Non-durable Consumption and Housing Net Worth in the Great Recession: Evidence from Easily Accessible Data*

Greg Kaplan Kurt Mitman Giovanni L. Violante

June 29, 2017

Abstract

In an influential paper, Mian, Rao, and Sufi (2013) exploit geographic variation in housing supply elasticities to measure the effect of the fall in housing net worth on household expenditures during the Great Recession. Their widely-cited estimates are based on proprietary house price and expenditure data. We revisit their study using alternative more easily accessible data on house prices and spending on a subset of *non-durable* goods. We re-affirm their findings in our data, and refine their analysis in several dimensions: (i) we provide a micro-foundation for their empirical specification as measuring a wealth effect in response to a change in house prices; (ii) we suggest a test for this interpretation by distinguishing the roles of lower house prices and initial leverage; (iii) we separate the impact of the shock on quantity consumed from that on prices; (iv) using the Consumer Expenditure Survey, we infer the implied elasticity of total non-durable expenditures in goods and services to housing net worth; and (v) we show how the elasticity is affected by controlling for contemporaneous changes in labor market conditions.

JEL Codes: E21, E32, R21.

Keywords: Non-durable Expenditures, House Prices, Consumption, Great Recession.

*Kaplan: University of Chicago and NBER, e-mail gkaplan@uchicago.edu; Mitman: IIES and CEPR, e-mail kurt.mitman@iies.su.se; Violante: Princeton University, CEPR, IFS, IZA and NBER, e-mail glv2@princeton.edu. We thank Emi Nakamura and Amir Sufi for detailed comments on an earlier version of the paper.

1 Introduction

One of the most distinctive features of the Great Recession was that the drop in household consumption expenditures was sharper, broader and more persistent than in other recent downturns. Virtually all components of consumption expenditures, not just durables, dropped substantially (Petev, Pistaferri, and Eksten, 2012). The leading explanation for these aggregate dynamics is the simultaneous extraordinary decline in housing net worth: aggregate real house price indexes fell by around 30 percent over this period, and have only partially recovered towards trend since.

To what extent was this drop in housing wealth responsible for the decline in the consumption expenditures of US households during the Great Recession? A reliable answer to this question is helpful in shaping the way that economists think about key issues such as consumption insurance, the sources of aggregate fluctuations, and the role of policy in mitigating the costs of business cycles.

In a widely influential paper, Mian et al. (2013) — hereafter referred to as MRS — exploit geographic variation in house price declines over the period 2006-2009 and household balance sheets in 2006, to estimate the elasticity of consumption expenditures to changes in the housing share of household net worth. Their estimates are derived from both OLS regressions and IV regressions that use the local housing supply elasticity index constructed by Saiz (2010) as a source of exogenous variation in the exposure of different geographical areas to a common aggregate housing shock.[1]

House price data in the MRS empirical analysis are obtained from CoreLogic. Their key source of expenditure data is the R.L. Polk dataset on new vehicle registrations, hence durable goods. They also report estimates for MasterCard data on credit-card spending based on a 5% random sample. From this latter source of data, they are able to quantify the effect of the shock on a range of nondurable goods and services.

All of the data used by MRS come from proprietary sources. This feature has restricted the possibility for other researchers to replicate their findings and verify the robustness of their estimates. Given the enormous influence that the MRS estimates have had on academia and policy, the main goal in this paper is to revisit the MRS

[1]There is disagreement in the literature about the validity of the Saiz (2010) housing supply elasticities as an instrument for regional differences in house price changes in the context of these regressions (Davidoff, 2013). We are agnostic about this debate since the role of the IV specifications in this note is limited to a replication and robustness analysis of an influential piece of academic research. Readers who are skeptical of the validity of the instrument can safely focus on the OLS estimates —none of our measurements necessarily require a causal interpretation in order to be of economic interest.

findings using alternative, publicly available, or more easily accessible data on housing net worth and expenditures.[2]

To construct our measure of local housing net worth, we use house price data from Zillow, which is freely available online. For expenditures, we use store-level sales of a subset of nondurable goods from the Kilts-Nielsen Retail Scanner Dataset (KNRS), a panel dataset on quantities sold and sale prices at the UPC (barcode) level for around 40,000 geographically dispersed stores in the US. Subscriptions to KNRS data are now held by nearly 100 academic institutions and are available for academic research for a non-prohibitive fee.

In spite of these differences, our findings are very reassuring. When we replicate MRS using our own data sources, we obtain an OLS estimate of 0.24 and an IV estimate of 0.36 for the elasticity of non-durable expenditures to housing net worth shocks. Based on MasterCard data on non-durables alone, MRS report OLS estimates of 0.34-0.38. Using the KNRS expenditure data together with a measure of the change in the housing share of net worth provided by MRS, we obtain an OLS estimate of 0.34 and an IV estimate of 0.37 – essentially the same elasticities that MRS find. Our lower baseline estimate can therefore be attributed to our use of Zillow house price data, which show a somewhat different cross-regional pattern of house price growth than the CoreLogic house price data. Overall, we find it encouraging that two very different measures of household spending yield such similar elasticity estimates.

After reaffirming the MRS findings with alternative, more easily accessible data on both the main dependent and independent variables, we offer five additional contributions to the academic and policy debate.

Our first contribution is theoretical. Following Berger, Guerrieri, Lorenzoni, and Vavra (2015), we are able to derive a structural interpretation to the reduced form regression specification proposed by Mian et al. (2013) and used in our empirical analysis as well. This microfoundation also provides a back-of-the-envelope calculation to verify, ex-post, the plausibility of the estimated value of the elasticity of consumption to housing net worth shocks. This simple calculation agrees with the micro estimates.

Second, we distinguish the overall elasticity of non-durable expenditures with respect to changes in house prices from the elasticity with respect to the housing net worth shock that arose from those price changes. This alternative regression specifi-

[2]There is a growing consensus among economists that transparency of empirical research should be an important goal for the profession, particularly when it concerns high profile, influential and surprising findings. See, for example Cochrane (2015) and Cochrane (2016), as well as Card, Chetty, Feldstein, and Saez (2010).

cation is useful because (a) it can be more easily compared with previous attempts to estimate the relationship between house prices and consumption; (b) it allows us to test, rather than to impose, that the household balance sheet channel was the most important transmission mechanism during this episode; and (c) it is valuable as a summary statistic that can be used to discipline quantitative structural investigations of the joint dynamics of consumption and housing during the Great Recession.[3] We find that the interaction between the fall in local house prices and the size of initial leverage has no (or at best, a weak) statistically significant effect on nondurable expenditures, once the direct effect of the fall in local house prices has been controlled for. A number of caveats apply to this result that we discuss in the paper.

Third, we separate the price and quantity components of the fall in nominal consumption expenditures. In the KNRS data, it is possible to observe the quantities of goods sold at the UPC level, together with the corresponding average price for every UPC for every store in every week. This unique feature of these data allows us to distinguish the effect of changes in house prices on the relative price of non-durable goods, from the effect on the quantities of goods purchased and consumed. The first component, as noted by Stroebel and Vavra (2014), can be interpreted as the outcome of demand shocks on local mark-ups. The second component measures the impact of changes in household wealth on the real demand for non-durable goods, including the substitution and income effects that result from the equilibrium change in prices. When we control for these changes in prices, we find an elasticity that is 20% smaller than our baseline estimates for nominal expenditures.

Fourth, since the KNRS bundle covers only a subset of nondurables and services, a plausible concern is that it is not representative of the reaction of total non-durables to house prices. For example, the data used by Mian et al. (2013) includes all in-store purchases using either debit or credit cards that are part of the MasterCard network and, as such, it features a better coverage. We deal with this concern by using the Diary Survey of the Consumer Expenditure Survey (CE) to estimate the elasticity of total nondurable goods and services to the CE counterpart of expenditures in the KNRS bundle.[4] We obtain an elasticity between 0.7 and 0.9, suggesting that the elasticities we estimated for KNRS expenditures should be reduced by roughly 20% when applied

[3] A number of recent papers study this topic: see, for example, Berger et al. (2015), Corbae and Quintin (2015) Favilukis, Ludvigson, and Van Nieuwerburgh (2015), Gorea and Midrigan (2015), Hedlund (2014), Huo and Rıos-Rull (2015), Jeske, Krueger, and Mitman (2013), Landvoigt (2015), Kaplan, Mitman, and Violante (2015).

[4] Attanasio, Battistin, and Ichimura (2005) argue that the Diary Survey is the component of the CE where household non-durable expenditures are best measured

Department	All Stores (2006)	Continuing Stores (2006-09)	Baseline OLS Sample
Dry grocery	37%	37%	37%
Frozen foods	8%	8%	8%
Dairy	8%	8%	8%
Deli	2%	2%	2%
Packaged meat	3%	3%	3%
Fresh produce	3%	3%	2%
Non-food grocery	13%	13%	13%
Alcohol	5%	5%	5%
Health and beauty aids	14%	14%	14%
General merchandise	8%	8%	9%
Number of stores	31,093	29,681	14,756

Table 1: Distribution of types of goods sold at stores in the KNRS sample

Notes: The baseline OLS sample restricts attention to continuing stores located in CBSAs for which we have house price data available.

to total non-durable goods and services.

Finally, we verify the robustness of our estimates by also controlling for contemporaneous changes in labor market conditions in the regressions. The coefficients on the labor market variables have the expected sign and increase the R squared of the regressions, but in no case are the estimates of the elasticity of consumption to the housing net worth shock statistically different from those without these additional controls.

2 Data sources

2.1 Expenditure data

We use data on store-level sales from the Kilts-Nielsen Retail Scanner Dataset (KNRS) as our measure of non-durable expenditures. The KNRS is a weekly panel dataset of total sales at the UPC (barcode) level for around 40,000 geographically dispersed stores in the United States. The survey records both quantities and prices. From this weekly-UPC level data we construct an annual store-level panel of total sales. We also aggregate sales across all stores in each Core-Based Statistical Area (CBSA) to obtain a measure of CBSA-level expenditures.[5]

Table 1 shows the breakdown of goods sold at stores in the KNRS sample by department code in 2006.[6] The KNRS bundle is mostly composed of non-durables

[5]See `https://en.wikipedia.org/wiki/Core-based_statistical_area` for a definition of a CBSA and its relationship to Metropolitan Statistical Areas (MSA) and Combined Statistical Areas (CSA).

[6]Department code is the first level in the product hierarchy, with UPC being the most detailed level of disaggregation.

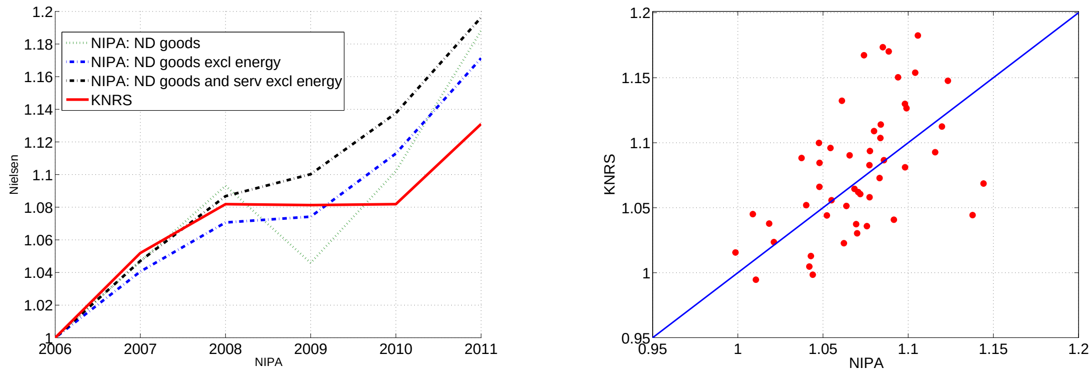

Figure 1: Left panel: Total annual sales for continuing stores in KNRS data vs various definitions of non-durable expenditures in NIPA, all normalized to 1 in 2006. Right panel: 2006-09 state-level sales growth for continuing stores in KNRS vs non-durable expenditures in NIPA

and is overweighted in food, but also contains non-food grocery (e.g., detergents and laundry supplies), health and beauty aids (e.g., cosmetics and drugs), and a residual category called general merchandise that includes some small household durables (e.g., cookware, electronics, office supplies).[7]

The type of goods covered by the data is unchanged when we restrict attention to stores present in both 2006 and 2009, and to stores located in CBSAs for which we have data on housing net worth (i.e. the sample used for our baseline estimates). According to the KNRS data manual (Kilts Center for Marketing, 2014), in 2011 the expenditures reflected in the raw data cover 53% of total sales in food, 55% of drugs, 32% of mass merchandise and 1% of liquor.

Retail sales in KNRS are a good proxy for non-durable spending in terms of aggregate time-series variation and geographic cross-sectional variation.[8] The left panel of Figure 1 shows a time-series plot of annual expenditures in the KNRS sample for the subset of stores that are always present in the data, together with various categories of consumption expenditures from the National Income and Product Accounts (NIPA, Table 2.3.5). Between 2006 and 2009, nominal growth in KNRS sales lies between growth in Personal Consumption Expenditures (PCE) in non-durables goods and growth in PCE in non-durables goods and services (excluding gasoline and other energy goods, whose price plummeted in the recession). Growth during 2010-11 is also

[7]Nielsen also collects information on goods that do not have UPCs (known as Magnet data). These goods are excluded from our analysis.

[8]In Section 6, we also show that the KNRS bundle accounts for roughly 40 percent of expenditures in non-durable goods and services.

aligned well with these measures. The only significant discrepancy occurs in 2009-10, when KNRS expenditures are flat whereas NIPA expenditures rise.[9]

The right panel of Figure 1 shows a scatter plot of the state-level 2006-09 change in expenditures in the KNRS data versus the NIPA data (state-level is the finest level of geographic aggregation for expenditure data that is published by NIPA). The correlation of these growth rates is 0.54.

Since we conduct our analysis using store-level changes in sales, we effectively control for changes in the composition of stores in a given region across years. None of our findings are affected by restricting attention to stores that are present for all intermediate years. There is still the concern that variation in entry and exit of stores —especially exit, over this recessionary period — differentially affects sales growth of continuing stores across areas and generates an attenuation bias in our estimates. For example, areas with the largest drop in house prices may be those with the sharpest rise in store exit which, in turn, mitigates the drop in sales in continuing stores as households shift their shopping towards surviving stores. To verify whether this is a serious concern, we have also repeated our analysis by aggregating store sales at the broader CBSA-level. None of our main results are affected.[10]

2.2 Housing and financial net worth data

The second important variable in our analysis is household net worth, which we construct for the years 2006-2010. We define household net worth in region i at date t as

$$NW_t^i = H_t^i + F_t^i - M_t^i - D_t^i$$

where H_t^i is housing wealth, F_t^i is financial assets, M_t^i is mortgage debt, and D_t^i is non-mortgage household debt.

We now describe the construction of these variables one by one. Each region i is a county, which we later aggregate into CBSAs.

Financial assets: We follow the corresponding calculation in MRS step by step. From the quarterly IRS Statistics of Income (SOI) data, we obtain the fraction of non-wage income (Adjusted Gross Income - wages and salaries) coming from interest

[9]Our analysis focuses almost exclusively on the period 2006-09, during which the trends in KNRS expenditures and NIPA expenditures are closely aligned.

[10]That our results are not affected by aggregating to the CBSA-level also mitigates any potential concerns about store-switching among continuing stores. Such store-switching would at worst lead to measurement error in the dependent variable which would affect the precision of our estimates but would not introduce additional bias.

and dividends for each county. Next we allocate total financial assets from the Flow of Funds (FoF) Balance Sheet of Households to each county/quarter based on the fraction of interest and dividends in each county/quarter. The implicit assumption is that the representative household in each county/quarter holds the market index for stocks and bonds.

Housing wealth: We compute the total number of houses by county from the American Community Survey (ACS) and generate housing wealth by multiplying them by the Zillow Home Value Index for All Homes. We verify that total housing wealth lines up well with its FoF counterpart for this period.

The Zillow data are publicly available from `http://www.zillow.com/research/data`. In constructing housing wealth, MRS use the CoreLogic house price price index, which is based on proprietary data. This is the most important discrepancy between our data sources and those in MRS for the construction of household net worth. The main difference between CoreLogic and Zillow is that the former is a repeat-sale index, whereas the latter is a hedonic price index that also includes sales of new homes. There are pros and cons to both approaches, as discussed in Fleming and Humphries (2013). The left panel of Figure 2 shows a strong time-series correlation between the two aggregate house price indexes, although the CoreLogic data show both a larger boom and larger bust than the Zillow data. The right panel of Figure 2 shows annual house price growth for selected US states according to the two price series for the year ending 2013. For some areas, there are sizable differences in price growth between the two series.

Liabilities: Our main source of data on household debt by county is the quarterly FRB-NY Consumer Credit Panel (CCP). The underlying source of these latter data is Equifax, which is the data source used by MRS, so this portion of the data construction is also very comparable. The CCP has information on levels of mortgage debt and non-mortgage debt (auto loans and revolving consumer credit) in each county. Since the CCP does not have data on student loans, we do an imputation for each county based on the aggregate fraction of total household debt represented by student loans from the FoF. We then define other debt (D_t^i) as the sum of auto and student loans and revolving consumer credit. Finally, we rescale debt in each county proportionately, so that the total in the CCP in each quarter equals the FoF total household liabilities.

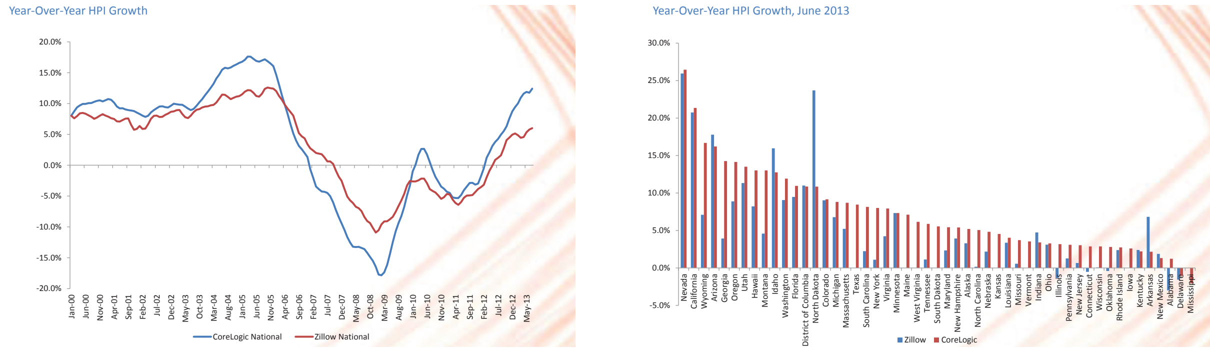

Figure 2: Left panel: CoreLogic vs Zillow house price growth, aggregate time series. Right panel: CoreLogic vs Zillow house price growth across states, June 2013. Source: Fleming and Humphries (2013)

3 Expenditure elasticities to housing net worth shocks

We follow closely the regression specification in MRS. In particular, we define the housing share of net worth as the ratio between housing wealth and household net worth, H_t^i/NW_t^i, and the log-change in this variable between date t and $t+\tau$ induced by changes in house prices – referred to as the 'housing net worth shock' – as $\Delta HNW_{t,t+\tau}^i = \Delta \log p_{t,t+\tau}^i \times (H_t^i/NW_t^i)$.

In our baseline model, we regress the three-year changes in store-level annual sales from 2006 to 2009 on the CBSA-level housing net worth shock over the same time-period. We focus on 2006-09 since this corresponds roughly to the period of the sharpest house price declines (Figure 2), and is the three-year period studied in MRS. However, since Figure 2 also shows that house prices were still rising in early 2006 and still falling in 2010, we also present results for other periods that exclude 2006 and include 2011.

Our OLS regression specification is

$$\Delta \log C_{06-09}^{s,i} = \beta_0 + \beta_1 \Delta \log p_{06-09}^i \left(\frac{H_{06}^i}{NW_{06}^i} \right) + \epsilon_{06-09}^{s,i}. \tag{1}$$

where the dependent variable is sales of KNRS goods in store s in CBSA i. The independent variable is the CBSA-level change in the housing share of net worth induced by changes in local house prices.[11] We weight observations by store-level sales in 2006 (alternative weighting possibilities have little impact on the results), and we cluster by CBSA when computing standard errors.

[11] In Section 5 we consider reasonable alternative ways of constructing the independent variable in this regression.

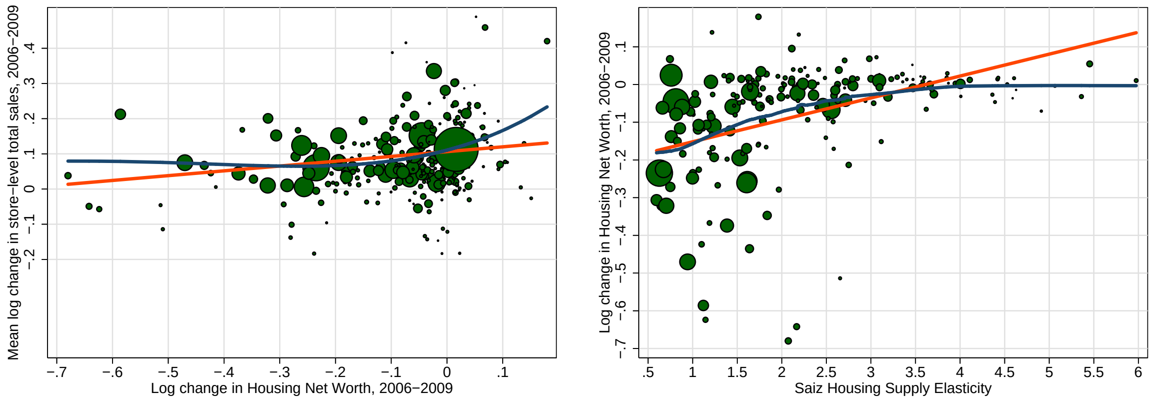

Figure 3: Left panel: Mean log change in store level sales from 2006-2009 versus log change in housing net worth. Right panel: Log change in housing net worth from 2006-2009 versus housing supply elasticity. Linear and non-linear fit lines. Size reflects CBSA-level sales in 2006.

In the left panel of Figure 3 we show a scatter plot of the change in CBSA-level sales, i.e. $\Delta \log C^i_{06-09}$ where $C^i = \sum_{s \in i} C^{s,i}$, against the CBSA-level change in the housing share of net worth, together with linear and non-linear fitted lines. The size of the circles reflects the weight of each CBSA in the regression. There is a clear positive slope that is strongest in areas which experienced the smallest declines in the housing share of net worth. Among areas with large declines (below -10 pct), there is essentially no relationship between spending and the log-change in the housing share of net worth.

Our IV regressions follow closely the specification in MRS. We use the estimates of housing supply elasticities from Saiz (2010) to instrument for the housing net worth shock. This instrument is provided at the CBSA level, and is not available for all of the CBSA's in which we observe store-level changes in expenditure because not all CBSAs are covered by the Saiz (2010) data. As a result, the OLS and IV samples differ.[12] The right panel of Figure 3 shows a scatter plot of the first stage of this regression, i.e. the change in the CBSA-level housing share of net worth against the Saiz (2010) instrument, together with linear and non-linear fitted lines. The figure is suggestive of a strong first stage, but also reveals a marked degree of non-linearity in the strength of the relationship. In particular, the relationship between changes in the housing share of net worth and the instrument is much stronger among low elasticity areas than among high elasticity areas. Given the poor fit of the linear specification, particularly for high

[12] This difference in samples has a negligible impact on the estimates.

	2006-09 CBSA			2006-09 County		2006-09 CBSA - MRS HNW	
	OLS	IV	IV (linear)	OLS	IV	OLS	IV
ΔHNW^i	0.239**	0.361**	0.405**	0.207**	0.192*	0.341**	0.286**
	(0.029)	(0.077)	(0.089)	(0.025)	(0.080)	(0.047)	(0.116)
N	14,756	12,701	12,701	21,226	16,748	22,945	19,513
Clusters	281	181	181	584	382	330	233
R^2	0.024	0.017	0.012	0.017	0.017	0.019	0.018

Table 2: Elasticity of non-durable expenditures to housing share of net worth

elasticity areas where the reduced form relationship is strongest (Figure 3, right panel), we use a quartic polynomial in the Saiz housing supply elasticity as our instrument. The use of a non-linear first stage represents another difference with MRS, who use a linear first stage, and sharpens the estimates without having a large impact on the value of the coefficients.[13]

Table 2 reports our main results. For the three-year period 2006 to 2009, we obtain a baseline elasticity estimate of 0.24 (0.03) using OLS, and of 0.36 (0.08) using IV (Table 2, first two columns). Both estimates are significant at the 1% level. The corresponding IV estimate with a linear first stage is 0.41 (0.09) (Table 2, third column).

In their Table III, MRS report an estimated elasticity of 0.63 (0.12) using OLS, and of 0.77 (0.24) using IV for their measure of durable spending (vehicle registration). These are larger numbers compared to ours, which is intuitive as durables are much more cyclical. In Table 2 of their Online Appendix not for publication, the authors report elasticity estimates using data on non-durable and services expenditure from MasterCard, a proprietary dataset of purchases using either debit or credit cards that are part of the MasterCard network. This proxy for expenditures is closer to ours and to a representative bundle of non-durable goods.

Their OLS estimates for non-durables vary from 0.34 (0.11) to 0.38 (0.10), depending on the exact definition, and are therefore remarkably similar to ours.[14] One possible concern may be that MRS use county-level data rather than CBSA-level data in these regressions. When we re-estimate (1) using the county-level equivalent of our measure of the housing net worth shock, we obtain an OLS estimate of 0.21 (0.03) and an IV estimate of 0.19 (0.08) (Table 2, fourth and fifth columns). Hence the difference in

[13]For all specifications, instrumenting with a quartic polynomial of the elasticity results in uniformly marginally lower IV estimates and smaller standard errors than the corresponding estimates that restrict to a linear first stage.

[14]MRS do not report the IV counterpart of these estimates for non-durable expenditures.

the level of geographic aggregation has only a minor effect on the OLS estimates. The impact on IV estimates is somewhat larger, but the discrepancy between county- and CBSA-level estimates is not statistically significant.

Although MRS are not able to make their county-level measures of the housing share of net worth available for other researchers due to the proprietary nature of their sources, they do make an analogous CBSA-level measure available that replaces Equifax data for debt with the FRB-NY CCP data (the same source that we use), and replaces CoreLogic data on house prices with the house price index produced by the FHFA. When we use this one as our independent variable, we obtain an OLS elasticity estimate of 0.34 (0.05) and IV of 0.29 (0.12) (Table 2, sixth column). It is very reassuring that two completely different sources of data on non-durable expenditures generate essentially identical estimates for the elasticity with respect to changes in the housing share of net worth, provided that the latter variable is consistently measured.

Our estimates of the elasticity of non-durable spending to changes in the housing share of net worth are relatively insensitive to the particular choice of time period. In Table 3 we report corresponding estimates for alternative time periods around the Great Recession. Since the decline in expenditures typically lagged the fall in house prices, including 2010 and/or 2011 leads to larger estimates (since house prices had mostly leveled off by 2010 but consumption was still declining), and excluding 2006 leads to smaller estimates (since house prices were falling through most of 2006 but the largest declines in consumption were still to come).

We conclude this section by noting that most of our regressions (and those in Mian et al. (2013) as well) yield IV estimates that are larger than OLS coefficients. This may appear puzzling, given that the role of the IV is to purify the housing net worth shock of a component that drives both house price movements and expenditures, for example, such as changes in income or unemployment. The presence of a common factor would lead to an upward bias in the OLS estimates. One possible explanation is that the Saiz instrument is also correlated with such a common factor and is thus invalid. Another possible explanation is that idiosyncratic variation in house prices is more transitory than is the variation in the common component of house prices that is isolated by the instrument. More persistent house price movements might be expected to have a bigger effect on expenditures.

	2006-10		2006-11		2007-09		2007-10		2007-11	
	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV
ΔHNW^i	0.263**	0.455**	0.274**	0.462**	0.208**	0.328**	0.244**	0.458**	0.260**	0.509**
	(0.024)	(0.099)	(0.023)	(0.090)	(0.027)	(0.088)	(0.031)	(0.131)	(0.027)	(0.133)
N	14,536	12,518	14,220	12,231	16,266	13,735	16,032	13,544	15,682	13,226
Clusters	281	181	281	181	338	183	338	183	338	183
R^2	0.028	0.015	0.032	0.021	0.015	0.009	0.02	0.006	0.021	0.005

Table 3: Elasticity of non-durable expenditures to housing share of net worth in alternative time periods

4 Structural interpretation of elasticity to housing net worth

Our empirical analysis has so far focused on the regression specification in equation (1), which is the specification originally proposed by Mian et al. (2013). In this section of the paper, we build on the analytical results derived in Berger et al. (2015) to show that this specification can be micro-founded from a simple model of optimal consumption and housing choices. When viewed through the lens of this model, the elasticity, which can be expressed in terms of observable objects, measures the strength of the wealth effect, which is given by the ratio of non-human (housing plus financial) wealth to total wealth. A simple back-of-the-envelope calculation yields values for the elasticity that are very much in line with the Mian et al. (2013) and our estimates.

Consider a household that solves the following problem:

$$\max_{\{C_{it},Q_{it},A_{it}\}} \sum_{t=0}^{T} \beta^t \frac{\left(C_{it}^{\alpha} Q_{it}^{1-\alpha}\right)^{1-\sigma}}{1-\sigma}$$
$$s.t.$$
$$C_{it} + p_t \left[Q_{it} - (1-\delta) Q_{i,t-1}\right] + A_{it} = Y_{it} + (1+r)\, A_{i,t-1}$$

where C_{it} is non-housing consumption expenditures, Q_{it} is the housing stock that yields some utility flow proportional to the stock, p_t is the aggregate house price, δ is housing depreciation, A_{it} is holdings of financial wealth, and Y_{it} is income.

We assume that $\beta(1+r) = 1$, the path of income is $\{Y_{it}\}$ is deterministic, there are no borrowing constraints or transaction costs for housing, and $p_t = p$ for every t. With these assumptions, it is straightforward to show that optimal consumption is given by

$$C_{it} = \alpha (1-\beta) \left[\sum_{\tau=t}^{T} (1+r)^{t-\tau}\, Y_{i\tau} + p(1-\delta)\, Q_{i,t-1} + (1+r)\, A_{i,t-1} \right]. \qquad (2)$$

The formal derivation of equation (2) is in Appendix A.

In what follows, to simplify the notation we denote human wealth as Υ_t^*, housing quantity net of depreciation as Q_t^* and financial wealth comprising current interests as A_t^*, i.e.

$$\begin{aligned}
\Upsilon_t^* &= \sum_{\tau=t}^{T} (1+r)^{t-\tau} Y_{i\tau} \\
Q_{it}^* &= (1-\delta) Q_{it} \\
A_{it}^* &= (1+r) A_{it}.
\end{aligned}$$

Differentiating equation (2) with respect to p, to capture the effect of an unexpected permanent change in house prices p on consumption gives

$$\frac{\Delta C_{it}}{\Delta p} = \alpha (1-\beta) Q_{i,t-1}^*,$$

and rearranging yields

$$\begin{aligned}
\frac{\Delta C_{it}/C_{it}}{\Delta p/p} &= \frac{\alpha (1-\beta) p Q_{i,t-1}^*}{C_{it}} \\
&= \frac{H_{i,t-1}^*}{\Upsilon_{it}^* + H_{i,t-1}^* + A_{i,t-1}^*},
\end{aligned} \tag{3}$$

where we introduced the notation $H_{i,t-1}^* = p Q_{i,t-1}^*$ for housing wealth. Switching to log notation and using the definition of the net worth shock in into (3) above yields:

$$\begin{aligned}
\Delta \log C_{it} &= \frac{H_{i,t-1}^* + A_{i,t-1}^*}{Y_{it}^* + H_{i,t-1}^* + A_{i,t-1}^*} \times \Delta \log p \left(\frac{H_{i,t-1}^*}{H_{i,t-1}^* + A_{i,t-1}^*} \right) \\
&= \frac{H_{i,t-1}^* + A_{i,t-1}^*}{\Upsilon_{it}^* + H_{i,t-1}^* + A_{i,t-1}^*} \times \Delta HNW_{it}
\end{aligned}$$

Therefore, the regression specification (1) has a structural interpretation where the elasticity of consumption of the housing net worth shock is given by

$$\beta_1 = \frac{H_{i,t-1}^* + A_{i,t-1}^*}{\Upsilon_{it}^* + H_{i,t-1}^* + A_{i,t-1}^*} \tag{4}$$

that is, the ratio of financial and housing net worth to total net wealth including human wealth.[15]

This structural interpretation provides a way to check ex-post whether the empirical estimates of this elasticity based on geographical variation are reasonable. Consider the following simple calculation. Average household financial and non-financial net worth in 2007 (from the Survey of Consumer Finances) was $\$500,000$ and average household labor income was around $\$70,000$. To compute human wealth, a rough back of the envelope calculation for a 45 year-old household with 20 years left in the labor market paying an average tax rate of 20% and receiving an earnings replacement rate of 0.4 in social security benefits over its residual lifetime of 15 years in retirement, all discounted at 3% per year, gives approximately $\$1M$ for the term $\Upsilon^*_{it} = \sum_{\tau=t}^{T} (1+r)^{t-\tau} Y_{i\tau}$ in (4).

This calculation, that does not rely on identifying empirically exogenous sources of variation for house prices, gives a value for β_1 of around $500/(1,000+500) = 1/3$, which compares well with our estimates in Table 2.

An important final point is that these equations are derived in a deterministic model without borrowing constraints, in which the consumption function is linear (not concave). A corollary of these assumptions (together with the Cobb-Douglas preferences over non-durables and housing services) is that the elasticity of expenditures to a change in *(*house prices) in this model is due to a wealth (or balance sheet) effect – it is large when the ratio of housing wealth to total wealth is large (see Berger et al. (2015)).

The elasticity of expenditures to *housing net worth* (β_1) is not informative about the concavity of the consumption function (due to, say occasionally binding credit constraints or other sources of precautionary saving). A large estimated value for β_1 could be entirely consistent with a linear consumption function, and yet sill the elasticity of expenditures to either house prices or housing net worth depend on the composition of the household balance sheet. For the former elasticity, the strength is determined by the ratio of housing wealth to total wealth; for the latter elasticity the strength is determined by the ratio of non-human wealth to total wealth.

[15] To complete the interpretation of equation (1) through the lens of this model involves aggregating across individuals within a region. It is then necessary to assume that the average ratio of non-human wealth to total wealth differs across regions (so that β_1) is a constant, but that the average composition of non-human wealth between housing and financial wealth differs across regions (in order to generate cross-regional variation).

5 Role of initial leverage

The MRS regression specification (1), which is justified by the simple model in Section 4 includes as the the independent variable the change in the housing share of net worth induced by the change in house prices, $\Delta \log p^i_{06-09}\left(\frac{H^i_{06}}{NW^i_{06}}\right)$. MRS focus on this specification because they are interested in investigating the strength of a particular transmission mechanism for the effects of changes in house prices on expenditures, which they label the household balance sheet channel. According to this narrative, one would expect two regions experiencing the same size decline in house prices to differ in their expenditure responses depending on how big of an impact the change in house prices has on their net worth, i.e. depending on the average balance sheet composition of households living in those regions. For example, more levered regions should respond more; regions where households have more financial wealth, over and above housing wealth, should respond less. Viewed through the lens of the model in Section 4 this is equivalent to assuming that regions differ in their response to changes in house prices because of differences in the housing share of non-human wealth, i.e. differences in initial leverage.

Richer models than the one in Section 4 would include both additional channels through which house price changes can affect consumption, which may not mediated by initial leverage, as well as possible effects of initial leverage on changes in expenditure that are not mediated by changes in house price. A useful way to investigate these possibilities is to recognize that the independent variable in (1) is an interaction between local house price changes $\Delta \log p^i_{06-09}$ and the initial share of housing in net worth $\frac{H^i_{06}}{NW^i_{06}}$. We can therefore test for the balance sheet channel by estimating the following more general specification of the regression:

$$\Delta \log C^{s,i}_{06-09} = \beta_0 + \beta_1 \Delta \log p^i_{06-09}\left(\frac{H^i_{06}}{NW^i_{06}}\right) + \beta_2 \Delta \log p^i_{06-09} + \beta_3\left(\frac{H^i_{06}}{NW^i_{06}}\right) + \epsilon^{s,i}_{06-09}. \tag{5}$$

Whereas estimation of (1) effectively assumes that house price changes are mediated by initial leverage (as implied by the simple model) and estimates the average strength of this effect, estimation of (5) instead tests for the balance sheet channel by measuring the extent to which the transmission of house prices to expenditure depends on initial leverage.

In Table 4 we report estimates of alternative specifications based on equation (5). The estimate of the elasticity of non-durable expenditures with respect to house prices,

	2006-09					
	OLS	IV	OLS	OLS	OLS	OLS
$\Delta \log p^i$	0.169**	0.260**	0.136*	0.176**	0.119**	0.116**
	(0.020)	(0.057)	(0.057)	(0.067)	(0.031)	(0.029)
$\Delta \log p^i \left(\frac{H^i}{NW^i}\right)$			0.050	-0.043		
			(0.097)	(0.128)		
$\frac{H^i}{NW^i}$				-0.053		
				(0.039)		
$\Delta \log p^i \left(\frac{H^i}{H^i - M^i}\right)$					0.036	0.037*
					(0.019)	(0.017)
$\frac{H^i}{H^i - M^i}$						-0.003
						(0.003)
N	14,756	12,701	14,756	14,756	14,756	14,756
Clusters	281	181	281	281	281	281
R^2	0.026	0.019	0.026	0.026	0.026	0.026

Table 4: Elasticity of non-durable expenditures to house prices: alternative specifications to assess the household balance-sheet channel

$\hat{\beta}_2$, is 0.17 (0.02) using OLS, and 0.26 (0.06) using IV. Since the housing share of net worth at the CBSA-level is typically far below unity, the elasticities with respect to house prices in the first two columns of Table 4 are necessarily smaller than the elasticities with respect to the house price induced changes in the housing share of net worth in the first two columns of Table 2. However, the R^2 from these regressions are slightly hire than the corresponding regressions in Table 2. When we include the interaction term as an additional regressor, either with or without the initial level of the housing share of net worth, the OLS estimate of β_2 is barely affected while the OLS estimate of β_1 is small and negative, and is statistically indistinguishable from zero (columns 3 and 4). When we instead use a measure of housing leverage ($\frac{H^i}{H^i - M^i}$), the interaction term has the correct sign and is statistically significant at the 5% level (columns 5 and 6).[16]

Discussion: The results in Table 4 do not offer much support to the view that the degree of leverage in 2006 was a key determinant of the elasticity of expenditures to the fall in house prices that occurred between 2006 to 2009. Here, a number of remarks

[16]We report only OLS estimates for the specifications that include both changes in house prices and the interaction between changes in house prices and initial net worth, since estimation by IV would require additional instruments.

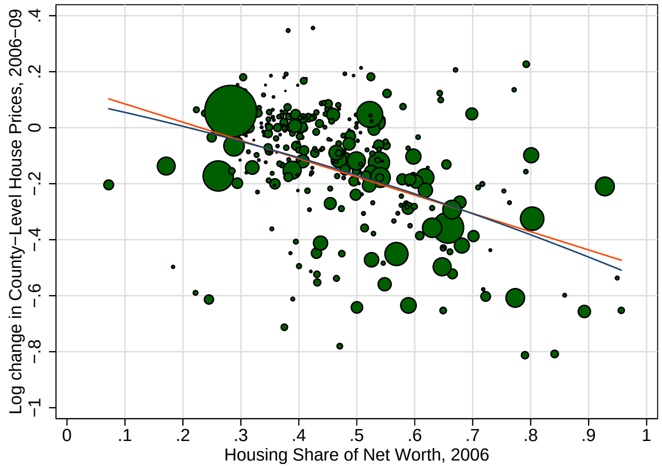

Figure 4: Scatterplot of Log change in housing net worth from 2006-2009 versus initial share of housing in household net worth across CBSAs.

are in order.

First, Table V in Mian et al. (2013) separates the effect of the change in house value from the interaction with initial housing leverage and finds that the latter is statistically significant in determining the response of auto spending.[17] Thus, our finding that the balance sheet effect may be weaker could be due to the fact that our dependent variable is *non-durable* expenditures.[18]

Second, the regressions in Mian et al. (2013) are specified in levels as opposed to logarithms as they are interested in measuring MPCs rather than elasticities. Elasticities have the advantage of being scale-invariant, whereas measuring MPCs requires correctly measuring the absolute level of spending, assets and liabilities in each geographical area. This can only be achieved by making certain proportionality assumptions about the relation between the level of each variable in the regional sample and the national aggregates from NIPA or Flow of Funds data.

Third, estimating the partial elasticitits β_1, β_2 and β_3 requires significant variation in house price growth and initial balance sheet position in the cross section of regions. One may worry that the lack of strong significance of the interaction terms in our regressions is due to the fact that there is not enough variation at the CBSA level. Indeed, in their Handbook chapter (see, in particular, Table 3), Mian and Sufi (2016)

[17]Also Baker (2017) finds that higher levels of leverage are significantly related to higher sensitivity of expenditures to income using a measure of total expenditures that includes durables.

[18]Mian et al. (2013) do not separate the direct effect of house prices from the interaction term in their regressions on MasterCard spending.

	2006-09			
	OLS	IV	OLS	IV
$\Delta \log H^i$	0.124**	0.183**		
	(0.019)	(0.038)		
$\Delta \log \left(H^i - M^i\right)$			0.072**	0.121**
			(0.011)	(0.025)
N	14,756	12,701	13,724	11,745
Clusters	281	181	229	171
R^2	0.021	0.017	0.021	0.012

Table 5: Elasticity of non-durable expenditures to gross and net housing wealth

show that these two variables are much less correlated at the ZIP-code level when including county fixed-effects and advocate that this type of empirical analysis be done at the ZIP-code level. Figure 4 contains a scatterplot of the change in house prices and initial housing share of net worth across CBSAs in our sample. The correlation between the two variables is strong, but there is a substantial amount of variation. In addition, the standard errors in Table 4 do not blow up, suggesting that the collinearity problem is not too severe in our data. In the end, whether our result that the effect of initial leverage is weaker for non-durables than for durables is a genuine finding or an artifact of excessive geographical aggregation (and hence low power) is a question that can only be properly answered with highly disaggregated data on non-durable expenditures, data that at the moment are not available.

Finally, for completeness, we conclude this section by reporting elasticities with respect to gross and net housing wealth directly, rather than to house prices or the housing share of net worth. Future quantitative studies of consumption and housing in the Great Recession may be interested in these elasticities. Which independent variable is most relevant will differ depending on the specifics of the mechanism being investigated. These alternative elasticity measures are contained in Table 5. The estimated elasticities with respect to gross housing wealth H^i are 0.12 (OLS) and 0.18 (IV) (first and second columns). These estimates are lower than the corresponding elasticities with respect to house prices in Table 3 — 0.169 (OLS) and 0.26 (IV)— because, during this period, the regions with the largest decline in house prices also had the largest fall in the quantity of housing (through lower investment).

	2006-09				2007-11			
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
ΔHNW^i	0.196**	0.290**			0.218**	0.444**		
	(0.026)	(0.085)			(0.027)	(0.135)		
$\Delta \log p^i$			0.140**	0.213**			0.128**	0.261**
			(0.018)	(0.061)			(0.015)	(0.077)
N	14,756	12,701	14,756	12,701	15,682	13,226	15,682	13,226
Clusters	281	181	281	181	338	183	338	183
R^2	0.016	0.012	0.018	0.013	0.013	0.001	0.013	0.002

Table 6: Elasticity of real non-durable consumption to housing net worth shocks

6 Housing consumption versus housing expenditures

Our results have so far focused on nominal consumption expenditure. This measure is of first order importance for understanding the transmission of house price shocks to fluctuations in aggregate economic activity. However, it provides an imperfect measure of the change in the real consumption of goods by households (e.g., Aguiar and Hurst (2005) document that changes in food expenditures are not equivalent to changes in food consumption). Since it is real consumption that matters for household welfare, understanding the effect of changes in housing wealth on the quantity of nondurable goods consumed is also of interest.

We construct a series for store-level real sales by aggregating product-level sales (at the bar-code level) using product-level prices common across stores at a fixed date. Our baseline measure uses the 2012 economy-wide average price for each product (or the last year for which we have price data for discontinued products) – and is thus close to a Paasche index – but we have tried various alternative choices and our results are barely affected.[19]

When using consumption as the dependent variable, the elasticity estimates are uniformly around 20% lower than the corresponding estimates in which expenditure is the dependent variable (Table 6). For the 2006-09 period the OLS estimate of the elasticity with respect to the housing net worth shock falls from 0.24 to 0.20, and the IV estimate falls from 0.36 to 0.29. For the elasticity with respect to house prices, the OLS estimate falls from 0.17 to 0.14, and the IV estimate falls from 0.26 to 0.21. For the 2007-11 period the estimates of the elasticity of consumption with respect to

[19]Our results are not affected by using a weighted (across stores) average price for each product, or by restricting attention to products that are present in every year of the Nielsen data.

the housing net worth shock are 0.22 (OLS) and 0.44 (IV), and with respect to house prices are 0.13 (OLS) and 0.26 (IV).

These findings suggest that a significant portion of the drop in consumption expenditures is due to equilibrium prices falling in response to the negative demand shock. This conclusion is in line with the earlier results presented in Stroebel and Vavra (2014) who argued that a decline in mark-ups is responsible for these price dynamics.

7 From Kilts-Nielsen to total non-durables

A possible concern throughout our empirical analysis is that our measure of household consumption expenditure obtained from the KNRS data may be rather narrow. One may worry that these categories could display different dynamics form total nondurable expenditures. In this section, we use the Consumption Expenditure survey (CE) to estimate the elasticity of total nondurables to a subset of expenditures that is as close as possible to the KNRS bundle. This number can be then used to rescale the various expenditure elasticities to changes in housing net worth estimated in the previous sections. Our aim is to estimate

$$\log c_{it}^{ND} = \mathbf{D}_t + \beta_0' \mathbf{X}_{it} + \beta_1 \log c_{it}^{KN} + \varepsilon_{it}, \tag{6}$$

where $\mathbf{D}_t$ are time dummies, $\mathbf{X}_{it}$ are a set of controls, and c_{it}^{ND} and c_{it}^{KN} are expenditures on non-durables and the KNRS bundle, respectively. The elasticity of interest is β_1.

Our starting point is the sample constructed from the Diary Survey (DS) of the CE by Attanasio et al. (2005).[20] The DS is a cross-section of consumer units asked to self-report their daily purchases for two consecutive one-week periods by means of product-oriented diaries. Each diary is organized by day of purchase and by broad classifications of goods and services. Compared to the more commonly used Interview Survey (IS), where households are retrospectively asked for their usual expenditure in the last quarter, the key advantage of the DS component of the CE is that expenditures on the goods we are interested in –specifically, the KNRS bundle which is the independent variable of regression – are much more accurately measured.[21] This is an important consideration for us, since the attenuation bias from measurement error

[20] We refer the reader to their paper for a description of the data. For an ever more detailed presentation, see Battistin (2003). We thank Erich Battistin for sharing the data.

[21] The insight of the Attanasio et al. (2005) paper is precisely that of using the DS measures for frequently purchased goods and the IS measures for more durable goods and services in order to more accurately measure changes in consumption inequality over time.

Dependent variable:	$\log c_{it}^{NDgoods}$	$\log c_{it}^{NDgoods\&serv}$
$\log c_{it}^{KN}$	0.905	0.679
	(0.003)	(0.004)
Other controls	Y	Y
N	37,893	37,893
R^2	0.81	0.54

Table 7: Elasticity of total non-durable expenditures to expenditures in the the KNRS bundle. Source CEX.

tends to artificially reduce the estimate of the elasticity of total nondurables to KNRS expenditures, the coefficient β_1 in (6).

The sample in Attanasio et al. (2005) covers a large set of items belonging to nondurable goods and services for survey years 1987-2001, i.e. the period preceding the boom-bust. Based on their detailed classification, we define KNRS consumption as the sum of food and non alcoholic beverages at home, alcohol, personal care, and housekeeping products. This definition is close to the aggregate of the items included in the KNRS data described in Section 2. For total non-durables we use two definitions. *ND goods* include, in addition to the KNRS goods, clothing and footware, tobacco, books, newspapers and magazines. This set of goods is close to the NIPA definition of nondurable goods, excluding energy (NIPA Table 2.4.5). The second variable we construct, *ND goods and services,* also includes food away from home, clothing services, entertainment, communication services, and transportation.[22]

In our DS sample, median monthly KNRS expenditures are \$840 and median spending in ND goods (ND goods and services) are \$1,160 (\$2,150). Thus, the fraction of our strict (broad) definition of nondurables accounted for by the KNRS bundle is 72 (39) percent. For comparison, the same calculation from the NIPA Table 2.4.5 for 2000 yields 70 (39) percent.

In the regression (6) we control for year dummies (which capture changes in the relative price of the KN bundle to total nondurables) as well as an equivalence scale, a polynomial in age, and indicator variables for family type, race, education and region. The results of these regressions are displayed in Table 7.

The estimates in Table 7 suggest that the elasticity of nondurable expenditures to

[22] With respect to the NIPA definition of total services, we therefore exclude expenditures on housing, health care, education, financial and insurance services. As pointed out by Attanasio et al. (2005) and Blundell, Pistaferri, and Preston (2008), these services have a more durable nature that closely assimilates them to investment and saving activities.

the KNRS bundle varies between 0.68 and 0.90, depending on how broad the definition of non-durable expenditures is. Thus the estimates in the preceding sections should be reduced by around 10 to 30 percent when interpreting them in terms of the effects on total non-durable expenditures of consumption.

8 Role of labor market conditions

In this section we examine the effect of controlling for contemporaneous changes in labor market conditions during the housing bust. Although there are clear endogeneity issues when including changes in labor market variables alongside changes in house prices, we nonetheless find it useful to know how the estimated elasticities are affected by introducing these controls.

Our regression specification is the same as in (1) except that we add the 2006-09 change in various measures of labor market conditions, which we label X^i:

$$\Delta \log C^{s,i}_{06-09} = \beta_0 + \beta_1 \Delta \log p^i_{06-09} \left(\frac{H^i_{06}}{NW^i_{06}} \right) + \beta_{LM} \Delta X^i_{06-09} + \epsilon^{s,i}_{06-09}. \tag{7}$$

We consider three alternative measures of labor market conditions: the log unemployment rate, the log employment rate, and log labor income per-capita. As in Table 2, we report both OLS estimates and IV estimates where we instrument for the change in house prices.[23]

The estimates are reported in Table 8. For comparison, the first two columns report the OLS and IV estimates from the baseline specification without labor market controls. The labor market variables have the expected effect — worse labor market conditions are associated with larger drops of non-durable consumption. The inclusion of any of the labor market variables leads to a lower measured elasticity of expenditures to housing net worth. This effect is largest for the unemployment rate. In no case are the estimates that control for labor market conditions statistically different from those that do not. Overall, the inclusion of changes in labor market conditions lowers slightly the elasticity estimates without changing the substantive conclusions.

[23]We do not attempt to construct an instrument for changes in labor market condition since our objective is to understand on the measured elasticity of expenditures to housing net worth is affected by controlling for labor market variables, not to measure the elasticity of expenditure with respect to labor market conditions *per se*.

	No Controls		Unemployment		Employment		Labor Income	
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
ΔHNW^i	0.239**	0.361**	0.173**	0.312*	0.226**	0.356**	0.213**	0.344**
	(0.029)	(0.077)	(0.028)	(0.126)	(0.028)	(0.078)	(0.026)	(0.078)
$\Delta \log U^i$			-0.096**	-0.047				
			(0.017)	(0.046)				
$\Delta \log E^i$					0.277**	0.224*		
					(0.084)	(0.109)		
$\Delta \log Y^i$							0.217**	0.157
							(0.056)	(0.088)
N	14,756	12,701	14,756	12,701	14,756	12,701	14,455	12,440
Clusters	281	181	281	181	281	181	279	179
R^2	0.024	0.017	0.030	0.023	0.026	0.019	0.027	0.020

Table 8: Elasticity of non-durable expenditures to housing share of net worth controlling for labor market conditions

9 Conclusions

Transparency in empirical work, and the ability to replicate and verify the robustness of widely influential results should be a pillar of applied economic research.

In this paper we have shown that, by combining public and easily accessible data, it is possible, for the first time, to reassess the findings of Mian et al. (2013) — findings that have been instrumental in guiding the academic and policy debate on the role of the collapse of housing in the Great Recession.

Our analysis largely confirms their results, therefore ruling out the possibility that the large co-movement of consumption and house prices in the Great Recession is due to peculiarities of their data sources.

We offer four additional contributions to this debate: (i) the reduced-form regression specification of Mian et al. (2013) can be microfunded by a simple extension of the permanent-income hypothesis, and thus the associated balance-sheet effect does not need to be associated with the concavity of the consumption function; (ii) after controlling for the drop in house prices, we do not find a significant independent effect of initial leverage on non-durable expenditures; (iii) real consumption drops approximately 20 percent less than nominal expenditures on the KNRS bundle, implying a sizable demand-induced fall in producer prices; (iv) when applied to total spending in nondurable and services, the estimated elasticity of expenditures on the KNRS bundle to changes in housing net worth should be reduced by roughly 20 percent; and (v) controlling for contemporaneous changes in labor market conditions leads to slightly lower elasticities but does not change any of the substantive conclusions.

A useful specification for computing an implied MPC out of housing wealth is the IV estimate of the elasticity of nominal nondurable KNRS spending with respect to housing equity, 0.121 (Table 5). To translate this value to an elasticity with respect to real total spending in nondurable goods and services, we need to reduce it by 40 pct, and thus we arrive at 0.073. Given an aggregate drop in housing equity of 50 percent (Flow of Funds), the implied drop in aggregate real spending in nondurables and services is 3.5%, or just below half of the observed total. Finally, to map this elasticity into an annual MPC, we multiply it by the ratio of aggregate expenditures to housing equity (around 0.37), which yields a value of 2.7 percent.

References

Aguiar, M. and E. Hurst (2005): "Consumption versus Expenditure," *Journal of Political Economy*, 113, 919–948.

Attanasio, O., E. Battistin, and H. Ichimura (2005): "What Really Happened to Consumption Inequality in the US?" *Hard to Measure Goods and Services: Essays in Honor of Zvi Griliches. Eds: Ernst R. Berndt and Charles R. Hulten. Chicago: University of Chicago Press.*

Baker, S. R. (2017): "Debt and the Consumption Response to Household Income Shocks," *Journal of Political Economy, Forthcoming.*

Battistin, E. (2003): "Errors in Survey Reports of Consumption Expenditures," Working Paper, Institute for Fiscal Studies.

Berger, D., V. Guerrieri, G. Lorenzoni, and J. Vavra (2015): "House Prices and Consumer Spending," Working Paper, National Bureau of Economic Research.

Blundell, R., L. Pistaferri, and I. Preston (2008): "Consumption Inequality and Partial Insurance," *The American Economic Review*, 98, 1887–1921.

Card, D., R. Chetty, M. S. Feldstein, and E. Saez (2010): "Expanding Access to Aministrative Data for Research in the United States," *American Economic Association, Ten Years and Beyond: Economists Answer NSF's Call for Long-Term Research Agendas. Eds: Charles L. Schultze and Daniel H. Newlon. American Economic Association.*

Cochrane, J. (2015): "Secret Data," [Online; posted 28-December-2015].

——— (2016): "Secret Data Encore," [Online; posted 5-January-2016].

Corbae, D. and E. Quintin (2015): ""Leverage and the Foreclosure Crisis." *Journal of Political Economy*, 123, 1–65.

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

Favilukis, J., S. C. Ludvigson, and S. Van Nieuwerburgh (2015): "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium," Working Paper, National Bureau of Economic Research.

Fleming, M. and S. Humphries (2013): "Home Price Indices: Appreciating the Differences," Working Paper, The Urban Institute.

Gorea, D. and V. Midrigan (2015): "Liquidity Constraints in the US Housing Market," Working Paper, NYU.

Hedlund, A. (2014): "The Cyclical Dynamics of Illiquid Housing, Debt, and Foreclosures," Working Paper, University of Missouri.

Huo, Z. and J.-V. Rios-Rull (2015): "Balance Sheet Recessions," Working Paper, Federal Reserve Bank of Minneapolis.

Jeske, K., D. Krueger, and K. Mitman (2013): "Housing, mortgage bailout guarantees and the macro economy," *Journal of Monetary Economics*, 60, 917–935.

Kaplan, G., K. Mitman, and G. L. Violante (2015): "Consumption and House Prices in the Great Recession: Model Meets Evidence," Working Paper, NYU.

Kilts Center for Marketing (2014): "Retail Scanner Dataset Manual," Working Paper, The Kilts Center Archive of The Nielsen Company.

Landvoigt, T. (2015): "Housing Demand during the Boom: The Role of Expectations and Credit Constraints," Working Paper, University of Texas Austin.

Mian, A., K. Rao, and A. Sufi (2013): "Household Balance Sheets, Consumption, and the Economic Slump," *The Quarterly Journal of Economics*, 128, 1687–1726.

Mian, A. and A. Sufi (2016): "Who Bears the Cost of Recessions? The Role of House Prices and Household Debt," *Handbook of Macroeconomics*, 2, 255–296.

Petev, I., L. Pistaferri, and I. S. Eksten (2012): "Consumption and the Great Recession: An Analysis of Trends, Perceptions, and Distributional Effects," *The Great Recession. Ed: David Weimer. NY: Russell Sage.*

Saiz, A. (2010): "The Geographic Determinants of Housing Supply," *The Quarterly Journal of Economics*, 125, 1253–1296.

Stroebel, J. and J. Vavra (2014): "House Prices, Local Demand, and Retail Prices," Working Paper, National Bureau of Economic Research.

A Appendix A: Derivation of equation (2)

The Lagrangian of the household problem, dropping the i subscripts, is:

$$\max_{\{C_t,Q_t,A_t\}} \sum_{t=0}^{T} \beta^t \left\{ \frac{\left(C_t^{\alpha} Q_t^{1-\alpha}\right)^{1-\sigma}}{1-\sigma} + \lambda_t \left[Y_t + (1+r)\, A_{t-1} + p(1-\delta)Q_{t-1} - C_t - pQ_t - A_t \right] \right\}$$

The FOCs are:

$$\begin{aligned}
\alpha C_t^{\alpha(1-\sigma)-1} Q_t^{(1-\alpha)(1-\sigma)} &= \lambda_t \\
\beta^t (1-\alpha)\, C_t^{\alpha(1-\sigma)} Q_t^{(1-\alpha)(1-\sigma)} &= \beta^t \lambda_t p - \beta^{t+1} \lambda_{t+1} (1-\delta)\, p \\
\beta^t \lambda_t &= \beta^{t+1} \lambda_{t+1} (1+r)
\end{aligned}$$

Rearranging, we obtain

$$\begin{aligned}
\left(\frac{1-\alpha}{\alpha}\right) \frac{C_t}{Q_t} &= p \left(1 - \frac{1-\delta}{1+r}\right) \qquad \text{(A1)} \\
\lambda_t &= \lambda_{t+1}
\end{aligned}$$

where the first condition sets the optimal shares of expenditures between the two goods and the second is the Euler equation. Both conditions use the assumption $\beta(1+r) = 1$.

Substituting these two conditions into the first FOC above yields

$$\begin{aligned}
C_t^{\alpha(1-\sigma)-1} Q_t^{(1-\alpha)(1-\sigma)} &= C_{t+1}^{\alpha(1-\sigma)-1} Q_{t+1}^{(1-\alpha)(1-\sigma)} \\
C_t^{\alpha(1-\sigma)-1} \left[\left(\frac{1-\alpha}{\alpha}\right) \frac{C_t}{p\left(1+\frac{1-\delta}{1+r}\right)} \right] &= C_{t+1}^{\alpha(1-\sigma)-1} \left[\left(\frac{1-\alpha}{\alpha}\right) \frac{C_{t+1}}{p\left(1+\frac{1-\delta}{1+r}\right)} \right] \\
C_t &= C_{t+1}
\end{aligned}$$

which is the perfect consumption smoothing result of the PIH with $\beta(1+r) = 1$. As a consequence of $(A1)$, we also have that $Q_t = Q_{t+1}$.

Note that since from $(A1)$

$$\left(\frac{1-\alpha}{\alpha}\right) \frac{C_t}{Q_t} = p \left(1 - \frac{1-\delta}{1+r}\right),$$

we have that

$$C_t + pQ_t \left(1 - \frac{1-\delta}{1+r}\right) = \frac{C_t}{\alpha}. \qquad \text{(A2)}$$

Iterating forward on the budget constraint:

$$\begin{aligned}
C_t + pQ_t &= Y_t + (1+r)\,A_{t-1} + p\,(1-\delta)\,Q_{t-1} - A_t \\
&= Y_t + (1+r)\,A_{t-1} + p\,(1-\delta)\,Q_{t-1} - \frac{1}{1+r}\left[C_{t+1} + pQ_{t+1} - Y_{t+1} - p\,(1-\delta)\,Q_t + A_{t+1}\right].
\end{aligned}$$

Using condition $(A2)$ yields

$$\frac{C_t}{\alpha} = Y_t + (1+r)\,A_{t-1} + p\,(1-\delta)\,Q_{t-1} - \frac{1}{1+r}\left[C_{t+1} + pQ_{t+1} - Y_{t+1} + A_{t+1}\right].$$

Exploiting the property that C_t is constant over the lifecycle and continuing iterating yields

$$\frac{1}{\alpha}C_t\left[1 + \frac{1}{1+r} + ...\right] = (1+r)\,A_{t-1} + p\,(1-\delta)\,Q_{t-1} + \sum_{\tau=t}^{T}(1+r)^{t-\tau}\,Y_\tau.$$

For T sufficiently large, we can approximate the geometric sum on the left-hand-side with the infinite sum and obtain

$$C_t \simeq \alpha\,(1-\beta)\left[\sum_{\tau=t}^{T}(1+r)^{t-\tau}\,Y_\tau + p\,(1-\delta)\,Q_{t-1} + (1+r)\,A_{t-1}\right],$$

where we use the fact that $(1-\beta) = \frac{r}{1+r}$. This is equation (2) in the main text.