Davenport, Alexander; Levell, Peter; Sturrock, David

Research Report

Why do wealthy parents have wealthy children?

IFS Report, No. R196

Supplementary appendices to ‘Why do wealthy parents have wealthy children?’

Appendix A. Definition of wealth measures

Detailed information on wealth is present in BHPS-USoc waves 5, 10, 15, 22 and 26 which correspond predominantly to the years 1995, 2000, 2005, 2012 and 2016.

We build up wealth from information on house values, outstanding mortgages, savings and investment assets, and unsecured debts. The BHPS-USoc does not contain sufficient information for us to accurately gauge pension wealth, and so this is left out of our wealth measure.

Housing wealth

We take reported house values for homeowners and divide this wealth equally between those who are reported as owning the home. For wealth measures in waves 5, 10 and 15, the same is done for the value of second homes. All of our parent wealth observations come from waves 10 and 15 and so account for second homes. Most child observations are from waves where second homes are not observed, but the prevalence of second homes is likely to be low at the ages at which children are observed.

Mortgages

Mortgages are calculated in the same way as housing wealth, i.e. split equally between homeowners of the house the mortgage is covering. Where these are missing in some cases in waves 22 and 26, the value of the mortgage in a different wave is used and the recorded interest rate used to impute the value of the mortgage. Where this interest rate is also not present and the mortgage is interest only, the interest rate is imputed from the growth rate of mortgage debt in other waves.

Financial assets

Financial assets are comprised of savings, investments and unsecured debts.

For savings and investment values, the data are structured differently in waves 5, 10 and 15 from in waves 22 and 26.

For investments in waves 22 and 26, we take the sum of investments in National Savings certificates/bonds, unit/investment trusts, stocks and shares, and other investments such as gilts or government/company. For investments in waves 5, 10 and 15, we take the sum of investments in the above categories plus premium bonds and PEPs, but individual amounts for these different investment categories are not split out.

For savings in waves 22 and 26, we take the sum of amounts held in standard savings/deposit accounts, National Savings accounts, cash ISAs, stocks and shares ISAs, premium bonds and other types of savings accounts. For savings in waves 5, 10 and 15, we take the sum of savings accounts, National Savings bank accounts and TESSAs/ISAs, but again individual amounts for individual investment categories are not split out.

For both savings and investments in all waves, if respondents indicate they have assets in a given class, they are asked how much these are worth. If they do not give an exact response, they are asked a series of bands, and we assign the mid-points of these bands. For example, if an individual says their amount held in cash ISAs is above £500 but under £1,000, we assign them a value of £750 for this asset class. Those with asset values above the top banded response have their amounts imputed from those giving exact responses above the top band amount in the same wave. We impute based on a linear regression of asset size on age, age squared and education, drawing a residual from the estimation sample at random when imputing.

For non-mortgage debts in all waves, we take the sum of hire purchase agreement debt, personal bank loans, catalogue order debt, Social Fund loans, loans from private individuals, overdrafts, student loans, credit card debts and other debts.

Some of these asset classes can be held in a sole name or jointly. Respondents are asked whether accounts are held in sole name, jointly or both. If they respond 'jointly', the amounts are split among those who hold the assets. If they respond both, they are asked how much is held in their sole name and we split the remainder by the number of joint owners. In some situations, the two partners give inconsistent responses on the amounts held jointly in various accounts; in these situations, we take the values reported by the individual of interest.

For waves 22 and 26, risky assets are the sum of the amounts for unit/investment trusts, stocks and shares, other investments, and stocks and shares ISAs. For waves 5, 10 and 15, since investments and savings are not split out, investments are classed as risky assets and savings as safe assets.

For waves 22 and 26, individuals are assigned a missing value for investments if more than two of the four classes of investments elicit a non-response, and they are assigned a missing value for savings if more than two of the six classes of savings elicit a non-response.

Total financial wealth is then the sum of savings and investments net of non-mortgage debts. However, we allow one of these categories of financial wealth to be missing, but if more than one is missing then financial wealth for the individual is listed as missing too.

Total wealth

Total wealth is the sum of housing wealth and financial asset wealth minus mortgage values and unsecured debt values. We also calculate net housing wealth as housing wealth minus mortgage values, and net financial wealth as savings and investment values minus unsecured debt. For parents, we add up wealth values across both parents and then take the average of this for waves 10 and 15, corresponding to 2000 and 2005. At this point, most of our parents are between 40 and 60 years old.

Earnings

Earnings are calculated as the sum of income from an individual's primary job and, where it exists, secondary job. For the children in our sample, we take the average income from the five waves up to and including the wave in which they most recently appear up to wave 26. For the parents in our sample, we create a panel with an observation for each wave above the age of 20 but before retirement. We then add up total parental income for a child within each wave. We then recover the parents' permanent income (in total for the child) using a fixed effects regression with individual fixed effects, controlling for age, age squared and interview year.

Saving

Individuals are asked whether or not they save in each wave and we record this for children. For parents, we note whether either of the parents saves in a given wave. In every other wave, individuals are also asked how much they save. For parental saving rates, we take the average saving amount for a parent in the years they are not retired but older than 20, and divide it by their average earnings from this period. We then average across the mother's and father's saving rates to get a parental saving rate for each child. For child saving rates, we divide saving by earnings in each wave where these are present and take the average of the five most recent values present for these.

Price index

We deflate all asset and earnings values to 2019 prices using the UK Consumer Prices Index.

Appendix B. Sample selection and parent–child linkages

Sample selection

In order to ensure children are present before turning 17, but can also reach 30 by our final wealth wave in 2016, we require children to be born between 1974 and 1986 to be in our sample. As can be seen in Table B.1, there are 3,809 members of the survey born between these years and observed at age 16 or below. Of these, 3,005 are ever present in one of the wealth module waves, and 813 of these are present in our survey at or above the age of 30.

Table B.1. Sample size progression under imposition of conditions for inclusion in sample

Condition imposed	Number not meeting this condition	Sample size after imposing this condition
Born between 1974 and 1986	118,969	27,560
Present aged 16 or before	23,751	3,809
Present in wealth waves	804	3,005
Reach age 30	2,192	813
Have ‘complete’ wealth information	103	710
Ever have ‘complete’ parental wealth	12	698
Have parental wealth in waves 10 and 15	64	634
Have ‘consistent’ parents	40	594

In order to get a good picture of an individual's wealth position, we require wealth observations to be non-missing for their house value, mortgage value and the value of their savings/investments. 103 of our sample do not meet this condition so are dropped, leaving us with 710 children.

Parent–child linkages

To have appeared in our sample aged 16 or younger, our sample members will be linked to at least one parent of some form, whether this is their biological parent or another 'parent' figure who could be a step-parent, grandparent or adopted parent. Since we are interested in the relationship between parent and child wealth, we require a child to have non-missing wealth information for the housing, mortgage and savings/investment values for at least one linked parent, and 12 members of our sample do not meet this condition, leaving us with 698 children. We additionally require this complete wealth information to be present for at least one parent in waves 10 and 15 so that we can take the mean of parental wealth over a set of waves that is consistent over all observations. 64 members of our sample do not meet this condition, leaving us with 634 children.

The survey follows only those who were in the household when it was originally sampled plus the children of those in the household when it was originally sampled. Therefore, if a parent who was not in the household when originally sampled (e.g. a step-parent not present when the household was sampled) was to subsequently leave the household, that parent would drop out of the survey. It is also possible that there will be attrition from the survey when one member of a couple leaves the household or dies. In the case where one of a child's parents is missing because of a parent not being followed by the survey design, or because of attrition, that child's parental wealth as observed in the survey will be incorrect.

In some cases, the remaining parent re-partners and so a (new) step-parent (where 'step-parent' could refer to grandparents or adopted parents as well as just a new partner of an existing parent) is recorded. In the cases where there is a new step-parent, we use the new step-parent's wealth. We only classify a new step-parent as valid to 'replace' a biological parent if they are present in the child's household for three or more years before the child turns 16. Where there is more than one such parent, we take the step-parent present for the largest number of waves; in the case of a tie, we go with whichever appears first. In cases where the original parent disappeared due to death, we take the surviving parent's wealth as accurately reflecting parental wealth. However, where a parent who was previously present leaves the survey before wave 10 and is not replaced (‘Parent present before wave 10, but missing in wave 10 or 15 and not dead’ in Table B.2), or where the parents used in waves 10 and 15 are not the same (‘Parents are not the same across waves 10 and 15, not due to death’ in Table B.2), we drop the child. Table B.2 shows the distribution of observations by parents present and the reason for parents being missing or changing across waves.

Table B.2. Distribution of observations by parents present and reason for parents being missing or changing across waves

Situation of parent	Mothers	Fathers
Biological parent present in waves 10 and 15	598	469
Same step-parent present in waves 10 and 15	10	28
Parent was present before wave 10 but died before wave 10	3	16
Parent died between waves 10 and 15	2	3
No parent ever present during survey	11	88
Total: ‘consistent’ parents	624	604
Parent present before wave 10, but missing in wave 10 or 15 and not dead	6	16
Parents are not the same across waves 10 and 15, not due to death	4	14
Total: ‘consistent’ and ‘inconsistent’ parents	634	634

The result of this is that 40 observations are dropped due to having inconsistent parents, 30 for inconsistent fathers and 10 for inconsistent mothers. 22 of these 40 are dropped for the parent disappearing without replacement or death before wave 10, and 18 for being inconsistent between waves 10 and 15 without death. Our main results are unchanged if we include these observations, however, so our decision to drop these observations is not consequential for our conclusions.

Table B.3 shows the resulting combinations of child–parent relationships for our remaining sample of 594.

Table B.3. Distribution of observations by combination of parent types

Combination of parents	Observations
Both biological parents	441
Biological mother and step-father	21
Biological mother only	104
Biological father and step-mother	2
Biological father only	14
Both step-parents	7
Two biological parents in wave 10, one dies by wave 15	5
Total	594

It is possible for two children in our sample to have the same combination of parents if they are siblings. Table B.4 shows how many parent combinations there are with each number of children. As can be seen, although most of the sample have a unique combination of parents, a third come from a parent combination where two children end up in our final sample, and a small number come from a combination where there are three or four children. When regression analysis is carried out, we cluster standard errors on parent combinations.

Table B.4. Distribution of observations by number of children

Number of children	Number of parent combinations	Number of children observations	Percentage of sample covered
1	346	346	58.2
2	95	190	32.0
3	18	54	9.1
4	1	4	0.7
All	460	594	100.0

We make one final sample selection. One child has a level of mortgage debt that is more than twice the value of their house. We drop this observation, leaving us with a final sample of 593 child observations.

Appendix C. Weighting

We are provided with survey weights which, when applied to the whole BHPS-USoc sample, would give a representative picture of the wider population. For outcomes that do not require linkage to parents and where non-response is not a significant issue, we can calculate statistics within the whole sample in the 1974 to 1986 birth cohort, for the relevant waves, weighted by the provided survey weights. This should give us the 'true' value for these outcomes within the population for those cohorts in the relevant waves. We do this, presenting statistics for selected outcomes in the first column of Table C.1. In the second column, we present the equivalent statistics for our sample of 593 observations, weighted by the baseline survey weights.

The table shows that our sample is indeed selected, having a lower percentage of low-educated children, higher levels of earnings and slightly lower rates of homeownership and financial wealth than the whole cross-section for the relevant waves.

To correct for our sample being selected, we adjust the sample weights using an inverse probability method. We obtain our adjustments by taking the whole sample for the relevant cohorts and waves and running a probit model where the outcome variable is a dummy variable for being in our sample and the explanatory variables are: being low-educated, region dummies, homeownership dummy, single-year-of-age dummies, earnings and a dummy for positive earnings. The probit is weighted using the survey weights. To obtain our created weights, we divide the survey weights by the predicted probability of being in the sample, as given by the estimated probit model. The final column of Table C.1 shows the summary statistics for our sample when using the adjusted weights. These adjusted weights are used to weight the sample in all of the analysis in the main body of the report.

Table C.1. Analysis of outcomes with survey weights and adjusted weights

Outcome	Whole sample, survey weights	Our sample, survey weights	Our sample, adjusted weights
Low-educated	26.1%	19.6%	24.1%
Earnings (£, 2019 prices)			
Mean	20,138	21,402	20,881
Median	16,700	20,040	17,980
25th percentile	0	2,248	0
75th percentile	31,865	33,152	32,115
Has positive earnings	69.9%	75.8%	70.0%
Homeowner	51.5%	44.3%	51.7%
Net financial wealth (£, 2019 prices)			
Mean	5,563	4,581	6,385
Median	0	0	0
25th percentile	–2,294	–2,070	–1,178
75th percentile	3,212	3,451	5,461
Observations	16,086	593	593

Source: BHPS-USoc, waves 15, 22 and 26.

