Income Inequality and Aggregate Saving

Niklas Amberg Erik Öberg*

Bachelor's Thesis
Stockholm School of Economics
Department of Economics
May 17, 2010

Tutor: Lars Ljungqvist

Abstract

We suggest that the prevalence of social norms of minimum consumption can result in consumption behaviour that differs from that predicted by the permanent income/life cycle hypothesis. We model such social norms as a belief constraint in the intertemporal consumption decision and show that if the resulting threshold effect is present at low income levels, aggregate saving will increase as income becomes more equally distributed. Evidence from US consumption data suggests that a threshold effect is prevalent at low income levels. Using data from the World Income Inequality Study, we are not, however, able to show any statistically significant relation between income inequality and aggregate saving in the OECD countries.

* Amberg: 21344@student.hhs.se. Öberg: 21381@student.hhs.se.

Acknowledgements

We would especially like to thank our tutor Lars Ljungqvist, professor at the Stockholm School of Economics and New York University, for his guidance during the entire course of writing this essay. His insightful comments as well as his general encouragement have certainly helped to improve our thesis. We also want to thank Markus Jäntti, professor at Stockholm University, for helping us find appropriate data, and David Domeij, associate professor at the Stockholm School of Economics, for helpful comments at the early stage of idea formulation. Finally, we want to thank Karin Hederos Eriksson, Dany Kessel, and Anna Sandberg for reviewing our thesis and also Ylva Ersvik and Tomas Öberg for lingual corrections.

All faults and flaws that remain are fully the responsibility of ours.

Table of Contents

I Introduction

Aggregate saving is a key variable in the macro economy. As a determinant of investment, saving affects output both in the short and the long run. Establishing what the determinants of saving are is therefore an important task for economists, but it has proven to be a difficult one. A specific puzzle is that saving rates differ substantially across countries. For instance, in 2006 the gross domestic saving rate in Sweden was 23 percent while in Ireland 38 percent (The World Bank Development Indicators 2010). Economists still struggle with explaining these differences.

Early in this field of research it was suggested that the distribution of income can affect the aggregate saving rate. This notion was based on two postulates put forth by Keynes in his *General Theory* (1936); from his own experience of human psychology, he believed that the marginal propensity to consume (mpc) is always greater than zero but less than one and also that the average propensity to consume (apc) declines with income.[1] This implies that a more unequal distribution of income will increase the aggregate saving rate. [2] Kaldor (1957) showed that the same result could be obtained by assuming that mpc declines with income. These assumptions also imply that as aggregate income increases, the aggregate saving rate will increase as well. Kuznets (1942) and Goldsmith (1955) showed, however, that long term data on US consumption did not support this, which shed doubt on the Keynesian theory of consumption and thus also on the hypothesis that income distribution affects the aggregate saving rate.

By exploring the micro foundations of consumption behaviour, Friedman (1957) and Modigliani and Brumberg (1954) showed that saving should not be understood as a function of current period income, as Keynes stipulated. According to these authors, the economic motives of saving are to smoothen life time wealth over the life cycle, earn interest and manage income risk. This implies that consumption and saving rates in each period are constant fractions of lifetime wealth (the permanent income/life cycle hypothesis). In this framework, the income distribution does not affect the aggregate saving rate.

The explanatory power and logical appeal of the Brumberg-Modigliani-Friedman framework (henceforth the BMF framework) have been widely acknowledged; it explains the findings of Kuznets and Goldsmith as well why we might observe higher saving rates among the rich in cross-sectional data. As a consequence, it has become the standard theory of consumption behaviour. This notwithstanding, there have been numerous attempts to relate income distribution to aggregate saving within the BMF framework. There is no consensus, however, about what the relation should be; some theories point to a positive relation between inequality and aggregate saving while others suggest the opposite effect. The empirical results are ambiguous as well (Leigh and Posso 2009).

While the BMF framework properly explains the consumption behaviour of the homo oeconomicus, it fails to recognise the motives of the homo sociologicus. On a general level, we might expect most economic behaviour to be influenced by social structures as well as strict

[1] The mpc is defined as the derivate of consumption on income. The apc is defined as current period saving divided by current period income.

[2] To see the direct relation between consumption and saving propensities, note that income I can be used for consumption C or saving S; $I = C + S$. As such, the marginal propensities of consumption and saving must sum to one, $\frac{dC}{dI} + \frac{dS}{dI} = \frac{dI}{dI} = 1$, and we find that $\frac{dS}{dI} = 1 - \frac{dC}{dI}$. Analogously, the average propensities of consumption and saving must sum to one: $\frac{C}{I} + \frac{S}{I} = \frac{I}{I} = 1$, and we find that $\frac{S}{I} = 1 - \frac{C}{I}$. Keynes two postulates is simply put forth as $0 < \frac{dC}{dI} < 1$ and $\frac{d}{dI}\left(\frac{C}{I}\right) < 0$.

economic motives. There is, however, a lack of research on how social structures might affect intertemporal consumption behaviour.

Attempting to fill this gap with a new theory of intertemporal consumption behaviour as well as with an unused and improved data set on income inequality, we hope to shed new light on the old and ambiguous question of if and how the distribution of income is related to aggregate saving.

The rest of this paper is organised as follows: Section II presents previous research on income inequality and aggregate saving. Section III presents a social norm hypothesis of consumption behaviour, and shows how this implies a relationship between income inequality and aggregate saving. Section IV presents results from our empirical investigations. Section V summarises.

II Previous Research

A number of attempts to explain saving as a function of income within the BMF framework have been made. Blinder (1975) made two important contributions. Firstly, he showed that mpc need not be constant within the BMF framework if bequests are incorporated into the utility function. Secondly, he emphasised that if mpc is a function of income, all analyses in the life cycle framework are done *ceteris paribus*, i.e., they do not consider how transfers of national income between different socioeconomic groups affect aggregate saving. If different socioeconomic groups behave differently as consumers and income is transferred between them, one cannot use such models for analysing the effects of income redistribution on saving. This was the first recognition in this field of research that saving might be affected by social structures. In the light of Blinder's study, Menchik and David (1983) estimated the marginal propensity to bequeath for different income groups, and found that it is greater for the rich than for the poor. This suggests that a more unequal distribution of income will increase aggregate saving. The effect was not, however, significantly present for redistribution of income within the four bottom income quintiles, but only when shifting incomes from the four bottom quintiles to the top quintile.

Musgrove (1980) further developed the Keynesian idea of an income-declining apc in a model that does not directly contradict the BMF framework. In his model, apc is dependent on how far an individual is from a threshold subsistence level, and he argued that the propensity to consume incomes above this threshold level declines to a constant fraction as income increases (in this sense, the model becomes consistent with the BMF framework at high incomes). As a consequence, it is only the distribution of income above the subsistence level that affects aggregate saving, and because of the income-declining property of apc, aggregate saving increases as the economy becomes more unequal.

Duesenberry (1949) proposed an alternative model in which individuals' consumption is a function of how large their income is relative to other's income. This was the first attempt to explore the possibility of social esteem motives in intertemporal consumption decisions. Given this, Duesenberry hypothesised that the more time individuals spend with people richer than themselves, the larger a fraction of their income they will consume. Hence societies with unequal income distributions, but nonetheless with substantial social interaction between the rich and the poor, will have high aggregate consumption rates and low aggregate saving rates.

With the development of endogenous growth theory, studies on how income distribution affects saving and growth through political redistribution processes have been undertaken

(Alesina and Rodrik 1994, Persson and Tabellini 1994). According to these studies, the median voter profits more from redistribution of wealth through capital taxation as income becomes more unequally distributed. A higher capital tax will, in turn, distort incentives and steer disposable resources from investment to present consumption, which leads to lower aggregate saving. Neither the relation between income distribution and redistributive processes nor the relation between income distribution and investment were investigated empirically, but both studies find a significant negative relation between income inequality and growth. This is, in our opinion, poor evidence for their theory, due to the problem of establishing causality.

Krusell and Smith (1998) have investigated how income heterogeneity can affect macroeconomic aggregates in a dynamic stochastic growth model. Allowing for shocks in aggregate productivity, they investigate how precautionary saving is related to income and wealth. Except for the poorest, they find that saving is almost completely independent of current period income, and since the poorest affect only a small part of the aggregates, they draw the conclusion that the macro economy can be described almost perfectly using only the mean of wealth distribution.

The early empirical studies on income inequality and aggregate saving (Blinder 1975, Della Valle and Oguchi 1976, Musgrove 1980) found negative, although insignificant, relationships in general. However, due to poor data, these results are probably unreliable. Deininger and Squire (1996) produced an enhanced and improved data set on income distribution which was used by Li and Zou (2004). They found a negative non-significant relationship when using all available countries in the data set, but a positive non-significant relationship when using only the subset of OECD countries. The Luxemburg Income Study has also provided improved data sets, and has been subject of research by Leigh and Posso (2009). They found a negative relationship by using the share of income of the top ten percent and top one percent earners, respectively, in the economy as a measure of inequality. Their results are not, however, robust to adding interest rates and country-specific effects as control variables.

III Theory

The BMF framework is based on three decisive assumptions. Firstly, that individuals only have economic motives in mind when making decisions. Secondly, that they have the ability to plan over their entire lifetime, and thirdly, that they can stick to these plans. However, sociologist Carol Walker (1994) summarised a number of surveys on social assistance recipients in the United Kingdom, and found that only twelve percent of the subjects considered themselves able to manage their economy. The rest explained that they were living on the edge, with no ability to make long term plans or to save money for unexpected expenses. Since these people have incomes far above the absolute subsistence level (in 2008, The World Bank estimated the global poverty threshold at purchasing-power-parity to be 1.25 USD a day), there are reasons to believe that there are motives other than pure economic ones that influence these people's intertemporal consumption decisions. Walker's study suggests that the BMF framework might not be suitable for analysing consumption behaviour of the relatively poor in high-income countries.

Which are these other motives? One could think that for some reason, poor people are less economically rational than the rich. Walker found that this is the general belief of claimants, social security staff and wider society. If the poverty were due to some form of economic shock

(such as sudden unexpected unemployment), one could understand temporary irrational behaviour and lack of planning ability, but since most of the social assistance recipients live on very limited resources over long periods of time and are forced adjust their way of living, such an explanation is probably not correct. Alternatively, we can imagine a massive time inconsistency among the poor people interviewed. The ability to stick to optimal consumption plans could be lower among the less educated. Since these are more likely to be found amongst the poor, one could expect some of the deviation to be explained by time inconsistency. However, it is difficult to motivate that the lack of ability to save anything at all is due to this.

Rather, to understand the deviating saving behaviour of the poor, we suggest that it is caused by a general constraint in beliefs that affects everybody. While it is possible to survive on incomes close to the absolute subsistence level, it cannot be considered humanely to be forced do so in a high-income country. The economic and social environment implies that you should wear a certain type of clothes, eat a certain type of food and so on. The mere awareness of a society's wealth can form a norm of minimum consumption. We suspect that Walker's findings can be explained by a belief of that it is not possible to live on less than a certain amount of money, say, 20-30 USD a day. While this belief constraint does not origin in the direct physical needs for survival, it works in the same way, since the prevalence of a minimum consumption norm makes it impossible to imagine an existence with less than this amount of consumption a day. Hence, people with incomes near this level will not make long term plans to protect themselves from unexpected expenses and income losses.

On which level this potential consumption norm exists is an empirical question, and can probably differ between countries. We suggest that the awareness of a society's wealth, or rather the awareness of other people's wealth, is an important factor in forming the consumption norm. As GDP grows, we thus expect the perceived minimum consumption level to increase. Moreover, from this perspective, the spread of information about consumption patterns among a society's members will probably also affect these consumption norms. Duesenberry's suggestion, that more interaction between the rich and the poor will lead to increased consumption among the latter group, can probably work in a similar manner here.

Moreover, not only the direct interaction, but also the exposure from media could strengthen these norms. There has been a number of studies suggesting that media is important in forming our picture of reality (Hadenius and Weibull 1993). Thus, if media informs about society's wealth and the consumption patterns of the rich, we are inclined to suspect that this helps form the notion of what is considered "a dignified lifestyle" or what is believed to be the minimum level of consumption required to be fully accepted in society. The relative importance of media in establishing these consumption norms can only be guessed here.

The notion of social structures affecting consumption decision of the poor is also related to the "Catching up with the Joneses" literature (see, for example, Ljungqvist and Uhlig (2000)). Similarly, we suggest that social factors cause the propensity to consume to be connected with location in the income distribution. In the "Catching up with the Joneses" utility function, individuals experience disutility by consuming less than the average consumer. This, however, still constitute a motive for smoothening consumption over time. Hence it cannot explain Walker's findings and therefore we do not add further preference variables to the utility function. Instead, we argue that the BMF framework is the proper model for understanding consumption behaviour, when a belief constraint of minimum consumption is added to the intertemporal consumption decision.

Whether the time inconsistency or the social norm explanation is correct, the phenomenon of poor people saving less can be related to the interpretation of the Keynesian hypothesis of income-declining mpc and apc. Although the idea of social norms cannot account for a declining effect across the entire span of income, it explains the Keynesian intuition of why the poor save less than the rich. To see the difference between the Keynesian hypothesis and our theory, consider the idea of Musgrove (1980). He assumed that saving will go to a constant fraction as income goes to infinity. However, if explaining differences in consumption between the poor and the rich by the prevalence of social norms, consumption should assume a constant fraction of income after a given income level.

The direct implication of the consumption norm hypothesis is that income distribution will affect the aggregate saving rate. If there is a social norm of what is considered minimum consumption in an economy, the saving behaviour of people with income near this minimum level will differ from those far above it. In order to keep up with the consumption norm, poor people will not save in order to protect themselves from income losses, such as retirement and unemployment, or unexpected expenses. A more unequal distribution will then increase the fraction of people restricted by this consumption norm, and thus decrease aggregate savings.

To see this more clearly, we present a simple model.[3] Assume every individual lives for two periods, $t \in \{1,2\}$ (interpret this as the individual's life as young and old, or as employed and unemployed), and act to maximise their expected lifetime utility U. Each individual earns an income in the first period but none in the second. Utility streams from consumption of a single good in a logarithmic utility function.[4] The consumer's problem can then be formulated as

$$max\ U = log(c_1) + \beta log(c_2) \tag{1}$$

where c_t consumption in period t and β the discounting factor. For simplicity, we will not assume β to be subjective, i.e. all consumers discount future consumption at the same rate.

Assume the consumer earns total lifetime assets y in period 1 and can save a fraction of this income to period 2. However, she will consume her entire income in period 1 if it is below a threshold value φ due to the prevalence of a social norm of minimum consumption.[5] We capture this in the budget constraint:

[3] Step-by-step derivations can be found in Appendix A

[4] The logarithmic utility function is chosen for simplicity. The result would be the same using whichever utility function with concave property.

[5] If instead modeling the deviation in intertemporal consumption behavior among the poor due to time inconsistency, simply assume β to go to zero when consumption in period 1 is below φ. The result will be the same.

$$
\begin{aligned}
c_1 + \frac{c_2}{1+r} &= y \\
c_1 &\geq \varphi \;\; if \;\; y > \varphi \\
c_1 &= y \;\; if \;\; y \leq \varphi
\end{aligned}
\tag{2}
$$

where r is the exogenous interest rate. Solving the maximisation problem yields the individual's average propensity to save (aps):

$$
aps = \begin{cases} 0 \;\; if \;\; y \leq \varphi \\ 1 - \dfrac{\varphi}{y} \;\; if \;\; \varphi < y \leq (1+\beta)\varphi \\ \dfrac{\beta}{(1+\beta)} \;\; if \;\; y > (1+\beta)\varphi \end{cases}
\tag{3}
$$

A schematic graphic illustration of the solution is provided below.

Graph 1. Average propensity to save in the intertemporal consumption model with social norms

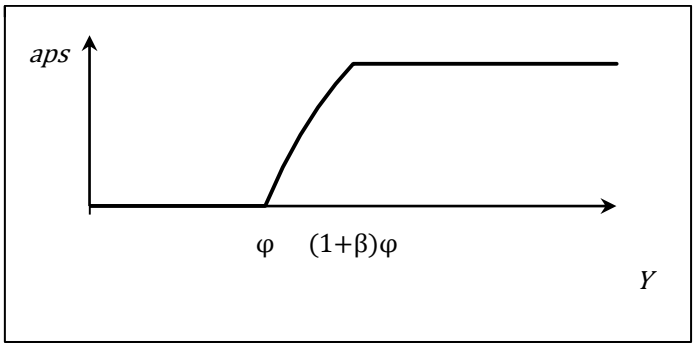

We find the aggregate average propensity to save (APS) by integrating individual aps over income

$$
APS = \int_{\varphi}^{(1+\beta)\varphi} \left(1 - \frac{\varphi}{I}\right)\frac{y}{Y}N(y)dy + \frac{1}{\bar{y}}\int_{(1+\beta)\varphi}^{\infty} \frac{\beta}{(1+\beta)}\frac{y}{Y}N(y)dy
\tag{4}
$$

$N(y)$ is the number of people with income y and Y is total GDP.

The economy consists of three income groups; those with income below φ who do not save anything at all, those with income between φ and $(1+\beta)\varphi$ with aps $(1-\frac{\varphi}{I})$ and finally those with income above $(1+\beta)\varphi$ with aps $\beta/(1+\beta)$. Since $(1-\frac{\varphi}{I})$ always is equal to or less than $\beta/(1+\beta)$ within the integral limits, we can draw two conclusions about how the distribution of income will affect the aggregate savings. Firstly, we see from the second term that redistributions of income above the numerical value $(1+\beta)\varphi$ do not influence APS. People with income above this level are not affected by the belief constraint and save perfectly according to their utility function. However, by comparing the terms we see that redistributions that imply movement

from the top to the lower income groups do affect aggregate savings, as well as redistributions within the group with incomes between φ and $(1+\beta)\varphi$.

As (4) shows, the critical parameter for analysing redistribution effects on saving is the threshold value φ. If the average income is less than $(1+\beta)\varphi$, more unequal economies will have higher aggregate saving rates, since such economies will have more people in the two bottom income groups. If average income is greater than $(1+\beta)\varphi$, the opposite effect will be present. Also, as φ gets smaller relative to average income, redistributions of income will have smaller effect on the economy, since it implies smaller movements between the income groups. Moreover, since the contribution of an individual's saving to APS is the product of her aps and relative income y/Y, the aps of the rich is more important to the aggregate. This implies that the marginal effect of redistributions increases with φ.

If one interprets φ as a physical subsistence level of consumption, φ is very small compared to the mean income in a developed economy, which implies that income redistributions should not affect aggregate saving much at all. If φ is a social norm of minimum consumption, however, and thus is a significant number in relation to the mean income, the model shows that distribution of income could affect aggregate saving significantly.

IV Empirical Results

Evidence of Threshold Effects in Consumption from US Data

To properly investigate potential threshold effects in consumer behaviour against the BMF framework, either data that follows individuals' income and consumption over time or cross sectional micro data on individuals' incomes and consumption, accompanied by demographic data, would be needed. The Consumer Expenditure Survey (CES), conducted by the Bureau of Labour Statistics, have collected the latter type of data from American households for about three decades, and has publicly available data from 1984 and forth. Alas, we have not had access to their micro data, and have therefore had to use data on the level of income quintiles.

A number of problems arise when one uses this data to estimate saving rates. Firstly, after-tax income is overestimated due to missing tax data. Secondly, consumption is underestimated since it is not measured as total outlays, but as total expenditure, which do not include payments on loans and credit cards. This means that calculating saving as the difference between income and expenditure will overestimate the saving rate for all income groups. The overestimation turns out to be quite large; in the CES data, the aggregate saving rate for the U.S. in 2006 is about twenty percent, while it actually was fourteen percent according to the World Banks' World Development Indicators. As we mainly are interested in differences in consumption behaviour between different income groups, and not primarily in the level of total saving, this data could still be valid for our purposes as long as the distortions are not spread too asymmetrically across the income groups. It seems reasonable that this is the case but we cannot confirm it. Most important is that the saving rates of the poorest are not disproportionally overestimated, and since credit card expenses and loan payments are more likely to affect the upper income quintiles, this seems to be the case.

The fact that we have to use data on the level on income quintiles means that we cannot improve the CES data set in a way that others have done when using it. For example, Dynan et al

(2004) excludes all households with incomes below 1,000 USD as well as those with invalid income or missing age data and those that did not participate in all interviews. This notwithstanding, their data, like ours, shows an extremely low saving rate (about -100% of income) for the lowest income quintile. According to Dynan et al, this must reflect a bias from measurement error in income and perhaps transitory income, since it is not possible for an income group as a whole to sustain such a high rate of dissaving in the long run. However, since people can and do move between different income quintiles, it is actually possible for a quintile as a whole to dissave for long periods of time, theoretically for infinitely long periods. This means that the very low saving rate of the bottom income quintile in the CES data does not have to be due to biases in the data, at least not in its entirety.

Given this, using the CES data is a bit problematic, but not to the extent that our results become invalid. However, to firmly establish the empirical results presented in what follows, similar tests ought to be run on a more extensive micro data set.

We have calculated the saving rates for the income quintiles as the aggregate after-tax income minus aggregate expenditure divided by aggregate after-tax income in each quintile. In Graph 2, saving rates for US households per income quintile are shown for the period 1984-2008. For reference, nominal average income levels are shown in Graph 3.

Graph 2. Saving rates per income quintile in the US.

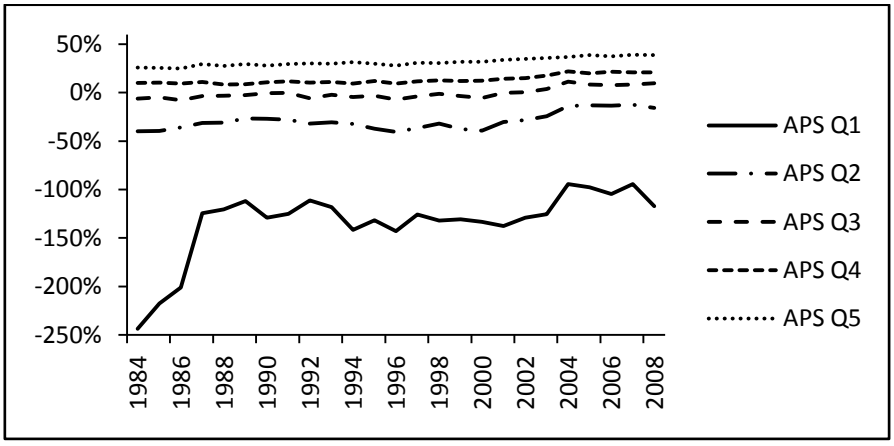

Graph 3. Nominal average yearly income per income quintile (in USD).

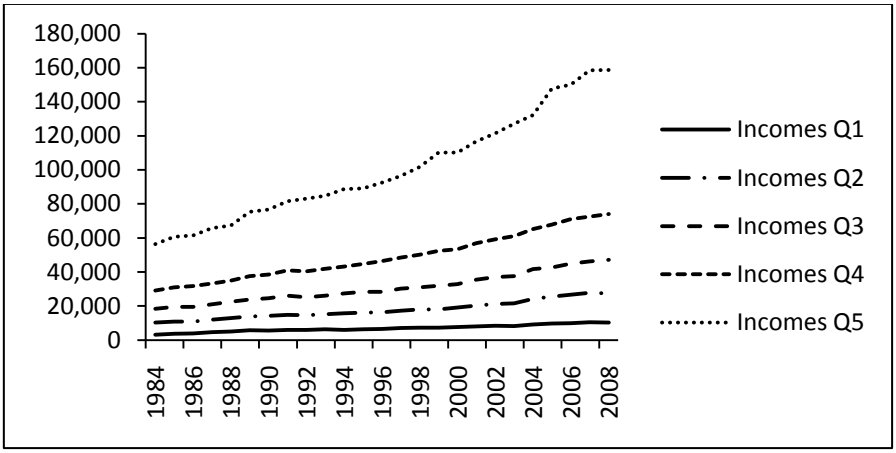

In Graph 2, we clearly see that the saving rates increase with income. We also notice that the gap between the first and second income quintile is much larger than between the other quintiles. This might be caused by the fact that apc declines with income as well as by a threshold effect between the first and second income quintile. To establish if this is the case, one has to control for the age distribution within the income quintiles, so as to be able to know if these differences are explained by the BMF framework. The CES provides data on the mean age of the reference persons and the share of people older than 65 years in the households for all income quintiles, which make such an investigation possible. Together, these two variables serve as a proxy for the age distribution in the income quintile. To test the explanatory power of the BMF framework against the hypothesis of income-declining apc and threshold effects in consumption, we use a linear regression model:

$$sav_i = \alpha + \beta_1 MEANage_i + \beta_2 Sh65pl_i + \varepsilon_i \quad (5)$$

sav_i is the average household saving rate in each income quintile, $MEANage_i$ is the mean age of the reference persons of the households in the income quintile and $Sh65pl_i$ is the average share of people over 65 in the households in the income quintile.

The regression was performed on two data sets in order to test if age distribution can account for the large difference in saving rates between the first and second income quintile. In the first regression, all observations across all income quintiles were included. In the second regression, only the four upper quintiles were included. The results are presented in Table 1.

Table 1. Regression of household saving on age distribution.

	(1)	**(2)**
MEANage	0.024	0.031*
	(0.027)	(0.004)
SH65PL	-6.786*	-4.145*
	(1.041)	(0.151)
Constant	-0.443	-0.948*
	(1.169)	(0.167)
Adjusted R^2	0.592	0.948
N	125	100

** Estimate significant on 1 % level. No other estimates are significant on any acceptable level.*

Both estimations yield a large, statistically significant coefficient on the $Sh65pl_i$ variable and an economically insignificant coefficient on the $MEANage_i$ variable. The second estimation, performed on the four upper income quintiles, shows that the age distribution explains 95 percent of the variation in household savings with significant estimates on the age distribution variables, which provides evidence in favour of the BMF framework and against the hypothesis of income-declining apc. When adding the first income quintile to the regression, the explanatory power decreases to 60 percent. Also the $MEANage_i$ variable becomes statistically insignificant. This suggests that the age distribution cannot account for the major savings gap between the first and second income quintile, which points to the prevalence of a threshold effect at this income level.

Since the threshold effect is prevalent substantially above the income needed for absolute subsistence, it cannot be interpreted as a result of the minimum consumption needed for survival. Rather, the data is consistent with Walkers' findings amongst the poor in United Kingdom, and as such, these results may be explained by the prevalence of a social norm of minimum consumption. This interpretation thus implies that the distribution of income will affect the aggregate saving rate in accordance with (4).

Bunting (1991) also used the CES data set to estimate the mpc of the income quintiles, and found that the propensities to save increase as the quintile share of total income increases. However, as shown, the differences in saving rates of the four upper income quintiles are explained by the BMF framework. Controlling for the age distribution, this data suggest that there is no support for the income-declining apc hypothesis and Bunting's conclusion.

Cross-Country Evidence on Income Distribution and Aggregate Saving

Now that we have indicative empirical support for the hypothesis that low-income earners save less than the rest of the population, we are ready to test if income distribution is related to aggregate saving. For this purpose, we propose a linear regression model:

$$GDS_i = \alpha + \beta_1 GINI_i + \beta_2 R_i + \beta_3 GDPgr_i + \sum_{j=1}^{n-1} \beta_{j+3} D_{ji} + \varepsilon_i \qquad (6)$$

GDS_i is gross domestic saving as a percentage of GDP, $GINI_i$ is the Gini coefficient, R_i is the real interest rate, $GDPgr_i$ is GDP growth, and D_{ji} country-specific dummy variables for the n countries in the data set.

The Gini coefficient is an aggregate measure of income inequality based on the Lorenz curve (the cumulative distribution function of income across households), where a larger number means a more skewed distribution.[6] 0 means perfect equality and 100 perfect inequality. We control for the real interest rate R since it is the price of future relative to present consumption, and is the only variable that affects the intertemporal consumption decision in the classical BMF framework with no risk. The real interest rate can affect saving in either direction, depending on the relative strengths of the income and substitution effects. The GDP growth variable is included to account for shock effects and income risk in the business cycle. Dummy variables are included to account for country-specific effects, such as tax rates, political structures, cultural differences in preferences etc.

The Gini coefficients have been obtained from the World Income Inequality Database (WIID), run by the World Institute for Development Economic Research. This data set is an improved and extended version of Deininger and Squire (1996), which summarises earlier studies on income inequality across the world. In the WIID, new data has been added from the Luxemburg Income Study, the Transmonee data from UNICEF/ICDC and specific research studies. The time span ranges differently for different countries, but from 1980 and onwards, there are regular observations for most OECD countries reported in the data set. The summarised data set in WIID has been harmonised in order to deal with differences in definition

[6] For a detailed description and discussion of different mathematical treatments, see Xu (2004)

of the Gini coefficients used when obtaining the data. Where multiple observations are provided, we use the mean of these in our data set.

Real interest rates and gross domestic saving rates have been obtained from the World Development Indicators database, provided by the World Bank.

In the estimation, all the 23 OECD countries that were classified as high-income countries in 1996 were used, with the exception of Iceland and the Republic of Korea, since they both had too few observations to be used.[7] The number of observations for each country included ranges from eight (Portugal and Luxemburg) to thirty-four (United Kingdom). No observations prior to OECD membership were used in the regression. Greece has an extremely large and unstable variation in GDS over time, between 1974 and 1988 it varied from 59, via -8, to 17 percent of GDP, and we thus suspect that it is an outlier. Summary statistics on Gini coefficients and GDS are presented in Table 2.

Table 2. Summary statistics on Gini coefficients and gross domestics saving (GDS) in the WIID/WDI set.

	N	**Mean**	**Std. Dev.**	**Min**	**Max**
GINI	675	33.3	7.45	16.6	69.2
GDS	589	24.5	6.65	-8	59
GDS (excl. Greece)	562	24.5	5.76	12	49

We run five regressions. In the first four, Greece is excluded from the sample. In the first regression, only the Gini coefficient is included as explanatory variable. In the following three regressions, interest rate, GDP growth and country-specific effect are added sequentially. For comparison, we include the Greece observations in the fifth regression. The results are presented in Table 3.

[7] The included countries are Australia, Austria, Belgium, Canada, Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Japan, Luxemburg, Netherlands, New Zealand, Norway, Portugal, Republic of Korea, Spain, Sweden, Switzerland, the United Kingdom and the United States.

Table 3. Regression of Gini coefficient on gross domestic saving (GDS).

	(1)	(2)	(3)	(4)	(5)
Gini coefficient	-0.205*	-0.247*	-0.316*	-0.024	0.199*
	(0.042)	(0.048)	(0.046)	(0.032)	(0.055)
Real interest rate	-	-0.286*	-0.234*	-0.256*	-0.417*
		(0.067)	(0.063)	(0.032)	(0.055)
GDP growth	-	-	0.779*	0.430*	0.560*
			(0.103)	(0.053)	(0.087)
Country dummies	-	-	-	. . .	. . .
Constant	30.986*	33.136*	32.739*	17.781*	8.737*
	(1.346)	(1.572)	(1.487)	(1.404)	(2.421)
Adjusted R^2	0.040	0.083	0.193	0.832	0.615
N	562	453	448	448	470

** Estimate significant on 1 % level. No other estimates are significant on any acceptable level.*

In the first regression, we see that the estimated effect of the Gini coefficient is significantly negative and large in magnitude; increasing Gini by five percentage points is associated with a one percentage point increase in GDS. Note that the explanatory power of Gini by itself is very small, with an R^2 of only 0.04. Adding real interest rate and GDP growth, the effect remains significant and large, but when including country-specific effects, the estimated Gini coefficient drops almost to zero and becomes insignificant. Though not reported in Table 3, all dummy variables are significant on the one percent level, except for Portugal and the United Kingdom in any regression. The estimated coefficients on the dummies are in general large (the mean effect is 7.96), implying that country-specific effects on saving are important. The estimated coefficients of the real interest rate R are negative and strongly significant, implying that the income effect outruns the substitution effect in this data set.

The adjusted R^2 decreases from 0.832 to 0.615 when Greece is added to the regression, which means that the aggregate saving rates of Greece are hard to explain with the used variables. Since OLS estimation is sensitive to outliers, including Greece probably makes the results unreliable for analysis.

The estimated regressions suggest that there is no significant relationship between income inequality and aggregate savings. We find no support for the hypothesis that potential social norms of minimum consumption affect the macro economy in a measurable magnitude. Moreover, the hypothesised positive relationship between income inequality and aggregate saving motivated by an income-declining mpc, as suggested by Kaldor (1957), Blinder (1975) and Musgrove (1980), is not supported in the WIID data, neither is the endogenous political effect proposed by Alesina and Rodrik (1994) and Persson and Tabellini (1994).

How do we explain such findings? The first possible explanation is that there is no relation between income distribution and the aggregate saving rate. It could be that Walker's (1994) findings and the support in the CES data are due to some form or error in measurement. Moreover, the hypothesis of income-declining mpc due to different propensities to bequeath has vague support in previous studies; the effect estimated by Menchik and David (1983) was only present when redistributing incomes from the four bottom quintiles to the top quintile. Similarly,

no causal relations have been established in the political redistribution models proposed by Alesina and Rodrik (1994) and Persson and Tabellini (1994).

A second explanation is that the data set used is too poor to serve our purposes. In the early estimations of Gini coefficients around the world, different studies often provided substantially different estimates for the same country a given year. However, most observations are from 1980 and onwards, a period in which the data shows much less discrepancies, which support its use.

A third explanation, and the one that we find most likely, is that even if there exist social norms of minimum consumption that break the saving pattern of the BMF framework amongst the poor, the saving of the poor constitutes too small a part of total aggregate saving for this to be visible in aggregate data. The social norm binds at too low a level of income for the deviating savings behaviour by those affected by it to influence APS. As seen in (4), the individual's contribution to APS is the product of her aps and relative income y/Y, making the saving of the rich more important than that of the poor. In the CES data, the income share of the first income quintile is only three percent of aggregate income, which means that even if the first income quintile saved only one half of the fraction saved by the rest of the population, aggregate saving would only decrease relatively by 1.5 percent.[8]

Our findings are consistent with those of Li and Zou (2004), who used the smaller data set from Deininger and Squire (1996), with a total of 180 observations for the OECD countries. The estimated coefficient on the Gini variable in their regression was larger in magnitude, but even less significant than what we found in the regression excluding Greece.

V Summary

Evidence from the CES on US consumers suggests that there exists a large threshold effect that breaks the smoothening mechanism in consumption at low income levels. We have suggested the interpretation that this is due to a social norm of minimum consumption. As shown in our simple model of rational agents, these results also imply that a more skewed income distribution will decrease aggregate saving.

We were not able, however, to find evidence of any relation between income distribution and saving on an aggregate level. We have presented three possible reasons for this. The one most likely is that, even if our suggested threshold effect exists and affects the saving behaviour of the poor in the way our model predicts, the share of total aggregate saving that is constituted by the poor is too small for this to be visible in aggregate data.

To further investigate the saving behaviour of the poor, micro data with extensive demographic information would be needed. Educational variables could help in investigating if the different saving behaviour of the poor can be explained by irrational economic behaviour (mainly time inconsistency). Further demographic variables would also be needed for investigating the explanatory power of BMF framework exhaustively. Moreover, our empirical findings could be deepened and nuanced by examining if and how social security affects the saving needs and thus the saving behaviour of the poor.

Our findings do not support a number of hypothesised relationships in the literature. The finding that age distribution explains the variation in saving rates in the CES data show that the

[8] If, for instance, the rest of the population on average saved 20 percent, aggregate saving would be 20*0.985=19.7 percent

conclusions of Bunting (1991) can be misguiding. The Keynesian theory of income-declining apc and mpc clearly has no support in the CES data, and its implication of income distribution affecting aggregate saving has no support in the WIID data. Moreover, the WIID data show no support of the endogenous political effect suggested by Alesina and Rodrik (1994) and Persson and Tabellini (1994).

References

Alesina, A.; Rodrik, D. (1994) "Distributive politics and economic growth", The Quarterly Journal of Economics, Vol. 109, No. 2: pp. 465-490

Blinder, A. S. (1975) "Distribution Effects and the Aggregate Consumption Function", The Journal of Political Economy, Vol. 83, No. 3: pp. 447-475

Bunting, D. (1991) "Savings and the Distribution of Income", Journal of Post Keynesian Economics, Vol. 14, No. 1: pp. 3-22

Deininger, K.; Squire, L. (1996) "A New Data Set Measuring Income Inequality", The World Bank Economic Review, Vol. 10, No. 3: pp. 565-591

Della Valle, P. A.; Oguchi, N. (1976) "Distribution, the Aggregate Consumption Function, and the level of Economic Development: Some Cross-Country Results", The Journal of Political Economy, Vol. 84, No. 6: pp. 1325-1334

Duesenberry, J. S. (1949) "Income, Saving and the Theory of Consumer Behaviour", Harvard University Press

Dynan, K. E.; Skinner, J.; Zeldes, S. P. (2004) "Do the Rich Save More?", The Journal of Political Economy, Vol. 112, No. 2: pp. 397-444

Friedman, M. (1957) "A Theory of the Consumption Function", Princeton University Press (for National Bureau of Economic Research)

Goldsmith, R. (1955) "A Study of Saving in the United States. Vol. 1", Princeton University Press (for National Bureau of Economic Research)

Hadenius, S.; Weibull, L. (1993) "Massmedier", 5th Ed., Bonnier Alba

Kaldor, N. (1957) "A Model of Economic Growth", Economic Journal, Vol. 67: pp. 591–624

Keynes, J. M. (1936) "The General Theory of Employment, Interest and Money", Macmillan

Kuznets, S. (1942) "Uses of National Income in Peace and War", National Bureau of Economic Research

Krusell, P.; Smith A. A. (1998) "Income and Wealth Heterogeneity in the Macroeconomy", The Journal of Political Economy, Vol. 106, No. 5: pp. 867-896

Li, H.; Zou, H. (2004) "Savings and Income Distribution" Annals of Economics and Finance, Vol. 5: pp. 245-270

Leigh, A.; Posso, A. (2009) "Top Incomes and National Savings", Review of Income and Wealth, Series 55, No. 1: pp. 57-74

Ljungqvist, L.; Uhlig, H. (2000) "Tax policy and Aggregate Demand under Catching up with the Joneses", The American Economic Review, Vol. 90, No. 3: pp. 356-366

Menchik, P. L.; David, M. (1983) "Distribution, Lifetime Savings and Bequests", The American Economic Review, Vol. 73, No. 4: pp. 672-690

Modigliani, F.; Brumberg, R. (1954) "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data" Post-Keynesian Economics, ed. Kenneth K. Kurihara, Rutgers University Press

Mugrove, P. (1980) "Distribution and the Aggregate Consumption Function", Journal of Political Economy, Vol. 88, No. 3: pp. 504-525

Persson, T; Tabellini, G. (1994) "Is Inequality Harmful for Growth?", The American Economic Review, Vol. 84, No. 3: pp. 600-621

Walker, C. (1993) "Managing Poverty" Routledge, London

Xu, Kuan: "*How Has the Literature on the Gini's Index Evolved in the Past 80 Years?*" Department of Economics, Dalhousie University. Available at [http://economics.dal.ca/RePEc/dal/wparch/howgini.pdf]

Appendix A: Step-By-Step Derivations

Solving the Individual Saving Problem (Eq. 1)

The individual saving problem means finding the interior solution of (1) on (2). This can be done by direct substitution of (2) in (1):

$$E(U) = log(c_1) + \beta log((1+r)(y - c_1))$$

The first order condition yields:

$$\frac{dE(U)}{dc_1} = \frac{1}{c_1} - \frac{\beta}{y - c_1}$$

$$\frac{dE(U)}{dc_1} = 0 \Rightarrow$$

$$c_1 = \frac{1}{1+\beta} y$$

To find the *aps* without further constraints, divide the difference between income and consumption in period 1 by total income:

$$aps = \frac{y - c_1}{y} = \frac{\beta}{1+\beta}$$

However, note that (2) constrain the solution by $0 \leq aps \leq 1 - \frac{\varphi}{y}$. We thus find (4):

$$aps = \begin{cases} 0 \;\; if \;\; y \leq \varphi \\ 1 - \dfrac{\varphi}{y} \;\; if \;\; \varphi < y \leq (1+\beta)\varphi \\ \dfrac{\beta}{(1+\beta)} \;\; if \;\; y > (1+\beta)\varphi \end{cases}$$

Aggregating the Individual Saving Function (Eq. 3)

The aggregate saving, S, is found by integrating (2) over income:

$$S = \int_0^{\infty} aps \cdot yN(y)dy$$

And we thus split up the integral to find (4):

$$S = \int_0^{\varphi} aps \cdot yN(y)dy + \int_{\varphi}^{(1+\beta)\varphi} aps \cdot yN(y)dy + \int_{(1+\beta)\varphi}^{\infty} aps \cdot yN(y)dy$$

$$= 0 + \int_{\varphi}^{(1+\beta)\varphi} (1 - \frac{\varphi}{y}) yN(y)dy + \int_{(1+\beta)\varphi}^{\infty} \frac{\beta}{(1+\beta)} yN(y)dy$$

$$= \int_{\varphi}^{(1+\beta)\varphi} (1 - \frac{\varphi}{y}) yN(y)dy + \int_{(1+\beta)\varphi}^{\infty} \frac{\beta}{(1+\beta)} yN(y)dy$$

Divide by national income Y to find APS:

$$APS = \int_{\varphi}^{(1+\beta)\varphi} \left(1 - \frac{\varphi}{I}\right) \frac{y}{Y} N(y)dy + \int_{(1+\beta)\varphi}^{\infty} \frac{\beta}{(1+\beta)} \frac{y}{Y} N(y)dy$$