in the ZTRAX database.

4 House price heterogeneity

Different components of house prices may have different effects on consumption. In order to assess this, we first need to understand the relative magnitude of volatility associated with each component. For example, idiosyncratic or neighborhood level house price movements may have very large effect on consumption, however they may make up a small proportion of the total variance in house prices.

Several recent papers have considered the size of idiosyncratic house price movements. For example, Landvoigt et al. (2015) investigate house price movements in San Diego between 1997 and 2008. They estimate that the standard deviation of idiosyncratic price movements is around 8% at the beginning of the housing boom, and up to 14% in the housing bust. Giacoletti (2016) studies house prices in in Los Angeles, San Francisco, and San Diego from 2000 to 2012, and finds that idiosyncratic risk varies from 7 to 15% over this period. In both papers, the authors suppose that the initial price of a house prior to a sale is a proxy for the initial quality of the house. This quality may be associated with unobserved features of the house, or possibly local amenities. However, beyond controlling for initial quality, these papers do not consider how house price risk might be distributed across individual houses, neighborhoods, and cities. For example, Giacoletti (2016) shows that there are differences in idiosyncratic risk across the three cities studied, but does not consider how that risk might be related across them.

Consider a simple model of house price movements that attempts to identify the various components of house price movements. House price variation is attributed to aggregate movements in house prices, observable characteristics of the individual houses, as well as CBSA, zip-code, and idiosyncratic price components. Denote the log-price of a house i sold at time t in zip-code z in CBSA m as $p_{m,z,i,t}$. Then a simple model of these components is:

$$p_{m,z,i,t} = \beta_t X_{i,t} + u_t + v_{m,t} + w_{m,z,t} + \varepsilon_{m,z,i,t},$$

where $X_{i,t}$ are observable characteristics of an individual house, u_t is an aggregate price component, $v_{m,t}$ is a CBSA price component, $w_{m,z,t}$ is a zip-code level component, and $\varepsilon_{m,z,i,t}$ is an idiosyncratic component. The notation should make clear that individual houses belong to a particular zip-code, and zip-codes belong to a particular CBSA.

The observable house characteristics can be interpreted as components of a hedonic pricing model, but are more important for our purpoes to account for the composition effect of different houses being sold at different times and in different locations. This allows us to interpret the variance of house prices of otherwise similar houses. Of course, location is itself a characteristic of the house, but this is captured in the CBSA and zip-code price components, $v_{m,t}$ and $w_{m,z,t}$ respectively.

This model structure is reffered to formally in the econometrics literature as a nested error components regression model (see Baltagi et al. (2001)), or a multi-dimensional random effects model (see Balazsi et al. (2016)). Recently, Kaplan et al. (2016a) used this model form to study price dispersion across goods within and between stores. As noted there, in very large panel data settings, estimating these econometric models with maximum likelihood or panel data techniques may not be feasible. Instead, they propose a multi-stage GMM approach, which I follow here.

First, note that since we are not interested in conducting inference on the coefficients β_t and

μ_t, these can be estimated consistently via OLS.[4] Define the residual error component as

$$\hat{p}_{m,z,i,t} = p_{m,z,i,t} - \beta_t X_{i,t} - u_t$$

The CBSA component is then estimated as the within-CBSA mean of the residual error:

$$v_{m,t} = \frac{1}{n_{m,t}} \sum_{z \in m} \sum_{i \in z} \hat{p}_{m,z,i,t}$$

where $n_{m,t}$ is the number of house sales observed in MSA m at time t. Now define the after-CBSA component residual as:

$$\hat{\hat{p}}_{m,z,i,t} = \hat{p}_{m,z,i,t} - v_{m,t}$$

. Then the zip-code component is estimated as the within-zip-code mean of this residual:

$$w_{z,m,t} = \frac{1}{n_{z,t}} \sum_{i \in z} \hat{\hat{p}}_{m,z,i,t}$$

where $n_{z,t}$ is the number of house sales observed in zipcode z at time t. Finally, the idiosyncratic component defined as the residual:

$$\varepsilon_{m,z,i,t} = \hat{\hat{p}}_{m,z,i,t} - w_{z,m,t}$$

Consider first a variance decomposition of the components over time. The cross-sectional means of $v_{m,t}$, $w_{m,z,t}$, and $\varepsilon_{m,z,i,t}$ in any year t are zero. The variance of the idiosyncratic component is straightforward. For the CBSA and zip-code components I compute observations-weighted variances. This gives more influence to CBSA and zip-code components with high numbers of sales. The variance decomposition is presented in Figure 2. The CBSA variance dominates the other two components, with idiosyncratic variance dominating the zip-code level variance. However, there is some time variation in the components. In particular, the MSA component falls significantly between 2007 and 2009, and the zip-code component variance doubles between 2006 and 2009, and remains elevated thereafter.

Observe that there is significantly more time-variation in the CBSA component variance than the other components. This reflects large movements in particular regions of the US in recent years. For example, Figure 3 shows the CBSA components for all CBSAs with complete data from 2000 to 2016. CBSAs in regions that experienced significant booms during the housing cycle, such as San Francisco and San Diego, have been persistently higher than other CBSAs. Whereas CBSAs in other parts of the country, such as Michigan, have persistently declined relative to other CBSAs.

Figures 4a, 4b, and 4c show the cross-sectional autocorrelation functions of each of the price components. That is, I compute the autocovariance matrix in the cross-section for each price component, and then construct the autocorrelation function from this autocovariance matrix.[5] The order zero ACF of any process is equal to 1, so the year in which the ACF is 1 in all figures is the start year. Reading left-to-right, we can see the autocorrelation between that start year, and every other year back to 2000. Note that while the ACF for the CBSA and zip-code components relies on the CBSA and zip-code components constructed in each year, the ACF

[4]The standard errors of these estimates are biased, however.

[5]See Appendix D for details.

Figure 2: House price variance decomposition

Data source: ZTRAX.

Figure 3: CBSA price components

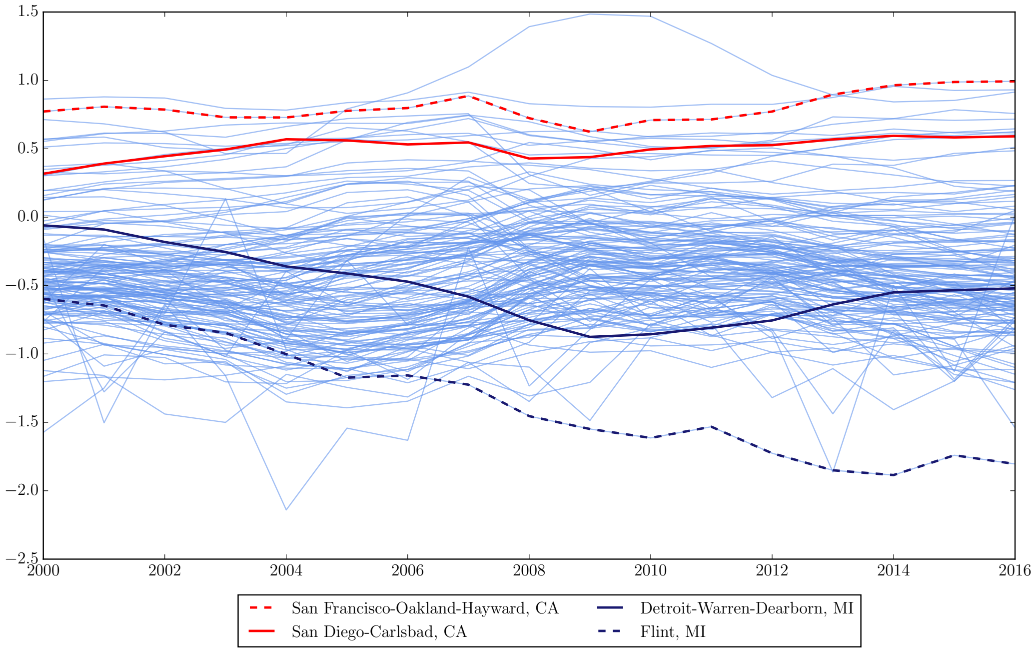

CBSA component of house prices for 150 CBSAs with complete data from 2000 to 2016. Data source: ZTRAX.

for the idiosyncratic component relies on repeat-sales. Different houses sell in different years, so I rely on an unbalanced panel of repeat sales to construct the covariance matrix in this case.

Figure 4a shows that the CBSA component is extremely persistent. The one-year autocorrelation in most years is above 0.99, while the three-year autocorrelation is often above 0.97. However, during the housing bust, the persistence of CBSA component fell dramatically. In 2008, the one- and three-year autocorrelations were just 0.93 and 0.86 respectively. This sudden change in persistence reflects a significant increase in the mean-reversion of the CBSA component during the housing bust. This helps explain the sudden fall in the variance of the CBSA component in 2008, as shown in Figure 2. The very high persistence of the CBSA component outside of the housing bust suggests it behaves like a random walk, which induces a large and volatile cross-sectional variance. However, the sudden drop in persistence during the housing bust induced a significant amount of mean reversion in CBSA level house prices, which dramatically decreased the variance of these prices.[6]

Figure 4b suggests the zip-code component is highly persistent, but less so than the CBSA component. And Figure 4c shows that there is very little persistence in the idiosyncratic component of house prices prior to the housing bust. However, from 2009 onwards, the slower decline in the ACF over several years suggests the appearance of some persistent component in idiosyncratic house prices. Giacoletti (2016) estiamtes a statsitcal model for idiosyncratic house price movements from 2000 to 2012 with persistent and transitory components, but finds no evidence of a persistent component. This would be consistent with the ACF of the idiosyncratic component prior to 2010, but seems less reasonable from 2010 onwards.

4.1 Statistical model of house price movements

Incomplete!

The next step in this research is to formally characterize and estimate a statistical process for each of the house price components. Given the previous findings, the statistical model ought to have the following properties. The CBSA component is clearly extremely persistent, perhaps even a random walk, but with periods of much lower persistence. Landvoigt et al. (2015) describe a model that allows for time-varying persistence, and find that the San Diego housing market faces periods of extremely high persistence, and other periods of mean reversion. This model can be written as:

$$v_{m,t} = \alpha_m + y^p_{m,t} \tag{1}$$

$$y^p_{m,t} = \rho_t y^p_{m,t-1} + \varepsilon_{m,t} \tag{2}$$

where α_m is a CBSA-specific fixed effect, $y^p_{m,t}$ is the persistent component, ρ_t is a time-varying persistence parameter, and $\varepsilon_{m,t} \sim \mathcal{N}(0, \sigma^2_m)$ is a shock to the persistent component.

The zip-code level component is less persistent, but its ACF function also displays periods of higher and lower persistence. Additionally, there is a sudden increase in the variance of the zip-code component during the housing bust. A decrease in persistence at the same time as an increase in variance suggest the presence of time-varying transitory shocks rather than time-varying persistence parameters. Following the persistent-transitory component models estimated for labor income (e.g. Blundell et al. (2008)), I write the zip-code component model as:

$$w_{m,z,t} = \alpha_z + y^p_{m,z,t} + \eta_{m,z,t} \tag{3}$$

$$y^p_{m,z,t} = \delta y^p_{m,z,t} + \varepsilon_{m,z,t} \tag{4}$$

[6]Note that this finding is somewhat in contrast to the results presented in Landvoigt et al. (2015), who find that for San Diego, the persistence of idiosyncratic risk increases significantly during the housing bust.

(a) CBSA autocorrelation function

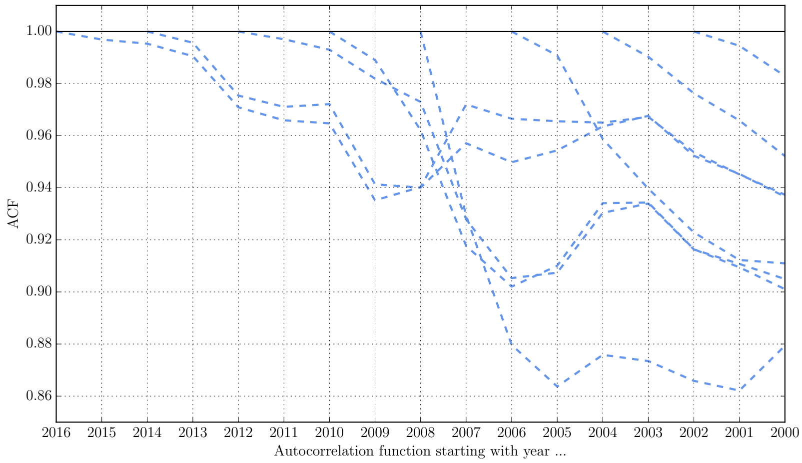

(b) Zip-code autocorrelation function

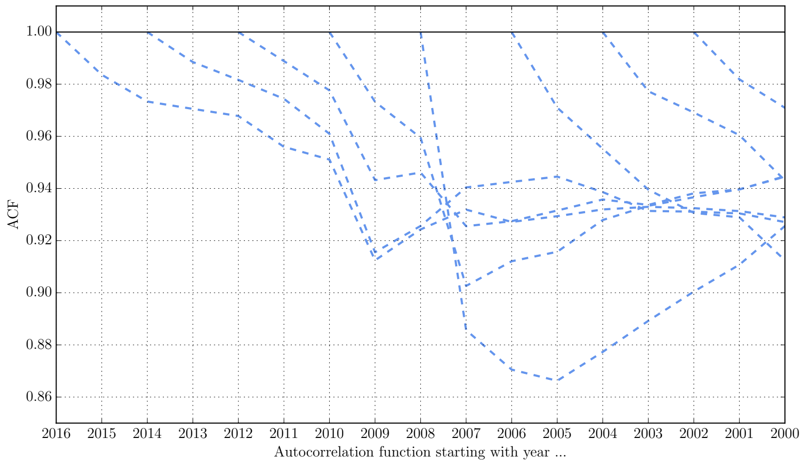

(c) Idiosyncratic autocorrelation function

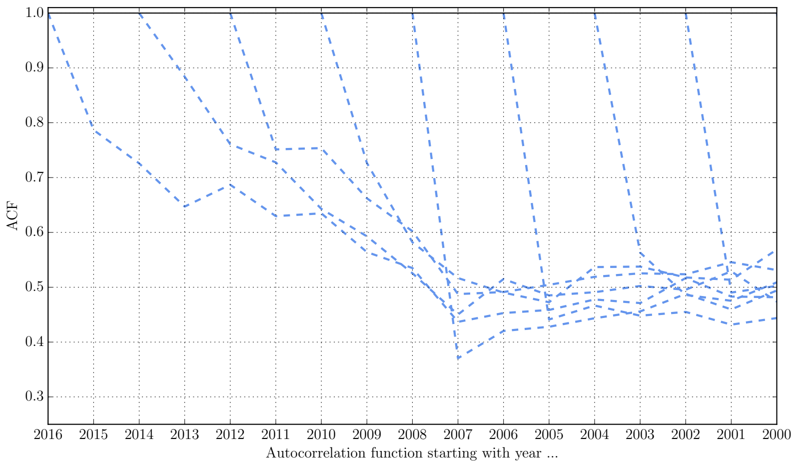

CBSA component of house prices for 150 CBSAs with complete data from 2000 to 2016. Data source: ZTRAX.

where α_z is a zip-code-specific fixed effect, $y^p_{m,z,t}$ is the persistent component, $\eta_{m,z,t} \sim \mathcal{N}(0, \sigma^2_{z,t})$ is a transitory shock with time-varying variance, δ is a persistence parameter for the persistent process, and $\varepsilon_{m,z,t} \sim \mathcal{N}(0, \sigma^2_{\varepsilon,z})$ is an IID shock to the persistent process.

Finally, the idiosyncratic component is largely not a persistent process, however the variance of the component suggests some time-variation in idiosyncratic housing risk. I write this process as

$$\varepsilon_{m,z,i,t} = \alpha_i + \eta_{m,z,i,t} \tag{5}$$

where α_i is a house-specific fixed effect, $\eta_{m,z,i,t} \sim \mathcal{N}(0, \sigma^2_{i,t})$ is an IID shock, with time-varying variance. Note that the parameters for this process are easily identified using the idiosyncratic component of prices associated with repeat sales. In particular the covariance between the same house sold in two different periods is given by $\sigma^2_{\alpha_i}$, while the cross-sectional variance of $\varepsilon_{m,z,i,t}$ is $\sigma^2_{\alpha_i} + \sigma^2_{i,t}$.

4.2 The effect of house price components on consumption

Many previous studies investigate the effect of house prices on consumption (e.g. Campbell and Cocco (2007), Mian et al. (2013), Kaplan et al. (2016b), Aladangady (2017)). Typically, studies make use of cross-sectional variation in house prices and consumption to identify the effect of house prices. The typical cross section for the US is zip-code, county, or CBSA/MSA level. Two identification problems are central to these studies. First, there may be reverse causation from consumption to house prices. This occurs if higher consumption in a location increases employment or income, which increases demand for housing and thus increases prices. Second, endogenous variables may affect both house prices and consumption. For example, a positive productivity shock to an MSA or region increases incomes which increase consumption, but the increase in incomes also increases house prices through demand.

The first identification problem does not affect regressions using household level panel data. Because households are small relative to the market in their location, we can assume that an increase in an individual household's consumption does not lead to an increase in house prices in that region. The second identification problem does affect regressions using panel data, however. So far I have not addressed this second problem with the findings presented here. A next step in the research is to provide an instrument for house price movements at the CBSA level. A typical approach is to make use of the Saiz (2010) housing supply elasticities for various MSAs. Aladangady (2017) is an example of this approach using household-level panel data from the Consumer Expenditures Survey. Campbell and Cocco (2007) construct a synthetic panel using repeated cross-sectional data, and instrument for house prices using changes in a proxy for unemployment, house price changes, and the second lag of changes in income.

I consider a model of the following form:

$$\log c_{m,z,i,t} = \alpha_i + \beta x_{i,t} + \delta q_{m,t} + \gamma_a \log p_t + \gamma_m \log p_{m,t} + \gamma_z \log p_{z,t} + \varepsilon_{m,z,i,t} \tag{6}$$

where $c_{m,z,i,t}$ is total consumption expenditure of a household i living in CBSA m and zip-code z, α_i are household level fixed effects, $x_{i,t}$ are household-specific observable characteristics, $q_{m,t}$ are CBSA-level observables, p_t is the aggregate house price, $p_{m,t}$ is the CBSA-level house price, and $p_{z,t}$ is the zip-code level house price. The CBSA-level observables are CBSA-level real personal income per capita, and the unemployment rate.[7] The observable household demographic characteristics are: household size, household income, the age of the household head, the education level of the household head, marital status, and race. All nominal variables are deflated by the CPI.[8][9]

At the quarterly level, a particular zip-code may report very few house sales. Dropping these zip-codes would result in a loss of information, in particular about household-level responses to aggregate and CBSA-level house prices. Instead, I report the results of weighted least squares regressions, where the weight for each household-observation is given by the multiple of the household-level projection factor provided in the Kilts Consumer Panel and the number of sales in the household's zip-code in that quarter. All standard errors are clustered by zip-code and quarter.

I run the regression using two sets of house prices. The first set of results, reported in 1 uses average house prices across the entire country, across each CBSA, and across each zip-code. To the extent house prices are correlated across CBSAs and zip-codes, the prices in these regressions are correlated. However, using these prices provides the simplest test of the effect of different house price components on consumption.

The second set of results, reported in table 2, uses the house price components computed in section 4. Thus the aggregate component is average house prices, controlling for the composition house sales in each year. The CBSA component is the average price deviation from aggregate house prices, in a given CBSA. And the zip-code component is the average price deviation from CBSA-level house prices, in a given zip-code. This set of results provides a first look at the effect of the various levels of house price shocks. Although these are preliminary results and do not yet control for endogeneity or provide a structural interpretation of the shocks, the are indicative of the intutition discussed in the Introduction.

In each of Tables 1 and 2, columns 1 through 3 report the relationship between house prices and consumption for each of the price components on their own. Column 4 reports the results when including all three price components. And column 5 presents results controlling for implied homeownership status.

Table 1 shows that aggregate house prices have a strong association with individual level consumption, even when other controls are included. Because the indepdent variable is household-level consumption, there is no effect of reverse causality in this relationship, although endogeneity remains a problem. Since national house prices are strongly correlated with the business cycle (particularly during the 2004-2015 period), aggregate shocks driving the business cycle affect both aggregate house prices and individual consumption. Including controls for CBSA-level unemployment and personal income mitigates some of this endogeneity, alhough aggregate shocks may influence consmption through many channels such as credit conditions,

[7]Real personal income per capita is available from the BEA Regional Income accounts. Unemployment is available from the BLS county-level unemployment statistics. Unemployment is aggregated up to CBSA-level using the cross-walk provided by the NBER. See the appendix for more data details. Each of these variables is only available at the CBSA-level at annual frequency. For the quarterly specification, I linearly interpolate these values across quarters.

[8]I use the seasonally adjusted CPI for all urban consumers. Source, FRED (code: 'CPIAUCSL').

[9]In ongoing work, it will be useful to control for other observables. For example, the Current Population Survey provides demographic information at zip-code level, such as age, education, race, income, and homeownership. Although I control for these variables at the individual level, neighborhood clustering on these characteristics may be a confounding factor which is only imperfectly controlled for with individual-level fixed effects and zip-code level standard error clustering. I can also include other aggregate-level variables to help control for endogeneity between aggregate the aggregate and CBSA-level house price moements and consumpton. For example, national unemployment, aggregate income, credit conditions, etc.

beliefs about future employment or income prospects, and so on.

Column 2 shows that the relationship between CBSA-level house prices and consumption is much weaker than for aggregate house prices. Moreover, columns 4 and 5 show that the relationship between CBSA-level house prices and consumption disappears once the other components of house prices are controlled for.

Column 3 suggests that although the relationship between zip-code level house prices and consumption is weaker than the aggregate relationship, it may be stronger than the CBSA-level relationship. Columns 4 and 5 show that the zip-code level house price effect remains even when controling for the other house price components. As discussed, this relationship may reflect the fact that households respond more strongly to neighborhood level house price movements than city level movements. As shown in Figure 1, households are much more likely to move within counties/across neighborhoods than they are to move across cities. As such, neighborhood level price movements may have stronger wealth effects on consumption than city-level price movements.

Although the results are weaker, Table 2 confirm the findings reported in in Table 1. Namely, aggregate price level movements have a stronger association with consumption than either CBSA or zip-code level price movements, and zip-code level price movements are more strongly associated with consumption than CBSA-level movements. Note, that these regressions examine the effect of price movements at each level that are exogenous to each other. That is, the CBSA-level house price movement is in addition to any aggregate house price movement. Similarly, zip-code level price movements are in addition to aggregate- and CBSA-level price movements. Thus the results lend stronger weight to the interpretation that zip-code-specific house price movements have their own effects on consumption, independent of aggregate or city-level prices.

Campbell and Cocco (2007) investigate the link between house prices and consumption in the UK. One of their regression specifications considers whether there is a link between regional house prices and consumption, over and above the effect of national house prices. In their OLS results, but not the IV specificiation, they show that that there is such a link, but that the relationship is weaker than the relationship between national house prices and consumption.

Table 1: Effect of house prices on houehold-level consumption: simple house price aggregates

Dependent variable:	log(real consumption expenditure)				
	(1)	(2)	(3)	(4)	(5)
$\log(p_t)$	0.355***			0.303***	0.295***
	(0.052)			(0.048)	(0.050)
$\log(p_{m,t})$		0.085***		−0.012	−0.020
		(0.024)		(0.021)	(0.022)
$\log(p_{z,t})$			0.097***	0.077***	0.093***
			(0.015)	(0.014)	(0.019)
$\log(Y_{m,t})$	−0.330***	−0.423***	−0.435***	−0.372***	−0.372***
	(0.059)	(0.052)	(0.055)	(0.054)	(0.054)
$\log(U_{m,t})$	−0.230	−1.819***	−1.660***	−0.044	−0.041
	(0.345)	(0.343)	(0.241)	(0.385)	(0.387)
$\mathbb{1}_{owner} \times \log(p_t)$					0.012
					(0.014)
$\mathbb{1}_{owner} \times \log(p_{m,t})$					0.012
					(0.020)
$\mathbb{1}_{owner} \times \log(p_{z,t})$					−0.023
					(0.016)
Observations	936,363	936,363	936,363	936,363	936,363
R^2	0.704	0.703	0.703	0.704	0.704

All specifications include household-level controls for size, income, age, education, marital status, and race. All specifications also include fixed effects at the household level. All specifications use regression weights computed as the multiple of household-level projection factors and quarterly zip-code sales. Standard errors, clustered by zip-code and quarter, are reported in parentheses. Significance at 1 (***), 5 (**), and 10 (*) percent levels.

Table 2: Effect of house prices on houehold-level consumption: house price components

Dependent variable:	log(real consumption expenditure)				
	(1)	(2)	(3)	(4)	(5)
$\log(p_t)$	0.079***			0.081***	0.080***
	(0.008)			(0.008)	(0.008)
$\log(p_{m,t})$		−0.040		0.043	0.052
		(0.033)		(0.034)	(0.035)
$\log(p_{z,t})$			0.077***	0.056***	0.073***
			(0.019)	(0.015)	(0.021)
$\log(Y_{m,t})$	−0.285***	−0.343***	−0.381***	−0.321***	−0.321***
	(0.056)	(0.052)	(0.056)	(0.052)	(0.052)
$\log(U_{m,t})$	0.365	−2.443***	−2.417***	0.442	0.444
	(0.223)	(0.293)	(0.270)	(0.285)	(0.287)
$\mathbb{1}_{owner} \times \log(p_t)$					0.001
					(0.001)
$\mathbb{1}_{owner} \times \log(p_{m,t})$					−0.012
					(0.016)
$\mathbb{1}_{owner} \times \log(p_{z,t})$					−0.025
					(0.023)
Observations	966,997	966,997	966,997	966,997	966,997
R^2	0.705	0.704	0.704	0.705	0.705

All specifications include household-level controls for size, income, age, education, marital status, and race. All specifications also include fixed effects at the household level. All specifications use regression weights computed as the multiple of household-level projection factors and quarterly zip-code sales. Standard errors, clustered by zip-code and quarter, are reported in parentheses. Significance at 1 (***), 5 (**), and 10 (*) percent levels.

5 Model

At this stage of the research, the model serves two purposes. First, we want to build a model that allows for house price hetereogeneity at multiple levels e.g. aggregate and neighborhood, or aggregate and idiosyncratic. Second, we want to explore the elasticities of consumption that the model generates with respect to different kinds of house price movements.

The model is a finite-horizon, partial equilibrium, life-cycle model. Households make decisions about consumption, liquid assets, rental services, housing assets, and mortgages, subject to fluctuations in two levels of prices. These prices are interpreted as some aggregate level from the household's perspective, e.g. economy-wide house price movements or CBSA level movements, as well as some lowe-level price movements, e.g. neighborhood or idiosyncratic movements.

Households live from the age of 21 and die with certainty at age 80. Households work until age 64, and then retire at 65. The model period is one year.

House prices

Households are subject to two house price components. First, there is an aggregate house price component, P_h. Households expect this component to remain constant forever. Unexpected shocks may move P_h, however after the shock households expect the new price to persist forever. The aggregate component of prices applies to all houses that are bought and sold by households. Rent prices are also a constant fraction of this aggregate component of house prices. Thus, when an unexpected shock hits the aggregate component, rents rise in lock step.

Second, there is an lower level (neighborhood/idiosyncatic) house price component, P_z. This component is assumed to follow an AR(1) over time. The persistent, but non-permanent, nature of these prices reflects the findings in the empirical section of the paper: CBSA-level house prices are often permanent, while neighborhood level components are persistent, but not permanent. Households are assumed to live in a neighborhood until they sell their house (renters pay a constant fraction of the aggregate house price, regardless of where they live). At this time, the sale of the house receives the current neighborhood house price component P_z. New house purchases are assumed to come from other neighborhoods. Households are likely to choose new neighborhoods similar to their current neighborhoods. Thus, house prices in new neighborhoods are likely to be (perhaps imperfectly) correlated with the house price in the current neighborhood. Thus, I assume that the new neighborhood price component is given by $P_z \exp(\eta)$, where $\eta \sim \mathcal{N}(0, \sigma_\eta^2)$ is an IID shock to the current neighborhood price. The larger is the standard deviation of the IID shock to neighborhood prices, the less correlated are neighborhood level hosue prices. I explore different degrees of correlation between neighborhoods that households are likely to purchase from.

When a household moves neighborhoods and experiences the IID price shock, this affects the future path of house prices that the household faces, since the shock enters the AR(1) neighborhood price process. Thus, the neighborhood house price process can be thought of as an AR(1) subject to a continuous normally distrbuted shock as well as a second, independent normally distributed shock with an arrival rate determined endogenously by the household's decision to move houses/neighborhoods.

Note that this mechanism for neighborhood level prices simplifies possible neighborhood/location-choice problem. Rather than have households choose from among many possible neighborhoods, each with their own characteristics, amenities, and so on, I assume that the only difference across neighborhoods is house prices, which are not known until the time of purchase.

Later, I want to add a choice or shock that determines whether to move neighborhoods or not. That is, households can always move house and perhaps stay within the same neighborhood, however circumstances (or choice) may force them to switch neighborhoods discretely. A "moving" shock that affects housing utility is one way to get at this.

Income

Income during working life consists of a deterministic function of age and a stochastic autoregressive process. The deterministic process is a quartic function of age. Loosely following Kaplan, Mitman, and Violante (2017), the deterministic component grows by a factor of 3 from age 21 to 50, and then declines slowly until retirement. From retirement at age 65, agents receive a pension equal to 40% of total income at age 64. This is a proxy for dispersal from retirement accounts accumulated during working life. Note, this also means that income during retirement is certain. Income during working life, m_j^w, can be expressed as:

$$\log m_j^w = \chi_j + \log y_j$$
$$\log y_j = \rho \log y_{j-1} + \varepsilon_{y,j}$$

The AR(1) process applies only during working life, and we assume that $\varepsilon_y \sim \mathcal{N}(0, \sigma_y^2)$ and the initial draw for income at age 21 comes from the stationary distribution, $y_0 \sim \mathcal{N}(0, \frac{\sigma_y^2}{1-\rho^2})$.

Bequests

Households leave bequests due to a warm-glow motive. Bequests are value by households, but are also luxury goods from their persepective. Since households are not attached to each other dynastically, I assume that in the initial period of life households receive liquid assets drawn from log-normal distribution, $\log(a_1) \sim \mathcal{N}(\mu, \sigma_a^2)$.

Consumption

Household consumption is a bundle of non-durable goods and housing services. Non-durable goods are the numeraire, while housing services can either be purchased as rental housing services or owner-occupied housing. For computational tractability, I assume that both rental services and owner-occupied housing are chosen from finite sets, $\mathcal{S}$ and $\mathcal{H}$. Note that since the rental housing choice is a static problem, it can be solved as a continuous choice variable in each period in which the household chooses to rent (see Gorea and Midrigan (2017)). However, when owner occupied housing is chosen from a finite set, the household often finds a continuous rental choice optimal if adjacent housing options are too far apart. A solution to this problem is to increase the size of the housing set, however this increases computational burden. Kaplan et al. (2017) choose the elements of the rental and housing sets such that the sets overlap and the rental options consistent of the smallest few housing options. This ensures that when households switch from rentals to housing, they purchase similar sized houses to the ones they recently rented. This means that shocks driving changes homeownership do not increase demand for the overall stock/size of housing, but simply change the composition of owners and renters.

House sales are subject to a sales cost proportional to the value of the house. Houses depreciate at rate $delta_h$, so households pay maintenance costs proportional to the value of the house lived in last period in order to maintain housing value.

Mortgages

Households may hold mortgage debt against the value of the house they own. Mortgages are long-term debt contracts. At origination, mortgages are subject to a maximum loan-to-value ratio constraint as well as a payment-to-income ratio constraint. The constraints are, respectively:

$$b' \leqslant \theta P_h h'$$
$$b' \leqslant \lambda_\pi m$$

The interaction between the two constraints is explored in Greenwald (2016).

Although there is no default in the model, I assume that high LTV mortgages are considered to be risky, and so a penalty rate applies above the a conforming mortgage LTV limit. Above the limit θ_c, for a mortgage of size b' the household receives funds qb', where $q \leqslant 1$ is the penalty price. The price is given by: $q = (1 + \gamma(\theta_c - z))$, where z is the LTV ratio.[10]

Mortgages are amortized over the remaining life of the household, with minimum payments required in every remaining period unless the mortgage is refinanced or the house is sold. The constant amortization formula for mortgage payments is given by:

$$d_j(b) = \frac{r_b(1+r_b)^{J+1-j}}{(1+r_b)^{J+1-j} - 1} b$$

which yields constant payments for the life of the loan.[11] Mortgage balances evolve according to $b' = (1 + r_b)b - d(b)$. Households can repay the mortgage more quickly than the schedule given by the constant amortization formula, however this requires refinancing which is costly. Refinancing may occur at any time, and is subject to the loan-to-value ratio above. Both new mortgages and refinancing of mortgages are subject to a fixed cost, reflecting origination costs.

Preferences

Household's maximize lifetime utility by choosing liquid assets, consumption, and rental services, and warm-glow bequests. Lifetime utility is:

$$\mathbb{E}\left[\sum_{j=1}^{J} \beta^{j-1} u(c_j, s_j) + \beta^J v(\alpha_J)\right] \tag{7}$$

[10]For now the mortgage penalty is linear in the LTV. Hedlund (2016) and Kaplan et al. (2017) explicitly microfound the mortgage penalty function using an individual borrower's default probability given current state variables.

[11]Note that power $J + 1 - j$ ensures that households make mortgage payments in every period of life, including the final period J. In the final period, then, the final mortgage payment is $(1 + r_b)b$, which is the entirety of remaining principal plus interest. If the power were expressed as $J - j$, households refinancing or originating new mortgages in period $J - 1$ would face no payments in period J and would repay the mortgage out of assets remaining at death.

The utility function is given by:

$$u(c,s) = \frac{(c^{\chi}s^{1-\chi})^{1-\sigma}}{1-\sigma} \tag{8}$$

where χ is the share of consumption in non-housing services. Housing services s depend on whether the household is renting or a homeowner. Renters receive housing services equal to the size of the rental unit. Homeowners receive housing services equal to ϕh with $\phi > 1$. This reflects the extra utility of owned housing services.

The bequest function is given by:

$$v(a) = \psi\frac{(\alpha + \underline{\alpha})^{1-\sigma}}{1-\sigma} \tag{9}$$

where ψ is the strength of the bequest motive, and $\underline{\alpha}$ reflects the extent to which bequests are a luxury good.

5.1 Household's problem

Households must decide between renting and owning a home. Conditional on owning a home, households may choose to keep their current housing stock, adjust the housing stock, or refinance the mortgage. The state variables for the household are liquid assets a, mortgage principal b, house size h, transitory income y, house price P_h. I express the problem in terms of those state variables here, however it is convenient to solve the model in the state variables: cash on hand x, loan-to-value ratio z, housing size h, transitory income y, and house price P_h. See the appendix for details. Denote the state vector $\mathbf{s} = \{x, z, y, P_h\}$.

The discrete choice of the household at age $j < J$:

$$V_j(\mathbf{s}) = \max\left\{V_j^S(\mathbf{s}), V_j^N(\mathbf{s}), V_j^A(\mathbf{s}), V_j^R(\mathbf{s})\right\}$$

In the final period of life, J, the household may only rent or keep the house they are in (i.e. they cannot adjust their housing size or refinance their mortgage):

$$V_J(\mathbf{s}) = \max\left\{V_J^S(\mathbf{s}), V_J^N(\mathbf{s})\right\}$$

Renter's problem

Household renters purchase rental housing services, consume non-durable goods, and save in liquid assets. If a renter enters the period owning a house, the house is sold and any outstanding mortgage is repaid. The household carries forward no housing assets or mortgage debt. The renters value function at age j is written as:

$$\begin{aligned}
V_j^S(\mathbf{s}) = \max_{c,s,a'} \; & u(c,s) + \beta\mathbb{E}(V_{j+1}(\mathbf{s}')) \\
\text{s.t.} \quad & c + a' + P_r s + (1 + r_b)b = m_j + a(1 + r) + (1 - \delta_h - F_s)P_h h - b(1 + r_b) \\
& a' \geqslant 0, b', h' = 0 \\
& s \in \mathcal{S}
\end{aligned}$$

Homeowner's problem

A homeowner that chooses not to adjust its housing stock will live in the same sized house, consume non-durable goods, save in liquid assets, and make a required mortgage payment. The value function for a household that does not adjust at age j is:

$$\begin{aligned} V_j^N(\mathbf{s}) &= \max_{c,a'} u(c,\psi h) + \beta \mathbb{E}(V_{j+1}(\mathbf{s}')) \\ \text{s.t.} \quad & c + a' + \delta_h P_h h + d_j(b) = m_j + a(1+r) \\ & a' \geqslant 0 \\ & b' = (1+r_b)b - \pi \end{aligned}$$

where $d_j(b)$ is the required payment on the mortgage.

A homeowner that chooses to adjust its housing stock will sell its current house subject to a proportional selling cost, repay the outstanding mortgage balance, purchase a new house, choose a new mortgage subject to the LTV and PTI constraints, consume non-durable goods, and save in the liquid asset. The value function for a household that adjusts at age j is:

$$\begin{aligned} V_j^A(\mathbf{s}) &= \max_{c,a',h',b'} u(c,\psi h') + \beta \mathbb{E}(V_{j+1}(\mathbf{s}')) \\ \text{s.t.} \quad & c + a' + P_h h' + b(1+r_b) + F_m = m_j + a(1+r) + P_h h(1-\delta_h - F_s) + qb \\ & a' \geqslant 0 \\ & b' \leqslant \theta P_h h', \quad d(b') \leqslant \lambda_d m_j \\ & h' \in \mathcal{H} \\ & z = b'/P_h h' \\ & q = \begin{cases} (1+\gamma(\theta_c - z)), & \text{if } z \geqslant \theta_{nc} \\ 1, & \text{otherwise} \end{cases} \end{aligned}$$

A homeowner that chooses to refinance its mortgage repay the outstanding mortgage balance, and choose a new mortgage subject to the LTV and PTI constraints, consume non-durable goods, and save in the liquid asset. The value function for a household that refinances at age j is:

$$\begin{aligned} V_j^R(\mathbf{s}) &= \max_{c,a',b'} u(c,\psi h) + \beta \mathbb{E}(V_{j+1}(\mathbf{s}')) \\ \text{s.t.} \quad & c + a' + b(1+r_b) + \delta_h P_h h + F_m = m_j + a(1+r) + qb \\ & a' \geqslant 0 \\ & b' \leqslant \theta P_h h, \quad d(b') \leqslant \lambda_d m_j \\ & z = b'/P_h h \\ & q = \begin{cases} (1+\gamma(\theta_c - z)), & \text{if } z \geqslant \theta_{nc} \\ 1, & \text{otherwise} \end{cases} \end{aligned}$$

The final period problem

In the final period of life, at age J, the household can either sell any housing it holds and become a renter, or hold onto the house to be liquidated in the following period. A household that chooses to rent solves the following problem:

$$
\begin{aligned}
V_J^R(\mathbf{s}) = \max_{c,s,a'} & \; u(c,s) + \beta\psi v(\alpha) \\
\text{s.t.} \quad & c + a' + P_r s + (1+r_b)b = m_J + a(1+r) + (1-\delta_h - F_s)P_h h \\
& a' \geqslant 0 \\
& s \in \mathcal{S} \\
& \alpha = a'(1+r)
\end{aligned}
$$

where the bequest is simply the household's remaining liquid assets.

If, instead, the household chooses to remain in its house, it solves the following problem:

$$
\begin{aligned}
V_J^N(\mathbf{s}) = \max_{c,a'} & \; u(c,\psi h) + \beta\psi v(\alpha) \\
\text{s.t.} \quad & c + a' + \delta_h P_h h + d_J(b) = m_J + a(1+r) \\
& a' \geqslant 0 \\
& d_J(b) = (1+r_m)b \\
& \alpha = a'(1+r) + (1-\delta_h - F_s)P_h h
\end{aligned}
$$

where the final mortgage payment clears any remaining principal and interest, and the bequest is final liquid assets and the resale value of the house.

Solving the lifetime problem: Now we solve the model recursively, iterating backwards from the final period. We can write the Bellman equation:

$$
V_j(a,y) = \max_{c_j,s_j,a'_j} u(c_j,s_j) + \beta\mathbb{E}(V_{j+1}(a'_j,y')) \tag{10}
$$

Notice that starting from the final period, we always have the solution to the value function of the following period. E.g. the final value function is known:

$$
V_J(a,y) = u(c_J,s_J) + \beta^J v(a'_J)
$$

Model parameters

Table 3 shows a preliminary parameterization of the model. The model period is one year. At this stage, these parameters are plausible initial values, with calibrated and estimated parameters to come with future research. The most important parameters for this model are those governing aggregate and neighborhood level house prices.

Although the model is only a partial equilibrium model, we can solve for a stationary distribution of the model. I do this using the histogram method of Young (2010). Figure 7 shows several cross-sectional moments of the model in the stationary distribution. Note that the model produces reasonable life-cycle profiles for income, consumption, housing size, homeownership, and housing leverage. It will be interesting to try and match the model to the number of movers in the CPS data, as shown in Figure 1.

Parameter	Symbol	Value	Source
Preferences			
Discount factor	β	0.91	
Risk aversion	σ	2	
Non-durables consumption share	χ	0.85	
Borrowing constraint	$\underline{a}$	0	
Housing			
Maximum LTV ratio, conforming	θ_c	0.8	
Maximum LTV ratio, non-conforming	θ_{nc}	0.9	
House sale fixed cost	F_s	0.06	Kaplan et al. (2017)
New mortgage cost	F_m	0.0385	Gorea and Midrigan (2017)
House depreciation	δ_h	0.015	
Rent-to-house price ratio	λ_{rtp}	0.075	Zillow
Interest rates			
Risk-free rate	r	0.015	
Mortgage interest rate	r_b	0.025	
Exogenous idiosyncratic processes			
Transitory income persistence	ρ_y	0.938	Gorea and Midrigan (2017)
Transitory income std. dev.	σ_y	0.20	Gorea and Midrigan (2017)
Transitory income initial std. dev.	$\sigma_{y,0}$	0.20	
Initial assets distribution, mean	$\mu_{a,0}$	- 1	
Initial assets distribution, std. dev.	$\sigma_{a,0}$	1	
House prices			
Neighborhood-level persistence	ρ_z	0.85	
Neighborhood-level std. dev.	σ_z	0.02	
Neighborhood-level IID moving shock std. dev.	σ_η	1e-5	

Table 3: Model parameters (preliminary)

Figure 5: Cross-sectional moments in the stationary distribution

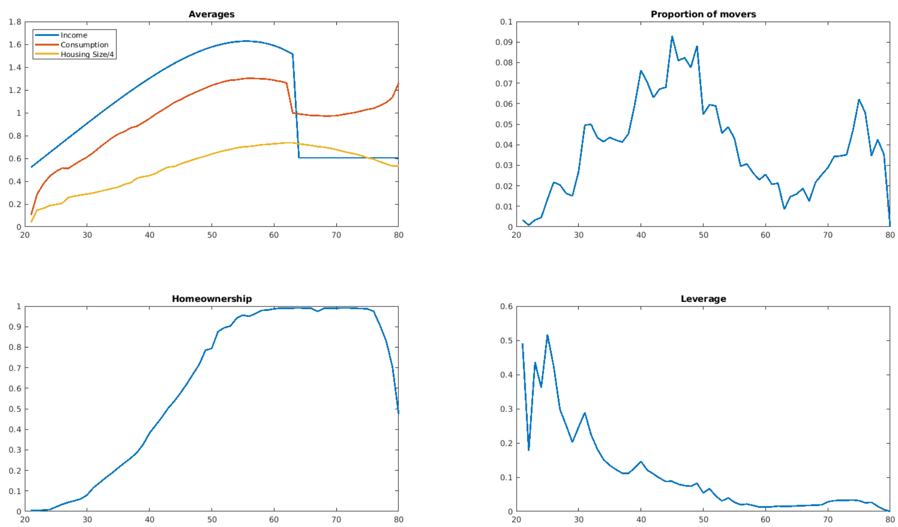

Figure shows several moments of the model across age, given in the stationary distribution of the model.

Consumption responses to house prices in the model

In this section I show in the model the elasticity of consumption with respect to different house price component movements. Consumption elasticities in the model are given by $\frac{\partial c_j(\mathbb{s};P_z,P_h)}{\partial P_z}$ and $\frac{\partial c_j(\mathbb{s};P_z,P_h)}{\partial P_h}$. I compute these elasticites in the model via a simulation. For a given household, I simulate two paths: one with the price held constant, and another with exactly the same set of other shocks, but with a shock to the price at time T. For each age j, I simulate 500 households with the price shock occuring at that age j. This yields a distribution of household ages and other state variables subject to the shocks. I then compute the elasticities by age and by loan-to-value ratio, by averaging over all agents of a given age or with a given loan-to-value ratio.

For the aggregate price shock, P_h, there are two interesting findings. First, the model shows that younger households tend to have weaker responses to the price shock than older households. This is consistent with the argument in Buiter (2008) that younger households who face higher lifetime housing costs following a price shock have small or negative responses, while older households who plan to downsize or sell and who have shorter horizons have larger responses. It is also consistent with the evidence in Campbell and Cocco (2007), who shows that younger households (and renters) in the UK have much smaller, often negative, consumption responses to house prices.

Second, high LTV households, in particular those very close to the non-conforming mortgage limit, have stronger responses to house price shocks. This is because households with loan-to-value ratios near the maximum LTV ratio are relatively constrained in that they cannot easily refinance their mortgage. An increase in house prices relaxes the LTV constraint for these households, which increase their ability to borrow and consume. The higher consumption response to house prices of these consumers resembles the collateral constraint effects discussed in the literature (see Kaplan et al. (2017), Berger et al. (2015), Greenwald (2016)).

Figure 6: Aggregate house price elasticity of consumption by LTV

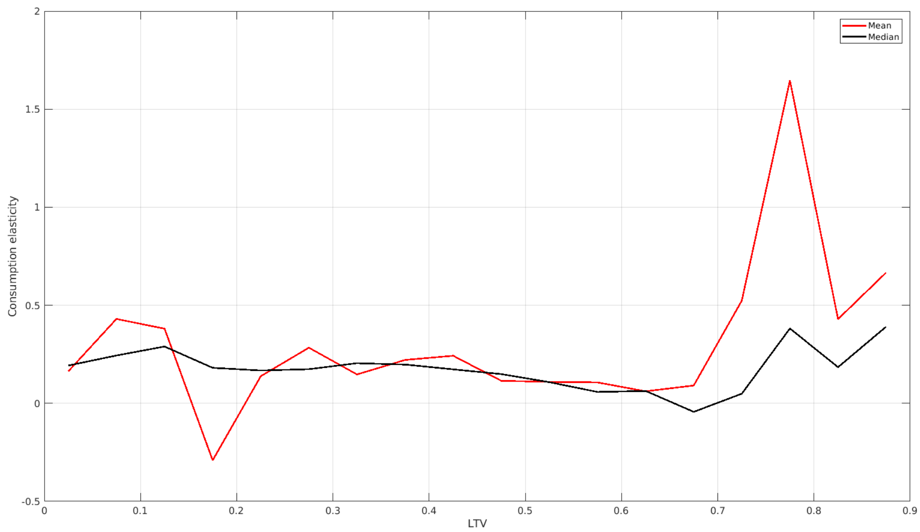

Figure 7: Aggregate house price elasticity of consumption by age

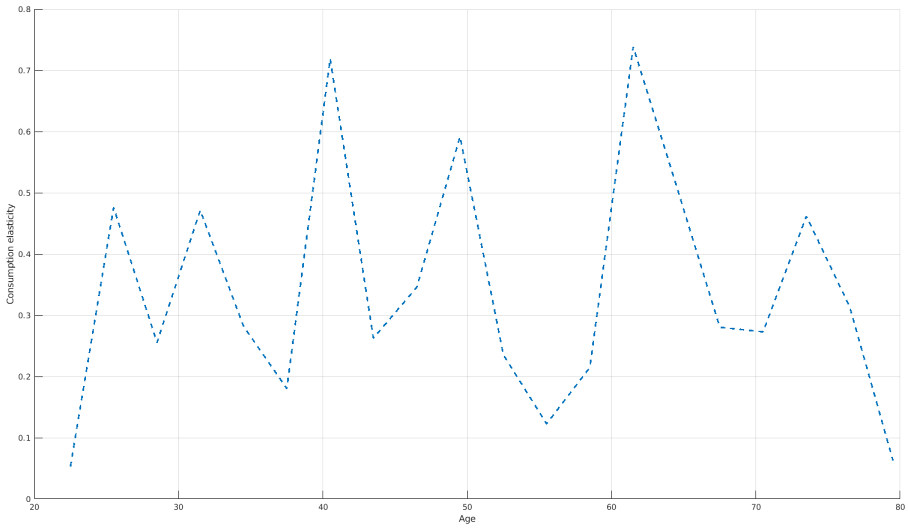

6 Work in progress

The paper is obviosly still very incomplete. This section provides a brief list of items left to complete:

- Estimate a statistical model of the various components of house prices.
- Relate this statistical model to household level consumption responses.
- Both of the above can be done in the framework of the consumption insurance literature (e.g. Blundell et al. (2008))
- Complete the structural lifecycle model of consumption responses to house prices.
- I still need to complete the consumption response to neighborhood level house price shocks, and I intend to add more detail to the cross- and within-neighborhood moving decisions.

7 Conclusion

In this paper, I investgiate the effect of movements in different components of house prices on consumption. Using a new dataset on individual house transactions from across the US, I show that city-, neighborhood-, and idiosyncratic-risk all contribute to hosue price movements. City-level effects tend to dominate the volatility of house prices, although the idiosyncratic component plays a large role, and the volatility of neighborhood-level prices increased sharply during the housing bust.

As a first empirical look at the effect of house prices on consumption, I match the aggregate, city, and neighborhood components of house prices to household consumption panel data. I find that consumption is most strongly associated with the aggregate comoponent of house prices, however consumption is also significantly associated with the neighborhood level component of house prices. In contrast, the city-level component of house prices is only weakly associated with consumption, and the effect is weaker than for the neighborhood level component.

These findings are consistent with the view that households respond more strongly to the component of house prices that affects their future wealth. Since households are less likely to move across cities than they are to move witin cities, the within-city component of house prices is more relevant for household wealth than the city-wide component of house prices.

In order to explore this intuition further, I build a strucutral life-cycle model in which agents face realistic housing decisions. In particular, households can own or rent; they can choose their housing size; they can take out long-term mortgages against the value of their house; they can refinance their mortgage; and they face both city-level and neighborhood level house price risk. In the model, I find that the average elasticity of consumption with respect to unexpected, permanent city level house price shocks is around 0.2. However, there is sigificant heterogeneity in this elasticity both over age and over the loan-to-housing value distribution.

A Data sources

This section documents the sources of data used in the paper.

- Household equity and income data is retrieved from the Panel Study of Income Dynamics at http://psidonline.isr.umich.edu/.
- Household equity and income data is also retrieved from the Survey of Consumer Finances at https://www.federalreserve.gov/econres/scfindex.htm.
- Real residential property prices and the CPI-U come from FRED at https://fred.stlouisfed.org/.
- Regional income and population data comes from the Bureau of Economic Analysis at https://www.bea.gov/regional/.
- County names within Metropolitan statistical areas can be found at https://www.bea.gov/regional/docs/msalist.cfm and the NBER's cross-walk data: http://www.nber.org/data/cbsa-fips-county-crosswalk.html.
- Regional unemployment data comes from the Bureau of Labor Statistics at https://www.bls.gov/lau/data.htm.
- Panel consumption data comes from the Kilts Neilsen HomeScan survey. Access to the data is given by Chicago Booth. Informaton can be found at:...
- House price data comes from Zillow's ZTRAX Assessment and Transaction Database. This data was provided by Zillow. For further information on access to this data, contact: ....

B ZTRAX summary statistics

Table 4: Observations by state

State	Total	Non-zero	State	Total	Non-zero
AK	87,481	0.6	AL	764,753	76.7
AR	525,394	74.7	AZ	3,870,789	77.0
CA	12,695,929	81.2	CO	2,465,752	83.8
CT	977,941	95.3	DC	153,775	90.8
DE	186,550	84.2	FL	1,0549,532	83.7
GA	2,946,447	75.9	HI	318,845	90.1
IA	343,278	83.0	ID	282,022	0.9
IL	2,756,664	78.8	IN	1,315,892	5.4
KS	436,669	5.6	KY	542,103	89.5
LA	73,519	86.4	MA	2,312,393	94.7
MD	3,332,529	82.4	ME	155,485	3.2
MI	1,990,535	72.3	MN	1,373,273	94.7
MO	1,264,435	31.1	MS	203,776	20.9
MT	89,388	4.6	NC	1,019,788	81.3
ND	59,523	47.1	NE	248,312	85.4
NH	359,665	92.0	NJ	802,287	90.9
NM	27,992	2.0	NV	1,722,843	83.5
NY	2,766,500	88.3	OH	3,644,017	78.6
OK	981,942	75.3	OR	1,166,150	88.9
PA	2,685,987	93.4	SC	1,207,735	83.8
SD	19,872	66.4	TX	7,651,403	1.6
UT	1,188,371	0.9	VA	2,088,157	93.1
WA	2,647,489	82.8	WI	443,308	87.7
WV	42,623	67.1	WY	36,137	5.7

Figure 8: House prices by state

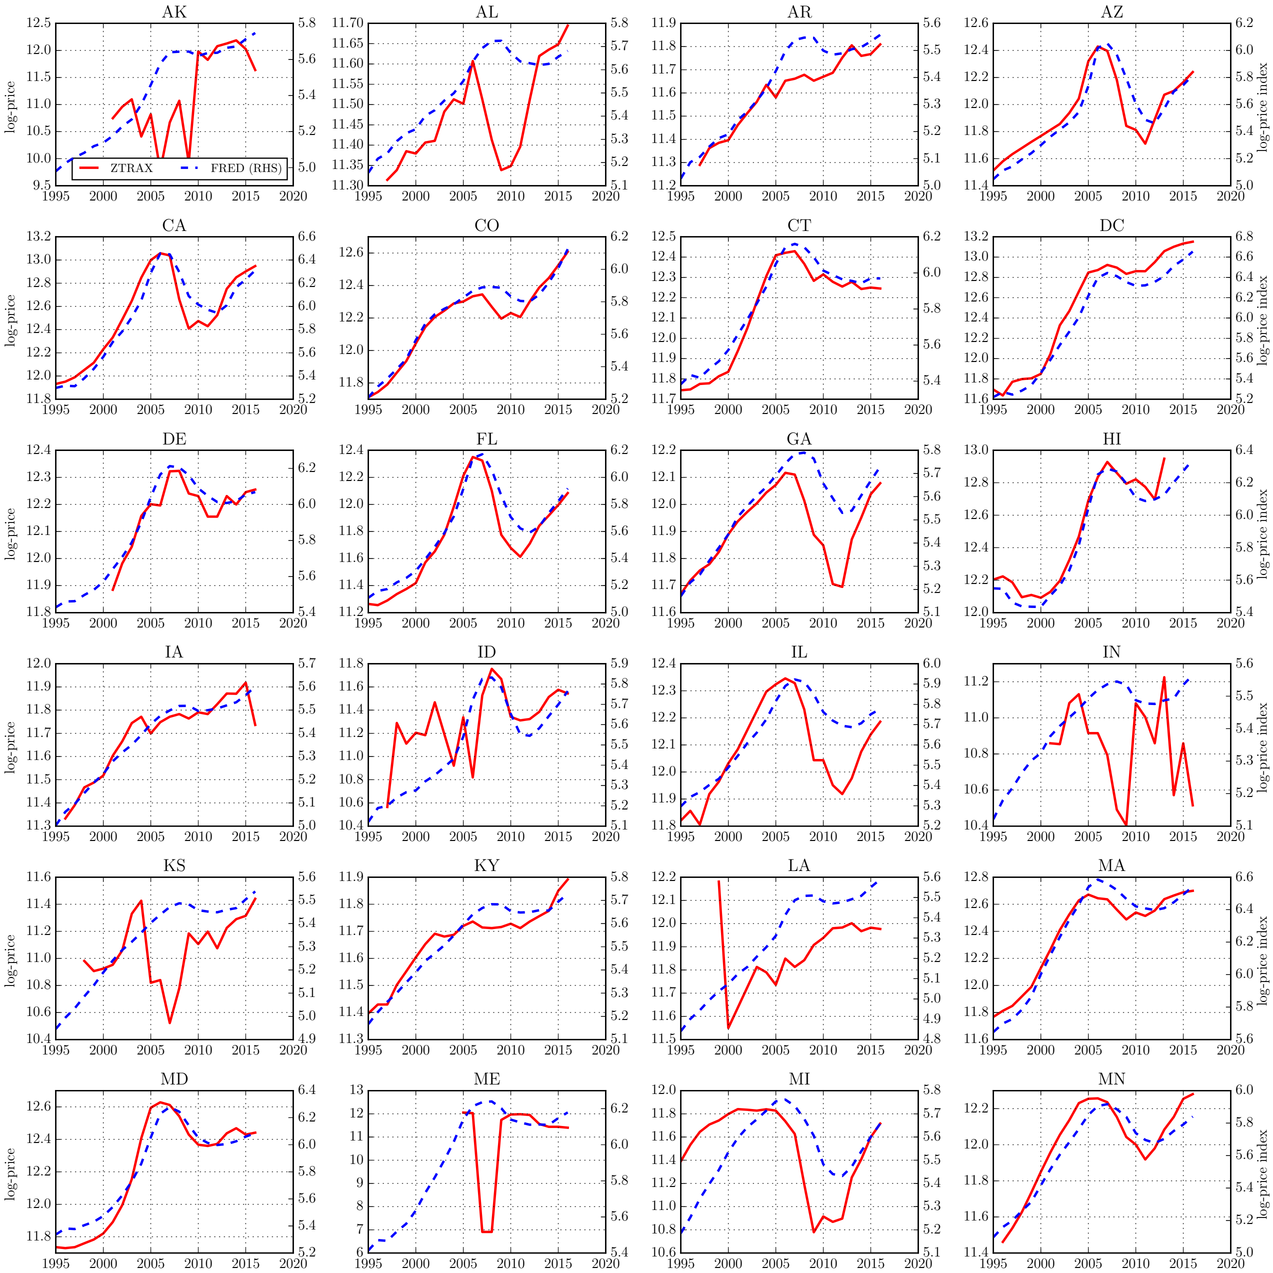

Source: Zillow, FRED.

Figure 9: House prices by state

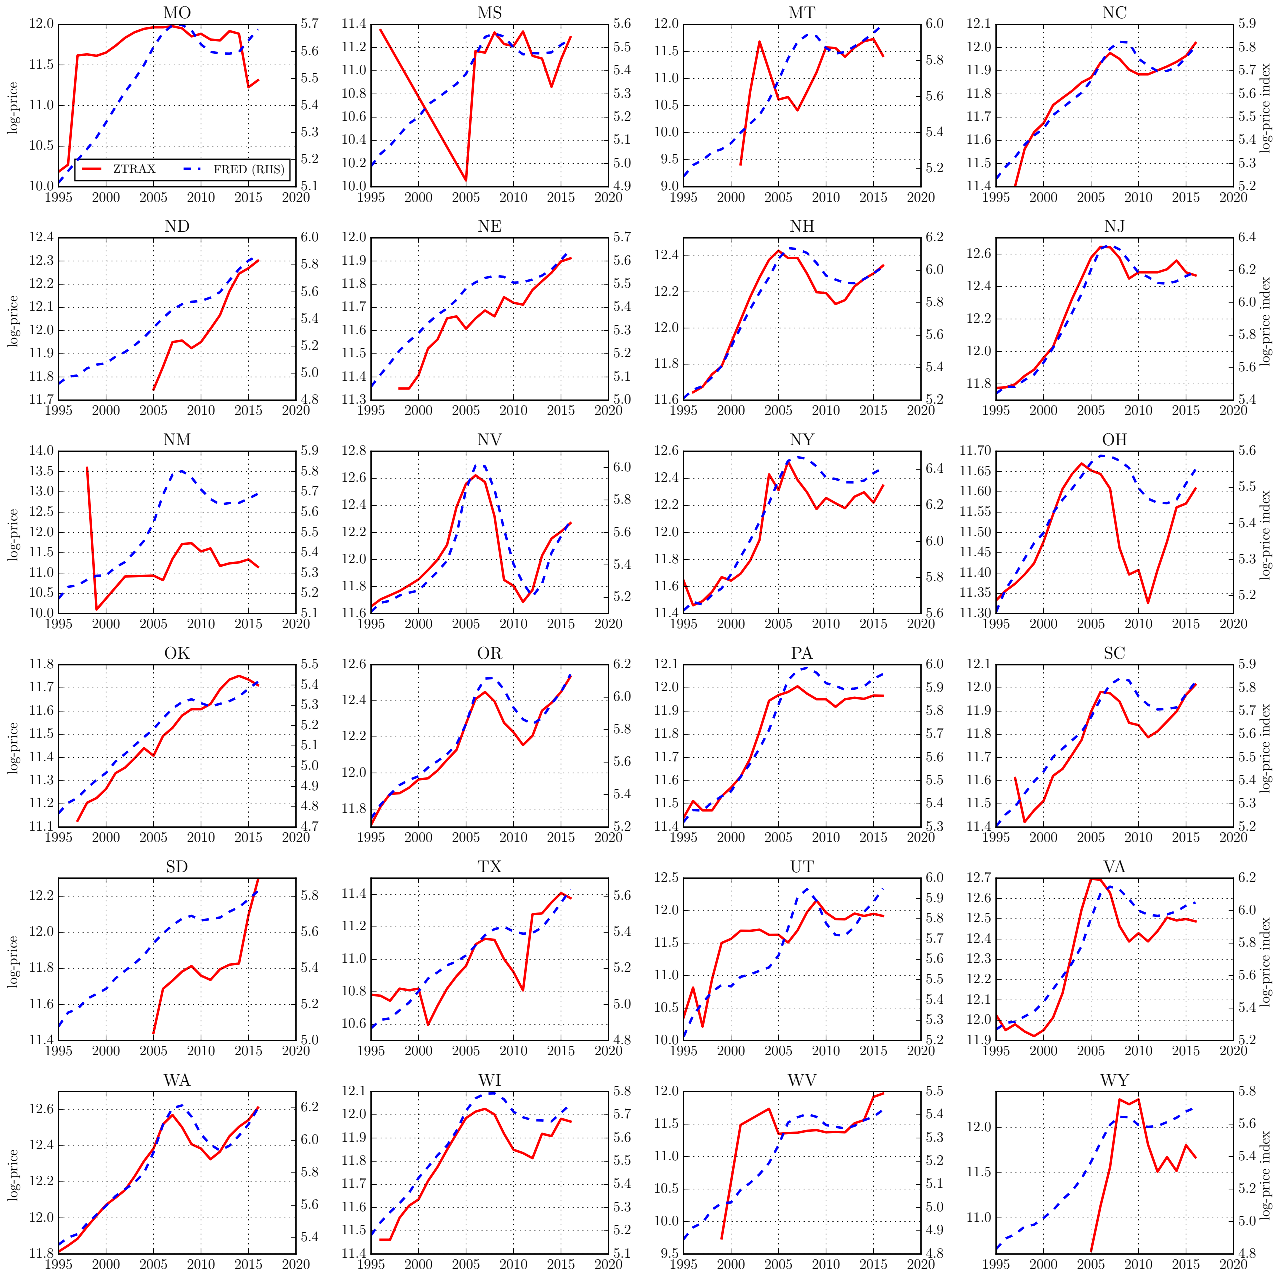

C Kilts Consumer Panel Data summary statistics

Year	Total Panelists	Remaining Panelists	Year	Total Panelists	Remaining Panelists
2004	39577	37331	2010	60658	56974
2005	38863	36605	2011	62092	58450
2006	37786	35627	2012	60538	57045
2007	63350	59302	2013	61097	57565
2008	61440	57539	2014	61557	57996
2009	60506	56734	2015	61380	57880

Table 5: Consumer Panel Data observations

D Estimating the house price process

In this section I describe the process for estimating the statistical house price processes for the CBSA, zip-code, and idiosyncratic price components.

D.1 Computing the autocovariance matrices

In order to estimate the autocovariance matrices for MSA, zipcode, and idiosyncratic house price movements, I modify the approach of Blundell et al. (2008), who estimate the autocovariance structure for income and consumption. Let j denote a unit of analysis at either MSA, zipcode, or individual house level. Let $\mathbf{p}_j = [p_{j,1}, \cdots, p_{j,T}]'$ denote a vector containing the observations of that unit across time. Note that in many cases, there will be missing data, in particular for idiosyncratic house prices since individual houses infrequently sell. Let $\mathbf{d}_j = [\sqrt{N}_{j,1}, \cdots, \sqrt{N}_{j,T}]'$ denote a vector of the square root of the number of observations for unit j in a given period. For the idiosyncratic component, $\sqrt{N}_{j,t}$ is one when the house is sold, and zero when it is not sold. The $T \times T$ autocovariance matrix M is then given by:

$$M = \left[\sum_{j=1}^{N_J} (\mathbf{p}_j \circledast \mathbf{d}_j)(\mathbf{p}_j \circledast \mathbf{d}_j)' \right] \oslash (W_1 - W_2 \oslash W_1)$$

where $\circledast$ is element-wise multiplication, $\oslash$ is element-wise division, $W_1 = \sum_{j=1}^{N_J} \mathbf{d}_j \mathbf{d}_j'$, and $W_2 = \sum_{j=1}^{N_J} (\mathbf{d}_j \mathbf{d}_j') \circledast (\mathbf{d}_j \mathbf{d}_j')$.[12] Note that we are here computing the weighted sample covariance matrix. This is required since different numbers of houses are sold in each MSA or zipcode across time. For the idiosyncratic component, there is only one observation per house sold and the covariance matrix reduces to that found in **BLUNDELL**:

$$M = \left(\sum_{j=1}^{N_J} \mathbf{p}_j \mathbf{p}_j' \right) \oslash \left(\sum_{j=1}^{N_J} \mathbf{d}_j \mathbf{d}_j' \right)$$

[12]This can be written in matrix notation as:

$$M = \left[(\mathbf{p} \circledast \mathbf{d})(\mathbf{p} \circledast \mathbf{d})' \right] \circledast \left[(\mathbf{dd}') \oslash ((\mathbf{dd}') \circledast (\mathbf{dd}') - (\mathbf{d} \circledast \mathbf{d})(\mathbf{d} \circledast \mathbf{d})') \right]$$

Let the vector of moments relevant for estimation be $\mathbf{m} = vech\{M\}$. Then the variance-covariance matrix of $\mathbf{m}$ relevant for inference is:

$$\mathbf{V} = \left[\sum_{j=1}^{N_J} ((\mathbf{m}_j - \mathbf{m})(\mathbf{m}_j - \mathbf{m})') \circledast (\mathbf{D}_j \mathbf{D}_j') \right] \oslash (\mathbf{D}\mathbf{D}')$$

where $\mathbf{m}_i = vech\{(\mathbf{d}_j \circledast \mathbf{p}_j)(\mathbf{d}_j \circledast \mathbf{p}_j)'\}$, $\mathbf{D}_i = vech\left\{\mathbf{d}_j \mathbf{d}_j'\right\}$, and $\mathbf{D} = vech\{W_1 - W_2 \oslash W_1\}$.

Not sure about this last covariance matrix at all. Need to check...

D.2 Minimum distance estimation

Let θ be the vector of parameters to be estimated. Then $f(\theta)$ is the half-vectorization of the model-implied autocovariance matrix. We can then form the minimum distance objective function as:

$$\min_{\theta} (m - f(\theta))' A (m - f(\theta))$$

where A is a diagonal weighting matrix, which I take to be $diag(V^{-1})$, as in **BLUNDELL**. The results are not sensitive to using an identity weighted matrix. Let the Jacobian be $G = \partial f(\theta)/\partial\theta|_{\theta=\hat{\theta}}$, then we can compute standard errors for the parameter estimates from the sandwich-form of the variance-covariance matrix of $\hat{\theta}$:

$$\widehat{\text{var}(\theta)}|_{\theta=\hat{\theta}} = (G'AG)^{-1}G'AVAG(G'AG)^{-1}$$

Note that the Jacobian requires us to compute the derivatives of the model moment vector $f(\theta)$. **SHOW WHAT THESE ARE...**

References

Aditya Aladangady. Housing wealth and consumption: Evidence from geographically-linked microdata. *American Economic Review*, 107(11):3415–46, 2017.

Laszlo Balazsi, Badi Baltagi, and Laszlo Matyas. Modelling multi-dimensional panel data: a random effects approach. Technical report, 2016.

Badi H Baltagi, Seuck Heun Song, and Byoung Cheol Jung. The unbalanced nested error component regression model. *Journal of Econometrics*, 101(2):357–381, 2001.

Martin Beraja, Andreas Fuster, Erik Hurst, and Joseph Vavra. Regional heterogeneity and monetary policy. Technical report, National Bureau of Economic Research, 2017.

David Berger, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra. House prices and consumer spending. Technical report, National Bureau of Economic Research, 2015.

Richard Blundell, Luigi Pistaferri, and Ian Preston. Consumption inequality and partial insurance. *The American Economic Review*, 98(5):1887–1921, 2008.

Willem H Buiter. Housing wealth isn't wealth. Technical report, National Bureau of Economic Research, 2008.

John Y Campbell and Joao F Cocco. How do house prices affect consumption? evidence from micro data. *Journal of monetary Economics*, 54(3):591–621, 2007.

Hui Chen, Michael Michaux, and Nikolai Roussanov. Houses as atms? mortgage refinancing and macroeconomic uncertainty. Technical report, National Bureau of Economic Research, 2013.

Jack Favilukis, Sydney C Ludvigson, and Stijn Van Nieuwerburgh. The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium. *Journal of Political Economy*, 125(1):000–000, 2017.

Fernando Ferreira and Joseph Gyourko. Anatomy of the beginning of the housing boom: Us neighborhoods and metropolitan areas, 1993-2009. Technical report, National Bureau of Economic Research, 2011.

Kilts Center for Marketing. Consumer panel dataset manual, 2016.

Marco Giacoletti. Idiosyncratic risk in housing markets. Technical report, Working Paper, Stanford Graduate School of Business, 2016.

Denis Gorea and Virgiliu Midrigan. Liquidity constraints in the us housing market. *Working Paper*, 2017.

Daniel L Greenwald. The mortgage credit channel of macroeconomic transmission. *Working Paper*, 2016.

Jonathan Heathcote, Kjetil Storesletten, Giovanni L Violante, et al. Quantitative macroeconomics with heterogeneous households. *Annual Review of Economics*, 1(1):319–354, 2009.

Aaron Hedlund. Illiquidity and its discontents: Trading delays and foreclosures in the housing market. *Journal of Monetary Economics*, 83:1–13, 2016.

Greg Kaplan, Guido Menzio, Leena Rudanko, and Nicholas Trachter. Relative price dispersion: Evidence and theory. Technical report, National Bureau of Economic Research, 2016a.

Greg Kaplan, Kurt Mitman, and Giovanni L Violante. Non-durable consumption and housing net worth in the great recession: Evidence from easily accessible data. Technical report, National Bureau of Economic Research, 2016b.

Greg Kaplan, Kurt Mitman, and Gianluca Violante. Consumption and house prices in the great recession: Model meets evidence. *Manuscript, New York University*, 2017.

Tim Landvoigt, Monika Piazzesi, and Martin Schneider. The housing market (s) of san diego. *The American Economic Review*, 105(4):1371–1407, 2015.

Atif Mian, Kamalesh Rao, and Amir Sufi. Household balance sheets, consumption, and the economic slump. *Quarterly Journal of Economics*, 128(4), 2013.

Albert Saiz. The geographic determinants of housing supply. *The Quarterly Journal of Economics*, 125(3):1253–1296, 2010.

Johannes Stroebel and Joseph Vavra. House prices, local demand, and retail prices. Technical report, National Bureau of Economic Research, 2014.

Eric R Young. Solving the incomplete markets model with aggregate uncertainty using the krusell–smith algorithm and non-stochastic simulations. *Journal of Economic Dynamics and Control*, 34(1):36–41, 2010.