HOUSEHOLD DEBT AND CONSUMPTION DURING THE FINANCIAL CRISIS

Asger Lau Andersen, Economics, Charlotte Duus and
Thais Lærkholm Jensen, Financial Markets

INTRODUCTION AND SUMMARY

Danish households' gross debt rose sharply in the years leading up to the financial crisis. A high level of debt may affect households' response when the economy is hit by a financial crisis. House prices dived during the most recent crisis. Combined with the high level of debt this meant that homeowners' loan-to-value (LTV) ratios reached an unusually high level. This may have created a need for consolidation among households, which has contributed to dampening private consumption and thus aggregate demand and economic activity.

This article presents an analysis of the relationship between the LTV ratios among Danish homeowner families before the most recent financial crisis and the families' consumption patterns during the crisis. The analysis is based on register-based microdata for Danish homeowner families in the period 2003-11. The methodology and results are described in detail in a recent working paper, cf. Andersen et al. (2014). Our analysis shows a clear negative relationship between a family's LTV ratio in 2007 and the change in its consumption in the following years. This also applies if a number of family-specific conditions are taken into account, including developments in the family's income and the value of its home(s). The negative relationship is broad-based and exists in all parts of Denmark and among homeowner families in all age groups and income brackets. Consequently, the analysis indicates that the high debt level among households contributed to amplifying the drop in private consumption during the financial crisis.

In previous analyses, Danmarks Nationalbank has investigated the impact of the high level of debt on the ability of households to service their debts and on financial stability in Denmark, cf. Andersen et al. (2012a, 2012b, 2013). The overall conclusion of these analyses is that the vast majority of Danish homeowner families have robust finances, and even in the event of considerable financial adversity they will be able to service their debts. Hence, the high level of debt among Danish households is not a serious threat to financial stability in Denmark.

This article addresses another aspect of the issue concerning the high level of debt among Danish households, which has not been the direct focus of the previous analyses, namely the consequences for consumption and *macroeconomic* stability in periods of financial turmoil. The results indicate that the high level of debt among Danish households contributed to making the Danish economy more vulnerable in the boom years just before the financial crisis, presumably causing more pronounced fluctuations in consumption and economic activity in the years around the financial crisis than would have been the case if household indebtedness had been lower.

The rise in household debt in the years before the financial crisis is closely linked to the development in house prices. Higher house prices improve homeowners' opportunities

Household housing debt, home equity and LTV ratio

Chart 1

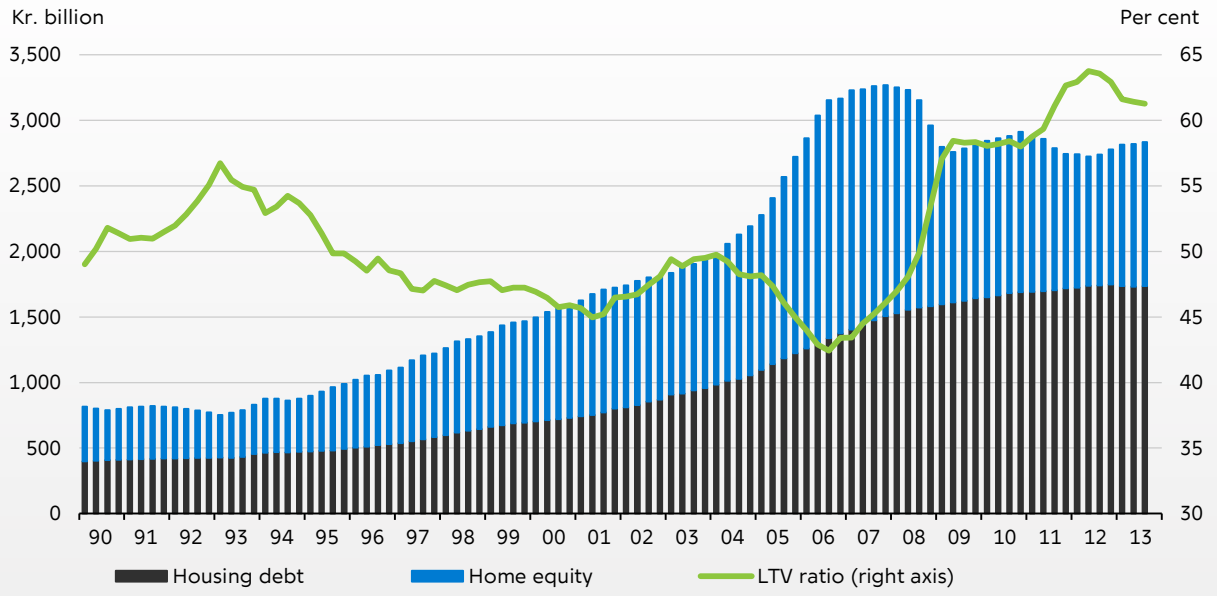

Note: Here, the LTV ratio is defined as the total housing debt of households as a percentage of the total value of their homes. Housing debt includes all debt raised against a home as collateral.
Source: Danmarks Nationalbank, Statistics Denmark and own calculations.

to raise loans against the home as collateral. For many families, a higher price level also means that they must take out larger loans to enter the housing market. It is therefore likely that the surging house prices in the pre-crisis years contributed to a strong increase in gross household debt during that period – thus also being a factor behind the high LTV ratios experienced by many families when house prices subsequently fell. In our assessment, based on the results in this article, more subdued house price developments in the years before and after the financial crisis would have resulted in a less pronounced fall in private consumption during the crisis. This emphasises that a stable development in house prices is a key condition for stable macroeconomic development.

MACROECONOMIC BACKGROUND

The total housing debt of households rose sharply over the past decade, cf. Chart 1. The rise was particularly pronounced in the pre-crisis years. But the rapidly increasing housing debt was offset by even greater price rises in the housing market during that period, and the LTV ratio, i.e. the ratio of mortgage debt to housing value, declined. Just before the outbreak of the financial crisis, the situation in the housing market reversed, however, and when the crisis erupted, the increases of previous years were replaced by large price drops. Combined with the already high level of debt this meant that homeowners' LTV ratios reached an unusually high level.

This article examines the relationship between a family's pre-crisis LTV ratio and its consumption response during the crisis. The LTV ratio reflects the *relative* relationship between the family's assets and debt.[1] Net wealth, i.e. the *absolute* difference between the value of assets and liabilities, is a closely related concept. A high LTV ratio will often coincide with low or negative net wealth. But it is also possible to increase the LTV ratio without changing net wealth . This can be done by increasing the holdings of both assets and liabilities, i.e. balance-sheet expansion.

1 The LTV ratio is calculated as the ratio of the family's total debt to banks and mortgage banks and its mortgage deed debt to the value of its home(s). The calculation is explained in more detail in the next section. Net wealth, on the other hand, is calculated as the difference between total assets and liabilities.

This article focuses on just that form of variation in the families' LTV ratios. The purpose of the analyses in this article is thus to compare the consumption developments during the crisis for families who had different LTV ratios immediately before the crisis, but who are otherwise comparable in many other dimensions, including net wealth. Box 1 presents a hypothetical example illustrating the question we seek to answer in this article.

There are sound theoretical arguments that the LTV ratio may affect the consumption of a household: Loans secured on the home are the most important source of credit for most homeowners. If the debt exceeds the value of the home, obtaining further credit via this channel is not possible, however, and this may cause the household to reduce its consumption. The risk of this situation arising may be sufficient to affect household consumption decisions. Hence, some households can be expected to want to insure themselves against unexpected events by creating a buffer with a suitable distance to any upper limits for the LTV ratio. Accordingly, a high LTV ratio may affect consumption negatively long before these upper limits are reached, cf. Carroll (1997).

There is also empirical evidence of a negative relationship between debt on the one hand and consumption and/or economic activity on the other. Such a relationship has been found in a number of different countries using very different methodologies and data sources as described in Andersen et al. (2014).

A hypothetical example Box 1

Take two families, Family A and Family B, with the same income. The two families are also identical as regards size and age, and they purchased their homes in the same year. Furthermore, the families have the same *net wealth*, i.e. the absolute difference between assets and liabilities is the same in both families. The balance sheet of Family A exceeds that of Family B, however, so Family A has larger *gross wealth*, but also larger *gross debt*. Family A may live in a home worth kr. 2 million, in which its debt amounts to kr. 1.5 million, while Family B lives in a home worth kr. 1 million in which it owes kr. 500,000. Family A thus has an LTV ratio of 75 per cent, while Family B's LTV ratio is 50 per cent.

Enter the financial crisis: House prices drop, credit standards are tightened, and uncertainty about the families' future income increases. This article seeks to answer the question of whether the difference in the LTV ratios of the two families in itself induces Family A to respond more strongly to the changed circumstances than Family B and reduce its consumption more substantially.

DATA

The data used in the analyses in the article consists of register-based individual-level data from Statistics Denmark for the period 2003-11. Information on income, wealth and debt comes from the personal income register. The main source of this register is information from the Central Tax Administration (SKAT). Information on e.g. age, residence and family relations is retrieved from the population register. Using the latter, all income, wealth and debt information is aggregated at family level, and measures of the family's consumption, LTV ratio and other background variables are constructed.

A key challenge for the analysis is that register-based information on consumption is not available at individual or family level. For this reason, consumption is calculated residually using information on disposable income and wealth, cf. Box 2.

The LTV ratio for a given homeowner family is calculated as the family's total bank and mortgage debt as well as its mortgage deed debt at the end of the year as a percentage of the total value of the family's properties at that point in time. It should be emphasised that this relates to the family's *gross debt*, so any holdings of financial assets, pension wealth, etc. have not been deducted. The value of the family's properties is estimated at approximated market prices by multiplying the public property valuations by a scaling factor reflecting the relationship between public valuations and actual sales prices for the relevant combination of the type of property, geographical area and year. The method is described in more detail in Andersen et al. (2012a). Unfortunately, the data does not allow us to distinguish between bank loans secured on the home and other bank debt, so all debt to banks is included in the calculation of the LTV ratio.

Imputing consumption from income and wealth data

Box 2

Statistics Denmark's personal income register contains information on each citizen's disposable income, assets and liabilities. This information can be aggregated at family level by adding up the data for the individual members of each family. Based on these variables, we construct an imputed measure of the family's consumption (excluding housing consumption) by observing the family's total disposable income in a given year and comparing it with the change in the family's net wealth.[1] This method follows the approach described in Browning and Leth-Petersen (2003) and later applied in Leth-Petersen (2010).

A family's consumption in year t, c_t, is imputed using the following accounting identity:

$$c_t = y_t - s_t$$

where y_t denotes the family's disposable income, and s_t denotes the family's net saving over the year. While information about the family's disposable income is available from the personal income register, no precise information is available about net saving. The latter is therefore approximated as the change in the family's net wealth from the beginning to the end of the year. The rationale behind the approximation is that if the amount of the family's net wealth is smaller at the end of the year than at the beginning, this indicates that the wealth reduction was used for consumption. Conversely, an increase in net wealth would indicate that part of the disposable income was saved and consequently not used for consumption.

The imputation results in a noisy measure of the family's consumption. The primary source of noise is that, as a main rule, changes in net wealth attributable to saving (i.e. changes in the physical holdings of assets and liabilities) are non-separable from changes attributable to positive or negative capital gains (i.e. changes in the prices of assets and liabilities).

The most important type of capital gains is changes in the families' housing wealth as a result of house price fluctuations. In this particular case, however, capital gains can be separated from saving. The reason is that it is possible, using register data, to identify the families that were involved in property transactions in a given year. These families are excluded from the analysis. For the remaining families, who were not involved in property transactions, changes in their housing wealth from the beginning to the end of the year must then be attributable to capital gains.[2] Hence, for these families changes in housing wealth do not reflect actual saving, and such changes are therefore not included in the imputation of consumption.

Fluctuations in stock prices are another important source of capital gains. Unfortunately, capital gains are non-separable from saving in this case, since the data used does not include information about the families' purchases and sales of stocks. Instead, we make use of a crude adjustment for capital gains on the families' stock portfolios. This is done by multiplying the value of the family's stock portfolio at the beginning of the year by the average stock price development over the year. The result is taken to reflect the capital gain from the family's stocks.

One of the most important ways for Danish families to save is via contributions to pension schemes. Here, we are favoured by having access to exact data for the saving component in the form of annual contributions to the schemes, so in this case no approximation is required. Contributions to employer-administered pension schemes are pre-deducted in the calculation of disposal income, so it is only necessary to deduct contributions to privately administered schemes in the consumption imputation.

Even after these adjustments, the imputed measure of consumption is fairly noisy. This is evident from the fact that for some families the method results in extreme values. To minimise the impact of such outliers we impose an additional restriction on our analysis sample: In each year, the families are ranked by the ratio of their imputed consumption to their disposable income, and the bottom and top 5 per cent of the families are then excluded.

1. See Andersen et al. (2014) for more details.
2. The changes may also be attributable to investments in a home that the family already owns, e.g. in the form of additions or renovations, but our data does not allow us to identify the families who make such investments. Using our method, the increase in value resulting from these investments will consequently be regarded as a capital gain, while the investment costs will generally be registered as consumption in the year they are incurred.

Initially, the sample comprises all homeowner families in Denmark that include minimum one person aged 15-99. We then add a number of restrictions: Families with members who are self-employed or non-taxpayers are excluded.[2] The same goes for families who have moved or been involved in property transactions during the period analysed, cf. Box 2. Another restriction is that the composition of adult members of the family must remain the same throughout the analysis period. This excludes families that have been exposed to divorce, a death in the family, etc. Finally, we exclude families whose estimated consumption is either extremely high or extremely low relative to their disposable income as described in Box 2. After these restrictions, the analysis sample comprises almost 800,000 families.

2 Non-taxpayers are persons who are not liable to pay tax in Denmark. They include e.g. foreigners with a provisional residence permit in Denmark.

Development in disposable income, housing wealth and consumption: families with high LTV ratios versus other homeowner families

Chart 2

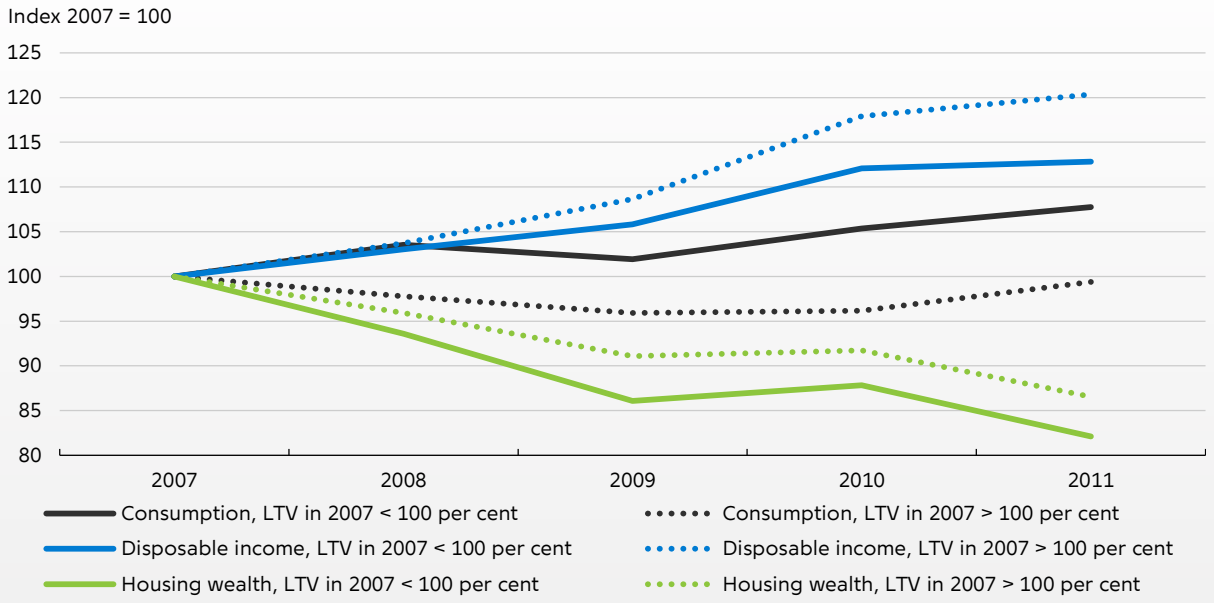

Note: The chart shows the nominal developments in disposable income, housing wealth and consumption, indexed relative to 2007, for two different groups: 1. Homeowner families with LTV ratios of less than 100 per cent in 2007 (solid lines). 2. Homeowner families with LTV ratios above 100 per cent in 2007 (dotted lines). Indexations are made for each family, i.e. relative to the 2007 level for the same family. The median value of the three indexed variables within each group is shown for every year. The chart only includes families that existed in every year of the period 2007-11 and were not involved in property transactions during the period under review.
Source: Own calculations based on register data from Statistics Denmark.

LTV RATIO AND CONSUMPTION RESPONSE DURING THE FINANCIAL CRISIS

The purpose of the analyses in this article is to illustrate how a high level of debt among Danish homeowner families has affected the families' consumption decisions in the years since the outbreak of the financial crisis. One way of answering this question is to divide the homeowner families into groups according to their LTV ratios at end-2007 – i.e. immediately before the financial crisis – and then to compare consumption developments since 2007 across those groups.

A SIMPLE COMPARISON

Chart 2 illustrates a simple version of such a comparison. Here, the homeowner families are divided into two groups: The first group consists of the homeowner families with LTV ratios of less than 100 per cent in 2007, while the second group consists of the highly indebted homeowner families with LTV ratios above 100 per cent at end-2007. The chart shows developments in disposable income, housing wealth and consumption (excluding housing consumption) since 2007 for the median family in each of the two groups.

The group of families with high LTV ratios experienced a stronger relative development in disposable income than the other homeowner families in the period 2007-11. One underlying factor is the drop in mortgage rates during that period, which has mainly benefited the highly indebted families. Housing wealth decreased in the period under review for both groups of homeowner families, reflecting the drop in house prices.[3] In relative terms, the drop was less pronounced for the group of highly indebted families than for the other families, however. Nonetheless, consumption showed a weaker development among the families with high LTV

3 Families that have been involved in property transactions during the period under review are not included in the chart. Hence, changes in the housing wealth of an individual family are attributable only to changes in house prices and any improvements to or wear and tear on the family's home(s).

Descriptive statistics Table 1

	LTV ratio in 2007				
	0-40 per cent	40-60 per cent	60-80 per cent	80-100 per cent	Over 100 per cent
Number of families	363,142	162,821	127,161	73,145	65,975
Number of children, mean	0.3	0.8	1.0	1.1	1.3
Age of oldest person, mean	64.6	52.5	47.9	45.1	44.0
Number of years since the family took up residence at the address, mean	29.2	15.9	12.2	10.3	9.6
Disposable income, mean, kr.	278,437	330,971	337,230	337,058	343,236
Net wealth, mean, kr.	2,450,028	1,337,761	763,072	319,295	-169,659

Note: Descriptive statistics for the families included in the analysis sample. The families are grouped according to their LTV ratios in 2007, and the variable means are calculated on the basis of data for 2007.
Source: Own calculations based on register data from Statistics Denmark.

ratios than among the other homeowner families. For the median family in the former group, consumption in 2009 was almost 5 per cent below the 2007 level, while the median family in the group of other homeowner families experienced consumption growth of just under 2 per cent from 2007 to 2009. The difference in consumption growth between the two groups widened further in 2010 and continued to exist in 2011, when consumption among the highly indebted families remained below the 2007 level.

Hence, the above comparison shows that homeowner families with high LTV ratios prior to the financial crisis reduced their consumption more than other homeowner families during the crisis, even though they experienced a more favourable development in both disposable income and housing wealth during the same period. This simple comparison does not take into account that families with high LTV ratios also differ from the other homeowner families in a number of other areas. For instance, families with high LTV ratios in 2007 are generally younger than families with low LTV ratios, cf. Table 1. At the same time, they have more children and higher incomes, whereas average net wealth, including pension wealth, is lower the higher the LTV ratio. These differences alone may have contributed to families with different LTV ratios in 2007 showing different consumption developments in the subsequent years, so it is important to take them into account.

ECONOMETRIC ANALYSIS OF LTV RATIO AND DEVELOPMENT IN CONSUMPTION

Andersen et al. (2014) present an econometric analysis of the relationship between the LTV ratio in 2007 and the subsequent development in consumption which controls for the above differences between families with different LTV ratios, among other factors. The econometric method is outlined in Box 3. The main findings are illustrated in Chart 3, which shows the estimated relationship between the LTV ratio in 2007 and the subsequent development in consumption in a given period. The first year of the period is 2007 in all the cases shown, while the time horizon of the period varies from 1 to 4 years. The chart shows no clear relationship between LTV ratios and the development in consumption as long as the LTV ratio in 2007 is below approximately 40 per cent. But at higher LTV ratios there is a clear negative relationship: The higher the LTV ratio in 2007, the weaker the development in consumption in the subsequent years. This applies whether the period under review is 1, 2, 3 or 4 years.

In terms of size, the differences are considerable. For an average family, an LTV ratio of 100 per cent in 2007 would, according to our esti-

Regression analysis of LTV ratio and subsequent change in consumption

Box 3

In Andersen et al. (2014) we set up and estimate an econometric model in which the development in a family's consumption since 2007 is modelled as a function of the family's LTV ratio in 2007. The model is inspired by the method in Dynan (2012) and is given as the following equation:

$$\Delta C_{i;07\text{-}s} = \alpha + F(\boldsymbol{\beta}, LTV_{i;07}) + \delta_1 \ln(Y_{i;07}) + \delta_2 NW_{i;07} + \delta_3 LA_{i;07} + \delta_4 \Delta Y_{i;07\text{-}s} + \delta_5 \Delta H_{i;07\text{-}s} + \delta_6 \Delta kids_{i;07\text{-}s} + \delta_7 \Delta C_{i;06\text{-}07} + \boldsymbol{\gamma X}_{i;07} + \varepsilon_{i;s}$$

where the variable on the left-hand side denotes the change in consumption from 2007 to an end-year s for family i. The end-year may be 2008, 2009, 2010 or 2011. To facilitate comparison across families, the change in consumption is measured as a percentage of the family's total income before tax in 2007.

The central explanatory variable on the right-hand side of the equation is the LTV ratio in 2007. To allow for any non-linear effects, this variable is represented by a parametric function F, which is assumed to be continuous and piece-wise linear. Specifically, it is assumed that the function is linear in each of the LTV ratio intervals [0;20],]20;40],]40;60],]60;80],]80;100],]100;120] and]120;∞[, but with the possibility of different slopes in different intervals. The slopes in each interval are given by the parameter vector β.

The family's disposable income in 2007, $Y_{i;07}$, net wealth in 2007, $NW_{i;07}$, and the family's holdings of liquid assets in 2007, $LA_{i;07}$, are included as control variables. Disposable income is measured logarithmically, while net wealth and the holdings of liquid assets are measured as percentages of the family's total income before tax in 2007. In addition, we control for the changes in disposable income, $\Delta Y_{i;07\text{-}s}$, housing wealth, $\Delta H_{i;07\text{-}s}$, and the number of children in the family, $\Delta kids_{i;07\text{-}s}$, from 2007 to the last year of the period. The changes in disposable income and housing wealth are measured as percentages of the family's income before tax in 2007.

The variable $\Delta C_{i;06\text{-}07}$ denotes the change in consumption from 2006 to 2007, given as a percentage of income before tax in 2007. This variable is included to allow for any extraordinary fluctuations in consumption in 2007. For instance, if the family bought a car in 2007, this will appear as a substantial increase in consumption in that year – and a subsequent major decline from 2007 to 2008. If the purchase of the car is loan-financed, it will also, all else equal, lead to a higher LTV ratio in 2007. If this is not taken into account, a negative correlation may be found between the LTV ratio and the subsequent development in consumption, driven solely by random, extraordinary fluctuations in consumption in 2007.

Finally, $\boldsymbol{X}_{i;07}$ denotes a vector of the family's characteristics in 2007 which may affect the subsequent development in consumption: the age of the family's oldest member, the age of the family's youngest child, the number of years since the family took up residence at the address, the number of pensioners in the family, the number of family members with higher education, and the family's municipality of residence. For a more detailed description of all variables and estimation results, see Andersen et al. (2014).

LTV ratio in 2007 and subsequent change in consumption

Chart 3

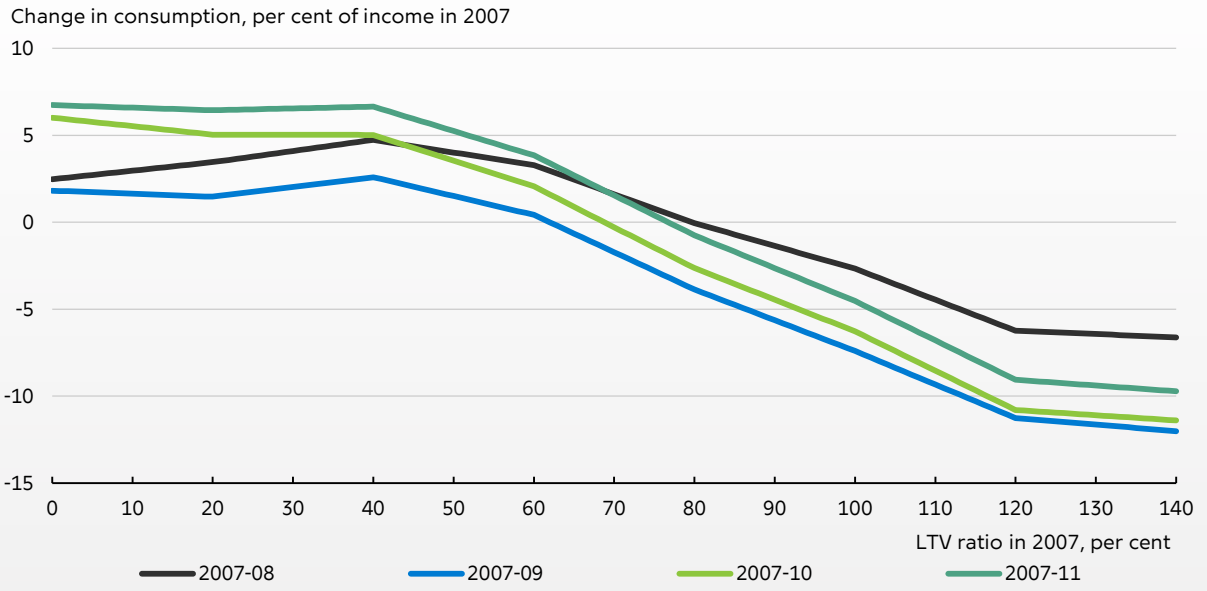

Note: The chart is based on an estimation of the econometric model described in Box 3. The curves indicate the average model-predicted changes in consumption over the time horizon stated for various values of the LTV ratio in 2007. For a given value of the LTV ratio in 2007, the average model-predicted change in consumption is calculated as follows: First, the model-predicted change in consumption at the relevant LTV ratio for each family in the sample is calculated, given the family's other characteristics. The average is then calculated for all the families in the estimation sample.
Source: Own calculations based on register data from Statistics Denmark.

mates, result in a fall in consumption from 2007 to 2011 of 4.5 per cent of its income before tax in 2007. Conversely, a comparable family with an LTV ratio of 60 per cent in 2007 would have experienced an *increase* in consumption from 2007 to 2011 of 3.8 per cent of its income before tax in 2007. Hence, the difference in consumption growth from 2007 to 2011 between a family with an LTV ratio of 100 per cent in 2007 and a comparable family with an LTV ratio of 60 per cent in 2007 is estimated at -8.4 per cent of its income in 2007. Accordingly, if a family with an LTV ratio of 100 per cent in 2007 had an income before tax of kr. 500,000 in that year, the change in consumption from 2007 to 2011 would, all else equal, be approximately kr. 42,000 lower than for a comparable family with an LTV ratio of 60 per cent in 2007.

HOW SHOULD THE RESULTS BE INTERPRETED?

The above results can be interpreted to mean that the marked indebtedness among Danish homeowner families before the financial crisis contributed to amplifying the drop in private consumption during the crisis. One reason may be that a sudden tightening of credit conditions forced the highly indebted families to reduce their consumption more than other homeowner families. Another reason may be that increased crisis awareness and uncertainty about future financial conditions induced the highly indebted families to reduce consumption on their own initiative. Whatever the exact mechanism, according to this interpretation it was the *combination* of a high level of gross debt and a financial crisis that led to the substantial reduction of consumption among the families with high LTV ratios. The high level of debt among Danish households thus contributed to amplifying the effects of the financial crisis.

This interpretation of the results is supported by the fact that the relationship between the LTV ratio and the development in consumption is much more pronounced in the years around the financial crisis than in the preceding years, cf. Chart 4. In the left panel of the chart we consider the relationship between the LTV ratio and the subsequent development in consumption in the period *before* the financial crisis. The chart is constructed in exactly the same way as Chart 3, only with 2004 as the base year instead of 2007. To facilitate comparison of the two charts, the chart with 2007 as base year is presented in the right panel of Chart 4 on the same scale as in the left panel. There are clear differences between the two panels: Firstly, there is a distinct level difference between the the two sets of curves, reflecting that consumption growth was higher in the period 2004-08 than in the period 2007-11. More importantly, however, the slopes of the curves differ considerably in the two panels. Both panels show a negative relationship between the LTV ratio and the subsequent development in consumption, but the relationship is much stronger in the financial crisis period 2007-11 than in the period 2004-08. This indicates that the negative relationship between the LTV ratio and the development in consumption in the financial crisis period is related to the special economic circumstances during that period. The development in house prices may have played a particularly important role in that perspective: Families with high LTV ratios already in 2007 must, all other things being equal, have had *even* higher LTV ratios in 2009 due to the subsequent pronounced drop in house prices. This probably contributed to amplifying the drop in consumption. But in the period 2004-08, the development in house prices had the opposite effect: For families with high LTV ratios in 2004, the subsequent price increases contributed to lower LTV ratios and thus a weaker impact on consumption.

The contrast between the right-hand and left-hand sides of Chart 4 may also be seen as an argument against a number of other interpretations of the negative relationship between the LTV ratio and the subsequent development in consumption. For instance, one alternative interpretation is that the observed divergence in consumption development between families with high and low LTV ratios, simply reflects ordinary consumption and saving patterns: Families that, for some reason, had relatively high levels of consumption in the period until 2007 typically had high LTV ratios in 2007 for the

LTV ratio and subsequent change in consumption, 2004 and 2007 Chart 4

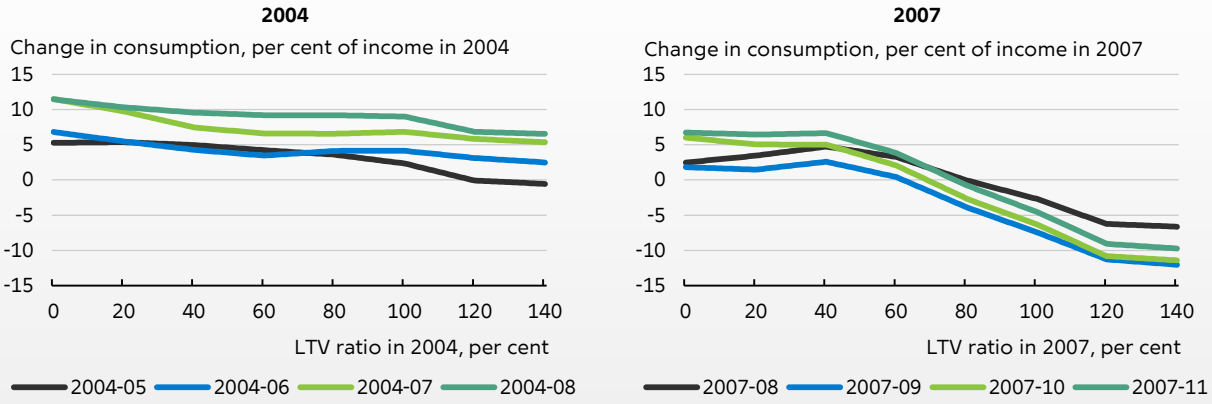

Note: See the note to Chart 3.
Source: Own calculations based on register data from Statistics Denmark.

same reason. After a period of high consumption financed by borrowing, their consumption has to decline again at some point (in relative terms), so the consumption growth of those families should typically be more subdued in the period after 2007 than that of other families. So according to this interpretation there is no direct linkage between the LTV ratio and the subsequent development in consumption. But if the observed relationship between the LTV ratio and the development in consumption were attributable to such normal, planned consumption patterns alone, such relationship should be just as pronounced in other time periods. Chart 4 shows that this is not the case.

DOES THE RELATIONSHIP EXIST AMONG ALL GROUPS OF FAMILIES?

The negative relationship between the LTV ratio in 2007 and the development in consumption in the subsequent years applies to all age groups, although it is somewhat less pronounced among the oldest families, cf. Chart 5 (left). A similar pattern is seen if families are grouped according to net wealth in 2007, cf. Chart 5 (right): Among the one fourth of the families with the largest net wealth, the relationship between the LTV ratio and the subsequent development in consumption is not quite clear. This group differs somewhat from the others, however: A relatively clear negative relationship is thus found among all other groups. It is particularly worth noting that the one fourth of the families with the smallest net wealth does not differ substantially from the other groups. This emphasises that the results in the preceding sections apply not only to the families that have borrowed so much that the value of their liabilities significantly exceeds the value of their assets. Similar conclusions are reached if the families are grouped according to e.g. geographical region of residence, income or holdings of liquid assets in 2007, meaning that the negative relationship is broad-based across various groupings.

This gives us some idea of the mechanism behind the observed relationship between the LTV ratio in 2007 and the subsequent development in consumption. If the weaker development in consumption among the highly indebted families were attributable solely to a tightening of credit conditions, the negative relationship should have been particularly evident among families with small holdings of liquid assets in 2007, since those families can be expected to be more affected by credit constraints than other families. That is not the case, however. The results therefore indicate that factors other than restrictions on the access to credit play a role. One possible explanation of the results is that events during the financial crisis led to widespread crisis awareness, inducing the highly indebted families to reduce their consumption to avoid further debt. Such crisis

LTV ratio and change in consumption 2007-11, broken down by age and net wealth in 2007 Chart 5

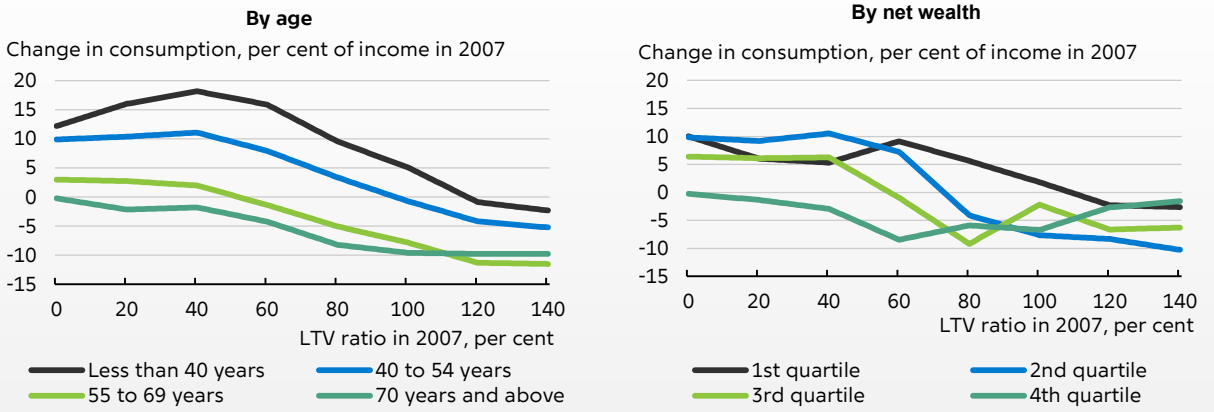

Note: The chart shows the average model-predicted changes in consumption over the period 2007-11 for different values of the LTV ratio in 2007 (see the note to Chart 3). In the left-hand side of the chart, the families are grouped according to the age of their oldest member in 2007. In the right-hand side of the chart, the families are divided into four groups of equal size, ranked according to the size of their net wealth in 2007. The families in the 1st quartile have the smallest net wealth, while the families in the 4th quartile have the largest net wealth.
Source: Own calculations based on register data from Statistics Denmark.

awareness may have affected a wide range of indebted families across differences in age, income and wealth, etc.

LTV RATIOS AND LEVELS OF CONSUMPTION

In the previous sections we have concentrated solely on the relationship between a family's LTV ratio immediately before the financial crisis and the *change* in its consumption pattern during the crisis. But it is also relevant to look at the relationship between the LTV ratio and the *level* of consumption before and during the crisis. Families with high LTV ratios at end-2007 generally had higher consumption-to-dispos-

Development in families' consumption ratios, broken down by LTV ratio in 2007 Chart 6

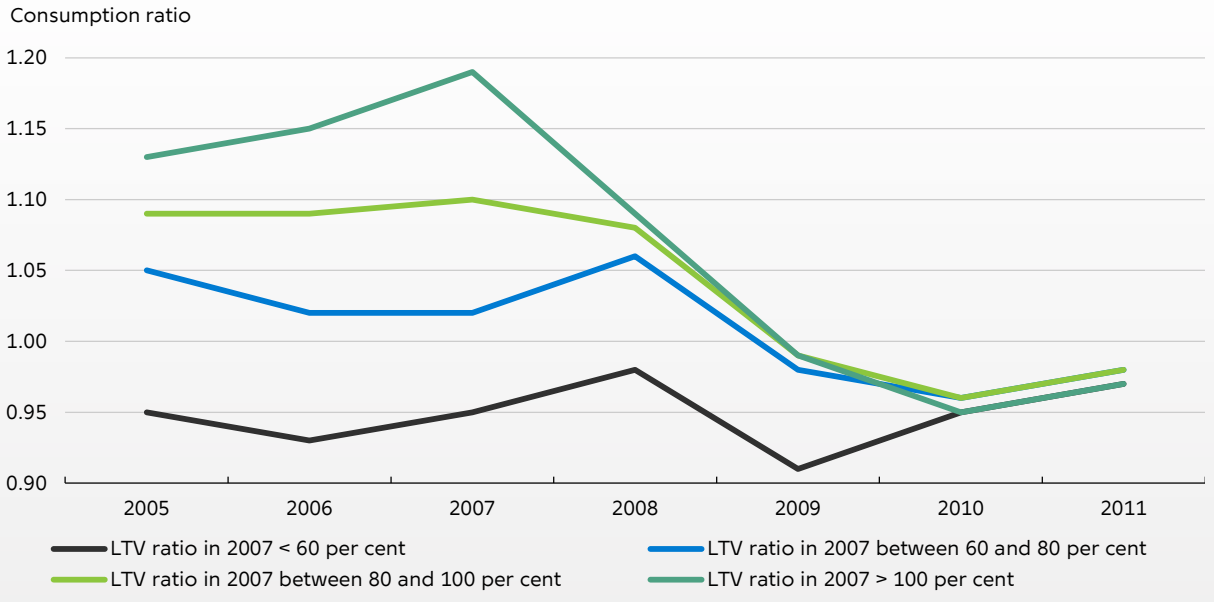

Note: The consumption ratio is calculated as consumption divided by disposable income. The chart shows the median consumption ratio within each group of families. The chart only includes families that existed in every year of the period 2005-11 and were not involved in property transactions during the period under review.
Source: Own calculations based on register data from Statistics Denmark.

able-income ratios before the crisis than families with lower LTV ratios, cf. Chart 6. As also shown in the previous sections, it was precisely those families who reduced their consumption the most when the financial crisis began. Looking at the level of the consumption ratio in 2010 and 2011, on the other hand, there is no clear correlation with the LTV ratio in 2007. In those years, the highly indebted families are at the same level as the other homeowner families.

So the results in this article only document a negative relationship between the LTV ratio before the financial crisis and the subsequent *change* in consumption. On the other hand, there is no direct basis for concluding that the *level* of consumption would have been higher in recent years if the indebtedness of Danish homeowner families had been lower before the crisis. Rather, the pattern in Chart 6 seems to indicate that the heavy indebtedness before the crisis contributed to a very high level of consumption in that period among some homeowner families. When the crisis erupted, those families reduced their consumption substantially, thus contributing to a marked decline in aggregate demand and activity in the Danish economy.

It is difficult to determine whether the observed reduction of consumption among the families with high LTV ratios should be viewed as an adjustment to an excessive level of consumption during the crisis, or whether it is in fact the lower level of consumption in the subsequent years that deviates from the norm. The analysis in this article provides no definitive answer to that question. One of the reasons is that the analysis only includes data from a single business cycle, thus making it difficult to predict the development of consumption among the highly indebted families in the coming years. The results in this article indicate, however, that a high level of debt among homeowners may contribute to larger fluctuations in consumption when the economy is hit by financial turmoil.

LITERATURE

Andersen, Asger Lau, Anders Møller Christensen, Nick Fabrin Nielsen, Sigrid Alexandra Koob, Martin Oksbjerg and Ri Kaarup (2012a), The wealth and debt of Danish families, Danmarks Nationalbank, *Monetary Review,* 2nd Quarter, Part 2.

Andersen, Asger Lau, Anders Møller Christensen, Charlotte Duus and Ri Kaarup (2012b), Danish families' financial robustness, variable rates and deferred amortisation, Danmarks Nationalbank, *Monetary Review,* 4th Quarter, Part 2.

Andersen, Asger Lau and Charlotte Duus (2013), Danish families in mortgage arrears, Danmarks Nationalbank, *Monetary Review,* 3rd Quarter, Part 2.

Andersen, Asger Lau, Charlotte Duus and Thais Lærkholm Jensen (2014), Household leverage and consumption during the financial crisis – evidence from Danish micro data, *Danmarks Nationalbank, Working paper,* No. 89/2014.

Browning, Martin and Søren Leth-Petersen (2003), Imputing consumption from income and wealth information, *The Economic Journal,* 113, No. 488.

Carroll, Christopher (1997), Buffer stock saving and the life cycle/permanent income hypothesis, *The Quarterly Journal of Economics,* 112, No. 1.

Dynan, Karen (2012), Is a household debt overhang holding back consumption?, *Brookings Papers on Economic Activity.*

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