Non-durable consumption and housing net worth in the Great Recession: Evidence from easily accessible data☆

Greg Kaplan [a,b], Kurt Mitman [c,d,e], Giovanni L. Violante [f,d,g,e,b,*]

[a] University of Chicago, United States of America
[b] NBER, United States of America
[c] IIES, Sweden
[d] CEPR, United Kingdom
[e] IZA, Germany
[f] Princeton University, United States of America
[g] IFS, United Kingdom of Great Britain and Northern Ireland

ARTICLE INFO

ABSTRACT

In an influential paper, Mian, Rao, and Sufi (2013) exploit geographic variation 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 proprietary expenditure data and therefore not easily replicable. We use alternative data on a subset of non-durable goods and on house prices, which are more easily accessible, to replicate their study. When estimating their same specification on our data, we obtain values for the elasticity of expenditures to the housing net worth shock that are virtually indistinguishable from theirs. However, our robustness analyses with respect to alternative model specifications yield more nuanced conclusions about the separate roles of house prices and initial housing exposure/leverage for the drop in expenditures. Moreover, the estimated elasticity is consistent, theoretically and quantitatively, with a simple calibrated model with wealth effects where leverage and credit constraints play no role.

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 et al., 2012). The leading explanation for these aggregate dynamics is the extraordinary decline in housing net worth that occurred during this period: aggregate real house price indexes fell by around 30%. However, other factors could have simultaneously affected expenditures over that same period, including a depressed labor market, a spike in economic uncertainty, and a tightening of consumer credit.

To what extent was the drop in housing wealth alone responsible for the decline in US household consumption expenditures during the Great Recession? A reliable answer to this question helps shape 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 in initial 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 variation in the exposure of different geographical areas to a common aggregate housing shock.

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, which reflects a durable good. They also report estimates using credit-card spending from a 5% random sample of MasterCard holders. From this latter source of data, they can quantify the effect of the shock on a range of non-durable goods and services.

All of the data used by MRS come from proprietary sources. This feature has impeded other researchers from replicating their findings and verifying the robustness of their estimates. A consensus is emerging in economics that empirical research should be as transparent and replicable as possible, especially in the context of important conclusions (Burman et al., 2010; Cochrane, 2015, 2016; Card et al., 2010). 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 findings.

The first part of our replication exercise consists of constructing alternative data on housing net worth and expenditures, which is either publicly available or more easily accessible, and re-estimating their same specifications on these new data. This exercise is as close as we can get to pure replication, in the language of Hamermesh (2007).

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 non-durable 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 worldwide and are available for academic research for a non-prohibitive fee.

Despite the data differences, our findings are incredibly reassuring. When we estimate the MRS specification using our 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. When we use 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. Using MasterCard data on only non-durables, MRS report OLS estimates of 0.34–0.38 – essentially the same elasticity that we find when using the same data on housing net worth. Our lower baseline estimate can, therefore, be attributed to our use of Zillow house price data, which shows a somewhat different cross-regional pattern of house price growth than the CoreLogic house price data. Overall, we find it very encouraging that two very different measures of household spending yield such similar elasticity estimates.

We conclude this first part by leveraging a key advantage of the KNRS expenditure data relative to transaction-level data, i.e., the ability to separate changes in price from changes in quantity. The price component, as noted by Stroebel and Vavra (2014), can be interpreted as the outcome of demand shocks on local mark-ups. The quantity component measures the impact of shifts 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. We estimate an elasticity of quantities to housing net worth that is 20% smaller than our baseline estimates using nominal expenditures.

In the second part of our exercise, we assess the sensitivity of the MRS findings to alternative empirical specifications. Namely, we propose specifications where we can distinguish the effect of changes in house prices from the effect of initial exposure to housing, i.e., the initial housing share of net worth, and initial leverage. Much of the narrative and the interpretation of the results in MRS is linked to the role of household debt and leverage. We find that the expenditure elasticity with respect to house prices is always statistically significant. In contrast, the expenditure elasticity with respect to initial housing exposure and initial leverage is at most weakly significant, after controlling for the direct effect of the fall in local house prices. There are several possible interpretations for this result that we discuss at length in the paper.

Empirical macroeconomics is grounded in theory. Replication in empirical macroeconomics should, therefore, also reassess the structural interpretation of the original specification used by the authors.

Following Berger et al. (2015) closely, we use conventional consumption theory to derive a structural interpretation of the reduced form regression specification proposed by MRS. Our theoretical analysis indicates that the MRS specification can be interpreted as a pure housing wealth effect, without any role for leverage or credit constraints. This micro-foundation also provides a back-of-the-envelope calculation for the value of the elasticity of consumption to housing net worth shocks. This alternate measurement yields values for the elasticity in agreement with the micro estimates.

The rest of the paper is organized as follows. Section 2 briefly discusses the issue of replication in empirical macroeconomics. Section 3 describes our data. Section 4 analyzes robustness of the MRS estimates with respect to alternative data and Section 5 with respect to alternative model specifications. Section 6 concludes.

2. Replication in empirical macroeconomics

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 findings that become influential in both academic and policy circles (Burman et al., 2010; Cochrane, 2015, 2016; Card et al., 2010). This view is in line with the recent recognition of the importance of replicability in the physical sciences and other social sciences (Begley and Ioannidis, 2015; Munafò et al., 2017; Nosek et al., 2015).

There are different meanings of replicability and various types of replication studies. Hamermesh (2007) draws a useful distinction between pure replication, in which the same data and methods are used to verify the findings of existing research, versus scientific replication, in which the key idea or conclusion put forth in existing research is re-examined using either different data or different methods.

With respect to scientific replication, Alm and Reed (2015, p. 141) emphasize the importance of reproducing the findings of the original study by using the same methods and data used by the original authors. This is a key component of replication because it gives readers confidence that any failure to confirm the results of the original study are not due to the inability of the replicating authors to implement the original authors' approach correctly. The well-known work of Young (2019) on the robustness of estimated treatment effects in RCTs is a prominent example of this type of exercise. After replicating the findings of the original study, Young explores robustness with respect to regression design, in particular the role of outliers.

Our exercise falls under the category of scientific replication, for two reasons. First, the data used in the original MRS study is proprietary, and thus a pure replication would be difficult, if not impossible, to implement. To circumvent this issue, we combine various sources of alternative data, all freely or easily accessible, and replicate the analysis of MRS on these alternative data. The data on house prices and expenditures, which we discuss in the next section, are of comparable quality, with some advantages and some disadvantages relative to the data used by MRS. One notable advantage of the KNRS expenditure data compared to the credit card transaction data used in MRS is that they allow disentangling price (e.g., mark-ups) vs. quantity effects, which is an important distinction for interpreting the findings.

Moreover, for research in economics — particularly those concerning macroeconomic issues — often the main findings concern the relative importance of different economic mechanisms, rather than the magnitude of particular estimated parameters. Such conclusions are often arrived at by combining empirical estimates with a particular theoretical framework adopted by the authors. Indeed, it is often the case that the chosen theoretical framework guides the choice of empirical methodology and the specification. For this type of research, we view scientific replication even more broadly. We view replication to include using alternative theoretical frameworks to interpret the same empirical estimates and to suggest different empirical specifications consistent with such alternative frameworks. In the context of our replication, we

argue that from the point of view of standard consumption theory, the MRS specification and the estimated values of the elasticity of expenditure to housing can be reconciled with an economic mechanism based on wealth effect, without any role for leverage and credit constraints which, instead, featured prominently in the narrative surrounding the MRS findings.

Finally, we note that one of the central issues in the debate around the MRS estimates is the validity of the Saiz (2010) housing supply elasticities as an instrument for regional differences in house prices. We abstract from this point in our analysis. Although we think that this is an extremely important issue, a number of existing papers have already discussed potential endogeneity problems (Davidoff, 2013).

3. Data sources

3.1. 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.[1]

Table 1 shows the breakdown of goods sold at stores in the KNRS sample by department code in 2006.[2] The KNRS bundle is mostly composed of non-durables 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).[3]

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.[4] The left panel of Fig. 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 aligned well with these measures. The only significant discrepancy occurs in 2009–10 when KNRS expenditures are flat whereas NIPA expenditures rise.[5]

The right panel of Fig. 1 shows a scatter plot of the state-level 2006–09 change in expenditures in the KNRS data versus the NIPA data (state-

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

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

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

level is the finest level of geographic aggregation for expenditure data that is published by NIPA). The correlation between 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 toward 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.[6]

3.2. From the KNRS bundle to total non-durables

A possible concern throughout our empirical analysis is that our measure of expenditures obtained from the KNRS data may be rather narrow. One may worry that these categories could display different dynamics from total non-durable expenditures. In this section, we use the Consumption Expenditure survey (CE) to estimate the elasticity of total non-durables to a subset of expenditures that is as close as possible to the KNRS bundle. This number can then be used to rescale the various expenditure elasticities to changes in housing net worth estimated in the previous sections. The implicit assumption we are making is that the unconditional correlation between total non-durables and the KNRS bundle that we estimate is close enough to the correlation conditional on the housing net worth shock.

Our aim is to estimate

$$\log c_{it}^{ND} = \mathbf{D}_t + \beta_0' \mathbf{X}_{it} + \beta_1 \log c_{it}^{KN} + \varepsilon_{it}, \qquad (1)$$

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).[7] 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

---

[1] 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).

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

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

[4] In Section 4.1, we also show that the KNRS bundle accounts for roughly 40% of expenditures in non-durable goods and services.

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

[6] 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.

[7] We refer the reader to their paper for a description of the data. For an even more detailed presentation, see Battistin (2003). We thank Erich Battistin for sharing the data.
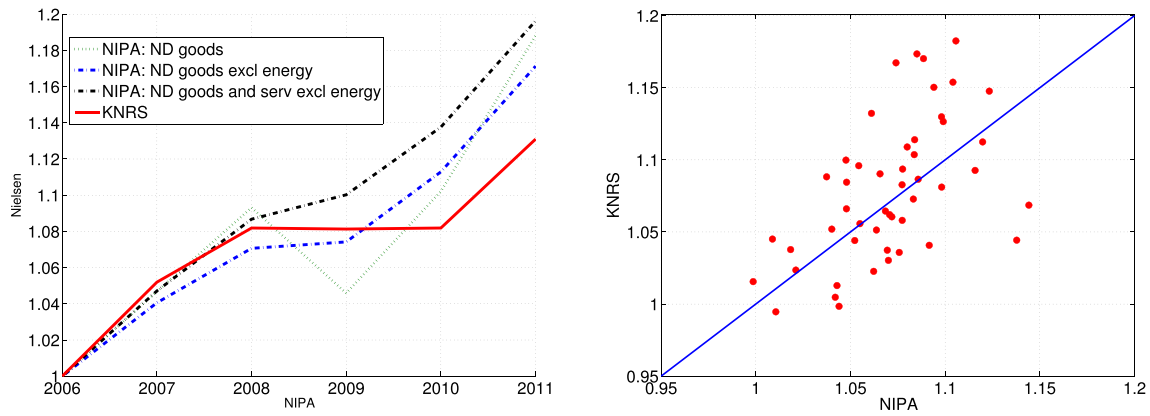
Fig. 1. Left panel: Total annual sales for continuing stores in KNRS data vs various definitions of non-durable expenditures in NIPA, all normalized to 1 in 2006. Right panel: 2006–09 state-level sales growth for continuing stores in KNRS vs non-durable expenditures in NIPA.

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.[8] This is an important consideration for us since the attenuation bias from measurement error tends to artificially reduce the estimate of the elasticity of total non-durables to KNRS expenditures, the coefficient β_1 in Eq. (1).

The sample in Attanasio et al. (2005) covers a large set of items belonging to non-durable goods and services for survey years 1986–2001, 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 non-durable 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.[9]

In our DS sample, median KNRS expenditures are 72 (39) percent of median spending in ND goods (ND goods and services). For comparison, the same calculation from the NIPA Table 2.4.5 for 2000 yields 70 (39) percent.

In the regression (1) we control for year dummies (which capture changes in the relative price of the KN bundle to total non-durables) as well as an equivalence scale, a polynomial in age, and indicator variables for family type, race, education, and region.

The estimates in Table 2 suggest that the elasticity of non-durable expenditures to the KNRS bundle varies between 0.68 and 0.90, depending on how broad the definition of non-durable expenditures is. Thus the estimated elasticities of expenditures to housing wealth obtained in the following sections should be reduced by around 10 to 30% when interpreting them in terms of the effects on total non-durable expenditures.

3.3. 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 explain how we construct each of these variables. Each region i is a county, which we later aggregate into CBSAs.

3.3.1. Financial assets
We follow the corresponding calculation in MRS. 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 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.

3.3.2. 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

---

[8] 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 measure changes in consumption inequality over time more accurately.

[9] 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 et al. (2008), these services are more durable in nature and more closely resemble investment and saving activities than non-durable expenditures.

Table 2
Elasticity of total non-durable expenditures to expenditures in the the KNRS bundle. Source CEX.

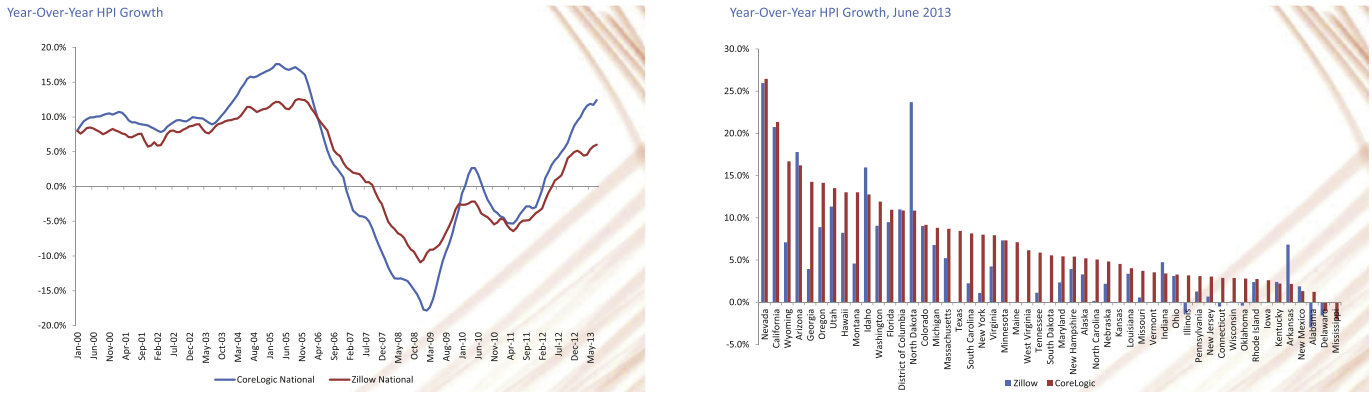| Dependent variable:	$\log c_{it}^{NDgoods}$	$\log c_{it}^{NDgoods\&serv}$
$\log c_{it}^{KN}$	0.905	0.679
	(0.003)	(0.004)
Other controls	Y	Y
N	37,892	37,892
R^2	0.81	0.54

Fig. 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).

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 Fig. 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 Fig. 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.

3.3.3. 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.[10]

4. Replication I: alternative data

In this section, we focus on the robustness of the MRS findings with respect to alternative data. In particular, we adopt the same regression specification as in MRS. 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 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 (Fig. 2), and is the three-year period studied in MRS. However, since Fig. 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 2010 and 2011.

Our OLS regression specification is

$$\Delta \log C_{06-09}^{s,i} = \beta_0 + \beta_1 \Delta \log p_{06-09}^i \left(\frac{H_{06}^i}{NW_{06}^i} \right) + \varepsilon_{06-09}^{s,i}. \tag{2}$$

where the dependent side variable is sales of KNRS goods in store s in CBSA i. The right-hand side 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.

In the left panel of Fig. 3 we show a scatter plot of the change in CBSA-level sales, i.e., $\Delta \log C_{06-09}^i$ 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 that 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.
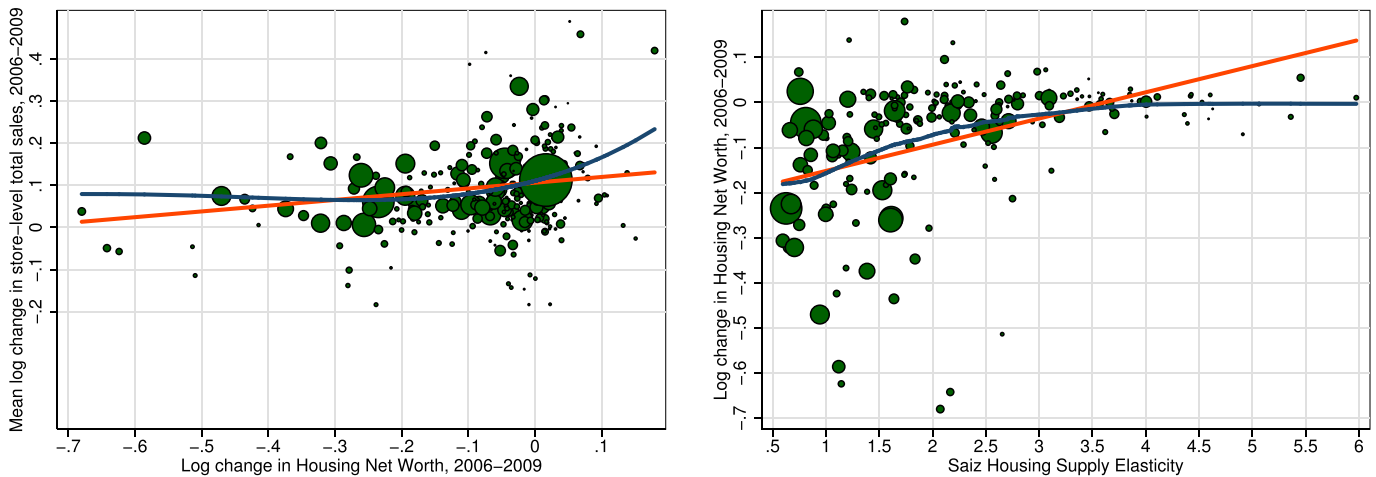
Our IV regressions follow the specification closely in MRS. We use the estimates of housing supply elasticities from Saiz (2010) to instrument for the housing net worth shock. This instrument is provided at the CBSA level and is not available for all of the CBSA's in which we observe store-level changes 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 Fig. 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 elasticity areas where the reduced form relationship is strongest (Fig. 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]

---

[10] The CCP data by county we use here was publicly available at the time of our study.

[11] In Section 5 we consider reasonable alternative ways of constructing the right-hand side variable in this regression.

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

[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.

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

Table 3 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 3, 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 3, 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 (2) 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 3, fourth and fifth columns). Hence the difference in 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 3, 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 the time period. In Table 4 we report corresponding estimates for alternative time periods around the Great Recession. Since the

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

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, such as changes in income or unemployment: the presence of a common factor would lead to an upward bias in the OLS estimates.

A possible explanation of why IV estimates are larger is that idiosyncratic variation in house prices is more transitory than the variation in the common component of house prices isolated by the instrument. More persistent house price movements should have a bigger effect on expenditures. Fig. 3 suggests a different, but related, interpretation. As we pointed out, the relation between changes in expenditures and housing net worth is weaker in areas where house prices collapsed. Perhaps, households there perceived a larger share of this drop to be temporary (an over-reaction). The presence of these observations lowers the OLS elasticity toward zero. The change in housing net worth predicted by the instrument is much smaller in those areas (see the right panel of Fig. 3), which increases the elasticity estimated by IV. Finally, the most pessimistic explanation for the downward bias in the OLS is that the Saiz instrument is also correlated with a number of other factors driving local housing demand and are thus invalid (Davidoff, 2016). We are open to this as a possible explanation.

4.1. Advantages of KNRS data: expenditures versus consumption

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 non-durable goods consumed is also of interest. One advantage of the KNRS data compared to, for example, the Mastercard data used by MRS is that we separately observe nominal sales and prices.

Table 3
Elasticity of non-durable expenditures to housing share of net worth.

| | 2006–09 | | | 2006–09 | | 2006–09 | |
| | CBSA | | | County | | 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

We construct a series for store-level real sales by aggregating product-level sales (at the bar-code level) using product-level prices that are common across stores at a fixed date. In our baseline measure, we use 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.[15]

When using real 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 5).

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 findings in Stroebel and Vavra (2014), who argued that a decline in markups is responsible for these price dynamics.

5. Replication II: specifications and interpretation

5.1. Alternative specifications

In this section, we discuss the sensitivity of the MRS results with respect to alternative regression specifications. In the original MRS regression specification (2), the right-hand side variable is the housing net worth shock, i.e., the change in the housing share of net worth induced by the change in house prices, $\Delta \log p^i_{06-09}(\frac{H^i_{06}}{NW^i_{06}})$. MRS focus on this specification because they want to emphasize that the initial heterogeneity in household balance sheets across geographical areas is a strong source of variation in the degree of exposure to the drop in house prices.

According to this narrative, one would expect two regions experiencing the same decline in house prices to differ in the size of the shock (and its impact on expenditures) depending on the initial share of housing net worth for households living in those regions: regions where the housing share is larger should experience larger shocks. Similarly, regions where households initially have higher housing leverage should experience larger shocks. This is what MRS call the *household balance sheet channel*.

Since the right-hand side variable in Eq. (2) is an interaction between local house price changes $\Delta \log p^i_{06-09}$ and the initial share of housing in net worth $\frac{H^i_{06}}{NW^i_{06}}$, one can look for evidence of the balance sheet channel by estimating the following more general specification:

$$\Delta \log C^{s,i}_{06-09} = \beta_0 + \beta_1 \Delta \log p^i_{06-09}\left(\frac{H^i_{06}}{NW^i_{06}}\right) + \beta_2 \Delta \log p^i_{06-09}$$
$$+ \beta_3\left(\frac{H^i_{06}}{NW^i_{06}}\right) + \varepsilon^{s,i}_{06-09}. \quad (3)$$

In Table 6 we report estimates of alternative specifications based on Eq. (3). The estimate of the elasticity of non-durable expenditures with respect to house prices, $\hat{\beta}_2$, 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 6 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 3. However, the R^2 from these regressions are slightly higher than the corresponding regressions in Table 3. 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 statistically indistinguishable from zero (columns 3 and 4). Interestingly, when we use a direct 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]

Taken at face value, the results in Table 6 offer at best weak support to the view that the degree of initial exposure to the shock in 2006 was a key determinant of the drop in expenditures that occurred between 2006 and 2009. However, several caveats 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 find that the interaction is statistically significant in determining the response of auto spending.[17] Thus, our finding of a weak balance sheet effect could be due to the fact that our dependent variable is *non-durable* expenditures.[18]

Second, estimating the partial elasticities 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) 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. Fig. 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 6 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 nondurables than for durables is a genuine finding or an artifact of excessive geographical aggregation (and hence low power) is a question that can

---

[15] We have tried various alternative choices, and our results are barely affected. Our results are also robust to using a weighted (across stores) average price for each product, or to restricting attention to products that are present in every year of the Nielsen data.

[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.

[17] Also Baker (2017) finds that higher levels of leverage are significantly related to a 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.

Table 4
Elasticity of non-durable expenditures to housing share of net worth in alternative time periods.

	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

only be properly answered with highly disaggregated data on non-durable expenditures, data that at the moment are not available.

Third, there could be greater measurement error in H/NW or $H/(H - M)$ compared to $\Delta \log p$. There are two potential sources of measurement error here: (1) The assumption that the representative household in each county/quarter holds the market index for stocks and bonds. Some support for this explanation comes from the last column of Table, where the interaction term measured as housing leverage, and hence without data on financial wealth, remains mildly significant. (2) Mismeasurement of the *level* of H. It is possible, for example, that changes in housing values may be more accurately measured than levels. Put differently, the fact that there is some amount of orthogonal variation in Fig. 4 doesn't help much if that orthogonal variation is primarily measurement error. The implied attenuation bias could explain our findings.

Finally, 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.

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 7. 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 4 − 0.17 (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).

5.2. Structural interpretation of different specifications

The discussion in the previous section demonstrated that, without any guidance from theory, it is unclear what the right regression specification is and whether one should expect initial exposure or initial

leverage to play any role. To make progress, in this section we build on the analysis of Berger et al. (2015) to show that the elasticity of expenditures to both house prices and the housing net worth shock can be given a structural interpretation and expressed in terms of observables. A simple back-of-the-envelope calculation yields values for these elasticities that are very much in line with those of Mian et al. (2013) and our estimates from the previous sections.

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. Note that the model is frictionless: there is no debt, no binding borrowing constraints, nor transaction costs for housing.

Denote human wealth as Υ_t^*, financial wealth including current interest as A_t^*, housing quantity net of depreciation as Q_t^* and housing wealth as H_t^*, i.e.

$$\Upsilon_{it}^* = \sum_{\tau=t}^{T} (1+r)^{t-\tau} Y_{i\tau}$$
$$A_{it}^* = (1+r)A_{it}$$
$$Q_{it}^* = (1-\delta)Q_{it}$$
$$H_{it}^* = p_t Q_{it}^*.$$

We assume that $\beta(1+r)=1$, the path of income $\{Y_{it}\}$ is deterministic and $p_t = p$ for every t. In Appendix A, we show that:

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

Eq. (4) states that the elasticity of expenditures with respect to house prices is simply the initial housing share of total wealth, including human wealth. Households that are "overweight" on housing should respond more to house prices.

Table 5
Elasticity of real non-durable consumption to housing net worth shocks.

	2006–09				2007–11			
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
ΔHNW^i	0.196**	0.290**			0.203**	0.455**		
	(0.026)	(0.085)			(0.026)	(0.143)		
$\Delta \log p^i$			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	0.012	0.001	0.013	0.002

Table 6
Elasticity of non-durable expenditures to house prices: alternative specifications to assess the strength of the household balance-sheet channel.

	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)
$\Delta \log p^i (\frac{H^i}{NW^i})$			0.050	−0.043		
			(0.097)	(0.128)		
$\frac{H^i}{NW^i}$				−0.053		
				(0.039)		
$\Delta \log p^i (\frac{H^i}{H^i - M^i})$					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

Dividing and multiplying Eq. (4) by $H^*_{i,t-1}/(H^*_{i,t-1} + A^*_{i,t-1})$ yields:

$$\frac{\Delta \log C_{it}}{\Delta HNW_{it}} = \frac{H^*_{i,t-1} + A^*_{i,t-1}}{\Upsilon^*_{it} + H^*_{i,t-1} + A^*_{i,t-1}}, \qquad (5)$$

which states that the elasticity of expenditures with respect to the housing net worth shock is the share of non-human wealth over total wealth.

In sum, both elasticities we estimated in Sections 4 and 5 have a structural interpretation. However, it isn't clear that there is any advantage to using the MRS specification (5) relative to (4). In particular, Eq. (4) also captures the role of initial exposure to housing net worth transparently. It does so not through the size of the shock as in specification (5), but through the degree of transmission to expenditures, i.e., the magnitude of the elasticity itself. There are, instead, some major disadvantages of using Eq. (5). First, house prices are measured directly, while the disaggregated housing net worth variable needs to be constructed using additional data by making strong assumptions. In addition, HNW_{it} might introduce spurious variation that has nothing to do with the housing market. Suppose that nothing at all happens to house prices. We might nevertheless measure a shift in HNW_{it} because of a stock market shock that moves A_{it} or a labor market shock that affects Υ_{it}. This makes the reduced form correlation between consumption and the housing value variable even more difficult to interpret. Finally, the housing supply elasticity gives an instrument for house prices, not for housing net worth, so that the IV specification relates far more directly to the specification in Eq. (4).

This structural model also provides a way to check ex-post whether the empirical estimates of the elasticities based on geographical variation obtained in Sections 4 and 5 are reasonable. Consider the following simple calculation. Average household net worth $H^* + A^*$ in 2007 (from the Survey of Consumer Finances) was around $500,000, the housing share of net worth was around 1/2, 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 $\Upsilon^*_{it} = \sum_{\tau=t}^{T}(1+r)^{t-\tau}Y_{i\tau}$.

Combining these pieces together gives a value for the elasticity with respect to house prices of around $250/(1,000 + 500)$=0.165. This calculation, which does not rely on identifying empirically exogenous sources of variation for house prices, compares well with our estimates in Tables 6.

Since the housing share of net worth is around 1/2, our stylized model suggests the elasticity with respect to the housing net worth shock should be roughly two times larger than that with respect to house prices, or 1/3. This second back of the envelope calculation also compares reasonably well with the estimates in Tables 3 and 4.
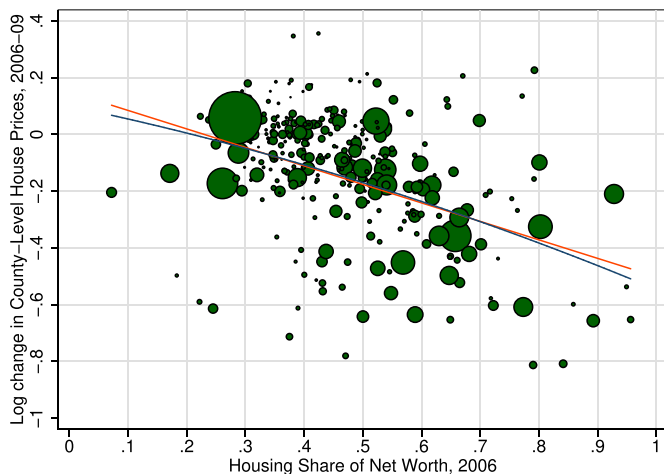
An important final point is that these specifications are derived from a deterministic frictionless model in which the consumption function is linear. Thus, elasticities of expenditures to housing net worth of the size estimated on micro data around the Great Recession are entirely consistent with a pure housing wealth effect and do not require any role for debt, leverage and credit constraints.

6. 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 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

Fig. 4. Scatterplot of Log change in housing net worth from 2006 to 2009 versus initial share of housing in household net worth across CBSAs.

Table 7
Elasticity of non-durable expenditures to gross and net housing wealth.

	2006–09			
	OLS	IV	OLS	IV
$\Delta \log H^i$	0.124**	0.183**		
	(0.019)	(0.038)		
$\Delta \log (H^i - M^i)$			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

in the Great Recession. Our replication exercise is developed in two parts.

First, we overcome the hurdle imposed by the fact that Mian et al. (2013) used expensive and hard to access proprietary data by combining public and easily accessible data that allow us to perform their same empirical study. Our analysis using these alternative data largely confirms their results, therefore ruling out the possibility that the strong co-movement of expenditures and house prices in the Great Recession is due to peculiarities of their data sources. In addition, we can estimate the effect of the collapse separately in housing wealth on prices and quantities: real consumption drops approximately 20% less than nominal expenditures, implying a sizable demand-induced fall in producer prices.

Second, we explore the robustness of the Mian et al. (2013) results to alternative specifications. One of our key conclusions is that, after controlling for the drop in house prices, we do not find much of an independent effect of initial housing exposure and initial leverage on nondurable expenditures. We offer several possible explanations for this result. We also show that it is possible to interpret the estimated elasticity, theoretically and quantitatively, as a pure housing wealth effect, with no role for leverage and credit constraints. This is, we believe, a useful observation given the strong emphasis put on debt and leverage in the narrative surrounding the findings of Mian et al. (2013).

Appendix A. Derivation of Eqs. (4) and (5)

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

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

The FOCs are:

$$\alpha C_t^{\alpha(1-\sigma)-1} Q_t^{(1-\alpha)(1-\sigma)} = \lambda_t$$
$$\beta^t(1-\alpha)C_t^{\alpha(1-\sigma)} Q_t^{(1-\alpha)(1-\sigma)} = \beta^t \lambda_t p - \beta^{t+1}\lambda_{t+1}(1-\delta)p$$
$$pt\beta^t\lambda_t = \beta^{t+1}\lambda_{t+1}(1+r)$$

Rearranging, we obtain

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

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

Iterating forward on the budget constraint:

$$C_t + pQ_t = Y_t + (1+r)A_{t-1} + p(1-\delta)Q_{t-1} - A_t$$

we obtain

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

Continuing iterating, using the optimality Condition (A1), and exploiting the property that C_t is constant over the lifecycle yields

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

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

$$C_t \approx \alpha(1-\beta)\left[\sum_{\tau=t}^{T}(1+r)^{t-\tau}Y_\tau + p(1-\delta)Q_{t-1} + (1+r)A_{t-1}\right], \qquad (A2)$$

where we use the fact that $(1-\beta) = \frac{r}{1+r}$..

Denote human wealth as Υ_t^*, financial wealth comprising current interests as A_t^*, housing quantity net of depreciation as Q_t^* and housing wealth as H_t^*, i.e.

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

Differentiating Eq. (A2) 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

$$\frac{\Delta C_{it}/C_{it}}{\Delta p/p} = \frac{\alpha(1-\beta)pQ_{i,t-1}^*}{C_{it}} \qquad (A3)$$
$$pt = \frac{H_{i,t-1}^*}{\Upsilon_{it}^* + H_{i,t-1}^* + A_{i,t-1}^*},$$

where we used Eq. (A2) in the denominator. This is Eq. (4) in the main text.

Switching to log notation (e.g., $\Delta C/C = \Delta \log C$) and multiplying and dividing Eq. (A3) by the housing net worth share yields:

$$\Delta \log C_{it} = \left(\frac{H_{i,t-1}^* + A_{i,t-1}^*}{\Upsilon_{it}^* + H_{i,t-1}^* + A_{i,t-1}^*}\right) \times \Delta HNW_{it}.$$

which is Eq. (5) in the main text.

Appendix B. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.jpubeco.2020.104176.

References

Aguiar, M., Hurst, E., 2005. Consumption versus expenditure. J. Polit. Econ. 113, 919–948.

Alm, J., Reed, W.R., 2015. The need for replications. Public Finance Review 43, 139–142.

Attanasio, O., Battistin, E., Ichimura, H., 2005. What really happened to consumption inequality in the US? In: Berndt, Ernst R., Hulten, Charles R. (Eds.), Hard to Measure Goods and Services: Essays in Honor of Zvi Griliches. University of Chicago Press, Chicago

Baker, S.R., 2018. Debt and the response to household income shocks: validation and application of linked financial account data. J. Polit. Econ. 126 (4), 1504–1557. https://doi.org/10.1086/698106.

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

Begley, C.G., Ioannidis, J.P., 2015. Reproducibility in science: improving the standard for basic and preclinical research. Circ. Res. 116, 116–126.

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

Blundell, R., Pistaferri, L., Preston, I., 2008. Consumption inequality and partial insurance. Am. Econ. Rev. 98, 1887–1921.

Burman, L.E., Reed, W.R., Alm, J., 2010. A call for replication studies. Public Finance Review 38, 787–793.

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

Cochrane, J., 2015. Secret Data [Online; posted 28-December-2015].

Cochrane, J., 2016. Secret Data Encore [Online; posted 5-January-2016].

Davidoff, T., 2013. Supply elasticity and the housing cycle of the 2000s. Real Estate Econ. 41, 793–813.

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

Hamermesh, D.S., 2007. Replication in economics. Canadian Journal of Economics/Revue canadienne d'économique 40, 715–733.

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

Mian, A., Sufi, A., 2016. Who bears the cost of recessions? The role of house prices and household debt. Handb. Macroecon. 2, 255–296.

Mian, A., Rao, K., Sufi, A., 2013. Household balance sheets, consumption, and the economic slump. Q. J. Econ. 128, 1687–1726.

Munafò, M.R., Nosek, B.A., Bishop, D.V., Button, K.S., Chambers, C.D., Du Sert, N.P., Simonsohn, U., Wagenmakers, E.-J., Ware, J.J., Ioannidis, J.P., 2017. A manifesto for reproducible science. Nat. Hum. Behav. 1, 1–9.

Nosek, B.A., Alter, G., Banks, G.C., Borsboom, D., Bowman, S.D., Breckler, S.J., Buck, S., Chambers, C.D., Chin, G., Christensen, G., et al., 2015. Promoting an open research culture. Science 348, 1422–1425.

Petev, I., Pistaferri, L., Eksten, I.S., 2012. Consumption and the Great Recession: an analysis of trends, perceptions, and distributional effects. In: Weimer, David (Ed.), The Great Recession. Russell Sage, NY.

Saiz, A., 2010. The geographic determinants of housing supply. Q. J. Econ. 125, 1253–1296.

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

Young, A., 2019. Channeling Fisher: randomization tests and the statistical insignificance of seemingly significant experimental results. Q. J. Econ. 134, 557–598.