

Hanging On To The Edges:

The worst thing about poverty is not having enough money

If the poor are poor due to bad choices or preferences, then providing them with additional income alone will not necessarily achieve any observable improvements in... outcomes

- Randall Akee and colleagues,
How does household income affect child personality traits and behaviors?

In his 2015 speech to the Conservative party conference, then-Prime Minister David Cameron vowed to use his remaining time in office to mount an all-out assault on poverty in the UK. A worthwhile aspiration, indeed; and not an aspiration we necessarily expect to hear from Mr. Cameron's side of the political spectrum. As it turned out, Mr. Cameron's remaining time in office was not to be very long. In less than a year, he had burnt his wings in the EU referendum and disappeared without trace. I want to talk about an interesting feature of his anti-poverty evangelism, though: central to his planned assault was the idea that poverty was not entirely, perhaps not even mainly, about money.

The intellectual work behind the Cameron approach to poverty was carried out in the preceding years, primarily by the Centre for Social Justice (CSJ) think-tank¹. The CSJ's analysis is, like the curate's egg, good in parts. The CSJ quite rightly stresses that low incomes are correlated with a whole raft of non-income problems. Low-income families are disproportionately likely to be affected by: addiction; alcoholism; family instability; criminality, anti-social behaviour; educational failure; and so on. So there is a manifold of social issues that cluster together, and make life unpleasant or difficult for certain parts of the population. The CSJ rightly argues that if you *just* raised poor people's incomes, whilst making no impact at all on the unequal burden of these other problems, you would not have cracked the problem of social disadvantage in its entirety.

The CSJ then proposes that we *measure* poverty, not just by the amount of money people have, but by a basket of indicators including all these other things like alcoholism, family instability, and so forth. This proposal has, as far as I can see, no merit whatever. It is one thing to acknowledge that poverty is correlated with all kinds of non-income issues. Maybe it is even causally connected to those other issues. But the best way to guarantee that you will never be able to tease out the linkages is to measure poverty in such a way that confuses it with the other issues at the outset. Let me take an example:

¹ See for example their 2012 policy paper 'Rethinking child poverty' (<https://www.centreforsocialjustice.org.uk/library/rethinking-child-poverty>) and their 2013 blog entry 'It's not all money, money, money' (<https://www.centreforsocialjustice.org.uk/csj-blog/its-not-all-money-money-money>).

suppose I am interested in how ocean temperature relates to coral bleaching. Because I feel these things are linked, I could propose to measure ocean temperature by a raft of different indicators including the extent of coral bleaching. The one thing I would now be unable to do is find out whether ocean temperature is related to coral bleaching. I have simply muddled them by assumption; having done so, it becomes impossible to study the relationship between them, because you can't even identify the two phenomena you wish to relate. Thus, whilst I and many social scientists would concur that *well-being* is not just about income, claiming that *poverty* is not just about money is a bit like saying that hyperbolas are not just about a plane intersecting both halves of a double cone. Isn't that, kind of, how you know you are talking about a hyperbola rather than something else?

If we set aside the CSJ's definitional peculiarities, though, we see that there is an interesting idea in there somewhere. Poverty, they say (presumably with the income definition of poverty in mind in this instance), is often 'a symptom of deeper social issues'. What do we mean when we say 'a symptom'? Typically, a symptom is: (a) one of a network of associated phenomena, as in 'symptoms include swelling, fever and rash'; and (b) by implication, not the one you want to go for if you want to causally manipulate the system, as in 'it's best to treat the cause rather than just the symptoms'. So really, the CSJ is making an empirical claim, namely: if you want to lessen the well-being burden due to the inter-related network of poverty, family breakdown, addiction, and so forth, then raising income is not the most effective strategy. Instead, we need to tackle the other nodes directly. Incomes will follow in turn, as better-functioning families get their lives into order and become more economically productive. In fairness to the CSJ, this is hardly a *laissez-faire* recipe for benign neglect of poor people. It gets the government off the hook in terms of the moral case for direct redistribution of cash. But for the government seriously to take on the mantle of responsibility for the family relationships, narcotic consumption, educational attitudes, and normative behaviour of every individual in the land is a mind-blowingly interventionist, not to mention very expensive, aspiration to hold.

The CSJ then, has put out there a big idea. No problem with that. It's just that there is a growing consensus in social science for the opposite view: if you want to deal with the manifold of social problems faced by poor people, both here in the UK and in developing countries, just giving people money is actually a pretty effective strategy. Accepting this opposite view does not come easily to me. I attained my political consciousness in a third world development movement which was pretty much predicated on the aphorism, 'give a man a fish, and he will feed himself for a day; teach a man how to fish, and he will feed himself for a lifetime'. It's hard for me to accept that just giving out fish can possibly be right. I am going to spend the rest of this essay reluctantly conceding that it could be.

§

Let's all admit, for the sake of argument, that low income, family breakdown and addiction are related to one another. I don't just mean that they are correlated. I mean that there are real causal linkages from each of them to both of the other two. Low income increases the likelihood of developing addiction, and of families breaking down; addiction increases the likelihood of family breakdown, and of losing income; and family breakdown increases the likelihood of losing income and of developing an addiction. It's a mutually reinforcing trio of problems: a dynamical system. Now let's say you want to make the world a better place. Where would you do best to put your dollar? You can choose between directly raising incomes; providing addiction treatment programmes; and providing family counselling.

One of the things you need to consider is the magnitude of the effect of changing one of the variables on each of the other two. For example, if you can reduce family breakdown, to what extent do income and addiction then improve? The CSJ hypothesis is, in effect, that the knock-on impact of reducing

family breakdown or addiction for income would be rather large, but the effect of raising income on the other two problems would be small. Perhaps it would even be zero, or negative, as previously poor people went out and frittered away their newly-acquired cash on social bads like drugs. *They wouldn't know how to use the money sensibly.* So naturally, the CSJ concludes that raising incomes alone is not the best approach.

The examples we are going to see suggests that they have it the wrong way around. Raising the incomes of poor people, even absent any other changes, can have a surprisingly large positive impact on all kinds of social and behavioural problems, and hence well-being. It does not eliminate all social problems, of course: nothing we know of does that. Nonetheless, it can do a lot to reduce the non-income wellbeing disparity between rich people and poor people, as well as, more obviously, the income disparity. It makes sense that, other things being equal, raising incomes is likely to be the most effective way of perturbing the dynamical system of social and family problems. That's because giving people cash is remarkably *efficient*, especially if you do it in some fairly non-bureaucratic way. There's a few cents in the dollar for administration and banking charges, but beyond that, the more money you transfer to poor people, the more their incomes go up. The efficiencies of family counselling and drug treatment programmes are likely to be much lower. I am not saying these initiatives don't work at all; I am sure they do. But you have to recruit and train up counsellors and staff. These people are typically much more middle-class than the people we are trying to target. They need decent compensation packages, and that costs a lot, typically much more than a poor family earns. For overseas development, they need to be flown in and housed. Then they have got to access the populations with the need. And even assuming they manage to do this, their help only has a certain degree of success; plenty of families go through family counselling and still break up anyway; plenty of addicts receive treatment but don't escape their addiction. So it would probably be fair to speculate that the efficiencies of non-income forms of aid directed at poor people are typically lower than that of direct income support.

§

The same people and places tend to have the lowest incomes, the poorest physical and mental health, the most crime, the lowest trust, more behaviour problems, and so on. However, this does not in itself help you decide on the best remedy for poverty. Both the CSJ and the cash-first hypotheses are consistent with there being a manifold of positive correlations of all the different kinds of life-crapness. If you want to get anywhere in adjudicating between the two hypotheses, you need something like the scientific experiment. In an experiment, you hold everything else constant, and perturb one variable (for example, income) in the absence of any other change. Then you see what effects follow on the outcomes that interest you. Hold on, you say, that's all very well. But social scientists can't do experiments. People's incomes never change without their education, culture, or other aspects of their behaviour changing too, in uncontrolled ways. Social, political and economic life just go on, and we social researchers are limited to documenting them and interpreting their fluxes.

The situation is actually not quite as bad as this. Sometimes one factor does get changed, pretty much independently of all the others, and for reasons that are largely exogenous to the system. Social scientists spend a great deal of time studying these situations, and the results come us close to a decomposition of causality as you could reasonably hope for. The gold standard situation is the *randomised control trial*, the true scientific experiment applied to a social policy innovation. More and more of these are now done. But even where randomised control trials have not yet proven possible, there are nearly-as-good sources of causal inference: *natural experiments* or *quasi-experiments*. These are situations where some change occurs that is outside the researcher's control (this is how it differs from a true experiment), but nonetheless alters just the variable of interest, and just for some people

but not for some other, comparable ones. When a social policy is introduced into one jurisdiction but not a similar neighbouring one, then as long as the reason for the introduction happening where it did is not reducible to any existing characteristic of the jurisdictions, then you have a kind of natural experiment. And social policy changes sometimes happen for the strangest and most random of reasons.

My favourite quasi-experiment comes from the Great Smoky Mountains study. This began as a fairly run-of-the-mill longitudinal study of psychiatric problems, addictions and problem behaviours amongst young people in parts of Western North Carolina, beginning in 1993, and continuing as the young people grew into adults. But it became something far from run-of-the-mill in 1996. A fair proportion of the participants were Native Americans from the Eastern Band of Cherokee. In 1996, a casino was opened on their reservation land (Native American reservations are outside state gaming laws). Some of the profits were put back into the Band community, and the mechanism chosen for doing this was basically a Universal Basic Income: all adult Band members received an equal portion, in the form of semi-annual cash payments, for which they did not have to do anything other than be themselves. Small at first, these payments had risen to \$9000 per person per year by 2006, enough to very substantially raise household incomes in that part of the world. And for Eastern Cherokee youth, there was a large lump sum to be held in trust and collected on their 18th birthdays.

It's important to appreciate that, before the payments began, the Eastern Cherokee had the usual poverty smorgasbord: as well as their incomes generally being low, there were lots of problems of addiction, anti-social behaviour, and family strife. It was classic CSJ stuff. And if the CSJ hypothesis were right, then the cash payments, which after all did nothing at all but lodge a cheque, would not have helped with all these other 'symptoms of something deeper'. Things could have even got worse. Suddenly having cash in the bank, and lacking the family stability and life skills to know what to do with it, you might have expected the newly cashed-up young people to drop out of school (who needs to work when you are given money for nothing?), and turn to drink, drugs and gambling. Nothing could be further from the truth.

There are several good studies of what happened to the Eastern Band of the Cherokee, so here I will focus a few of the most noteworthy. Elizabeth Jane Costello and colleagues systematically compared young men and women from Cherokee families with non-Cherokee of the same age from the Great Smoky Mountains cohort². These non-Cherokee were effectively the control group. Not a very good control, you might say, since the non-Cherokee were bound to differ from the Cherokee in many non-income ways. However, the researchers could turn here to the fact that they had data from Cherokee of different age cohorts. The oldest cohort had benefited rather little from the casino scheme—the lump sum payable at 18 only started to cumulate in 1996, so those turning 18 in 1998 got only a very modest amount, and had not benefited from increased parental income for very long either. So the differences between the oldest cohort of Cherokee and oldest cohort of non-Cherokee tells you something about the *status quo ante casino*. By contrast, the youngest Cherokee, turning 18 in 2002, received \$35,000 on their birthday, besides which their parents had had quite large sums coming in for all of their teenage years. So if cash does anything good for non-income outcomes, you should see the youngest cohort of Cherokee doing better relative to their non-Cherokee peers than earlier cohorts of Cherokee had done. This is a variant of what is called a 'difference in differences' study design, because any causal impact of the money is going to change the differences between Cherokee

² Costello, E. J. et al. (2010). Association of family income supplements in adolescence with development of psychiatric and substance use disorders in adulthood among an American Indian population. *Journal of the American Medical Association* 303: 1954-9.

and non-Cherokee outcomes between the oldest cohort (not much casino cash), and the youngest cohort (lots of casino cash).

And the differences were indeed different. Looking at the oldest cohort, by the time of study, 41% of Cherokee had experienced some kind of psychiatric disorder, against 31% of non-Cherokee. Much of this was made up of or included some kind of substance dependence (35% of Cherokee, against 29% of non-Cherokee). The rates of diagnosed 'behavioural disorder' (which is often a catchall for minor criminality and anti-social behaviour) were five times higher in the Cherokee than the non-Cherokee. But remember these were the Cherokee cohort who had benefited only marginally from the coming of the casino. In the youngest cohort, who had benefited very substantially from casino money, not only had the Cherokee caught up with their non-Cherokee brethren, but they had surpassed them. The differences were all in the opposite direction: any psychiatric disorder: 31% Cherokee versus 37% non-Cherokee; substance dependence: 23% Cherokee against 35% non-Cherokee; behavioural disorders three times higher in the *non*-Cherokee than the Cherokee.

In related work, Randall Akee and colleagues looked at involvement in criminal activity, and at school performance, whilst the members of the study were still minors³. Again by comparing those who received different amounts of casino transfer and those who received none, they were able to estimate that an additional \$4000 per year of unearned income per year reduced the likelihood of ever getting involved in minor crime by 22% (for a 16 or 17 year old); and, moreover, that it increased the average amount of formal schooling completed by a whole year. Prior to the casino, Cherokee youth had worse rates of minor criminality and lower rates of high school completion than non-Cherokee youth. Over the first few years of the casino, they not only closed the gap, but gone beyond: now they were *more* likely to finish high school, and *less* likely to commit minor crime, than non-Cherokee youth in the area.

Akee and colleagues were able to do two other important things. First, rather ingeniously, they established that what mattered for the beneficial effect of the casino scheme on a household was not how far it was geographically from the casino, which might have been the case had the mechanism for the behavioural changes been, say, meeting lots of morally improving outside role models who had come to the area to use the casino facilities. (No, I don't think that's very plausible either, but the good thing about science is that you can try to test these possibilities against the data). No, what mattered for the beneficial effect was just *how much money* came in to the household (you got more over time, remember, and the family got more the more registered Cherokee persons there were living in the household). The other thing the researchers were able to show was that a big part of the beneficial effect operated through creating better relationships within the household. The parents did not work any fewer hours as the free money increased (there's one for the Universal Basic Income advocates). As their financial situations improved, though, they reported higher quality relationships with one another and with their children. And the harmony was not achieved by trading in their feckless spouses for new models, either: they just got on better with whoever they were already with. This makes sense: put people under less strain, and it's easier for them to get along well. And a big way of taking the strain off is through the pocket book.

Akee, Costello and colleagues have one further set of results worth highlighting. They recently delved back into the questionnaires and evaluations supplied by the parents and, for some variables, the

³ Akee, R.K.Q. et al. (2010). Parents' income and children's outcomes: A quasi-experiment. *American Economic Journal: Applied Economics* 2: 86-115.

children, of Cherokee and non-Cherokee families⁴. They found, confirming previous analyses, that receiving the cash payments reduced symptoms of emotional disorders (basically anxiety and depression), and of behavioural disorders (basically being antisocial). Moreover, the researchers measured three of the 'Big Five' personality traits, Agreeableness, Conscientiousness, and Neuroticism. The former two are important for how you get on in life: Agreeableness describes the tendency to be cooperative and get along well with others, whilst Conscientiousness describes the propensity to be hard-working and organized. The classic successful bourgeois is pretty Agreeable, and highly Conscientious. The prisons are full of people who rate low on both traits. And guess what: the arrival of income payments was associated with large increases in Agreeableness and Conscientiousness amongst young Cherokee. No comparable personality changes over time were seen in the non-Cherokee members of the study, who were living through the same general social period, but not getting the income increase.

In summary, then, the Eastern Cherokee casino income was 'helicopter money', a large increase in income that descends from on high with no skills training, no family counselling, no conditionality, and no prior logic. The CSJ hypothesis predicts that its arrival shouldn't have solved the whole network of social problems. But by looking at how the Eastern Cherokee compare in social outcomes to their non-Cherokee neighbours both before the money arrives, and after, we can make pretty clean causal inferences about what raising incomes does. The evidence tells us unambiguously: relationships in families improve, kids stay on in school, kids become less likely to get involved in minor crime and antisocial behaviour, addiction goes down, and even, most remarkable of all, *people's personalities change*. Not bad for treating the symptoms rather than the cause, eh?

§

Our second example comes from the developing world, from Kenya's Cash Transfer for Orphans and Vulnerable Children programme. This programme was a response to the fact that, due among other things to the AIDS epidemic, Kenya had a huge number of young people whose parents had died or were dying, and these young people needed supporting. Kenya could have spent its money in various ways: skills-training programmes, counselling, orphanages, and so on. It chose another path, a completely unconditional regular cash payment to the household in which the young orphan was living: helicopter money. So again, we have a nice clean test of whether cash alone does much for poor people with a manifold of different social problems. And better still, we have a proper randomised control trial of the programme. It was impossible to roll the programme out simultaneously in all of Kenya. Thus, districts were randomised to receive the programme immediately (the experimental group), or in a later wave (these districts served as the control group at the time of the evaluation, when their orphans had not yet received anything).

Comparing the two groups showed that the cash transfer improved school attendance rates, particularly when school was costly to attend (e.g. when the school was far away); and particularly when the children were older (which is the time when there is an opportunity cost of going to school, instead of generating money directly or looking after the household)⁵. This is an important finding given that there was no conditionality in the programme whatever: the cash would continue to appear regardless of whether the young person stuck with school or not.

⁴ Akee, R.K.Q. et al. (2016). How does household income affect child personality traits and behaviors? *NBER Working Paper No. 21562*.

⁵ Kenya CT-OVC Evaluation Team (2012). The impact of Kenya's Cash Transfer for Orphans and Vulnerable Children on human capital. *Journal of Development Effectiveness* 4: 37-41.

What I want to focus on, though, is a nice, revealing study of household expenditure in the control and the experimental groups, before and after the programme⁶. Experimental households spent more money on nearly everything once the cash started to roll in. This should not surprise us: after all, they had more money to spend. Here the researchers were able to do some clever econometric stuff, though. They used the pre-intervention spending data to construct a model of how household expenditures of different types in this population scale with income. This allowed them to make predictions: if this household behaved like a typical household in this population, then when you increase its income by 1500 shillings a month, how much more should we expect it to spend on food, how much more on healthcare, how much more on alcohol, etc.? The actual observed changes when the programme kicked in could then be compared to these predictions. In effect, the researchers asked, does the programme allow households to satisfy more completely the priorities they had anyway, or does it change their priorities?

When they performed this comparison, the researchers found that a number of categories of expenditure went up by less than expected, for example food. Within foods, expenditures on cheap tubers went down relative to expectation; it was only spending on high-quality foods that went up. Expenditures on alcohol and tobacco actually went down. Expenditure on healthcare went up relative to expectation, and households also saved and invested more than when the programme began. In short, as the cash landed, households shifted their preferences away from hedonic gratification (alcohol and tobacco) and immediate subsistence (tubers), and towards looking after their long-term health, and making investments. This is the riposte to the CSJ 'won't it be bad to just give poor people money when they don't have the skills to know how to spend it wisely' type of argument. They do seem to spend it wisely. Perhaps they are smart, and can figure out how to do so for themselves. Perhaps they are as smart as you, me or a development expert, but have had worse luck until now.

Both the Kenya example and the Cherokee one bring to light a very interesting conundrum: when you give people more money, their expenditure on narcotic substances goes down. Given abundant resources, it seems, most people don't value these things very highly. But, if people don't value them very highly, then why, when money was short (i.e. before the cash helicopter landed), were they spending anything on them at all in the first place? If something has a low value, then surely it would get crowded out when money was tight, and only perhaps creep in when the money supply gets looser? This conundrum really gets to the heart of the matter. The CSJ view looks at the behaviours of poor people, such as their proneness to use narcotics, and sees a disposition. It then says: shovelling cash on to this disposition won't do any good, and may even do harm. It's the disposition, stupid. The cash-first view, on the contrary, looks at the behaviours of poor people and says: that's a response to a situation. Change the situation (add money), and all kinds of decisions will follow suit.

This brings to mind classic animal research on addictive substances. Rats or mice living alone in small barren cages will self-administer morphine or cocaine enthusiastically, if given the chance. It turns out that the very same animals living in spacious, enriched stimulating environments, will do so significantly less, even when the drug is easily available⁷. In other words, the motivation to use

⁶ Kenya CT-OVC Evaluation Team (2012). The impact of the Kenya Cash Transfer Program for Orphans and Vulnerable Children on household spending. *Journal of Development Effectiveness* 4: 9-37.

⁷ Alexander, B.K., Coombs, R.B., and Hadaway, P.F. (1978). The effect of housing and gender on morphine self-administration in rats. *Psychopharmacology* 58: 175-179; Chauvet, C. et al. (2009). Environmental enrichment reduces cocaine seeking and reinstatement induced by cues and stress but not by cocaine. *Neuropsychopharmacology* 34: 2767-2778; Solinas, M et al. (2009). Reversal of cocaine addiction by environmental enrichment. *Neuropsychopharmacology* 34:1102-11.

rewarding narcotics is not a biologically inflexible drive in these creatures; it's a way of coping with adverse environmental contexts, the lack of alternative sources of reward, and it spontaneously though not completely fades away as those contexts improve.

This doesn't completely deal with the argument, in the human case, that goes as follows: why can't poor people just spend less on alcohol and tobacco in the first place? If they did so, it would be as if they were giving themselves a cash transfer programme. They could start to climb the ladder towards a better life using the money they saved, without having to wait for a casino to come along. This is a hard argument, but I expect the answer might be along the following lines. When you are poor, you occasionally do find yourself with a little left over, but this surplus from the requirements of subsistence is small and temporally unpredictable. Thus, you can't plan to use it in any calculated escape route from poverty. When it does come, quite understandably, you treat it as a break in the weather, a small moment to medicate yourself from the awfulness and difficulty of life. But what poor people do when they get a small and unpredictable surplus in an otherwise bleak existence is not indicative of what they would do if we gave them a large and predictable permanent surplus. Both Kenya and the Cherokee case show us this. When you give poor people large, reliable, long-term surpluses, they start to behave just like people who have been lucky enough to have large, reliable, long-term surpluses in the first place. They don't need teaching; they don't need conditionality or monitoring. *They just need the money.*

§

The evidence we have reviewed has implications for two sets of things I care about: politics, and our view of human nature. First, the politics. Many cynical commentators suspected that the CSJ's de-emphasis of the purely financial aspects of poverty was a smokescreen for regressive policies. After all, if the bad thing about poverty is not the lack of money, then there is no compelling case for income redistribution. The proposal to measure poverty using non-income indicators makes it much easier to enact financially regressive policies and not have anyone notice that is what you are doing. In support of this point, the tax and benefit changes enacted by the Cameron government and its successor between 2010 and now have clearly been regressive: the households in the lower deciles of the income distribution have seen their incomes eroded much more sharply (even in absolute terms, never mind proportional ones) than those at the top⁸. So this is what the promised assault on poverty actually looked like: redefine poverty as being not about money, and then take money away from poor people. There is, I suppose, a kind of logic to it.

In my view, the political implications of cases like the Eastern Band of Cherokee are clear: we need to redistribute income and wealth more, from the people at the top of the distribution to whom it brings little benefit, to the people at the bottom for whom, as we have seen, it makes a profound difference⁹. So impressive are the Cherokee results that redistribution becomes not just a moral obligation, but a simple matter of pragmatics. A thought experiment: Say I tell you I've developed an intervention programme that costs a bit, but has been shown to reduce crime, increase educational attainment, reduce drug and alcohol dependency, promote family stability, and make people nicer and more conscientious. That sounds great you say, imagining perhaps some combination of turmeric and mindfulness-based cognitive therapy, that intervention should be funded nationally. Now I tell you what my intervention is: it's called making sure people aren't poor. You might suddenly feel less keen,

⁸ Equality and Human Rights Commission (2018). The cumulative impact of tax and welfare reforms. Downloaded from www.equalityhumanrights.com.

⁹ See *Why inequality is bad*.

even if I tell you what could well be true: it's cheaper and more effective than the alternatives in the long run.

Now, what about our view of human nature? The evidence from cash transfers strikes quite strongly against any view in which individual differences in behaviour are the result of some fixed inner essence that is obtained early in life and inflexible thereafter (disposition, taste, personality, culture, or whatever you want to call it). If such fixed-essence views were correct, then the CSJ would be basically right: helicoptering cash to people would just lead them to perpetuate whatever they are doing already, but under looser constraints. The alternative to the fixed-essence view places situation and context more centrally as drivers of people's immediate behaviour; sees people as highly plastic in their tastes, strategies and decisions; and emphasises that a big part of where we end up in life is due to (reversible) luck.

One influential version of the fixed-essence view has genetic differences doing a lot of the work. You might think that your personality, for example, is largely due to your genetic inheritance¹⁰. If your lineage is disagreeable and not conscientious, it drifts down the social ladder. If you are lucky enough to have agreeable and conscientious genes, you climb up. But this cannot be the whole story: One of the reasons I love the Cherokee example so much is that it shows personality, that most fixed and genetic-feeling of things, changes in response to helicopter money. So what is going on? Are psychologists wrong about personality being fixed and heritable?

They are not entirely wrong, but some nuance is needed. There seem to be heritable influences on personality, as evidenced from studies of twins; and two people facing the same environment of poverty can respond to it quite differently, which might well have something to do with dispositional differences between them. But it is too great a leap to move from 'genetic differences explain some of the variation in personality between individuals facing similar environments' to 'group differences in personality are best explained by variation in genes between those groups'.

Let me propose the following analogy. Imagine you grow corn on a field uniformly rich in fertilizer. All of your corn plants will be tall, but some of them will be a little taller than others. The differences in height between your plants will probably be mainly due to genetic variation between them (after all, they all developed in the same benign environment). If you did a study of your plants at this point, you would conclude that height is highly heritable, mainly a matter of genetics. Now you take away the fertilizer from half the field. The plants in that half of the field grow much less tall in the next year. This is entirely caused by the change in inputs. So although you had concluded that the individual differences in plant height were heritable when the environment was good, you have also proved that a group difference (between one half of the field and the other) has nothing to do with genetics, and everything to do with environmental factors. And so, I think, with people: there may well be genetic variations, but where we really see their importance is in explaining the residual variation given a constant environmental context. The environmental context for different social groups is nowhere near constant, though, and that's a much more relevant explanatory principle for differences in average level across groups.

Another form of fixed-essence thinking puts culture in the place of genes. The argument here is that culture is a pseudo-genetic inheritance system: you absorb a system of behaviours via social learning in your childhood, and thereafter you are pretty much stuck with it. Social change, when it happens, is a matter of cultural mutations that gradually change in frequency over the course of generations,

¹⁰ The evidence is reviewed in Nettle, D. (2007). *Personality: What Makes You The Way You Are*. Oxford: Oxford University Press.

faster than genetic change, perhaps, but still slower than the individual lifetime¹¹. This kind of thinking can't really explain the Cherokee case. The cash changed behaviour patterns massively within a single cultural generation, and without (as far as we know) changing who learned their culture from whom. Whatever people were doing, it was not internalising and persisting with the behaviours they had been exposed to in childhood. Instead, you have to see people as strategic agents who change their decisions and dispositions more or less in real time as their environments (and their information about their environments) change. This does not preclude roles for norms and social transmission in an account of human behaviour; but it does warn us against pushing the analogy between cultural and genetic inheritance too far¹².

§

In spite of everything, I still find it hard to accept that the best thing I could do to help poor people is just to give them my money. I know that many other people feel the same. At the bottom of this, I think, is some kind of *illusion of the validity of our expertise*. The idea of an illusion of validity comes from classic work by Daniel Kahneman and Amos Tversky¹³. Kahneman served in the Psychology branch of the Israeli army. This branch developed expert methods for predicting who would make a good officer. Kahneman followed up and found that the expert methods for officer selection were, in fact, quite useless: you might just as well have chosen every third recruit and stuck a badge on them. Nonetheless, the Psychology-branch experts remained convinced that their expertise was valid, and continued self-importantly to deploy it. As well as over-valuing our own expertise, I think we are all prone to under-value the expertise of people we consider unlike us, in this instance, poor people.

We seem to feel sure that we have a good analysis of how poor people could make their lives better, so sure that we are not shy of coming up with advice, diagnoses, intervention strategies, and training programmes. On average, these probably do less good than just transferring the equivalent amount of money to low-income households. That's hard to accept, especially for a professional academic whose day job is having something expert-sounding to say. It's hard to accept because conceding the value of just transferring money is tantamount to admitting that my expertise in how to fix things is low. I might think I am brilliant, but in truth I would probably be really bad at living on the bread-line: I have not developed the skills. Poor people, on the other hand, are generally going to be more expert at coping with that context. Therefore, on average, they are probably going to make better decisions about how to navigate the shoals than I am. Cash transfer takes micro-allocation decisions out of the hands of people who don't really know what they are doing (like me), and into the hands of people with expertise (the recipients). Accepting the case for cash transfer then, is really about accepting that poor people are in general cognitively equivalent to rich people, but on average more skilled; and therefore trusting them to make their own decisions. To do this requires letting go of the intuitions that give us paternalism and the idea that poor people are deficient in decision-making capacity. And those intuitions, I suspect, are more deeply embedded even amongst progressive academics than it is comfortable to admit.

¹¹ This view is particularly clearly articulated by Mathew, S. and C. Perreault (2015). Behavioural variation in 172 small-scale societies indicates that social learning is the main mode of human adaptation. *Proceedings of the Royal Society, B*. 282: 20150061.

¹² See *The cultural and the genetic*.

¹³ See Kahnemann, D. and A. Tversky (1973). On the psychology of prediction. *Psychological Review* 80: 237-51.