Beyond 'The Spirit Level'

Diarmid Weir

djgw@blueyonder.co.uk

Introduction – Selling Equality to the Rich

'The Spirit Level' is a book-length distillation of the work of Richard Wilkinson and Kate Pickett on the statistical relationship between standard measures of economic inequality and various social ills, such as ill-health, lack of social trust and crime. Its importance, and the controversy surrounding it, derives from its apparent 'scientific' justification of what might otherwise be instinctive or 'moral' beliefs in the desirability of equality.

The significance of the work is in the prospect it raises of a new dimension to the clash between rich and poor – clear-cut evidence showing that social factors affecting the whole population (rich and poor alike) are beneficially affected by a smaller spread between the highest and lowest incomes. These benefits are come over and above the individual advantage of a higher rather than lower income.

While not necessarily disputing their conclusion, in this essay I point out the potential weaknesses that cast doubt on their particular analysis, and so render it a less than potent political weapon for egalitarians. At the same time I emphasise the importance of tackling overall inequality, not just inequality of income and wealth, for the benefit of the vast majority who would certainly gain.

Inequality as 'Social Pollution'

From the point of view of those with high incomes, if greater equality improves their absolute welfare no intrinsic social conscience is necessary to see this as desirable. This clearly broadens the potential appeal of egalitarian policies that involve the redistribution of incomes and wealth. We can illustrate the argument diagrammatically.

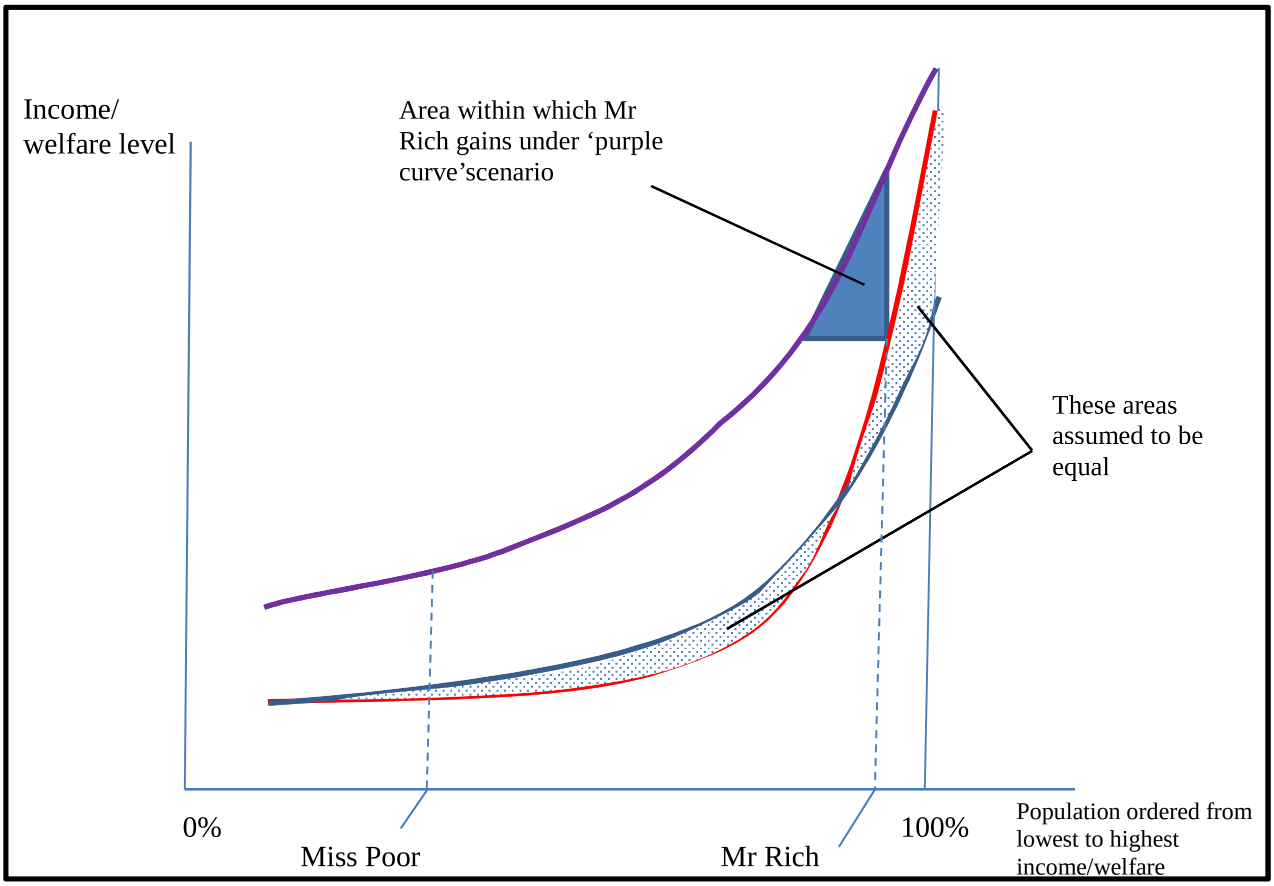

Figure 1 General effects of inequality

The population is ordered along the horizontal axis from poorest to richest, and their individual level of income indicated on the vertical axis. It is assumed that welfare is directly related to income, so their initial distributions can be illustrated by the same curve. Taking any group represented by a segment of the horizontal axis, their total income/welfare can be deduced from the area between that segment and the distribution curve above. The red curve shows an unequal distribution with a much greater share of total income going to the richest. The blue curve shows a more equal distribution of the same total income. Since there is no general benefit from a reduction in inequality, while most (as represented by Miss Poor) are better off, the richest (as represented by Mr Rich) are worse off. The purple curve shows what might happen if there were significant general benefit (to rich and poor) from the same redistribution – the 'Spirit Level' case. In this case there is effectively an increase in total welfare, so that even the richest can be better off despite having a smaller share (Miss Poor and Mr Rich both experience higher levels of welfare).

Figure 1 above shows the share of incomes (and assuming a direct relationship between the level of income and welfare, the latter also) plotted on the vertical axis, against the population ordered from lowest to highest income on the horizontal axis. The red curve shows the income level rising slowly from person to person for those with lower incomes, then much more rapidly for those with higher incomes. The total income is given by the *area* under the curve, and the total income received by any group is indicated by the area under the part of the curve on which that group is situated. Under the red curve the area belonging to the poorer half of the population (on the left of the graph) is much less than that belonging to the richer half (on the right of the graph). This red curve therefore represents a high level of inequality of income, and consequently of welfare.

The blue curve on the diagram represents a more equal distribution of the same total income. Assuming no *direct* effect of the reduction of equality, the welfare levels would *also* be distributed according to the blue curve. While the majority of the population are therefore better off with the blue curve – since their welfare levels lie above what they would have been with the red curve - there are a small number of the rich who are now worse off. This is likely to be a block to motivating political support for inequality reduction, firstly because the economic power of the rich can be translated into political power in disproportion to their numbers, and secondly perhaps because of aspirational beliefs of some of those who would *currently* be better off but calculate according to a positive probability of *one day* being among the rich. The extent to which these beliefs are fostered by organisations and media outlets controlled by the rich serves their current interest in maintaining the red curve.

If however, as Wilkinson and Pickett argue, a lowering of inequality is associated with a general benefit throughout society, in the sense of reducing some sort of ‘social pollution’, then an income redistribution from rich to poor may result in a welfare distribution following not the blue curve, but the purple curve on the diagram. In this case everyone *rich and poor* is to some extent better off. Redistribution of income has improved *everyone’s* welfare, because it has also increased the total quantity of welfare to be shared out. The self-interested rich should rationally support *this* redistribution to the extent that the net result of the redistribution and the overall gain leaves them better off. While they may have a smaller percentage of the total welfare than before, they can be better off because they have a smaller percentage of a *larger* total.

If we take the case of Miss Poor on our diagram, under a shift from the red curve to the blue curve – redistribution but with no general effect of inequality - she gains from getting an enhanced share of *unchanged* total welfare. She will gain *even more* if there is a general effect of decreased inequality so that her post-distribution welfare level is determined by the purple curve. By contrast, under redistribution Mr Rich *loses* from a shift from the red curve to the blue curve, since he is getting a *lower* share of the same welfare total. He will gain, however, from a shift to the purple curve as long as his movement down the *slope* of the purple curve (from his reduced share of the total) is offset by the higher *level* of the purple curve. This gain results for him as long as his post-redistribution position on the diagram remains within the marked blue triangle.

Why 'The Spirit Level' is not a Clincher

Individual Income and Individual Welfare

If Wilkinson and Pickett's interpretation of the data, as demonstrating a generalised effect of inequality that impacts on rich and poor alike, is correct this is an important finding. However it is a subtle one to understand, and its clinching of political arguments for greater equality is not given. For example, those at the top of the income curve may turn out to value their relative position more highly than expected. This may imply a translation from high incomes to high welfare great enough to offset the general effect from any feasible redistribution – at least for the wealthiest (and most powerful).

Moreover, their statistical findings are open to doubt from two important sources. Firstly, there is a statistical peculiarity with much of the effects they elicit. It is not enough to show a correlation between inequality and, say, average life-expectancy across countries and regions. At the least, other plausible explanations for this correlation must be ruled out. One plausible explanation lies in the nature of the relationship between individual income and individual life-expectancy, before the general impact of inequality has been factored in. An individual's life-expectancy is probably subject to diminishing returns in his or her own income – going from a very low to a moderate income has a much greater impact than going from a moderate to a very high income. We can again illustrate this with a diagram.

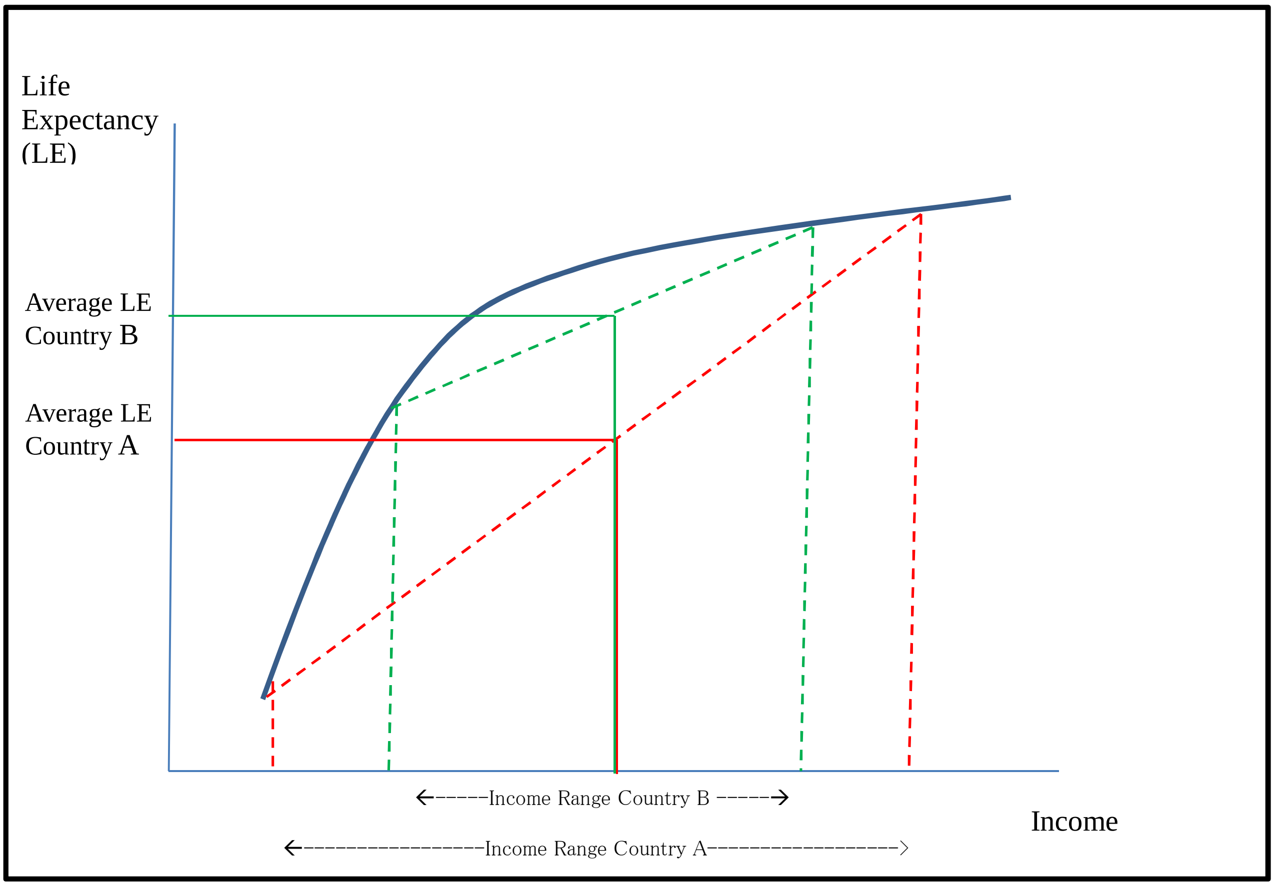

Figure 2 Diminishing returns to income

The diagram assumes the existence of two countries, A and B, with an even spread of incomes across their populations. The solid blue curve shows how individual life expectancy increases with income. For the poor, a small increase results in a big improvement. For the rich the increase per unit of income is much smaller. This means that the narrowing of the range of incomes (from country A to country B) has a much greater effect on the poor than on the rich, resulting in an increase in average life expectancy even whilst the average income remains the same. Yet the same level of individual income produces the same life-expectancy in country A as in country B. If the same income-life-expectancy relationship exists in most countries, then a correlation between inequality (the range of incomes) and average life-expectancy will be observed, despite there being no general causal effect of inequality. Despite the correlation between individual income and individual life-expectancy, countries A and B have very different average life-expectancies despite having the same average per capita incomes. This illustrates that a lack of correlation between average per capita national incomes and average national life-expectancy does not tell us much about the true relationship.

Figure 2 above shows a curve that illustrates the general relationship between an individual's income and his or her own life expectancy. Poor people have lower life-expectancy than rich people even within developed countries. But there is a marked difference as to how an increase in income affects a poor person as opposed to a rich person. Because of the slope of the relationship between income and life-expectancy, an increase in income for a poor person has a much greater beneficial impact on their health than the same increase for a rich person.

Consider two countries, A and B. As indicated on the graph, Country A has a broad range of incomes from very low to very high – it is extremely unequal. (We have assumed for the sake of illustration that this wide range of incomes is evenly spread – so there are approximately the same numbers at the top of the income distribution as at the bottom and at the intervening levels.) The poor, as anticipated, have low life expectancy as a consequence of their low income and so Country A's average life expectancy is as shown. Country B has a narrower spread of incomes – the poorest of Country B are better off than the poorest of Country A, but the richest of country A are less wealthy than the richest of Country B. (These differences between A and B are depicted here as being the same at the top and bottom of the distribution.) Country B's average life-expectancy is as shown, and is considerably higher than that of country A. In this case, increasing the income of the poor from country A to country B *both* reduces inequality and increases average life expectancy. Yet we have not assumed any general beneficial effect of reducing inequality. For every level of individual income, life-expectancy is exactly the same in country A and in country B. The higher average life-expectancy in country B is not a *result* of reducing inequality, but simply of *increasing* the incomes of the poor and the large impact this has on their life-expectancy.

Of course, from a humanitarian or utilitarian viewpoint, whether the increase in average life expectancy is a consequence of reducing *inequality* or reducing *poverty* may seem irrelevant. But the *political* importance is that without a general effect of inequality the rich of country A are better off than the rich of country B – their individual life expectancy is (albeit marginally) actually *reduced* as a consequence of any reduction in inequality that involves a reduction in their incomes. Were a politician in country A to point to country B's income distribution and life expectancy to show the benefits of redistribution to the poor, the rich would be likely to reject such a policy on the basis that they would only see their incomes reduced with no benefit to them. A hard sell indeed. Note the other lesson from this diagram – that even though there is a clear relationship between *individual* income and *individual* life-expectancy, the *average* incomes of country A and country B are more or less identical despite a large difference in *average* life-expectancy. This is another feature of the relationships between incomes and health that make statistical interpretation a trap for the unwary.

Of course, there are many different distribution patterns that can exist – an even spread of incomes is most improbable – but Figure 2 demonstrates that a shift from the red line of Figure 1 to the purple line (also directly benefitting the wealthy) is not so easily distinguished from a shift from red to blue (not directly benefitting the wealthy) as Wilkinson and Pickett's work might suggest (or at least as easy as those reading it might have thought).

Common Factors

The second potential problem for Wilkinson and Pickett is that even when individual income is statistically accounted for, and there is a remaining definite association of average welfare with income inequality over and above this – as is the case for most regional analyses of the United States - there may be some third factor which tends to cause both inequality and the poor outcome (such as low average life expectancy). Particularly for the US, racial divisions are a strong candidate here. In this case, income inequality (as opposed to income itself) and low average life expectancy will be correlated in that they tend to occur together, but there is no necessary direct causation from inequality to low average life expectancy. We would still not therefore expect to see an improvement in the life-expectancy of the better off as a direct consequence of a public policy of redistribution, but only as a result of tackling this third factor.

These complex statistical relationships, and the failure of some of the earlier work on inequality and health to address them, has rather muddied the water on an issue which is inevitably controversial. Wilkinson and Pickett themselves were probably a little naïve not to expect the barrage of (often, but not always, partisan and skewed) criticism that has come their way since 'The Spirit Level's' publication. In particular, given the points made above, it was always going to be easy to attack the rather simplistic use of correlation diagrams in their book.

Studies that have attempted to carefully distinguish between the direct effects of individual poverty and of generalised inequality have produced mixed results depending on the method used and the consequence studied. The current state of opinion is probably in favour of a small but significant correlation of inequality with life-expectancy and selfreported health over and above individual income effects. Certain causes of death, such as homicide, show clearer correlations than others. What remains unclear is exactly how much of this correlation is a result of common additional factors. Race and education may account for some, but not all, of the correlation within the United States.

Clearly similar arguments may apply to several of the other correlates with income inequality identified by Wilkinson and Pickett, such as illegal drug use, poor educational attainment and dropouts, teenage pregnancy rates, imprisonment and levels of trust. In particular, if poverty is correctly viewed as a 'threshold' concept, so that those below the threshold are much more likely to suffer these problems than those above it, then once again the correlation may have little or no direct significance for those at the top of the income scale. The sophistication of the statistical methods required to tease out the causal elements means that clarity and certainty at the level of the average citizen and the mainstream press are probably elusive. Using this particular idea as a central peg to support egalitarian policies is probably unwise.

Prospects for the Majority

Some of the reaction to 'The Spirit Level' from well-funded right-wing thinktanks and maverick social scientists suggests a clear 'debunking' agenda that has not been aimed at clarifying the debate, but obfuscating it. This in itself suggests that to some, high relative status is more important than direct welfare considerations; that even if direct inequality effects do exist the very rich will always believe they can insulate themselves from them; and that there is major ideological blinkering.

The probable reality is that those at the very top of the income and wealth scale cannot be engaged in any reasoned debate about the merits of policies for the 'common good' even when that common good appears likely to include their own. Whatever the merits or otherwise of the discrepancy in resources between the rich and poor *for the rich*, it is most certainly bad for the poor, and almost certainly bad for the rest of us. So the focus should not be on the rich, but on the overwhelming majority of poor to average wealth who will unequivocally gain from a different distribution of resources.

Most are now aware of the disparity of wealth and power that exist – as they become aware of self-seeking behaviour at the top in terms of the press, the banking sector and politicians. At the bottom those on benefits and low wages are squeezed and even demonised, while in the middle even those who felt comfortably off are finding it difficult to cope, to the extent that 1 in 4 families are below what has been calculated by a team from Loughborough University and the Joseph Rowntree Foundation as the minimum income needed to be fully part of modern British society. For the young, even graduates, adequately-paid employment is increasingly scarce.

Of course we do not expect that everyone will rise to the top of the earnings ladder or the top professions (or even that all of our children will). We do all have different abilities and talents (although the extent to which this is innate is exaggerated) formative experiences and luck. And we generally accept that those with some talents in abundance may command higher incomes on the basis of low supply and high demand.

What we generally do not accept is that those with potential who can contribute their talent to society and receive that level of reward appropriate to that contribution are not prevented from doing so through circumstances apart from that potential, such as race, parental circumstances and place of birth. Clearly society benefits from matching individuals to their best role of contribution because this both maximises their value to society and minimises the support/defensive expenditure required to deal with the negative effects to the individual of poor satisfaction/fulfilment of potential/low access to resources.

The latter suggests the issue of 'social mobility' – the extent to which an individual can move according to their 'talent' and independently of other circumstances. There are various measures of this but the basic problem is that equal *opportunity* is pretty meaningless when divorced from concerns about overall inequality. A more equal share of a small probability of advance is little consolation.

Incentivising the Rich

The income/wealth inequality debate therefore continues to hinge on the issue of incentives. It is always possible to argue theoretically that reducing the potential for the rich to earn or accumulate wealth will reduce the effort applied to initiating and/or enhancing economic activity which creates employment, incomes and the provision of new goods and services or the more efficient provision of existing goods and services.

According to this argument limiting the wealth of the richest and transferring it to the poor will not benefit the middle/lower income groups – indeed may even make them worse off. Yet there must surely be some limit to this. It's a morally dubious state of affairs, as the political philosopher Jerry Cohen pointed out. Effectively the rich are blackmailing the rest of us; to continue to work for the common good they demand an ever greater share of total wealth. Given that it's unlikely this additional share is necessary for them to provide this extra effort or to compensate them for the uncongeniality of their work, were this demand made explicitly, face to face, it would be hard to justify. And as Cohen argues '…we do not make policy together if we make it in the light of what some of us do that cannot be justified to others'.

Conclusion – It's the Politics…

Even if the incentive argument is accepted, the evidence is overwhelming that the very rich and their proxies are doing more than simply demanding payment over the odds for their role in generating economic activity. They are now gaming markets for their own profit, manipulating the political process and spreading disinformation through the media outlets they overwhelmingly control. As the Scottish-born Princeton economist Angus Deaton, himself a 'Spirit Level' sceptic, points out,

> The very rich have no need of national health insurance, of disability or income support schemes, of public education, or of public policy that will limit the inheritance of deprivation from parents to children. They do not wish to pay taxes to support such schemes, and their immense wealth and political influence provides them with a potent weapon to prevent them having to do so.

The distribution of political wealth, unlike economic wealth, is zero-sum. If one person gains influence everyone else must lose some. The time has come to accept that there is unlikely to be a 'soft' solution to the problem of widening inequality. Political courage of a high order is required if the accelerating shift of resources, welfare and power from poor and middle-earners to the rich and very rich in both public and private spheres is to be halted and then reversed.