

Welfare Reform: A Brief Summary

At its simplest, the welfare state can be thought of as an institution that redistributes the allocation of resources from what would be the 'free market' outcome. This redistribution might take the form of money through the tax system (both direct and indirect taxes), through tax credits, through welfare benefits but also in the form of goods (notably, education and health). For example, if healthcare is provided free of charge and the costs met through general taxation then it is likely to be redistributive in its impact.

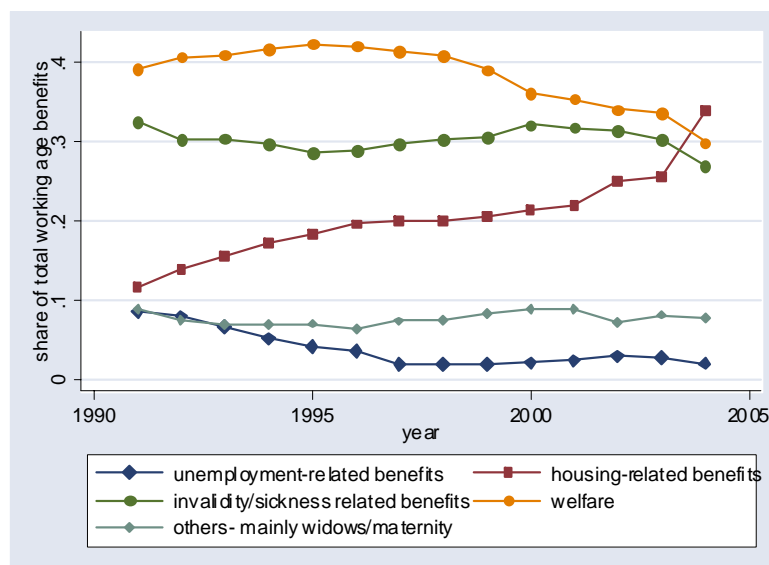
Why do welfare states exist at all? There are three broad arguments

- they provide insurance against 'bad luck' e.g. sickness/unemployment
- they are the outcome of a political process in which the median voter has income below the mean level (because of the nature of the market distribution of income)
- they are 'just' – the free market outcome leads to 'too much' inequality for voters tastes.

Types of Redistribution

There are many forms which this redistribution takes. For example, some of the readings discuss 'welfare' (in the classic American terminology) that provides cash to those who would otherwise have low incomes, 'unemployment insurance' that provides cash to those who would have no income because they are unemployed, and 'workers compensation' that provides cash to those who have no income because they are sick or disabled. These different types of programmes target different groups but the underlying principles needed to understand their likely effects are all very similar so I will concentrate on discussing a generic welfare programme.

To get some idea of the relative importance of these different types of welfare benefit the following figure shows the shares of broad types for the UK



There are several points worth noting that are common to many countries:

- unemployment-insurance related benefits are now often relatively unimportant especially in countries where unemployment is much lower than it used to be
- sickness and invalidity benefits have grown and remain stubbornly high
- benefits to lone parents are also very large

The Objectives of the Welfare State


Policy-makers have struggled for many years to accomplish three goals with welfare programmes:

- raise the living standards of low-income families
- provide incentives to work
- keep costs low

The conflict between these three objectives is sometimes called the “iron triangle” of welfare reform as it is difficult if not impossible to improve outcomes in one dimension without worsening them in another.


At different times, emphasis has been placed more on some of these objectives than others. In the 1960s the focus was poverty reduction and no-one worried too much about adverse effects on incentives to work. In the 1980s the emphasis was on cost reduction, currently it is more on work incentives.

Lets try to understand the trade-offs in the “iron triangle”. Consider the following diagram that relate market income to final income:



In the absence of a welfare state the relationship between market and final-income will be

a 45 degree line, represented by the dotted line in Figure 1. But with a welfare state, we will have something different say something like the solid line in Figure 2:



Lets understand some aspects of the programme:

Costs

Everyone whose final income is above their market income will be a net recipient from the programme – these are people with final income above the 45 degree line. Everyone with final income below market income will be a net contributor. If the programme is to break-even or if there is some target amount to be spent on the programme then we need to think about the net cost as the difference between the areas above and below the 45 degree line. The net cost depends on the distribution of market income even for a given programme – to keep things simple just imagine that market income is uniformly distributed up to the vertical line in Figure 2. Then the net cost will be area A minus area B so, in this example the programme does not break-even.

I am going to discuss policies that “break-even” but some discussions (e.g. the paper by Blank, Card and Robins on the reading list) talk about a given budget for a programme and designing subject to this ‘budget constraint’.

It is useful to think in terms of the average tax rate defined as:

$$\text{average tax rate} = \frac{\text{market income} - \text{final income}}{\text{market income}} \quad (1)$$

This is a measure of how much an individual gains/loses from the programme. The vertical distance between the programme line and the 45 degree line as a proportion of the height to the 45 degree line is one minus the average tax rate.

Incentives

This is concerned with how much incentive people have to increase their market income. A natural measure of the strength of this incentive is the derivate of final income with respect to market income – this is also one minus the marginal tax rate:


$$1 - \text{marginal tax rate} = \frac{\partial \text{final income}}{\partial \text{market income}} \quad (2)$$

It is not just the marginal tax rate that is likely to affect incentives to work but the average tax rate. The reason is that economists think that leisure is a normal good so an increase in the average tax rate increases work. But increases in the marginal tax rate reduce it.

Poverty reduction


There are many ways to measure poverty and the impact of the programme on poverty but the simplest way is perhaps to think about the level of income provided to those with no market income – this is the height of the programme line on the vertical axis.

Now lets understand the “iron triangle”. The programme as drawn in Figure 2 does not break-even so lets try to make it. The simplest way to do so is to simply lower the programme line until the two areas are equal: we might have something like the situation in Figure 3:



The programme now breaks even and the marginal tax rates faced by everyone are the same. But there is a cost – the programme now provides a lower income for those with zero market income so its effectiveness in reducing poverty has diminished.

So lets try to make the programme of Figure 2 break-even while keeping the point on the vertical axis fixed. We then might end up with something like the change depicted in Figure 4:



This now breaks-even and provides the same level of income for those with no market income but there is a cost – incentives to work have been reduced at the lower levels of market income.

So, both ways of reducing costs have some disadvantages – either we reduce the effectiveness in reducing poverty or we reduce work incentives. This is the “iron triangle” in action. Perhaps you can play around with it and try to come up with some design of programme that avoids the problem – but I think you will struggle.

Differences across Countries


Different countries at different points in time have chosen different ways to try to resolve the problems faced by the iron triangle. Perhaps the most noteworthy difference is between the US and Europe in which, very crudely:

- the US does less distribution so expenditure on ‘welfare’ is lower
- the minimum level of income guaranteed is lower in the US
- there is less effect on incentives to work in the US.

The following picture gives one comparison: between the UK and US for support to lone parents:

FIGURE 7

Net Income as a Fraction of Female Median Earnings for Lone Parents with Two Children, US and UK



Notes: As Figures 2 and 3. Median gross weekly female earnings in 2000 were \$491 and £338 (sources: Bureau of Labour Statistics (2000) and Office for National Statistics (2000) respectively).
Source: Author's calculations.

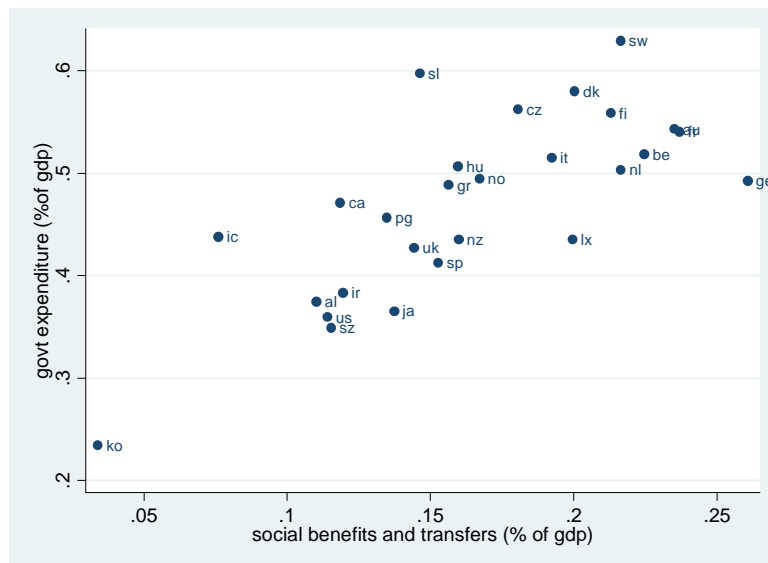
Source: M. Brewer "Comparing In-Work Benefits and the Reward to Work for Families with Children in the US and the UK, Fiscal Studies, 2001.

This is scaled in relation to average earnings as that makes the figure most comparable.

One can see all the points made above in this diagram but often the comparison will be more extreme:

- the UK does not have a particularly generous welfare state by European standards
- the US does 'guarantee' an income for families with children but those without could conceivably get nothing whereas they would get something in the typical European system.

To get some idea of this consider the following graph which shows averages for the period 1992-2003 inclusive



Source: OECD

Note that much of the differences in the share of government in gdp is ‘explained’ by differences in the share of social benefits and transfers - the location of health and education in the private or public sector is the other big source of differences. The rise in the share of government expenditure in gdp in many countries is also largely the result of the rise in social transfers.

The Poverty Trap

One characteristic of the welfare programmes I have drawn so far is that there is typically a high marginal tax rates at low levels of income – the poverty trap. In real-world systems this is often a feature (even though many people seem to think the highest marginal rates of tax are those on higher incomes) and there are two main reasons for it:

First, there is ‘stupidity’. Recipients often get many different types of welfare benefits e.g. for food, housing etc which are withdrawn at different rates as incomes increase. In the past, the sum of the marginal tax rates was sometimes more than 100% making people actually worse off if they earned more money. This type of thing is much rarer now.


The second reason is the cost of improving incentives to work. Suppose we reduce the marginal tax rate at some income level but keep it the same at all other income levels. This is a reduction in the average tax rate for everyone with incomes above that level. So a fall in the marginal tax rate at low incomes is very costly, while at high incomes it is less costly.

Innovation

In recognition of the problems posed by the “iron triangle”, policy-makers have been trying to innovate around it i.e. to find ways in which the trade-offs can be eased using methods which are hard to represent in the diagram we have used so far.

First, the programme may be targeted on specific groups. Typically this is going to reduce the net cost making it more affordable as it reduces the number of recipients. For example, since 1997, the UK has been especially concerned with children in poor households and poor pensioners, in the belief that this is particularly bad (see the Brewer-Clark paper on the reading list). One can then go for a programme with a larger net cost per recipient but a lower number of recipients. The disadvantage is that the programme will do nothing to reduce poverty among those outside the target group.

Secondly, one can make the receipt of benefits conditional on some activity. Increasingly benefits have been targeted on the working poor e.g. through the use of tax credits. These are only paid out if hours are above some level (often 16 or 30 hours per week). This induces a discontinuity in the programme outcome line as shown:



Typically the discontinuity will be at 16 or 30 hours of work, so the position will vary with the hourly wage, being at lower levels of market income, something that is probably a good thing. This is typically less costly than a programme without a discontinuity that wants to provide the same minimum level of income and the same incentive to work 30 hours – this would be represented by the dotted line in the above figure.

Or receipt of benefit might be made contingent on actively seeking work or being ‘job-ready’ having received some kind of training or other assistance. For example, in 1996 the UK replaced Unemployment Benefit with Jobseekers’ Allowance in which tighter checks are made on recipients to make sure they are actively looking for work.

This approach has been quite successful in getting the unemployed off welfare but it has proved harder to reach the sick, partly because the receipt of these benefits is often based on them “being too sick to work”.

Empirical Evidence

Non-experimental Evidence

There are many studies that use non-experimental data to look for evidence on the impact of welfare systems. Many are summarized in the papers on the reading list by Meyer and Blank-Card-Robins. What are the sources of identification used? You will be able to find examples of many of the techniques that have been discussed this term and I am not going to make an exhaustive list.

To estimate the impact of welfare systems one needs variation in the welfare systems. This is hard in many European countries where there is a single national system so the only exploitable variation is over time when it is hard to disentangle the effect of welfare reform from the effect of other aggregate effects.

For this reason much of this research comes from the US and exploits the fact that, as Judge Louis Brandeis said, “it is one of the happy incidents of the federal system that a single courageous state may, if its citizens choose, serve as a laboratory; and try novel social and economic experiments without risk to the rest of the country”. So there are many studies that exploit state variation in welfare programmes using methods like “differences-in-differences”.

Or one exploits the fact that some welfare programmes apply to some sort of people and not others or there are discontinuities. For example, many unemployment-related benefits are time-limited – they expire after a maximum duration. Sometimes this maximum duration is changed. One can then see whether there are differences in job-finding rates around the time of expiration of benefits.

There are good, bad and indifferent studies based on non-experimental evidence. But there are increasingly experiments as well.

Experimental Evidence

Some of the earliest social experiments in economics were conducted around welfare systems. These were the negative Income Tax Experiments of the late 1960s and early 1970s. There were many variants of this but they all replaced a complicated welfare system with a guaranteed income level (quite high by US standards to get people to participate) and a tax rate. These typically found reductions in labour supply. Robins in the *Journal of Human Resources*, 1985 has a good summary.

Since then there have been many experiments conducted. Some of them are reviewed in the Blank-Card-Robins paper on the reading list. I will summarize just one – the Canadian Self-Sufficiency Project (SSP) which operated in the late 1990s.

This was a randomised experiment with a number of different treatments. It targeted individuals who had been on welfare for a year and giving them a 3-year earnings supplement if they found full-time work within a year. This is targeted on a very specific

group (to minimize windfall beneficiaries) and has non-standard features – the ‘clock’, the requirement to work FT. The earnings supplement was generous – if someone worked full-time at the minimum wage then the combination of earnings and supplement would, on average, double their income. There was also a treatment group that received employment-related services in addition to the financial incentive.

SSP was found to be effective in increasing employment rates and annual earnings, and reducing poverty rates among the target groups. But, it did cost money, though arguably represented good value for money.

Although one can get a good estimate of the treatment effect there are some caveats. Notably, entry effects. SSP was targeted on long-term welfare recipients but has the effect of increasing the incentive to remain on welfare a long time. So, does it reduce the incentives for those on welfare 6-12 months to get a job? – this is not clear.

Does this evidence and other evidence from related studies mean the way in which welfare programmes should be designed has been solved? The biggest question mark is about how effective these programmes will be in a recession. In the 1990s in the US, UK (and Canada), arguably anyone who was ‘job-ready’ could get a job relatively easily. But in a bad recession this is not the case and one may see many of the old problems re-emerging.