Economics Letters 192 (2020) 109230

Contents lists available at ScienceDirect

Economics Letters

journal homepage: www.elsevier.com/locate/ecolet

A method for decomposing the impact of reforms on the long-run income distribution, with an application to universal credit

Mike Brewer [a,b], Robert Joyce [a], Tom Waters [a,*], Joseph Woods [c]

[a] Institute for Fiscal Studies, 7 Ridgmount Street, London, WC1E 7AE, UK
[b] University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, UK
[c] HM Treasury, 1 Horse Guards Rd, Westminster, London, SW1A 2HQ, UK

ARTICLE INFO

Article history:
Received 9 January 2020
Received in revised form 11 May 2020
Accepted 12 May 2020
Available online 16 May 2020

JEL classification:
D31
H23
H24
I38

Keywords:
Inequality
Redistribution
Income mobility
Lifetime
Taxes
Benefits

ABSTRACT

Income inequality, as well as the impact of tax and benefit reforms on it, has typically been evaluated with respect to 'snapshot' incomes, measured over short periods such as one week or year. But longitudinal data allows long-run measures of income to be used, which will be of interest to policymakers interested in persistent, rather than only temporary, poverty. We show that the long-run distributional impact of a reform is the combination of three effects: a 'static' effect, which would be observed if individuals' circumstances were consistent throughout their life; an 'income dynamics' effect, resulting from individuals moving around the income distribution over time; and a 'tagging' effect, resulting from the reform affecting individuals differently according to whether they have a characteristic predictive of long-run income conditional on current income. We propose a simple method to decompose these three effects for any inequality, poverty, or distributional statistic. We use the method to examine the distributional impact of the introduction of 'Universal Credit', the most important reform to the UK benefit system in decades. We show that Universal Credit is less regressive on a long-run basis than a snapshot one, partly because of income dynamics but also because it reduces entitlements for (or 'negatively tags') those who are more likely to find a period of low income to be temporary, rather than persistent.

1. Introduction

Analysis of the distributional impacts of tax and transfer policies has traditionally examined the effects on incomes measured over short periods (a 'snapshot'), typically a week or year (e.g. Hoynes and Rothstein, 2019; Figari et al., 2015; Congressional Budget Office, 2019). But researchers are increasingly using longitudinal data to evaluate distributional impacts on long-run measures of income (e.g., Huggett and Parra, 2010; Brewer et al., 2012; Levell et al., 2016; Haan et al., 2017; Bartels and Neumann, 2018; Roantree and Shaw, 2018). Both of these sorts of analyses tell us something important. The snapshot impact is useful because part of the function of the welfare state is to protect families against short-term hardship or more generally redistribute towards periods of low income. But the long-run impact provides an important complement, as consumption smoothing behaviour can make short-term incomes a poor reflection of true living standards, and policymakers may care more about persistently low living standards than temporarily low ones.

A key finding of this literature is that – even when governments base tax liabilities and transfer entitlements only on current circumstances (e.g. incomes, assets, family composition) – the distributional consequences of reforms can be significantly different when analysed on a long-run basis. A very common result is that the long-run impacts of reforms are substantially more distributionally neutral than the snapshot impacts (e.g. Levell et al., 2016; Roantree and Shaw, 2018). This is because people move around the (snapshot) income distribution over time, attenuating the long-run differences in how people are affected by a reform.

We show that the long-run distributional impact of a reform is the combination of three effects. First, a 'static' effect: the distributional impact that would occur if individuals' circumstances never changed. Second, an 'income dynamics' effect: peoples' incomes change over time, so (for example), a reform which only affects the poor in the snapshot will affect some who are richer on a long-run measure. This is why long-run distributional effects of reforms are typically attenuated compared to the snapshot. Third, the snapshot impact of a reform may be correlated with an individual's long-run income conditional on their snapshot income. Borrowing a phrase from the optimal tax literature (Akerlof, 1978), we term this effect 'tagging': a reform 'tags' those

* Corresponding author.
E-mail addresses: mbrewer@essex.ac.uk (M. Brewer), robert.joyce@ifs.org.uk (R. Joyce), tom.waters@ifs.org.uk (T. Waters), joseph.woods@hmtreasury.gov.uk (J. Woods).

https://doi.org/10.1016/j.econlet.2020.109230
0165-1765/

with a certain characteristic if it affects them differently to others with the same income but without the characteristic. Like optimal tax tagging, here the government makes transfers conditional on something observable and predictive of something unobservable (in our case, long-run income; in optimal tax, ability). It is different to optimal tax tagging in that it is not necessary that the characteristic be immutable. This form of tagging creates an additional wedge between the snapshot and long-run distributional impacts of a reform.

We propose a method to decompose the long-run progressivity of a reform into these three components. Disentangling these impacts is useful because (absent behavioural response) tagging is the only way policymakers can change the long-run distributional consequences of a policy package for a given snapshot effect. Simply comparing the snapshot and long-run distributional impacts of policy options will often obscure that part of the comparison that policymakers can independently vary (the tagging effect).

To highlight the usefulness of this approach we analyse several policy reforms, including the transition in the UK to 'Universal Credit' (UC), the most important reform to the UK benefit system in decades, and one with sizable tagging effects.

2. Method

Standard snapshot distributional analysis uses a cross-sectional dataset of individuals. Net incomes (hereby 'incomes') are calculated under two tax and benefit systems (pre- and post-reform). Typically these analyses ignore any possible resulting behavioural changes, which we also do throughout this paper. This gives pre- and post-reform income distributions, from which inequality and poverty statistics can be calculated. In long-run distributional analysis, the same exercise is conducted using a panel dataset, with every observation of every individual run through the same two tax and benefit systems. Under each system, incomes are averaged across observations for every individual to compute their 'long-run income'. Again various statistics can be calculated with these long-run distributions. Our approach for decomposing the impact of a reform on long-run income does not assume any particular length of panel or income process. However, clearly average income over several periods is more interesting to the extent that it captures some long-run measure of income. This is more plausible the longer the panel and the better behaved the underlying income process.

Our method for decomposing the impact of a reform on the long-run distribution is as follows. Suppose we have a balanced panel, with T periods and N individuals. $y_{i,t}$ is the income for individual i at time t in the pre-reform system; $z_{i,t}$ the post-reform. Long-run incomes are defined as the individual's average income over the T periods, and are given by y_i and z_i. The pre-reform (snapshot) income distribution at t is a vector of N observations of $y_{i,t}$, $\mathbf{y_t} = [y_{1,t}, \ldots, y_{N,t}]$. $\mathbf{z_t}$ is the analogous post-reform distribution, and $\mathbf{y}$ and $\mathbf{z}$ are the distributions of long-run incomes.

For any inequality, poverty, or distributional statistic I we denote the value of I applied to long-run incomes I_l. We decompose the impact of the reform on I_l as follows:

$$\Delta I_l = \underbrace{\Delta I_s}_{\text{static}} + \underbrace{(I(\mathbf{h}) - I(\mathbf{y}) - \Delta I_s)}_{\text{income dynamics}} + \underbrace{(I(\mathbf{z}) - I(\mathbf{h}))}_{\text{tagging}} \quad (1)$$

ΔI_s is the impact of the reform on I applied to short-run incomes. We derive this by pooling the snapshot distributions across time periods, and calculating the static as:

$$\Delta I_s = I(\mathbf{z_1}, \ldots, \mathbf{z_T}) - I(\mathbf{y_1}, \ldots, \mathbf{y_T})$$

$\mathbf{h}$ is the distribution of income resulting from a counterfactual reform, which has the same snapshot distributional impact as the actual reform but where an individual's long-run income and the proportional snapshot impact of the reform on them are conditionally independent, given their snapshot income, i.e., $\left(\left(\frac{h_{i,t}}{y_{i,t}}\right) \perp y_i\right) \mid y_{i,t}\ \forall\, i, t$. Calculating this distribution is useful because it helps us distinguish income dynamics from tagging: because the effect of this counterfactual reform is conditionally independent of an individual's long-run income, by construction it has no tagging effects.

We construct $h_{i,t}$ incomes as follows. First, we pool observations across periods and separate them into Q quantiles based on their pre-reform income, indexed by q. The set of individuals who at time t are in quantile q is given by $\gamma_{q,t}$.[1] The ratio of pre- and post-reform snapshot incomes, across all T periods, in snapshot quantile q is given by:

$$\alpha_q \equiv \frac{\sum_{t=1}^{T} \sum_{i \in \gamma_{q,t}} z_{i,t}}{\sum_{t=1}^{T} \sum_{i \in \gamma_{q,t}} y_{i,t}}$$

For each i and t, $h_{i,t}$ is given by:

$$h_{i,t} = y_{i,t} \cdot \alpha_{q(i,t)}$$

Note that the α parameter depends on the snapshot quantile (q) that i is in at t. Thus, the impact of the counterfactual reform on an individual is *only* dependent upon their snapshot income, and, conditional on that, independent of their long-run income. Hence, the counterfactual reform has no tagging effects. However, it has the same snapshot distributional impact as the actual reform, since the α parameters are derived from the snapshot impact of the actual reform. From the set of individual-period counterfactual incomes $h_{i,t}$ we derive the counterfactual long-run distribution, $\mathbf{h}$.

We label the first term in Eq. (1) the 'static' effect because it is the effect that the policy would have on I_l if individuals' circumstances never changed. This is of course equivalent to the traditional snapshot effect. The second term we call the 'income dynamics' effect because it measures how the snapshot and long-run effects would differ if there was no tagging, i.e. it is the consequence of the income processes that people with different incomes (but not different characteristics) experience. For a given population, the snapshot effect uniquely determines the income dynamics effect, with the mapping a function of individuals' income processes. The final term is the 'tagging' effect because the $\mathbf{z}$ and $\mathbf{h}$ distributions differ only insofar as the reform treats people with the same snapshot income – but different long-run incomes – differently.[2]

By construction, income dynamics and tagging contributions sum to zero. In principle the results are sensitive to the number of quantiles Q used. In practice a large Q will often be sufficient, since within a narrow quantile the effect of the reform typically will not differ much between the top and bottom of the quantile. Testing with different Q can indirectly confirm this.

[1] Quantile bounds are defined by pooling all periods, rather than within period.

[2] Because here we use pre-reform incomes to put individuals in q quantiles when deriving α_q, the tagging component should be interpreted as showing the impact of the reform differently affecting those with the same snapshot *pre-reform* income but different long-run *pre-reform* incomes. For this reason, it is most natural to examine the effects of the reform on statistics (I) which rank individuals by pre-reform incomes (such as the effect on average incomes in a particular pre-reform quintile, the statistic we analyse below). To examine the effect on statistics which involve re-ranking – such as percentiles of the distribution – it would be more natural to calculate α_q by putting individuals in q quantiles using a measure of income that is invariant to the tax and benefit system, such as gross (pre-tax and benefit) income.

Table 1
Proportional impact (%) on snapshot incomes among those in bottom snapshot quintile.

Reform	Long-run quintile		
	Bottom	Top four	All
Inc. support for OOW families	1.6	0.6	1.3
Abolish asset tests	0.8	2.1	1.1

3. Analysis of reforms

We use this method to analyse reforms to the UK tax and benefit system. We use 2009–10 to 2016–17 data from 'Understanding Society', the UK Household Longitudinal Study (UKHLS), a panel survey of UK households.[3] Our sample is individuals observed as adults in all eight waves.[4] We uprate financial variables in the data such as earnings in line with their respective average growth rates, (e.g. we uprate earnings from past years in line with average earnings growth since that year) so that they are in a common year's values.[5] We simulate tax liabilities and benefit entitlements using TAXBEN, the IFS microsimulation model, to calculate household (equivalised) incomes.[6] We set $Q = 100$, though results are little changed with other large values. The statistic we examine (ΔI) is the proportional impact of the reform on average (pre-reform) incomes by quintile.

3.1. Increasing out-of-work benefits vs. removal of asset tests

To illustrate the method, we compare two reforms with the same fiscal impact: increasing out-of-work benefit entitlements for claimants with children by £25.50 per week[7]; and 'abolishing asset tests', so benefit entitlement is not affected by financial assets.[8]

Both reforms largely benefit those who are low income in the snapshot. But the snapshot impact among such individuals varies significantly according to whether the individual is low income in the long-run too. Table 1 shows the impact of the policies on the snapshot incomes of those in the bottom snapshot quintile, split by whether they are temporarily or persistently poor (top four or bottom long-run quintile respectively). Those who find their period of low income to be temporary gain much more in the snapshot from abolishing asset tests, while those who find their period of low income to be persistent gain more from the increase in support for workless families. This is driven by the fact that having assets is, conditional on currently having a low income, predictive of higher long-run incomes; while being out-of-work with children is, conditional on currently being low income, predictive of low long-run incomes.

Fig. 1 shows the long-run distributional consequences of these reforms. The static contributions of these policies (equivalently, the snapshot effects) are very similar. The income dynamics components are also very similar (this follows by definition since, for a given population, the snapshot effect uniquely determines the income dynamics effect).

But the total long-run distributional consequences are rather different. While both reforms are less progressive in the long-run than in the snapshot (compare the 'static' bars with the 'total' dots), this is particularly true for abolishing asset tests. This is a consequence of tagging: increasing support for out-of-work families slightly 'positively' tags – allocates a higher proportion of aggregate income to – those who are long-run poor, while abolishing asset tests 'negatively' tags the long-run poor. This is not surprising given the results in Table 1.

The key point is that two reforms with a very similar snapshot distributional picture can have relatively different long-run pictures, if they differentially tag long-run incomes by conditioning support on characteristics predictive of them (such as having significant assets). This approach allows policymakers to adjust the long-run distributional consequences of a policy package *for a given* snapshot distributional effect.

3.2. Universal credit

The UK is in the process of a radical benefit reform, with a single integrated payment – Universal Credit (UC) – replacing the six main means-tested working-age benefits, and changing entitlements for 76% of those entitled (Brewer et al., 2019).

Relative to the system that it replaces, UC's impact is regressive, both in the snapshot and long-run (as shown below), but by less in the latter. This is partly due to the attenuating effects of income dynamics.

However, various features of UC determine not only its snapshot impact but also the relationship between its snapshot and long-run impact. UC on average reduces entitlements for those with significant assets, self-employment income, or who own their own home. As Brewer et al. (2019) show, all three groups are particularly likely to find that a period of low income is temporary rather than persistent. The effect is that UC reduces incomes more for the temporarily poor than the persistently poor: among those in the bottom snapshot quintile, those who are in the bottom long-run quintile see an average loss of 0.8% of their snapshot income, while those in the top four long-run quintiles see an average loss of 2.4%.

That means that UC 'positively' tags those who are long-run poor. This can be seen in Fig. 2, which shows that the tagging effect redistributes from the second quintile to the poorest. Absent such an effect, the long-run impact of the policy would reduce incomes in the bottom quintile by 1.0%, and in the second by 0.4%. But the tagging effect means that the two quintiles are similarly affected (0.8% and 0.7%). Decomposing the tagging effect from the income dynamics and the static effects allows us to see this, highlighting that UC is significantly less regressive in the long-run than one would expect given its snapshot impact — thanks to its policy design, and not only income dynamics.

4. Conclusion

Tax and benefit reforms have impacts on families' snapshot and long-run incomes. Both of these effects matter from a policy standpoint — the effect on the snapshot is informative about how the system protects against short-term hardship, whereas the effect on the long-run shows how the reform affects those with persistently high or low incomes, which may be more informative about their living standards.

The long-run distributional consequences of tax and benefit reforms differ from the snapshot impact because of both income

[3] University of Essex, Institute for Social and Economic Research (2018).

[4] Further methodological details are as described in Brewer et al. (2019).

[5] This is equivalent to modelling the same reform in each period without fiscal drag — so, for example, tax thresholds remain a constant fraction of average earnings.

[6] For a description of TAXBEN, see Waters (2017).

[7] This is achieved by increasing the basic allowances in UC for families with children by £25.50 per week and reducing the 'work allowance' – the amount a claimant can earn before UC starts to be withdrawn – by an offsetting amount such that the income of anyone earning over the work allowance is unaffected. This also benefits those in work but earning under the work allowance: a relatively small group who by definition are not earning much, and so we just refer to this reform as affecting 'out-of-work families'.

[8] Under UC, benefits are tapered away for those with assets over £6000, and removed completely for those with assets over £16,000.

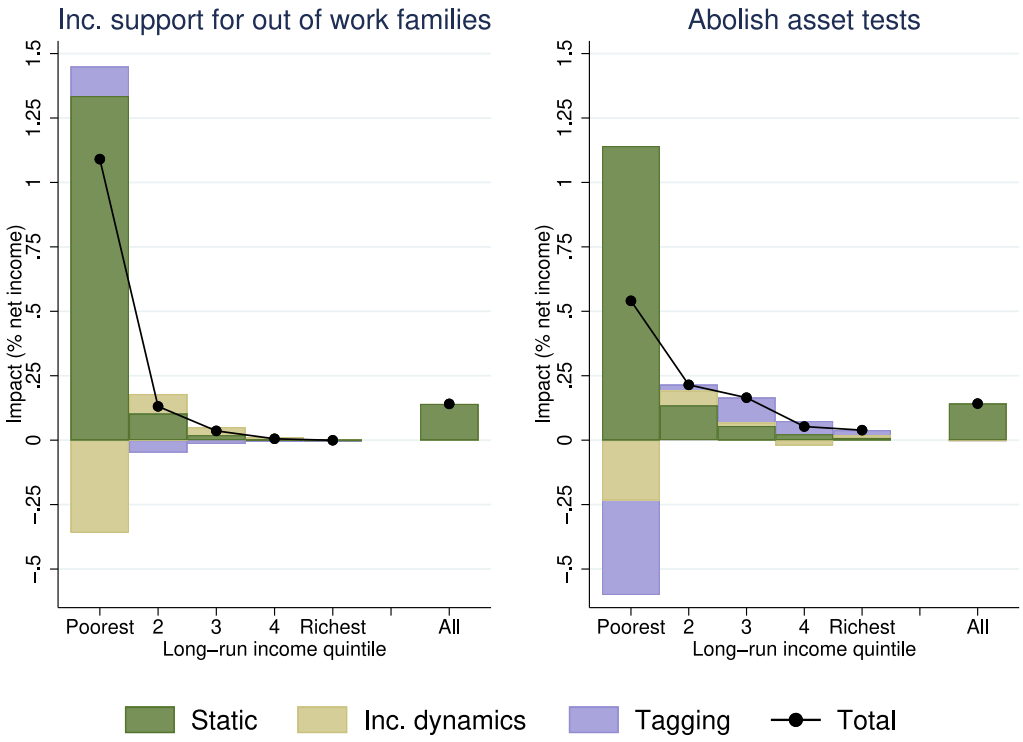

Fig. 1. Long-run distributional impact of increasing support for out-of-work families with children and removing asset tests.

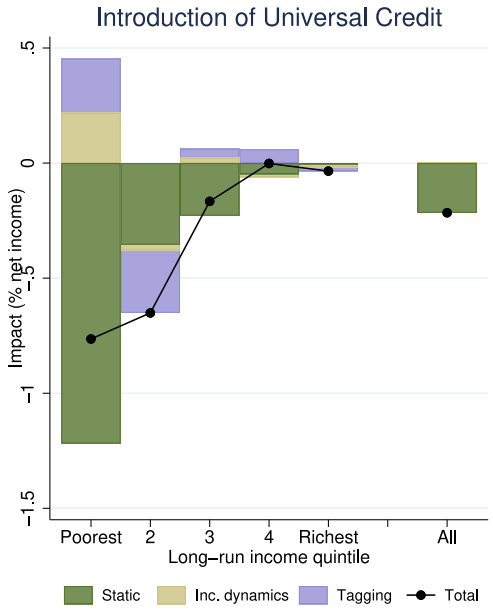

Fig. 2. Long-run distributional impact of Universal Credit.

dynamics and a 'tagging' effect, where the reform differently affects those with characteristics predictive of their long-run income conditional on their snapshot income. These tagging effects can be large, and considerably reduce the regressivity of the long-run distributional consequences of Universal Credit.

Policymakers are often constrained to base tax liabilities and benefit entitlements on current information. Adopting a 'tagging' approach – making transfers conditional on characteristics predictive of long-run incomes – allows them to nonetheless affect the long-run distribution independent of the snapshot; and even if they make no explicit attempt to do so, this form of tagging will often be a feature of their policies in practice. Decomposing the long-run impact allows the tagging component to be isolated and compared across alternative policies.

Acknowledgements

This work was supported by Economic and Social Research Council (ESRC), United Kingdom funding as part of Understanding Society, Waves 6 to 8 (ES/K005146/1). M. Brewer gratefully acknowledges funding by the ESRC, United Kingdom through the Research Centre on Micro-Social Change (MiSoC) at the University of Essex, grant number ES/L009153/1. We are grateful to Stuart Adam, Tom Crossley, Paul Fisher, Peter Levell, and Paul Johnson for advice and comments on this work, as well as seminar participants at the IFS, RES conference, Understanding Society conference, and Essex workshop on longitudinal and administrative data for distributional analyses. Co- funding from the ESRC, United Kingdom-funded Centre for the Microeconomic Analysis of Public Policy (ES/M010147/1) at the IFS is also gratefully acknowledged. Understanding Society is an initiative funded by the ESRC and various government departments, with scientific leadership by the Institute for Social and Economic Research, University of Essex, and survey delivery by NatCen Social Research and Kantar Public. The research data are distributed by the UK Data Service. This paper expresses only the views of the authors and HM Treasury bears no responsibility for its contents.

References

Akerlof, George A., 1978. The economics of "tagging" as applied to the optimal income tax, welfare programs, and manpower planning. Am. Econ. Rev. (ISSN: 0002-8282) 68 (1), 8–19.

Bartels, Charlotte, Neumann, Dirk, 2018. Redistribution and Insurance in Welfare States around the World. SOEPpaper 985.

Brewer, Mike, Dias, Monica Costa, Shaw, Jonathan, October 2012. Lifetime inequality and redistribution. IFS Working Paper W12/23.

Brewer, Mike, Joyce, Robert, Waters, Tom, Woods, Joseph, April 2019. Universal credit and its impact on household incomes: The long and the short of it. Institute for Fiscal Studies Briefing Note BN248.

Congressional Budget Office, 2019. The Distribution of Household Income, 2016. Technical Report.

Figari, Francesco, Paulus, Alari, Sutherland, Holly, 2015. Chapter 24 - microsimulation and policy analysis. In: Atkinson, Anthony B., Bourguignon, François (Eds.), Handbook of Income Distribution. In: Handbook of Income Distribution, vol. 2, Elsevier, pp. 2141–2221. http://dx.doi.org/10.1016/B978-0-444-59429-7.00025-X.

Haan, Peter, Kemptner, Daniel, Prowse, Victoria L., 2017. Insurance, redistribution, and the inequality of lifetime income. SSRN Electron. J. (ISSN: 1556-5068) http://dx.doi.org/10.2139/ssrn.3109312.

Hoynes, Hilary W., Rothstein, Jesse, February 2019. Universal Basic Income in the US and Advanced Countries. Working Paper 25538, National Bureau of Economic Research.

Huggett, Mark, Parra, Juan Carlos, 2010. How well does the U.S. social insurance system provide social insurance? J. Political Econ. (ISSN: 0022-3808) 118 (1), 76–112. http://dx.doi.org/10.1086/651513.

Levell, Peter, Roantree, Barra, Shaw, Jonathan, September 2016. Mobility and the lifetime distributional impact of tax and transfer reforms. IFS Working Paper W16/17.

Roantree, Barra, Shaw, Jonathan, 2018. What a difference a day makes: Inequality and the tax and benefit system from a long-run perspective. J. Econ. Inequal. (ISSN: 1573-8701) 16 (1), 23–40. http://dx.doi.org/10.1007/s10888-017-9362-x.

University of Essex, Institute for Social and Economic Research, 2018. Understanding society: Waves 1–8, 2009–2017 and harmonised BHPS: Waves 1–18, 1991–2009. In: UK Data Service. SN: 6614, [Data Collection]. 11th Edition.

Waters, Tom, 2017. TAXBEN: The IFS tax and benefit microsimulation model. https://www.ifs.org.uk/publications/12858.