The distributional impact of Universal Credit: a long-run perspective

Robert Joyce, Tom Waters and Joseph Woods

Institute for Fiscal Studies

[PRELIMINARY RESULTS: DO NOT CITE]

Abstract

Most distributional analysis of tax and benefit reforms, including the shift to Universal Credit (UC), is based on a 'snapshot' of household incomes and circumstances. However, the impact of reforms on living standards may be better measured as the average effect over a longer time horizon. We use longitudinal Understanding Society data together with a tax and benefit microsimulation model to compare the impact of UC on incomes measured in the 'snapshot' against incomes measured over a 24 year period. We find that UC distributes more to the lifetime poor, and look to explain this by examining various aspects of the UC system that treat lifetime rich and poor differently to the 'legacy' system if replaces. For example, UC on average gives less to those with financial assets than the legacy system. It also gives more to working renters, and less to working homeowners. Since ownership of financial assets and homes is correlated with higher lifetime incomes, UC on these margins distributes support more to the lifetime poor and less to the snapshot poor. Finally, we show that while hypothetical UC reforms which target in-work claimants are less progressive than those with target out-of-work claimants in the snapshot, they are more similar in the long-run. This emphasises how being out of work is often a temporary state for many of the poorest adults in the snapshot.

JEL classification: H24

Keywords: Inequality, measuring living standards, redistribution, life cycle

1. Introduction

Universal Credit (UC) is the largest change to the UK benefit system in several decades, integrating six means-tested benefits and tax credits for working-age families (the 'legacy' system) into a single monthly payment. Typically, distributional analysis of UC is examined on a snapshot basis – taking individuals' circumstances at a point in time – and is shown as regressive. This is to be expected; UC represents an aggregate cut to working-age benefits, and these benefits disproportionately go to lower income households. More interestingly (and the focus of this paper), UC also embodies a series of decisions about how to 'reprofile' welfare support between different kinds of people. Because so many aspects of the benefits system are changed, UC represents an excellent case study on how specific design choices of the benefits system – and the subsequent reprofiling of support between different groups – affects targeting of resources towards the poor and different demographic groups.

Furthermore, there are reasons to believe that these impacts – both the overall impact on household incomes and the reprofiling between groups – can be different when circumstances are measured over multiple years. For example, UC considerably reduces support for self-employed people with low earnings at a given point in time – but these people tend to enjoy relatively high lifetime incomes. These considerations mean that the distributional impact of UC over a longer-term horizon is far from clear without a careful modelling exercise and high quality longitudinal data.

In this paper, we provide the first estimates of the distributional analysis of UC over a longer-run perspective. Over a lifetime peoples' circumstances change in systematic ways – such as the family formation and the age-profile of earnings – and transitory ways, such as unemployment. This paper connects to a wider literature which highlights how extending the time horizon of analysis over several years can generate a very different picture of the impact of benefit reform. We find that, over the long-run, UC looks less regressive than the snapshot picture and elements of UC actually work to shield the poorest households from the reduced spend on welfare. The overall net impact on household incomes is categorised into an "aggregate" effect, reflecting that total spend in UC is lower than the legacy system it is replaces, and a 'reprofiling' effect which captures how UC allocates that total spend.

Continuing with this concept of 'reprofiling', we then isolate specific features of UC and examine how design choices of a benefit system lead to divergences in the targeting of support between individuals with snapshot hardship and individuals considered long-term poor. We find that harsher treatment of assets and self-employed incomes look regressive in conventional snapshot analysis, but reprofile support towards the poorest individuals in the long-run. Likewise, relatively generous tapering of awards for in-work renters has relatively little impact on the distribution of support in the snapshot but looks progressive in the long-run. Conversely, reducing support for those with severe disabilities reduces targeting of resources to the poor in both the snapshot and long-run.

Finally, we look at hypothetical UC reforms – prevalent in the current policymaking debate – and show that over the long-term, as people cycle in and out of employment, reforms which target in-work claimants alone can generate similar long-run distributional impacts to those which also support the out-of-work, though they are less progressive when incomes are measured in a snapshot.

This paper makes two main contributions to the literature on reforms to benefit policy and their associated impacts across the income distribution. The first is the 'lifecycle' literature on the longer-term impacts of benefit reforms on living standards. Corroborating previous research which found that a far greater share of individuals receive support from the welfare system over a long horizon (Roantree and Shaw 2018), we find that nearly half of adults live in households which receive some UC over a 24 wave period, and nearly 1 in 5 adults in the top lifetime income decile receive UC at some point. This highlights the mobility of individuals between different groups and economic circumstances across time.

The second contribution is that we show how specific design features of a welfare reform can be used to influence the trade-off between targeting short-run or long-run poverty. Previous work has shown how specific policy changes which look progressive in a snapshot sense may not effectively target long-run poverty (Levell, Roantree and Shaw 2015). We show that UC is more efficient at channelling resources to the lifetime poor, and that this is driven by certain design features in the system.

Understanding both the contemporaneous and long-term impacts of welfare reform is important in order for policymakers to make informed decisions on tackling short-term hardship or long-run low income. Understanding how specific features of UC redistribute resources between groups with varying propensities of short-term and long-term poverty is essential for designing effective benefit systems.

The rest of this paper proceeds as follows. Section 2 describes the data, while section 3 presents results of both snapshot and long-run distributional analysis. Section 4 isolates specific features of the benefit system and studies their redistributive impact. Section 5 discusses hypothesized reforms to UC. Finally, Section 6 concludes.

2. Data and policy context

2.a British Household Panel Survey and Understanding Society

We exploit the longitudinal properties of the British Household Panel Survey (BHPS) and Understanding Society survey data (USoc) to follow a panel of households across 24 waves of data. The BHPS is a panel survey of approximately 5,500 households in the UK, running for 18 waves between 1991 and 2008.[1] It collected a range of information on socio-economic indicators including incomes, employment status, education and training, household composition, and disability.

USoc, like the BHPS, contains a plethora of questions designed to capture valuable socio-economic information. This study superseded BHPS in 2009 (eligible BHPS sample members continue as part of the USoc sample from 2010), and follows a panel of approximately 40,000 UK households, with the latest available wave covering 2015. For the purposes of this paper, we use the harmonised BHPS data series which has been created to facilitate the combined use of BHPS and USoc.

[1] BHPS and USoc waves cover two calendar years (with sequential waves having one overlapping calendar year). For brevity, we refer to these throughout using the earlier year.

We select observations that are present for all 24 waves of the study.[2] This gives us a final sample size of 1,580 adults (spanning 36,555 individual observations).

To account for non-random attrition, longitudinal weights from the most recent USoc survey wave are used to weight the analysis. Whereas the traditional 'snapshot' refers to a population cross-section in a specific year, the 'snapshot' in this paper fixes the population at the 1,580 non-attrited individuals, but treats each wave observation independently. Consequently, we also use longitudinal weights in the snapshot analysis as opposed to cross-sectional weights.

For longitudinal analysis, the constant unit of observation across time is the individual person. Therefore our analysis is conducted at the person level. Individuals must be aged 16 or over to be included.

We compute, for each adult, their net household equivalised income in each wave and the average across all waves.[3] This equivalised income is used to rank adults in the snapshot and long-run income distribution. Where a household is entitled to benefits, we assume that the benefits are equally shared among adults in the household.

Finally, we are interested in isolating the impact of specific features of the benefit system so use TAXBEN – the Institute for Fiscal Studies' tax and benefit microsimulation model for the UK –to calculate adults' disposable income, tax liabilities and benefit entitlements.[4] Benefit receipt is calculated on an entitlement basis – in other words, assuming that everyone takes up the benefits to which they are entitled. Individuals' nominal earnings, rents, mortgage payments, savings/dividend incomes, and pension incomes are uprated using aggregate indices to put them in 2017–18 prices. Given these uprated values we calculate what individuals' incomes, tax liabilities, and benefit entitlements would be given the 2017–18 tax and benefit system. This allows us to model a world in which the sample are exposed to the same policy environment for the full 24 waves of data, rather

[2] Sometimes an individual appears for all 24 waves, but in some waves they, or a member of their household, fail to give enough information to allow for robust microsimulation of their incomes and benefit entitlements. We include such individuals in the sample so long as we have sufficient information on them and their households in at least 20 waves (but ignore the waves where such information is lacking).
[3] We use the modified OECD equivalence scale.
[4] For a brief description of TAXBEN, see Waters (2017).

than the different policy conditions they experienced in reality between 1991 and 2015. Implicitly we are therefore taking individuals' choices (for example, over employment) as given, assuming no change in behaviour in response to a change in policy.

Error! Reference source not found. captures some descriptive statistics of the snapshot and long-run samples. There are three things to note from the Table. First, around half of individuals are at some point in a household receiving legacy benefits (a result aligned with previous literature (Roantree and Shaw 2018)), but only 14% are in such a household at a given point in time. Second, the Table shows that circumstances change considerably over time. For example, more than three times as many people at some point in the 24 year period are in a workless household than are at any given point in time, and twice as many people have dependent children. Thus it is likely that the distributional consequences in the 'snapshot' will differ to the average effect over the long-run.

Table 1. Descriptive statistics for the snapshot and long-run samples.

	Snapshot	Long-run
N	36,555	1,580
Average age (years)[5]	49.5	-
% male	48	-
% in receipt of legacy/ever in receipt	14	54
% in receipt of UC/ever in receipt	11	47
% disabled/ever disabled	5	18
% in workless household/ever in workless HH[6]	15	53
% with dependent children/ever with children	30	58
% with a partner/ever with a partner	76	90

2.b Universal Credit and the 'legacy' system it replaces

UC is a payment from the government to help with living costs. It integrates six means-tested legacy benefits for the working-age. These are:

- Income Support (IS);

[5] It is worth noting, that while the mean age is well aligned to other surveys, our final dataset under samples young adults and older pensioners; these two groups are more likely to attrit.
[6] An individual is counted as in a workless household if all working-age people in the household are out of work. If there are no working-age people in the household, the individual is not included in this statistic.

- Income-related Employment and Support Allowance (ESA);
- Income-related Jobseeker's Allowance (JSA);
- Housing Benefit (HB);
- Working Tax Credits (WTC) and Child Tax Credits (CTC).

While Universal Credit is paid to claimants in a single payment, this actually comprises of different 'elements' which determine a claimant's *maximum amount*. These elements are:

- A standard allowance for single or couple claimants;
- A child element for the first two children;
- A childcare costs element;
- A housing costs element related to rental costs[7];
- An element for claimants assessed as having limited capability for work (LCW) or limited capability for work-related activity (LCWRA).

This *maximum amount* is then assessed against various tests and thresholds in order to determine a final UC *award*. First, the 'work allowance' – which varies by family type and whether the claimant is a renter or not – sets the amount a claimant can earn before their UC award is affected. Second, above the work allowance UC is tapered at a rate of 63% for (net of tax) earned income, and at a rate of 100% for unearned income. Third, for those with capital/savings between £6,000 and £16,000, their award is reduced by £4.35 per month for every £250 in this range. Claimants with capital/savings in excess of £16,000 are not entitled to UC.

In the legacy system, the six different benefits being replaced have their own thresholds, tests and taper rates in order to determine final awards. For instance, Tax Credits are tapered at a rate of 41% for income in excess of the relevant income threshold, whereas Housing Benefit is withdrawn at a rate of 65%. Claimants can be on multiple tapers at once, which can result in very high effective marginal tax rates.

[7] Those with a mortgage can also get a loan to pay mortgage interest when on UC. Since this is a loan rather than a transfer, we do not include this in our analysis.

UC is being rolled out gradually across the UK. For those who are transferred across to UC and find their maximum amount is less than their legacy equivalent, UC includes an additional 'transitional protection' element to make up this difference. Transitional protection is removed after a change of circumstances (e.g. moving into work). For the purposes of this analysis, we have assumed a steady state world with completed roll out, and therefore no transitional protection.

It is worth noting that in both UC and the legacy system, benefit entitlement is almost entirely dependent upon your current circumstances. Unlike in many other benefit systems across the world, there is very little in the way of time limiting of benefits or the use of labour market history in calculating benefit entitlement. This means that there is no *mechanical* way in which UC's effects on snapshot incomes will differ to its effects on long-run incomes.

3. Distributional analysis of Universal Credit

In this section we examine the distributional analysis of UC on both a snapshot basis and across a long-run perspective. We also introduce the concept of 'reprofiling' support between different groups.

Figure 1 shows the marginal impact on household incomes from the introduction of UC, on a snapshot basis, breaking individuals into (snapshot) income deciles. This ignores transitional protection, and can therefore be thought of as the long-run effect of introducing UC after transitional protection has expired.

Figure 1. Impact of UC as a percentage of net household income, by snapshot income decile.

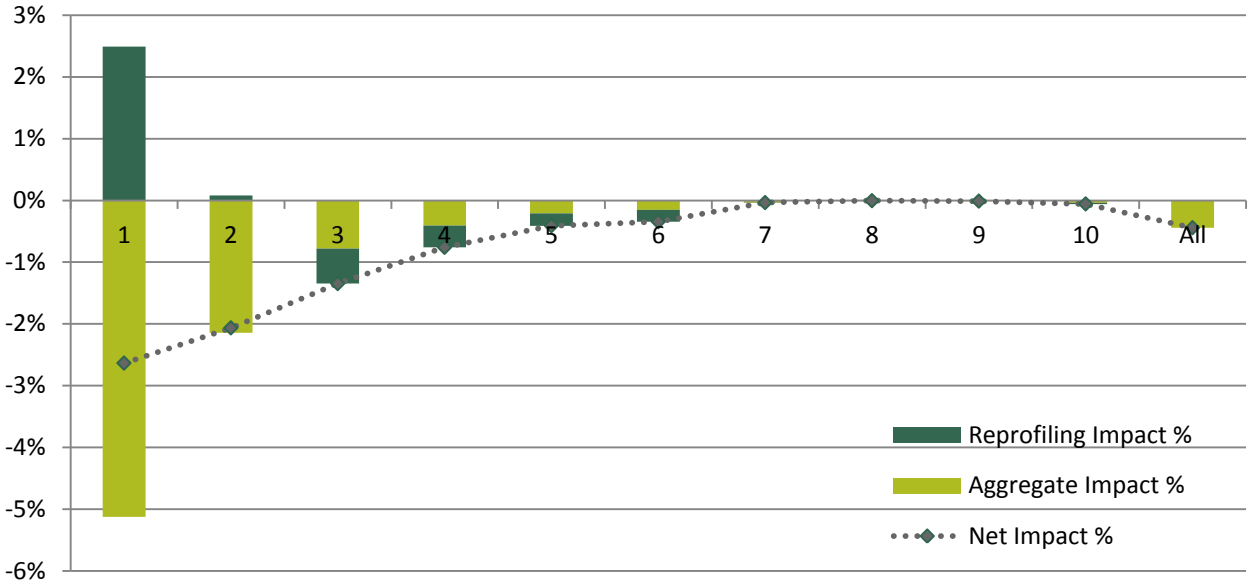

Figure 2. Long-run impact of UC as a percentage of average net household income, by long-run income decile.

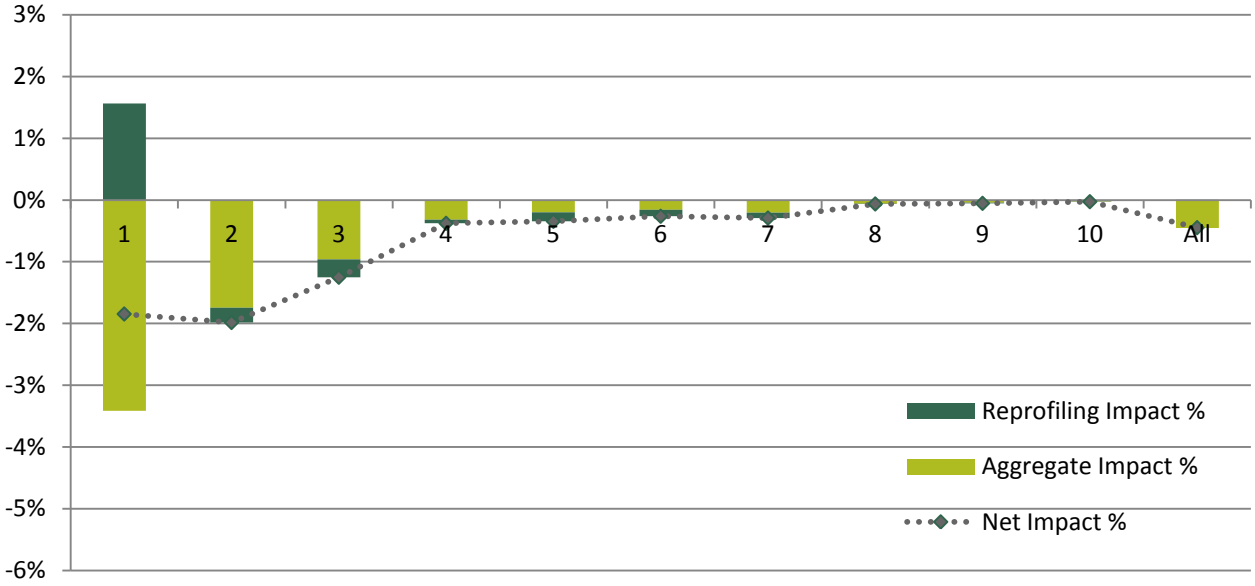

UC represents, on average, a loss to households (as reflected in the dashed line). On average, adults in the poorest 10% of households lose most from the introduction of UC (just under 3% of net income) as the reform represents an overall cut to entitlements, and benefit receipt makes up a larger fraction of their total income; it is this mechanism which generates the regressive profile. Note however that this net impact has been decomposed into:

- An “aggregate” effect (light green bars) – this represents a scenario in which individuals’ existing legacy awards are scaled down, so that total benefit spend in this imagined legacy scenario equals that in UC.
- A “reprofiling” effect (dark green bars) – this represents the difference between the ‘scaled’ legacy award and UC entitlement. This captures the impact of the change in how UC and legacy target their support. By definition, this effect sums to zero across all adults.

While, on average, adults in the poorest households lose out from the introduction of UC, this is driven entirely by the aggregate cuts to welfare spending. The positive dark green bar shows that features of the UC system actually help negate some of these losses. To phrase differently, for a given amount of benefit spend, UC allocates (‘reprofiles’) more support to adults at the bottom of the income distribution. For instance, while roughly only 1 in 10 adults in the poorest decile see a rise in benefit entitlement from the introduction of UC, 1 in 3 adults in this decile gain from the reprofiling of support.

However, circumstances change over time. There are reasons to believe that the “aggregate” and “reprofiling” effects of UC, and hence the distributional impact of the UC reform, to be different when circumstances are measured over a longer-run perspective.

- One would expect the “aggregate” effect to change as individuals in receipt of benefits in a snapshot year do not receive benefits in each year of the long-run horizon.
- Similarly, one would expect the “reprofiling” effect to change as membership to groups which are positively/negatively reprofiled against changes across time.

Figure 2 plots the impact of the introduction of UC, but this time on a longer-run (24 year) basis. Individuals are ranked by average net equivalised household income. The Figure shows that the reform is more progressive when measured on a long-run basis than on a snapshot basis. Net losses for the poorest 10% of adults fall from 3% to 2% of income while, for the rest of the income distribution, the pattern is broadly unchanged. There are two things to note: first, the predominant

driver behind the reduced long-run impact in decile 1 is a much smaller 'aggregate' effect. This pattern is to be expected from any long-run analysis of a cut to means-tested benefits. Over a long period, many of the poorest adults in the snapshot picture will experience prolonged periods of employment and spells where they are less reliant on benefit income. Second, while the reprofiling effects at the bottom of the income distribution are smaller in the long-run analysis, they are still positive, meaning that UC acts to redistribute a greater share of its resources to the individuals who are poorest in the long-run.

For the remainder of the paper we will focus on this concept of benefit system 'reprofiling' – in other words, we will analyse the way in which UC allocates its support, rather than the reduction in total support that the reform implies. One way to capture this is to compare the share of total support going to adults at different points on the income distribution. This abstracts from the overall level of spend in the benefit system, but does capture how it is allocated. Table 2 divides the population into mutually exclusive cells based on their combination of snapshot and lifetime income, and shows the percentage of total benefit spend going to each of these cells in legacy and UC.

Table 2. Proportion of total support allocated to adults in legacy and UC, split by snapshot and long-run non-means tested benefit income quintile.

		UC			
		Snapshot Quintile			
		1	**2**	**3-5**	**All**
Long-run Quintile	**1**	57.4%	7.5%	1.5%	*66.5%*
	2	13.7%	3.3%	2.3%	*19.3%*
	3-5	7.5%	2.8%	3.8%	*14.2%*
	All	*78.7%*	*13.7%*	*7.6%*	*100.0%*

		Legacy			
		Snapshot Quintile			
		1	**2**	**3-5**	All
Long-run Quintile	**1**	51.7%	9.3%	2.2%	*63.2%*
	2	13.3%	4.5%	3.0%	*20.8%*
	3-5	8.1%	3.3%	4.6%	*16.0%*
	All	*73.1%*	*17.1%*	*9.8%*	*100.0%*

Note: non-means tested benefit income refers to private incomes, plus incomes from the state pension, child benefit, attendance allowance, disability living allowance, and personal independence payment.
Source: Authors calculations using TAXBEN microsimulation model.

There are two things to note from this Table. First, and consistent with the Figures shown above, UC is more successful in targeting individuals in the poorest households in both the snapshot and long-run: 78.7% of support goes to the poorest snapshot quintile in UC, compared to 73.1% in legacy, and 66.5% of support goes to the long-run poorest quintile UC, compared to 63.2% in legacy.

Second, the Table shows the effects of the interaction of being snapshot and long-run poor. Focusing on the support which goes to those in the snapshot *and* lifetime bottom quintile (top left cells), we can see that 57.4% of overall UC support goes to 'snapshot poor-lifetime poor' individuals, compared to 51.7% in legacy. The implication here is that not only does UC allocate a greater fraction of its total spend to those who are snapshot poor, but also those who are long-run poor too.

In the next section we look to explain this result: what are the design features of UC which generate this finding that it is more successful in focusing support on those adults who are not only snapshot poor, but also lifetime poor?

4. Design features of the benefit system, and their redistributive impact

In this section, we explore how specific features of UC and the legacy system focus support more or less towards the snapshot poor and lifetime poor.

4.a Treatment of assets and unearned incomes

UC treats those with financial assets differently to the legacy system. Under the legacy system, families could receive tax credits regardless of their level of financial assets. Under UC, people with assets between £6,000 and £16,000 see their support tapered, and those with assets over £16,000 receive no UC at all. This is true for all would-be UC recipients, including those who under the legacy system would have qualified for Tax Credits.

Figure 3. Average long-run income centile, by snapshot income decile. Adults categorised by household asset level.

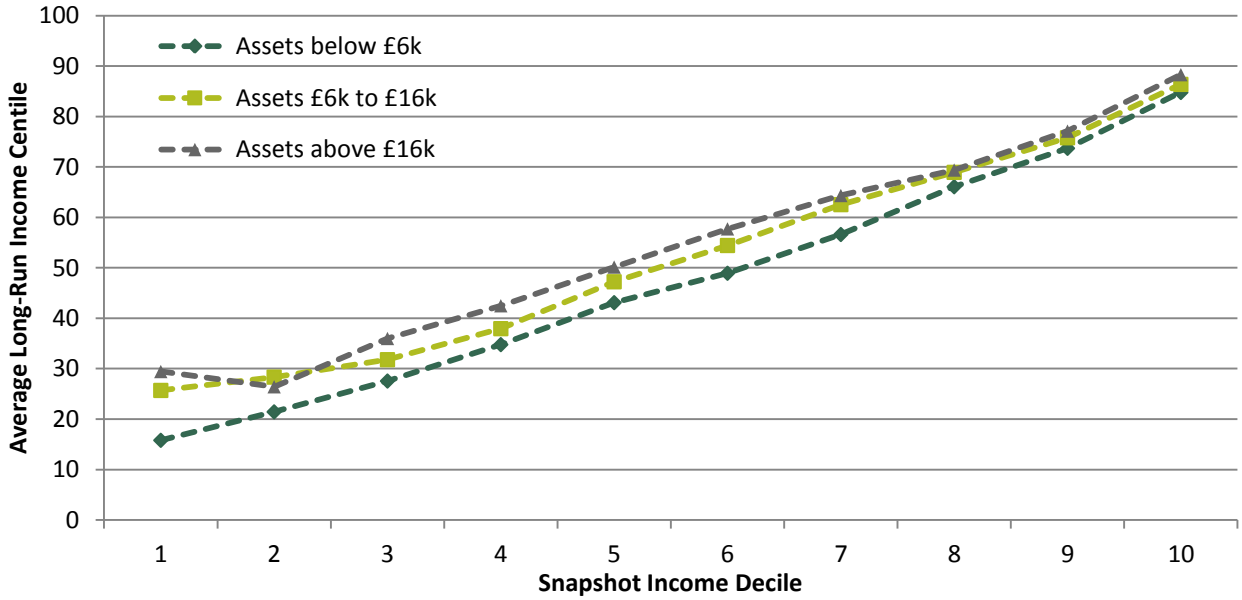

Figure 3 allocates individuals to snapshot income deciles, and computes their average position in the long-run income distribution (measured by their lifetime income centile). It does this for those in different (snapshot) asset groups. The Figure shows that, on average, those with assets tend to experience higher long-run incomes than those without assets, even after conditioning on snapshot income levels, i.e. if one were to select two people with the same level of income at a particular point in time, the person with the greater stock of assets is more likely to experience higher incomes in the future, or have experienced higher incomes in the past. For example, adults in the lowest decile with household assets below £6,000 end up in the 16th centile of the lifetime income distribution on average. This compares to the 29th centile for adults in the same snapshot decile but with household assets in excess of £16,000.

Given that we have established a link between asset ownership and lifetime income, even for given snapshot income, we now turn to quantifying how this harsher treatment of assets leads to the reprofiling of support towards the lifetime poor. Figure 4 shows how the harsher asset rules in UC affects the share of support going different parts of the snapshot and long-run income distribution.

Figure 4. Effect of asset rules in UC on the share of total support going to different parts of the income distribution.

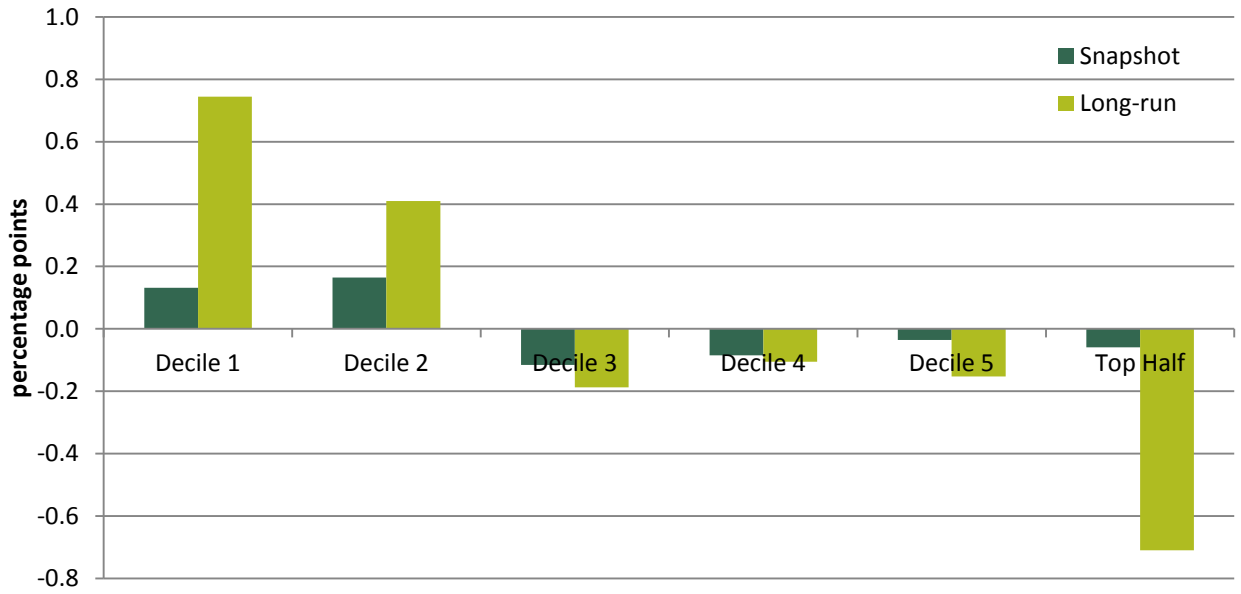

Figure 4 attempts to capture how much more regressive (or progressive) the asset rules in UC are relative to legacy.[8] A positive value means that the tougher asset rules under UC reprofile more support to these individuals. Conversely, a negative value means that that the tougher UC asset rules reprofile support away from these individuals. For instance, the introduction of tougher asset rules in UC increase the share of benefit spend going to decile 1 by 0.7ppts in the long-run. What Figure 4 makes clear is that the tougher asset rules in UC allocate slightly more of the total UC spend to lower deciles in the snapshot, but considerably more when incomes are measured on a long-run basis.

This snapshot result can be explained by comparing the equivalent asset rules. The tougher asset rules in UC affect those who, under legacy benefits, would have been receiving Tax Credits – but does not affect receipt of the out of work benefits (Income Support, Job Seekers' Allowance, Employment and Support Allowance), nor housing benefit. Tax Credits recipients are disproportionately in work – and so tend to be concentrated in deciles 3, 4 and 5, Thus, in the snapshot, the tougher asset tests reduce the share of support going to those deciles, and increase it in the bottom two. In the long-run, this effect is exacerbated by the fact that asset owners tend be higher up the lifetime income distribution (as shown in Figure 3).

[8] Having asset rules rather than no asset rules – for both legacy and UC – is regressive in the snapshot, as they reduce the share of support going the poorest individuals. In the long-run, asset rules increase this share.

4.b Treatment of self-employed incomes

A second area whereby UC potentially generates an interesting divergence in reprofiling between snapshot and long-term poor is in the treatment of self-employed incomes. A notable feature of UC – absent from the legacy system – is the adoption of a Minimum Income Floor (MIF), where self-employed individuals with earnings less than the MIF are treated as though they earn the MIF. For most people the MIF is equivalent to 35 hours at the National Minimum Wage. Thus, low income self-employed people are usually worse off in UC than in legacy – sometimes considerably so. For example, a single, self-employed person who currently has no earnings will typically see their entitlement in UC reduced by £8,250 per year as a result of the MIF.[9]

Figure 5. Average long-run income centile, by snapshot income decile. Adults categorised by highest household earnings source.

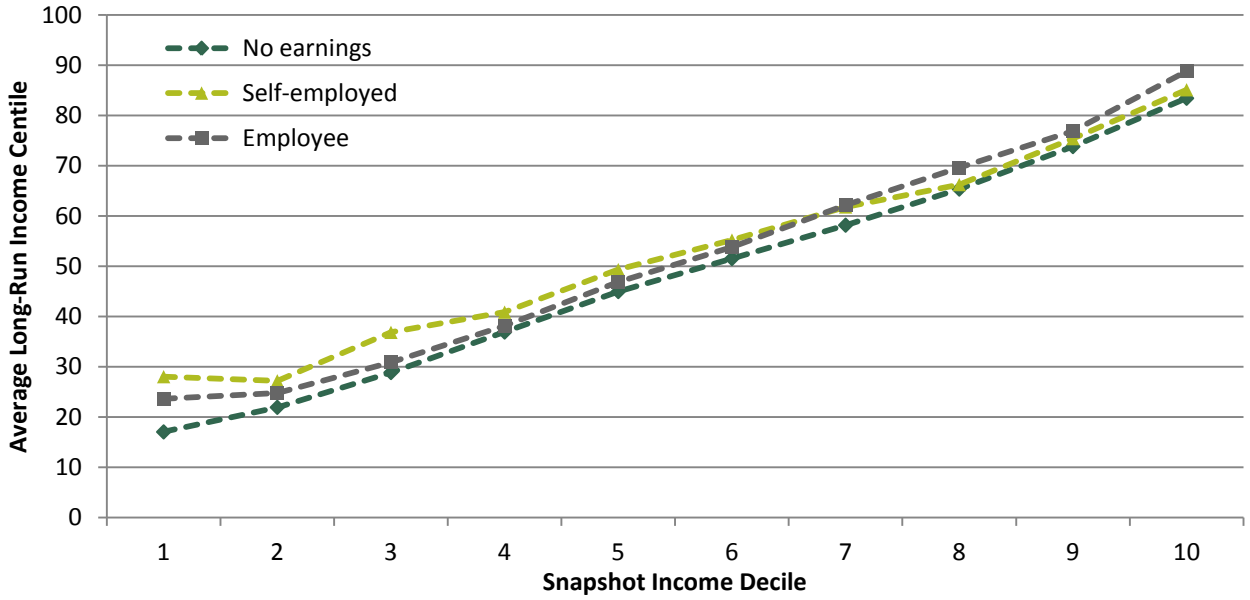

Repeating the same exercise as in section 4.a, Figure 5 shows the interaction between snapshot and long-run incomes but split by household earnings status (defined as the highest source of income in the household). Focusing on the bottom half of the income distribution, we can see that, conditional on snapshot income, adults in households with predominantly self-employed earnings are more likely to have higher lifetime incomes than either households without any earned income, or households

[9] For the purposes of choosing the appropriate National Minimum Wage, we have assumed the individual is older than 25.

with predominantly employee incomes. Amongst self-employed adults in snapshot decile one, the average centile ranking in the lifetime income distribution is 4 centiles higher than for adults in households where employee earnings are the dominant source.

We now turn to analysing the impact of the MIF on the distribution of benefit support by both snapshot and long-run income.

Figure 6. The reprofiling impacts of the MIF in UC, across the income distribution.

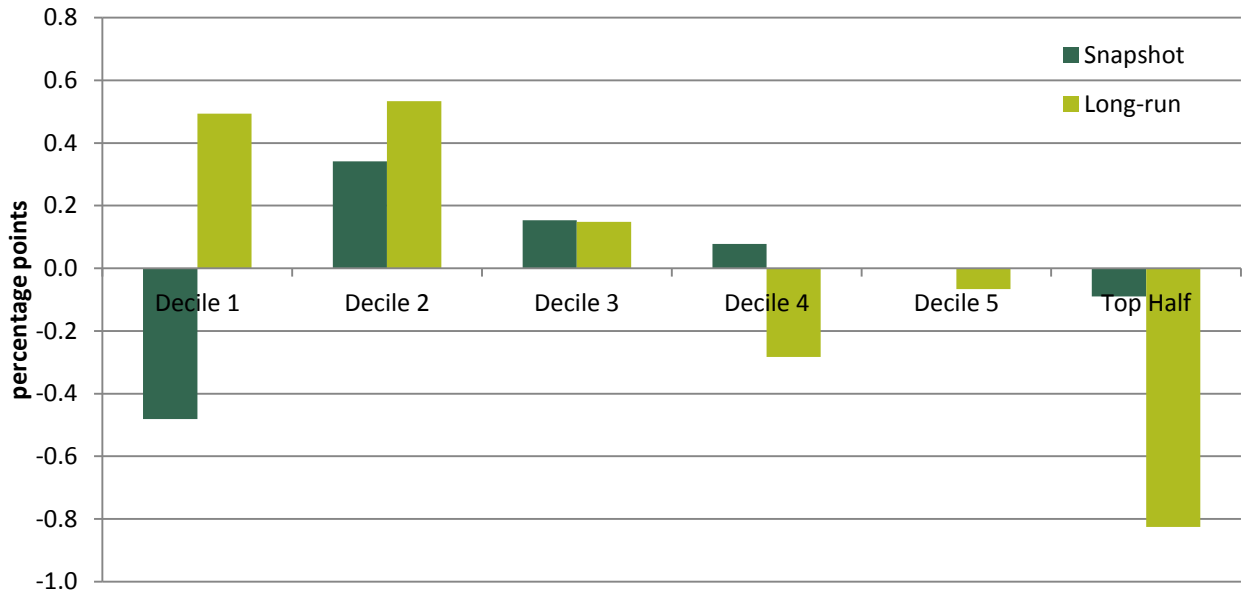

Figure 6 shows how the MIF reprofiles support to different parts of the snapshot and long-run income distribution. The MIF reduces the share of support going to the snapshot poorest decile by 0.5ppts, but increases the share going to deciles 2 and 3 by 0.3ppts and 0.2ppts, respectively. When adults' incomes are measured over 24 years however, the share of support going to the bottom 3 deciles uniformly increases. The reduction in support going to the snapshot poorest is because the MIF specifically reduces benefit entitlement for lowest-earning self-employed individuals (who are largely in the bottom income decile). Conversely, while deciles 2 and 3 do contain low-earning self-employed individuals, this group is a far smaller proportion of the overall UC claimant group, and so their UC, as a proportion of the total UC spend, increases. In the long-run, low income self-employed adults receiving benefits in the snapshot still have relatively high long-term incomes (as seen in Figure 5), meaning that the policy reduces incomes in higher long-run deciles. This suggests that designing a

system which is harsher on self-employed incomes (even those with the very lowest snapshot incomes) will be more effective at targeting available support to the lifetime poor.

4.c Withdrawal of support for in-work renters

A third feature of the UC system is the relative generosity to in-work renters, by getting rid of the overlapping tapering of Housing Benefit and Tax Credits. For a renter paying National Insurance and income tax, transitioning to UC could reduce their total effective marginal tax rate (EMTR) from around 90% to 75%. For homeowners, the level of work allowances in UC means that they receive less in-work support than they receive under Tax Credits. Previous research has shown how UC, on average, increases benefit entitlement for in-work renters while reducing support for (in work and out of work) owner occupiers (Browne, Hood and Joyce 2016). To the extent that renting is a proxy for lower lifetime income (i.e. owning your home is a proxy for higher lifetime income), designing a system which favours renters in the snapshot will be more effective at reprofiling support to the long-term poor than if it favoured homeowners with the same current level of income.

Figure 7. Average long-run income centile, by snapshot income decile. Adults categorised by household tenure.

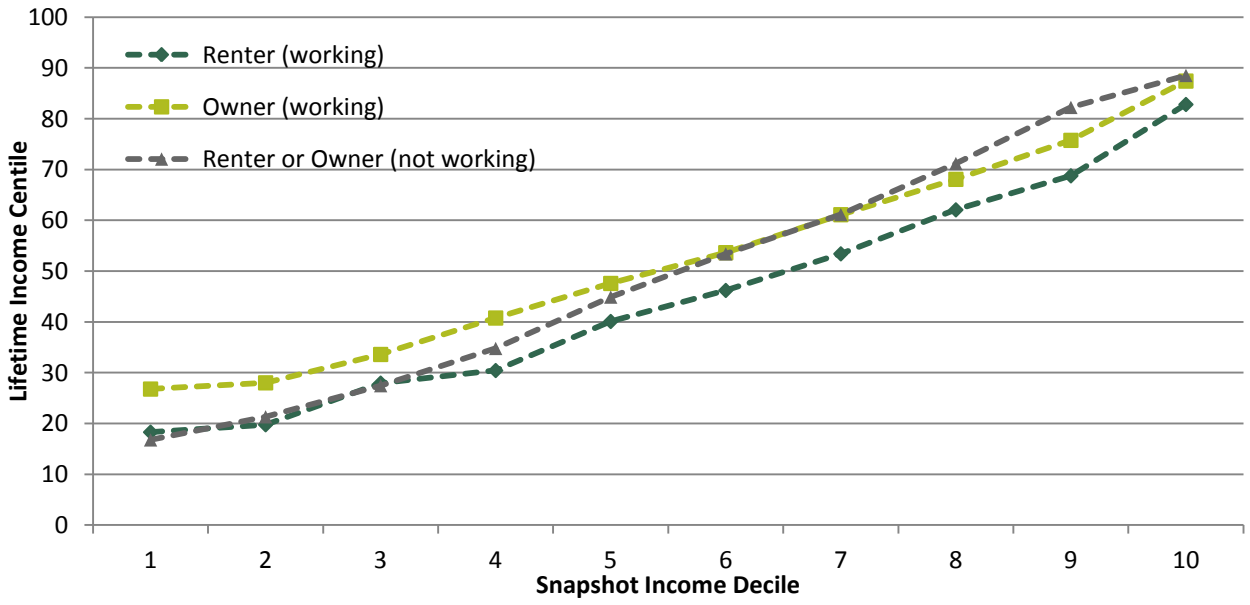

Figure 7 shows that, across the income distribution, being a renter is associated with lower lifetime income even conditional on snapshot income. Conditional on adults being in the poorest 10% of

households at a point in time, those who rent, on average, end up 8 centiles lower in the lifetime income distribution than those who own their home. This divergence is greatest in decile 4 where renters end up 10 centiles lower in the lifetime income distribution. Given that UC reprofiles away from working home owners and towards working renters, this again has implications for the share of support going to the lifetime poor.

Figure 8. The reprofiling impact of more generous renter treatment in UC, across the income distribution.

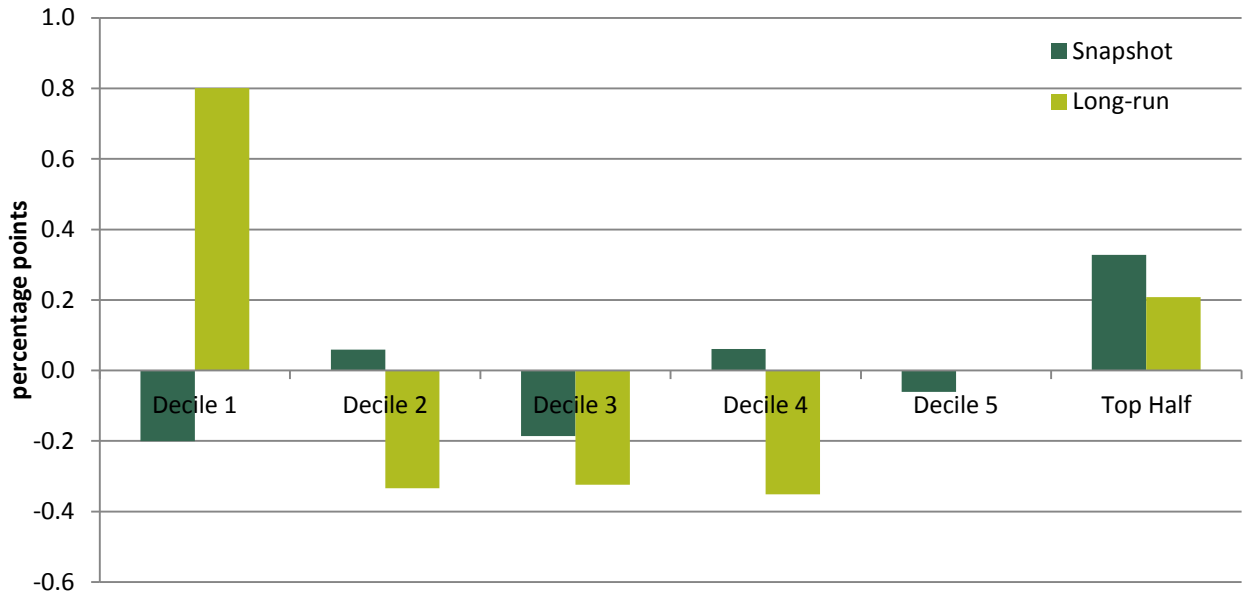

As before, we estimate the impact that this design feature has on the distribution of support. We can approximate this reprofiling effect (of UC's more generous treatment of working renters and less generous of working homeowners) by constructing a counterfactual UC system in which in-work renters enjoy the same benefit entitlement in UC as they did under legacy. This is an imperfect measure of the effect of UC's more generous treatment to working renters because it also captures any *other* differences between legacy and UC which affect working renters' entitlements. However, it should give an indication of the sort of effect that this design feature has.

Figure 8 shows that, in the snapshot, reprofiling support towards in-work renters has little impact on the distribution of support. In the long-run however, the more generous treatment of working renters clearly increases support to the lifetime poor; with the percentage of support going to the poorest

10% increases by 0.8ppts. This is because most people are in work over much of their lifetime, and renting is good proxy for lower lifetime earnings.

Figure 8 also shows increases in the share of support going to the top half of the income distribution in both the snapshot (0.3ppts) and long-run (0.2ppts). This captures the fact that, because of the lower taper rates discussed in section 2.b, for some groups UC entitlement extends further up the income distribution. While these people are not necessarily receiving large UC awards, it is still more than their legacy equivalent.

4.d Reducing premia for severely disabled

The final feature of the UC system examined in this section is the available support for disabled claimants. There are winners and losers from the change in how UC treats disability, but it tends to reduce entitlement for those deemed as more severely disabled.

One would expect a benefit system with reduced entitlement for the severely disabled to be less effective at targeting support to those with low lifetime incomes, and therefore have the opposite impact to the other features examined hitherto.

Figure 9. Average long-run income centile, by snapshot income decile. Adults categorised by disability status.

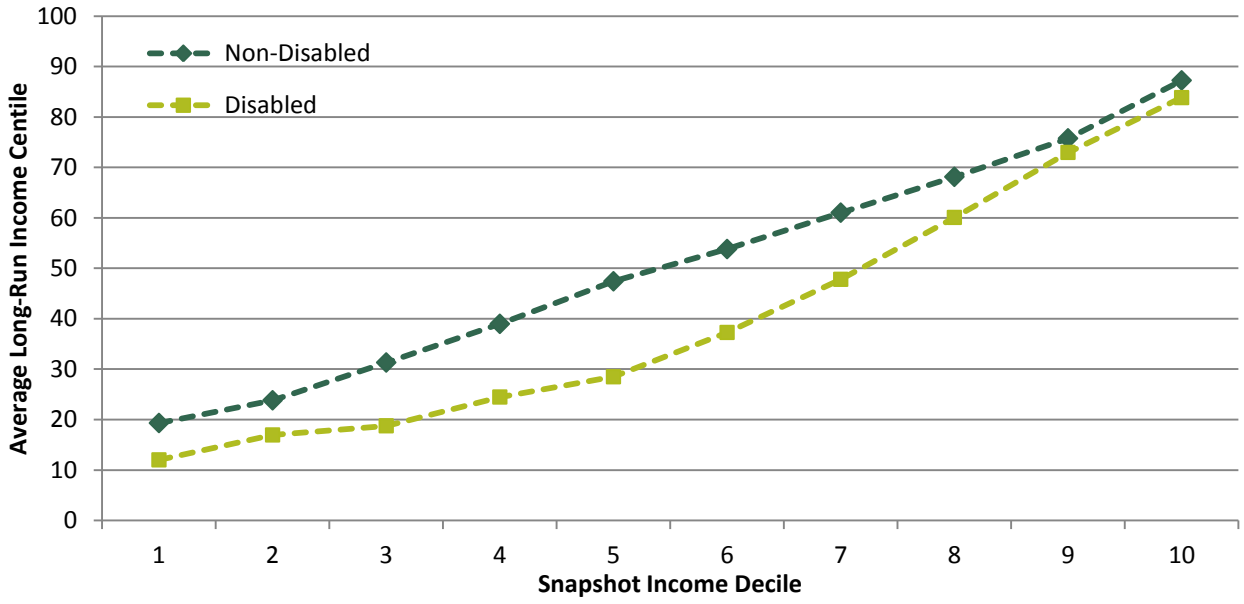

Figure 9 splits adults by reported receipt of a disability benefit, and plots snapshot income against lifetime income. It shows a clear pattern between disability and lower lifetime income across the entire distribution, even conditional on current income. Conditional on being in the snapshot poorest 10% of households, disabled adults can, on average, expect to end up in the 12th centile of the lifetime income distribution. This compares with the 19th centile for non-disabled adults. This result is most pronounced in decile 5, where the difference in average lifetime income ranking is 19 centiles.

Figure 10. The reprofiling impact of UC's treatment of disability, across the income distribution.

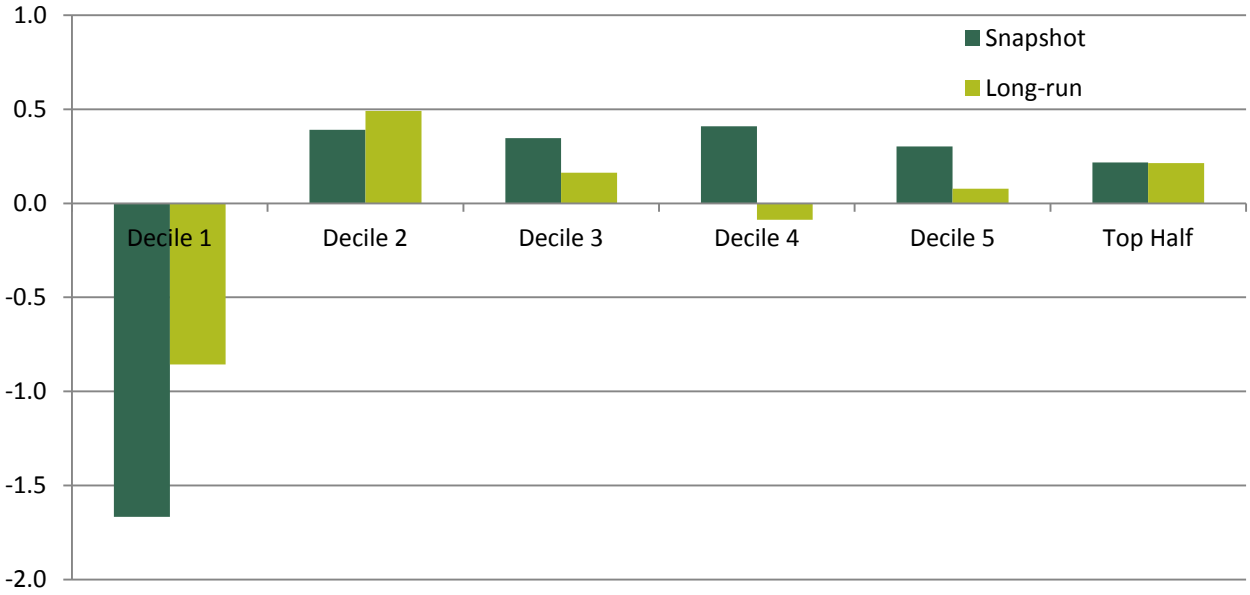

Figure 10 shows that the change introduced by UC regarding disabilities reduces the share of support going to the very poorest adults in the snapshot. This is not surprising given that, at any point in time, these adults will predominantly be out of work. Furthermore, this impact persists in the long-run. This is to be expected given that severe disability is likely to persist over several years.

4.e Summary

These results have important implications:

- The introduction of UC represents an aggregate cut to welfare entitlement and is a regressive reform. Overall welfare spend is forecast to be lower in UC, and Figure 1 and
- Figure 2 showed how the losses are, on average, felt most by the poorest adults for whom benefit receipt is a more important income source.

- In section 3 we also showed that, for a given amount of benefit spend (and therefore ignoring the effects of the aggregate cuts), UC actually prioritises and reprofiles support towards the poor. Of the support going to the snapshot poor, UC is better targeted at those who are also lifetime poor.
- While harsher treatment of assets, self-employed incomes, and home ownership reprofiles support away from the poorest in the snapshot, the people affected are relatively likely to be better off on a long-term basis. Hence these features actually allow for greater targeting of the lifetime poor.
- However, the treatment of disability in UC moves in the opposite direction. Reducing the generosity of disability premia for the severely disabled weakens targeting of support to the poorest both in the short and long-run.

5. Impact of hypothetical reforms to UC

In this section we turn our attention to UC-specific reforms, prevalent in the current policymaking debate, and discuss their respective merits in the trade-off between alleviating short-run hardship and targeting the long-term poor.

We consider four alternative policies and compare their snapshot and long-run impacts. These alternatives have been designed to generate comparable snapshot cost consequences, and can be categorised into two groups: reforms which target in-work claimants, and reforms which span both in- and out-of-work claimants. They are:

- A reduction in the UC taper rate (in-work);
- An increase in the monthly Work Allowance (in-work);
- An increase in the Child Element (in-work and out-of-work); and
- An increase in the Standard Allowance (in-work and out-of-work).

Figure 11. Snapshot impact of UC reforms as a % of net household income, by income decile.

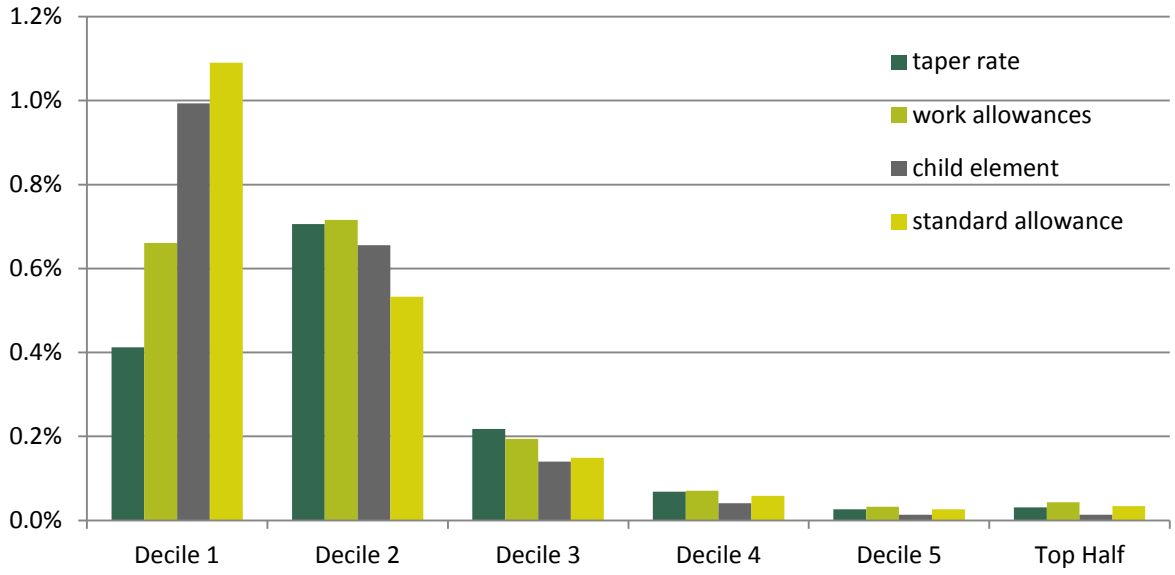

Figure 12. Long-run marginal impact of UC reforms as a % of net household income, by income decile.

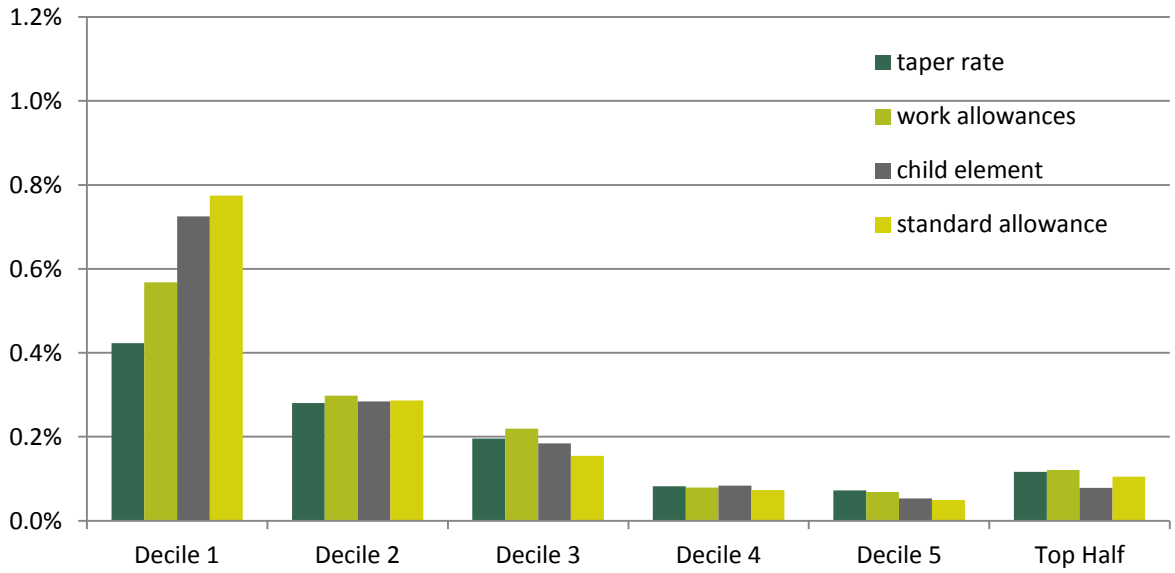

Figure 11 shows the snapshot marginal impact of each of the reforms. Note that here we show the impact of reforms on incomes (as a percentage of net household income), rather than the impact on the distribution of total support . There are two things to note from the Figure. First, the gains across all four policies are concentrated in the bottom two deciles. This is unsurprising given that recipients of UC will be heavily concentrated lower down the income distribution. Second (and also unsurprising), the two most progressive reforms are the options which span both in-work and out-of-

work adults: the rise in the standard allowance generates the largest gains (1.1%) amongst the poorest adults. This is followed by the boost in the child element (1.0%). This can be explained by the fact that the bottom deciles contain the largest shares of non-working individuals, who are more likely to be on benefits. While the taper rate reduction and work allowance increase do still represent giveaways to low income benefit recipients, the largest gains are seen in decile 2 as opposed to the very poorest. Approximately 20% of the total spending giveaway in the taper rate option goes to the poorest snapshot decile, compared with nearly half of the standard allowance giveaway.

Over the long-run however, the picture changes and is shown in Figure 12. While the in-work options were somewhat progressive in the snapshot, they become noticeably more progressive in the long-run. The largest gains now accrue to the poorest lifetime adults, and the overall net impact profile is much closer to the policies which also span out-of-work claimants. This is explained by the fact that, over the long-run, the poorest individuals are still likely to spend prolonged periods of time in work where they will benefit from the lower taper rate/ higher work allowance. This supports previous research which found similar lifetime impacts between in-work and out-of-work reforms within the legacy benefit system (Levell, Roantree and Shaw 2015).

For policymakers, this emphasises the need to be clear on what the intended purpose is and who the intended targets are of benefits reform. If the intention is to tackle long-run poverty, then Figure 12 suggests that reforms which target in-work claimants can have similar distributional impacts to those which also target out-of-claimants. These policies have the additional bonus of improving work incentives. However, if the intention is to increase incomes during temporary periods of low income, increasing the standard allowance is more effective.

6. Conclusion

In this paper, we have estimated the long-run distributional impact of the introduction of UC and shown that, while benefit systems tend to become less effective at targeting poor over long-run horizons, UC is better equipped to target support to the poorest adults. We have also studied a number of design features of UC to identify how specific aspects of a benefit system contribute to this increased targeting of support to the lifetime poor. We find that harsher treatment of assets and of self-employed incomes, as well as more generous treatment of renters, reprofiles support towards those who experience low incomes over the long-run. Conversely, reduced support for the severely disabled means that, along this margin, UC is actually less effective at targeting support to both the snapshot and long-term poor.

Turning our attention to potential, UC-specific reforms, we find that changes to the taper rate and work allowances – which directly target in-work claimants – are markedly less progressive in the snapshot than alternatives which increase support to out-of-work UC recipients. Over the long-run however, this gap in progressivity is narrowed.

There are several reasons why this research should be of interest to policymakers. First, while UC represents an overall cut to welfare entitlement – and is therefore regressive – we have shown that, for a given level of benefit spend, it is more effective at targeting the snapshot and lifetime poor. Second, previous research has highlighted the importance of considering the long-run impacts of welfare policy but this paper goes one step further by examining how specific features of benefit systems are effective tools to reprioritise support between those with temporary low incomes, and those in long-run poverty. These considerations are important when governments are looking for ways to redesign benefit systems in a manner which minimises losses for the poorest. Third, supporting previous research, the comparison between the distributional impacts of in-work and out-of-work UC reforms highlights the significance of not restricting policy making to the conventional snapshot mindset of either being employed or unemployed.

References

Browne, J., Hood, A., and Joyce, R. (2016), 'The (changing) effects of universal credit', Institute for Fiscal Studies, Green Budget 2016 Chapter 10.

Levell, P., Roantree, B., and Shaw, J. (2015), 'Redistribution from a lifetime perspective', Institute for Fiscal Studies, Working Paper W15/27.

Roantree, B., and Shaw, J. (2018), 'What a difference a day makes: inequality and the tax and benefit system from a long-run perspective', *The Journal of Economic Inequality*, 16, 23-40.

Waters, T., (2017), 'TAXBEN: The IFS tax and benefit microsimulation model', Institute for Fiscal Studies – https://www.ifs.org.uk/publications/12858

Appendix 1: Simulating legacy and UC without asset rules

We use TAXBEN to create two hypothetical counterfactual tax and benefit systems:

- A legacy counterfactual in which the capital rules and unearned income disregard in Income Support and Tax Credits are turned off.
- A separate UC counterfactual in which the capital rules regarding tariff income and capital limits are turned off.

In other words, we consider what would happen if neither the UC nor legacy systems penalised assets or unearned incomes.

By examining and contrasting the percentage of support going to each decile between: i) the legacy counterfactual with actual legacy; and ii) the UC counterfactual with actual UC, we can isolate and compare the reprofiling impacts of the relevant asset rules of the two benefit systems.