CPB Netherlands Bureau for
Economic Policy Analysis

Disentangling the effect of household debt on consumption

We estimate the contemporaneous relationship between household debt and consumption during the crisis. We find that average consumption of households with high debt has decreased much more than that of other households. We disentangle this into an effect through the availability of credit for direct consumption and through household debt overhang. On the micro level, the decline is sharpest for the households who are less able to finance one-off high consumption with new debts.

On the macro level the drop in consumption of households who have negative home equity has the biggest impact on macro consumption, because their number sharply increased during the crisis. Precautionary savings motives appear to contribute most to the decline.

CPB Discussion Paper

Kan Ji, Rutger Teulings, Bram Wouterse

April 2019

Disentangling the effect of household debt on consumption

Kan Ji[a] & Rutger Teulings[a] & Bram Wouterse[a,b,*]

[a]CPB Netherlands Bureau for Economic Policy Analysis, The Hague, The Netherlands

[b]Erasmus School of Health Policy & Management, Erasmus University Rotterdam, Rotterdam, The Netherlands

February 4, 2019

Abstract

We estimate the contemporaneous relationship between household debt and consumption for the period 2006 to 2015. Using Dutch administrative data, we find that the average consumption of households with high debt has decreased much more during the crisis than that of other households. We disentangle this into an effect through the availability of credit for direct consumption and an effect through household debt overhang. On the micro level, the consumption drop is the sharpest for the households who are less able or willing to finance one-off high consumption with new debts after the crisis. On the macro level, however, the drop in consumption of households who have negative home equity for a longer period had a much bigger impact on macro consumption, because their number sharply increased during the crisis. Our results suggest that precautionary savings motives among the highly indebted households contributed most to the consumption decline during the crisis.

Key words: Consumption, precautionary savings, spending normalization, financial crisis, administrative data.

JEL classification: D12, D14, E21.

*Contact details. Address: Centraal Planbureau, P.O. Box 80510, 2508 GM The Hague, The Netherlands; e-mail: k.ji@cpb.nl, r.m.teulings@cpb.nl, b.wouterse@cpb.nl.

We would like to thank Nicoleta Ciurila, Albert van der Horst, Clemens Kool, Rob Luginbuhl, Mauro Mastrogiacomo, Remco Mocking, Bert Smid, Benedikt Vogt, Karen van der Wiel and Lu Zhang for their valuable comments, improving our paper.

1 Introduction

The build-up of high levels of household debt before the financial crisis is often seen as an important explanation why the crisis has had such severe consequences for household consumption. Dutch households' debt-to-income (DTI) ratios over the period 2006 to 2015 were on average 2.7, ranking second among all OECD countries (OECD data). Figure 1 shows that consumption decreased by 7 percentage points from the start of the financial crisis in 2008 to the end of the crisis in 2014, based on national account data. This suggests that the large drop in consumption of Dutch households might indeed be related to their relatively high level of debt. However, micro evidence for this link is still lacking. The main contribution of this paper is to fill this gap.

Figure 1: The change in consumption with respect to 2007 (index)

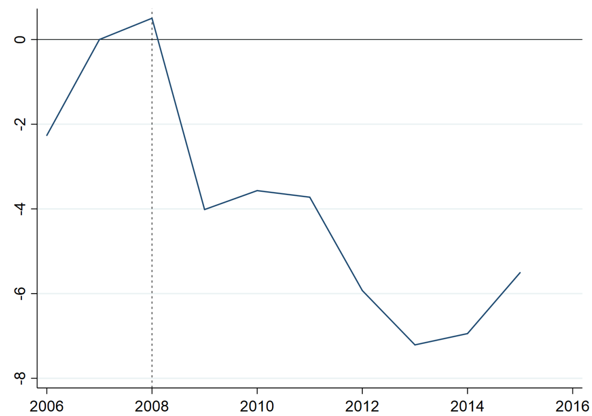

Based on national accounts data by Statistics Netherlands (CBS).

We estimate the contemporaneous relationship between household debt, measured as DTI, and household consumption for the period 2006 to 2015. The analysis is based on a large random sample from an administrative data set containing income, financial assets, debt and other household characteristics, for all Dutch households.

We disentangle two channels through which the financial crisis might have had an effect on the relation between debt and consumption. First, households with high debt overhang might have responded more strongly to financial and macroeconomic shocks than other households, either because they ran into credit constraints or out of precaution. In our paper we call this the stock effect of debt. Second, the crisis might have restricted the availability of credit for households to finance high, one-off, consumption. We call this the flow effect of debt. We analyze these two channels by estimating the consumption of different subgroups within households with high debt, for which either the flow or the stock effect can be expected to be more important[1].

On the micro level, we find that the strongest decline in individual consumption is related to the flow effect of debt. Before the crisis, households, who move to a new house, seem to use part of their debt to consume, as they consume substantially more than other households. After the crisis, the difference in consumption between movers and other households is substantially smaller.

On the macro level, however, the stock effect of debt seems to dominate. The number of highly indebted home owners with negative home equity increases substantially during the crisis. Although the drop in consumption of this group on a micro level is smaller than the consumption reduction of movers, the large size of this group leads to a bigger impact on the macro level. Our results suggest that increased precautionary savings among already highly indebted households played an important part in the consumption decline during the crisis.

The remainder of the paper is structured as follows. Section 2 places the paper in the literature. Next, we describe the data and construction of the consumption proxy in Section 3 and in Section 4 we discuss the estimation method and counterfactual analysis. In Section 5 we describe the results and in Section 6 we perform several robustness checks. Finally, Section 7 concludes.

2 Literature

When households do not face any borrowing constraints, standard life cycle theory (Ando and Modigliani, 1963) finds that household consumption is only driven by the level of a household's net wealth and lifetime income; the level of debt should not play a separate role. However, in practice households do face frictions and constraints, and these might have become stricter during the financial crisis.

There are two potential channels through which the crisis could have affected the relationship between debt and consumption. First, high levels of debt might have affected the willingness or ability of households to consume. Second, the crisis might have decreased the amount of credit available for, one-off, household consumption. We call the first channel the 'stock effect' and the second the 'flow effect'. We borrow these terms from the literature looking at the relationship between the level and flow of debt and economic growth at the macro level (Bezemer et al., 2016; Biggs et al., 2009). There, it is used to distinguish between the positive short-term effect of increase in the flow of credit, which

[1] Our study focuses on the debt side of the household balance sheet. See Zhang (2019) for an analysis of the role of house prices (the asset side) on consumption

leads to more liquidity and transactions, on growth, and, the possibly negative, long-term effect of an increase in the stock of debt.

Most of the theoretical literature on household consumption and debt seems to focus on the impact of debt overhang (the stock effect). For these households, consumption can be affected through a balance sheet effect: during a crisis, highly indebted households are suddenly faced with binding constraints, forcing them to decrease their consumption substantially (Eggertsson and Krugman, 2012). Cooper (2013) shows that this channel is an important driver in explaining consumption patterns for homeowners in the US. A second reason why highly indebted households might reduce their consumption is precaution: a crisis might increase (perceived) uncertainty, and this induces risk averse households to increase savings, even if they are not (yet) faced with binding credit constraints or actual shocks in their income or wealth (Carroll et al., 2012; Challe and Ragot, 2016; Guerrieri and Lorenzoni, 2017). As the potential consequences of these risks are worse for highly indebted households, the precautionary motive can be expected to be strongest for this group.

Also the empirical research has mostly focused on the consumption of households with high debt overhang. It has been well established for the U.S. (Dynan et al., 2012; Baker, 2018) and the U.K. (Bunn and Rostom, 2015) that households with severe debt before the financial crisis reduced their consumption more strongly than other households during and after the financial crisis. These studies tend to focus on the overall effect of household debt on consumption. In contrast, Mian et al. (2013) try to identify the household balance sheet effect by estimating the effects of drops in house prices during the crisis in the U.S., using the variation in the housing supply elasticity across counties as an instrument. They find that falling housing prices negatively impact consumption, and this impact is especially strong for highly leveraged households.

The empirical evidence on the role of household debt on consumption for European countries is more mixed. Van Beers et al. (2015) and Bijlsma and Mocking (2017), who exploit variation in house prices within regions in the Netherlands, find much smaller effects of housing wealth on consumption than Mian et al. (2013), and they do not find that highly leveraged household consistently react more strongly than households with lower levels of debt. In a study for Estonia, Kukk (2016) finds that household with high debt, and especially a high debt service ratio, reduced their consumption more during the crisis than others.

The impact of household debt on consumption during the financial crisis through the availability of credit for one-off high consumption (the flow-effect) has been less investigated. Households might use debt in year t to finance one-off expensive consumption, so called durable consumption (e.g. furniture, cars), in year t. Thus we can expect that the contemporaneous relationship between (new) debt and consumption is positive (Andersen et al., 2016) . We want to investigate in this paper, whether the effect of the financial crisis on the relationship between debt and consumption might also go through this channel: the crisis might have led to a decreasing ability of households to use debt for consumption, which, next to the stock effect of debt, may have contributed to the overall decrease in household consumption.

To our knowledge, Andersen et al. (2016) are the only ones who explicitly investigate the flow effect of household debt using microdata. They use it to show that it is not sufficient evidence for a household balance sheet effect to just show that households, with high debt before the crisis, decreased their consumption more than other households during and after the crisis. Because when households use debt to finance one-off high consumption in one year, consumption in the following years will automatically be lower as these households return to their 'normal' level of consumption. Using administrative data from Denmark, they indeed find that the relation between high debt before the crisis and decreases in consumption in the following years is driven by this 'normalization' effect. What they do not do, is to explicitly investigate whether the crisis had an effect through the flow effect itself: did the crisis lead to a smaller contemporaneous effect of debt on consumption? In this paper, we do investigate this channel.

3 Data

3.1 Construction of the sample

We use administrative micro-data for the Dutch population, collected by Statistics Netherlands (CBS). To keep our analysis computationally manageable, we randomly select 10% of the Dutch households in the year 2010. Within these selected households, we randomly select one of its household members. Next, we construct a panel by following these individuals over the time period 2005-2015 and include all households to which these individuals belong during the sample period. Households of which the household head is younger than 21 are excluded. We do include households of whom the main earner is self-employed, although this is not common in the literature[2].

[2]We include self-employed because the share in total employment increased rapidly during our sample period up to 17%. This is even the strongest increase among all OECD countries. One of the reasons is that employers converted their employees in self-employed to allow for more flexibility during the crisis. In a robustness check, we do exclude households of whom the main earner is self-employed. The main results do not change.

3.2 Data on income, wealth and household characteristics

We have data on income and wealth from the Dutch tax authority, and municipality administration data on personal and household characteristics. We link these data and construct a data set comprising of income, wealth and household characteristics.

Disposable household income is the income definition we use. This consists of all sources of income from production factors and fiscal transfers summed over all household members. Taxes on income and wealth, unemployment and health insurance premiums, and paid income transfers to other households (e.g. alimony) are subtracted. Statistics Netherlands winsorizes the resulting income measure at one million euros.

Wealth data consists of ten different categories: the balance on all (savings) accounts, bonds, stocks, substantial shareholding of a single firm, assets of privately owned firms, value of the house, value of other properties, other assets, mortgage and other debt. We remove any observations in which a household has non-zero substantial shareholding of a single firm (this is only 1% of all observations), because the accuracy of measurement changes during the sample period.[3] We exclude the wealth category "other debt" in the construction of the consumption measure and DTI, as this category is not consistently measured over time.[4] We do not omit the corresponding observations because this would result in a loss of 13% of the observations, with relatively more observations in later years.[5] We have no reliable data on household's accumulated pension wealth.

We construct DTI ratios using the real mortgage[6] of each household and divide it by its average real disposable household income, where we use the consumer price index from Statistics Netherlands to deflate both series. We drop all observations for which the DTI ratio is above 10.[7]

Finally, we use the following household's characteristics: the age of the main earner,

[3] Our main conclusions do not change if we do include these observations.

[4] In the Dutch tax system one only pays taxes over the part of wealth that is above 24 thousand euro. Debts (excluding mortgages) are deductible from this amount. So people only have an incentive to report debts when they have wealth of more than 24 thousand euros. Therefore, debt is hardly reported in the data, especially small amounts. In later years, the data collection method improves and there is a gradual increase in the number of reported debt.

[5] If we remove all observations for which other debt is non-zero we find similar results. Similarly, our main conclusions do not change, if we include the category other debt into the construction of the consumption measure and in our DTI ratios.

[6] Mortgages can be artificially high for some households, because of the so called "savings mortgages". In these types of mortgages households maximize their mortgage rate deduction, a tax scheme, by saving their amortization on a separate savings account instead of reducing their mortgage directly. Unfortunately there is no way to identify households with these mortgages in the data.

[7] We have also experimented with a cut-off of 5. The results are qualitatively the same, but the impact of high DTI is reduced substantially since we remove 16% of all observations all of whom have high debt. So the impact of high debt becomes smaller.

the household-type (single, couple, etc.), number of children, whether a household rents or owns a house and the household's address. The latter is used to determine whether a household moves in a specific year or not.

3.3 Construction of the consumption measure

We construct a consumption measure based on income and wealth, using an accounting identity: household's consumption equals household's disposable income minus its savings. However, we do not observe household savings, instead we observe household wealth. So, we proxy savings by the change in the wealth stock and, following Browning and Leth-Petersen (2003), Leth-Petersen (2010) and Browning et al. (2013), use it to construct consumption for household i at time t:

$$C_{i,t} = Y_{i,t} - \Delta W_{i,t}, \tag{1}$$

where $C_{i,t}$ is household consumption, $Y_{i,t}$ is disposable household income and $\Delta W_{i,t}$ is the change in household wealth. This measure captures both durable and non-durable consumption.

This proxy has one main disadvantage. The change in household wealth does not only capture savings but also capital gains or losses due to price changes, the so called price effect. To limit the sensitivity to this kind of measurement error, it is common in the literature to exclude households that own wealth categories that are typically the most susceptible to the price effect, such as stocks and bonds, and households that move to a different address (e.g. Andersen et al. (2016)). However, this does not only result in the loss of a substantial amount of observations (around 20%), but it also creates a selection bias. In fact, the excluded households, especially the movers, are the ones where we might expect the flow effect of debt on consumption to be most relevant. So, we do not exclude households that move or own stocks and bonds, instead we correct for the price effect as much as possible to mitigate the measurement error.[8] In Appendix A.2 we describe for each wealth category how we adjust for the price effect.

3.4 Quality of the consumption measure

Similar to Kreiner et al. (2015) and Bound and Krueger (1991), we compare the consumption proxy measure with self-reported survey data. We use the Dutch 2015 budget

[8] As robustness check we remove observations in which households have stocks and bonds. This leads to results comparable to the baseline, even though we remove a substantial share (19%) of the total sample. Removing all observations in which households move, also does not influence the baseline results. However, we are no longer able to analyze the impact of high DTI on consumption for movers.

survey of Statistics Netherlands[9]. A representative sample of 11,035 households is asked to report all their spending during a 4-week period in a diary. The respondent also fills in a questionnaire about all large expenses and monthly fixed costs in the previous year.

Figure 2 shows the relationship between the budget survey and the administrative measure (both measured as a share of household's average disposable income). If both consumption measures capture true consumption, all points would be on the 45-degree line, and the slope of a regression line would be 1. However, both the administrative and survey measures are known to suffer from different types of measurement error. The red line displays the best local polynomial fit. Between 0.3 and 0.8 the fit is reasonable with a slope coefficient between 0.40 and 0.57, but in the tails the measures deviate substantially from each other.[10] So both measures are less comparable than in for example Kreiner et al. (2015). As both measures have errors, we cannot solely attribute this deviation to measurement error in the administrative data (as in Kreiner et al. (2015)). Still, there is a clear positive relationship between both measures. In Appendix B we also compare the performance for different subgroups, showing that there is no difference in the quality of the consumption proxy across groups. Thus, we do not expect our group-specific estimates to be biased because of quality differences.

3.5 Descriptive statistics

We clean the consumption measure, as described in Appendix A.2.1, resulting in an unbalanced panel of 6,102,480 observations ranging from 2006-2015 with 715,804 unique households, where 60% of the households can be followed over the full ten years and 76% over at least nine years.

In Table 1 we display the mean and median for the low and high DTI subgroup for the three most important series: DTI, real disposable income and real consumption scaled by real average disposable income. Here the high DTI subgroup refers to the households with a DTI ratio in the top quantile of distribution, and the low DTI subgroup contains all the other households.

These descriptive statistics shows that the high DTI group has a considerable higher disposable income, but also spends more of it on consumption and the pattern is fluctuating more strongly over time. The low DTI group has a considerably lower DTI ratio, on average around 0.71 compared to 5.61 in the high DTI group. In Appendix C we display the distributional plots of both DTI and loan-to-value (LTV) ratios, also indicating the

[9]The survey is also performed in other years for different cross-sections. However, for these years it is not possible to link the survey results to the wealth and income data.

[10]If we use a linear fit we find a slope coefficient of 0.48.

Figure 2: Budget survey versus administrative proxy

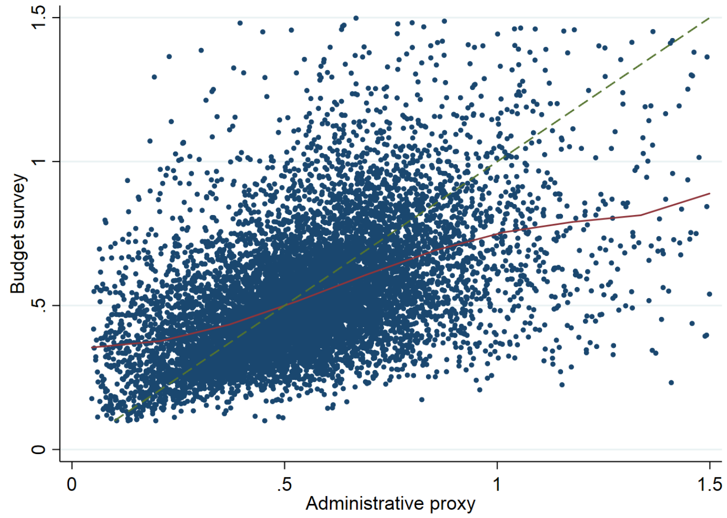

This figure reports real consumption over average real disposable income based on survey data (y-axis) and the administrative proxy (x-axis). The red line is the best local polynomial and the green dashed line is the 45 degree line.

location of the top quantile of the distribution.

We further break down the low and high DTI household groups into subgroups in Figures 3a and 3b. This reveals two relevant differences. First, in the high DTI group the percentage of households that move declines considerably over time while it stays relatively constant in the low DTI group. This suggests that, during and after the crisis, the number of movers declined because fewer households were willing or able to finance a house using debt. This might also have affected the use of (mortgage) debt to finance consumption (furniture, renovations). Second, the composition of the high DTI group is changing over time into younger households with children who have less financial wealth and more often have negative home equity. These compositional effects might influence the relationship between debt and consumption of households, both on a micro and macro level.

4 Methods

We now introduce our econometric specifications to analyze the relationship between household debt and consumption for the period 2006-2015. We first explain the baseline analysis. Then, we turn to an alternative specification that helps us to identify whether our patterns are driven by a flow effect. Last, we explain a counter-factual analysis to analyze the effect of high debt on consumption at the macro level.

Table 1: Descriptive statistics for high and low DTI

Subgroup	High DTI		Low DTI	
Statistic	Mean	Median	Mean	Median
DTI	5.68	5.42	0.71	0.00
2006	5.67	5.37	0.76	0.00
2007	5.72	5.44	0.74	0.00
2008	5.73	5.46	0.72	0.00
2009	5.75	5.49	0.70	0.00
2010	5.75	5.51	0.67	0.00
2011	5.73	5.48	0.68	0.00
2012	5.68	5.44	0.70	0.00
2013	5.60	5.36	0.72	0.00
2014	5.58	5.33	0.73	0.00
2015	5.61	5.35	0.70	0.00
Real Disp. income	44	41	35	29
2006	44	40	37	32
2007	45	41	37	32
2008	45	41	37	31
2009	45	41	36	30
2010	44	40	34	28
2011	44	40	34	28
2012	43	40	34	28
2013	43	39	33	27
2014	44	41	34	27
2015	49	46	37	30
Cons./Avg. disp. inc.	1.10	1.01	0.99	0.96
2006	1.16	1.02	0.99	0.95
2007	1.17	1.03	1.02	0.98
2008	1.18	1.04	1.06	1.00
2009	1.08	0.99	0.99	0.97
2010	1.10	1.01	1.03	0.99
2011	1.04	0.99	0.96	0.96
2012	1.06	0.99	0.98	0.96
2013	1.05	0.98	1.00	0.94
2014	1.01	0.97	0.94	0.92
2015	1.15	1.09	0.98	0.95
N	1549598		4552882	

Debt to income (DTI) is defined as real mortgage debt over average real disposable income. Real disposable income is measured in thousand 2015 euros. Real consumption is scaled by average real disposable income.

Figure 3: Relative size of the subgroups

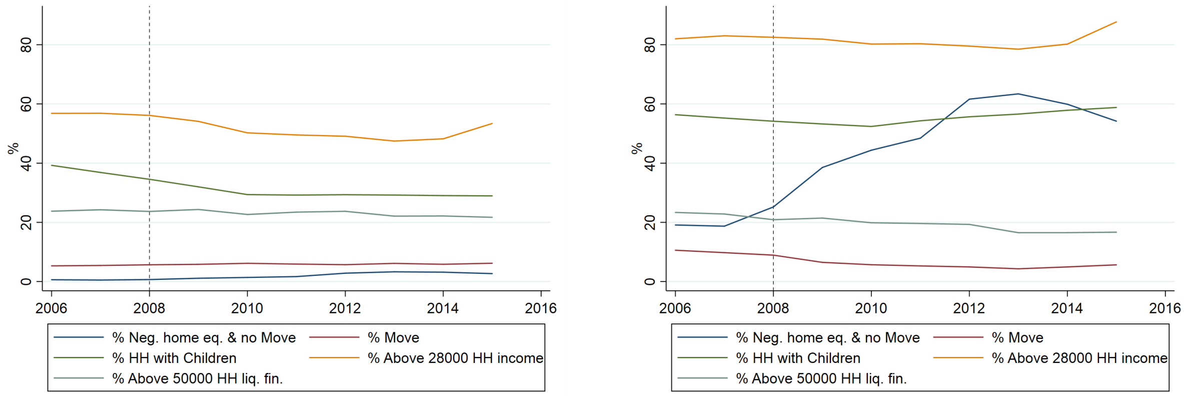

(a) Within the low DTI group

(b) Within the high DTI group

Five subgroups include: 1) households with negative equity who do not move, 2) households who move, 3) households with children, 4) households with net household income above 28,000 euro, and 5) households with liquid financial assets above 50,000 euro. The vertical dashed line indicates the start of the crisis.

4.1 Baseline analyses

We examine the relationship between household debt and consumption in year t using the following pooled regression:

$$\begin{aligned} \frac{c_{i,t}}{\bar{y}_i} =& \alpha + \beta_{1t} DTI^{high}_{i,t} + Sub'_{i,t}\beta_{2t} + Sub'_{i,t} \cdot DTI^{high}_{i,t}\beta_3 \\ & + X'_{i,t}\gamma_q + \theta_t + \varepsilon_{i,t}, \end{aligned} \tag{2}$$

where the dependent variable is the level of real household spending $c_{i,t}$ for household i in year t scaled by household's real disposable income averaged over time, $\bar{y}_i$. The latter is to ensure comparability across households with different income levels, following Andersen et al. (2016).[11]

The key explanatory variable is the dummy: $DTI^{high}_{i,t}$. It takes the value 1 if household i belongs to the group of highly indebted households (i.e. the top quantile of the DTI distribution) in year t. The difference in spending between households with high debt and households with low debt is captured by the coefficient β_{1t}.[12] Because this difference may evolve over time, we let this coefficient vary by year.

The estimate of the $DTI^{high}_{i,t}$ effect on consumption, β_{1t}, captures both the stock- and

[11] Andersen et al. (2016) scale spending by gross income in a fixed year (2007). However, household gross income in years 2006 to 2015 is volatile, so scaling relative to income in a single year can produce misleading results. Results are similar if we scale by average gross income.

[12] As robustness check we also remove all observations in which households have a DTI between the bottom and top 25th percentile of the distribution. So, we only compare the households with the highest DTI versus the lowest DTI. This does not change the result.

flow-effect.[13] The sign and development of the overall effect over time already tell us something about the importance of both channels. For instance, in the absence of a flow-effect, we would not expect high debt to have a positive effect on consumption in the pre-crisis years. More importantly, we estimate separate effects for different subgroups for which we can expect either the flow- or stock-effect to be relevant. For example, households that move to a different address are most likely susceptible to the flow effect, while for households with negative home equity the stock effect is more relevant.

The vector $Sub_{i,t}$ includes the following sub-groups: $Move_{i,t}$, $OwnHouse_{i,t}$, $NegHomeEq_{i,t}$ and $MoveToDTIh_{i,t}$. The variable $Move_{i,t}$ equals 1 if household i moves in year t. Next, $OwnHouse_{i,t}$ equals 1 if household i owns its house in year t. Furthermore, $NegHomeEq_{i,t}$ equals 1 if the equity value of the house in year t is negative, that is, when the total outstanding mortgage is larger than the value of the house. Finally, $MoveToDTIh_{i,t}$ equals 1 if household i moves from the low into the high DTI quantile in year t. We let the impact of these four subgroups, indicated by β_{2t}, vary over time. Furthermore, we are also interested in the additional impact of these subgroups when households have high debt, therefore we include an interaction term with $DTI_{i,t}^{high}$.

We control for other household characteristics in the vector $X_{i,t}$. We include household i's real total assets and real financial wealth, where the latter is defined as the sum of the savings accounts, stocks, bonds and other assets. We scale both variables by household's average real disposable income. The household's real disposable income is also included as a separate variable. We further control for the age of the oldest household member and the number of children. To allow for potential nonlinearities, we use an approach similar to Andersen et al. (2016) and let the impact of the control variables γ_q vary over the four quantiles (q) of the distribution of each control variable. So we employ a flexible functional form.

Finally, we include time fixed effects (FE) θ_t to capture common variations over time, that is, macro-economic fluctuations on the baseline group, the lowly indebted households. But they also capture possible changes in measurement of the income and wealth data over time (see Section 3). We assume that these measurement errors affect our consumption measure equally for lowly and highly indebted households and thus does not affect the estimates of β_{1t}, allowing us to identify the effect of high debt. [14]

We estimate the model using OLS with clustered standard errors at the household

[13] The endogeneity (debt at t affects consumption at t, but also the other way round), is thus on purpose. As a robustness check, we also ran a version of Equation (2) using lagged values of $DTI_{i,t}^{high}$. The results from this regression, that should only pick up the stock-effect, confirm that the effect of DTI we estimate in Equation (2) is indeed the composite of the stock- and flow-effect.

[14] We also considered household FE. However, including them did not change the results except for a shift in the means of different subgroups.

level. Furthermore, we perform several robustness checks. The most notable are reported in Section 6: the cut-off value of highly indebted households increases to 10% of the sample, and DTI is replaced by the LTV ratio.

Our main focus is on the question whether declining consumption is driven by households being less able to use debt to finance one-off consumption (the flow effect), or by increased savings by the already high-indebted households (the stock effect). In both cases households might increase savings or reduce borrowing because they actually run into binding constraints, or out of precaution. We will try to distinguish between these two, to some extent, by performing sub-group analysis, where we rerun our regression for households with low and high liquid wealth separately.

4.2 A different look at the flow effect

As an additional way to identify whether the flow–effect of debt plays an important role, we use a different model also used by Andersen et al. (2016). If households use high debt to finance one-off high consumption (durables), we would expect their consumption in the next years to return to their normal level. To test this, we estimate a variant of (2), in which we replace the dependent variable, the level of spending at year t, by the change in spending between year t and $t+2$:

$$\begin{aligned}\Delta\left(\frac{c_{i,t+2}}{\bar{y}_i}\right) =& \alpha + \phi_{1t} DTI^{high}_{i,t} + Sub'_{i,t}\phi_{2t} + Sub'_{i,t}\cdot DTI^{high}_{i,t}\phi_3 \\ &+ X'_{i,t}\gamma_q + \theta_t + \varepsilon_{i,t}.\end{aligned} \tag{3}$$

Together with the results from the estimates for the levels, this regression can show whether flow effects are indeed important for explaining spending differences between high and low DTI. In that case, we would expect that consumption of households with high debt in year t is higher than that of other households (and thus β_{1t} should be positive), but high debt households decrease their consumption back to 'normal' levels in the following years (and thus ϕ_{1t} should be negative). This is the normalization pattern as identified by Andersen et al. (2016).

The differentiation between subgroups again also helps to distinguish between explanations: if the crisis has led to restricted possibilities for households to use debt to finance one-off consumption, we would expect a decrease in the contemporaneous positive effect of debt in t on consumption in t (so a smaller β_{1t}), and a smaller adjustment in consumption during the next years (so a smaller ϕ_{1t}). For the credit constraints or precautionary savings explanation, we would expect consumption in t of households with longstanding debt in t to be lower than that of other households during the crisis years and we would not expect an (upward) adjustment of their consumption during the following 2 years.

4.3 Counter-factual analyses

The estimates from the pooled regression provide insights in the relationship between household debt and consumption on the micro level; they show how individual households within particular groups behave. We are also interested in how these effects translate into average consumption over income for the whole population (the macro consumption over income). To get some insight in the latter, we follow Bunn and Rostom (2015) and perform a counter-factual analysis. We take 2007 as base year and consider the average (macro) consumption over income in the years 2006 to 2015 that would have materialized, if debt would have had the same influence on spending in each year as it did in 2007. We use 2007 as base year because it is the last year before the financial crisis.

To do so, we make two predictions. First, we predict consumption over income for each household in each year using the estimated year-specific coefficients from (2). Calculating a weighted average over all households by year gives the actually realized macro consumption over income, where we weight by average disposable income, $\bar{y}_i$. Second, we predict consumption over income for each household in each year, setting the year-interaction effects, for high DTI and for the subgroups, at their 2007 levels. Calculating the weighted average by year over this second set of predictions, gives the counter-factual macro consumption over income. Subtracting the first prediction from the second gives the loss in consumption over income due to the changing effect of high debt on consumption.

5 Results

5.1 Spending of households with high DTI

We first compare the spending of households with a high DTI ratio in year t to that of households with a low DTI ratio in year t. For this purpose, we estimate a restricted baseline version of (2), where we restrict the effects of the subgroups to be constant over time ($\beta_{2t} = \beta_2$), and exclude interaction effects between subgroups and the high DTI ratio ($\beta_3 = 0$). The regression results can be found in Model 1 of Table 2. Figure 4a shows the average predicted consumption (as a share of average net income) for households in both groups. We have taken out the yearly deviation from the average level effect of the year dummies[15], as we do in all predicted series in this paper, but include the interaction effects between year and high DTI.

[15]The remaining variation in the average prediction of the low DTI group is due to changes in the covariate values across years.

Table 2: Regression results for (2) and (3) with and without subgroup interactions

Model		1	2	3	4
Dep. variable		$\frac{c_{i,t}}{\bar{y}_i}$	$\Delta\left(\frac{c_{i,t+2}}{\bar{y}_i}\right)$	$\frac{c_{i,t}}{\bar{y}_i}$	$\Delta\left(\frac{c_{i,t+2}}{\bar{y}_i}\right)$
$\beta_{2t} = \beta_2$ & $\beta_3 = 0$		Yes		No	
$DTI^{high}_{i,t}$	2006	0.110*	−0.070*	0.119*	−1.479*
		(0.002)	(0.008)	(0.013)	(0.066)
	2007	0.111*	−0.066*	0.094*	−1.337*
		(0.002)	(0.006)	(0.013)	(0.065)
	2008	0.100*	−0.021*	0.055*	−1.245*
		(0.002)	(0.007)	(0.013)	(0.066)
	2009	0.070*	0.013*	0.051*	−1.246*
		(0.001)	(0.005)	(0.013)	(0.065)
	2010	0.050*	0.036*	0.017	−1.281*
		(0.002)	(0.006)	(0.013)	(0.066)
	2011	0.060*	−0.015*	0.062*	−1.407*
		(0.001)	(0.005)	(0.013)	(0.065)
	2012	0.034*	0.049*	0.035*	−1.333*
		(0.002)	(0.006)	(0.013)	(0.066)
	2013	0.007*	0.149*	−0.003	−1.29 *
		(0.002)	(0.006)	(0.013)	(0.066)
	2014	0.027*		0.030*	
		(0.001)		(0.013)	
	2015	0.121*		0.081*	
		(0.002)		(0.013)	
$\cdot Move_{i,t}$				−0.214*	0.629*
				(0.004)	(0.018)
$\cdot NegHomeEq_{i,t}$				−0.015*	0.098*
				(0.002)	(0.008)
$\cdot OwnHouse_{i,t}$				0.024	1.315*
				(0.013)	(0.065)
$\cdot MoveToDTIh_{i,t}$				0.750*	−2.291*
				(0.009)	(0.046)
R^2		0.082	0.043	0.086	0.046
N		5,928,890	4,171,964	5,928,890	4,171,964

Clustered standard errors on a household level are between brackets. * indicates significance at the 5% level. This level is used throughout the paper.

Figure 4: Average predicted consumption for households with low and high DTI in the restricted baseline model using two different dependent variables

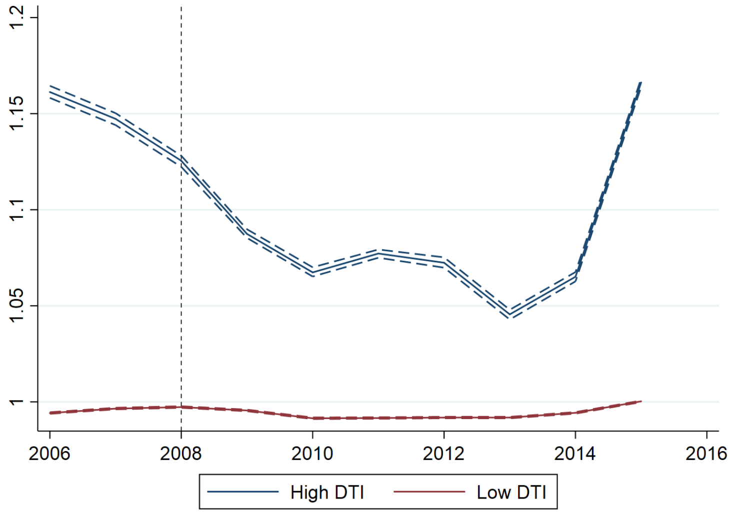

(a) Level of consumption: $\widehat{\frac{c_{i,t}}{\bar{y}_i}}$

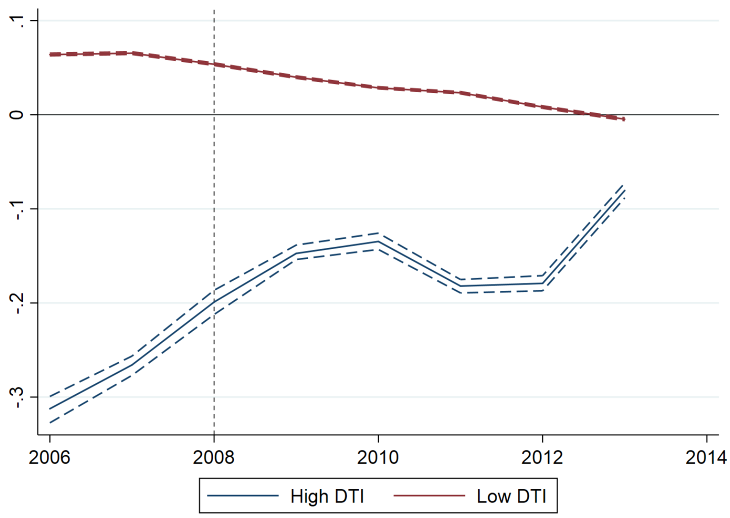

(b) Difference of consumption: $\Delta\widehat{\left(\frac{c_{i,t+2}}{\bar{y}_i}\right)}$

From all predicted consumption measures we subtract the yearly deviation from the average level effect of the year dummies. The vertical dashed line indicates the start of the crisis. We report 95% bootstrapped confidence bands based on re-sampling the observations and using 200 repetitions.

The average predicted consumption of households with high debt has indeed decreased more strongly during the crisis years (2008-2013) than that of household with low debt. However, in all years average predicted consumption remains higher for the high debt group than for the low debt group. Figure 4a shows that in 2006, households with high DTI consume 17 percentage points of their income more than households with low debt. In 2013, this is only 5 percentage points. In 2015, the difference in consumption increases again to 17 percentage points. We find similar patterns if we look at the estimated coefficients in Model 1 of Table 2 or in Figure 5, which shows the marginal effects of having high DTI over the years.

Figure 5: Marginal effects of $DTI_{i,t}^{high}$ in the restricted baseline model

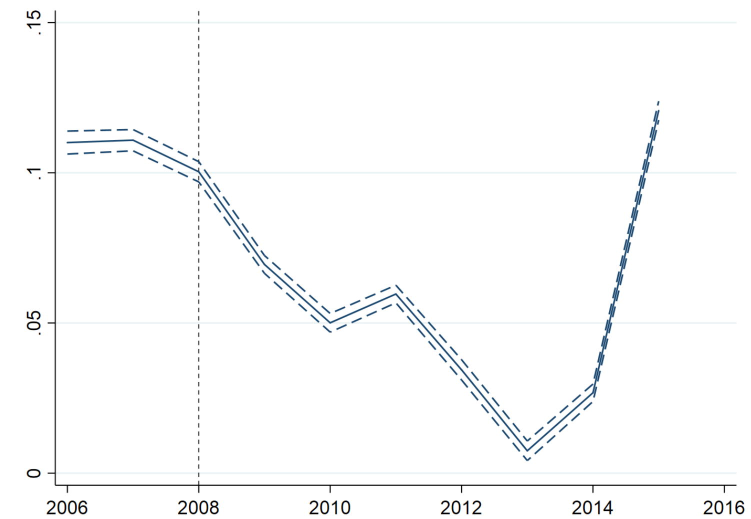

This figure displays the estimated marginal effects $\hat{\beta}_{1t}$. The 95% confidence bands are based on clustered standard errors.

The fact that consumption of households is higher for households with high debt across the whole time period lends support to the importance of the flow effect: households seem to use high debt to finance one-off high consumption. If this is indeed the case, then we would expect that the consumption in the consecutive years after year t would decrease for households with high debt in year t. This is indeed what Figure 4b shows. We plot the average predicted change in consumption from year t to $t + 2$ for households with high and low debt in year t based on estimating (3). Households with low DTI in year t tend to consume slightly more in year $t + 2$ than in t (except for 2013). Households with high DTI in year t always have substantially lower consumption in $t + 2$ than in t. In addition, the negative difference between consumption in $t + 2$ and t becomes smaller over time. This indicates that the flow effect becomes less important over time; as high debt in t is less strongly related to high consumption in t, there is also a less strong decline of consumption in subsequent years. This finding is supported by the estimated coefficients in Model 2 of Table 2, which show a similar pattern for the high DTI group.

Although the results for the high DTI group as a whole support the hypothesis that the flow effect is an important driver of the spending pattern for households with high DTI, this does not necessarily mean that the stock effect is not important as well. In the next section we therefore zoom in on different subgroups of households with high DTI.

5.2 Heterogeneity across groups of households with high DTI

We estimate the model from Equation (2), where we include subgroup effects heterogeneous over time and two-way interactions between DTI and subgroups. The subgroups are: movers (moved to another residence in year t), home owner, having negative home equity (the loan to value ratio of the house is above 1), and new in the high DTI top quantile (having high DTI at time t, but not at $t-1$). The estimated coefficients between the subgroups and high DTI are displayed in Model 3 of Table 2.

To be able to distinguish more clearly between the stock- and the flow-effect, we will focus on two specific groups. Figure 6 shows the average predicted consumption for households with low DTI and high DTI, and two groups within the high DTI group, namely households who (1) move to another house in year t, and (2) who own a house and have negative home equity, and who do not move in period t. Results for the first group should be driven by the flow effect, and for the second group by the stock effect.[16].

Households who move to a different house in year t have a substantially higher level of consumption in that year than other households with high DTI, especially in the years 2007 to 2009. Average consumption for this group sharply declines during the crisis and the years after. The positive effect of debt on consumption for this group lends support to the presence of the flow-effect. The decline of the effect over time seems to suggest that it became harder to use debt for on-off consumption during and after the crisis.[17]

The group of home owners with negative home equity (who do not move) also have higher average consumption than other households with high debt (e.g. home owners

[16]These two groups are based on combinations of the four subgroups in the regression model. The group of movers includes all households with high DTI who move in a particular year, with no restrictions on membership of other subgroups. There is substantial overlap (around 60% over time) between this group and the households who are new in the high DTI top quantile, since these households most likely use their new acquired debt to finance the moving. The second groups consists of all household who did not move in year t, own a house, and have negative home equity

[17]The initial sharp decline in consumption by households that move does not coincide with a reduction of the property transfer tax from 6 to 2% halfway 2011. So this cannot explain our result.

without negative home equity) before the crisis, but the difference with other households is much smaller than for movers and disappears over time. Similarly, the level of consumption also drops during the crisis, especially in the first few years, but less strongly than the consumption reduction of the movers. The marginal effects for households that move to a different house or have negative home equity (Figure 7) show time patterns that are very similar to those of predicted consumption in Figure 6.[18]

Figure 6b displays the change in consumption from time t to $t+2$. Households who have high DTI and move have by far the largest (negative) change in consumption. This again supports the idea of a flow effect, for this group, as these households most likely finance one-off consumption in t and then have to return to their normal consumption level in the consecutive years. Also in support of the other findings is that the decrease in consumption becomes smaller over time, suggesting that the use of debt for immediate consumption decreased due to the crisis.

5.3 The role of liquidity

To investigate to what extent binding liquidity constraints play a role for our findings, we compare two sub-samples: households with low liquidity, being households in the lowest decile of the distribution of real financial wealth over household's average real disposable income, versus households with high liquidity, being households in the top quantile of the distribution.[19] If tighter liquidity constraints are the main driving force that leads to a larger decrease in the consumption of households with high leverage, we would expect that highly indebted households with low liquid wealth reduce their consumption more strongly than highly indebted households with high liquid wealth.

Figure 8a and 8b display the average predicted consumption for households with low and high liquidity, respectively, by estimating (2) for the different sub-samples. The underlying estimated coefficient of both models are displayed in Model 5 and 6 of Table D.1, respectively. All subgroups in both figures exhibit similar patterns compared to Figure 6a. So liquidity constraints do not seem to play an important role.

[18]The marginal effect (ME) of each subgroup is calculated by adding the time heterogeneous effect of the high DTI group as a whole, the subgroup and finally the interaction effect between the subgroup and the high DTI dummy; so, using the notation of (2), this can be written as $\widehat{ME}_t^s = \hat{\beta}_{1t} + \hat{\beta}_{2t}^s + \hat{\beta}_3^s$, where the superscript s indicates that it is subgroup specific.

[19]The choice of these two sub-samples ensures that the behavior of the two groups are sufficiently distinguishable from each other. As a robustness check we have also investigated alternative groups of high-liquidity households, namely households that fall between the 10th to 90th, 10th to 75th, 25th to 75th and 25th to 90th percentile of the distribution. We find similar results. In Figure C.2 we display the distribution of real financial wealth over household's average real disposable income.

Figure 6: Average predicted consumption for subgroups of households with high DTI versus low DTI based on (2) using two different dependent variables

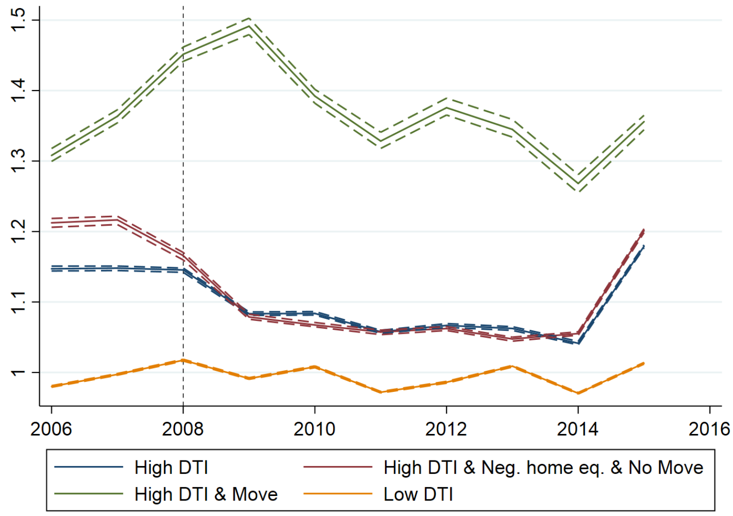

(a) Level of consumption: $\widehat{\frac{c_{i,t}}{\bar{y}_i}}$

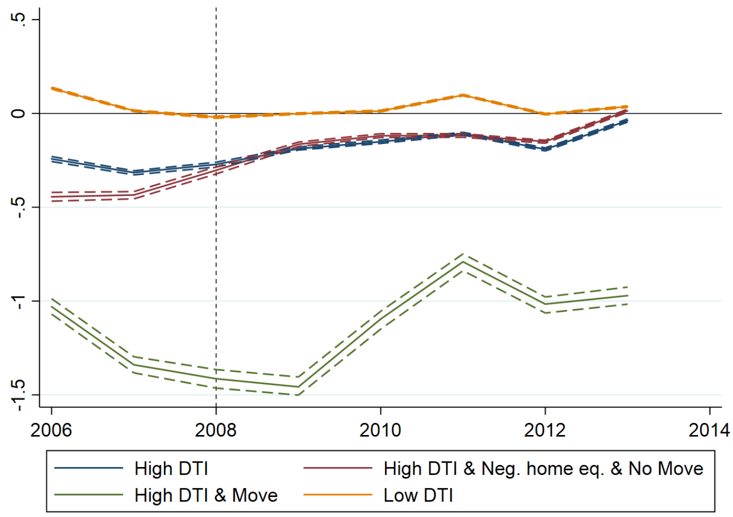

(b) Difference of consumption: $\Delta \widehat{\left(\frac{c_{i,t+2}}{\bar{y}_i}\right)}$

See the note of Figure 4 for more information.

Figure 7: Marginal effects of different subgroups

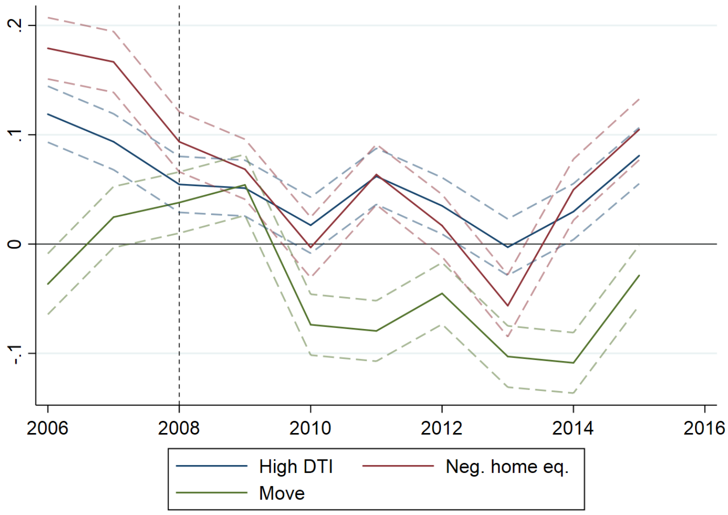

Based on the notation of (2), the marginal effects (ME) are constructed using $\widehat{ME}_t^s = \hat{\beta}_{1t} + \hat{\beta}_{2t}^s + \hat{\beta}_3^s$, where the superscript s indicates that this are the subgroup specific parameters. The 95% confidence bands are based on the covariance matrix of the underlying parameters.

5.4 Counterfactual analysis: effects on macro consumption over income

The results in Section 5.2 show that over time the drop in consumption over income was largest for households who move to a different house, but this does not necessarily mean that these groups also have the largest effects on macro consumption. As Figure 3 shows, movers only form a small and declining share of the households with high DTI. On the other hand the share of households with negative home equity and high DTI is much larger and increases over time. So the composition of the households with high DTI is changing.

To get some insight in how the changing effect of high debt across subgroups and the composition of highly indebted households has affected macro consumption over income, we perform a counter factual analysis, as explained in Section 4.3. Figure 9 shows the loss in macro consumption over income due to the changing effect of high debt across years. Each panel splits up the total effect into the effect for two exclusive subgroups: households who move to a different house versus households who do not move, or households with negative home equity versus no negative home equity. The groups are constructed in such a way that the sum of the effect for both subgroups always equals that of the impact of high DTI (the blue line).

At its most extreme point, 2014, the weighted average of consumption over income, for the whole population, would have been 5.9% of income higher than it actually was, if households with high debt would have behaved like they did in 2007. This implies that

Figure 8: Average predicted consumption for subgroups of households with high DTI versus low DTI based on (2) for different samples

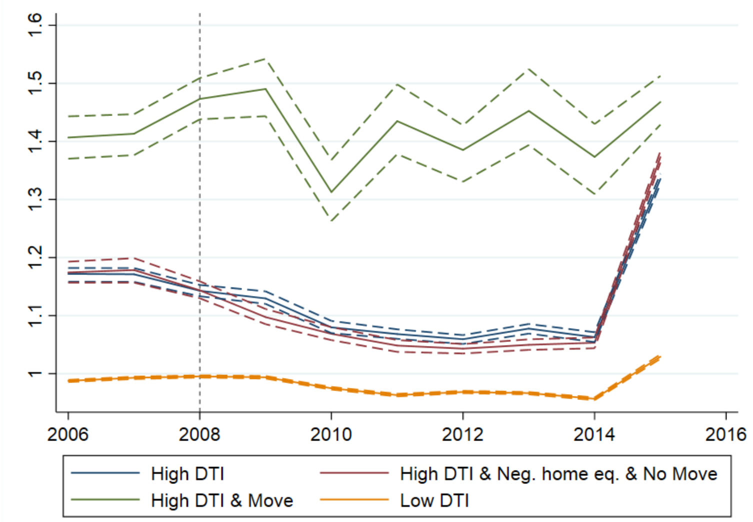

(a) Households with low liquidity (the lowest decile)

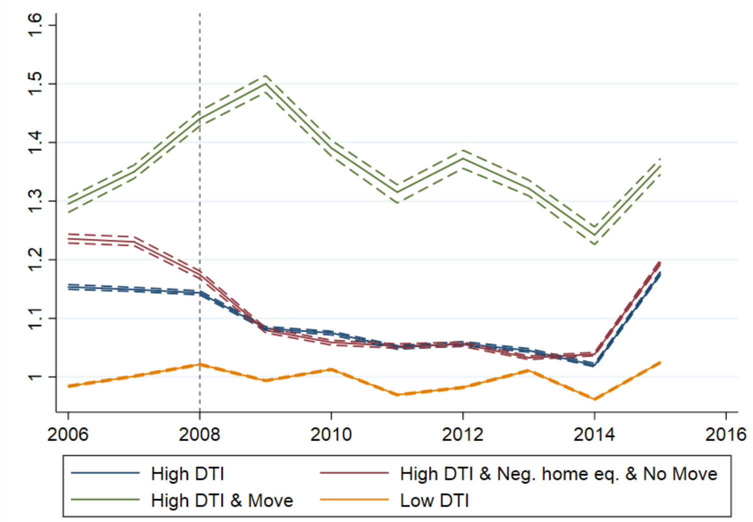

(b) Households with high liquidity (above 25th percentile)

See the note of Figure 4 for more information.

Figure 9: Change in macro consumption over income due to changing impact of high DTI

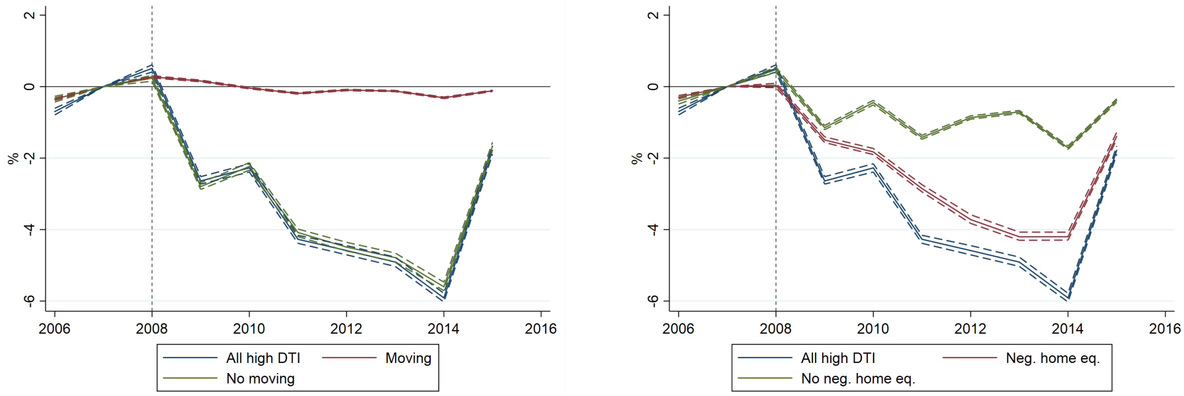

(a) Movers

(b) Negative home equity

The red and green line are constructed such that they together add up to the total impact of the high DTI group (blue line). See the note of Figure 4 for more information on the confidence bands.

the drop in consumption over income we observe in Figure 1 is mostly explained by highly indebted households.[20] In 2015, however, these households behave already quite similar to 2007 again and consumption is almost restored to pre-crisis levels. Strikingly, the effect on macro consumption seems to be driven largely by the highly indebted home owners with negative home equity. So although the drop in individual household consumption is most striking for groups affected by the flow-effect, the groups affected by the stock-effect are so much larger in size that this dominates the macro results.

6 Sensitivity analysis

6.1 Redefining the high DTI group as the top 10%

We investigate the robustness of our results for choosing a higher cut-off value when a household is highly indebted. We raise the cut-off value from the top 25 to the top 10 percentile of the sample, and rerun our analysis. In Figure 10 we display the predicted consumption for different subgroups. Again the results show a similar pattern compared

[20]If we compare this result with the change in consumption based on our own consumption proxy, we find a similar pattern as in Figure 1. The maximum decrease in consumption is 12 percentage points in the period 2008-2014. This difference might be driven by deviations in the sample under consideration. For example, the national accounts data also counts non-profit institutions for the benefit of consumers as households or because the average age in our sample is in later years slightly higher than the average age of the Dutch population.

to our results above and our main conclusions still hold. However, the differences between the different subgroups are more pronounced. This is because the top decile has a higher average DTI ratio. Therefore, the above described effects are exacerbated. For example, at the start the difference between households with high DTI and low DTI is much bigger, now they consume 24 percentage points more of their income instead of 17. This declines to 8 percentage points in 2013. Both the initial difference and the drop during the financial crisis are substantially higher compared to the results in Figure 6a. These results are confirmed by the estimation results of Model 7 in Table D.1.

Figure 10: Subgroups of high DTI household versus low DTI households defining the top 10% of the sample as high DTI

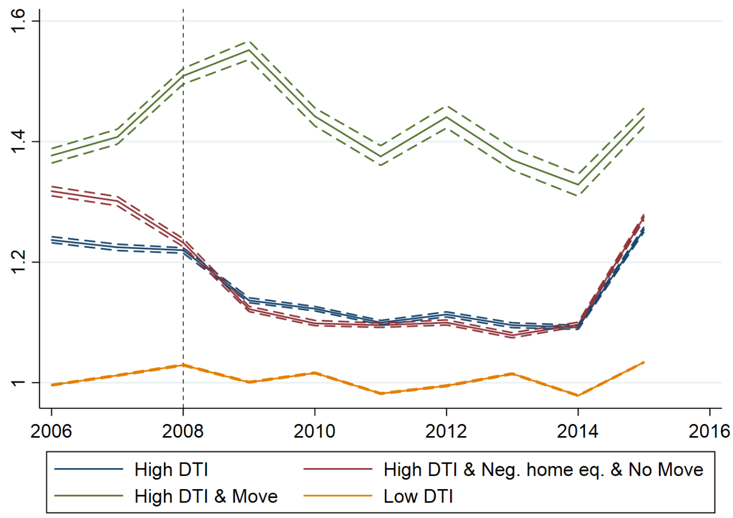

See the note of Figure 4 for more information.

6.2 LTV instead of DTI

We replace debt-to-income (DTI) by loan-to-value (LTV), to see whether the results are sensitive to the type of debt measure we use. Just like DTI, LTV is an often used measure to identify high financial risk, both at the individual and the macro level. Instead of looking at the amount of debt compared to income, LTV looks at the relationship between the amount of debt and the value of the underlying assets. The distribution of LTV is shown in Figure C.1. We rerun the pooled regression from Equation (2), but we replace $DTI^{high}_{i,t}$, by a dummy for being in the top quantile of LTV.[21] The underlying coefficient estimates

[21]Similar to Section 6.1 we also define the dummies for being in the top decile of LTV as a sensitivity analysis. The results are consistent.

are listed in Model 8 of Table D.1.

Figure 11 displays the predicted spending for subgroups of households with a high and low LTV ratio. The results confirm the leverage-spending relationship found in Section 5.1. Households with high leverage reduce their spending more strongly during the crisis period (2008-2013) than those with low leverage.[22] Nevertheless, households with high debt always spend more during the sample period. Also concerning the leverage-spending relationships across different subgroups, the results are similar to Figure 6. Households who move to a new house have substantially higher consumption than other households with high LTV, but also have the strongest decline of consumption during the crisis. Still, compared to Figure 6 there are two notable differences. First, households that move have a consistently higher level of consumption and also exhibit a smoother declining pattern between 2006 and 2012. Second, households with negative home equity have a considerably lower level of consumption.

Figure 11: Subgroups of high LTV household versus low LTV households

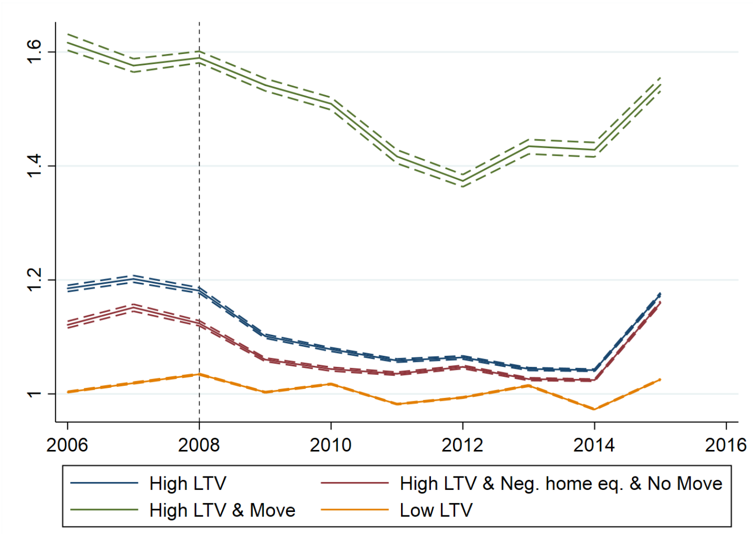

See the note of Figure 4 for more information.

[22]Starting from 2013, the government has tightened the maximum LTV ratio for households from 106% to 100% by steps of 1% each consecutive year. This does not coincide with the timing of the decrease in spending we observe and therefore cannot explain our results.

7 Discussion

We have estimated the relationship between household debt and consumption for the Netherlands in the period 2006-2015. We find that the financial crisis indeed had a stronger negative effect on consumption for highly indebted households than for others.

That households with high debt reduced their consumption more strongly during the crisis than others has been well established. However, Andersen et al. (2016) show that this pattern in itself is not necessarily evidence that the crisis had a larger impact on highly indebted households. It could also be explained by a normalization pattern that can also be found in other years before or after the crisis: households use high debt to finance one-off consumption, and return to their normal level in the following years. In this paper, we have tried to disentangle the relation between household debt and consumption by focusing on two possible channels: the use of debt to finance consumption (flow effect) and the effect of debt overhang on consumption (stock effect).

Just like Andersen et al. (2016), we find that households with high debt in one year have, on average, higher consumption in that particular year and lower in the subsequent years. However, we find that the normalization pattern does not fully explain the drop in consumption for highly indebted households after the crisis. First of all, we also find a drop in consumption among groups for which the stock effect of debt instead of the flow effect is important. Highly indebted home owners with negative home equity also reduced their consumption during the crisis, most likely due to precautionary motives. In fact, this group is largely responsible for the drop in consumption at the macro level. Although the reduction in consumption on the individual level was small for this group compared to other high debt groups, its relative size is large and increased sharply during the crisis.

Second, we consider a mechanism that Andersen et al. (2016) do not explicitly look at: the crisis can also have affected the willingness and ability of households to use high debt to finance one-off expensive consumption (the flow effect). At the individual level, the drop in consumption is largest for households who move to another house. Before the crisis, these households indeed seem to have used part of their mortgage debt to finance one-off high consumption (e.g. durables such as a kitchen) and, as a result, their consumption in the year of moving was higher than that of other households. During and after the crisis, however, households were less able or willing to use new debt for consumption, and the consumption of movers became more in line with that of other households.

The time patterns for highly indebted households with high and low liquid wealth are very similar, which suggests that households became less willing to use debt for one-off consumption mainly out of precaution and not because of binding restrictions. Changes in government policies and regulations might also have played a role. However, the time patterns in consumption that we observe predate the most important policy changes, which suggest that these are, at least, not solely responsible for our findings.

References

Andersen, A. L., Duus, C., and Jensen, T. L. (2016). Household debt and spending during the financial crisis: Evidence from Danish micro data. *European Economic Review*, 89:96–115.

Ando, A. and Modigliani, F. (1963). The "life cycle" hypothesis of saving: Aggregate implications and tests. *The American Economic Review*, 53(1):55–84.

Baker, S. R. (2018). Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data. *Journal of Political Economy*, 126(4):1504–1557.

Bezemer, D., Grydaki, M., and Zhang, L. (2016). More mortgages, lower growth? *Economic Inquiry*, 54(1):652–674.

Biggs, M., Mayer, T., and Pick, A. (2009). Credit and economic recovery. DNB Working Paper 218, DNB.

Bijlsma, M. and Mocking, R. (2017). The impact of house price shocks on the savings of Dutch homeowners and renters. CPB Discussion Paper 346, CPB Netherlands Bureau for Economic Policy Analysis.

Bound, J. and Krueger, A. B. (1991). The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right? *Journal of Labor Economics*, 9(1):1–24.

Browning, M., Gørtz, M., and LethPetersen, S. (2013). Housing Wealth and Consumption: A Micro Panel Study. *Economic Journal*, 0:401–428.

Browning, M. and Leth-Petersen, S. (2003). Imputing consumption from income and wealth information. *The Economic Journal*, 113(488):F282–F301.

Bunn, P. and Rostom, M. (2015). Household debt and spending in the United Kingdom. Bank of England working papers 554, Bank of England.

Carroll, C. D., Slacalek, J., and Sommer, M. (2012). Dissecting saving dynamics: Measuring wealth, precautionary, and credit effects. CFS Working Paper Series 2012/10, Center for Financial Studies (CFS).

Challe, E. and Ragot, X. (2016). Precautionary saving over the business cycle. *The Economic Journal*, 126(590):135–164.

Cooper, D. (2013). House Price Fluctuations: The Role of Housing Wealth as Borrowing Collateral. *The Review of Economics and Statistics*, 95(4):1183–1197.

Dynan, K., Mian, A., and Pence, K. M. (2012). Is a household debt overhang holding back consumption? *Brookings Papers on Economic Activity*, pages 299–362.

Eggertsson, G. B. and Krugman, P. (2012). Debt, deleveraging, and the liquidity trap: A Fisher-Minsky-Koo approach. *The Quarterly Journal of Economics*, 127(3):1469–1513.

Guerrieri, V. and Lorenzoni, G. (2017). Credit crises, precautionary savings, and the liquidity trap. *The Quarterly Journal of Economics*, 132(3):1427–1467.

Kreiner, C. T., Lassen, D. D., and Leth-Petersen, S. (2015). Measuring the Accuracy of Survey Responses Using Administrative Register Data: Evidence from Denmark. In Carroll, C. D., Crossley, T. F., and Sabelhaus, J., editors, *Improving the Measurement of Consumer Expenditures*, volume 74 of *Studies in Income and Wealth*, chapter 10, pages 289–307. Chicago: University of Chicago Press.

Kukk, M. (2016). How did household indebtedness hamper consumption during the recession? Evidence from micro data. *Journal of Comparative Economics*, 44(3):764–786.

Leth-Petersen, S. (2010). Intertemporal consumption and credit constraints: Does total expenditure respond to an exogenous shock to credit? *American Economic Review*, 100(3):1080–1103.

Mian, A., Rao, K., and Sufi, A. (2013). Household balance sheets, consumption, and the economic slump. *The Quarterly Journal of Economics*, 128(4):1687–1726.

Van Beers, N., Bijlsma, M., and Mocking, R. (2015). House price shocks and household savings: Evidence from Dutch administrative data. CPB Discussion Paper 299, CPB Netherlands Bureau for Economic Policy Analysis.

Zhang, L. (2019). Do house prices matter for household consumption? CPB discussion paper, CPB Netherlands Bureau for Economic Policy Analysis.

A Data cleaning procedures

In this appendix we describe the different data cleaning procedures. First, we describe how we clean income and wealth data. Second, we describe how we correct for the price effect when using the change of wealth data as a proxy for household savings. Finally, we discuss the cleaning procedure of the administrative consumption proxy.

A.1 Cleaning income and wealth data

To start we have to make the timing of wealth and income data consistent with each other. Income is measured by the tax authority on the last day of the year, while the stock of wealth is measured at the first day of the year. To make the timing consistent we chose to set the wealth back by one day, effectively lagging wealth by one year.

We drop all observations for which we do not observe disposable income. Furthermore, we delete all observations for which net household income is below a threshold value: 75% of the yearly social welfare level for a single household in 2009 (5760 euro).

Wealth data consists of ten different categories.[23] In all wealth categories we set missing values to zero. Furthermore, depending on the wealth category, we set negative values to zero if it is clear that this category cannot have negative values. In four wealth categories (i.e. value of the house, mortgage, value of other properties and assets of privately owned firms) we use linear interpolation if observations are missing, since the values of these categories are very stable over time and we can thus infer missing observations from surrounding years. For more information on cleaning these wealth categories we refer to Bijlsma and Mocking (2017). Lastly, we drop any observation for which the value of the house is below 50 thousand and above 2.5 million euros.

Finally, we want to remark that the distribution of the savings account data changes over time due to improved measurement of the small amounts (below 500 euro).[24] Because almost all households have a savings account, we do not omit this wealth category from the data. Instead we do not clean for it because it manly concerns small values and assume that the resulting measurement error is captured by the time FE.[25]

A.2 Accounting for the price effect in wealth data

Before we can construct the administrative consumption proxy we need to correct for the price effect in the change of wealth that we use as proxy for savings, as explained in Section

[23]The wealth data of 2015 is from a new data set provided by the CBS, with slightly different wealth categories. We change (as much as possible) all wealth definitions such that they coincide with the definitions in earlier years. As robustness check we omitted 2015 from our sample, this did not change the results.

[24]During our sample period the Dutch tax authority gradually got better access to information on all savings accounts by obtaining the data directly from banks. This mainly leads to an increase in the observed small amounts on savings accounts.

[25]Another solution is to truncate the savings account distribution each year, such that it does not change over time. This yields similar results.

3.3. Here we will describe how we correct for the price effect in each wealth category.[26]

We assume that the price effect on savings accounts, assets of privately owned firms and other assets can be neglected. For the first asset-type this assumption is not too far besides the true, since the interest rate in our sample period is close to zero. For the second there is most likely a strong price effect however it is not easy to correct for this and the number of households that privately own firms are small.[27] The third type consists of a wide range of assets, both liquid (e.g. cash) and illiquid (e.g. art), and therefore correction for the price effect is not possible.[28]

For stocks and bonds we correct for the price effect by using the national mutation in stocks and bonds due to financial transactions (savings) and due to changing prices (the price effect) from the national account data of the CBS. These series allow us to calculate the share of the mutation due to changing prices with respect to the total mutation in each year. Under the assumption that for each household the change in prices is equal to that on the national level, we can calculate which share of the change in their stocks and bonds portfolio is due to investment. This approach is more accurate than calculating the excess return of a household's portfolio compared to a benchmark portfolio based on the Amsterdam Exchange (AEX) index in year t and assume that the difference equals the additional investment in stocks in that year (see for example Andersen et al. (2016)).[29] The latter leads to substantial deviations compared to macro level mutations in stocks, especially during the financial crisis.

We only take into account a change in housing wealth when a households moves to a

[26]Inter vivos transfers, that is, wealth transfers from one household to another, are also a problem. These transfers are mostly from parents to their children. The consumption pattern of both the parents (they dissave and appear to consume more) and the children (they save and appear to consume less) are distorted. However, there are only about 5 thousand transfers each year and the monetary value is typically small because of tax treatment. So this will not distort the consumption proxy too much. However, in 2013 and 2014 the tax-free amount parents could transfer to their children was temporarily increased to 100 thousand euros, as long if it was invested in their own house. This leads to a strong jump in the amount of transfers. However, we do not think it drives our results since all the main patterns already occur before 2013. For more information about inter vivos transfers see the report by the national auditors: www.rekenkamer.nl/publicaties/rapporten/2017/12/06/schenkingsvrijstelling-eigen-woning.

To see whether large inter-vivos transfers affect our results we have omitted all observations in which households reduced their mortgage with more than 10%. So we still have small reductions in mortgages due to annual amortization in our sample, but exclude large reductions in mortgages potentially due to inter vivos transfers. The main results are not affected.

[27]As robustness check we omitted all observations in which households have non-zero assets of privately owned firms. This did not influence the main results.

[28]As robustness check we omitted all observations in which households have non-zero other assets. This did not influence the main results.

[29]However, if we use the AEX index to correct for the price effect, the main results stay qualitatively the same. Although the impact on consumption is stronger in the earlier years of the crisis compared to our baseline results. This is because the AEX index underestimates the price effect during the crisis, making it look like households save more and consume less.

new address, because only in this case the change in housing wealth is due to investment in housing, i.e. savings, instead of a price effect. In all other cases it is a price effect. However, we do not know the exact selling price of the old house and the purch asing price of the new house for each household that moves. Furthermore, we only observe the value of the house at the end of each year. However, we do observe the month that a households changes its address, so we impute the value of the old and new house to this specific month using the housing price index of the CBS and assume that this is the actual selling and buying price of the old and new house, respectively.

We do not observe what other real estate is comprised of. This category consists of second houses, but also holiday homes. We correct for the price effect by assuming that a household only invests in real estate if the year-on-year growth rate is more than 15%.[30] All smaller changes are assumed to be caused by the price effect. The choice of this benchmark is based on the Dutch national housing price index that varies between 5% and −7%. Because the growth in real estate prices varies widely between regions, we let the benchmark be substantially higher than the maximum absolute change in the price index.

A.2.1 Cleaning the administrative consumption proxy

We clean the constructed administrative consumption proxy.[31] We drop all observations for which we do not observe any spending. On a household level, we drop all households that have an aggregate consumption that is negative or if their average consumption over time is lower than 5760 euro, that is, 75% of the social welfare in 2009. Furthermore, we drop households for whom their average consumption is 120 thousand euro larger than their disposable income. Lastly, we winsorize consumption at the bottom, at 5760 euro, and at the top, at one million. The latter threshold is chosen because disposable income is also winsorized at one million by the CBS.

Next, we drop all observations in which a major change in the household-type occurs: from single to a couple or vice versa. We do this because in the year of this event household income and wealth change substantially since wealth is split between or aggregated over the two individuals in the household, respectively.

Finally, we construct our dependent variable $c_{i,t}/\bar{y}_i$. To prevent outliers, most likely caused by measurement errors, from influencing our results, we winsorize our dependent variable at the top at 2.5 and at the bottom at 0.2. This is a conservative approach, where we winsorize 4% and 5% of the observations, respectively.

[30] We have also experimented with a 10% and 20% cut-off value, this did not influence the results.

[31] see Section 3.3 for a description of its construction

B The quality of the consumption measures

Figure 2 shows a scatter plot of the budget survey and the administrative proxy of consumption (both scaled by household's average real disposable income). In Section 3.4 we discuss that between 0.4 and 0.8 both are relatively similar, but that they deviate substantially in the low and high intervals. This indicates that there is substantial amount of measurement error in one or both measures.

Figure B.1 shows two similar scatter plots for all the households with high DTI (Figure B.1a) and with low DTI (Figure B.1b). Both groups have a linear slope coefficient close to that of the full sample (0.48). This indicates that for both subgroups the data is of comparable quality. Looking at even smaller subgroups is difficult because it reduces the number of available observations, making the slope coefficient more sensitive to outliers. Although both measures show similar patterns across subgroups, there is still room for improvement in reducing the difference between both. We leave this for future research.

Figure B.1: Budget survey versus administrative proxy for high and low DTI

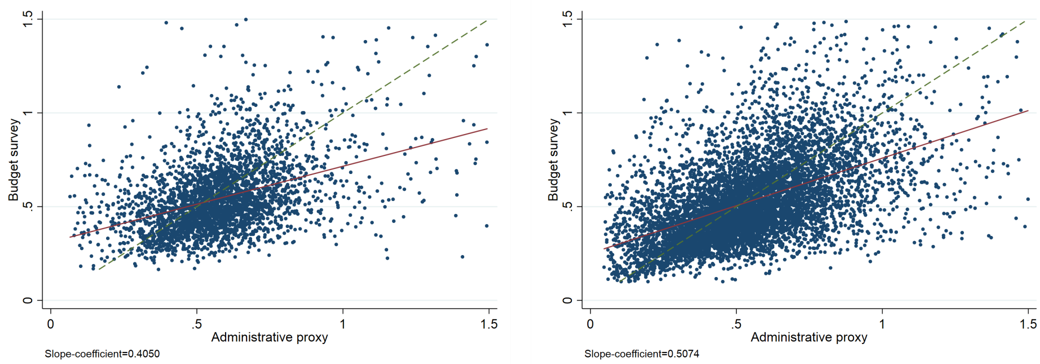

(a) High DTI (b) Low DTI

See the figure note of Figure 2.

C Additional descriptive statistics

We display three additional histograms of DTI, LTV and real financial wealth scaled by household's real disposable income averaged over time, where we define financial wealth as in Section 4: the sum of the savings accounts, stocks, bonds and other assets. In Figures C.1a-C.1b we display the position of the top quantile and decile, which we use in our regression analysis to define $DTI_{i,t}^{high}$ or $LTV_{i,t}^{high}$ (see Section 4). In Figure C.2 we display the bottom decile and quantile, which we use in Section 5.3 for the sample split between households with low and high liquidity, respectively.

Figure C.1: Histograms for DTI and LTV

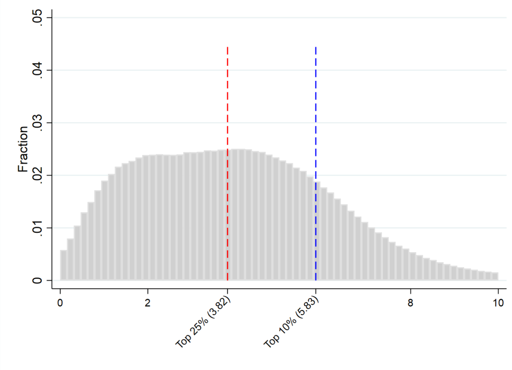

(a) DTI

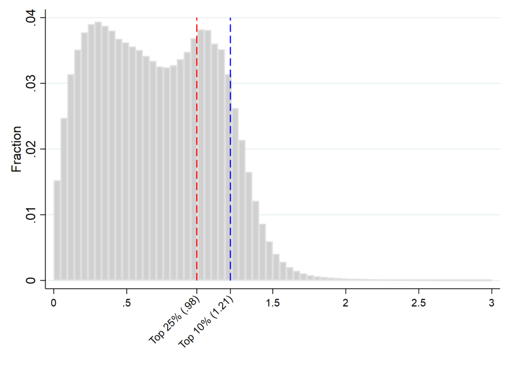

(b) LTV

In both figures the dashed red and blue line indicate the top quantile and decile, respectively. For the purpose of display we exclude all observations for which DTI and LTV equal zero and for which the LTV ratio is larger than 3. Remark that our data is constructed such that the DTI ratios are capped at 10 (see Section 3.2).

D Underlying estimation results

Table D.1: Regression results for (2) using different model specifications

Model		5	6	7	8
Description		Low liq.	High liq.	q=10	LTV
$DTI^{high}_{i,t}$ / $LTV^{high}_{i,t}$	2006	0.044	0.122*	0.020	−0.061*
		(0.040)	(0.017)	(0.017)	(0.007)
	2007	0.065	0.095*	−0.004	−0.141*
		(0.040)	(0.017)	(0.017)	(0.007)
	2008	0.0174	0.057*	−0.031	−0.196*
		(0.040)	(0.017)	(0.017)	(0.007)
	2009	−0.018	0.052*	−0.064*	−0.251*
		(0.040)	(0.017)	(0.017)	(0.007)
	2010	0.005	0.018*	−0.079*	−0.263*
		(0.041)	(0.017)	(0.017)	(0.007)
	2011	−0.008	0.065*	−0.053*	−0.129*
		(0.040)	(0.017)	(0.017)	(0.006)
	2012	−0.023	0.041*	−0.062*	−0.225*
		(0.040)	(0.017)	(0.017)	(0.008)
	2013	−0.015	−0.002	−0.084*	−0.164*
		(0.040)	(0.017)	(0.017)	(0.007)
	2014	0.024	0.028	−0.059*	−0.03 *
		(0.039)	(0.017)	(0.017)	(0.005)
	2015	0.127*	0.082*	−0.005	0.012*
		(0.040)	(0.017)	(0.017)	(0.005)
$\cdot Move_{i,t}$		−0.136*	−0.236*	−0.338*	0.106*
		(0.016)	(0.004)	(0.005)	(0.004)
$\cdot NegHomeEq_{i,t}$		−0.020*	−0.013*	−0.057*	–
		(0.009)	(0.003)	(0.002)	
$\cdot OwnHouse_{i,t}$		0.075	0.019	0.127*	–
		(0.039)	(0.017)	(0.017)	
$\cdot MoveToDTIh_{i,t}$		0.624*	0.777*	0.788*	–
		(0.040)	(0.010)	(0.010)	
R^2		0.166	0.079	0.074	0.064
N		598,753	4,436,631	5,928,890	5,928,890

The dependent variable is $c_{i,t}/\bar{y}_i$. In Model 5 and 6 the sample is split based on the financial wealth distribution (see Figure C.2). Model 5 uses all households that are in the lowest decile, Model 6 use all households above the 25th percentile. In Model 6 the high DTI group is defined as the top decile of the sample. Finally, in Model 7 DTI is replaced by LTV. In this model the last three interaction terms are omitted to prevent multicollinearity. For more information see the table note of Table 2.

Figure C.2: Histogram of real financial wealth scaled by household's average real disposable income

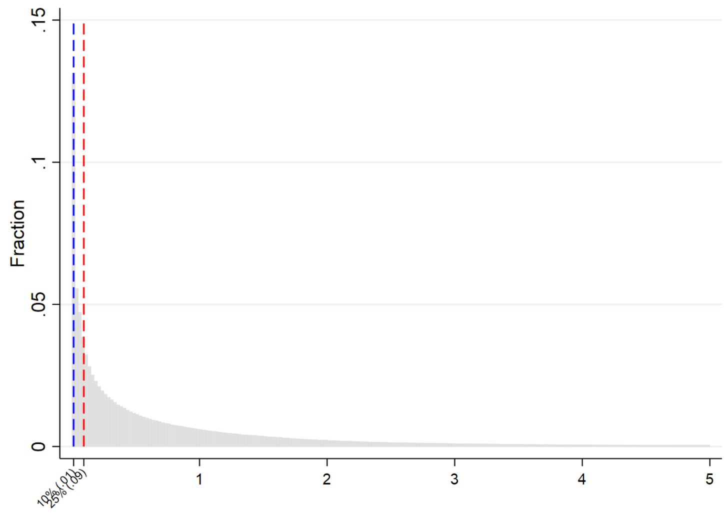

The dashed blue and red line indicate the bottom decile and quantile, respectively. For the purpose of display we exclude all observations for which financial wealth is equal to zero.