Netspar

Network for Studies on Pensions, Aging and Retirement

Netspar THESES

Matthijs Triep

Losses in the Great Recession

Dutch Households under Fire

MSc Thesis 2013-016

Losses in the Great Recession:

Dutch households under fire

MATTHIJS TRIEP (1626345)

Master Thesis

7th of July, 2013

Abstract

This thesis attempts to quantify relative lifetime consumption expenditure losses of five Dutch household cohorts of different birth years in the Great Recession. Using a calibrated partial equilibrium model, where households are subjected to age-specific labour income shocks and general asset (wealth) shocks, recession scenarios and a normal economy benchmark are simulated and compared. As our five birth cohorts (selection is based on their age in the pre-recession year 2007) feature different expected wage paths, wealth balances and remaining life spans, distinct dynamics drive their economic performance. Our results show the oldest cohort, 65+ in 2007, to be robustly off the worst in a temporary recession of ten years, while the youngest cohort, younger than 35, comes in second. Intermediate cohorts suffer the least losses in various scenarios. For a permanent recession we find the youngest cohort to have the largest relative consumption losses. The oldest households in this case are harmed the least.

MSc. Economics (academic year 2012/13), VU University Amsterdam

Supervisors: dr. S. Hochguertel & prof. dr. B. Hobijn

Acknowledgements

I would like to thank dr. Stefan Hochguertel for his supervision during the production of this thesis, as well as his supervision during the research project that formed the basis for this work. His suggestions in the early stages of planning and for dealing with the datasets were very valuable, just like the feedback I received. I am also grateful that he contacted CentERdata for me, so I could access the data early.

My gratitude also goes out to prof. dr. Bart Hobijn, my second supervisor, for creating and programming the basic model, a feat I definitely would not have been able to do at this moment. Many of his suggestions during the skype calls we had also proved useful and have made their way into this thesis, for which I also thank him.

I would also like to thank CentERdata for allowing me access to the DHS database with its unique micro data, which my thesis contributed from substantively.

Finally, my thanks go out to my fellow students, with whom I have spoken about this thesis, who have helped me or kept me company while I was making this thesis. I would like to mention two students specifically: Roy Breurken, who also worked with the DHS data and who I could share data issues and code with, and especially Max Hikke, who made the long long summer of 2012 more bearable, when we were both working on our theses.

Introduction

Since the start of the Great Recession, the wealth distribution has once again become a topic of interest. Large individual and aggregate shocks to the economy have impacted household wealth and income.[1] From 2008 to 2011, average real household wealth in the Netherlands has dropped on average by 5.8%,[2] while annual disposable income on average decreased by 1.2% in real terms for 2009-2011.[3] Compared to a growth of 2.3% for the three years before, this is a 3.5% decrease.

But these figures only concern aggregate amounts, i.e. for all households at the same time. When we consider age cohorts, based on the age of the main earner, the picture looks a bit different. We find for the period 2008-2011 decreases of 14.3% of average real wealth for those aged 25-45, compared to only 2.5% for households of retirees (aged 65 or older). For real annual disposable income growth we can for the same household groups compute the difference between 2009-2011 and 2006-2008. These differences are -3.9% and -2.3%, which again shows dispersion.[4]

Households obviously can be divided differently, based on various characteristics. When we divide households according to wealth quintiles, we find mutations of mean real household wealth of 2% (1st quintile, the least wealthy), -20.9% (2nd quintile), -14.6% (3rd quintile), -7.8% (4th quintile) and -4.3% (5th quintile). Household wealth in different provinces can be measured as well. The largest drop, -9.5%, occurred in Flevoland, compared to -4.2% in Zeeland.[5]

Overall, these indicators are but an indication, but the message should be clear. Households cannot be treated alike. Naturally, they are the result of different characteristics, which co-determine economic performance. Retired households are not directly impacted by unemployment, which might explain their minor hit, as opposed to young working households that could be hurt by these effects, even for their whole lifetime. Economically weak groups, such as uneducated households could be expected to be hurt firmly, for a lack of quick adjustment possibilities. And of course household indicators are correlated with industries in their respective provinces and as these are hurt differently, so are households.

Birth year is a special characteristic though, because it is, together with gender, fixed for the whole lifetime. This means adjustment is special, for it will have to run through other channels. When considering heterogeneity with respect to age or birth year, recession implications for households can be expected to be different, both in the short and long run. Young households have little wealth to

[1] These include for example drops in stock prices, increased unemployment and a fall of real estate prices, which is by far the largest asset in Dutch household portfolios.
[2] Wealth here is defined as total assets minus total debt. Average wealth was calculated using aggregate wealth data from the Central Bureau of Statistics (CBS). Inflation adjustments were made using annual CPI from the CBS.
[3] Disposable income is taken from the CBS and is defined as income from labour (including business income) or wealth, minus certain premia and taxes on income and wealth. Inflation adjustments were made using CPI data from the CBS.
[4] Changes of the indicators for the other groups that the CBS provides in their basic wealth statistics, under 25 years old and 45-64, are between these extremes.
[5] Both of these subdivisions of households were also computed using the CBS aggregate wealth data.

smooth their short-run consumption, in contrast to older households. However, due to large asset price drops, younger households can benefit through cheap asset prices and future price increases. Labour income shocks form another issue, as growth or employment possibilities could depend on age. Finally, the remaining lifetime could feature cohort-specific dynamics. Young households have more time to smooth out recession effects, as their remaining lifetime is longer. These arguments are in no way the complete story, as there are also indirect effects, such as constraints to liquidity and portfolio choice. Household asset portfolios for example are highly dependent on age as younger households generally do not own, say, expensive houses, as they have not accumulated enough wealth for financing them yet. Portfolios may also impact household consumption decisions, through the returns it accumulates.

The different impact of the recession on households is highly relevant from the point of stabilization policy, which has once again become relevant and has recently attracted further attention in the media and Dutch politics. Who needs it the most is at this moment unclear. But this is not the only aspect, as stabilization costs have to be borne by someone, providing another intergenerational aspect.[6] This highlights the fact that quantifying the costs of the Great Recession are important.

To study household losses from the Great Recession we select five birth cohorts in 2007, the last year before a recession began to unwind, and simulate their remaining life from then on. Using a calibrated partial equilibrium lifetime household maximization framework we compute lifetime consumption statistics for recession scenarios with cohort-specific labor income and general asset (wealth) shocks and compare this with a benchmark non-recession alternative. Robustness of these results will be judged by considering two recession scenarios and various sets of household model parameters.

We find for a temporary recession of ten years the largest consumption expenditure losses for the cohort 65+, with a first place in more than half of the scenarios. The youngest cohort, below 35 years old, comes in second. The best off are middle-aged households at 45-54, while the cohorts in the intervals 35-44 and 55-64 are in the middle. When a lasting recession is simulated, these results change. We find the oldest cohort to perform the best, while the youngest cohort is now consistently the worst off. The other cohorts are now in the middle in different orders. For these results we find the length of the remaining lifetime and labour income shocks to be quantitatively important.

In Section 1 we will briefly review some literature that is relevant for analyzing welfare implications and household differences. Section 2 will provide some relevant stylized facts for a characterization of the household birth cohorts. Section 3 explains our model of household decisions and section 4 the possibilities for calibration. Section 5 discusses the simulations, after which in section 6 a conclusion and discussion of our results will be provided.

---

[6] This aspect is relevant due to the theoretically different implications of borrowing and taxation.

1. Literature review

1.1 Earlier attempts at quantifying welfare losses from business cycle fluctuations

Whereas the wealth distribution has only recently returned in the spotlight, wealth implications from recessions and contractions have received more attention. A fundamental paper has been the approach by Lucas (1987). He analyzes the impact of removing pure variability of consumption by modeling a lifetime utility maximization framework of the form:

$$\sum_{t=0}^{T} \beta^t \frac{(C_t^{1-\theta} - 1)}{1-\theta}, \qquad \theta > 0 \tag{1}$$

Here β is the discount factor, θ the coefficient of relative risk aversion, C_t annual consumption at time t and T is the moment of death. As the CRRA specification of the utility function of equation 1 induces the consumer to smooth his consumption, an assumption of the lifecycle hypothesis, there is a gain from eliminating all variability of consumption. This gain can be shown to be equal to $\frac{\theta}{2}\bar{C}^{-\theta-1}\sigma_C^2$, where $\bar{C}$ is average annual consumption and σ_C^2 the variance of annual consumption (Romer, 2012, p. 529).[7] Based on realistic figures for the coefficient of relative risk aversion (between 1 and 4), the gain from removing variability for the average US consumer would be between 0.05% and 0.2% of annual consumption (Lucas, 2003), which seems negligible.[8] This is supported by the calculation of Imrohoroglu (2008), who computed the stabilization gains would amount to approximately 29 dollars per person per year in the United States.

The negligible estimates of stabilization gains that have been found by Lucas may be tempting to conclude that the relevance of stabilization policy is minimal. Household losses would then be insignificant over the business cycle, i.e. nothing more than a temporary fluctuation that brings along minor costs. Lucas (1987) himself, in fact, considers these results to be evidence that avoiding large policy mistakes is what matters most, given the small upside potential.

But what does this finding say about the costs for heterogeneous cohorts? First of all, Lucas's analysis is concerned with aggregates and therefore concerns average households. Could household cohorts be analyzed in the same way as the aggregate household? Theoretically, the answer is yes. As long as we can extract the required household parameters of cohorts, as well as the distribution of consumption, similar analysis could be used to calculate costs.

However, there is another point of interest. The assumptions of the Lucas model are in fact very strict. As has been summarized by Imrohoroglu (2008), these assumptions are important for the

---

[7] A formal proof is provided in appendix A.

[8] For comparison, Lucas (2000) has shown that a reduction of the growth rate of consumption of one percent point (conditional on 3 percent points growth) and 10 percent points inflation bring costs of 20% and 1%, respectively, of annual consumption for a representative US consumer.

magnitude of the costs of business cycle fluctuations that have been found. According to her, violation of several assumptions of the Lucas model can be shown to open up the possibility for indirect effects to impact the costs. When we consider heterogeneous household cohorts these indirect effects may have even more effects on welfare costs as indirect effects may be cohort-specific, potentially biasing the gains of stabilization.

Lucas assumes for example that all consumers can perfectly insure/smooth away fluctuations due to the existence of perfect capital markets. Insurance markets in general have been found to be incomplete for most consumers due to a lack of access to capital markets and/or liquidity constraints (Imrohoroglu, 1989). This means we can expect welfare costs to be higher in the absence of perfect insurance markets. Heterogeneity is also an aspect here, as different possibilities to insure, among household cohorts, may not be unlikely. Wealthy households may be able to reduce their costs through their large amount of wealth, but for poor consumers, who are constrained in their liquidity or restricted from investing in certain assets, such shocks may be large. As older household cohorts are generally more wealthy than younger cohorts, as we will see, this could be highly relevant for dispersion of welfare losses. Other indirect effects can be expected to materialize. We could basically think about any determinant of income and wealth that is not unique to cohorts. Examples are education level (high level of education gives more flexibility when getting unemployed), household composition (households with more earners or less people to feed are better off), but also geographical aspects (living in economically significant areas) and more importantly, the persistence of shocks.

Several academics have estimated welfare costs using Lucas style setups that account for several aspects that are ignored by the original, but results are mixed. Some researchers find higher costs for certain poor groups, though differences are generally small, whereas others find the costs of business cycles in general to be more significant compared to what Lucas finds (Imrohoroglu, 2008). These costs may subsequently carry over for heterogeneous cohort case, i.e. that they are also significant, although perhaps different.

1.2 More refined techniques for quantifying welfare losses

The models we have considered until now are static, they largely ignore the influence of markets and thus household interaction. The Lucas argument for stabilization benefits thus does not deal with equilibrium effects in markets. Any effects on prices or demand and supply are basically exogenous for the consumer. Although such an assumption is convenient for analysis, realism is affected significantly. If consumers massively sell their assets at depressed prices for smoothing, prices will go down. Similarly, lower consumption impacts businesses and therefore also the labour market. Financial markets and government may show an additional response. Higher order effects will happen due to the successive shocks in all markets and modeling a recession should in that case be done using a general equilibrium framework. Such a setup is not only better suited for such a purpose, it also is from the point of the indirect effects for households that were discussed. Cohort specific effects can theoretically be tracked much better from this point of view. However, this interdependence might lead to effects being hard to separate.

One attempt at using such a framework was made by Glover, Heathcote, Krueger and Rios-Rull (2011). Using a calibrated general equilibrium framework (with overlapping generations), featuring American households with heterogeneous age profiles (using cross sectional data), a labour market and financial markets, they have simulated aggregate shocks. Their models centers around two main sources of heterogeneous household elements. Firstly, wealth is concentrated among older households and secondly, labour income shocks are more severe for young cohorts. This gives larger possibilities to smooth for cohorts that either own more or earn more, generally older households. Secondly, recession exposure varies with respect to both asset shocks and income shocks. As the model concerns a general equilibrium, several risks can be shared with other generations. However, aggregate shocks have an asymmetric effect as it is impossible to share risk with future or past generations (Glover et al., 2011).

By simulating they find using several setups and recession and non-recession regimes that older household cohorts suffer from higher welfare (utility) losses than young households for a fierce, but temporary recession. In their model shocks to asset prices were found to be larger than shocks to aggregate output (this depends on the strength of the smoothing motive), so assets were sold at large discounts. As older households owned disproportional amounts of assets, wealth losses were consistently increasing in the age of households. Younger households conversely suffered from disproportional losses in earnings, but benefitted from asset fires sales through future price increases, as the economy returned to normal times. However, they explain the total effect is not merely due to asset prices, as the differences in remaining lifetime are also a factor. Separate setups indicate that both asset prices and remaining lifetime effects are quantitatively important (Glover et al., 2011). Their simulated scenarios have thus shown the use of a general equilibrium framework; the fundamental aspect of interdependence within the economy of households and other markets and sensitivity with respect to these forces.

1.3 More academic work

Other academics have also taken the Great Recession as an opportunity to study channels of impact on the wealth distribution. Perri and Steinberg (2012) have looked at the inequality arising from the recession and the result of government redistribution policies for US households. They find the losses to have been distributed unequally; the bottom 20% of households in terms of earnings has never been as far behind the median in the post-war period. However, they claim that in terms of consumption or disposable income the position of this 20% has not changed significantly, primarily due to government policy. Still, Perri and Steinberg (2012) consider the wealth losses that have occurred next to the earnings drop for these households a future source of vulnerability, i.e. for the longer term. Higher order effects could thus further impact those in specific parts of the wealth distribution.

Among the other research about income is the work of Guvenen, Ozkan and Song (2012) about counter-cyclical income risk in the United States. They state that large wage increases become less likely during recessions, while large drops become more likely. Additionally, they find median shocks to be rather stable, i.e. they do not depend on the business cycle. These findings will no doubt impact certain cohorts differently, while also implying further distributional consequences.

Consumption implications have also received a deal of attention. Consumption is important as it is both a measure of wellbeing, but also a large determinant of GDP and thus important for higher order effects. For individual households higher order effects of course also exist. Petev, Pistaferri and Saporta Eksten (2012) have found short-term redistributive effects in consumption. They find a decline in consumption inequality, as well as a decline in mobility of the consumption distribution. The sources they attribute this to are the disproportional drop in assets for richer households, as well as increased uncertainty. Similarly, De Nardi, French and Benson (2011) have found the negative wealth effect to be the explanation for the drops in consumption that have been found, though they also report consumer income expectations (obviously related to uncertainty) .

Summarizing, the quickly increasing volume of academic work about the implications of the Great Recession and other recessions point to the interrelatedness of household characteristics and the effects these may have on the wealth distribution. We see that welfare costs of business cycles are generally considered to be related to the household budget and how it is spread out over the lifetime, i.e. according to its preferences. The budget however is not fixed though as it is dependent on asset shocks, wage shocks and household portfolios, among possibly other effects, which may all have higher order effects for the general economy. Expectations of these shocks have their own dynamics, for households do not know with certainty what the future will bring. For young households this uncertainty is the largest, for they have a longer period to go, whereas old cohorts, in turn, have less flexibility. Based on the research of previous academic work, differences between cohorts may thus arise, but ultimately who bears most costs depends on which effects dominates.

2. Data

2.1 Intro

In this section, we will provide some stylized facts about the endowment of Dutch households in the context of the household maximization problem and the Great Recession. Although the aspects of lifecycle savings and subsequent principles related to consumption, such as smoothing, are well recognized in the literature, we do not know with certainty how relevant it is for the Dutch household cohorts we wish to study. We are mainly interested though if we can find the shocks that Glover et al.

(2011) used as their main drivers (wealth and labour income shocks) for analyzing wealth effects are present for household cohorts in the Netherlands.

2.2 Data sources

For data analysis we have the availability of two sources of household data. On the one hand we have cross-sectional aggregate wealth and income data from the CBS, which allows separation into households cohorts based on age of the head of the household. On the other hand we have a detailed micro dataset from the DNB Household Survey (DHS). This is an unbalanced panel survey (roughly 2000 households), organized and managed by the research institute CentERdata at the university of Tilburg, and has been held annually since 1993. The panel structure of the data allows us to track households, and thus cohorts, over time, although the sample size is very limited.

Our two data sources complement each other in various ways. The aggregate CBS data has a smaller sampling error due to its large sample size, but suffers from issues, besides being cross-sectional. First, several types of assets and debt of households are not represented, or only partially (CBS, 2010). As the CBS explains this has to do with the methods used to extract household data, which is mainly from tax returns. The DNB micro data has a only a limited sample size, but it includes several types of assets and debts that are not included in the CBS data, because figures are being reported by households themselves.

For our analysis we have selected five household birth cohorts, based on the age of the head of the household. The choice of birth cohorts allows us to make a verdict with respect to who bears the largest total costs based on the remaining lifetime of cohorts. This is relevant as these cohorts enter the recession with different ages due to their birth date. We choose for the following cohorts: younger than 35, 35-44, 45-54, 55-64 and 65+, all taken from our base year 2007. In terms of birth years we thus deal with cohorts that are born before 1943, between 1943-1952, 1953-1962, 1963-1972 (cut-off points are inclusive) and after 1972. The reason we opt for choosing these categories are related to the availability of data. Ideally, we wish to space our cohorts equally, i.e. all 10 years in length, 20-29, 30-39, etc., but this is problematic. In the micro dataset young (below 30) and old (70+) households are underrepresented, which would give comparability issues. This selection reduces this problem. Other advantages of these intervals are comparability with the CBS data, which can be split according to 5 year intervals (25-29, 30-34, etc.), and the fact that a composite cohort with both retirees and working households is avoided.[9]

[9] We assume households retire at 65, the official retirement age for our sample period.

2.3 Wealth and income statistics

We will provide some basic indicators of household wealth and income for our base year 2007. For wealth, we would ideally compute an indicator that comprises everything the household owns, including the value of future income, human wealth. Measuring this, however, is practically impossible, for valuation principles require extensive knowledge about uncertainty with respect to discounting future values. Therefore we will provide marketable wealth, i.e. what can be sold at this moment. Defined as the total value of assets that can be sold in the market minus the debts, it is the most commonly used version of net wealth (Davies and Shorrocks, 1999, pp. 606-607 & 629). Our measure of wealth will therefore comprise of real property, financial assets and so on.

To compute marketable wealth we aggregate per household in the DNB data set all the asset accounts that are available and subtract all debt accounts. This ranges from various savings accounts, financial assets (stocks, bond, options, etc.) and real property to assets like cars and other savings. For an approximation of total debt we aggregate all debt accounts, which consists of mortgage debt, study loans, consumer credit and other debts. A detailed summary of specific asset and debt accounts is provided in appendix B and data construction for the indicators in appendix C.

For annual labour income we consider the micro dataset too noisy to compute a reliable household average (missing values, unclear tax accounts, household aggregation issues, etc.). Therefore we take the cross-sectional income data from the CBS as an approximation for our cohorts (i.e. cross sectional data but for households of the same age as our households of interest). From disposable income (includes taxes and premia) we subtract CBS asset income (incomplete as it lacks capital gains (CBS, 2013)), and we call this labour income. For a better approximation of asset income, we will use the micro data set.[10] We aggregate per household income from real estate, interest income, income from stocks and lower this by the amount of total paid interest. Appendix C contains a small description of data construction issues. For capital gains we take a price index of residences from the CBS and acquire, per year, the monthly (first day of the month) stock price of the AEX index from Datastream. We average these stock prices into a yearly average.[11] The annual change of these price indices are multiplied by the household stocks of real estate and equity. The sum of asset returns and capital gains is then called asset income.

Table 1 reports panel data statistics about mean and median household wealth values for our birth cohorts, as well as the composition in terms of assets and debts for 2006-2011. Although our analysis is concerned with 2007 and onwards, we have also included 2006 so we can compare changes of the recession period with that of normal times. Ideally, we would provide statistics for a longer timeframe,

---

[10] We will use asset income (shocks) and wealth income (shocks) interchangeably. Technically wealth income consists of asset minus debt income, but asset income is much more intuitive due to its interpretation as investments. Debt is then interpreted as a negative asset investment.

[11] This is preferred to using the stock price of the first trading day in January as extremes are averaged out this way.

Table 1: Descriptive wealth statistics (panel) of household indicators 2006-2011

This table gives an overview of important summary statistics about wealth for various Dutch household birth cohorts for 2006-2011, calculated using micro panel data from the DHS survey. Age is based on the age of the head of the household in 2007. Figures are in thousands of euro and are expressed in constant 2011 prices. Price levels have been deflated using CPI data from the CBS. Appendix B provides a detailed summary of assets and debts that have been included, appendix C deals with data construction issues.

Age	Indicator	2006	2007	2008	2009	2010	2011
Total	Wealth (mean)	208	217	221	211	216	212
	Wealth (median)	143	148	150	151	150	161
	Assets	278	297	297	293	299	288
	Debts	70	80	76	82	83	76
<35	Wealth (mean)	61	54	45	54	83	85
	Wealth (median)	22	23	22	19	26	40
	Assets	141	162	167	182	229	220
	Debts	80	108	122	128	146	135
35-44	Wealth (mean)	131	136	143	131	131	129
	Wealth (median)	81	92	96	101	109	93
	Assets	223	243	248	261	253	240
	Debts	92	107	105	130	122	111
45-54	Wealth (mean)	187	202	194	194	185	183
	Wealth (median)	147	165	137	148	139	138
	Assets	275	303	281	280	277	263
	Debts	88	101	87	86	92	80
55-64	Wealth (mean)	282	304	304	266	284	278
	Wealth (median)	214	229	225	217	219	205
	Assets	343	365	369	328	347	338
	Debts	61	61	65	62	63	60
65+	Wealth (mean)	302	316	297	295	289	275
	Wealth (median)	254	258	236	253	253	254
	Assets	341	356	335	341	330	315
	Debts	39	40	38	46	41	40

such that multiple business cycles are included for a more stable long-term pattern, which has less sensitivity to annual idiosyncrasies. The reason for starting from 2006 is due to comparability issues. Some of our birth cohorts in the micro dataset have only 200-300 household observations, while representing over a million of households (source: calculated from CBS statistics). For this reason we cannot ignore to check for robustness by comparing with other sources, in our case the CBS data. The micro dataset has data starting from 1993. Since the CBS data is only available from 2006 and onwards we have chosen to match this timeframe for the DHS data.

As wealth is generally concentrated among a small amount of rich households, the mean could be misleading. Medians are therefore also provided. The principle of the concentration of wealth is supported by the fact that in table 1 medians are lower than means for all cohorts. The sample sizes for the specific cohorts range between roughly 200 and 300 observations, totals comprise 900-1300 observations, both of which are low. Therefore we provide a table of comparable CBS cross-sectional data as well for sensitivity, which is table 3.

Table 2: Descriptive income statistics of households, 2006-2011

This table reports mean values of annual household income for age cohorts for 2006-2011, calculated using micro panel data and aggregate cross-sectional CBS data. Labour income is calculated by subtracting CBS asset income from CBS disposable income figures. Asset income consists of household asset return and capital gains, calculated with the micro data and price indices of real estate and stocks. Age categories are based on the head of the household, but do not coincide with our selected cohorts perfectly (except for 2007). Figures are in thousands of euro and are expressed in constant 2011 prices. Price levels have been deflated using CPI data from the CBS.

Age	Type of income	2006	2007	2008	2009	2010	2011
	Labour	31.9	32.8	34.3	34.6	34.7	35.0
Total	Asset	9.9	8.7	4.9	-14.4	-2.8	-5.8
	Total	41.8	41.5	39.2	20.2	31.9	29.2
	Labour	27.0	27.9	28.6	28.6	28.1	28.0
<35	Asset	2.0	0.9	-1.4	-10.4	-9.1	-8.7
	Total	29.1	28.8	27.2	18.2	19.0	19.3
	Labour	38.1	39.9	41.5	41.8	42.5	42.7
35-44	Asset	4.7	3.5	1.3	-13.0	-7.0	-8.2
	Total	42.8	43.3	42.9	28.8	35.5	34.6
	Labour	40.3	41.7	43.4	43.4	43.5	43.8
45-54	Asset	8.2	8.1	3.6	-13.8	-3.3	-6.4
	Total	48.5	49.8	47.0	29.6	40.1	37.4
	Labour	33.3	33.9	35.9	36.5	36.8	37.5
55-64	Asset	13.1	12.1	7.1	-14.1	-3.0	-5.2
	Total	46.5	46.1	43.0	22.5	33.8	32.3
	Labour	21.8	22.3	23.7	24.0	24.7	25.4
65+	Asset	15.5	13.7	8.5	-17.5	2.5	-3.7
	Total	37.3	36.0	32.2	6.5	27.2	21.6

When we compare cohorts by year in table 1, we immediately see the large dispersion in terms of wealth. The oldest two cohorts of our selection are by far the most wealthy, though it is not obvious where the peak is located. In the empirical literature, this peak is generally located at the cohort that is closest to retirement (Davies & Shorrocks, 1999, pp. 12), i.e. 55-64 in our case. Except for the uncertainty of the peak, all annual wealth averages show that wealth is increasing in the age of the cohorts. Younger households have likely not been able to build up large stocks of wealth, which is shown in this table. These findings highlight the main reason for the thesis. By looking at total average household wealth, the representative household would be located somewhere within the cohort aged 45-64 (likely somewhere close to 50-55). This ignores the younger generations, especially the youngest households within the younger than 35 cohort.

Now what does the table tell us about the effects of the Great Recession? First of all, four of the five cohorts have seen their wealth drop below their 2008 value in subsequent years when we look at the mean wealth levels. The youngest cohort is in this sense an outlier, for it shows strong growth in wealth. For the years 2010 and 2011, years of small recovery in terms of GDP growth,[12] we would normally expect an increase of wealth compared with 2009. Yet, only two cohorts show this to hold.

[12] Source: CBS economic growth accounts.

Table 3: Descriptive statistics of household indicators, 2006-2011 (CBS)

This table gives an overview of the magnitude of average wealth held for various Dutch household age cohorts for 2006-2011, calculated using aggregate CBS data. There are two types of debt reported, which has to do with the shortcomings of the CBS data. The CBS has only been able to extract data about debt for a subset of the households for which they have data about assets. The average debt for these households is represented by Debts II. Extrapolation of these averages to the full sample for calculating wealth has not been performed by the CBS and instead, their average wealth figures have been calculated by dividing aggregate wealth by the number of households. To keep the notion of wealth as assets minus debts we provide the indicator Debts I, the difference between mean wealth and mean assets. Age categories are based on the head of the household, but do not coincide with our selected cohorts perfectly (except for 2007). Figures are in constant 2011 prices and have been deflated using CPI data from the CBS.

Age	Indicator	2006	2007	2008	2009	2010	2011
Total	Wealth	173	185	190	178	170	159
	Assets	261	281	288	280	273	261
	Debts I	88	96	99	102	103	102
	Debts II	183	190	196	201	206	204
<35	Wealth	35	34	37	29	22	18
	Assets	120	126	130	123	114	103
	Debts I	86	91	93	93	91	85
	Debts II	208	215	218	223	223	215
35-44	Wealth	127	137	136	124	106	93
	Assets	266	280	289	282	269	255
	Debts I	138	143	153	158	163	162
	Debts II	219	221	234	240	247	246
45-54	Wealth	202	217	219	210	200	174
	Assets	313	340	349	341	335	311
	Debts I	111	123	130	131	135	137
	Debts II	179	195	206	209	216	219
55-64	Wealth	283	302	305	282	266	257
	Assets	370	396	397	378	365	356
	Debts I	87	94	92	96	99	99
	Debts II	159	168	164	172	178	175
65+	Wealth	236	251	264	252	253	245
	Assets	263	280	293	286	290	284
	Debts I	27	29	29	33	37	39
	Debts II	109	111	110	119	127	131

The median wealth level does not complement these observations as the median wealth levels for four of the cohorts have risen above their 2008 value, the only exception being those aged 55-64. If we look at change from 2009 to 2010 and 2011 we have a very much mixed image. The youngest cohort has growing wealth holdings, 45-54 and 65+ drop monotonically, while 55-64 and 35-44 only drop in 2011. The volatile and contradictory median wealth levels may be explained by two reasons. One, there are large swings in the distribution that is behind the aggregate statistics or two, data quality issues. Of these, the latter is more plausible, for increases in median wealth indicate large groups of household to have increased their wealth. In a recession this seems unrealistic.

If we see these patterns (though unclear) in terms of the lifecycle principle, we could interpret decreases of wealth as smoothing by households, i.e. they have sold asset or borrowed for stable consumption. An increase in debts in fact is noticeable for three out of five households when we look at the difference between 2008 and 2009 and/or 2010. For 2011 all debt stocks have decreased again, which could be consistent with smoothing being less of an issue, following some recovery. Asset stocks however are inconsistent if we look at 2009 figures as they have mostly increased whereas this was a severe year in terms of GDP growth (i.e. -3.7% on annual basis).[13]

Now if we look at household income in table 2, we see that during the Great Recession the asset income part of total income has decreased substantially for all cohorts, to the point that losses have since 2008 persistently occurred. From the micro data we can see that the largest part of these losses are related to the main asset for households, their residence. Labour income in turn has not decreased and is still growing, except for the youngest cohort. Although the evidence is small, this suggests that young households are subjected to larger labour income shocks. For total income, we can see the same pattern as asset income dominates shocks, due to the relative stability of labour income.

Overall, the wealth statistics we have computed from the micro data do not show any clear pattern of movements in the Great Recession, but wealth and income shocks can be identified. This could be explained in two ways. Firstly, if households are perfect smoothers, the dynamics of the household budget (shocks to assets) and smoothing motives provides a volatile pattern. If we look at the income data we see declines in income, which are a reason to smooth. The asset stock patterns are not fully consistent with this however as shocks occur at odd moments, i.e. contradictory. Secondly, the quality of the indicators is simply too low due to the small sample size and perhaps other reporting issues (can households provide a reliable asset estimate themselves?). To see whether our micro data is at least somewhat reliable we compare the patterns to those from table 3, which contains the cross-sectional CBS data.

From the mean wealth values for the age groups in 2008-2011 in the CBS data we see a more clear pattern of wealth changes. For all of these cohorts (they largely coincide with our cohorts, but not perfectly) it has in fact decreased monotonically. Assets show the same pattern, a (near) monotonic decrease for all households. Debts have moved to the other side, they have increased since 2008.

Figure 1 shows for both of our sources the pattern of wealth holdings for our households, with 2007 being indexed at 100. The averages of the micro data show that the youngest cohort has enjoyed serious wealth growth in the years after 2008, whereas the other cohorts have not reached their 2007 levels (except the cohort 35-44). Relative wealth losses for the other cohorts provide not much of a clear pattern, although the three oldest cohorts have slightly higher losses. From the CBS data we see that relative wealth losses are in fact increasing in age. Although the levels of indicators from the CBS data are not completely reliable due to data being unavailable and not matching our cohorts perfectly, their statistics are probably a better indicator of wealth holdings of our household cohorts, due to the much higher sample size. For the few years after the start of the reduced growth period from 2008, younger cohorts have in the short term a bit larger relative losses, though the micro data provides

[13] Source: CBS economic growth accounts.

Figure 1: Wealth changes during the Great Recession

This figure contains real mean wealth values for five birth or age cohorts, based on age of the head of the household. The left panel shows the aggregate of the DHS micro data (birth cohorts), the right panel a constructed average of aggregate CBS data (age cohorts). Wealth values have been indexed at 100 in 2007. For the CBS data the cohorts are not completely the same as the cohorts we have selected as our cohorts of interest, but they are very much correlated. The micro data shows declines in wealth after 2008, though not monotonically. The youngest cohort mysteriously enjoyed spectacular growth. We expect this to be explained by a low sample size. Therefore the CBS data is also provided. It shows wealth losses for all cohorts following 2008, but the youngest households are hurt the most.

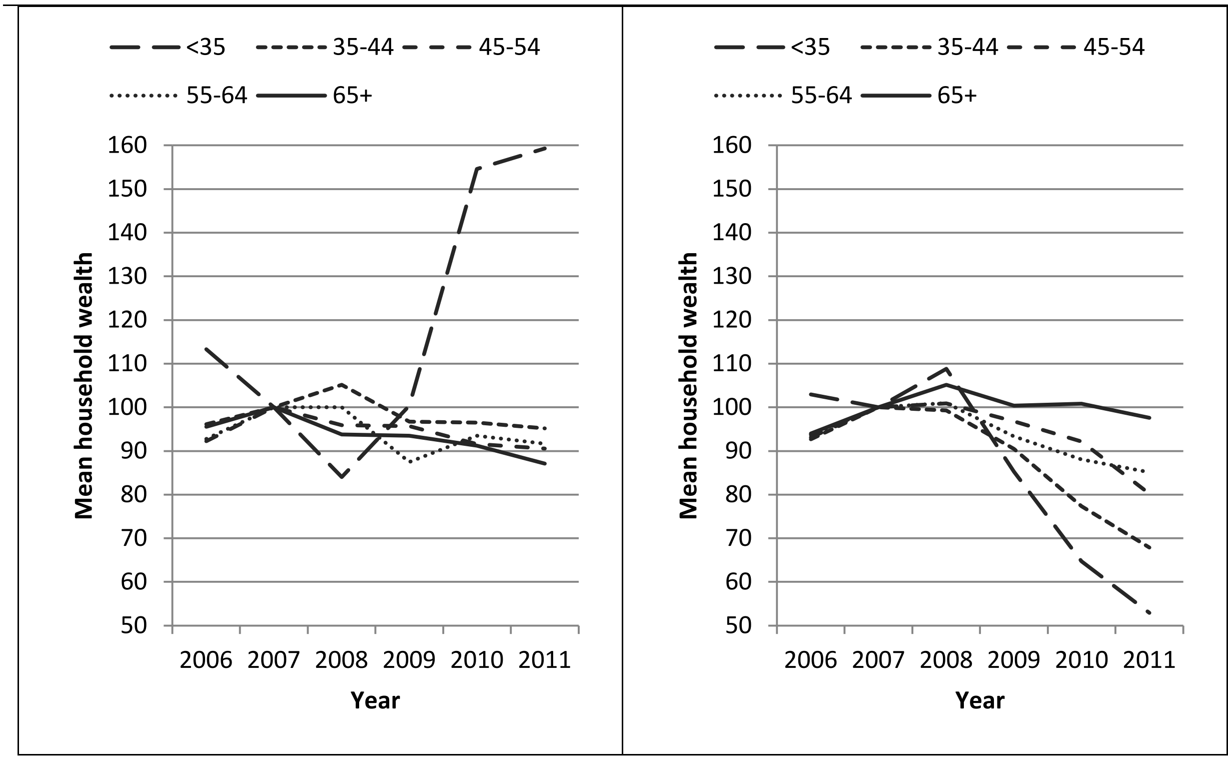

conflicting evidence. These losses are only for the short term though. For total lifetime effects we need simulations, which is what will be discussed in the next sections.

3. Model: Household maximization problem

3.1 Mathematical household problem

For analyzing costs to household cohorts we need to be able to quantify losses. For this we will use a standard household maximization problem, where agents maximize their lifetime utility, given their budget restriction. To keep the analysis relatively simple, we model the framework as a partial equilibrium model. Such a model is not able to capture effects of several markets simultaneously, such as the labour market and the asset market, but avoids distortion by a multitude of effects.

We assume that the household is driven by either consumption expenditures and/or leisure time. This means that consumption expenditures and labor supply will have to be set jointly to account for the fact that the budget for consumption is dependent on the amount of hours worked. For analyzing the effects of the recession, labour supply is not really necessary, for we could look at pure consumption effects and assuming full-time employment. Labor supply can be used though for proxying unemployment effects or a dislike for working. This is especially interesting because our households are of different ages, meaning shock absorption by working more cannot be used equally. Valuation of consumption and leisure time are assumed to be determined by the instantaneous utility function:

$$U(C_t, L_t) = \frac{\sigma}{\sigma-1}\left[C_t^{\theta}(1-L_t)^{1-\theta}\right]^{\frac{\sigma-1}{\sigma}} \qquad \text{for } \sigma \neq 1 \tag{2}$$

Here $C_t > 0$ concerns annual consumption expenditures at time t, $0 \leq L_t \leq 1$ represents labour supply as a fraction of total available time at time t, $\sigma > 0$ is the intertemporal elasticity of substitution (IES) and $0 < \theta \leq 1$ the relative preference for consumption. For $\sigma = 1$ it can be shown that the instantaneous utility function is equal to log utility:

$$\ln\left(C_t^{\theta}(1-L_t)^{1-\theta}\right)$$

The motivation for choosing this type of utility function has to do with its flexibility. From equation 2 we see it is basically a Cobb-Douglas function, nested within a CRRA type function. The Cobb-Douglas type function allows us to vary the strength for consumption and labour supply and does so by increasing the marginal effect through θ. By letting $\theta \rightarrow 1$ we can eliminate the labour supply from the utility function, effectively allowing us to study pure consumption decisions. The CRRA properties are required to force the household to smooth its consumption. The strength of the smoothing motive is partially governed by σ as it co-determines the shape of utility curves, along with θ. More concave shapes will make the household prefer more stable paths of consumption.

The household is now expected to maximize lifetime discounted utility given that it cannot spend more on consumption than its lifetime budget and has to allocate it over the remaining lifetime. Therefore we assume utility is additively separable and that a household dies after T periods. The household maximization problem can then be written as

$$\max\ V_0 = E_0\left[\sum_{t=1}^{T} \beta^t U(C_t, L_t)\right] \qquad \text{with respect to } C_t, L_t \qquad t = 1, \dots, T$$

and subject to

$$\sum_{t=1}^{T} \frac{P_s C_s}{(1+E[r])^t} + \frac{A_T}{(1+E[r])^T} = \sum_{t=1}^{T} \frac{W_t L_t}{(1+E[r])^t} + A_0 \tag{3}$$

$$C_t > 0 \qquad t = 1, \dots, T$$

$$0 \leq L_t \leq 1 \qquad t = 1, \dots, T$$

V_0 refers to the present value of lifetime expected discounted utility, β represents the discount factor of utility and is constant. P_t and W_t are the price level of a unit of consumption and nominal wage per unit of labour supply at time t, respectively. A_0 concerns initial wealth and A_T is terminal wealth. For simplicity we assume that the nominal wage and price levels are exogenous and deterministic, so the household knows them advance. $E[r]$ is the expected return as return is assumed to be uncertain. The budget constraint is flexible in the sense that initial wealth A_0 and terminal wealth A_T can be set to 0 if we think this is desirable, just like the price and wage levels can normalized. As we assumed that households are only interested in consumption and leisure, bequests do not have to be modeled and therefore we will set $A_T = 0$.

Alternatively, the budget constraint from equation 3 can be written in recursive form as:

$$A_{t+1} = (1 + E[r])(A_t + W_t L_t - P_t C_t)$$

By writing the budget constraint as a recursion between periods (intertemporal), we can gain additional insight compared to simply stating that lifetime consumption expenditures cannot be larger than lifetime income. A_t is the stock of net wealth at time t that is left from the previous period and can be used for consumption in the current period. Positive values for A_t means the household has positive wealth, negative means the household is at present indebted overall. We assume that at time t it saves and invests an amount A'_t for an uncertain return r, for the next period. This gives an intra-temporal budget constraint of:

$$A'_t = A_t + W_t L_t - P_t C_t \quad (4)$$

The choice of A'_t now determines, together with consumption expenditures, C_t, and labor supply, L_t, the future path of A_t.

For the return for the savings of the households we assume a discrete probability distribution. Returns from this distribution are drawn from a conditional density function that determines the probability of a certain return:

$$f(A_{t+1} | A'_t) \quad (5)$$

Equation 4 exposes a problem though, since the household could theoretically borrow an infinite amount, supporting infinite consumption and therefore utility. We wish to rule out such decisions, since they are unrealistic, for no lender would lend such an amount. This means that we need to add a credit constraint. We assume that households cannot borrow more than the present value of their future labor income. Therefore we have to introduce the constraint:

$$A'_t \geq \mu_t = -\sum_{s=t+1}^{T} \beta^{s-t} W_s \quad (6)$$

Here we assume that the interest rate is set equal to $\frac{1}{\beta}$ and we will also enforce this in the general model. We also wish to restrain households from dying indebted, since this is unrealistic (nobody would lend money to someone who dies and does not pay back) and therefore we also set $A_T' \geq 0$. For indebted households, we fix the interest rate such that $\frac{1}{1+r} \equiv \beta$, so the cost of shifting values through time is equal to the subjective cost of shifting utility through time. Households will then consume their budget in constant installments. This closes the model for indebted household, but since our analysis is focused on households with positive wealth (all cohort means are positive) and that smooth when negative shocks occur, we will not consider households with negative wealth.

3.2 Solution to the household maximization problem

We let V_t be discounted lifetime utility and a function of A_t, wealth at time t. We can then state the maximization problem using a Bellman equation:

$$V_t(A_t) = \max_{C_t, L_t, A_t'} \left[U(C_t, L_t) + \beta E_t[V_{t+1}(A_{t+1})]\right] \tag{7}$$

subject to

$$A_t' = A_t + W_t L_t - P_t C_t$$

$$0 \leq L_t \leq 1 \qquad t = 1, \dots, T$$

$$A_t' \geq \mu_t = -\sum_{s=t+1}^{T} \beta^{s-t} W_s$$

$$C_t \geq 0 \qquad t = 1, \dots, T$$

$$A_T = 0$$

$$A_T' \geq 0$$

From equation 7 we can read that finding the solution now comes down to finding the values of C_t, L_t and A_t' that maximize the value function, given that the values at time t imply a specific transformation for time $t+1$. That is, they jointly determine the expected (hence the expectation operator) value A_{t+1}. To determine this expected value, we use the conditional return distribution to weigh the utility of the future states. Mathematically this is taken care of by integrating the future value function over all possible states, which depend on the probability distribution of returns.

$$E_t[V_{t+1}(A_{t+1})] = \int f(A_{t+1}|A_t') V_{t+1}(A_{t+1}) dA_{t+1}$$

The maximization problem however can be simplified by eliminating a variable by conditionally solving for its optimal value. Using basic calculus we can find the optimal intra-temporal trade-off between consumption and leisure and solve it conditional on A'_t. We find for interior values of L_t:[14]

$$L_t = \theta + (1-\theta)\frac{A'_t - A_t}{W_t}$$

The restrictions $0 \leq L_t \leq 1$ imply that:

$$L_t = 1 \text{ if } \frac{A'_t - A_t}{W_t} \geq 1$$

$$L_t = 0 \text{ if } \frac{A'_t - A_t}{W_t} \leq -\frac{\theta}{1-\theta}$$

We summarize these three equations as $L_t = L(A'_t)$.

We now substitute the labor supply function and the budget constraint, both expressed as a function of A'_t, into the utility function to simplify the maximization problem, since A'_t is now the only remaining decision variable. The reduced problem is:

$$V_t(A_t) = \max_{A'_t}\left[U\big(C(A'_t), L(A'_t)\big) + \beta \int f(A_{t+1}|A'_t)V_{t+1}(A_{t+1})dA_{t+1}\right] \quad (8)$$

subject to

$$A'_t \leq W_t + A_t \quad (9)$$

$$A'_t \geq \mu_t = -\sum_{s=t+1}^{T} \beta^{s-t} W_s$$

$$A_T = 0$$

$$A'_T \geq 0$$

Equation 9 restricts the household from saving more than its budget, which in turn is equivalent to the restriction that consumption cannot be negative: $C_t \geq 0 \qquad t = 1, \dots, T$

3.3 Solution implementation

After calibrating the model properly, we can solve for optimal household decisions. To solve the household maximization problem we could use first order conditions of the household maximization problem, supplemented with conditions for inequalities, and solve for optimal values using algebra. However, this is only valid for simple problems with deterministic processes for asset return and wages. For other cases no closed form solution can be derived and solving the system will have to be

[14] We take the total differential of the utility function (i.e. to C_t and L_t), set it equal to 0 and solve for L_t.

done numerically using dynamic programming. We will use several tricks for making the algorithm a workable algorithm that allows the values to be tracked over time conveniently. Although the tricks are technical and have no economic meaning, we will discuss them instead of relegating it to an appendix, for calibration of the model is impacted. After implementing these tricks, our results will still be approximately true to the point that the results are not changed fundamentally.

Since we are solving numerically we need to have a finite amount of outcomes to check for the amount of savings put aside and wealth level, for there are infinite real numbers and iterations would go on indefinitely. Therefore, we define one grid of 100 positive possible values for household wealth and a zero value. Household wealth, A_t, and invested wealth, A'_t, will then only be able to contain values from this grid and the grid will stay constant over time.[15] After this choice, the return on assets that are invested is taken from five values, each having a certain probability of occurring (equation 5).[16] Given a value of A_t, a certain choice of A'_t will thus have five possible values for A_{t+1}. For $A'_{t,i}$ it can then move to $\left(1+r_j\right)A'_{t,i}$ with probability p_j for $j = 1,2,3,4,5$. The value function of this value for A_{t+1} will then be weighted by the probabilities and the algorithm will pick the value of A'_t with the highest value of the value function V_t according to equation 8.

To force the values A_{t+1} to be on the grid, we need to make the values of the grid scaled versions of each other. Therefore we assume $r_j = (1+r)^{j-3}$ with probability p_j for $j = 1,2,3,4,5$. If we now choose $A'_{t,3}$ as the value of wealth to be invested, the algorithm can move to either $(1+r)^{-2}A'_{t,3}$, $(1+r)^{-1}A'_{t,3}$, $(1+r)^{0}A'_{t,3}$, $(1+r)^{1}A'_{t,3}$ and $(1+r)^{2}A'_{t,3}$. As we chose the grids for A_t and A'_t to be equal these values are equal to $A_{t,1,}$, $A_{t,2}$, $A_{t,3}$, $A_{t,4}$and $A_{t,5}$! In the next period the algorithm can thus move from one of these points again and whatever A'_t it will choose, it will be on the grid again.[17, 18]

Now, what about the value r? It is now basically a growth rate of the values on the A_t (and A'_t) grid, but it determines the return probabilities. For calibration we thus have to choose this value jointly with the conditional probabilities p_j and the value A_1, as these determine the grid values, and its probabilities.

---

[15] This is chosen such that it can be implemented conveniently in a spreadsheet type of program and has no economic meaning. The basic point is that they need to be predetermined and it is impractical to have a time varying grid.

[16] The algorithm does not randomly draw a return, it simply weights the utility value of its next five states (these belong to the return percentages) by the probabilities of the returns.

[17] The step from $A_{t,0}$ (the zero value) to $A_{t+1,1}$ or $A_{t+1,2}$ (both positive numbers) is obviously not an increase of $(1+r)$ or $(1+r)^2$, as this would still be 0, but we just assume that it will move from 0 to the value $A_{t+1,1}$ or $A_{t+1,2}$ with a certain probability.

[18] We make a slight adjustment for the special cases $A_{t,0}$ (the zero value), $A_{t,1}$, $A_{t,100}$ and $A_{t,101}$. From $A_{t,0}$ we can go only to $A_{t,0}$, $A_{t,1}$ and $A_{t,2}$ and the probability that is normally assigned to the down movements is added to the conditional probability of the future value $A_{t,0}$. That is, the probability of going from $A_{t,0}$ to $A_{t+1,0}$ is thus $p_1 + p_2 + p_3$, whereas the probabilities of going to $A_{t+1,1}$ and $A_{t+1,2}$ are p_4 and p_5 respectively. The probabilities then sum to 1 again. For $A_{t,1}$ only $A_{t,0}$, $A_{t,1}$, $A_{t,2}$ and $A_{t,3}$ can be reached and the odds of the large down movement is added to future value $A_{t,0}$. The odds of going from $A_{t,1}$ to $A_{t+1,0}$ is thus $p_1 + p_2$. The largest values in the grid have to be adjusted analogously, but the up-movements are now capped in a similar way. These adjustments ensure we stay on the grid and that probabilities sum to 1.

4. Calibration

To make the model suitable to analyze Dutch household cohorts we calibrate the household characteristics to make them look like them. To simulate recession-type circumstances and normal times, we will create two regimes that we can choose from. We will now describe the possibilities.

4.1 Normal and recession regimes

To calculate costs of the Great Recession we will compare two extreme cases, which are basically two different regimes that can be imposed in the model. One case concerns what would have happen had there been no recession at all. This is our benchmark case. The method of calculating the results for the remaining lifetime of household cohorts will depend on what we observed before the Great Recession. It might not be realistic to assume that the economy will return to such an economic regime, as the recession might have changed things forever, but it is a simple but reasonable approximation. The alternative to extrapolating the pre-recession trend would mean we have to use complex forecasting techniques, which is beyond the scope of this thesis. To average out good and bad years we take data from 2001-2007, which spans approximately a full business cycle (CBS, 2009).

The alternative case will be what happens after the recession (2009 and further) broke out. Similar data for this period, despite the small sample interval (only three years), will be used to extract recession regime trend growth.

4.2 Household characteristics

For simplicity we fixed the lifespan of households. To approximate reality, we will assume that a household lives for 60 periods at maximum. If we assume that actual households live to about 80 years, which is roughly correct, this means life starts at age 20, probably not far from the average age at which households start working. However, the household cohorts we selected did not all start their working life when the recession broke out. Any utility, consumption or leisure time households have already enjoyed pre-recession is basically a sunk cost type value. Economically what matters are the actual future costs. We will therefore enter each of our cohorts into the model at their place in the lifecycle of 80 years. For example, for households in the category 65+ this means we enter them at $T = 46$ and we track them for 15 years, i.e. until $T = 60$. The quantity V_0 should thus be interpreted as the expected discounted lifetime utility from 2007 until their death.

The IES parameter σ contains information about substitution between periods and can therefore be seen as a determinant of the strength of the savings motive of a household. The marginal utility of consumption or leisure time is increasing in the IES, which means that a high value of the IES induces households to consume instead of save.

In the academic literature, many authors have attempted to estimate the IES, but results vary. According to Guvenen (2006), dynamic macroeconomic models generally require an IES of about 1, but studies that look at co-movement between aggregate consumption and asset returns find it to be close to 0. He attributes these differences to the use of representative consumer models as heterogeneity between the type of models does not allow valid comparisons. As such neither is apparently wrong.

A different approach to estimating the IES runs through its relation with the coefficient of relative risk aversion, which are each other's inverses. If we are emphasizing this interpretation of the parameter, we could use survey questions about risk taking as an approximation for this parameter when calibrating the model. Kapteyn, Kleinjans and van Soest (2009) have used psychological questions in the DHS survey to estimate the IES for the Netherlands and found a mean value of about 0.2, reasonably consistent with the earlier estimates (see literature review).

As heterogeneity is an important driver of estimation of the IES parameter, we would ideally assign different values to our household cohorts. This is troublesome though, as it would likely change over time and beyond the scope of this thesis. We will therefore set the IES at the same level for all household cohorts and keep it fixed over the lifetime. We will either use 0.2, the figure that Kapteyn et al. (2009) found, or 1, which implies log utility and report results for both. As this is a household characteristic we will not distinguish between normal and recession times.

The discount factor β determines how strongly households value utility through time. Due to the different remaining life spans of the household cohorts we are dealing with, the way utility is discounted is important. Future benefits, such as for younger households, may mean less if they are heavily discounted, but short term losses are another factor.

Estimation of discount factors for household cohorts suffers from the same issues that estimation of the IES has to deal with. It may not be constant over time, so distinguishing over cohorts would be painstaking. For this reason we will keep it constant for all households. A suitable discount factor could be extracted from survey questions about time preferences.

Another approach, which is simple to implement, would be to set the discount rate equal to the real interest rate, which was introduced in the model section. For a long-run equilibrium in a closed aggregate economy, lending and borrowing should be equal, a condition that is met when the real interest rate equals the average subjective discount rate. But what discount factor would be appropriate for these uncertain consumption streams (asset return is not fixed)? We avoid this issue and will consider two values for the discount factor for the impact. The values we pick are 0.97 and 0.99, roughly 3% and 1% for the discount rate. As this is a household characteristic we will again not distinguish between normal and recession times.

The relative preference for consumption θ determines the tradeoff between either consumption and leisure. Ideally this parameter would have to be set such that both consumption and working time matches the observed values reasonably. However, the CBS only has aggregate household consumption data, not for cohorts, whereas the DHS data is too noisy and unreliable for income and consumption data, so this is basically impossible to compare. Therefore we consider two cases; a value of 0.75, i.e. 75% preference for consumption, and 1, only consumption matters. As this is a household characteristic we will not distinguish between normal and recession times.

4.3 Income and prices

In our model wages and prices determine the level of consumption and labor supply, but for the households we will make it exogenous and deterministic for simplicity. A deterministic path of wages and prices is especially suitable in case we wish to impose a temporary recession, because we could force the levels or growth rates to return to normal times easily.

Prices are unimportant in this model and we could do several things to avoid our analysis being blurred due to price effects. The demand functions for consumption and leisure time can be shown to be homogeneous of degree 1, since scaling the initial wealth levels, prices and labor income by the same factor has no effect on consumption expenditures. We can therefore choose to let labor income equal real labor income, which normalizes the price level to 1 and eliminates distortion by prices.

For labor income, we use the income data for households from the CBS. We take disposable income, i.e. total income minus various premiums and taxes, and eliminate CBS wealth income from these figures, such that we are left with income that can be spent and that is due to working (equal to labour income from table 2). Wealth income will run through the returns on assets, for which we will use the crude approximation using the micro data.

We then need to find a way to impose recession shocks to the wage level. Wages increase over the working lifetime, but for the last decade generations have on average been richer than their predecessors were at the same age, i.e. an intergenerational difference. Now, if we take the income levels of the group 60-64 in the CBS data, divide by the income of the group with age under 25 and compute the geometric growth rate, we get the rate at which the current wage level of the group under 25 would have to grow annually to reach the level of income that 60-64 year old households enjoy now when they reach that age (given that there is no additional growth). If we actually compute this rate for the last 10 years, we see that it is rather stable at about 1.8%. But when we look at the average intergenerational growth rate of our non-retired cohorts, i.e., the rate at which successive cohorts become richer than their predecessors on annual basis, we see a drop in the recession period (2009-2011) compared to the years before (2001-2007). This difference ranges from 1.3% for the oldest cohort to 2.8% for the youngest. From the point of simulating labour income shocks this is interesting with respect to simulating labour income shocks. Earlier on we mentioned disproportional shocks for the youngest cohorts, which had been implemented in for example the article by Glover et al (2011). Young cohorts are hurt this way for a long period over their lifetime, as a shock in the growth rate of wages is persistent.

What we will do for the wage levels is this. We take the roughly constant growth rate of 1.8% and this rate will apply, independent of the regime. To simulate wage shocks we determine, per age group in the disposable income data (which match the intervals of our cohorts), the intergenerational rate, in both normal and recession times.[19] For a normal regime, a household will enjoy the normal intergenerational growth rate on top of the 1.8% constant rate. This rate is age-specific and as soon as a household moves into a new age interval, the rate will shift to the rate that belongs to the age group it is then part of. In a recession regime, the intergenerational growth rate will shift downwards to its post 2008 rate, as long as there is a recession. Rates will still shift if a household passes into a new age group while there is a recession, but towards the recession rate of that group. This rate may also shift back upon a regime shift. The rates that we find for our cohorts using the groups in the CBS data for normal times are roughly equal to 0.99%, 1.65%, 0.75% and 1.18% (young to old). In recession times these rates shift to -1.78%, -0.63%, -1.29% and -0.13%. Taking differences shows the losses being cohort-specific and disproportional for younger cohorts, as we get rates of -2.77%, -2.28%, -2.04% and -1.31% for our successive cohorts.[20]

For some intuition of this technique for setting up the growth of labour income, consider the example of a household of our youngest cohort, 20 years old. Its starting wage equals 27899 euro per year. If it would never encounter a recession, it would earn $27899 * 1.018^{45} * 1.0099^{15} * 1.0165^{10} * 1.0075^{10} * 1.0118^{10} \approx 103.000$ in the year before retirement. Now, consider a lasting recession. It would then earn $27899 * 1.018^{45} * 0.9822^{15} * 0.9937^{10} * 0.9871^{10} * 0.9987^{10} \approx 38.700$. This difference is extreme, but intermediate cases would show the same thing, that recessions permanently reduces labor income for all periods.[21]

One adjustment still has to be made, which is taking care of retirement. After retirement, real wage growth is unrealistic. Therefore, we set the wage of age 66 at a constant fraction of the wage at age 65. According to Duval (2003), this fraction, commonly called the retirement replacement rate, equals on average 62% in the Netherlands. When a household turns 66 it receives the retirement replacement rate of the wage at age 65 until age 80, without any growth. To impose this we assume that for these retired households the instantaneous utility function depends only on consumption and for convenience reduces to:

$$U(C_t,) = \frac{\sigma}{\sigma-1} C_t^{\frac{\sigma-1}{\sigma}} \qquad \text{for } \sigma \neq 1$$

[19] We make one small modification though, we set the intergenerational growth for the cohort that is younger than 35 at the rate that we find for the age group 30-35. The growth rate between generations younger than 30 is much lower and if we would impose an average, including these rates, labour wages for our youngest cohort would at age 35 not even exceed the current level of the cohort 35-44, even for normal times, which we deem unrealistic.

[20] An important point to realize is what we are actually picking up in the data by choosing and using these numbers. Wages cannot be lowered easily by the employer. The actual drop in the numbers is therefore likely the effect of unemployment of only certain members of the group. This may not be ideal, but we will still use this technique of simulating wage income shocks due to a lack of a better method.

[21] For a ten year recession we would find $27899 * 1.018^{45} * 0.9822^{10} * 1.0099^{5} * 1.0165^{10} * 1.0075^{10} * 1.0118^{10} \approx 78.000$ euro's.

For $\sigma = 1$ we get $U(C_t,) = \ln C_t$

4.4 Distribution of returns

Expected return on households assets consist of two elements, the level of the returns and probability of observing said return. We can expect shocks to occur in both of these aspects. First, the overall level of returns is lower in recessions, but the a shift in the probability distribution may be an important driver as well. Lower returns may be simply occur less frequently.

There are several options for calibrating the return distribution. It could be realistic to allow all household cohorts to invest in the same assets. This would mean that the return distribution is not cohort-specific, but simply dependent on whether there is a recession or not. In reality household portfolios are not the same for all cohorts though, which has to do with portfolio preferences (some household prefer to own a residence for example) and liquidity restrictions (what can be afforded at what point in life).

To account for this we could allow a cohort-specific portfolio and thus return distribution. This is what Glover et al. (2011) have done through scenarios with endogenous portfolio selection. For our model we would have to enforce it as an exogeneity. The downside of such a decision would be preventing households from investing in good investment opportunities and forcing them to hold potentially bad investment assets, while other cohorts can enjoy better returns. Additionally we would have to create a portfolio for them without a good guideline what to base it one. Historic portfolios are an option, but force the households to own what their predecessors owned when they were at that age, as if things never change. Therefore we deem a general portfolio (i.e. for all cohorts) with a single asset a better option for realism.

Due to the tricks we had to implement to make solving the problem workable we can only set r, the growth rate of the grid values for wealth. The five returns will then be $(1+r)^{-2}$, $(1+r)^{-1}$, $(1+r)^{0}$, $(1+r)^{1}$ and $(1+r)^{2}$ and we can set the probabilities. We basically have to find a reasonable value for r and adjust the probabilities to match the actual distribution for the five returns percentages. To derive an approximation for the possible wealth returns, we take household asset income levels (which were used to calculate the mean values in table 2). We then divide by the total stock of return generating assets (real estate, equity and savings (bank accounts)), which gives us asset returns by household. We do this for every year and pool the rates of return over all households for the years 2001-2007 for normal times. We repeat this for 2009-2011 for recession rates of return on assets.

Besides many good observations, this also yields many outliers (see also appendix C), so we cap the returns at 10% in absolute value. Figure 1 shows a histogram of the distribution of returns in normal times. We see that in normal times a major part of the mass is concentrated on the right of the 0% barrier, i.e. households are becoming richer by investing. The likelihood of losses is very small,

Figure 2: Histogram of asset returns in normal times (2001-2007)

This histogram shows the distribution of asset returns in normal times (2001-2007). Returns have been computed using the DHS dataset, capped at ±10% for realism and pooled over the years 2001-2007. The bin size on the horizontal axis has been set at exactly 1%, which means the mass of each category equals 5%. The distribution shows a large mass to the right of 0%, i.e. increases of assets.

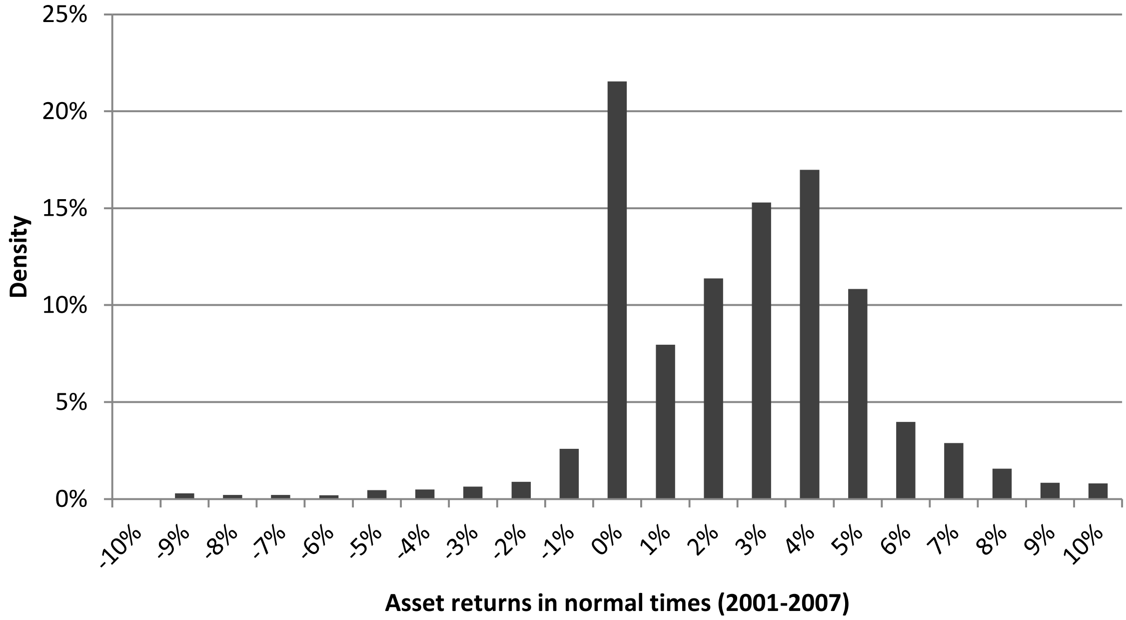

roughly 10%, and the losses themselves as well. There is also a large volume of households that approximately earn no return overall, which we see from the spike the next to 0% return. Larger returns than 5% have a small probability of occurring, it is as likely as losing a couple of percent in terms of assets. When we look at the histogram of return in the case of the recession years (figure 2) we see a large swing. On average households have lost over 2% in terms of assets and from the table we see that more than 75% of the households did not gain anything. Positive returns are realized for between 10 and 25% percent of the households, but for only 10% of households this return exceeded 0.45%, a very low rate. Losses in excess of 5% are not unlikely as well and at least 10% of the households saw their assets decrease by more than 6%. The explanation behind the shift of the distribution is motivated by the decrease in real estate prices, which for most households forms a major part of their assets.

Selecting a return percentage based on the graphs and the data is arbitrary, so we pick a return of 2% for the r value, as this seems convenient, because we get as approximate return values -4%, -2%, 0%, 2% and 4%. These percentages then approximate a large drop in wealth, a normal drop in wealth, no change, a small increase and a large increase. To assign probabilities for these approximate returns, we take the distribution of the return on assets and include the mass (based on the percentiles we find in our data) that is within 1% of this return for the three middle return percentages and for the outer two returns we include the mass up to ±10%. -4% return on assets will thus get the mass of asset returns from -10% to -3%, -2% will be assigned the mass between -3% and -1% and so on. For normal

Figure 3: Histogram of asset returns in recession times (2001-2007)

This histogram shows the distribution of asset returns in recession times (2001-2007). Returns have been computed using the DHS dataset, capped at ±10% for realism and pooled over the years 2009-2011. The bin size on the horizontal axis has been set at exactly 1%, which means the mass of each category equals 5%. The distribution shows a large mass to the left of 0%, i.e. decreases of assets.

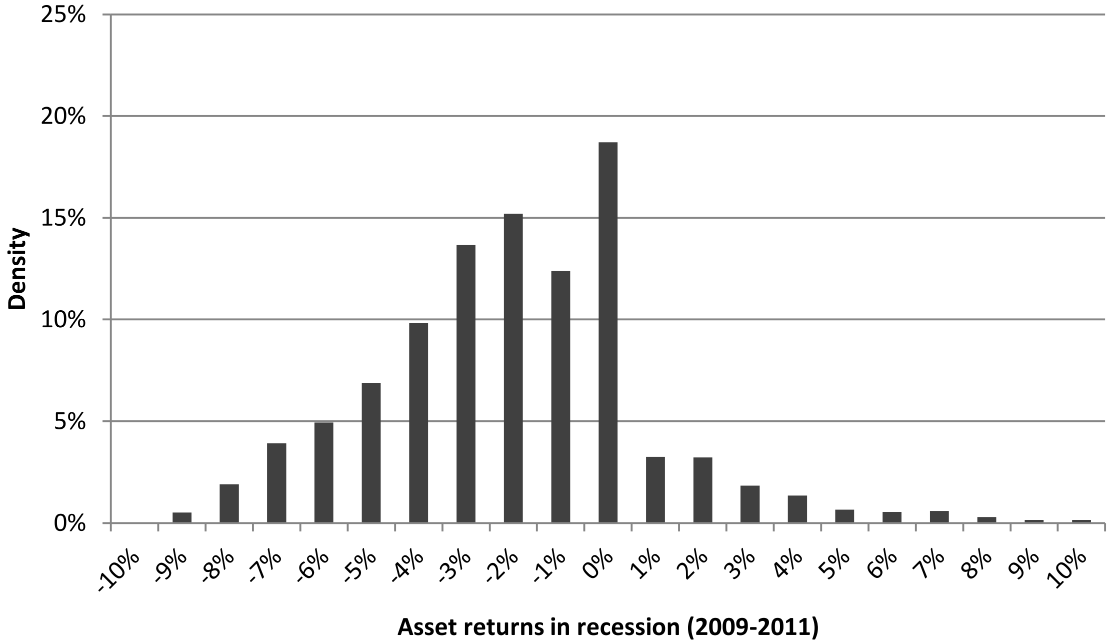

times we therefore assign the probabilities 2%, 3%, 30%, 27% and 38%. For recession times we end up with 42%, 28%, 22%, 5% and 3%. The shift of the distribution may be large, but it is in accordance with table 2 that asset income has moved into negative territory.

Now we only have to enter the wealth levels of our cohorts in some way. To do this we force the A_t' grid values to be close around the mean wealth for our cohorts from table 1, but such that there are a sufficient number of A_t' values higher and lower than the mean value.[22, 23] We unfortunately cannot test for any other values (median, percentiles, etc.) due to model rigidities.

[22] The story is a technical, but related to making the solver suitable to solve these type of problems. For some intuition, take cohort 5, which has a mean value of 316 thousand euro's. If we let the A_t' values range from 10.000 to about 100.000, this would not be realistic, as the algorithm would pick 100.000, despite the fact that a realistic value would be way higher. Setting 200.000 as the maximum would still cause the same problem. A better choice would in that case be from 10.000 to for example 400.000, such that there is enough variation in the grid values to pick. Now since A_t can go down if large losses occur for many periods, lower A_t' values are also realistic. To account for this we also need to include some low values (or the algorithm would pick the lowest value simply because lower values do not exist in the grid). These do not have to be close to zero though, as the odds of dropping to such levels is negligible, despite the 70% odds of losing wealth in a recession when going from period to the next one (would require many periods of losses in a row and thus low odds). Analogously, for non-recession times we also need to have high enough values to account for spectacular wealth growth (possibly 60 periods straight for the youngest cohort, less for the others). But, like for the low value case, we do not need to include as high values due to the probabilities. In the end we just need to balance these effects. Therefore we have chosen as our values of A_2': 10000, 25000, 35000, 50000 and 50000 (young to old). Using these values we find for neither of our scenarios any optimal A_t' value that is the result of picking the minimum of maximum value in absence of a better value that is available. This is exactly what we wanted, so the selected values are appropriate.

[23] We know from section 2 these indicators are not perfect and reliable, but neither is the level of wealth from the CBS data.

5. Simulations

5.1 Indicators of interest

Now we will describe scenarios that deal with heterogeneity of the cohorts and run calibrated simulations. Utility is just a decision tool to allocate consumption and labour supply, but lacks an interpretation, in addition to being only ordinal. This means we will use the present value of lifetime expected consumption as our unit of interest. We will calculate losses of discounted lifetime consumption expenditures for the mean of household wealth by cohort, based on the DHS micro data, which was reported in table 1.

5.2 Simulations

We will now simulate three scenarios. The first is our benchmark, the normal economy, which has the normal regime constantly. The second one is a lasting recession, where households are constantly subjected to the recession regime with lower wage growth rate and a shifted return distribution. For a bit more realism, we also simulate a recession that lasts for only ten years. For five years now has there been a period of reduced growth without a sign of serious and persistent recovery. This could thus be a very reasonable scenario and should be interpreted as a fierce long recession. The lasting recession should be then be interpreted as a hypothetical and extreme case. In the temporary recession scenario, there is a regime change back to normal times after the ten recession years.[24, 25] The losses we find for our two recession scenarios will then be compared to the normal economy to determine which cohorts are hurt the most.For some intuition of the impact of labour income shocks for our households we first provide graphs of the labour income paths. Figure 4 shows the various wage paths for cohorts under the three scenarios. A lasting recession is particularly harmful for mainly the young cohorts, as these suffer from both lower growth for more periods and disproportional growth rates shocks in the early part of working life. For older households a recession is still a relevant factor for their labour wage, but less so. Note that cohort 5, 65+ is not affected in their labour wage, since they are already retired. A temporary recession will only impact the growth rates temporarily, which we see in the graphs, but the damage to the overall level of the labour income is already done and permanent. For cohort 4 this temporary recession has no different effect from a lasting recession, as their last

[24] This can be implemented, since the dynamic programming algorithm starts at the final period and iterates backwards. We can then simulate with a normal regime at the end of life, insert expected utility values (and other indicators) and let it iterate with a recession regime from then on and shift it back again.

[25] With respect to the recession scenarios we have chosen, we need to realize our setup is static. For longer recessions the impacts may be so severe that in reality other effects would be triggered. We could think of drastic political actions, social movements, emigration, spillover to foreign economies and so on.

Figure 4: Wage paths

This figure contains the wage paths for the four youngest birth cohorts for the three different scenarios in our simulations; a normal economy, a lasting recession and a ten year long recession. Wages levels are given from 2007 and due to age differences among cohorts they start at different ages. Wage levels are in real terms and expressed in 2011 constant prices. Price levels have been deflated using CBS CPI data. The fifth cohort, 65+ is ignored, since their labour income is fixed in all three scenarios. For all cohorts we see that the difference between normal times and a lasting recession is growing larger and larger until retirement due to consistently lower growth rates. This is particularly harming to the youngest cohorts, but still relevant for the older cohorts. The ten year recession paths show that a temporary recession will also impact the labour path for the whole remaining lifetime. For cohort four there is no difference between this and a lasting recession, as their final period of labour growth coincides with the recession. Retirement income is also impacted through a lower level for all cohorts.

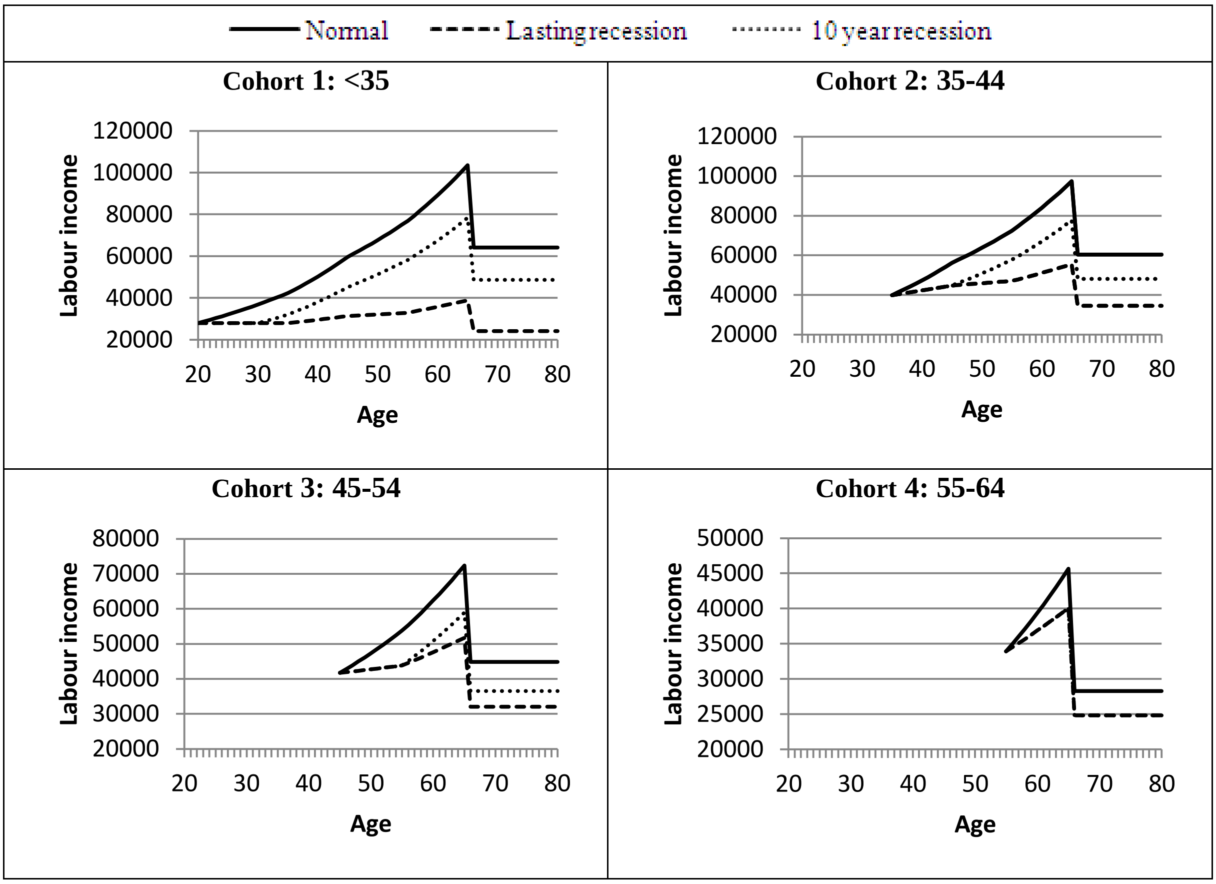

period of wage growth coincides with the recession in both cases.

As our base case we consider the parameter set $\beta = 0.97$, $\sigma = 1$ and $\theta = 1$. θ equal to one has the implication of effectively eliminating disutility of working in the model, so only consumption is important. This assumes full-time working households, which is convenient as it isolates consumption shocks. For sensitivity to labour shocks, important as our differently-aged households cannot compensate equally through working more, we will relax this assumption and test for $\theta = 0.75$. To keep things simple, we will not consider unemployment risks through setting labour supply at zero at (random) times. For σ we prefer to use 1 in our base case over the value of 0.2 that Kapteyn et al. (2009) found as it is within the interval of between 0.25 and 1 that has been found in the literature before (see section 1). The log utility function that results is concave and assumes the household is risk-averse to variable streams of consumption. Setting $\sigma = 0.2$ would further increase this concavity and make households even more risk-averse. We can therefore use this case for sensitivity of our results when we wish to see the effect of increased risk aversion. $\beta = 0.97$ is set arbitrarily, as we have deemed looking for an appropriate discount rate beyond the scope of this thesis. Sensitivity analysis will be performed with $\beta = 0.99$.

5.3 Results

From the model we can easily compute the expected present value of consumption expenditures and the expected lifetime labour income (product of labour supply times wage level). We can then use the budget constraint to back out lifetime asset income by taking lifetime consumption and subtracting labour income and initial wealth.[26] These are the asset income figures given in tables 4-7. We then divide the present value for each type of income by the present value of consumption to derive shares of consumption. These will be useful for analyzing shocks to consumption as in the end the total shock to consumption is simply the sum of a couple of weighted shocks; to labour income, to asset income and initial wealth (this is zero though by default). When we mention percentages, we mean percentage points.

Table 4 provides results for present values, losses and budget shares of consumption for our base case with $\beta = 0.97$, $\sigma = 1$ and $\theta = 1$. If we compare a 10 year recession with the normal economy we can see a pattern. The lowest losses are found for cohort 3 at 16.1% of consumption and losses are increasing if we move up (down) to a younger (older) cohort. This suggests that for cohorts different forces dominate for cohorts when dealing with a ten year recession.

The oldest cohort has the largest losses in terms of consumption at 23.1%. As their retirement income is fixed, they do not incur losses on this part of their income portfolio. Their total loss is therefore the result of either losses on investments in the ten recession years, foregone opportunity costs (could have had a better return had there been no recession) and a lost compounding effect (i.e. returns being invested again). If we compare their losses on asset income with the other cohorts, we see that their relative losses stand out as well at 36.2%. So what may cause this high loss? We have imposed $\theta = 1$, so cohorts cannot vary their labour supply, only consumption. Yet, all cohorts are dependent on the same investment environment. The explanation of the larger costs are therefore likely caused by their shorter remaining lifetime. During the recession years they can avoid investing some of their large wealth (to the point where smoothing is no longer cost-effective for maximizing lifetime utility), but when the regime shifts back there are only five years remaining to get a good return on their reduced savings balances. In other words, the relative importance of one year of investment possibilities is much higher than for other cohorts.

[26] As we solve numerically, this will only be approximately right.

The other interesting cohort in this table is cohort 1, the households that have just started working. As we have seen from figure 4 their labour income losses are potentially larger than for the other cohorts, especially for a permanent recession. As shown by table 4, even this temporary recession of “only” ten years brings them the largest loss among cohorts, at 22.3% of labour income. The explanation for this is that labour income growth rates are the highest in the early stages of working life, so large damage is done. As labour income is their main source of income (59% in normal times), this also translates into a large consumption shock. The labour shock however is accompanied by a relatively low shock to asset income (-15.3%), in fact the smallest among all cohorts. Multiple aspects contribute to this fact. One concerns their lower initial wealth, they only have 54 thousand euro to start with. This means that the wealth they invest and that is subjected to the negative investment environment is only small, which results in small losses. Secondly and more importantly, they have a long recession-less life ahead of them. After a recession of ten years, fifty periods long these households can invest their wealth at favorable opportunities. The largest absolute decrease in labour income (cannot be invested) was apparently not able to offset this effect. To put the asset losses of the cohort in perspective, we can look at the consumption shares of cohorts for the ten year recession. While cohorts 2 to 5 have decreased the share of asset income by 1 to 11 percent points, this cohort sees its relative share rise by 2 percent points. Overall though, the cohort still loses wealth compared to normal times, due to opportunity costs (also from labour income) and lost compounding effects.

When we look at cohorts 2-4, the effects we have discussed so far are more blended together. If we look at labour income, we find them to be decreasing monotonically in age, which is explained by the lower growth rate losses for older households. Although we found asset income shocks to be the smallest (largest) for the youngest (oldest) cohort, they do not increase monotonically. Cohort 2 has losses of 20.5%, cohort 3 only 18.4%, but cohort 4 has larger costs of 29.3%. Which effects dominate for these specific cohorts cannot be separated easily. The older cohorts among these have the largest effect of the shorter remaining lifetime, but also lower absolute losses from labour income (which would have been invested) and higher initial wealth (will be subjected to recession regime). When we combine the asset income and labour income shocks, we find cohort 3 to have the lowest consumption shock, before cohort 2 (-18.6%) and cohort 4 (-21.1%).

Now, when we look at a lasting recession we see the pattern of consumption shocks from the temporary recession shift in some way. The oldest cohort now incurs the lowest losses among all households at 37.1%. This loss itself is the result of pure asset shocks of -58.3%. When we look at the other cohorts, we find cohort 3 to lose the least in terms of consumption at 45.0%, before cohort 4 (49.3%) , cohort 2 (49.4%) and cohort 1 (50.7%). When we look at the magnitude of labour shocks, we find them to be higher than in the ten year recession case for cohorts 1-3, which makes sense, and also decreasing in age. The dispersion of these shocks have become larger as well, which is also not surprising. They are explained by the full persistence of compounded labour shocks, i.e. through reduced growth rates, which we saw from figure 4. The asset income shocks among all cohorts however are much larger, partially due to the increased labour income shocks, but of course also through the investment potential; the worsened environment has to be dealt with for a pro-longed time.

Previously only cohort 5 did not enjoy a labour shock, but if we compare the two recession scenarios, now also cohort 4 was not affected. Comparing these two therefore makes an interesting case. Cohort 4 has losses of 73% for asset income, almost two and a half times the loss in the ten year recession and the highest among the cohorts, while the losses of cohort 5 did not even double at 58.3%. The remaining lifetime argument now turns against the younger of these cohorts, as it has to endure a recession for its remaining 25 years, instead of only 15 for cohort 5.

For cohorts 1-3 the effects are more difficult. Cohort 1 has only one argument in favor of it, compared to the others, its lower initial wealth that can suffer from losses. Besides that, it has to endure the recession regime longer and features the largest labour shock in both absolute and relative terms. In that sense it is even surprising that they do not have the largest relative loss in asset income. Cost effective utility gains were apparently available through smoothing consumption instead of investing and seeing it decrease to losses on the portfolio. For cohorts 2 and 3 not much can be said about specific effects from this table. They suffer from comparable relative shocks to asset income, but as their exogenous labour shocks are smaller, so are their total consumption shocks, compared to cohort 1. Cohort 3 is also better off compared to cohort 4.

Now we will look at the sensitivity of our results. Table 5 provides results for $\sigma = 0.2$, i.e. households are more risk averse and will prevent dispersed streams of consumption over the lifetime. The first thing that sticks out for the ten year recession is that we find another pattern of increasing costs for very young and old households. The "curve" has sort of flipped around, as cohort 4 is now the cohort with the least losses, while cohort 1 has the largest costs. As θ is still equal to one, there are no changes for labour income. For asset income we find a pattern that looks like the $\sigma = 1$ case, though the youngest cohort sticks out as the second largest loss taker.

So, what could explain these patterns? As the households wish to smooth their consumption path, more wealth will be reallocated to the bad years, relative to the $\sigma = 1$ case, i.e. away from the good years. This leaves less budget for investments for positive returns, which decreases asset income drastically through higher order effects. Young cohorts already have a large labour income shock, which is now combined with these extra effects, and suffer from a 19.3 percent point increase in asset income losses, compared to the $\sigma = 1$ case. For cohort 2 and 3 we find excess losses of 11.2% and 8.9%, compared to only 1.2% and 1.9% for the oldest two cohorts. Apparently, these effects are less of an issue for the oldest cohorts, whose losses are dominated by initial wealth losses and relatively few good investment years. Their small increase in asset income losses is therefore probably related to higher order effects. Overall, these cohorts also have the lowest losses of consumption. This is related to an overall decrease in dependence on asset income, which we see from the budget shares, and the small labour income shock. The younger cohorts suffer from equal labour income shocks, but since they incur larger asset income shocks, their overall balance is more negative. For the lasting recession we see a small change in the pattern of consumption losses. For $\sigma = 1$ they were found to follow a "curve" with a minimum for cohort 3, but this time losses are decreasing in age. When we compare the losses to $\sigma = 1$ we see that the consumption losses of the cohorts have increased by 1.6% for cohort 1 and decreased by 2.9%, 5.7%, 16.7% and 13% for cohorts 2-5 respectively. The most plausible explanation is the relative dependence on the types of income. Labour income is more important for the cohorts according to the consumption shares, which is explained by the increased smoothing motive. The increase in dependence thereby compensates for the higher asset income shocks, resulting in lower final shocks of consumption.

For households that no longer have to work full-time ($\theta = 0.75$), though they can, we expect their labour supply to decrease for all scenarios relative to $\theta = 1$ due to the disutility of working. This prediction is true if we look at the results. From the normal scenario we can see reduced labour income dependence and the present value of lifetime labour income has been slashed for all working cohorts, but different magnitudes arise. By using table 4 we can approximate the aggregate decrease in labour supply in normal times to be 38.3%, 40.4%, 50% and 29.7% for cohorts 1-4 respectively.[27] Cohort 5 is already retired, so for these households we find the same results as in table 4 and ignore them here.

Once a ten recession has been simulated we find the largest consumption losses for the oldest households and cohort 3 again has the lowest costs. When we look at labour supply losses for the working cohorts, we find cohort 4 to have increased labour supply, i.e. through working a tiny bit more in aggregate, but by only 0.1 percentage points. For cohorts 1, 2 and 3 we find losses (12.5%, 2.5% and 0.8%), indicating they have reduced aggregate labour supply, which is counter-intuitive as working can relief some recession shocks. When we look at the asset shocks, they provide interesting results. Households of cohorts 1 and 3 have minor (below 10%) asset income shocks, while cohort 2 and 4 have lost more than 20% of asset income. When we compare with table 4, we find all four cohorts to have increased asset income dependence, relative to $\theta = 1$. The most likely explanation for the decreased labour supply of cohorts 1, 2 and 3 is in that case probably rooted in the cost-effectiveness for utility maximization. As the labour income losses concern aggregates, i.e. over the whole lifetime, it is plausible that the true pattern is varying over the lifetime. The budget for smoothing would then largely be acquired in the good years through asset income, but lower labour supply. During the recession years labour supply would be increased to absorb part of the more severe (compared to labour income) shocks from the recession. If we look at the actual expected paths we find this proposition to hold. While cohort 4 has large initial wealth as a possible partial explanation for the higher losses, cohort 2 does not, so what explains their large asset income shock is unclear due to the multiple effects.

[27] We compared the present values of lifetime labour income for $\theta = 1$ and $\theta = 0.75$ from tables 4 and 6 for these figures. For our recession scenarios this cannot be used as the wage level also changes due to labour income shocks. To derive aggregate labour supply changes we use the approximation $\frac{\Delta WL}{WL} \approx \frac{\Delta W}{W} + \frac{\Delta L}{L}$.

When we turn to the permanent recession we find the households to have reduced labour income shocks compared with the $\theta = 1$ case. For cohort 4 we find the same change by construction, for their labour income is not impacted during their retirement. For cohorts 1, 2 and 3 it is larger (4.9%, 6.8% and 1.4%), but at least not larger than the nominal wage shock, implying an increase in aggregate labour supply. The share of consumption for labour income has also increased, indicating that asset losses are worse overall. Overall, for consumption losses we find the same pattern as for $\theta = 1$, it even has the exact same ordering. If we look at asset income, we see that in general they are comparable to the scenario of $\theta = 1$. Only the reduction of costs of 12.3 percent points in terms of asset income for cohort 3 is a more serious change, but what causes this is unclear.

When we change β to 0.99 there are a couple of effects. Future values receive more weight, so households will be inclined to save more for future consumption, but monetary values are also inflated. From table 7 we can see that the effect compared to $\beta = 0.97$ is only minor though. For the ten year recession, losses are slightly lower compared to the case of $\beta = 0.97$, which is explained by households being more focused on the post-recession period as this is less heavily discounted. The order of losses among cohorts is basically unchanged (cohort 1 and 4 switched orders) though. For labour income shocks are still decreasing in age, with a slight increase in the dispersion and in the level of shocks. The largest labour losses occur in the later years of life and are thus more important for the household. For asset shocks the reverse is true, they have decreased slightly, but the order is identical to the $\beta = 0.97$ case. As they occur early, their relative weight is slightly reduced.

For a full recession the pattern we found for the case with a lower discount factor also still holds (only cohort 2 and 4 switched order). The gap between the oldest cohort and the others has become larger though, just like the difference between cohort 1, the leader in terms of costs. Labour income shocks are the reason for the latter observation, while the former is related to the remaining life effect.

6. Concluding remarks and discussion

The period of stagnated growth that started in the late 2000’s, sometimes known as the Great Recession, has impacted the Dutch economy significantly, including its households. Yet, who bears the costs has so far been unclear. For stabilization policy this is important, for this piece of information is essential for allocation issues.

Using detailed micro data we have divided the Dutch working households into five cohorts, based on the age of the head of the household in 2007 (equivalent to selection on birth year). With a calibrated household lifetime maximization framework in partial equilibrium we have generated cohort-specific labour income and general wealth shocks to simulate scenarios of recessions and a normal economic benchmark. As these recessions occur at different points in life, cohort specific effects were expected to happen. By comparing the present value of lifetime consumption expenditures for four different sets of household parameters we have subsequently determined which age cohorts have been hurt the most by the Great Recession.

We first simulated a temporary recession with a regime change back to normal times after ten years. With a first place in three (two) out of four (three) sets of parameters, the oldest cohort, 65+, bears the large costs. This matches the results of Glover et al. (2011) who, for a temporary but fierce recession, found their oldest cohort to be consistently suffering the largest losses. The youngest cohort, consistently among the first three spots, including a first place, comes in second. The cohort of households aged 45-54 bears the smallest costs with three fifth places and a fourth. Further analysis of losses indicates this pattern, minimum costs for middle-aged households and increasing for both younger and older households, to be explained by different dominating factors. For young cohorts, who are more dependent on labour income, persistent labour income shocks dominate shocks to wealth, although these are not completely separate. For older cohorts this is reversed. High dependence on asset income, coupled with large initial balances and a short remaining lifetime to compensate, cause the worse investment environment to drive their losses. Intermediate cohorts feature a mixture of the labour income and wealth shock effects. Only when we model the cohorts as extremely risk-averse optimizers, do these results change slightly. In that case younger households are the worst off through reduced possibilities of good investments, which is less of an effect for the older three cohorts. A reduced discount rate and endogenous labour supply (relative preference of 75% for consumption) do not impact results differently. These results are fairly consistent with what Glover et al. (2011) found; that the oldest cohort has the largest relative losses. While their results were mainly driven by demand and supply forces of the asset market (general equilibrium), they did also find the remaining lifetime effect to be quantitatively important.

For a lasting recession results shift significantly. The oldest cohort is now consistently best off with a fifth place in all four (three) scenarios. Further analysis suggests the shorter remaining life (less bad periods to endure) is the main explanation for this fact. The youngest cohort comes in first in all four cases in terms of relative losses of consumption. Dominating labour income shocks and higher order effects seem to form the main contribution. The costs for the other three cohorts are between the youngest and oldest cohorts, but in mixed compositions over our scenarios. As we find this for all scenarios, these cohorts are robustly not the worst off.

Despite this effort, much about the distribution of costs is still unknown. By studying partial equilibrium we have cut some corners. Equilibrium forces due to demand and supply, government policy impacts, financial markets, labour markets; we could not consider them. The indirect effects we found from previous research (liquidity restrictions, portfolio choice and especially unemployment risk) have not been tested. Besides these, household characteristics, such as education level, geographical location, household composition, to name a few, may be just as important quantitatively. Taking these into account in a general equilibrium will not be easy, for they have to be internalized in an already complex model. Finally, improved data about wealth and income for households will be required to uncover recession implications for households. In that sense, this is just the beginning.

7. References

Centraal Bureau voor de Statistiek (Dutch Central Bureau of Statistics) (2009). *Toelichting Conjunctuurklokindicator*. Retrieved at 19th of June 2012 at <http://www.cbs.nl/NR/rdonlyres/BAD85D71-4525-485E-BE4A-BF1E4DA81C07/0/toelichtingconjunctuurklokindicator.pdf>

Centraal Bureau voor de Statistiek (2010). Procesbeschrijving statistiek – Vermogens huishoudens. Claessen, J., Den Haag.

Centraal Bureau voor de Statistiek (2013). *CBS – Begrippen*. Retrieved at the 9th of April 2013 at <http://www.cbs.nl/nl-NL/menu/methoden/begrippen/default.htm?ConceptID=366>

Davies, J.B. & Shorrocks, A.F. (1999). The distribution of wealth. In Atkinson, A. B. and Bourguignon, F., editors, *Handbook of Income Distribution*, 605-675. Elsevier, Amsterdam.

De Nardi, M., French, E.B. & Benson, D. (2011). Consumption and the Great Recession. NBER Working Paper No. 17688.

Duval, R. (2003). The retirement effects of old-age pension and early retirement schemes in OECD countries. OECD Economics Department Working Papers. OECD Publishing.

Glover, A., Heathcote, J., Krueger, D., & Ríos-Rull, V. (2011). Intergenerational Redistribution in the Great Recession. NBER Working Paper No. 16924.

Guvenen, F., Ozkan, S. & Song, J. (2012). The Nature of Countercyclical Income Risk. NBER Working Paper No. 18035.

Imrohoroglu, A. (1989). Costs of Business Cycles with Indivisibilities and Liquidity Constraints. Journal of Political Economy, 97(6), 1364-1383.

Imrohoroglu, A. (2008). "welfare costs of business cycles." The New Palgrave Dictionary of Economics. Second Edition. Eds. Steven N. Durlauf and Lawrence E. Blume. Palgrave Macmillan, 2008. The New Palgrave Dictionary of Economics Online. Palgrave Macmillan. 06 October 2008 <http://www.dictionaryofeconomics.com/article?id=pde2008_W000130>doi:10.1057/9780230226203.1825(available via http://dx.doi.org/)

Kapteyn, A., Kleinjans, K.J. & van Soest, A. (2009). Intertemporal Consumption with Directly Measured Welfare Functions and Subjective Expectations. *Journal of Economics Behavior & Organization*, 72(1): 425-437.

Lucas, R.E. (1987). *Models of Business Cycles*. Yrjö Jahnsson Lectures. Oxford: Basil Blackwell Inc.

Lucas, R.E. (2000). Inflation and Welfare. *Econometrica*, 68(2), 247-274.

Lucas, R.E. (2003). Macroeconomic Priorities, *American Economic Review*, 93(1), 1-14.

Perri, F. & Steinberg, J. B. (2012). Inequality and Redistribution during the Great Recession. Economic Policy Paper. Federal Reserve Bank of Minneapolis.

Petev, I., Pistaferri, L. & Saporta Eksten, I. (2012). Consumption and the Great Recession: An Analysis of Trends, Perceptions, and Distributional Effects. In Grusky, D., Western, B. & C. Wimber, editors, *Analyses of the Great Recession.* Updated version of 2011.

Romer, D. H. (2012). *Advanced Macroeconomics Fourth Edition*. New York, NY: McGraw-Hill.

8. Tables

Table 4: Simulation results for $\boldsymbol{\beta = 0.97, \sigma = 1, \theta = 1}$

This table summarizes our simulation results for our household birth cohorts (the age of the head of the household is mentioned in the cohort column) using various setups and various parameter values. Initial wealth, labour income, asset income and consumption are all expressed in thousands of 2011 euro and concern lifetime present value aggregates for 2007. Labour income and consumption were taken from the simulations, asset income was computed by subtracting labour income and initial wealth from consumption. Figures have been discounted using β. Labour supply is fixed at one (full time employment) here as $\theta = 1$.

Household lifetime indicators

		Normal economy			**10 year recession**						**Lasting recession**					
Cohort	**Initial wealth**	**Asset income**	**Labour income**	**Consumption**	**Asset income**		**Labour income**		**Consumption**		**Asset income**		**Labour income**		**Consumption**	
	PV	PV	PV	PV	PV	Change in %	PV	Change in %.	PV	Change in %	PV	Change in %.	PV	Change in %.	PV	Change in %
<35	54	941	1411	2406	797	-15.3%	1097	-22.3%	1948	-19.0%	336	-64.3%	795	-43.7%	1185	-50.7%
35-44	136	1866	1489	3491	1484	-20.5%	1223	-17.9%	2843	-18.6%	580	-68.9%	1049	-29.6%	1765	-49.4%
45-54	202	1678	1090	2970	1369	-18.4%	921	-15.5%	2492	-16.1%	546	-67.5%	885	-18.8%	1633	-45.0%
55-64	304	1624	585	2513	1148	-29.3%	532	-9.1%	1984	-21.1%	438	-73.0%	532	-9.1%	1274	-49.3%
65+	316	1019	265	1600	650	-36.2%	265	0.0%	1231	-23.1%	425	-58.3%	265	0.0%	1006	-37.1%

Share of consumption

		Normal economy			**10 year recession**			**Lasting recession**		
Cohort	**Initial wealth†**	**Asset income**	**Labour income**	**Consumption**	**Asset income**	**Labour income**	**Consumption**	**Asset income**	**Labour income**	**Consumption**
<35	2-5%	39%	59%	100%	41%	56%	100%	28%	67%	100%
35-44	4-8%	53%	43%	100%	52%	43%	100%	33%	59%	100%
45-54	7-13%	56%	37%	100%	55%	37%	100%	33%	54%	100%
55-64	12-24%	65%	23%	100%	58%	27%	100%	34%	42%	100%
65+	19-32%	64%	17%	100%	53%	22%	100%	42%	26%	100%

†: range of values for the three scenarios. Actual values can be computed per scenario by subtracting lifetime asset income and lifetime labour income from lifetime consumption.

Table 5: Simulation results for $\boldsymbol{\beta = 0.97, \sigma = 0.2, \theta = 1}$

This table summarizes our simulation results for our household birth cohorts (the age of the head of the household is mentioned in the cohort column) using various setups and various parameter values. Initial wealth, labour income, asset income and consumption are all expressed in thousands of 2011 euro and concern lifetime present value aggregates for 2007. Labour income and consumption were taken from the simulations, asset income was computed by subtracting labour income and initial wealth from consumption. Figures have been discounted using β. Labour supply is fixed at one (full time employment) here as $\theta = 1$.

Household lifetime indicators

		Normal economy			10 year recession						Lasting recession					
Cohort	**Initial wealth**	**Asset income**	**Labour income**	**Consumption**	**Asset income**		**Labour income**		**Consumption**		**Asset income**		**Labour income**		**Consumption**	
	PV	PV	PV	PV	PV	Change in %	PV	Change in %.	PV	Change in %	PV	Change in %.	PV	Change in %.	PV	Change in %
<35	54	523	1411	1988	342	-34.6%	1097	-22.3%	1493	-24.9%	99	-81.1%	795	-43.7%	948	-52.3%
35-44	136	713	1489	2338	487	-31.7%	1223	-17.9%	1846	-21.0%	71	-90.0%	1049	-29.6%	1256	-46.3%
45-54	202	571	1090	1863	415	-27.3%	921	-15.5%	1538	-17.4%	44	-89.7%	885	-18.8%	1131	-39.3%
55-64	304	475	585	1364	330	-30.5%	532	-9.1%	1166	-14.5%	83	-82.5%	532	-9.1%	919	-32.6%
65+	316	567	265	1148	351	-38.1%	265	0.0%	932	-18.8%	290	-48.9%	265	0.0%	871	-24.1%

Share of consumption

		Normal economy			10 year recession			Lasting recession		
Cohort	**Initial wealth†**	**Asset income**	**Labour income**	**Consumption**	**Asset income**	**Labour income**	**Consumption**	**Asset income**	**Labour income**	**Consumption**
<35	3-6%	26%	71%	100%	23%	73%	100%	10%	84%	100%
35-44	6-10%	30%	64%	100%	26%	66%	100%	6%	84%	100%
45-54	10-22%	31%	59%	100%	27%	60%	100%	4%	78%	100%
55-64	22-33%	35%	43%	100%	28%	46%	100%	9%	58%	100%
65+	28-37%	49%	23%	100%	38%	28%	100%	33%	30%	100%

†: range of values for the three scenarios. Actual values can be computed per scenario by subtracting lifetime asset income and lifetime labour income from lifetime consumption.

Table 6: Simulation results for $\boldsymbol{\beta = 0.97, \sigma = 1, \theta = 0.75}$

This table summarizes our simulation results for our household birth cohorts (the age of the head of the household is mentioned in the cohort column) using various setups and various parameter values. Initial wealth, labour income, asset income and consumption are all expressed in thousands of 2011 euro and concern lifetime present value aggregates for 2007. Labour income and consumption were taken from the simulations, asset income was computed by subtracting labour income and initial wealth from consumption. Figures have been discounted using β.

Household lifetime indicators

		Normal economy			10 year recession						Lasting recession					
Cohort	**Initial wealth**	**Asset income**	**Labour income**	**Consumption**	**Asset income**		**Labour income**		**Consumption**		**Asset income**		**Labour income**		**Consumption**	
	PV	PV	PV	PV	PV	Change in %	PV	Change in %.	PV	Change in %	PV	Change in %.	PV	Change in %.	PV	Change in %
<35	54	960	870	1884	894	-6.9%	571	-34.4%	1519	-19.4%	328	-65.8%	533	-38.7%	915	-51.4%
35-44	136	1755	888	2779	1366	-22.2%	707	-20.4%	2209	-20.5%	572	-67.4%	686	-22.7%	1394	-49.8%
45-54	202	1539	545	2286	1414	-8.1%	456	-16.3%	2070	-9.4%	689	-55.2%	450	-17.4%	1341	-41.3%
55-64	304	1485	411	2200	1105	-25.6%	374	-9.0%	1783	-19.0%	445	-70.0%	374	-9.0%	1123	-49.0%
65+	316	1019	265	1600	650	-36.2%	265	0.0%	1231	-23.1%	425	-58.3%	265	0.0%	1006	-37.1%

Share of consumption

		Normal economy			10 year recession			Lasting recession		
Cohort	**Initial wealth†**	**Asset income**	**Labour income**	**Consumption**	**Asset income**	**Labour income**	**Consumption**	**Asset income**	**Labour income**	**Consumption**
<35	3-6%	51%	46%	100%	59%	38%	100%	36%	58%	100%
35-44	5-10%	63%	32%	100%	62%	32%	100%	41%	49%	100%
45-54	9-15%	67%	24%	100%	68%	22%	100%	51%	34%	100%
55-64	13-27%	68%	19%	100%	62%	21%	100%	40%	33%	100%
65+	19-32%	64%	17%	100%	53%	22%	100%	42%	26%	100%

†: range of values for the three scenarios. Actual values can be computed per scenario by subtracting lifetime asset income and lifetime labour income from lifetime consumption.

Table 7: Simulation results for $\beta = 0.99, \sigma = 1, \theta = 1$

This table summarizes our simulation results for our household birth cohorts (the age of the head of the household is mentioned in the cohort column) using various setups and various parameter values. Initial wealth, labour income, asset income and consumption are all expressed in thousands of 2011 euro and concern lifetime present value aggregates for 2007. Labour income and consumption were taken from the simulations, asset income was computed by subtracting labour income and initial wealth from consumption. Figures have been discounted using β. Labour supply is fixed at one (full time employment) here as $\theta = 1$.

Household lifetime indicators

		Normal economy			10 year recession						Lasting recession					
Cohort	**Initial wealth**	**Asset income**	**Labour income**	**Consumption**	**Asset income**		**Labour income**		**Consumption**		**Asset income**		**Labour income**		**Consumption**	
	PV	PV	PV	PV	PV	Change in %	PV	Change in %.	PV	Change in %	PV	Change in %.	PV	Change in %.	PV	Change in %
<35	54	1488	2575	4117	1288	-13.4%	1980	-23.1%	3322	-19.3%	427	-71.3%	1317	-48.9%	1798	-56.3%
35-44	136	2536	2290	4962	2072	-18.3%	1863	-18.6%	4071	-18.0%	672	-73.5%	1548	-32.4%	2356	-52.5%
45-54	202	2281	1505	3988	1906	-16.4%	1262	-16.1%	3370	-15.5%	592	-74.0%	1210	-19.6%	2004	-49.7%
55-64	304	2045	733	3082	1533	-25.0%	663	-9.5%	2500	-18.9%	488	-76.1%	663	-9.5%	1455	-52.8%
65+	316	1186	310	1812	784	-33.9%	310	0.0%	1410	-22.2%	487	-58.9%	310	0.0%	1113	-38.6%

Share of consumption

		Normal economy			10 year recession			Lasting recession		
Cohort	**Initial wealth†**	**Asset income**	**Labour income**	**Consumption**	**Asset income**	**Labour income**	**Consumption**	**Asset income**	**Labour income**	**Consumption**
<35	1-3%	36%	63%	100%	39%	60%	100%	24%	73%	100%
35-44	3-5%	51%	46%	100%	51%	46%	100%	29%	66%	100%
45-54	5-10%	57%	38%	100%	57%	37%	100%	30%	60%	100%
55-64	10-21%	66%	24%	100%	61%	27%	100%	34%	46%	100%
65+	18-28%	65%	17%	100%	56%	22%	100%	44%	28%	100%

†: range of values for the three scenarios. Actual values can be computed per scenario by subtracting lifetime asset income and lifetime labour income from lifetime consumption.

9. Appendices

Appendix A

To prove that the gains from stabilization gains are equal to $\frac{\theta}{2}\bar{C}^{-\theta-1}\sigma_C^2$, we take as utility function

$$U(C) = \frac{C^{1-\theta}-1}{1-\theta}, \qquad \theta > 0$$

Using a second order Taylor approximation around $C = \bar{C}$ we get:

$$U(C) \approx \frac{\bar{C}^{1-\theta}-1}{1-\theta} + C^{-\theta}(C - \bar{C}) - \frac{\theta}{2}\bar{C}^{-\theta-1}(C - \bar{C})^2$$

We now wish to approximate the average utility around the mean of consumption, so we choose $\bar{C} = E[C]$ and take expectations on both sides of equation 2.

$$E[U(C)] \approx E\left[\frac{\bar{C}^{1-\theta}-1}{1-\theta} + C^{-\theta}(C - \bar{C}) - \frac{\theta}{2}\bar{C}^{-\theta-1}(C - \bar{C})^2\right]$$

$$E[U(C)] \approx \frac{\bar{C}^{1-\theta}-1}{1-\theta} + C^{-\theta}(E[C] - \bar{C}) - \frac{\theta}{2}\bar{C}^{-\theta-1}E[(C - \bar{C})^2]$$

$$E[U(C)] \approx \frac{\bar{C}^{1-\theta}-1}{1-\theta} + C^{-\theta}(\bar{C} - \bar{C}) - \frac{\theta}{2}\bar{C}^{-\theta-1}\sigma_C^2$$

$$E[U(C)] \approx \frac{\bar{C}^{1-\theta}-1}{1-\theta} - \frac{\theta}{2}\bar{C}^{-\theta-1}\sigma_C^2$$

This means eliminating all variation, i.e. setting $\sigma_C^2 = 0$, would raise average utility by $\frac{\theta}{2}\bar{C}^{-\theta-1}\sigma_C^2$.

Appendix B – Aggregate wealth statistics

Table 7: Descriptive panel statistics of household indicators 2006-2011

This table gives an overview of the asset and debt accounts that were used to construct aggregate wealth statistics. All accounts are taken from the DHS survey data and aggregated per household.

Assets	Debts
Savings accounts:	**Mortgage debt:**
Checking accounts	Primary residence
Employer-sponsored savings accounts	Secondary residence
Savings or deposit accounts	Other real estate
Deposit books	
Savings certificates	
Single-premium annuity insurance policies	
Savings/endowment insurance policies	
Financial assets (basic):	**Other debt:**
Mutual funds	Private loans
Bonds	Extended lines of credit
Equity (includes stocks from substantial holdings	Finance debts
Options	Credit card debts
	Other credit
Other-financial assets:	Loans from family and friends
Business equity	Study loans
Money lent out to family and friends	Debts not mentioned before
Cash value of life insurance on mortgages	
Other savings and investments	
Real estate:	
Primary residence	
Secondary residence	
Other real estate	
Non-financial assets:	
Cars & motorbikes	
Boats	
Caravans/trailers	

Appendix C – Data construction

Wealth data

In the DHS data files, figures for all asset and debt components, besides real estate, are reported by all members of the household. This includes children, who often have many zero values and only a few small non-zero values for savings accounts. Aggregation by household is simple, per account we add the values for all household members.

For missing values for adults we have no choice to omit the household as ignoring these missing values would constitute to assuming zero values. For this there needs to be a good reason, which we do not have. For missing values in the accounts of children (i.e. non-adults) we can assume zero values in case of missing, as assets and debts for them are generally insignificant. We find however almost no (single digit values per year) missing values for children, so this is a very small issue. For adults we do find missing values, sometimes for more than one per household. For statistical safety we omit these households, even if it they only have a missing for one account. Although this may sound highly restrictive, since one missing value in only one asset or debt account for a single household member would invalidate a complete household, this is not much of an issue. Missing values do not occur for single accounts in the files. If a household member has a missing value, it has them in more accounts, generally all of them.

For real estate this approach will not work, for questions about residences have only been asked to one adult. For the relevant accounts (value of residence, mortgage value and cash value of life insurance on mortgages) we therefore take those values as the household aggregate. Some households have proper missing values in these accounts, these are again omitted. Other households have erroneous mortgage values reported due to data artifacts from the processing of data by CentERdata. Some of these values can be imputed however by looking at values in adjacent years. In case of zero values in these years, we assume a zero can imputed, i.e. the value is interpolated. For non-zero value none of the values could be realistically imputed, as the adjacent years were too different. These households were in that case also omitted.

The "other real property" account in the database suffers from another problem. Questions about possession of other real property besides the primary and secondary residences were answered by more than one person per household. We sometimes find identical non-zero values among this variable. This is suspicious, since one of these values is likely a duplicate. For common figures (50 thousand, 10 thousand, etc.) we left these in, but for non-standard figures, we decided to drop one of the values to avoid counting assets twice.

Finally, households with duplicate members were omitted completely, for we could not reliably uncover the nature of this problem. We do not know whether a member was simply duplicated and added or whether values of another member were completely overwritten by values of another household member.

After these selections and omissions we are on average per year left with almost 75% of the existing households. When we switch to panel statistics we lose almost 10% of the original households, so we are left with 65% on average per year for our wealth statistics.

Income data

For the income data we broadly use the same procedure as for the wealth data, we aggregate our variables of interest by household. As income data is only reported for adults, we do not suffer from incidental missing in children.

For asset returns we sum various accounts related to interest income and income from stocks, income from real estate and subtract interest payments. Mortgage interest should be reported by only one member of the household, as questions about mortgages and real property were only posed to one member (see wealth data). Therefore we set mortgage interest payment duplicates (within a household) to zero. Mortgage interest payments also suffer from data processing artifacts, i.e. values have been coded into an artificial value. In case we can infer from other variables and other years whether the household does not pay any mortgage interest, we recode these to a zero values. In other cases we have no choice but to drop the household. With respect to the income from real property and interest income we do the same if they have been coded into an artificial value; we recode them to zero if we cannot find any way to impute them by using other variables. Interest payments do not suffer from this issue, so we can then create asset return by household.

To compute capital gains we require household balances of real estate and stocks. We merge the wealth and income data and subsequently approximate the change in price from real estate and stocks, using only the households that have complete wealth data. We then add household asset return and capital gains.

Asset return percentages

For rescaling asset income to percentages we use the sum of return generating assets as the denominator. Return generating assets are composed of (bank) savings, the total amount of stocks and the amount of real estate. Some household observations in this case yield unrealistic percentages, thousands or even millions of percentage points. This cannot be avoided, for the data is not good enough for a completely consistent picture. Therefore, for a realistic image, we cap returns at ±10%.