HOUSEHOLD INEQUALITY AND THE CONSUMPTION RESPONSE TO AGGREGATE
REAL SHOCKS

Gene Amromin
Mariacristina De Nardi
Karl Schulze

Household Inequality and the Consumption Response to Aggregate Real Shocks
Gene Amromin, Mariacristina De Nardi, and Karl Schulze
NBER Working Paper No. 24073
November 2017
JEL No. D15,E21

ABSTRACT

To what extent does household inequality affect the response of aggregate consumption to aggregate real shocks? We first review two state-of-the-art papers with household heterogeneity and aggregate uncertainty. They teach us that having a larger fraction of poor and borrowing constrained households, who have a high marginal propensity to consume, amplifies the drop in aggregate consumption in response to a negative aggregate real shock. We then move on to the Panel Study of Income Dynamics (PSID) and Equifax data to quantify the fraction of people that are constrained in their consumption choices and to study how that fraction has changed before and after the Great Recession. We argue that the role of constraints cannot be adequately captured by only having a large share of households with no wealth before a recession. We find that, for all of the measures that we consider, the fraction of households that are borrowing constrained has drastically increased since the onset of the Great Recession and that it has remained high, or even increased, all the way through 2012, the last year for which we currently have PSID data. Thus, it is not surprising that aggregate consumption has experienced such a large drop and remained depressed for a long time.

Gene Amromin
Federal Reserve Bank of Chicago
230 South LaSalle Street
Chicago, IL 60604-1413
gamromin@frbchi.org

Mariacristina De Nardi
Federal Reserve Bank of Chicago
230 South LaSalle St.
Chicago, IL 60604
and University College London and Institute For Fiscal Studies - IFS
and also NBER
denardim@nber.org

Karl Schulze
Federal Reserve Bank of Chicago
kschulze@frbchi.org

Figure 1: U.S. real per-capita GDP, PCE, and their linear trends from 1985 to 2007. The values in year 2007 are normalized to 100 (Bureau of Economic Analysis, Census Bureau, Haver Analytics).

2007 = 100
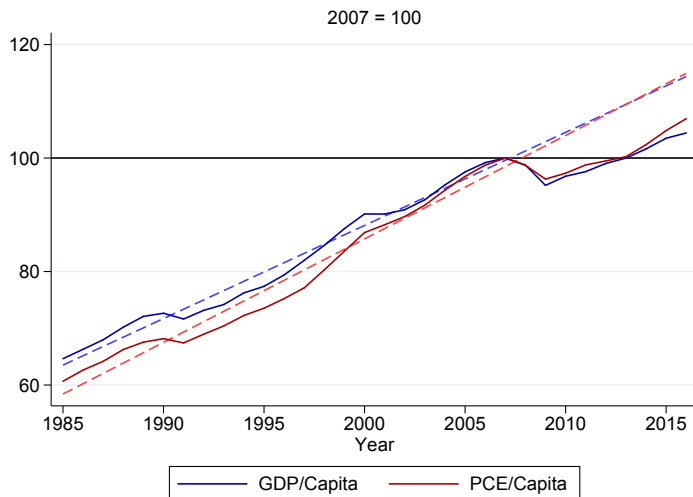
1 Introduction

The drop in output and consumption that occurred during the Great Recession has been large and prolonged. Figure 1 displays per-capita U.S. real gross domestic product (GDP) and personal consumption expenditures (PCE) between 1985 and 2016 and highlights the large drop in both consumption and output that occurred starting in 2007 and its parallel shift compared to the previous trend. In this paper, we ask why consumption has dropped so much and has been recovering so slowly. We also ask to what extent household inequality before and after the Great Recession interacted with the recession itself to generate such a large and persistent drop in consumption.

To reach our goal, we first review two key papers, Krusell and Smith (1998) and Krueger, Mitman, and Perri (2016), and summarize what we understand from quantitative macro models with aggregate uncertainty about wealth inequality, borrowing constrained households with heterogeneous marginal propensities to consume (MPCs), and the response of aggregate consumption to an aggregate shock. The key message from macroeconomic theory is that the extent to which households are borrowing constrained (which in these models means how many people have low wealth holdings) and how that changes both over time and with aggregate shocks are cru-

cial in explaining the aggregate economy's response and the speed of its recovery. That is, versions of these models with a realistic fraction of poor people can generate larger consumption drops in response to a total factor productivity (TFP) shock than models with fewer poor people. Hence, inequality does affect aggregate consumption dynamics in response to shocks in an important way.

While these papers provide very important lessons and are much more realistic than previous models, we argue that these newer models still abstract from important changes that happened during the Great Recession and are thus likely to understate the size (and potentially the persistence) of the consumption drop that occurred and the role of household inequality in generating it. First, they do not incorporate the large wealth losses (due to drops in house prices and in stock market values) that occurred during the Great Recession and have depressed aggregate consumption due to a negative wealth effect. Second, they take the view that there is only one asset and that borrowing constrained people are largely people with low amounts of this composite asset. However, the composition of a household's portfolio is important to determine the extent to which the household is borrowing constrained. For instance, in a model with liquid and illiquid assets that are subject to adjustment costs, people might have significant holdings of illiquid wealth and still be borrowing constrained (Kaplan and Violante, 2014). Thus, these models understate the fraction of households that are borrowing constrained and hence the size of the associated consumption drop. Third, these models assume that financial frictions are constant, including during a large downturn. In contrast, credit standards tightened considerably during the Great Recession, which also tends to reduce consumption (Guerrieri and Lorenzoni 2017). Fourth, these models assume that only unemployment risk changes over the cycle while there is significant evidence that earnings risk conditional on being employed is asymmetric: large earnings drops are much more likely than large earnings increases and this asymmetry becomes more pronounced during recessions (Guvenen, Ozkan, and Song 2014). If a household perceives higher earnings risk during a recession, it would likely cut consumption and increase savings, thus generating an even larger drop in aggregate consumption.

Thus, all four of these factors likely increase the size of the consumption drop that is related to household inequality because they increase the fraction of borrowing constrained households that have high marginal propensities to consume (MPC) and whose consumption reacts more strongly to a negative aggregate shock. Interestingly,

these factors do not apply symmetrically to all households. For instance, the wealth drop is more important for households in the top half of the wealth distribution, while credit market access plays a larger role for low-wealth households who lack alternative means of smoothing consumption. Accounting for these factors can thus improve the ability of the current models to match cross-sectional patterns of consumption responses.

One important implication of the four features that are missing in these macro benchmark models is that we should think more carefully how to best measure the fraction of households that are borrowing constrained and how this fraction changes over time and the business cycle. It is important to stress that the fraction of households that are borrowing constrained is not just the fraction of low wealth households, but also depends on their portfolio composition, exposure to earnings risk, and, potentially, on what non-discretionary expenses they expect (Campbell and Hercowitz 2016). We thus need a better measure of borrowing constrained households.

To help shed light on these important issues, we use the Panel Study of Income Dynamics (PSID) and credit bureau panel data (Equifax) to examine the interaction between consumption, wealth inequality, and borrowing constraints during the Great Recession. We find that, consistent with the models' implications, the poor and borrowing constrained households have a larger consumption expenditure rate as a fraction of their income and that their expenditure rate dropped by more during the Great Recession. We then use various measures of borrowing constraints to get a sense of how the fraction of constrained households has changed over time. We document that, for all of the measures that we consider, the fraction of households that are borrowing constrained has drastically increased since the onset of the Great Recession and has remained high, or even increased, all the way to 2012, the last year for which we currently have PSID data. Thus, it is not surprising that aggregate consumption has experienced such a large drop and remained depressed for a long time. Before turning to our data analysis, we now discuss what we know from macroeconomic theory in more detail.

4

2 Krusell and Smith "Income and Wealth Heterogeneity in the Macroeconomy," 1998

To better discuss the forces at play, let us remind ourselves of the crucial features of Krusell and Smith's (KS) model. The consumers are infinitely-lived, provide labor inelastically, and optimally choose how much to save and consume. They are ex-ante identical but get hit by an unemployment shock. Markets are incomplete: there is only one asset, aggregate capital, and consumers cannot borrow. The model induces endogenous wealth inequality: although households face the same stochastic process for an endowment shock, they receive different sequences of the unemployment shock and are thus ex-post heterogeneous in their wealth holdings.

There is an aggregate shock that affects total factor productivity of both capital and labor (TFP) in the aggregate production function and the aggregate unemployment rate. A recession in this framework is thus characterized by lower aggregate wages and higher unemployment. The equilibrium interest rate is the marginal product of capital. The equilibrium wage is the marginal product of labor. Due to the aggregate shocks, both the interest rate and the wage fluctuate over time. Thus, because of aggregate uncertainty, consumers have to forecast future prices (both the wage and the interest rate). In principle, the whole distribution of wealth, together with the aggregate shock, is a state variable to forecast future prices.

We learn two main lessons from this paper. The first one is that today's average capital, or net worth, is enough to forecast future prices for given aggregate shock and that we thus need to keep track of only one moment of the wealth distribution. The second main lesson from this paper is that the first result can still imply that the response of the aggregate time series depends on the distribution of wealth and the fraction of people that are constrained. In fact, the second moments of the aggregate time series in models with incomplete markets are different from those with a representative agent. Especially in models with more realistic wealth inequality, there is a larger correlation between consumption and income.

The key insight in these models with incomplete markets is that different saving (and consumption) propensities are associated with low levels of wealth. As wealth increases, the marginal propensity to save converges to one (permanent income behavior) and thus aggregates up across agents. Because the agents with lowest and most heterogeneous propensity to save (and highest to consume) are associated with low

5

wealth levels, they have very little effect on aggregate capital and prices, as opposed to the high-wealth agents whose policy functions aggregate. In contrast, because the poor have a high marginal propensity to consume, they account for a significant amount of aggregate consumption. Moreover, because they have little wealth to insulate their consumption from aggregate shocks affecting wages and income, they also have the largest consumption fluctuations over the cycle.

3   Krueger, Mitman, and Perri "Macroeconomics and Household Heterogeneity," 2016

Krueger, Mitman, and Perri (KMP) extend the KS's framework with preference heterogeneity to include a stylized life-cycle structure with constant probabilities of aging or dying in each sub-period (working period and retirement), labor productivity shocks conditional on employment, and unemployment insurance.

Their key finding is that, in these economies, the decline in aggregate consumption to an aggregate shock is larger in an economy populated by more wealth-poor households because their consumption responds more strongly to aggregate shocks, which are characterized by a reduction in TFP and increased unemployment risk.

Figure 2: Policy functions and wealth distributions, Krusell and Smith basic model (left panel) and Krueger, Mittman, and Perri (right panel). Graphs from KMP
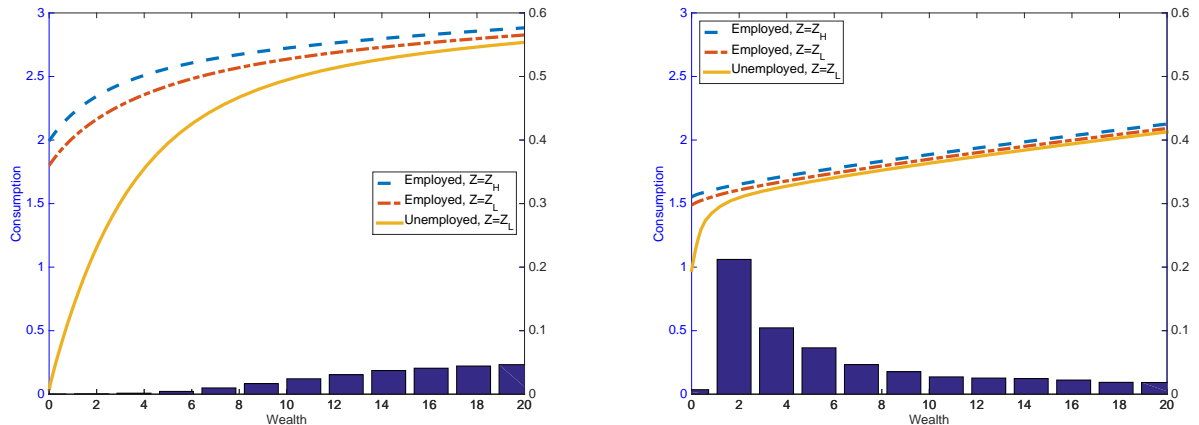
Figure 3: Consumption recessions in various versions of the model. Graph from KMP
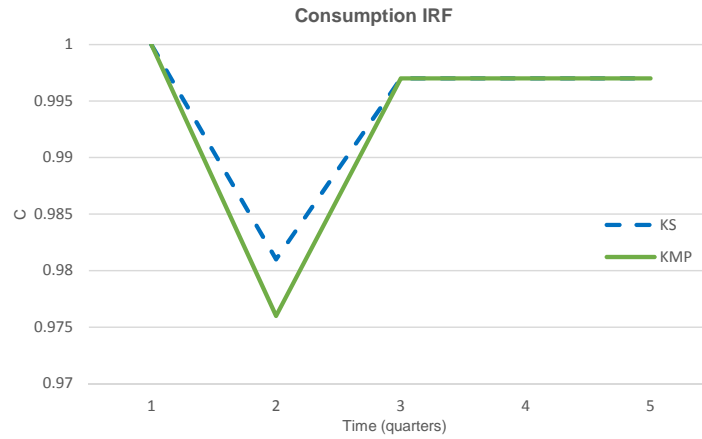

Figure 2 displays the consumption functions (plotted against individual wealth on the x-axis) and the pre-recession wealth distributions for the original KS economy (left hand side) and the KMP economy (both graphs are from KMP). The lines represent the consumption functions of the employed in an expansion, the employed in a recession, and the unemployed in a recession. Thus, for a given wealth level, the vertical difference between the consumption functions for the employed in the good aggregate state and the employed in the bad aggregate state gives the consumption drop in a large recession, conditional on the aggregate state switching but on not losing a job. In the same way, the vertical difference between the consumption function of the employed in an expansion and the consumption function of an unemployed in a recession shows the consumption drop for a household that experiences a recession and unemployment at the same time.

The figure reveals several interesting features. First, for a given level of wealth, the drop in consumption is larger in the original KS economy and especially so for households with little to no assets. Second, while there are almost no people in this situation in the original KS economy, there are many more households in those wealth quantiles in the KMP economy. It should be noted that average wealth is the same in the two economies. As it turns out, the KMP economy displays a larger consumption drop than the KS economy precisely because there are more people at low wealth holdings. Figure 3 (also from KMP) highlights this result.

4 Results from the models, a discussion leading to our data analysis

These two papers stress the importance of modeling heterogeneity and constrained households in understanding aggregate consumption and output dynamics. However, these models abstract from features of the Great Recession that are likely important in understanding the behavior of aggregate consumption during and after that recession. For example, while the additional heterogeneity in KMP deepens the initial consumption drop, it recovers at the same rate as the KS model. The richer dimensions of heterogeneity we document could potentially help future models match both the steep initial decline and slow recovery in consumption as observed in the Great Recession.

The features we consider include, first, changes in the value of wealth holdings (and thus the fraction of poor households) due to drops in house prices and financial asset valuation. Second, we study the implications of the observation that the asset structure is much richer than just one asset. A household's portfolio allocation and its liquidity can have crucial implications for who is constrained (the wealthy-hand-to-mouth). Third, changes in credit constraints that occurred, due to changes in credit conditions during the downturn, are likely important. Fourth, we examine changes in earnings risk for the employed in recessions and expansions. We know that the distribution of earnings displays more negative skewness in recessions, for instance. It is also possible that the persistence of earnings might depend on the status of the aggregate economy and one's earning level.

We now turn to discussing key aspects of the data and the strengths and weaknesses of the models in explaining empirical consumption dynamics after large shocks.

5 Data

5.1 Panel Study of Income Dynamics

We utilize the Panel Study of Income Dynamics (PSID) to analyze the relationship between household inequality and consumption dynamics. The PSID has two important advantages from the standpoint of the questions that we are asking. First, it provides information on earnings, income, consumption, and wealth for a sample

that is representative of the U.S. population. Second, it has a panel dimension that allows us to track households over time in addition to performing cross-sectional investigations. The latter allows us to check how wealth and consumption have changed for the same households.

We use data from the 2001 to 2013 PSID waves (covering the years 2000 to 2012), with particular focus on 2004, 2006, and 2010. Our baseline sample in these years consists of families with household heads aged 25 to 60. We also exclude a few households whose total annual consumption is implausibly small, less than $1000 in 2006 dollars. Our wealth quintiles are defined over this working-age sample. When analyzing changes over time, we make these restrictions only in the base year. Much of our analysis focuses on three variables: total income, consumption expenditures, and net worth. All dollar values are deflated to 2006 using the Consumer Price Index unless otherwise stated.

Total income includes all pre-tax sources of household income from all family members, including Social Security income and other transfers, as well as a family's income from assets and businesses. Transfer income includes unemployment compensation, transfers from other family members, and welfare benefits, among others. Total income also includes retirement and pension income. Net worth is composed of net housing equity, net value of other real estate, net value of farm or business, net value of vehicles, value of annuities and retirement accounts, value of non-retirement investments and savings, and the net value of other assets less the value of other liabilities. After 2010, the value of other liabilities is subdivided into credit card debt, student loans, medical bills, legal bills, and other family loan debt. In the analysis that considers illiquid housing wealth, we only look at the net equity in the primary home and other real estate. Net worth does not include the present value of defined-benefit pension income.

While the PSID harmonizes the variables for total family income and net worth, we must construct consumption expenditures using disparate expenditure categories. This is created from expenditures on transportation, vehicles purchased in the current year, food, clothing, education, child care, entertainment, recreation, vacation, utilities, rent, imputed rents for homeowners, property tax, household repairs, household durables, homeowner's insurance, out-of-pocket medical care, and health insurance premiums. As in KMP, the imputed rental value for homeowners is calculated by multiplying the total value of a household's home by an interest rate of 4%. Due to

9

the limited availability of variables, this definition of consumption is only available for 2004 onwards as expenditures on clothing, entertainment, recreation, vacation, household durables, and household repairs are not collected prior to this. All expenditure variables are adjusted to an annual frequency.[1]

5.2 Equifax

We also make use of the Equifax-FRBNY Consumer Credit Panel (CCP), which is a longitudinal dataset containing quarterly credit bureau records for a nationally representative 5% random sample of individual borrowers. The CCP database also includes records of all household members of the primary sample borrower. These data include information on all aspects of individual and household-level credit and debt, which includes credit cards, auto loans, student debt, as well as first- and second-lien mortgages. The data also include individual credit scores and the number of credit applications, which allow us to track household creditworthiness and demand for credit over time.

6 Important dimensions of heterogeneity in the PSID

We begin by discussing several key empirical facts from the PSID.[2] Table 1 reports the marginal distributions of earnings, income, consumption expenditures, and net worth in 2006, on the eve of the recession. All of these measures of household well-being are very unevenly distributed, but some more so than others.

Labor earnings are heavily concentrated at the top, with the top quintile accounting for 47% of all earnings and the top 5 percent of the distribution receiving 21%. In comparison, total income, which includes transfer income, is somewhat more evenly distributed. The distribution of consumption expenditures, while still skewed, is the least extreme – whereas households in the bottom two quintiles of labor earnings account for 11.9% of the total, the bottom two consumption quintiles make up 18.3% of

---

[1]We calculate similar results to KMP. See their paper for a comparison between aggregates in the PSID and National Income and Product Accounts.

[2]Tables 1, 2, and 4 are calculated by the authors from the PSID following the computations by KMP.

total consumption. It is the inequality in wealth holdings that is by far the starkest –
the top 5% of the distribution account for nearly half of the aggregate wealth, while
the bottom 40% of the population hold almost no net worth at all.

Table 1: Means and marginal distributions in 2006 (PSID).

	Labor earnings	Total income	Cons. exp.	Net worth
Mean (2006)	66425	80277	47526	324973
Q1	2.1	4.4	6.1	-1.5
Q2	9.8	10.9	12.2	1.2
Q3	16.2	16.3	17.2	6.0
Q4	24.9	23.4	23.4	14.7
Q5	47.1	45.0	41.0	79.6
90-95%	9.1	9.0	9.0	11.0
95-99%	13.8	13.6	11.3	23.4
Top 1%	7.2	7.1	5.8	26.3
Gini	0.45	0.46	0.39	0.76
Sample Size	6231	6231	6231	6231

For the moment, we will take net worth as a measure of household constraints,
since less wealthy households are likely less able to respond to shocks. How much
do wealth-poor households contribute to aggregate consumption? To answer this
question, Table 2 reports the distribution of several variables, conditional on wealth
quintile in 2006. It reveals that asset-poor households make up a sizeable measure of
aggregate consumption expenditures. In fact, the bottom two wealth quintiles account
for 30% of expenditures in 2006. Not surprisingly, households in these quintiles spend
much higher shares of their total income on consumption (expenditure rates).

Table 2: Shares and means by net worth quintile in 2006 (PSID).

	% Share of aggregate:				Mean (000s $2006):				
	Net worth	Labor earnings	Total income	Cons. exp.	Net worth	Labor earnings	Total income	Cons. exp.	Exp. rate (%)
Q1	-1.5	11.6	11.5	14.6	-15.5	29.4	35.0	26.8	76.8
Q2	1.2	14.2	14.0	15.3	13.3	37.8	44.9	29.6	66.0
Q3	6.0	20.3	19.5	19.6	71.1	57.4	66.5	40.2	60.5
Q4	14.7	20.8	20.6	20.9	213.2	73.1	86.8	53.2	61.2
Q5	79.6	33.1	34.4	29.7	1344.4	134.6	168.4	87.8	52.2
Corr. w/ NW	0.31	0.38	0.21						

Table 3 shows that the low-wealth quintiles are populated by households that are younger, less educated, have low propensities to own homes or financial assets, and are less likely to hold full-time employment throughout the year, as evidenced both by the number of hours worked and prevalence of unemployment spells. The data in these tables corroborate the essential ingredients of KMP: (1) a large fraction of households have no assets but account for a significant part of aggregate consumption, and (2) these households are vulnerable to shocks, as their low asset positions are compounded by lower levels of human capital, whether acquired through schooling or work experience.

Table 3: Household characteristics by net worth quintile in 2006. Education is expressed in years (PSID).

	Age	Educ.	Worked last year	Hours if Worked	Unemployed last year	Own home	Own fin. assets
Q1	39.2	12.7	83.4	1906.2	17.3	12.7	9.9
Q2	40.3	12.5	89.0	2107.1	10.7	36.4	14.3
Q3	42.3	13.2	92.3	2219.8	5.5	80.2	33.6
Q4	46.2	13.5	95.0	2222.4	3.1	87.7	54.9
Q5	48.8	14.8	94.7	2316.1	3.3	93.6	82.7

How did these household groups respond to the Great Recession? Table 4 reports annualized nominal changes in select variables before and during the Great Recession, conditional on initial wealth quintile. These changes are computed holding initial quintile composition fixed. For instance, to compute the change in consumption expenditures for households in the bottom wealth quintile between 2004 and 2006,

we take all households in that quintile in 2004 and compute their mean consumption expenditures in 2004 and in 2006. The annualized change in these means is reported in the table. Changes in expenditure rates are constructed similarly but are reported as percentage point differences.

The top row of Table 4 shows aggregate growth rates that tell a familiar tale of the Great Recession. After expanding at a pace of 14.1% during the boom years, aggregate net worth declined at an annual rate of 3.4% between 2006 and 2010. The growth rate of aggregate total income also slowed down from 2.6% to only 0.8% per year. However, the slowdown in aggregate consumption proved to be much faster: after increasing at an average rate of 5.5% per year, consumption fell at an average rate of 0.8% between 2006 and 2010. These relative movements in growth rates of income and consumption are reflected in absolute declines in the aggregate expenditure rate.

A convenient metric for measuring the effect of the Great Recession is the change in nominal growth rates between the 2004-06 and 2006-10 periods. The Great Recession depressed total income growth for all wealth groups but in a non-monotonic fashion. Households in the middle of the wealth distribution (quintiles 3 and 4) experienced the greatest slowdown, while for those at the lowest quintile total income growth decelerated much less (-0.6 percent per year). Consumption growth rates also decreased for all wealth groups, with the slowdown again being most severe in the middle of the wealth distribution. Low wealth households experienced a slowdown in consumption growth that was barely above those in the highest wealth quintile and well below those in quintiles 3 and 4. This finding is consistent with results in Meyer and Sullivan (2013), which documents a decline in consumption inequality during the Great Recession. They show that the fall in asset prices (both housing and financial) had a disproportionate effect on households with higher ex-ante consumption levels.

Importantly, slowdowns in consumption growth were larger than drops in income growth for all wealth groups. As a result, the share of income spent on consumption (expenditure share) declined for all wealth groups, but the poorest households experienced the largest decline. This empirical fact is consistent with the main intuition of the KMP model: when unemployment risk rises, low-wealth households cut back consumption because they fail to accumulate wealth that can be used to smooth consumption fluctuations. This happens whether they actually become unemployed or not. In the latter case, they cut consumption to build up their precautionary savings given the heightened unemployment risk during a recession.

Table 4: Annualized nominal changes in financial variables over 2004, 2006 net worth quintiles. Changes in expenditure rates are expressed as percentage point changes. Changes in net worth, income, and consumption expenditures are expressed as percentage changes (PSID).

	Net worth			Total income			Cons. exp.			Exp. rate		
	04-06	06-10	Δ	04-06	06-10	Δ	04-06	06-10	Δ	04-06	06-10	Δ
All	14.1	-3.4	-17.6	2.6	0.8	-1.8	5.5	-0.8	-6.4	1.6	-0.9	-2.5
Q1	.	.	.	6.0	5.4	-0.6	6.4	0.8	-5.5	0.2	-3.1	-3.3
Q2	69.1	16.1	-53.0	4.5	2.5	-2.0	5.9	1.8	-4.2	0.9	-0.4	-1.3
Q3	26.9	3.9	-22.9	4.9	1.6	-3.3	9.0	0.4	-8.6	2.3	-0.7	-3.0
Q4	15.4	1.1	-14.3	4.6	1.7	-2.9	6.0	-1.2	-7.2	0.8	-1.7	-2.5
Q5	11.0	-5.6	-16.6	-0.4	-1.9	-1.4	3.2	-2.9	-6.1	1.7	-0.6	-2.3

7 Matching heterogeneous patterns in the data

As shown in figure 3, the KMP model delivers a larger response to an aggregate recession shock than a model with a smaller wealth dispersion. It thus matches important aspects of the data and shows that an economy with borrowing constrained agents can generate larger drops in aggregate consumption. As shown in Table 5, the KMP model does well in matching heterogeneity in declines in disposable income. It also captures the qualitative ordering of recession-driven changes in expenditure rates, as low-wealth households react more to the recessionary shock than high-wealth households.

Table 5: Difference in annualized growth rates between recession period and normal times (PSID and KMP model results).

	Net worth (%)		Disp. Y (%)		Cons. exp. (%)		Exp. rate (pp)	
	Data	Model	Data[3]	Model	Data	Model	Data	Model
Q1	.	-20	-0.7	-2.3	-5.5	-2.2	-3.3	0.0
Q2	-53	-18	-2.6	-2.8	-4.2	-2.4	-1.3	0.3
Q3	-23	-12	-3.3	-4.0	-8.6	-2.7	-3.0	1.4
Q4	-14	-5	-3.3	-4.5	-7.2	-2.8	-2.5	2.0
Q5	-17	-4	-3.0	-5.4	-6.1	-2.9	-2.3	3.2
All	-17.6	-4	-2.9	-4.4	-6.4	-2.6	-2.5	1.3

---

[3]KMP calculate disposable, post-tax income and use this in their model. Since our analysis uses only pre-tax income, we report both their data and model results for this field. All other data results report our own computations.

However, the Great Recession was associated with large wealth losses, tightening of borrowing constraints, and increased earnings uncertainty. These forces would take the model even closer to the data. They could further reduce aggregate consumption growth and amplify the consumption response relative to income drops. They could further contribute to a slow recovery in consumption. Because wealth losses fall unevenly across the distribution, the implied reductions in consumption growth rates would be heterogeneous.

8 Additional elements of the Great Recession

8.1 Shocks to wealth

Unlike the baseline model, the Great Recession was characterized by substantial declines in household wealth. Figure 4 demonstrates the leftward shift in distribution of household wealth between the 2006 and the 2010 PSID surveys. The share of households with negative net worth jumped from 16% to 24%, and the mean household net worth plummeted from $324,973 to $197,780 in 2006 dollars. In the context of the KMP model, this drop in household net worth could reinforce the negative consumption dynamics by altering the wealth distribution in Figure 2. Put differently, a recession that brings about a sizeable leftward shift in the wealth distribution elicits an additional consumption response, especially among households who bear the brunt of this shock.

Figure 4: Cross-sectional distributions of net worth in 2006 and 2010 (PSID).
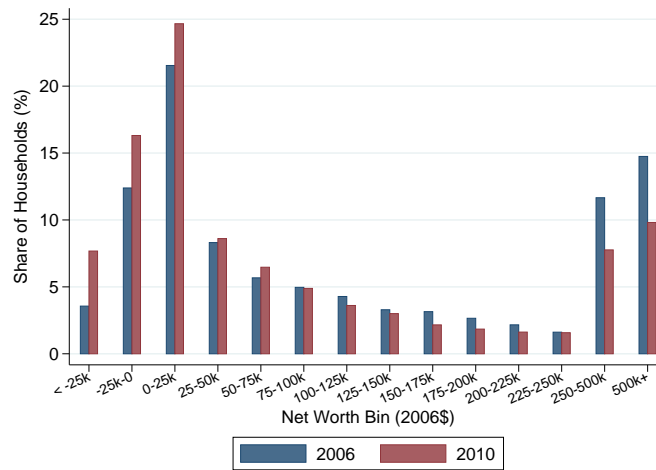
We can get some additional insight into these consumption dynamics by exploiting the panel structure of the PSID and linking changes in wealth with changes in consumption at the household level. Table 6 presents the shift in the distribution of household wealth between 2006 and 2010 as transitions across fixed net worth quintiles, defined by their 2006 threshold values. Each row of the table shows the distribution of households in a given wealth quintile in 2006 across wealth groupings in 2010. For example, 43.8% of households in the middle quintile of 2006 net worth distribution remained in the same group by 2010. However, more than 1/3 of households in this quintile (12.5%+22.9%) lost enough wealth by 2010 to drop into the bottom two quintiles. The table also captures the consumption growth rates for each of these groups, located in the brackets below the transition probabilities. Looking at households that moved from the middle quintile in 2006 to the bottom quintiles in 2010 shows that the deterioration in their wealth position was accompanied by massive reductions in consumption of about 3% per year. Table 6 shows that, in general, downward movements in wealth (the lower diagonal part of the transition matrix) are associated with negative consumption growth rates, shown in red. The lower subpanel of the table again underscores the shift in the wealth distribution: the bottom quintile in 2006 (< $2,500 in net worth) by construction accounted for 20% of households, but 23.7% of households found themselves with less than $2,500 in net worth in 2010.

Table 6: Transitions and annualized nominal consumption changes (%) across 2006 net-worth quintiles, 2006-2010 (PSID).

Limits:	2006Q	2010					Total:
		Q1	Q2	Q3	Q4	Q5	
≤ $2.5k	Q1	64.6	24.6	8.1	2.1	0.5	
		[-0.5]	[2.5]	[3.9]	[5.5]	[4.6]	[0.9]
$2.5-37.4k	Q2	30.2	46.2	18.3	4.6	0.7	
		[1.1]	[1.3]	[2.0]	[9.0]	[5.4]	[1.8]
$37.4-133k	Q3	12.5	22.9	43.8	16.9	4.0	
		[-2.8]	[-3.3]	[0.9]	[5.1]	[3.1]	[0.4]
$133-371k	Q4	7.6	8.2	26.6	42.9	14.7	
		[-2.9]	[-4.8]	[-3.7]	[-0.1]	[1.9]	[-1.2]
> $371k	Q5	3.3	1.4	7.2	20.0	68.2	
		[-26.8]	[-12.2]	[-5.9]	[-1.2]	[-1.3]	[-2.9]

|		2010					
	Total	23.7	20.6	20.8	17.3	17.6	
		[-3.3]	[-0.6]	[-0.8]	[1.0]	[-0.7]	

8.2   Shocks to composition of wealth

Housing equity represents the single largest wealth component for the majority of households, although it is illiquid. During the run up to the Great Recession, financial innovations made it much easier for households to use housing equity to support consumption, either by extracting it directly or by borrowing against it (Bhutta and Keys 2016). Case, Quigley, and Shiller (2013) estimate propensities to consume out of different types of wealth by analyzing changes in per capita consumption expenditures, income, financial wealth, and housing wealth in a quarterly panel of U.S. states spanning 1975 to 2012. They find much higher elasticities of consumption with respect to housing compared with financial wealth and attribute this result to changes in tax laws and the above-mentioned financial innovations. Thus, a drop in house values could further contribute to the drop in consumption.

Figure 5 again exploits the panel structure of PSID to show the distribution of individual changes in real housing wealth between 2006 and 2010, expressed in annualized, real percentage terms. The smaller figures show the unconditional distribution in each 2006 wealth quintile, while the top figure shows the overall distribution, which is the sum of the smaller histograms. The overall distribution of housing wealth changes

(top panel) was strongly asymmetric as many households experienced losses in real terms. On average, housing wealth declined by an average of 6.4% per year over this period and as many as 17.8% of PSID homeowners lost all of their housing wealth, whether through foreclosure (the -100% group) or by owing more than their house was worth (the "underwater" or <-100% group). The rest of the panels show the unconditional distributions of changes for each of the 2006 wealth groups. In these panels, non-homeowners and those with little housing wealth in 2006 are omitted and thus the sum over all the bins in each panel will roughly equal the homeownership rate for that quintile. Consequently, the histograms for the bottom two wealth quintiles suggest that housing losses there were limited because of low ownership rates. However, the rest of the wealth distribution was hit hard. In the middle quintile, complete loss of all housing equity represented the likeliest outcome. Although foreclosures were not nearly as common in the top wealth deciles, the share of households facing losses was even greater.

Figure 5: Distribution of household-level annualized, real percentage changes: illiquid net worth (PSID).
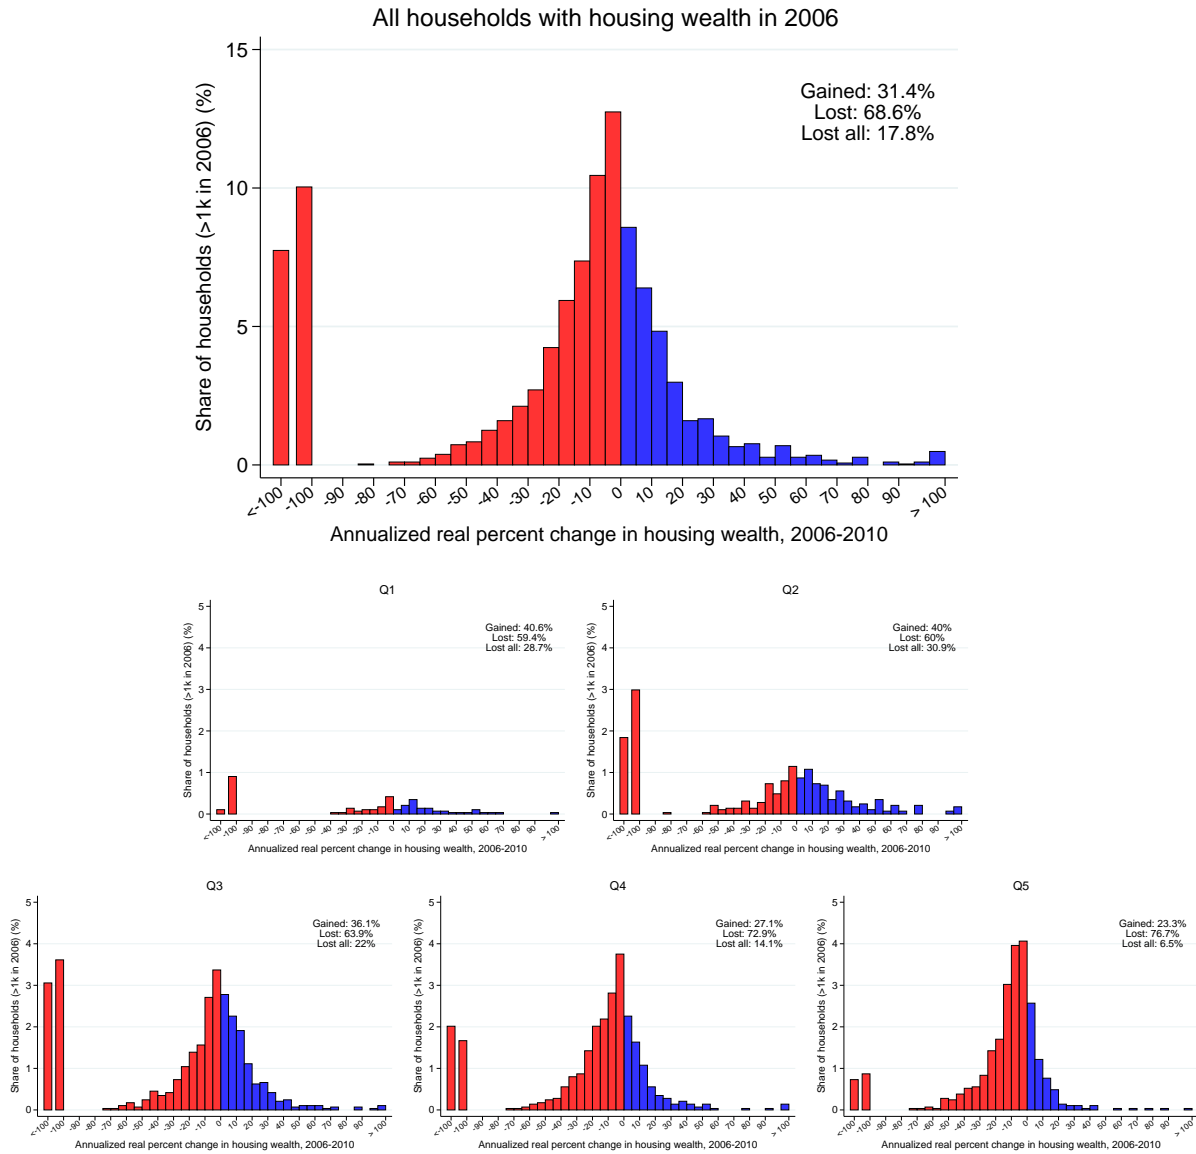


8.3 Changes in credit accessibility

The figures above describe erosion in housing wealth, which is commonly used as collateral. However, in addition to experiencing negative wealth shocks, households were also faced with a more adverse credit environment with the onset of the Great Recession. Lenders tightened their credit standards for home mortgages and other credit products. Moreover, as delinquency rates rose across credit markets, household

FICO scores deteriorated, which further contributed to difficulties in accessing credit.

To get some sense of magnitudes of the shock to credit access, we turn to the Equifax panel data. Figure 6 shows the distribution of FICO scores in a representative sample of households in 2006 and for the same sample in 2010. Unlike household wealth, the distribution of credit scores did not shift uniformly to the left. Rather, by 2010, the FICO score distribution became much thinner in the middle range of creditworthiness (690-790) and more concentrated at the very top. On net, the fraction of households that remained above the informal mid-point of the prime credit score range (FICO = 720) was nearly unchanged.

Instead, it was the tightening of minimum credit standards, shown by the dashed vertical lines, that pushed a large fraction of households out of credit markets. As shown in Figure 7, a FICO score of 720 roughly corresponded to the median score of a borrower approved for a home mortgage by Fannie Mae or Freddie Mac prior to the Great Recession. Over this time period, households with credit scores between 680 and 780 – the "prime" group in Figure 8 – experienced a success rate above 80% in getting their credit applications approved. Applying a FICO score of 720 as a notional cutoff for ready credit access to the distribution in Figure 6 places 47% of households in this category. However, by 2010, the same criteria for ready credit access – obtaining mortgage credit from the GSEs or having undiminished success rates for credit applications – required a credit score of about 780. This shift in credit requirements, shown by the red dashed line, shrank the fraction of households with easy credit access to just under 30%.

Figure 6: Changes in FICO scores and credit cutoffs, 2006-2010 (Equifax Credit Panel).
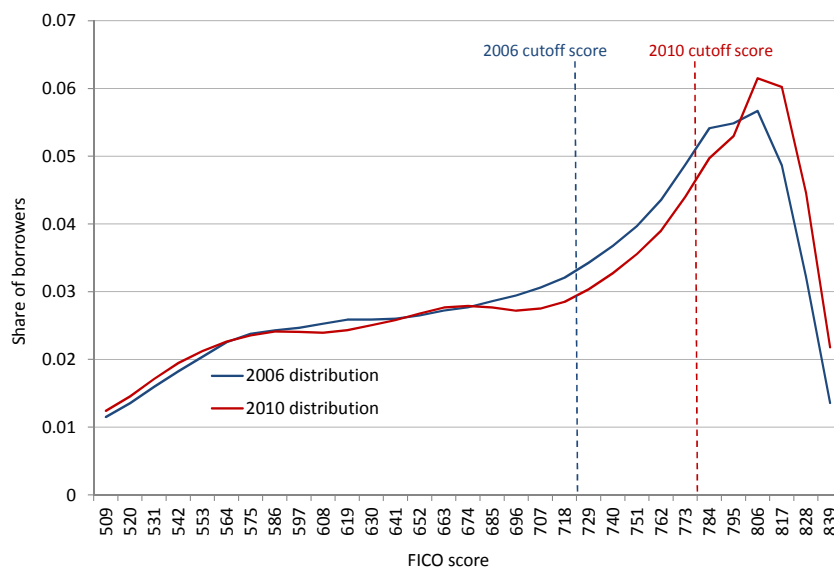
Evidence of credit tightening during the Great Recession can be seen in the Equifax data on median FICO scores for newly originated mortgages depicted in Figure 7. Stricter credit requirements are apparent both for mortgages backed by Fannie Mae and Freddie Mac as well as for mortgages for first-time homebuyers explicitly guaranteed by the U.S. Government through the Federal Housing Agency (FHA) or the Veterans Administration (VA). The same phenomenon is also illustrated by a sizable dip in success rates of credit applications in Figure 8, which includes not only mortgages but also auto loans and credit cards. This decline is especially pronounced for the lowest FICO score borrowers, who tend to be young and/or lower-income. Altogether, this negative shock to the accessibility of credit likely contributed to the observed patterns of substantial slowdowns in consumption growth and expenditure rates, especially at the bottom of the wealth distribution.

Figure 7: Median FICO score at mortgage origination (Equifax Credit Panel).
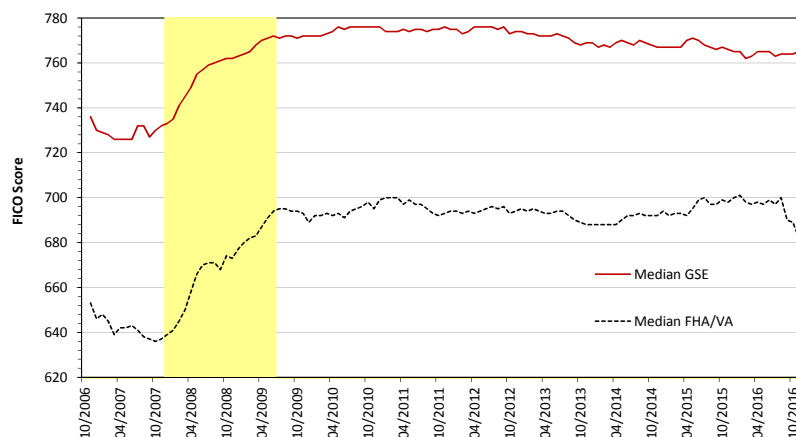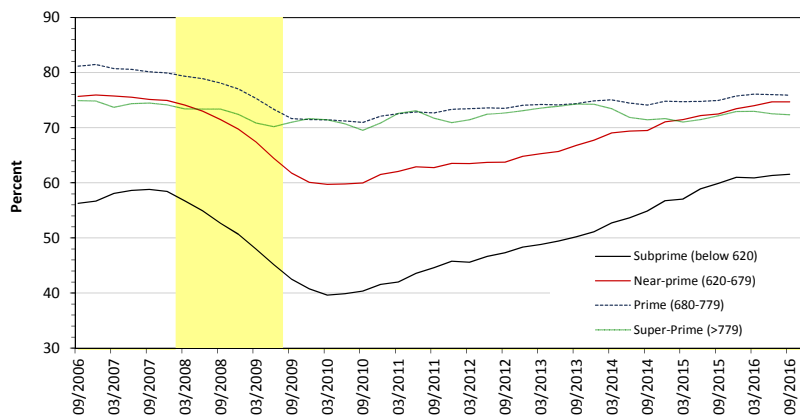

Figure 8: Success rate in obtaining credit (Equifax Credit Panel).

The persistence of various measures of credit tightening mirrors changes in household constraints, as shown in Figure 9 for several metrics. The top-left panel presents absolute measures of household wealth below a certain threshold. The top-right panel depicts measures of household wealth relative to their earnings or income. The bottom panel captures the share of households whose FICO scores are below the threshold commonly employed by mainstream financial institutions. All of these measures paint the same picture and tell us that the share of constrained households increased dramatically with the onset of the Great Recession and has showed no signs of decline through 2012. Since current models do not explain the slow recovery of consumption

22

with simply a large fraction of poor households, the persistence of these constraints means that this channel could hold particular promise in bringing existing models in line with realistic consumption dynamics.

Figure 9: Fraction of constrained households over time (PSID, Equifax Credit Panel).
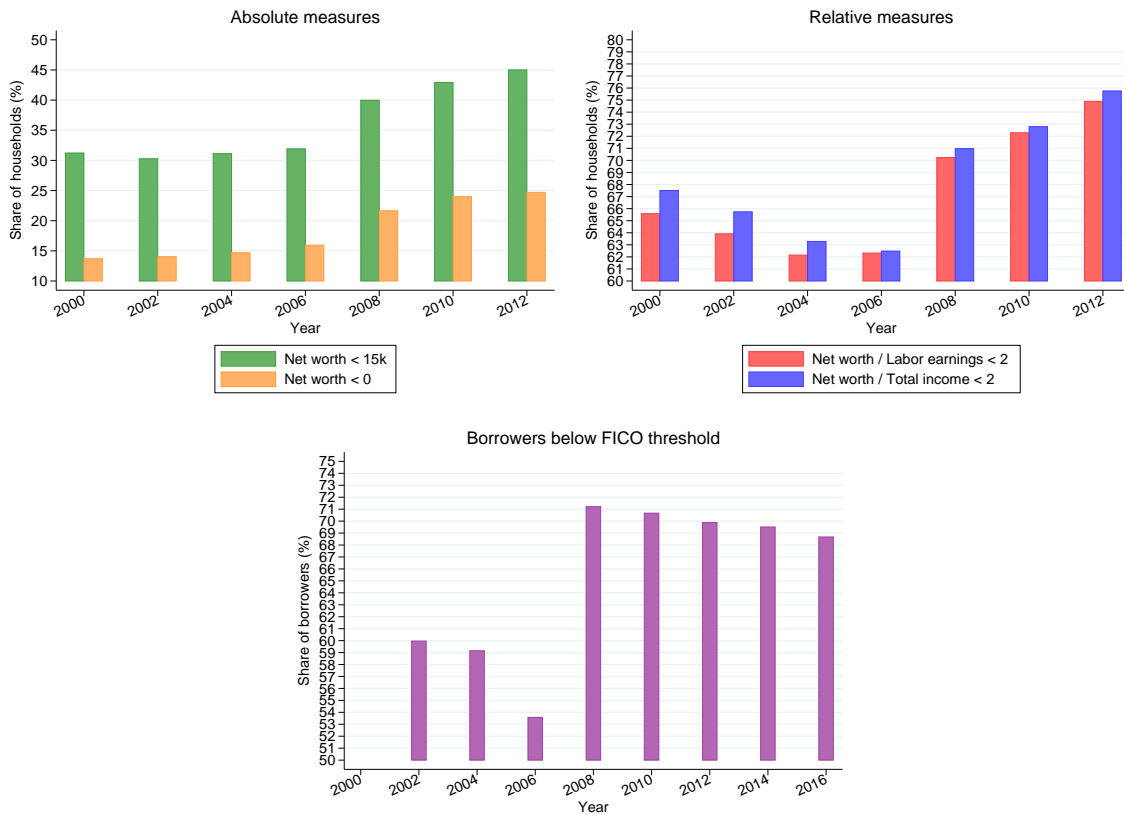
8.4 Changes in earnings risk

Table 7 discusses changes in some statistics of the labor market by net worth quintiles. First, it shows that unemployment increased during the Great Recession and especially so for those in the bottom wealth quintile. Second, it highlights that there were also large changes in hours and the fraction of people working and that both are U-shaped in wealth quintiles. In the baseline models, unemployment is the only labor market dynamic that changes (exogenously) during the cycle and does so independently of one's wealth level. Thus, the labor market dynamics of the model abstract from important features of labor supply.

Table 7: Annualized changes in employment variables over 2004, 2006 net worth quintiles. Changes in working status are expressed as percentage point changes. Hours conditional on working are expressed as percentage changes (PSID).

	Worked last year			Hours if worked			Unemployed last year			Head earnings if worked		
	04-06	06-10	Δ	04-06	06-10	Δ	04-06	06-10	Δ	04-06	06-10	Δ
All	-0.8	-1.4	-0.7	0.0	-1.7	-1.7	-0.7	1.1	1.8	4.5	1.1	-3.4
Q1	-0.9	-1.9	-0.9	1.7	-0.4	-2.1	-3.3	0.5	3.8	7.2	5.8	-1.3
Q2	-1.4	-1.3	0.1	0.4	-2.0	-2.4	-0.2	1.2	1.4	4.4	2.7	-1.8
Q3	-0.7	-1.1	-0.4	-0.7	-1.5	-0.9	0.7	1.0	0.4	4.2	1.5	-2.6
Q4	-0.4	-1.3	-0.8	-0.3	-1.9	-1.6	0.5	2.0	1.5	5.3	1.5	-3.8
Q5	-0.5	-1.8	-1.3	-0.7	-2.5	-1.9	-1.2	0.9	2.1	3.8	-1.3	-5.0

Other studies also found that earnings conditional on employment exhibit different dynamics over the business cycle. For instance, Guvenen, Ozkan, and Song (2014) find that the left-skewness of earnings shocks is strongly countercyclical: during recessions, large upward earnings movements become less likely whereas large drops in earnings become more likely. Both KS and KMP do not account for these important dynamics. The fact that negative earnings shocks become more likely in a recession is another force reducing consumption and increasing savings and might help bring the consumption response to aggregate shocks in the model more in line with the data.

This asymmetric increase in negative earnings risk during downturns holds promise in explaining the persistently languid consumption growth after the Recession, particularly among the poorest households, in addition to the initially steep decline. For example, Pistaferri (2016) runs regressions on changes in wealth and disposable income using pre-2008 aggregate data to predict consumption responses and then ex-

trapolates these for the post-2008 period. He finds that consumption has recovered significantly slower than would be expected based on the historical data (a finding corroborated in our Figure 1), but that accounting for household leverage and consumer confidence appears to explain the entire gap between observed and extrapolated trends. While the influence of the deleveraging process on reduced consumption has probably softened, he finds that factors influencing consumer confidence have not. In particular, pessimism is strongest amongst the lowest quartile of the income distribution who report worse expectations and increased uncertainty concerning financial conditions. These households also exhibit a permanent rise in the expected probability of job loss, not unlike the long-lasting aversion to financial risk documented among households who experienced the macroeconomic shock of the Great Depression (Malmendier and Nagel 2011). This indicates that the Recession could be seen as being characterized by both permanent negative shocks to income and persistent increases in income uncertainty (De Nardi, French, and Benson 2012). The prevalence of this pessimism amongst the poorest households underscores the importance of accounting for earnings dynamics in models of household heterogeneity.

9   Conclusion

This paper highlights a number of theoretical and empirical reasons established by earlier literature for the importance of household heterogeneity in understanding macroeconomic responses to shocks. In particular, we focus on identifying households who are constrained in their consumption choices. These households typically have a high marginal propensity to consume. Their prevalence in the economy and the nature of their constraints have a sizable impact on consumption dynamics.

While incorporating a realistic degree of wealth heterogeneity is crucial for generating plausible consumption responses, we argue that the role of constraints cannot be adequately captured by only having a large share of households with no wealth before a recession. This is particularly true in the case of slowdowns like the Great Recession, which delivered a strong negative shock not just to earnings, but also to wealth itself. We show that accounting for shocks to different types of wealth may also help explain consumption responses in the middle of the wealth distribution. Also, it is important to take into account supply-side changes in the availability of credit, which may further amplify the magnitude of the consumption response, especially among younger and less wealthy households. Finally, changes in the nature of earnings risks during booms and expansions should be taken into account.

References

[1] Bhutta, Neil and Benjamin J. Keys. 2016. "Interest Rates and Equity Extraction during the Housing Boom." *American Economic Review* 106(7), pages 1742-74.

[2] Campbell, Jeff and Zvi Hercowitz. 2017. "Liquidity Constraints of the Middle Class." *Federal Reserve Bank of Chicago Working Paper 2009-20.*

[3] Case, Karl E., John M. Quigley, and Robert J. Shiller, 2013. "Wealth Effects Revisited 1975-2012." *Critical Finance Review* 2(1), pages 101-128.

[4] De Nardi, Mariacristina, Eric French, and David Benson, 2012. "Consumption and the Great Recession." *Federal Reserve Bank of Chicago Economic Perspectives* 36(1), pages 1-16.

[5] Guerrieri, Veronica, and Guido Lorenzoni. 2017. "Credit Crises, Precautionary Savings, and the Liquidity Trap." *The Quarterly Journal of Economics* 132(3), pages 1427-1467.

[6] Guvenen, Fatih, Serdar Ozkan, and Jae Song. 2014. "The nature of countercyclical income risk." *Journal of Political Economy* 122(3), pages 621-660.

[7] Kaplan, Greg and Gianluca Violante. 2014. "A Model of the Consumption Response to Fiscal Stimulus Payments." *Econometrica* 82(4), pages 1199-1239.

[8] Krusell, Per and Tony Smith. 1998. "A Model of the Consumption Response to Fiscal Stimulus Payments." *Journal of Political Economy* 106(5), pages 867-896.

[9] Krueger, Dirk, Kurt Mitman, and Fabrizio Perri. 2016. "Macroeconomics and Household heterogeneity." *Handbook of Macroeconomics,* forthcoming.

[10] Malmendier, Ulrike and Stefan Nagel. 2011. "Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking?" *Quarterly Journal of Economics* 126(1), pages 373-416.

[11] Meyer, Bruce D., and James X. Sullivan. 2013. "Consumption and Income Inequality and the Great Recession." *American Economic Review, Papers and Proceedings*, 103(3), pages 178-183.

[12] Panel Study of Income Dynamics, public use dataset. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI 2017.

[13] Pistaferri, Luigi. 2016. "Why Has Consumption Remained Moderate after the Great Recession." *Working Paper.*