The long-term effects of inherited wealth on social equality

James T. Allen

Sandtable, 69 Old Street, London, EC1V 9HX
james@sandtable.com

Abstract

It is challenging to assess the long-term effects of policies related to the inheritance of wealth, due to the long timescales involved and the importance of family-level dynamics. However, agent-based models (ABMs) now provide a useful tool for simulating the effects of alternative policies, allowing a quantitative comparison to be made.

In this paper we describe an ABM simulating the long-term (up to 150 years) development of the UK population in terms of births, deaths and the passage of wealth through the generations. The parameterisation of the model is driven by publicly available data on the demographics of the UK, primarily from the Office for National Statistics. We compare two situations: first where wealth is inherited directly by the children of the deceased; secondly where wealth is placed in descendants' trusts for the benefit of the second and third generations that follow.

The simulations show a marked reduction in the Gini coefficient for the case where wealth is placed in descendants' trusts, relative to direct inheritance. This result implies that, overall, trusts of this sort act to reduce the level of wealth inequality across society. The same effect is seen whether examining the adult population as a whole or a single generation.

Keywords— Simulations; agent-based models; inheritance.

1 Introduction

1.1 Inherited wealth and social inequality

In the UK it is common for wealth to be passed down within families, with children inheriting all or most of their parents' estate. Policy-makers must balance the desire of people to leave their wealth to their children against the entrenchment of inequality that this system can produce. Policies such as inheritance tax address this issue by preferentially taxing people with the largest estates.

Aside from direct taxation, there are many other policies that could influence the effects of inheritance. In [1], David Halpern argued for the promotion of "descendants' trusts", in which the bulk of an estate is placed in a trust for the benefit of the grandchildren and great-grandchildren of the deceased, rather than their children. His argument is based on the observation that the correlation between your wealth and your descendants' wealth becomes weaker as the number of generations separating you gets greater. So an individual's chance of having one wealthy grandparent or great-grandparent is greater than their chance of having wealthy parents, and hence the descendants' trusts spread more wealth to those who would otherwise have little.

A further argument for descendants' trusts is that they improve the timing of inheritance, in that they cause people to receive inheritances earlier in their lives. The money they receive can then be used, for example, to support their education or help them buy a first house. Indeed, a majority of people expecting to inherit from their parents already plan to pass some or all of that straight down to the next generation, whether or not the will specifies this, in order to provide them with this assistance [2].

1.2 Agent-based models

The benefits or otherwise of descendants' trusts are very difficult to assess using traditional techniques. The long-term nature of the issue makes pilot studies impractical, while the strong network effects involved make any top-down analysis extremely challenging. Instead, a bottom-up approach is required; this can be provided by an agent-based model (ABM).

An ABM is a simulation of a population of agents that can interact with each other and with their shared environment. In the case of inherited wealth, each agent will represent an individual person, and the simulation must explicitly encode the behavioural rules that determine key aspects of their life such as their earnings, their savings, when they have children, who with, and when they die.

While ABMs have existed for many decades – Schelling's model of segregated housing being a famous early example [3] – recent advancements in computing power and the availability of data have greatly increased the number and variety of problems that can be analysed using ABMs. In particular, ABMs are increasingly being used in economics and social sciences to study the behaviour of systems such as the UK housing market [4] or financial markets [5].

In this paper we describe the use of an ABM to simulate the inheritance of wealth in the UK population over the next 150 years. By modifying the behavioural rules we are able to assess the long-term effects of alternative policies, providing valuable evidence to policy-makers. We describe the details of the model in the following section.

2 Model components

In order to test different approaches to wealth inheritance, we have built an ABM with three key components:

1. A population of agents who form relationships, have children, grow old and die.
2. An economic model to assign earnings and investments to each agent.
3. A method of passing wealth to an agent's descendants following their death.

Each of these components is described in the following subsections.

The ABM as a whole starts with a population of agents that represents the 2016 UK population, then evolves this population in a series of "ticks", or steps, each representing one year in the lives of the agents. Because of the stochastic nature of the simulation we run each setup 30 times and average the results.

2.1 Population

The demographics of the agent population are primarily set by data from the Office for National Statistics (ONS). We use an initial population of 10,000 agents, although this number varies during the simulation according to the birth and death rates.

2.1.1 Initial population

The initial population is assigned ages and sexes according to the fractions recorded in the 2011 census [6], and they are matched into relationships following the fraction of men or women at each age who are in a cohabiting relationship [7]. Education levels, encoded as one of five values from "no qualifications" to "undergraduate degree or higher", are also set following the census data as a function of age and sex [8, 9].

2.1.2 Relationships

In each tick, single agents have a chance of forming a relationship. The simulation's matching process is female-led, due to the greater importance of the mother's age in determining fertility. First, the single women that will form relationships are chosen according to the marriage rate for women as a function of age [10], multiplied by a factor f. The value of f is set at $100.0/52.3 = 1.91$ to account for children born outside of marriage: currently 52.3% of children are born to married and 47.7% to unmarried parents [11].

Having selected the female partners, we compare them to single men according to their age, wealth and education level. The level of difference, δ_{ij}, between potential partners i and j is given by

$$\delta_{ij} = \left(\frac{a_i - a_j}{3}\right)^2 + \left(\frac{s_i - s_j}{0.5/m}\right)^2 + \left(\frac{w_i - w_j}{25000/m}\right)^2 \qquad (1)$$

where a_i is the age of agent i in years, s_i their level of education on a 0-4 scale, w_i their wealth in pounds, and m is a dimensionless parameter that sets the strength of assortative mating – the tendency of people to choose partners similar to themselves – in the model. We set the value of m at 0.25 to reproduce the correlation between partners' education levels observed in [15], but also show the results of varying this parameter in Section 3.3.

After calculating the difference between each potential pair of agents, their compatibility, c_{ij}, is then

$$c_{ij} = \exp\left(-\delta_{ij}/2\right) \qquad (2)$$

Each female partner then draws one match from the eligible male partners, with probability proportional to their compatibility. If two or more female partners draw the same male partner all but one, chosen randomly, redraws until all matches are unique.

Existing relationships also have a chance of ending in each tick. The separation rate is taken from the observed divorce rate as a function of the length of the relationship and the year in which it started [12], assuming that in future years the divorce rate for relationships of a given length is equal to the most recently measured rate for relationships of that length.

2.1.3 Births and deaths

During each tick the female agents that are in a relationship have a chance of having a child, given by the fertility rate for married women as a function of age [11]. As described above, the number of relationships in the model is increased above the marriage rate such that the overall birth rate matches the true rate.

At the point of birth each agent is assigned a future educational attainment, although it has no effect until they begin earning in adulthood. The child's attainment is determined primarily by that of their father [16], with some random scatter applied, and adjusted to reproduce the current distribution of attainment levels.

Finally, each agent has a probability of dying during the tick, given by the ONS projected mortality rates as a function of age and sex [13].

2.2 Economics

The simulation has a simple economic model to assign wealth and income to each agent. Each agent is characterised by four factors: their age, sex, level of education and a "boost" variable drawn from a normal distribution with mean 0.0 and standard deviation 0.7. The boost is used to account for all factors that influence somebody's income beyond the variables explicitly included in the model.

The earning potential of each agent is calculated by interpolating the mean income values as a function of age, sex and education, then dividing by the mean of these values across all agents and finally adding the boost value. The resulting potential value is used to put the agents in order of their expected income, from which they can be mapped across to the true income distribution in the UK [14], the same procedure was applied to observed distributions of wealth, using the same boost value as for income.

During the simulation, each agent is assumed to save or invest 5% of their earnings. This fraction represents direct savings as well as processes such as paying off a mortgage. Any wealth they have additionally grows at a rate of 3% a year. While these assumptions are clearly simplistic compared to actual personal finances, they suffice for the purposes of the simulation.

On reaching the age of 65, each agent immediately retires and no longer earns an income. Half of their wealth at that point is removed from the simulation, to represent the purchase of an annuity. In later years their remaining wealth decreases by 1% a year.

2.3 Inheritance

In order to test the effect of descendants' trusts we run the simulation twice, once in which all agents use direct inheritance and once in which all agents use descendants' trusts.

For direct inheritance, when an agent with no current partner dies their wealth is immediately split equally between their surviving children. If they have a surviving partner, the share per person is equal for each child and the partner. The share for children of previous partners is paid to them immediately, while that for both the current partner and their children is paid to the partner (on the understanding that it will later be passed to the children).

When using descendants' trusts, the share for any surviving partner is calculated in the same way and immediately paid to them (not including the amount allocated for their children). The remaining amount is split into two equal parts, and used to create trusts for the grandchildren and great-grandchildren, respectively, of the dying agent. The trusts earn interest at the same rate as for wealth held by agents. When each of these descendants reaches the age of 18, or immediately if they are already 18 or above, they withdraw from their generation's trust an amount equal to the current value of the trust divided by the number of people still expecting to benefit from it.

In either case, if an agent dies with no partner or descendants then their wealth is removed from the simulation.

3 Results

A key advantage of ABMs is that they allow a consistent analysis of both individual-level (agent-level) and population-level results. We examine each of these levels in the following subsections.

3.1 Individual-level results

Figure 1 shows the life history of an agent from the simulation in which all wealth passes by direct inheritance, i.e. to the children of the deceased. For the sake of discussion we will call her "Mary". The figure shows some key events in Mary's life: her birth, entering and leaving a relationship, the births of her children, and finally her death. On top of this is her wealth, which gradually increases throughout her working life before dropping on retirement. When Mary is 70 years old the second of her parents dies and she inherits about £100,000, but she herself dies just a few years later.

In contrast, Figure 2 shows the life history of an agent (who we shall call "Paul") from the other simulation, in

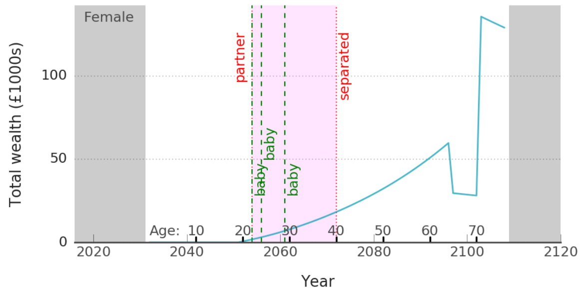

Figure 1: *The life history of "Mary", an agent in the direct inheritance simulation. Grey shaded regions mark the times before her birth and after her death. The pink shaded region marks the time she was in a relationship, and green dashed lines show the years in which her children were born. The blue curve shows her total wealth throughout her life.*

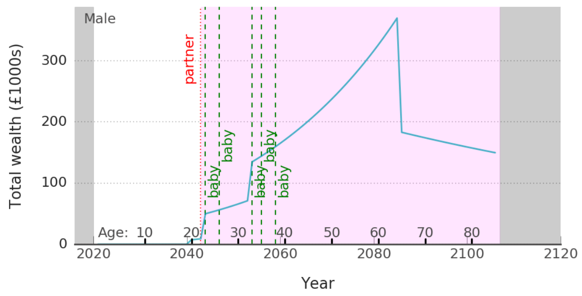

Figure 2: *The life history of "Paul", an agent in the descendants' trusts simulation. Colours and symbols are the same as for Figure 1.*

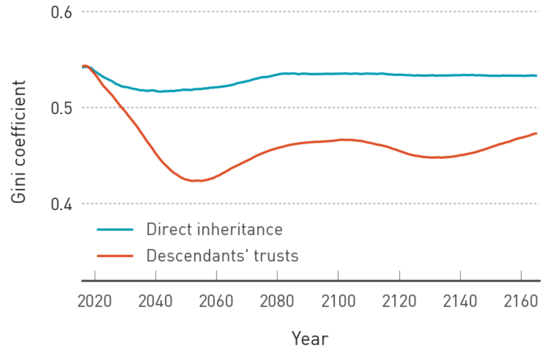

Figure 3: *The Gini coefficient for wealth measured across all adults, comparing the behaviour for direct inheritance (blue upper line) and descendants' trusts (red lower line).*

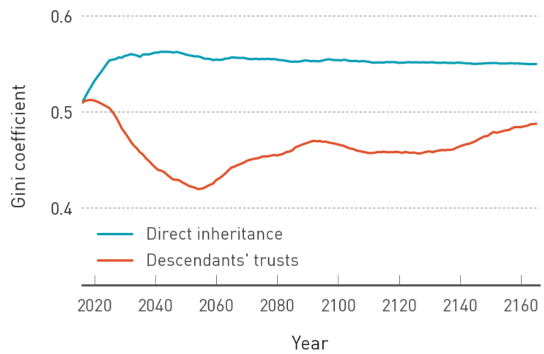

Figure 4: *As Figure 3 but including only adults aged 20–39.*

which all wealth is placed in descendants' trusts. The history of Paul's wealth shows a series of smaller jumps (up to £40,000) during his 20s and 30s as he inherits from his grandparents and great-grandparents.

Paul and Mary have similar histories in terms of the total amounts they earn and inherit during their lives, but a simple comparison of Figures 1 and 2 shows that throughout their adult lives Paul is significantly wealthier than Mary. This is an immediate result of having inherited earlier in his life, and illustrates the advantage of descendants' trusts in terms of injecting wealth when people are younger and arguably more in need.

3.2 Population-level results

We assess the success of the descendants' trusts by comparing the long-term behaviour of the Gini coefficient, a common measure of inequality, between the two simulations.

In Figure 3 we show the projected Gini coefficients over the next 150 years for each of the two cases. The Gini coefficient here is measured from the accumulated wealth values for all adults in the simulation. While the Gini coefficient stays approximately constant around 0.54 when direct inheritance is used, descendants' trusts produce a rapid decrease in the coefficient, reaching a low of 0.425 in the 2050s. The Gini coefficient then begins to increase again, with some oscillation, reaching 0.47 by the end of the simulation in the 2160s. The oscillation is caused by the sudden wholesale switch to descendants' trusts, which leaves one generation (people whose grandparents have already died but whose parents haven't) with little or no inheritance.

There are two factors that can contribute to the decrease in the Gini coefficient shown in Figure 3: a redistribution of wealth between generations and a redistribution within individual generations. To isolate the second factor, Figure 4 shows the Gini coefficient for adults aged 20–39 only, so at any point a single generation is being measured. While the difference between the two simulations is no longer as great as for the entire adult population, there is still a large decrease in the Gini coefficient brought about by the descendants' trusts.

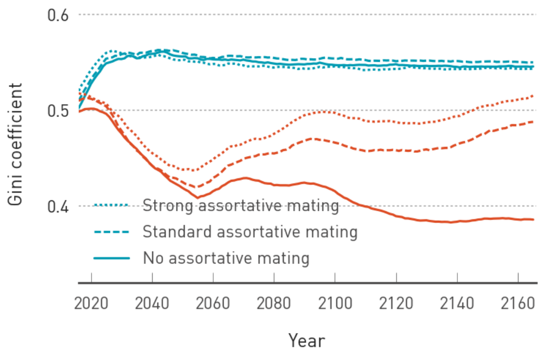

Figure 5: *As Figure 4 but showing three different levels of assortative mating. The dashed lines in this figure correspond to the lines in Figure 4.*

3.3 Assortative mating

The strength of the effect discussed above is strongly dependent on the level of assortative mating in the population. If agents always choose to partner with others who have very similar levels of education and wealth to themselves, there is less variation in wealth within a family and hence less advantage to the descendants' trusts. This effect is illustrated in Figure 5, which shows the Gini coefficient for adults aged 20–39 under three different levels of assortative mating. The "standard" level is that used in the previous figures, corresponding to $m = 0.2$ in equation 1. "Strong" and "no" assortative mating correspond to $m = 0.4$ and $m = 0.2$, respectively.

When there is no assortative mating mating the effect of descendants' trusts becomes much stronger, with the Gini coefficient falling below 0.4. Conversely, increasing the strength of assortative mating decreases the effectiveness of the trusts. However, even in the case of the "strong" assortative mating there remains a significant difference between the two simulations after the 150 years covered.

Figure 5 shows essentially no difference in the direct inheritance case when the assortative mating is varied, because only adults aged 20–39 are included. These agents are very unlikely to have received any direct inheritance, and the model includes no link between their income and their mother's, so assortative mating is not able to have any effect.

4 Conclusions

We have constructed an ABM that simulates the lives of a representative population of the UK, and used it to analyse the societal effects of inheriting wealth through descendants' trusts. The principal results are:

- Descendants' trusts produce a significant decrease in the Gini coefficient, measured for wealth across the adult population, when compared to direct inheritance.
- A similar decrease is seen when measuring across adults aged 20–39 only, indicating wealth is reallocated within a generation as well as between generations.
- The strength of the decrease is a function of the level of assortative mating in the population, with strong assortative mating producing a weaker effect.

More generally, this study provides an example of the application of ABMs to understand the implications of policy decisions. Particularly when looking at long-term effects that could not be measured from a pilot study, ABMs provide an invaluable tool to help guide policy-makers in their decisions.

Acknowledgements

I thank my colleagues at Sandtable for their contributions to this study, in the form of discussions relating to the model itself and the provision of infrastructure to help develop and run such models.

References

[1] Halpern, D., "The hidden wealth of nations", 2010, Polity Press

[2] Royal London, "Will harassed 'baby boomers' rescue Generation Rent?", 2017, Royal London Policy Paper 12

[3] Schelling, T. C., "Dynamic Models of Segregation", 1971, Journal of Mathematical Sociology, 1, 143

[4] Baptista, R., et al., "Macroprudential policy in an agent-based model of the UK housing market", 2016, Bank of England Staff Working Paper No. 619

[5] Buchanan, M., "Economics: Meltdown modelling", 2009, Nature, 460, 680

[6] Census table DC1117EW, https://www.nomisweb.co.uk/census/2011/dc1117ew

[7] Office for National Statistics, "Population estimates by marital status and living arrangements, England and Wales"

[8] Census table LC5102EW,
https://www.nomisweb.co.uk/census/2011/lc5102ew

[9] Census table LC5107EW,
https://www.nomisweb.co.uk/census/2011/lc5107ew

[10] Office for National Statistics, “Marriage statistics, cohabitation and cohort analyses”

[11] Office for National Statistics, “Births by parents’ characteristics”

[12] Office for National Statistics, “Divorces in England and Wales: Age at Marriage, Duration of Marriage and Cohort Analyses”

[13] Office for National Statistics, “Mortality Rates (qx), Principal Projection, United Kingdom”

[14] HMRC, “Percentile points from 1 to 99 for total income before and after tax”

[15] Blanden, J., “Intergenerational Mobility and Assortative Mating in the UK”, 2004, Centre for Economic Performance

[16] Office for National Statistics, “How do childhood circumstances affect poverty and deprivation as an adult?”, 2014