Household Balance Sheets, Consumption, and the Economic Slump

Atif Mian
Kamalesh Rao
Amir Sufi

APPENDIX

1. Data

Here is the list of sources for all of the data used in our analysis.

- County-level housing net worth shock: We cannot make the exact county-level housing net worth shock used in our analysis publicly available given restrictions in our agreements with Equifax and CoreLogic. However, we can make a CBSA-level data set available that replaces Equifax debt with county-level debt from the Federal Reserve Bank of New York, and replaces CoreLogic house price growth from 2006 to 2009 with FHFA house price growth from 2006 to 2009. At the CBSA level, a regression of the housing net worth shock used in the study on the publicly available version yields an R2 of 0.94 and a coefficient of 0.87. The estimated elasticity of spending growth on the housing net worth shock is almost identical using the public version of the data. The main disadvantage of the public version is a smaller sample. The public version of the data is available here:

  http://faculty.chicagobooth.edu/amir.sufi/data-and-appendices/MianRaoSufiQJE_housingnetworthshock.dta
  http://faculty.chicagobooth.edu/amir.sufi/data-and-appendices/MianRaoSufiQJE_housingnetworthshock.txt

- Zipcode-quarterly data from Equifax: These data are proprietary and cannot be shared without data provider approval. Please contact Linda Mueller at the Equifax and mention that you are interested in the zip-code quarterly aggregate CreditTrends data used by Mian and Sufi: linda.mueller@equifax.com.

- Zipcode-yearly data from HMDA: These data are available at: http://www.ffiec.gov/hmda/orderform.htm

- House price data: CoreLogic. These data are proprietary and cannot be shared without data provider approval. We use CoreLogic's indices available at the zipcode, county, and state level at a monthly frequency. Please contact Ryan Meyers at rmeyers@corelogic.com.

- CBSA-level housing supply elasticity from Saiz (2011): These data are available at Albert Saiz's website: http://real.wharton.upenn.edu/~saiz/

- Zipcode-monthly auto sales data from R.L. Polk: These data are proprietary and cannot be shared without data provider approval. Please contact Rob Sacka at R.L. Polk and ask him for the zip code-monthly level auto sales data used by Mian and Sufi: robert_sacka@polk.com.

- County-monthly data on household spending from MasterCard Advisors: These are not available to the public at this time.

- Zip level data on income: These data are from the IRS and the Stata version can be downloaded at this link: http://www.irs.gov/taxstats/indtaxstats/article/0,,id=96947,00.html

- 2000 decennial census zip level data: These data are from the census and can be downloaded at this link: http://www.census.gov/main/www/cen2000.html

- County level population from the census: These data are from the census and can be downloaded at this link: http://www.census.gov/popest/data/historical/index.html

- Zip code level fraction of homeowners underwater as of 2011. These data were provided by Zillow. Please email: Katie Curnutte [katiec@zillow.com].

- County-year data on construction employment share from County Business Patterns. See Mian and Sufi (2012) for more information and data availability.

Here is a table summarizing the data used in this study, and the most disaggregated availability

Appendix Table 1
Data Availability

Level of aggregation:	Zip Code	County	CBSA	Publicly available?
Data sets used in this study				
R.L. Polk auto sales	X			N
CoreLogic house price growth	X			N
Equifax (debt, defaults, credit scores)	X			N
HMDA (refinancing volume)	X			Y
IRS Income from tax returns	X			Y
Zillow homeowners underwater	X			N
Decennial Census 2000	X			Y
MasterCard spending		X		N
Saiz housing supply elasticity			X	Y
Public data sets				
FRBNY debt data		X		Y
FHFA house price growth			X	Y

2. MasterCard retail sales versus Census retail sales

The MasterCard data we utilize is based on a 5% sample of in-store purchases. It includes purchases using either debit or credit cards that are part of the MasterCard network. MasterCard also has data available on purchases made out of store (e.g., internet purchases) but we do not utilize these data as they cannot be matched to the geographic location of the purchaser. The inclusion of debit cards is particularly important given the increase in their use over the last decade (see Schuh and Stavins (2011)).

The MasterCard data serve two important purposes in our analysis. First, we use the share of total MasterCard purchases in a given county to allocate total retail sales in that county. Second, we use the growth in MasterCard purchases from 2006 to 2009 to project forward the total spending in a county in 2009.

How representative are purchases using MasterCard of total household spending? Unfortunately, we are quite restricted in our ability to relay specific information on the share of MasterCard purchases of total retail sales, given competitive concerns of MasterCard. Nonetheless, we can provide enough information here that should give the reader confidence that the MasterCard data is reflective of total aggregate household spending.

First, as background, purchases using debit and credit cards over our sample period (2005 to 2009) represented a very large fraction of total purchases. Foster, Meijer, Schuh, and Zabek (2010) show that in 2008 52% of all transactions by households were done via credit or debit card (although this is not a dollar weighted figure which is unavailable). Further, publicly available data from Card Hub suggests that as of 2010, MasterCard had a market share of 27% of all purchase volume (http://www.cardhub.com/edu/market-share-by-credit-card-network/). While we cannot confirm or deny this share, this estimate suggests that we are capturing a significant fraction of household spending.

Second, the MasterCard data does an excellent job tracking the aggregate patterns over time in the Census Retail Sales data. The Census Retail Sales data are only available at the aggregate level, so we can only do comparisons on a national basis.

In Appendix Figure 1 below, we show a scatter plot of total household spending from the Census Retail Sales and the MasterCard data. Each dot represents one of the years from 2005 to 2009. We cannot label the years and we cannot label the axes (given that this would give away market share of MasterCard).

Even with these restrictions, it is easy to see that the two data sets track each other very closely over time. If all dots were on an exact straight line, that would imply that MasterCard kept an exactly constant share of total expenditures through the 2005 to 2009 period. As the figures show, this is almost true. In fact, a regression of total Census Retail Sales by year on total MasterCard sales yields an R^2 of 0.97 for total sales.

In particular, while the share of total retail sales represented by our MasterCard data varies slightly over the sample period, it is very similar in 2006 and 2009, the two points in time that are most important in our regression specifications.

The bottom line on the aggregate analysis is that the MasterCard data do an excellent job of tracking the aggregate retail sales data from the Census.

Appendix Figure 1
MasterCard Total Purchases and Census Retail Sales Purchases, by Year

This figure plots total MasterCard purchases and total Census Retail Sales purchases by year. Purchases of autos (NAICS = 441) are excluded. We are not allowed to show the axis labels or axis values because they would reveal MasterCard's share of the total. We are also not allowed to label the years (each dot in the graph represents a year). Nonetheless, as these graphs show, there is a strong correlation between the two. For the years when Census Retail Sales are low, MasterCard purchases are also low.

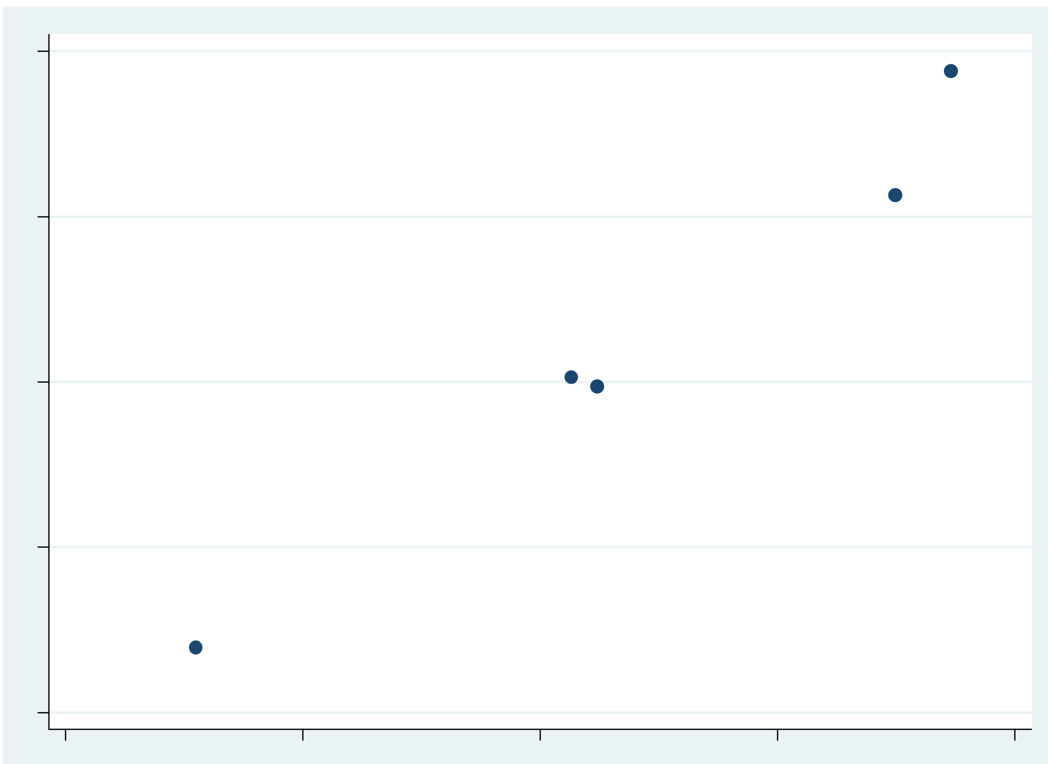

3. Cross-Sectional Issues

Another issue with our measurement of total spending is based on the cross-section. We find strong evidence in the study that counties with negative housing net worth shocks cut spending the most. Does our measurement of total spending bias our results toward finding such an effect?

In columns 1 through 4 of Appendix Table 2 below, we report specifications relating the growth in auto sales from R.L. Polk, and the growth in expenditures on other durables, groceries, and other non-durables from MasterCard. These specifications use growth rates directly from the raw data instead of any aggregation procedure. As expected, the elasticity of spending with respect to the housing net worth shock is strongest for durables and weaker for non-durables. This helps ensure that our aggregation procedure is not responsible for our findings. The strong relation between net worth shocks and spending is seen in the raw data.

Another concern is specific to the Master Card data: Perhaps counties experiencing a large negative housing net worth shock experience a relative decline in the fraction of transactions done via credit or debit card from 2006 to 2009. Or in other words, we may mistakenly find a reduction in household spending in these counties when in fact people are switching into cash disproportionately in these areas. Of course, given the evidence in Appendix Figure 1 that the aggregate share is similar in 2006 and 2009, this would imply that counties avoiding the housing collapse are seeing an increase in their card purchase to total purchase ratio.

We cannot directly test this hypothesis given that we do not have county-level data tracking the fraction of transactions done with a credit or debit card. However, there are three reasons to be skeptical of this argument. First and foremost, the auto sales results from R.L. Polk are completely independent of this concern. They represent all auto purchases, and we see a strong relative decline in high leverage counties. In fact, the estimates for auto sales are actually larger than the estimates for other durables from MasterCard, and much larger than for groceries and other non-durables. Given this fact, it seems unlikely that this bias could be driving our results.

Second, our MasterCard data include debit purchases which are a significant fraction of all credit purchases. According to Foster, Meijer, Schuh, and Zabek (2010), debit transactions make up 60% of all card transactions. The concerns on credit card availability do not apply to debit cards.

Third, we can use state level information on state sales tax revenues to measure household spending as a cross-check on the results. These are the data used in the work of Zhou and Carroll (2012). There are some disadvantages of state sales tax revenues to measure household spending. For example, sales tax rates may change in some states from 2006 to 2009 (Zhou and Carroll (2012) do an excellent job adjusting the data for this concern), and the information is not available at the county level. Despite these shortcomings, as we show in Appendix Table 1, the results are very similar if we use state-level regressions related the growth in sales tax revenue to the 2006 state-level debt-to-income ratio. For comparison purposes, we report in column 5 the state-level coefficients using total spending from our data. If anything, the coefficient is slightly larger using the state sales tax data.

The similarity of the results using our measure of total spending from R.L. Polk and MasterCard and the state sales tax revenue data convinces us that there is no significant data-specific bias in the cross-sectional analysis.

Appendix Table 2
Checking for Cross-Sectional Bias due to Data

In columns 1 through 4, we use growth rates from R.L. Polk and MasterCard directly. In column 6, we use state level data on sales tax revenues.

Dependent variable	(1) Auto sales growth, 06-09	(2) Other durables growth, 06-09	(3) Grocery spending growth, 06-09	(4) Other non-durables growth, 06-09	(5) Total spending growth, 06-09	(6) Sales tax revenue growth, 06-09
Housing net worth shock, 06-09	1.315**	0.829**	0.382**	0.341**	0.785**	0.806**
	(0.127)	(0.131)	(0.098)	(0.112)	(0.156)	(0.141)
Constant	-0.366**	-0.129**	0.139**	0.055**	-0.017	0.083**
	(0.028)	(0.023)	(0.013)	(0.016)	(0.014)	(0.016)
Data source	R.L. Polk	MC	MC	MC	As in study	State sales tax data
Level of analysis	County	County	County	County	State	State
N	944	944	944	944	50	50
R^2	0.392	0.167	0.074	0.093	0.547	0.426

**,* Coefficient statistically different than zero at the 1% and 5% confidence level, respectively

4. Construction of net worth as of 2006

The basic strategy we use to construct zip code level net worth is to use aggregate data from the Federal Reserve Flow of Funds, and to allocate the aggregate data to zip codes based on the fraction of variables we have available to us in the zip code. We need three key variables: total debt in a zip code, total housing wealth in a zip code, and total financial wealth in a zip code.

For total household debt, we take from Equifax the zip code level fraction of aggregate debt in Equifax, and then we allocate household debt from the Flow of Funds based on this fraction. So if a zip code has 5% of aggregate Equifax debt, we will allocate that zip code 5% of the Flow of Funds aggregate debt. This provides us total debt in a zip code as of 2006. We could alternatively use the Equifax data directly with little change.

The denominator contains housing assets and financial assets. For housing assets, we use the zip code level 2000 median house value from the Census. We then scale this up by house price growth from 2000 to 2006 from CoreLogic. We then have the median home value in a zip code as of 2006. To get the value of all homes in the zip code, we need an estimate of the number of homeowners in the zip code as of 2006. To obtain this, we use the total number of homeowners in the zip code as of 2000 from the census, and we grow this number by national population growth and the increase in homeownership from 2000 to 2006. We apply the national population and homeowner growth numbers equally to all zip codes.

Financial assets are much harder to construct given the lack of zip code level measures of financial wealth. We instead utilize flows from the IRS Statistics of Income. The SOI has four potential variables one could use to measure flows to financial assets: interest income, dividend income, realized capital gains, and all non-wage income. Obviously the last category includes the first three.

As we show below in Appendix Table 3, these four variables are all highly correlated. We choose to use all non-wage income to approximate financial income, but could use any of the four and get very similar results.

Appendix Table 3
Non-wage income correlations across zip codes

This table presents the correlation matrix of income per capita variables for zip codes. All variables are in per household terms. All correlations are significant at the 1% level of confidence.

	Non wage income	Interest income	Dividend income	Capital gains income
Non-wage income	1.000			
Interest income	0.929	1.000		
Dividend income	0.884	0.681	1.000	
Capital gains income	0.872	0.851	0.784	1.000

Using the non-wage income variable as a proxy for financial asset income, we calculate for each zip code the fraction of all aggregate financial asset income earned by the zip code. We then use this fraction to allocate the aggregate financial wealth from the Federal Reserve Flow of Funds. So if a zip code has 5% of aggregate financial asset income, we allocate to that zip code 5% of aggregate financial wealth.

This procedure produces a mean leverage ratio across zip codes of 0.21 and a housing wealth to (housing wealth+financial wealth) ratio of 0.26. From the flow of funds, the aggregate measures are 0.18 and 0.33, respectively. So in other words, this procedure produces aggregate numbers that are in line with the Flow of Funds data.

Of the three variables used to construct net worth, the measure of financial wealth is likely to have the most error. However, it is important to understand that by construction the total amount of financial wealth across zip codes must add up to total financial wealth in the economy. Therefore, the primary concern would be whether the use of financial income to allocate financial wealth gets the *ordering* of the zip codes substantially wrong.

We have no way of directly testing how much measurement error is in our net worth measure. But in Appendix Table 4 below, we provide some correlations that give us confidence that we are accurately capturing differences in net worth across zip codes. As it shows, our measure of net worth per household is strongly correlated with income per capita and education levels.

Appendix Table 4
Net Worth Correlations

This table presents regressions relating our measure of net worth per household as of 2006 to other variables. The level of analysis is zip codes. In column 4, the omitted category is the fraction with less than a high school education.

	(1)	(2)	(3)	(4)
Dependent variable	Net worth per household, 2006			
AGI per household, IRS, 2006	12.587**			
	(0.313)			
Wage income per household, IRS, 2006		21.072**		
		(0.453)		
Median HH income, Census, 2000			17.499**	
			(0.517)	
Fraction with high school education				293.802**
				(59.999)
Fraction with some college				167.085
				(85.258)
Fraction with exactly college				772.139**
				(160.902)
Fraction with more than college				3,872.301**
				(274.248)
Constant	-253.518**	-368.322**	-344.668**	-161.306**
	(17.217)	(16.882)	(21.332)	(31.525)
N	6,437	6,437	6,437	6,437
R^2	0.852	0.590	0.362	0.471

**,* Coefficient statistically different than zero at the 1% and 5% confidence level, respectively

5. The Timing of the Consumption Response in Large Negative Housing Shock Counties

As we discuss at the end of Section 3 of the text, one concern is that an income or employment shock drives both the housing and consumption response in inelastic housing supply counties. Appendix Table 5 presents evidence inconsistent with this view. As it shows, the housing net worth shock was larger in inelastic housing supply counties even from 2006 to 2007. In contrast, income growth and employment growth from 2006 to 2007 are uncorrelated with housing supply elasticity. Income growth and employment growth become positively related to housing supply elasticity only in 2008 for income and 2009 for employment. These results suggest that the housing net worth shock happened first in inelastic housing supply counties, and then income and employment responded.

Appendix Table 5
Timing of Shock

This table presents a series of univariate regressions relating the housing net worth shock, employment growth, and income growth in a county to the housing supply elasticity. Each row reports a separate univariate regression. The goal is to show that the housing net worth shock happened in inelastic housing supply counties before the relative decline in employment or income.

	Housing supply elasticity	Constant	R2	N
Housing net worth shock:				
2006-2007	0.013**	-0.046**	0.124	540
	(0.002)	(0.006)		
2006-2008	0.036**	-0.142**	0.162	540
	(0.006)	(0.016)		
2006-2009	0.046**	-0.174**	0.190	540
	(0.006)	(0.017)		
Employment growth:				
2006-2007	-0.001	0.004*	0.004	540
	(0.001)	(0.002)		
2006-2008	0.001	-0.016**	0.000	540
	(0.002)	(0.005)		
2006-2009	0.004	-0.060**	0.007	540
	(0.003)	(0.007)		
Income growth:				
2006-2007	-0.001	0.056**	0.003	526
	(0.002)	(0.004)		
2006-2008	0.009**	0.079**	0.050	526
	(0.003)	(0.005)		
2006-2009	0.014**	0.021*	0.050	526
	(0.004)	(0.009)		

**,* Coefficient statistically different than zero at the 1% and 5% confidence level, respectively

6. Housing net worth shock and 2006 debt to income ratio

In an earlier version of this study, we used the 2006 county-level household debt to income ratio as our main right hand side variable. Below, we show that the housing net worth shock from 2006 to 2009 is highly correlated with the 2006 debt to income ratio. We are exploiting the same underlying variation as the earlier version.

Appendix Figure 2
Correlation of 2006 Household Debt to Income Ratio and Housing Net Worth Shock

7. Accounting for defaulted debt in net worth calculations

As we discussed in the paper, the change in net wealth calculations ignored the effect of default. One can potentially view debt default as a "gain" from the perspective of the defaulting household. However, the net gain of default is considerably smaller since the defaulting homeowner also loses the house they were living in. Nonetheless, since we have information on debt default in our data, we show the robustness of our key results to the inclusion of defaulted debt as a gain for defaulting households.

This is done in Appendix Table 6 below. We first estimate the elasticity of consumption with respect to the net wealth change due to housing, where the net wealth calculation incorporates the value of defaulted debt as a gain. Column (1) replicates column (1) of Table 5 using our default wealth definition for comparison. Column (2) uses the wealth change definition inclusive of defaulted debt. The elasticity estimate increases slightly.

We then estimate the MPC coefficient using the new definition. Column (3) replicates column (1) of Table 7 for comparison. Column (4) estimates MPC using wealth change definition that is inclusive of defaulted debt. The estimated MPC hardly changes under the alternative definition.

Appendix Table 6
Net Worth Shock and Consumption Growth, 2006 to 2009
Robustness to Defaulted Debt

	(1)	(2)	(3)	(4)
Dependent variable:	Total spending growth, 2006 to 2009		Change in spending, 2006 to 2009 $000	
Housing NW shock, 2006-2009	0.634**	0.77**		
	(0.125)	(0.145)		
Change in home value, 06-09			0.054**	0.054**
			(0.009)	(0.0092)
Constant	-0.034*	-0.042**	-0.830	-0.828
	(0.015)	(0.014)	(0.536)	(0.536)
N	944	944	944	944
R^2	0.298	0.31	0.362	0.362

**,* Coefficient statistically different than zero at the 1% and 5% confidence level, respectively