



Resolution Foundation

Inequality, debt and growth

Report for the Resolution Foundation

Paolo Lucchino (NIESR) and Salvatore Morelli (University of Oxford)

May 2012

© Resolution Foundation 2012



Foreword

Like most economists, I internalised very early on in my career the concept of the "equity-efficiency tradeoff"; the idea that policy measures, such as progressive taxation or means-tested welfare benefits, that reduced income inequality would, because of their negative impact on incentives, be likely to reduce labour supply, effort and ultimately economic growth. At a theoretical level, this follows very clearly from the basic axioms of microeconomics. And even if we don't believe that people make their decisions according to those axioms, there is a kernel of truth to this idea: genuinely confiscatory taxation would reduce overall output because of the incentive impact. And since, in a civilised society, we don't want to let people starve if, for whatever reason, they don't have any money, society does need to recognise that there are both costs and benefits to policies that redistribute income.

However, along the spectrum of what is politically and socially feasible in a developed economy, it has been clear for some time that empirically there is no necessary correlation, let alone causation, between income inequality and growth at a macroeconomic level. Countries with high or low levels of inequality can sustain higher or lower growth. This in itself is an interesting puzzle, given the very clear predictions of microeconomic theory. And more recently, there has been considerable interest in the idea that for various reasons inequality could actually reduce growth. This line of inquiry was brought into sharp focus by the financial crisis. The sharp, although differentiated, increase in inequality observed in most developed countries over the last three decades has gone hand in hand with both the increasing importance of the financial sector and of course the explosion in earnings for many of those working in that sector; given the role of the sector in the crisis, it is at least plausible to hypothesise that there is a relationship between inequality, macroeconomic and financial developments, and growth and the crisis.

Both the rise in the importance of the financial sector and the rise in inequality were particularly marked in the UK. And the crisis and subsequent recession have been particularly severe here. So to understand why the crisis occurred, why the UK was so hard hit, and what policy should do now, it is essential to understand these interactions. This paper, prepared by NIESR for the Resolution Foundation, is one of the first contributions to our understanding in this area. While it does not claim to explain what happened, the evidence presented here does suggest that there may well have been a connection between the rise in income inequality in the years preceding the crisis and the rise in household borrowing, particularly for those on lower incomes. Of course, on its own this does not either explain the last crisis or tell us what policy measures will both put us on a path to sustainable growth and avoid a recurrence of the financial crisis. But it does I believe tell us that income inequality and income distribution matters for macroeconomics; when making policies that will affect the distribution of earnings and income, government and policymakers should not just be concerned with the microeconomic impacts, positive or negative, but with the broader long-term impacts on the sustainability of growth.

Jonathan Portes



Acknowledgment

The authors would like to thank Jonathan Portes (NIESR), Angus Armstrong (NIESR) and Tony Atkinson (University of Oxford) for their comments. They also thank Justin van de Ven (NIESR) for advice on the data. All errors and omissions remain the responsibility of the authors.

Contents

Introduction.....	1
Part 1 - The role of inequality in the run up to the recent crisis	2
Part 2 - UK context in the run up to the crisis	5
Part 3 - Limits to consumption and impact on the UK recovery	12
Conclusion	17
Annex 1 - Measurement of saving in the FES	18
Annex 2 – Simulation methodology	24

Introduction

The ongoing financial crisis and economic slump have re-opened the question of the link, if any, between inequality and growth. A large part of this literature has focused on the impact of inequality on consumption. In the past, inequality was seen as reducing growth through its dampening of aggregate demand as income moved away from the 'spending classes'. Attention has now turned to the alleged relationship between inequality and financial crisis, though consumption continues to play a role. In these new models, inequality leads to financial instability via an unsustainable surge in the level of households' indebtedness within the economy. As even a brief overview illustrates, the mechanisms at work in the inequality-household finance-growth nexus are therefore complex, though the attention they have attracted clearly reveals their importance.

This paper surveys the debate on how inequality can affect growth, focusing on the UK, both with respect to the origins of the recent crisis and to the policy debates on how to navigate through the current fragile stages of the recovery. Household indebtedness grew sharply in the pre-crisis years and arguably played a part in exacerbating, if not causing, the recent slump. In this paper, we carry out one of the first explorations into how inequality may have been a driver behind this rise in unsustainable debt. Going forward, high levels of inequality might reduce consumption and therefore growth. This argument is of particular relevance given that consumption made up around two thirds of UK growth in demand over the past two decades (BIS/DFID, 2011). It has therefore been argued that a more equal or broad-based income growth would have positive implications for the sustainability and strength of the UK economic recovery. In this paper, we carry out a simulation exercise to gauge for the likely quantitative extent of this effect

The paper is structured as follows. In Part 1, we present recent hypotheses and literature on the possible role of inequality in leading to the recent crisis. We then critically examine the long-term trends in inequality, debt, saving and growth in UK data to explore whether these hypotheses are relevant to the UK context (Part 2). In Part 3, we consider the challenges ahead of us and explore the extent to which continuing inequality and tight credit conditions may limit the growth of aggregate demand and the strength of the recovery.

Part 1 - The role of inequality in the run up to the recent crisis

The ongoing financial crisis and economic slump have spurred a new wave of research on the effects of inequality on growth. Given the specific traits of the recent crisis, this has tended to focus on the alleged relationship between inequality and financial crisis. A number of hypotheses on the mechanisms through which inequality may have increased the risk of financial crisis in the US have been put forward, somewhat informally, by Stiglitz (2009), Milanovic (2009), Fitoussi and Saraceno (2009, 2010) and Rajan (2010), among others, and formalised within standard modelling frameworks (Khumof and Ranciere, 2010). These hypotheses all see inequality as a driver behind an unsustainable surge in the level of household indebtedness. Such propositions therefore re-emphasise the role of household balance-sheets in macro-economic instability, an aspect which had previously been neglected in modern macroeconomic modelling frameworks (as pointed out in Jappelli, Pagano and Di Maggio, 2010). Some very recent work also attempts to empirically test the proposed link between inequality and crisis, as discussed in Atkinson and Morelli (2010 and 2011) and Bordo and Meissner (forthcoming). However, the hypotheses differ in the emphasis they place on whether inequality has led to an increase in debt via household demand for credit or via political economy motivations leading to an excessive supply of credit. In the following sections, we critically present each side in turn.

i) Demand-side arguments

Among the demand-side explanations, Stiglitz (2009) suggests debt surged among households in the lower part of the distribution as families sought to maintain a rising standard of living in the face of stagnating real incomes. A similar argument is presented in Khumof and Ranciere (2010). This borrowing later proved unsustainable, leading to default and pressure on over-extended financial institutions.

There are several ways in which trends in inequality may explain such an increase in the demand for credit. For example, increased inequality, as measured at a point in time, may be driven by income patterns becoming more erratic over time. In response, households increasingly resort to the credit market in order to smooth consumption, thereby increasing their stock of debt. Indeed, Iacoviello (2008) finds a clear link between the US wage inequality and the ratio of household debt to disposable income from 1960s to recent years and is able to explain both the trend and the cycle of households' debt level.¹ However, this is likely to be only part of the story, as it fails to provide a compelling justification for the striking increase in housing debt among US households before the crisis. Indeed, the evidence suggests that increases in income inequality were largely driven by increases in the dispersion of permanent incomes across households (see Jenkins, 2011, for the UK). Frank, Levine and Dijk (2010) attempt to explain some of the mechanisms through which such changes in permanent income inequality may lead to increasing household debt. Drawing on Duesenberry's (1949) relative-income hypothesis and the findings in the behavioural literature on demonstration effects, the model shows how over-indebtedness may arise as individuals emulate the consumption habits of their slightly wealthier peers, thereby triggering 'expenditure cascades'. They write: "...people do not exist in a social vacuum. ... the rich have been spending more ... Their spending shifts the frame of reference that shapes the demands of those just below them, who travel in overlapping social circles. So this second group, too, spends more, which shifts the frame of reference for the group just below ..." (Frank et al., 2010). Crucially, housing consumption is explicitly taken into account in their model. They also find empirical support for their theory, with evidence of the change in county-level Gini coefficient between 1990 and 2000 being positively associated with all their proxy measures of financial distress. In the context of this theory, the stronger gains that accrued to the top of the distribution during the past decades become a crucial feature, as they generate a ripple effect that can cause a decline in saving and an increase in debt among population groups lower down the income distribution. Therefore, as savings fall at each 'cascade' in the income distribution, inequality leads to a decrease in the aggregate saving rate. Note that this, as well as the model in

¹ Note that Iacoviello's paper is concerned with inequality of individual gross wages and not household disposable incomes.

Iacoviello (2008), contrast with the hypothesis made by Stiglitz (2009) that over-consumption is concentrated at the bottom of the distribution.

ii) Supply-side arguments

A second set of arguments posit that a demand-side argument is per se insufficient to explain the destabilising role of income inequality. For example, Khumof and Ranciere suggest that sustained gains going disproportionately to the top of the income distribution are reinvested within the financial market under the form of loans to those at the bottom of the distribution. Therefore higher inequality also influences the supply of credit. Other arguments proposed, for example by Rajan (2010) and Fitoussi and Saraceno (2009, 2010), introduce political economy arguments as drivers of increasing credit supply.

In this second set of propositions, inequality is the causal driver behind a form of 'political pressure' leading to policy choices which in turn lead to unsustainable household debt levels. For example, Fitoussi and Saraceno (2009, 2010) argue that increased inequality limited aggregate demand and prompted monetary policy to maintain low interest rates. This, in turn, allowed private debt to increase beyond sustainable levels. Similarly, Milanovic (2009) argues how the decades-long stagnation in middle class incomes gave rise to political pressure for redistribution. Rajan (2010) argues that the US Government responded by using regulatory tools to facilitate low income households' access to credit, particularly mortgages. In this set of theories, the increase in inequality at the bottom of the income distribution therefore has a causal effect in a chain reaction that, via financial deregulation, leads to increased financial instability.

However, there were clearly other sources of political pressure for financial deregulation, in particular from high-income individuals best placed to profit from it. The combination of deregulation and the search for investment returns may have led to the emergence of bubbles (Fitoussi and Saraceno, 2009, 2010). In other words, in this case it is increased inequality at top, rather than at the bottom, that may be driving the link between inequality, debt and instability.

iii) Inequality as coincidental with crisis

A number of theorists have argued that the trends in inequality and the occurrence of crisis are not causal but merely coincidental. The common experience of crises and inequality may be due to a third causal mechanism, often identified in the political arena (Acemoglu, 2011). For example, Krugman (2010) argues that both financial liberalisation and rising top incomes are a common result of a rightward shift in political thinking that has led to both rising inequality and a more vulnerable financial system. Atkinson and Morelli (2011) discuss other such hypotheses, such as financial sector remuneration policy (drawing on Cahuc and Challe, 2009) or the retrenchment of the welfare state.

iv) Empirical evidence

Efforts to empirically validate the posited relationship between inequality and crisis have been so far inconclusive. Atkinson and Morelli (2010 and 2011) carried out one of the first systematic empirical investigations into the hypothesis that inequality leads to crisis (broadly defined as including systemic banking crises and per-capita real GDP and consumption collapses). Using different measures of economic inequality for 25 countries and a long time-horizon, the authors did not find conclusive evidence supporting the hypothesis that a growing level of economic disparities leads to higher macroeconomic instability. Bordo and Meissner (forthcoming) use a more sophisticated empirical specification on a very similar group of countries, confirming the findings using growth rates of the top income shares and banking crises only. On the other hand, Berg and Ostry (2011) find that countries with more equal income distribution experienced longer growth spells.

Similarly, the link between inequality measures and household indebtedness levels has also been tested empirically. This appears particularly strong in Iacoviello (2008) and Frank et al. (2010) but not in Jappelli, Pagano and Di Maggio (2010) and Bordo and Meissner (forthcoming). However, it should be noted that the link appears weak in studies based on aggregate measures of debt and strong in studies


using evidence from micro-data. The latter is arguably more appropriate to the question at hand, and further such microeconomic investigation is therefore desirable.

Overall, the recent debate on the possible link between inequality and crisis is still in its infancy. While many hypotheses have been put forward, these have yet to be placed under the thorough lens of academic and empirical scrutiny. It is therefore still too soon to know where the weight of the evidence lies. The theories discussed above differ, for example, in the extent to which the increase in inequality is driven by the top of the distribution, the bottom, both or by changes across the distribution. If the mechanisms emphasised by Stiglitz, Kumhof and Rancier, Fitoussi and Saraceno have been at work, we would expect to see increasing debt levels and declining saving rates at the low end of the income distribution. Alternatively, hypotheses put forward by Frank and Iacoviello would be consistent with increasing debt and decreasing saving rates across the whole distribution. Which of these trends have occurred, both in the US and in other countries, and what effects they may have determined would be a welcome focus of further research.

Part 2 - UK context in the run up to the crisis


As an initial contribution to the above-mentioned research programme, in this section we explore whether similar trends as those presented in Part 1 have played out in UK data and whether the hypotheses presented might therefore also speak to the UK context.

We start by exploring the trends in inequality in the UK, as presented in Figure 1. The UK has seen a dramatic increase in inequality since the '70s, with much of this occurring in the '80s where both inequality and poverty saw rapid rises. Similarly, the secular wealth redistribution implied by declining top wealth shares also came to a halt in the '80s. It is interesting to note that the increase in income inequality and poverty in the 1980s is proportionately much larger than the increase in earnings dispersion, suggesting other factors, for example changes in household composition, employment patterns, and to the tax and benefit system, have been at work. Since then, there has been little increase in overall inequality (as measured in the Gini coefficient). However, incomes at the top continued to rise throughout the last two decades. Top income shares have more than doubled since 1979 and the ratio of top decile to median earnings has steadily increased over the same period.




Reproduced from Atkinson and Morelli (2012)

We use the Family Expenditure Survey (replaced by the Expenditure and Food Survey in 2001 and by the Living Costs and Food Survey in 2008) to explore trends in household finances since 1971 in more detail. An exploration of disposable income trends across income deciles confirm the above findings. Indexing the income level of each income decile at 100 in 1971 in Figure 2, we clearly see how the income levels of the top two deciles have risen faster than those of the other deciles, thereby confirming the increase in top income inequality over the period. On the other hand, the bottom two deciles track the fifth decile closely until the eve of the crisis, indicating their relative positions were unchanged over this period.



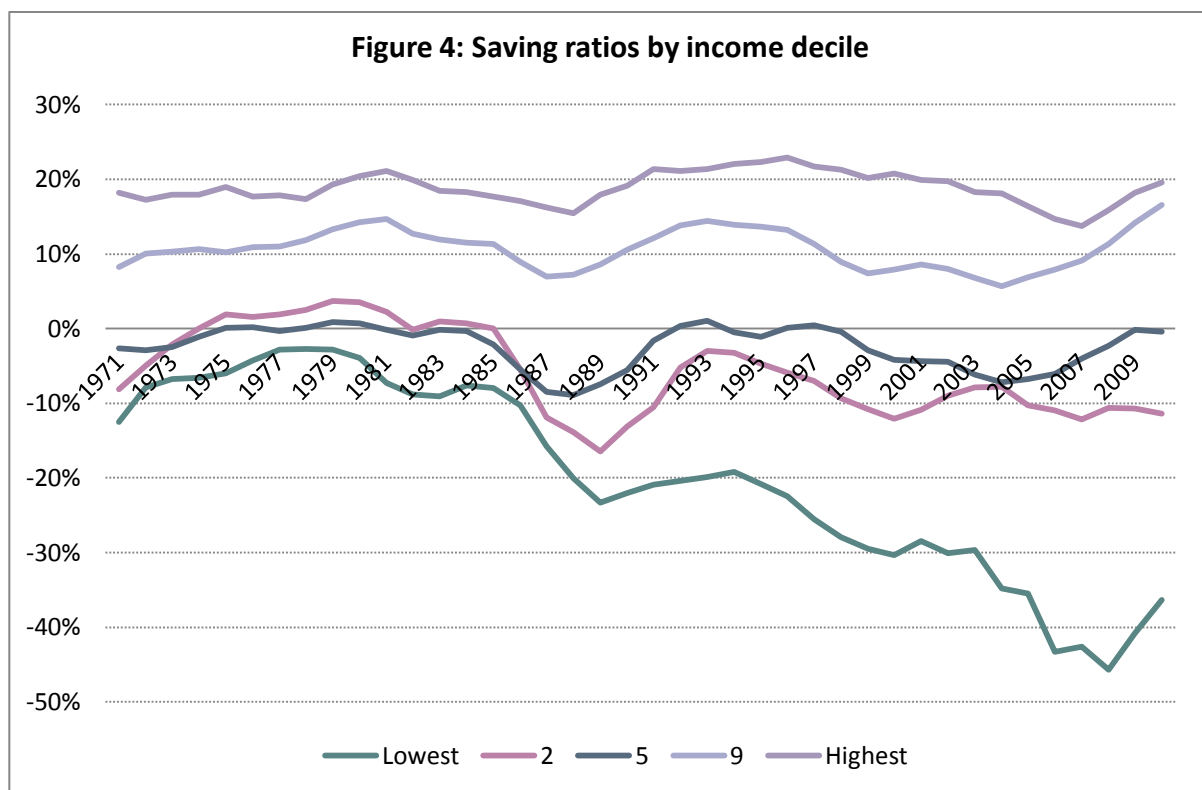
Source: Family Expenditure Survey and successive datasets, 1971-2010

The charts confirm that inequality in the UK has increased significantly in the recent decades. We now proceed to explore any correlations between this increase and trends in household finances, particularly household debt and saving. Figure 3 compares trends in inequality with measures of aggregate household debt-to-income ratios taken from National Accounts. The data clearly indicate that household debt has increased dramatically since the mid '90s. At the same time, top income inequality as measured by the share of the top 1% of incomes and the ratio of top to median earnings also increased while overall inequality, as measured by the Gini coefficient, also showed a mild upward trend from an already record-high level. Therefore, a link between inequality and debt cannot be ruled out, although deeper analysis of this relationship would require using survey data to obtain better evidence on which families have been most prone to accumulating debt and whether this can be traced back to trends in inequality. Unfortunately, this was not possible using the Family Expenditure Survey as it does not record data on debt.



Source: Atkinson and Morelli (2012), EUROSTAT

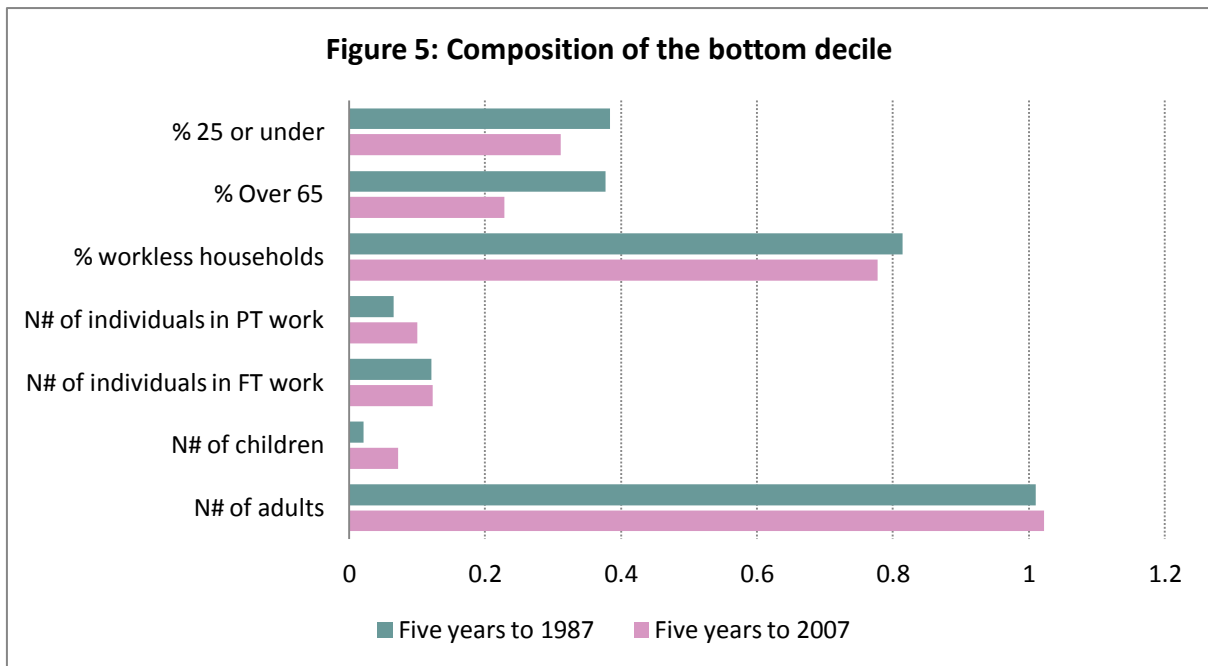
We can however explore saving trends since 1971 using Family Expenditure Survey data (see Annex 1 for methodological details and caveats). This survey gives detailed information on household expenditure and can therefore be used to measure saving defined as income minus consumption. Our income measure includes earnings as well as income from other sources such as benefits (excluding housing-related support), pensions and investments. Consumption is based on the FES total expenditure variable. This covers a comprehensive spectrum of expenditure categories including housing, mortgage and rent costs (net of any amount received in housing-related benefits). Defining saving as income minus consumption carries some limitations, in that it does not capture wealth trends and therefore the value of any capital gains. In other words, we measure a household's active saving, not any passive saving (see Annex 1 and Crossley and O'Dea, 2010 for details). Trends in active saving will to some extent be influenced by trends in other types of savings (e.g. the value of one's pension savings) and wealth trends more generally (e.g. housing or equity). In this sense, focusing on active saving only may give an incomplete picture of changes in saving behaviour. Nevertheless, it is still useful to explore if this data displays evidence of savings trends that are compatible with the hypotheses presented in Part 1.




Source: Plutocratic saving ratios, estimated from the FES, 1971-2010. See Annex 1 for details.

Figure 4 above plots household saving ratios by income decile. The FES data is proportionally adjusted to match estimates from National Accounts and offset the declining survey coverage of aggregate income and expenditure over the years (see Annex 1 for details). Overall, there appears to be a mild general decline in saving rates across deciles from the early '90s to the eve of the crisis.² The lowest decile is an exception to this, and exhibits a much more dramatic decline in saving rates compared to other deciles. Part of this decline could be driven by a change in the composition of households within the decile. However, with the possible exception of individuals over 65, we find no evidence of a substantial compositional change over this time period, as shown in Figure 5.

² Note however that this result is dependent on the adjustment to National Accounts. See Annex 1 for methodological details and further caveats.



Source: FES 1983-1987, 2003-2007



The same point can be expressed by comparing the total growth rates of consumption with the growth in disposable income over a given time period. In line with the preceding graph, Figure 6 highlights how households across the income distribution experienced some degree of 'over-consumption' in the ten years before the crisis. While consumption growth exceeding income growth does not necessarily mean one is living beyond their means, it does imply a deterioration in their saving rate. The result is less marked if one looks over the whole period since 1971, in Figure 7. However, there is evidence of a long-term trend of consumption growing faster than income, thereby determining a deterioration of household finances, in the bottom decile and to the top decile.

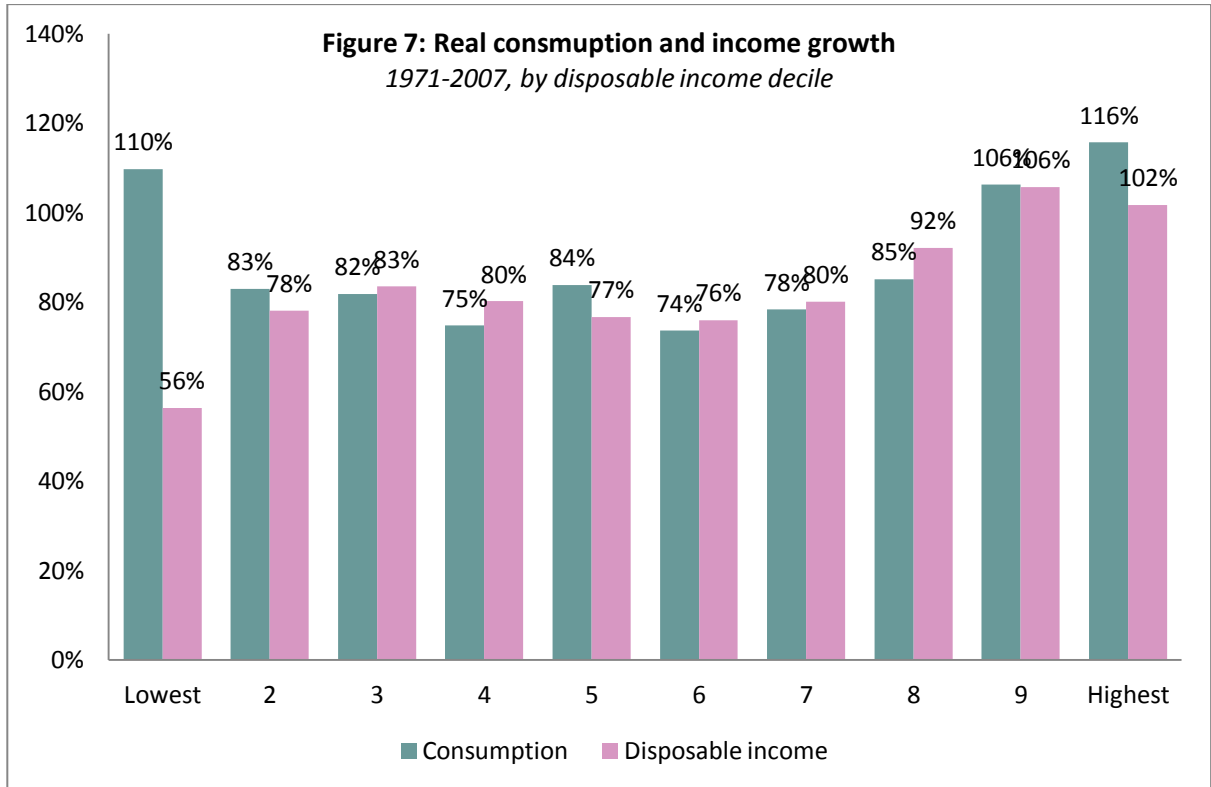
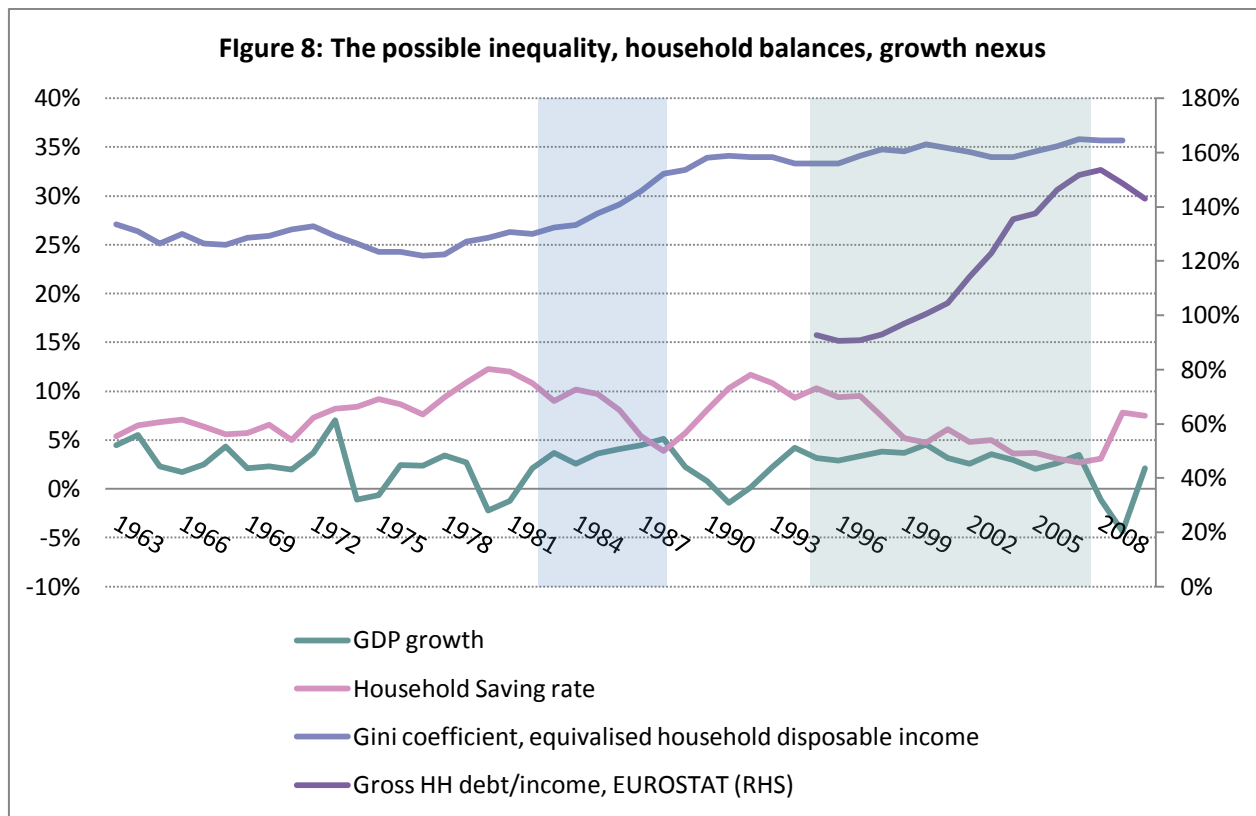


Figure 8 below brings together the tentative evidence presented in this section. Plotting the Gini coefficient alongside aggregate debt and saving ratios and GDP growth shows how the recent crisis has been preceded by a period of high and slowly rising inequality and rapidly deteriorating household debt and saving. Similarly, the early '90s recession was also preceded by rising inequality and falling saving ratios. On the other hand, the 1970s recessionary periods in the '70s were not preceded by such trends.




The data considered here allows some very preliminary comment on which of the theories from Part 1 may be most relevant to the UK context. The fact top income shares have increased substantially may lend more support to the thesis that the behaviour of high-earners was a determining factor behind financial deregulation and investment bubbles (Fitoussi and Saraceno) rather than to the thesis that low-earning families took on credit to maintain declining living standards (Stiglitz). In the former scenario, financial deregulation would nevertheless mean that low-earners accumulate debt and reduce saving more than households in the other parts of the distribution as more credit is made available to them. Indeed, we find evidence for this over the period since the '70s. However, our evidence is not clear cut, as we find that there is evidence of 'over-consumption' across the whole distribution during the ten years before the crisis. This would be consistent with individuals spending to maintain their relative living standards compared to their peer group (Frank et al.). Ultimately, however, testing the validity of each theory requires econometric analysis of survey data, which are beyond the scope of this paper.

As in the US and other countries, the UK has seen a sharp rise in household debt and a decline in saving in the decade preceding the crisis. While this period was not characterised by a marked increase in overall inequality, it was one of stable record-high levels of inequality following the previous dramatic step upwards in the '80s, as well as a continued growth in incomes at the top. This suggests that some of the mechanisms being highlighted in recent theories linking inequality, household finances and crisis may also have been at work in the UK, though more research is warranted to validate this hypothesis. In particular, the role of top income shares and that of persistent high levels of inequality appear to be the most promising avenues for further investigation.

Part 3 - Limits to consumption and impact on the UK recovery

Five years from the beginning of the recent downturn, the UK recovery is still very weak. Figure 9 shows this clearly: this has been the deepest recession since the Great Depression and, following bleak growth over the last two years, economic activity has now recovered considerably less (in relation to its peak) than it had done in the early '30s.



Source: NIESR

At the same time, the distribution of incomes and balance sheets of households continue to show the same alleged risk factors that were evident pre-crisis. As in other advanced economies, inequality in the UK did not fall during the recession.³ Inequality, according to all key measures, therefore remains at record high levels. In contrast, the household saving rate, which often increases during recessions, jumped from a pre-recession 3% to 8% in 2010, indicating some evidence of household rebalancing and increases in precautionary saving. However, in historical perspective, the current levels are no higher than the long term average. As a consequence, efforts by households to rein back their debt have only had a relatively moderate impact. In the two years to 2010, aggregate household debt to income ratio did fall 11 percentage points to 143%, but is still comparable to figures seen in the '00s and still well above levels seen in the '90s. Therefore, as the UK progresses on the current fragile recovery path, it is instructive to build on recent as well as more established literature to explore the ways in which inequality, and its effect on household finances, may or may not affect the prospects for sustainable recovery and growth.

In particular, one concern in the current context is that income inequality may limit consumption growth and therefore aggregate demand. As mentioned previously, this argument is of particular relevance given that around two thirds of UK growth in demand over the past two decades was accounted for by consumption (BIS/DIFID, 2011). Distribution and demand were at the centre of Post-Keynesian models. These originally analysed the effects of the functional distribution of income (i.e. between wages and profits) on consumption. Assuming that the propensity to consume out of wages is higher than the one out of profit, a redistribution that favours the latter would reduce overall level of consumption-led economic growth. This is the often referred to under-consumption theory of economic stagnation. Brown (2004) departs from such an arguably rigid class distinction and develops a model on the more flexible Keynesian view that the marginal propensity to consume is lower at higher levels of disposable income. In

³ The Gini coefficient, calculated for equivalised household disposable income, remained substantially unchanged from 2007 to 2009. Saez (2012) and Morelli (2012) discuss the impact of the financial crisis on top income shares in the context of the US. Their argument is that a decline in inequality can be difficult to achieve without strong changes in policy scenario (e.g. market regulation and fiscal policies).

this context, a redistribution of income away from poorer segments of the population is still consistent with a stagnation of aggregate demand.

This literature would appear to be at odds with the theories presented in Part 1. If the latter have emphasised how inequality increases the pressure to consume and leads to deterioration of household finances, this more Keynesian literature instead argues that inequality reduces the consumption possibilities of households and therefore the aggregate propensity to consume. Perhaps this contradiction may be assuaged by emphasising how the relationship between inequality and consumption can be seen as articulated across three dimensions: how inequality affects what households can afford to consume; how it shapes their perception of what they would like to consume; and how it affects the extent credit is made available to close this gap. Loosely speaking therefore, Keynesian theories might be regarded as concentrating on the affordability of consumption. The above-mentioned demand-side arguments are centred around how inequality frames households' desired living standards. Finally, the above-mentioned political economy theories concentrate on how inequality affects the role and functioning of credit. The expected relationship between inequality and consumption will therefore depend on which of these mechanisms is more prominent in the economic context under consideration. Indeed, the dampening effect of inequality on demand described in Keynesian theories assumes that households cannot access additional credit. Indeed, as presented in some of the other theories, the stagnation in aggregate demand can be counterbalanced by greater availability of credit in the economy. In the current post-crisis context of tight credit constraints however, this may no longer be possible. Furthermore, in light of the experience before the crisis, it may also not be advisable. In this context, it is possible that high levels of inequality could therefore weigh down on growth through their effect on aggregate demand. A number of commentators have therefore argued that a more equal or broad-based income growth would have positive implications for the sustainability and strength of the UK economic recovery.


Estimating the extent to which income distribution can have an effect on aggregate demand is essential to evaluate the importance of this posited mechanism. Replicating and extending the methodology in Brown's (2004) on US data, we carry out a static simulation on UK data to test whether changes in the income distribution have a substantial impact on aggregate demand, assuming the economy is correctly described a Keynesian-type model. Annex 2 provides a detailed technical presentation of the simulation procedure, however the following provides a flavour of it. The approach is broadly divided in two phases. Firstly, we construct a simple Keynesian model of consumption behaviour and parameterise it to match actual data as closely as possible. We characterise the household sector as being populated by ten households, representing the ten income deciles. We assume their consumption behaviour follows a Keynesian-type consumption function such that a) there is a minimum socially accepted level of consumption which each household would exhibit even when it has no income and b) the marginal propensity to consume falls as income increases. We choose such parameters such that the predicted consumption level matched actual consumption levels in FES data by deciles over the period 1971 to 2010 as closely as possible. In the second phase, we exogenously alter the dispersion of incomes and use the model to explore the effects of this change on consumption. As mentioned above, this approach is static and rests on a number of possibly restrictive assumptions. It cannot therefore purport to capture long-term and dynamic effects of the sort that we have presented in Part 1. However, it should allow to go some way towards obtaining an approximate sense of scale of the possible effect of income distribution on consumption levels, particularly in the short-run.

We simulate the effect of three hypothetical reductions in inequality. The first represents a relatively mild reduction in inequality, in the order of a 1 to 2 percentage point reduction in the Gini coefficient. We see this representing the likely effect of a policy intervention, for example through taxation, to reduce inequality in the short term. Accordingly, we called this scenario 'Policy intervention'. Our second scenario assumes a much more substantial reduction in inequality of around 10 percentage points in the Gini coefficient. This would equate bringing the current levels of inequality back to those seen in the '70s, and we therefore refer to this scenario as 'Back to the '70s'. Finally, in line with Brown (2004), we present the extreme scenario where the incomes of all households are equalised. We call this scenario 'full equality' and, by definition, would imply a Gini coefficient of zero.

Across all these scenarios, our interest is evidently forward-looking: how such changes in inequality would affect consumption in the future. However, we address this question indirectly by looking backwards: the extent to which consumption patterns would have been different had the income distribution been more

equal. We can then derive the aggregate propensity to consume in each scenario and compare it to the actual historical series to extract the impact of the changed income distribution.

The results from the three scenarios can be summarised in Figure 10. The graph shows the actual and simulated series for the aggregate propensity to consume, while the table below the chart reports the average difference between the simulated series and the actual series in each of the three scenarios considered.

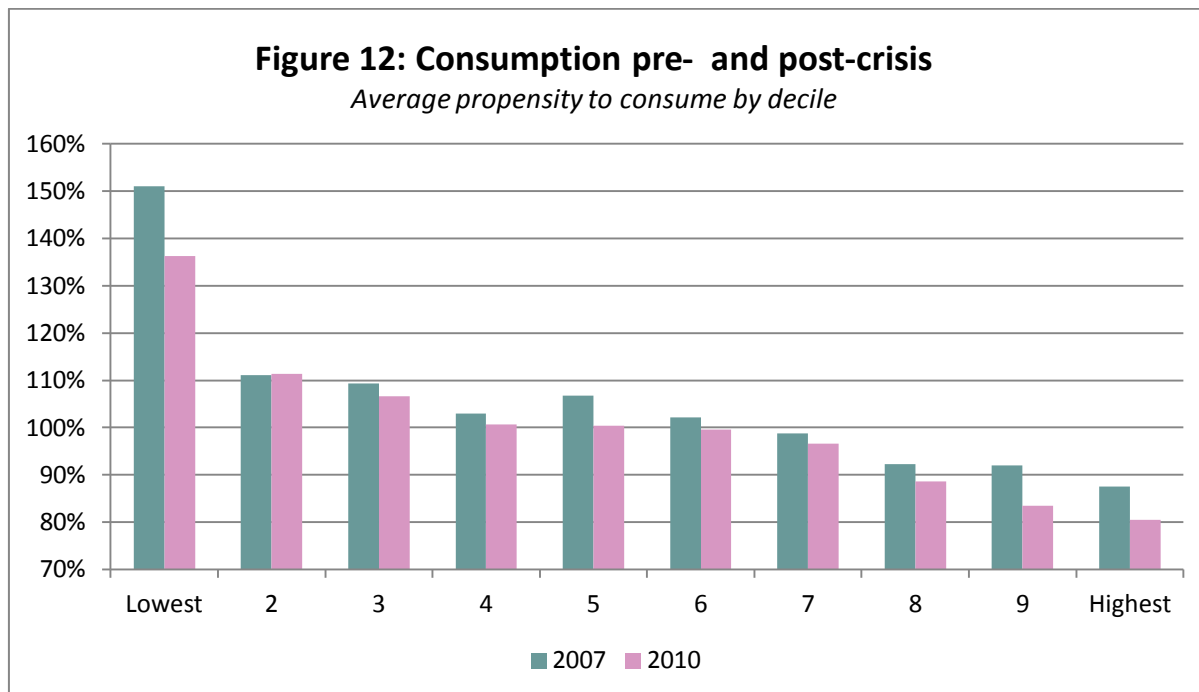


	Policy intervention	Back to the 70's	Full equality
ppt change in APC	0.1	0.3	0.6

The results indicate, in line with findings in Brown (2004) that even a substantial income distribution would not significantly affect the aggregate propensity to consume of households. Specifically, our estimates indicate that even the extreme scenario of making everyone's income the same is associated with only around a 0.6 percentage point increase in the aggregate propensity to consume. Considering that household consumption in the UK is around two thirds of GDP, such an increase would translate into less than half a percent of GDP. This perhaps surprising result is driven by the fact that the marginal propensity to consume we estimate from the data is not substantially different over the income distribution. The model therefore indicates that influencing the income distribution would not appear to be an effective way to boost to consumption.


A possible criticism to the above simulations is that they take actual consumption in previous years as their starting point, but that baseline consumption patterns going forwards may be significantly different from past trends. In particular, one could argue that the current tighter credit constraints compared to the pre-crisis credit environment will weigh heavily on consumption patterns over the medium term. As pointed out above, in such context, a reduction in inequality may deliver larger impacts on consumption by providing additional income to otherwise liquidity constrained low-income individuals. We address this point in two ways. Firstly, we explore the data to validate whether in fact pre- and post-crisis consumption patterns have been systematically different. Figure 11 below provides some insight into this by plotting the average propensity to consume of each decile in 2007 and 2010. Comparing the

data across the two years, we see how consumption does appear to have retrenched somewhat, particularly at the bottom and at the top of the distribution (interestingly, the same deciles where consumption had grown faster than income over the last 40 years). This is likely to reflect tighter credit conditions, as well as increases in precautionary saving in response to increased uncertainty. However, perhaps surprisingly, the comparison does not reveal a dramatic difference in consumption patterns before and after the crisis. Furthermore, the poorest three deciles continued to consume above their incomes, suggesting a continued access to liquidity through credit or dissaving.



Secondly, and abstracting from the point just made, we re-run the simulations imposing strict liquidity constraints in our model. Specifically, we impose a condition on the consumption function such that consumption cannot be higher than income (a condition which 2010 data in Figure 11 above indicates is likely to be far stricter than that occurring in reality). We therefore can then recalculate a new, lower, historical series for the baseline aggregate propensity to consume, and compare how simulated changes in the income distribution affect aggregate consumption levels when strict liquidity constraints are in place. In essence, the simulations reveal the extent to which consumption patterns would have been different in each scenario and when households can only consume out of current income. One would expect these simulations to identify a larger effect of falls in inequality on consumption, as liquidity constrained households benefiting from redistribution would be likely to spend virtually all of any additional income.

Results from this second set of simulations are presented in Figure 12. As expected, the simulated changes under strict liquidity constraints have effects on consumption that are several times larger than those obtained when liquidity constraints are absent. However, the economic magnitude of the effects remains small, especially considering the strictness of the liquidity constraints imposed. In this context, altering the income distribution to the highly extreme 'Full equality' scenario would be associated with only a 3.0 percentage point increase in the aggregate propensity to consume. The somewhat more plausible scenario of restoring the levels of inequality seen in the '70s would, similarly, lead to only a 1.2 percentage point increase in the aggregate propensity to consume. Overall then, even when liquidity and credit constraints are binding, the simulations suggest that influencing the income redistribution would not have significant effect on aggregate consumption.



	Policy intervention	Back to the 70's	Full equality
ppt change in APC	0.2	1.2	3.0

The above results may be seen as broadly in line with Blinder (1975), which finds that “*equalizing the income distribution will either have no bearing on, or (slightly) reduce, aggregate consumption.*” Indeed, even under a standard Keynesian framework, which is expected to predict an increase in the aggregate consumption, our model finds only modest effects. Moreover, if one introduces relative income considerations in one's behavioural model (such as the ones discussed in Part 1), equalisation might lead to decreases in consumption as the pressure to keep up with peers is eliminated. This is implied by Blinder’s conclusion to his work, when discussing increases in inequality: “*If the kinds of “demonstration effects” stressed by Duesenberry are at all important, disequalization can conceivably lead to more rather than less consumption.*” Overall, the effects, if any, of the two forces described above pull the aggregate propensity to consume in opposite directions, with the final result ultimately depending on their relative strength.

Conclusion

This paper has provided an overview of the theory and evidence on the interaction between inequality and growth, focusing on the impact on household consumption behaviour. After surveying this emerging literature, we present one of the first empirical investigations of these issues in the UK context. The initial evidence we present, while only suggestive at this point, does hint at a possible relationship between inequality and the recent crisis. As in the US and other countries, the crisis in the UK emerges after a sharp, decade-long rise in household debt and a decline in saving. At the same time, this period was characterised by record levels of inequality following a dramatic step upwards in the '80s, as well as by continued growth in incomes at the top. Detailed analysis of micro-data presented here indicates that there has been a particularly dramatic decline in saving among households at the lower end of the income distribution, which is consistent with the view that increasing demand for, and supply of, credit to these households may have reduced the sustainability of their debt burden and hence increased the risk of crisis. Mechanisms linking inequality, household finances and financial crisis may therefore have been at work in the UK, though more research is warranted to validate this hypothesis. In particular, the role of top income shares and that of persistent high levels of inequality appear to be the most promising avenues for further investigation.


The analysis of pre-crisis developments helps frame our understanding of the current economic situation. Inequality, and its effect on household finances, is also likely to affect the recovery path. However, the effect of inequality on consumption behaviour may differ markedly depending on the economic context. Continued income inequality can have limiting effects on consumption growth and therefore aggregate demand, particularly in the current context of tight liquidity constraints and household rebalancing. We explore the quantitative dimension of this hypothesis through simulations that build on previous analysis conducted on the US. Our results indicate that even a substantial income redistribution may not significantly affect the aggregate propensity to consume of households. Such results would indicate that, if one's objective is to sustain the UK recovery through additional demand, more equal or broad-based income growth are unlikely to be the most effective and immediate channel. This suggests that while reduced inequality and more broad-based income growth might help reduce the probability of a future crisis - and hence make recovery more sustainable - it is unlikely to have a substantial impact on the short-term prospects for growth.

Annex 1 - Measurement of saving in the FES

Part of the evidence presented in the paper is based on the analysis of survey data from the Family Expenditure Survey (FES), which was replaced by the Expenditure and Food Survey (EFS) in 2001 and by the Living Costs and Food Survey (LCFS) in 2008. These surveys, which we refer collectively in the paper as Family Expenditure Survey, collect information on household and personal incomes and a detailed expenditure record for a representative sample of UK households. Using the cross-sectional surveys from 1971, we are therefore able to construct a detailed dataset on which to conduct our analysis, particularly in relation to trends by deciles. There are however two main caveats worth pointing out in relation to the analysis of saving behaviour in the FES. Firstly, the information available provides an important but incomplete picture of individual behaviour. Secondly, aggregate estimates from the FES have increasingly diverged from National Account aggregates over time. We discuss each of these in turn.

The FES includes detailed information about income sources and expenditure categories at household level. In particular, it measures net disposable income and total household expenditure. Income is here defined as including earnings as well as income from other sources such as benefits (excluding housing-related support), pensions and investments. Consumption is based on the FES total expenditure variable. This covers a comprehensive spectrum of expenditure categories including housing, mortgage and rent costs (net of any amount received in housing-related benefits). These measures allow us to construct a measure of saving, defined as disposable income minus expenditure. This is a specific measure of saving, sometimes referred to *active saving*, and closely matches the popular understanding of saving as 'how much of your income do you put aside each month'. A broader definition of saving would include the accrual of capital gains that are not realised and spent. These add to one's stock of savings, and are rightly considered a form of *passive saving* (see Crossley and O'Dea, 2010 for an extensive discussion). Capturing a measure of passive saving is not possible with the FES as it does not include information on household wealth and capital gains. Trends in active saving will to some extent be influenced by trends in other types of savings (e.g. pensions) and wealth trends more generally (e.g. housing or stocks). In this sense, focusing on active saving only may give an incomplete picture of changes in saving behaviour. However, the FES dataset still provides important information on saving behaviour that can be useful in the context of the present study.

Estimates from the FES have increasingly diverged from National Account aggregates over time. Given the representative nature of the dataset, estimates of per capita income and expenditure or the household saving ratio should, in principle, closely track estimates from National Accounts. However, as Figure A1.1 below shows, there has been a long-term decline in the proportions of the National Account aggregates that are accounted for by the FES over the period under observation.



Reproduced from Van de Ven (2011)

There does not appear to be a clear consensus on how one should deal with this issue. Crossley and O’Dea (2010) argue that this implies one should be careful interpreting trends from the FES as necessarily being caused by genuine changes in household savings behaviour. Instead, van de Ven (2011) proposes a pragmatic solution. Noting there is currently no detailed understanding for this decline in the FES, he suggests proportionally adjusting the FES cross-sections to match the aggregates reported in the National Accounts. This is the solution we adopt in this paper. Nevertheless, as a robustness check, we calculate both unadjusted and adjusted saving rates and point out when results are dependent on adjustment in the paper.

We provide here some detail on the methodology used to calculate saving ratios before presenting the results. Firstly, we group benefit units by disposable income deciles defined over the distribution of disposable income in each year. Note that there may be more than one benefit unit in each household. This requires apportioning household expenditure and disposable income in proportion to the number of adults and children in each benefit unit. A benefit unit’s saving ratio is then defined as the saving (measured as disposable income minus consumption) over disposable income. Drawing on Crossley and O’Dea (2010), we construct three different types of saving ratio for each decile: the simple mean of all saving ratios across benefit units in a given decile (democratic measure); the median saving ratio across benefit units in a given decile (median measure); and finally the ratio of aggregate saving to aggregate disposable income in a given decile (plutocratic measure). The later definition is the one used in National Accounts to calculate the aggregate saving ratio. Due to a handful of observations having implausible responses on income, some saving ratios emerge as clear outliers. We therefore also calculate each saving ratio measure dropping the lowest and highest 1% of saving ratio. The resulting measures are presented in the graphs below.

Figure A1.2: Democratic saving ratios by selected deciles




Figure A1.3: Democratic saving ratios by selected deciles
(adjusted to match national aggregates)

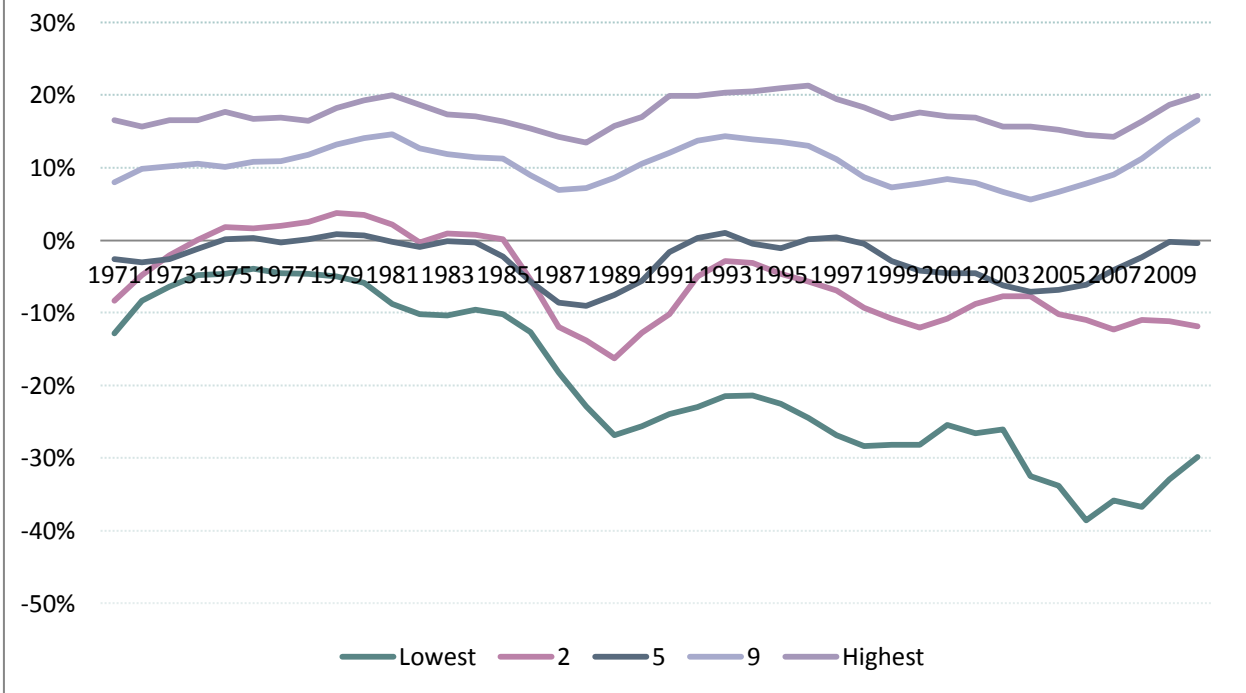


Figure A1.4: Median saving ratios by selected deciles




Figure A1.5: Median saving ratios by selected deciles

(adjusted to match national aggregates)




Figure A1.6: Plutocratic saving ratios by selected deciles



Figure A1.7: Plutocratic saving ratios by selected deciles
(adjusted to match national aggregates)



Figure A1.8: Plutocratic saving ratios by selected income deciles
(excluding individuals over 65, adjusted to match national aggregates)



Annex 2 – Simulation methodology

This Annex details the methodology used in the simulations, which replicates and extends that used in Brown (2004) on US data. The approach is broadly divided in two steps. Firstly, we construct a simple Keynesian model of consumption behaviour and recover preference parameters from survey data. Next, we use the preference function to model consumption responses to exogenous alterations in the dispersion of incomes.

Our household sector is characterised as being populated by ten households, representing the ten income deciles. We assume the consumption behaviour any household i in year t is determined by the consumption function described by the following equation:

$$C_{it} = \varphi_t + Y_{it}^{\alpha_t} \quad s.t. \quad 0 < \alpha_t < 1 \quad (1)$$

Equation (1) captures two Keynesian-type intuitions. Firstly, there is a time-varying minimum socially accepted level of consumption which each household would exhibit even when it has no income (φ_t). Secondly, the restrictions on α_t are such that the marginal propensity to consume falls as income increases. Furthermore, such curvature parameters are allowed to vary over time. Considering that one's evaluation of what should be the socially accepted minimum is likely to be influenced by the living standards displayed by others, we define the minimum level in relation to the distribution of income at a given point in time, as follows:

$$\varphi_t = \rho \times \bar{Y}_t \quad (2)$$

Equation (2) implies that the minimum socially accepted level of consumption will be a fixed proportion (ρ) of the average income at time t . Note that, for simplicity, ρ is fixed over time, meaning that while we allow for the minimum standard to rise as average living standards improve, the social preference as to what constitutes an acceptable distance between minimum and average is constant over time. Finally, in some of the simulations considered in the paper, we impose strict liquidity constraints. This is achieved by imposing the condition that one cannot consume more than their total income:

$$C_{it} = Y_{it} \quad if \quad \varphi_t + Y_{it}^{\alpha_t} > Y_{it} \quad s.t. \quad 0 < \alpha_t < 1 \quad (3)$$

We proceed to choose parameters of the unconstrained consumption function to match actual data as closely as possible. We obtain baseline income data by decile from the FES over the period 1971 to 2010. Next, for a given ρ , we solve for the set of α_t such that the ratio of aggregate simulated consumption over aggregate income in a given year matches that in the UK National Accounts for the same year. In line with Brown (2004), these are found to be in the range of 0.95-0.98, implying that the marginal propensity to consume is not substantially different across the income distribution. Finally, as we hold data from the FES on actual consumption, we can evaluate the extent to which the estimated consumption matches the true data, and use this information to choose the optimal value for ρ . Specifically, we ran the above procedure for all possible integer percentage values for ρ , re-evaluating the set of α_t each time, and chose the value which minimised the sum of the squared differences between simulated and actual consumption over all deciles and years. This yielded value of 12%, which implies a minimum socially accepted consumption level of around £65 per week in 2010. It is interesting to point out how this figure is strikingly close to the adult single rate of Job Seeker's Allowance of £65.45 per week in 2010/11.

Having determined a model for consumption behaviour as above, we can estimate baseline consumption levels and explore how these would be affected by hypothetical exogenous changes in the dispersion of incomes. We simulate the effect of three hypothetical reductions in inequality. As the most extreme scenario, we hypothesise the full equalisation of all incomes. We define the remaining scenarios as partial moves in that direction. In line with this, the simulations are achieved applying the following transformation to the income of household i in year t :

$$Y_{it} \rightarrow (\bar{Y}_t/Y_{it})^{\frac{1}{n}} \times Y_{it}$$

This transformation represents going $(1/n)^{\text{th}}$ of the way towards full equality. Note, in fact, that, if $n = 1$, all household incomes are transformed to the average at time t , and that, for all $n > 1$, applying the transformation n times achieves full equality. The transformed income values are then adjusted proportionally to ensure their average is equal to that in the baseline scenario, and therefore φ_t is unchanged. We chose n to be 20, 4 and 1 to define the 'Policy intervention', 'Back to the '70s' and 'Full equality' scenarios respectively. Calculating the Gini coefficients in the baseline and transformed income distribution in each year, reveals that these scenarios are equivalent respectively to a 1 to 2 percentage point drop in, a 10 percentage point drop in and reduction to zero of the Gini coefficient.

References

- Acemoglu, D, 2011, "Thoughts on inequality and the financial crisis", AEA meeting, Denver, 7th January
- Atkinson, A B and Morelli, S, 2010, "Inequality and banking crises: A first look", report for the International Labour Organisation.
- Atkinson, A B and Morelli, S, 2011, "Economic crises and Inequality", Human Development Research Paper 2011/06. Research Background paper for the Human Development Report 2011, UNDP, UN.
- Atkinson, A B and Morelli, S, 2012, "Chartbook of Economic Inequality: 25 countries 1911-2010", discussion paper
- Berg, A D, and Ostry, J D, 2011, "Inequality and Unsustainable Growth: Two Sides of the Same Coin?", IMF Staff discussion note, SDN/11/08
- BIS/DFID, 2011, "Sources of Economic Growth", Trade and Investment Analytical Papers Topic 6 of 18
- Blinder, A S, 1975, "Distribution Effects and the Aggregate Consumption Function", *Journal of Political Economy*, Vol. 83, No. 3 (Jun., 1975), pp. 447-475
- Bordo, M D, and Meissner, CM, (forthcoming), "Does Inequality Lead to a Financial Crisis?", *Journal of International Money and Finance*.
- Brown, C, 2004, "Does Income Distribution Matter for Effective Demand? Evidence from the United States" *Review of Political Economy*, Volume 16, Number 3, 291–307
- Cahuc, P, and Challe, E, 2009, "Produce or Speculate? Asset Bubbles, Occupational Choice and Efficiency," CEPR Discussion Papers, n.7602
- Crossley, T and O'Dea, C, 2010, "The wealth and saving of UK families on the eve of the crisis", Institute for Fiscal Studies
- Duesenberry, J, 1949, *Income, Saving, and the Theory of Consumer Behavior*, Harvard University Press, Cambridge, MA
- Fitoussi, J-P and Saraceno, F, 2009, "How deep is a crisis? Policy responses and structural factors behind diverging performances", OFCE Document de Travail, 2009-31.
- Fitoussi, J-P and Saraceno, F, 2010, *Inequality and Macroeconomic Performance*, OFCE/POLHIA N2010-13.
- Frank, R H, Levine, A S and Dijk, O, 2010, "Expenditure cascades", working paper available at SSRN.
- Iacoviello, M, 2008, "Household Debt and Income Inequality, 1963-2003", *Journal of Money, Credit and Banking*, Volume 40, Issue 5, pages 929–965, August 2008
- Jappelli, T, Pagano, M, and Di Maggio, M, 2010, "Households' Indebtedness and Financial Fragility", CSEF Working Paper N. 208
- Jenkins, S, 2011, "Has the Instability of Personal Income Been Increasing?", *National Institute Economic Review*, N. 218, October 2011.

Krugman, P, 2010, "Inequality and Crises", New York Times blog 'The Conscience of a liberal' (krugman.blogs.nytimes.com), 28th June

Kumhof, M and Ranciere, R, 2010, "Inequality, Leverage and Crises," IMF Working Papers 10/268, International Monetary Fund.

Milanovic, B, 2009, "Two Views on the Cause of the Global Crisis – Part I", Yale Global Online (yaleglobal.yale.edu), 4 May.

Morelli, S, 2012, "Banking crises and stock market crashes in the US: the response of top income shares in historical perspective", draft (September), University of Oxford.

Rajan, R G, 2010, Fault lines: How hidden fractures still threaten the world economy, Princeton University Press, Princeton.

Saez, E, 2012, "Striking it Richer: The Evolution of Top Incomes in the United States (Updated with 2009 and 2010 estimates)", working paper (elsa.berkeley.edu/~saez/)

Stiglitz, J E, 2009, , "Joseph Stiglitz and why inequality is at the root of the recession", Next Left (www.nextleft.org), posted 9 January 2009.

Van de Ven, J, 2011, "Expenditure and disposable income trends of UK households: evidence from micro-data", National Institute Economic Review, No. 218,

The Resolution Foundation

The Resolution Foundation is an independent research and policy organisation. Our goal is to improve the lives of people with low to middle incomes by delivering change in areas where they are currently disadvantaged. We do this by:

- undertaking research and economic analysis to understand the challenges facing people on a low to middle income;
- developing practical and effective policy proposals; and
- engaging with policy makers and stakeholders to influence decision-making and bring about change.

For more information on this Briefing Note contact:

James Plunkett

Secretary to the Commission on Living Standards

james.plunkett@resolutionfoundation.org

020 3372 2956

