

Inequality in an Equal Society: Theory and Evidence

Jochen O. Mierau* James Rockey
University of Groningen University of Leicester

3rd October 2016

Abstract

We consider empirically the degree of wealth and income inequality that would prevail in a society in which, other than differences in age, everyone was equal. Theory suggests that life-cycle factors will still lead to substantial ‘natural’ inequality. Analysing cross-national data from the recently collected National Transfer Accounts we find that societies with no other source of inequality will exhibit substantial concentrations of income and wealth – Gini coefficients of circa 0.45 are common for income, purely due to these life-cycle effects. Using a modified Gini coefficient we find a substantial fraction of observed inequality can be attributed to this natural inequality. Finally, we show that if relative cohort sizes were equal to their long run values inequality would increase further.

KEYWORDS: Income Inequality, Wealth Inequality, National Transfer Accounts.

JEL-CODES: D31, J10

1 Introduction

The most equal society will exhibit a substantial degree of income and wealth inequality. Even in the absence of differences in talent, individuals approaching retirement will be substantially wealthier than individuals still well away from retirement. Moreover, experience and seniority mean that older workers will have higher wages than their younger colleagues. Jointly, such life-cycle aspects of income and wealth give rise to a degree of inequality that is natural in all societies – even in societies where each individual over the course of the life-cycle is exactly the same as any other individual.

*We thank Viola Angelini, Tony Atkinson, Gianni De Fraja, David de la Croix, Martin Jensen, Melanie Krause, Mariacristina De Nardi, David Rojo-Arjona, Richard Suen, and Chris Wallace as well as conference participants at the Royal Economic Society Annual Meeting and at the Annual Meeting of the European Economic Association for useful comments. Any errors are ours alone. Email: j.o.mierau@rug.nl and james.rockey@le.ac.uk

An early version of this argument was made by Atkinson (1971), who suggested that the distribution of wealth should be expected to be unequal solely due to differences in accumulated savings over the life-cycle. In a related contribution Paglin (1975) uses an argument similar to Atkinson's to suggest that popular measures of inequality such as the Gini coefficient should be corrected for the age structure inherent in income and wealth profiles. While Paglin's suggestion for a correction has not been without controversy,¹ the core of his argument – that inequality measures should be adjusted for the underlying life-cycle structure – has survived unscathed. Indeed, Huggett (1996), employed it as a basis of his inquiry into wealth inequality in the United States. More recently Mierau and Turnovsky (2014) have used the life-cycle argument to establish the theoretical value of wealth inequality that would prevail in a society where differences in wealth holdings are solely generated by different ages. Christening the ensuing value the “natural rate of wealth inequality”, they go on to show, in simulations for the US, the Gini coefficient describing this natural rate is substantial – around 0.35.

The central contribution of this paper is to take the life-cycle argument to the data. In doing so we document how much income and wealth inequality is due solely to life-cycle effects and by implication how much reflects other factors. To this end, we use the recently released National Transfer Accounts (NTA). Uniquely, these provide precise life-cycle income and savings profiles using consistent conventions for a range of countries (see, Lee and Mason, 2014). Using the NTAs in combination with various other data sources we show that even in the absence of any inequality between individuals of the same age group, the societies we focus on exhibit substantial degrees of income and wealth inequality. In particular, we show that natural and actual levels of income inequality for most of the population, as measured by Gini coefficients, are very similar. Importantly, this is less true for wealth. By contrast, using measures focused on the tail of the distribution – such as the income or wealth of the richest 1%– we find that actual income and wealth distributions are substantially more skewed than the natural rate suggests.

Our findings suggest that the difference between natural and empirical levels of inequality is itself useful as a measure of inequality. Building on Formby and Seaks (1980) we formalise this notion as a Life-cycle Adjusted Gini coefficient and calculate the degree of excess inequality. We find that, for most countries, excess income inequality is low but excess wealth inequality is often substantial. One reading of these results is that the scope for redistribution of wealth is much greater than that of income. This is because, if most income inequality is due to life-cycle factors then

¹See the numerous rounds of comments and replies generated by his paper.

additional redistribution of income will involve redistribution from the old to the young. Whilst, there is no reason this cannot be done, there are few examples of such a policy being implemented.

By documenting the relationship between the demographic structure and the natural rate of inequality we also contribute to the literature on trends in inequality. To get a feel for how demography affects natural inequality – and thereby how it may change over time – we assess the impact of the disproportionate size of the baby-boom generation on natural inequality. To this end, we redeploy our Life-cycle Adjusted Gini coefficient to compare current natural rates of inequality to those that would prevail if the demographic structure is in its long-run equilibrium. This exercise suggests that the bulge on the demographic pyramid generated by the baby boom is depressing the natural rates. Hence, in the future, as the demographic pyramid settles into its long-run equilibrium, wealth and income inequality will increase. Perhaps worryingly, this process will accelerate further the trend of increasing wealth inequality documented by the seminal contributions of Piketty (2003), Piketty and Saez (2003) and more recently, Atkinson et al. (2011), Piketty and Saez (2014), and Saez and Zucman (2016). In that sense, our paper contributes to the extant literature on inequality trends by highlighting that demographic forces will further contribute to the observed disparity in income and wealth.

Our focus on the level of inequality due solely to life-cycle factors is directly related to the prominent literature that studies the determinants of the distributions of earnings and wealth. For example, Huggett et al. (2011) consider how shocks received at different life stages affect lifetime income. The distribution of wealth is studied by Cagetti and De Nardi (2006) who study a quantitative model of occupational choice with the potential for entrepreneurship and study the role bequests and restrictions on investment play in determining wealth inequality. See also Neal and Rosen (2000) for a review and Huggett et al. (2006) for a more recent example attempting to match the extent to which more or less sophisticated life-cycle models can explain observed income-inequality. In this class of models life-cycle inequality is determined by the choice of parameters, often calibrated to US data, and the form of the model. As in Cagetti and De Nardi (2006), this approach allows for sophisticated analyses of the interaction of different features of an economy but any estimates depend on how well the model corresponds to reality and how precisely the parameters are chosen. Our approach is different, we use new data to study the empirical importance of life-cycle inequality for income and wealth without recourse to additional assumptions. One way we contribute to this literature is by providing empirical evidence as to the extent to which income and wealth inequality should be attributed to life-cycle effects in this type of model.

The remainder of the paper proceeds as follows. The next section fixes ideas by providing a formalization of the above arguments using a stylized model of the individual life cycle. It also introduces the Life-cycle Adjusted Gini. Section 3 takes the notion of natural inequality to the NTA data. It reports the natural rates of inequality for both income and wealth and contrasts these natural rates with the level of inequality observed empirically. The following two sections employ the natural rates to empirically assess the Life-cycle Adjusted Gini and to study the sensitivity of the natural rates to a return to the demographic steady state, respectively. We close with a brief conclusion. The Appendix collects details of the model and simulations and provides some further discussion on the Life-cycle Adjusted Gini.

2 Natural Rates of Inequality

To fix ideas we start with a stylized exposition of the levels of income and wealth inequality that would prevail if the only difference between individuals is that they are a different stage of their life cycle. Starting with income inequality, consider the following process of labour income:

$$W(v, t) = E(t - v)w(t), \tag{1}$$

where $W(v, t)$ is the income at time t of an individual born at time v , $w(t)$ is the economy wide wage rate and $E(t - v)$ is an individual scaling factor that creates a life-cycle pattern in labour income. $E(t - v)$ can be driven by many factors, which, for the sake of brevity we do not model separately. Indeed, for the current purpose it suffices to acknowledge that $E(t - v)$ can contain experience effects by which more senior workers earn more than junior workers but also institutional factors such as a social security system that redistributes income from workers to retirees. Panel A of Figure 1 exhibits a typical life-cycle pattern of non-asset income.²

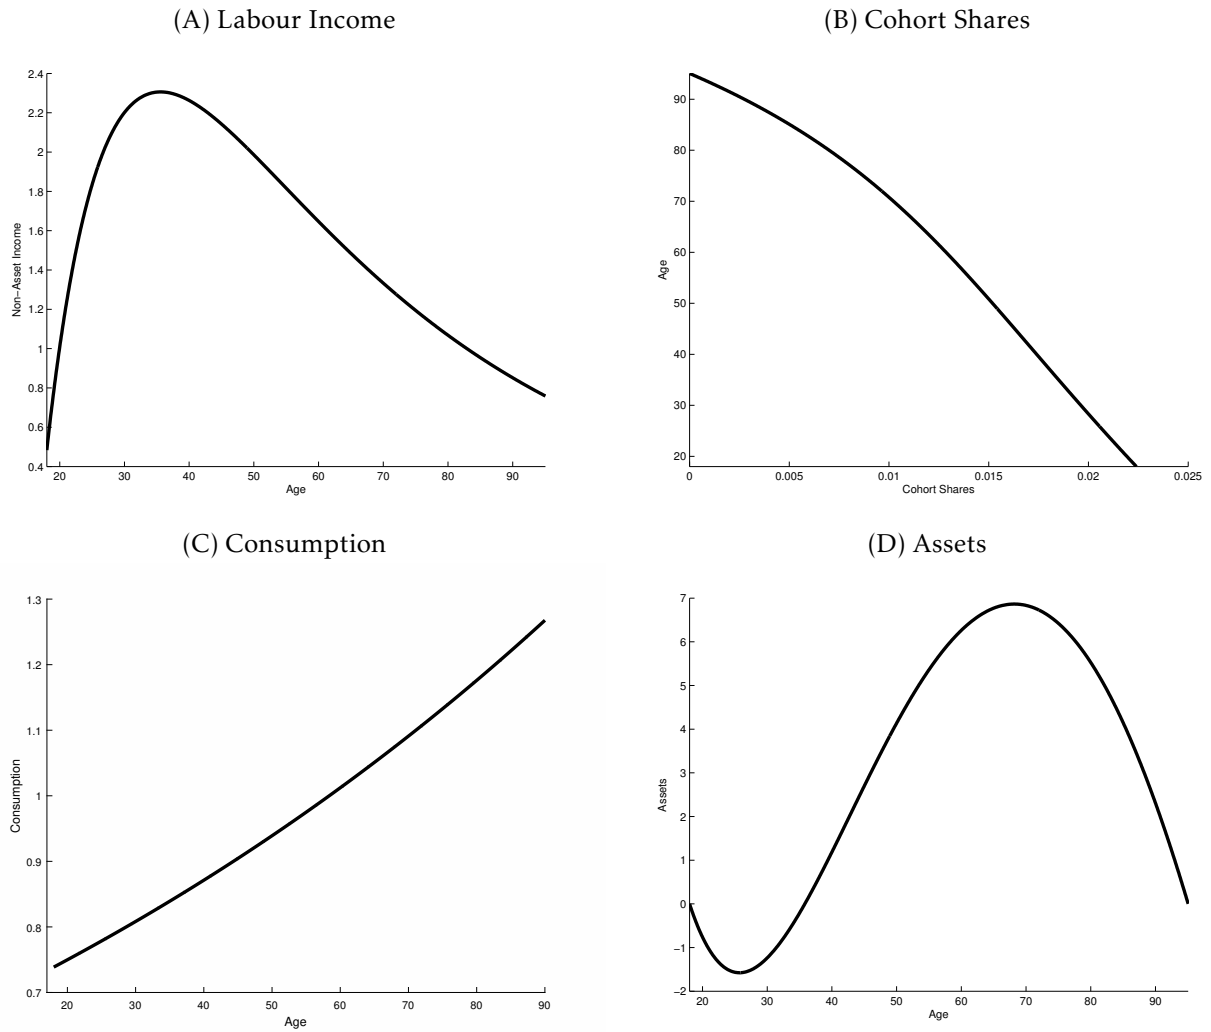
Panel B of Figure 1 displays the relative size of each cohort for a demographic structure that is in its steady state (see, Lotka, 1998).³ This provides the triangular shape that makes up the typical population pyramid.

Combining the life-cycle pattern of income of Panel A with the structure of cohort shares in Panel B implies that even if each individual has the same income at the same age, the fact that individuals of different ages coexist impose a level of income inequality into the economy. The particular

²The graph is a smoothed version of the life-cycle human capital pattern of Hansen (1993), which has been extrapolated to create a stylized social security benefit in old age. See Appendix C for details.

³The cohort shares are based on the mortality structure of the 2006 United States cohort under the assumption of constant fertility. See Appendix C for more details.

Figure 1: Life-cycle Profiles



Notes: Panel (A) displays the life-cycle profile of non-asset income. Panel (B) shows the relative size of each age group. Panel (C) exhibits the life-cycle consumption profile generated by the model in Section 2. Panel (D) displays the associated asset profile as also derived from the model.

example given above, for instance, implies a Gini coefficient of 0.11. While not large, it certainly vitiates the idea of a natural level of income inequality.

In keeping with the reasoning set out by Atkinson (1971) we express the arguments in the current section in terms of the life-cycle of an individual. But, when on the economic and demographic equilibrium paths, age and cohort levels of income and wealth are equivalent. That is, in equilibrium every cohort, say aged 37, is equivalent to any other cohort when aged 37. This means, that the trajectories in Figure 1 can be equivalently interpreted as describing an individual cohort's life-cycle or as a cross-sectional snap shot of the income and wealth levels of different cohorts.

To focus on wealth inequality we need to add more structure to our model. In particular, following the typical textbook life-cycle model, let the discounted life-time utility of an individual born at time v be given by:⁴

$$\Lambda(v) = \int_v^{v+D} U(C(v,t)) e^{-\rho(t-v)-M(t-v)} dt \quad (2)$$

where $C(v,t)$ is consumption at time t of an individual born at time v , ρ is the pure rate of time preference, $M(t-v)$ is the cumulative mortality rate and D is the maximum attainable age. While D is the maximum age, an individual has an increasing probability of death, $\mu(t-v) = M'(t-v)$, at any age before D . We assume an iso-elastic utility function $U(C(v,t))$, specifically:

$$U(C(v,t)) = \frac{C(v,t)^{1-1/\sigma} - 1}{1 - 1/\sigma} \quad (3)$$

where σ is the intertemporal elasticity of substitution. The budget constraint of the individual is:

$$\dot{A}(v,t) = r^A(t)A(v,t) - C(v,t) + W(v,t), \quad (4)$$

where $A(v,t)$ is the stock of financial assets, $\dot{A}(v,t) = \frac{\partial A(v,t)}{\partial t}$, $W(v,t)$ is non-asset income and $r^A(v,t)$ is the interest rate received, both of which depend on age.

Following Yaari (1965), we assume that individuals have recourse to a perfect annuity market. Hence, the return on assets becomes a composite of the market rate of interest ($r(t)$) and the annuity premium - which, due to competition between annuity firms, equals the probability of death ($\mu(t-v)$):⁵

$$r^A(v,t) = r(t) + \mu(t-v). \quad (5)$$

⁴In principle D can also be infinite, for instance if the specified survival function is of the Gompertz-Makeham form.

⁵In practice life-insurance markets need not be perfect, leading to a load factor on the annuity premium (see, Mitchell et al. (1999)). This can easily be accommodated in the current framework by following Hansen and İmrohoroğlu (2008) and Heijdra and Mierau (2012) and letting the life-insurance premium equal $(1-\lambda)\mu(t-v)$, where $\lambda \in [0,1]$ is the load factor. For sake of argument we focus on the case that $\lambda = 0$ but in Appendix F we also display the (hump-shaped) consumption profile arising in the case of $\lambda = 0.5$.

Individuals maximize 2 subject to 4, which provides the familiar consumption Euler equation:⁶

$$\frac{\dot{C}(v,t)}{C(v,t)} = \sigma(r - \rho), \quad (6)$$

where we have assumed that the economy is in its steady-state as is indicated by the absence of a time index on r . Solving 6 forward from time v allows us to write consumption at time t as:

$$C(v,t) = C(v,v)e^{\sigma(r-\rho)(t-v)}. \quad (7)$$

To obtain $C(v,v)$ (i.e., consumption of a new-born individual) we substitute 7 into 4 and solve the ensuing differential equation, which provides:

$$\tilde{C}(v,v) = w \frac{\int_v^{v+D} E(t-v)e^{-r(t-v)+M(t-v)} dt}{\int_v^{v+D} e^{-((1-\sigma)r+\sigma\rho)(t-v)+M(t-v)} dt}, \quad (8)$$

where $\tilde{C}(v,v)$ indicates the equilibrium value of $C(v,v)$, which we can then use to trace out the life-cycle consumption profile in Panel C of Figure 1.⁷ With $\tilde{C}(v,v)$ in hand we can solve 4 from any point in time forward to obtain the life-time path of assets:

$$A(v,t) = e^{r(t-v)+M(t-v)} \left(w \int_v^t E(t-v)e^{-r(t-v)+M(t-v)} dt - \tilde{C}(v,v) \int_v^t e^{-((1-\sigma)r+\sigma\rho)(t-v)+M(t-v)} dt \right). \quad (9)$$

For the special case that income is constant over the life-cycle (i.e., $E(t-v) = E$) it is straightforward to prove that $A(v,t)$ follows a hump-shape (see, Mierau and Turnovsky, 2014). For more elaborate income profiles, a proof is more involved, but the fact that $A(v,t)$ differs over the life-cycle is unambiguous. Using the parametrisation set out in Appendix C, we display the life-cycle profile of assets in Panel D of Figure 1. As can be seen it exhibits the commonly documented hump-shaped structure.

As above we can combine the life-cycle asset profile with the cohort shares to derive a measure of inequality that is natural to the society. In the case of wealth the Gini coefficient equals 0.54, which is reasonably large given that observed values in developed countries range from 0.55 in Japan and 0.80 in Denmark (see, Davies et al., 2011). In any case, the current argument clearly suggests the need for inequality measures to be corrected for the life-cycle structure of income and wealth – as suggested by Atkinson (1971) and Paglin (1975).

⁶Details of the derivation are contained in Appendix B.

⁷We return to the issue of the structure of the consumption profile (hump-shaped vs. upward sloping) in footnote 13 and Appendix F.

To formalize the reasoning developed above we summarize the main conclusions from the model in the following theorem.

Theorem 1. The Gini Coefficient of Income (Wealth) is positive in the presence of a non-flat life-cycle income (wealth) profile.

Corollary 1.1. Perfect income (wealth) equality implies a flat income (wealth) profile over the life-cycle.

The proof works by writing the Gini Coefficient as a product of the standardised variation of income, and the correlation of income with its rank, and noting that both of these terms are only zero when income is constant for all ages. The proof itself is in Appendix A.

Considering that observed inequality is generated by a host of factors, it seems appropriate to view natural inequality as a benchmark, deviations from which are useful as indicators of life-cycle adjusted measures of inequality.⁸ Figure 2 reproduces the conventional graph defining the Gini coefficient, but with an additional Lorenz curve. The thick curved line is the life-cycle Lorenz curve – the Lorenz curve associated with the natural rate – and the dashed line is the actual Lorenz curve. A indicates the area between the line of equality and the life-cycle Lorenz curve and B and B' indicate the areas under the life-cycle and actual Lorenz curves, respectively. The natural rate Gini can be expressed as: $\theta^{NR} = 1 - 2B$, similarly the non-adjusted or conventional Gini coefficient can be expressed as: $\theta^U = 1 - 2B'$. Using the graph we can also define the Life-cycle Adjusted Gini as: $\theta^{LA} = \frac{B-B'}{B}$, which can be derived from the above Ginis as:


$$\theta^{LA} = \frac{\theta^U - \theta^{NR}}{1 - \theta^{NR}}, \quad (10)$$

implying that a society with only natural inequality will have $\theta^{LA} = 0$, while a society exhibiting inequality in excess of natural inequality will take positive adjusted values.

Focusing on the Paglin (1975) debate about how to properly correct for age factors in inequality, we can observe that what we call the natural rate comes closest to what he calls the A(ge)-Gini, which was not the source of controversy. Importantly, however, while he derives the A-Gini from the average wages (or wealth holdings) per age group, we will derive it as generated from an economy in which everybody is the same at the same stage of the life cycle and in which there are no frictions such that everyone is paid their marginal product. Hence, while we note that our θ^{LA} is akin to the Modified-Paglin Gini suggested by Formby and Seaks (1980), they differ in that theirs is bounded

⁸Our life-cycle adjusted inequality indicator is related to the measure of *unfair* inequality of Almås et al. (2011), which aims at correcting inequality measures for much more than age factors alone. For the current context the advantage of our measure is that it can be derived without having recourse to detailed microdata thereby enabling us to compare excess inequality internationally.

Figure 2: The Life-Cycle Adjusted Gini Coefficient



strictly between 0 and 1, while ours may be negative.⁹ Negative values will occur if market frictions are such that they homogenize the age-earnings profile by a sufficient amount to more than offset the additional inequality due to cross-sectional variation in earnings. This will mean that there is less inequality in a society than that suggested by the natural rate.¹⁰ In that case Lorenz dominance of the solid over the dashed line of Figure 2 will be violated. Hence, in principle, $\theta^{LA} \in [-\infty, 1]$, but in practice the range will be substantially smaller. This is clear if we consider the approximate range of values found below, if $\theta^{NR} = 0.35$ (conversely, $\theta^{NR} = 0.5$) and $\theta^U = 0.5$ ($\theta^U = 0.35$) then $\theta^{LA} = 0.23$ ($\theta^{LA} = -0.18$).¹¹

In sum, taking inspiration from Atkinson (1971) and Paglin (1975) this section has shown that a stylized economy populated by individuals who are equal to each other at every stage of the life-cycle displays a substantial degree of income and wealth inequality. Moreover, we have seen that this minimum can be used to calculate a Life-Cycle Adjusted Gini coefficient. In the remainder of this paper we takes these insights to the data.

⁹The modification of Formby and Seaks (1980) of the Paglin (1975) measure amounts to redefining the the denominator of θ^{LA} as B and not A + B.

¹⁰For this reason we can also not decompose observed inequality into a pure age component, an intra-cohort component and an residual term representing the interaction between the two as in Mookherjee and Shorrocks (1982) and Lambert and Aronson (1993).

¹¹Further discussion of the properties of θ^{LA} , including some notes on the small sample properties, may be found in Appendix E.

3 Inequality in an Equal Society

Having established that a society where individuals are equal at each stage of the life-cycle exhibits a substantial degree of inequality, this section empirically assesses the quantitative importance of this inequality. To this end we start by describing the data and sketch it against the life-cycle profiles generated by the model above. After that we turn to the core of our analysis in which we calculate various inequality measures for a range of western countries. Importantly, by sourcing the trajectories of income and wealth straight from the data we assure that potential caveats of the life-cycle model (see, Browning and Crossley, 2001, Attanasio and Weber, 2010) do not affect our empirical results derived below. In particular, as we will see, the empirical data gives similar, but not equivalent, profiles of income and wealth. An important series of papers, including Cagetti and De Nardi (2006, 2009), seeks to develop models that can more fully explain the wealth distribution. But, our analysis depends only on the data and the simple model we presented above is purely to fix ideas.¹²

National Transfer Accounts

Understanding the relationship between the life-cycle and inequality requires accurate measurement of income and wealth disaggregated by age and documented with consistent conventions around the world. Until recently such data were unavailable but in a major advance the National Transfers Accounts (NTA) project of Lee and Mason (2014) addressed exactly this gap. As they write:

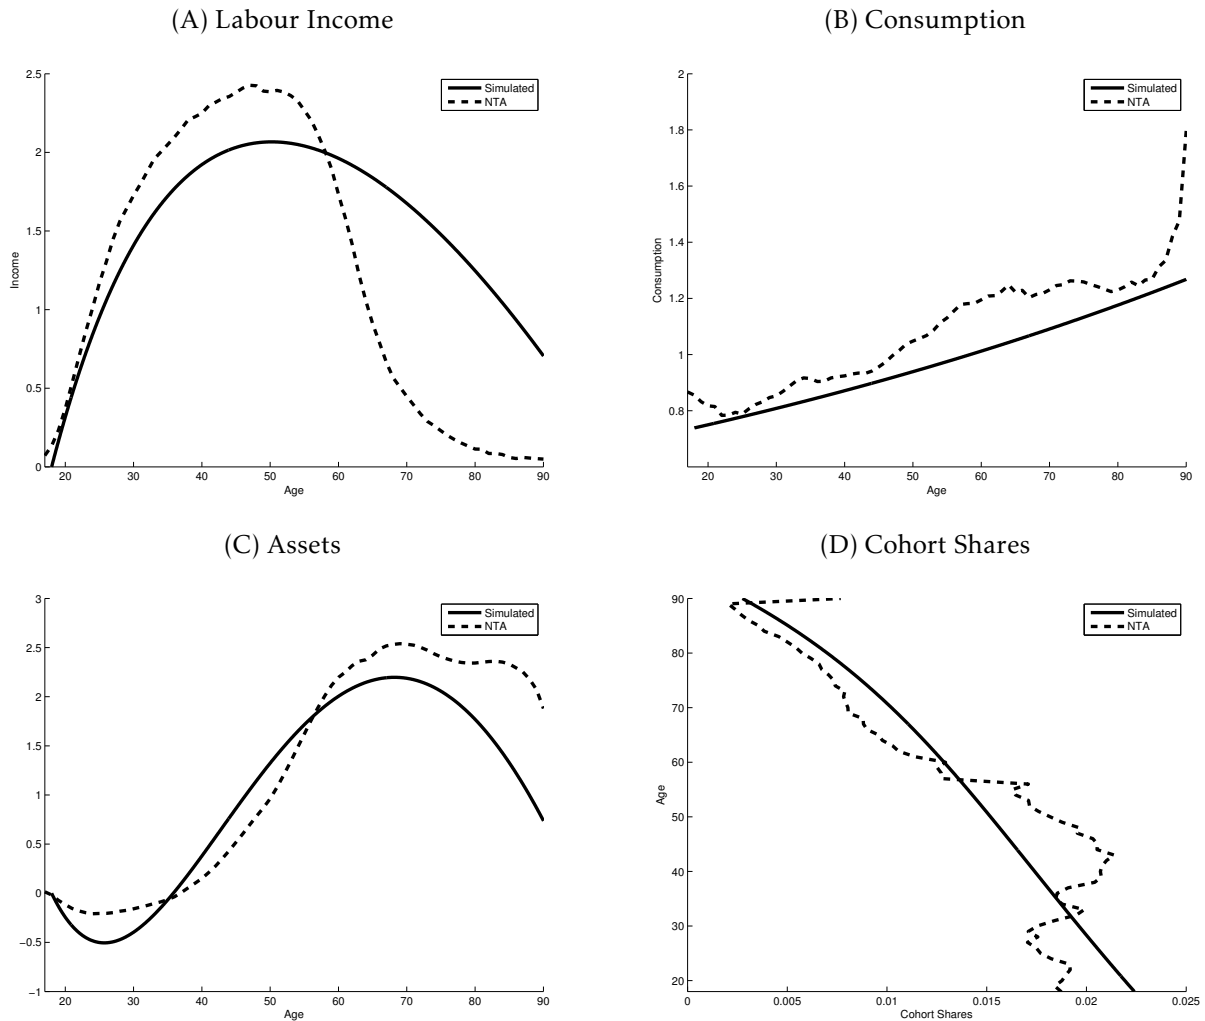
The National Transfer Accounts constitute a complete, systematic and coherent accounting of economic flows from one age group or generation to another, [...] in a given calendar year.

For the current purpose the NTAs provide precise data on labour and self-employment income by age for a broad set of western countries. As described in detail in UN (2013) the income profiles of the NTAs are constructed using household surveys or administrative data as main input. The raw, individual, data derived from such sources are then averaged by age, leading to a large - yet noisy - collection of data points. In consequence, the data are smoothed to create a continuous data series. Importantly, due to the smoothing, the income series are not simply age-group averages but represent what we consider as the profile of income that would prevail if at every point of the life-cycle everybody would be the same - as required by the arguments developed above. We refer the interested reader to UN (2013) for additional details on data definitions, coding conventions and smoothing procedures.

¹²See Cagetti and De Nardi (2006) for an anthology of approaches attempting to match life-cycle models to wealth inequality data.

We display the life-cycle profile of income from the NTAs for the United States in Panel A of Figure 3. For comparison we also plot with a solid line the income profile obtained with the model above. While both the simulated and the actual data display a clear life-cycle structure, we note that the actual income profile has a more pronounced hump than the theoretical one, implying that the natural rate of income inequality will be higher when determined using actual data than when calculated with simulated data – a finding borne out in the analysis below.

Figure 3: National Transfer Accounts



Notes: The various panels of Figure 3 compare the life-cycle profiles generated by the model of Section 2 to United States data from the National Transfer Accounts. In each panel the solid line is the simulated data and the dashed line is actual data. Panel (A) shows the profile of labour and self-employed income. Panel (B) displays the life-cycle consumption profile. Panel (C) is the asset profile. Panel (D) exhibits the actual US cohort shares and the cohort shares that would prevail in the demographic steady state.

The second core piece of data collected as part of the NTAs is individual consumption data by age group, which is again derived from household surveys. In contrast to income data, consumption data is often only known at the household level and, therefore, observed consumption is transformed into individual quantities using equivalence scales. Within the context of the NTAs consumption is

measured very broadly and also includes, for instance, the imputed rent of owner occupied housing. We display the life-cycle consumption in Panel B of Figure 3. As with the assets in Panel A, our stylized model captures the general pattern of consumption well, with the main difference being that actual consumption is less smooth. Both, however, show a clear upward trend over the life-cycle.¹³

Complementing the consumption data, the NTAs collect data on savings, which when combined provide the asset profile over the life-cycle. The ensuing profile is given in Panel C of Figure 3. Alongside the NTA data we trace out the profile suggested by the model and observe that the life-cycle structure of assets is captured well by our model.

The final piece of information that we draw from the NTAs is the size of the age cohorts. We display these for the United States in Panel D of Figure 3. Clearly there is a substantial difference between the smooth profile of cohort sizes suggested by the model and the actual cohort shares. This is because of the so-called Baby Boom and subsequent fertility bust which is reflected by the bulge in the demographic pyramid. We exploit this difference between the actual and the stylized demographic structure in Section 5 below where we analyse the role of the Baby Boom for the development of income and wealth inequality.

The NTAs are a cross-section and thus in constructing life-cycle profiles from them we are assuming that the (mean) life-cycle profile has been broadly consistent over time. In particular, we are assuming that the shares of consumption and savings in income have been broadly consistent for people of a given age over time. Therefore, we can calculate net wealth as an aggregate of past saving and the associated returns. We test this assumption in Appendix G and show that the life-cycle distribution of wealth calculated in the NTAs corresponds closely to that in the Survey of Consumer Finances.

While we have based the foregoing discussion of the data on the profiles for the United States, we may observe that the NTAs provide data for many more countries. For our current purpose we use a subset of these countries, in particular those countries for which we also have detailed mortality (Human Mortality Database, 2013) and fertility data (Human Fertility Database, 2013), allowing us to perform the analysis in Section 5. These countries are: Austria, Finland, Germany, Japan, Slovenia, Spain, Sweden, Taiwan, and the United States. Whilst limited in number, they are representative of a broad range of western countries.

¹³In light of the discussion on the hump-shaped consumption profile (see, for instance, Gourinchas and Parker (2002)) it serves to say that the consumption concept from the NTA displayed in the graph contains private consumption as well as health care expenditures and educational expenditures provided by the government. Focusing solely on private consumption (which is also available in the NTA) reveals the typical hump-shaped profile; see Appendix F.

Natural Inequality Around the World

Combining the income and wealth data with the demographic data from the NTAs we can now turn to the core of our analysis in which we calculate the empirical natural rates of wealth and income inequality for the nine countries in our sample. In line with the recent literature on long-run inequality (for example, (Atkinson and Piketty, 2010)) we complement the Gini coefficients used above with measures of wealth and income concentration among the top 1% and top 5% of the income and wealth distribution, respectively.¹⁴

Table 1 reports the natural rate of income and wealth inequality for each country in our data. The first thing to note, looking at the results for income in the top panel, is that the results are relatively consistent across countries. The Gini coefficient of income is in a narrow range of 0.46 – 0.5 for all countries. The same is true of the concentration indices, the top 1% consistently receives 3% of total income. There is more, but not much more, variation in the 5% index. Shares are now between 9% and 15%. This broad consistency is as expected; income profiles and the demographic structure are relatively similar across these countries.

There are three features of the income data of particular interest. Firstly, the natural Gini coefficients are relatively high. For comparison, Table 2 reports empirical Gini coefficients for the countries in our sample, before and after redistribution.¹⁵ We see that the predicted natural rate is around, or even just above, the current pre-redistribution level. One interpretation of this is that other forces than redistribution currently hold down inequality, for instance, non-market wage determination.

Secondly, the distribution of the natural rate shows much less variation than the empirical data - suggesting that it is not demographic factors that drive cross-country variation in inequality. Moreover, that the pre-redistribution data are also more dispersed suggests again that it is not just differences in government redistribution levels driving this variation. Below, we use the adjusted Gini coefficient to understand this variation in the degree of excess inequality better.

Finally, the concentration indices are consistently relatively low. This suggests that although the natural rate of inequality is relatively pronounced, it does not give rise to the extremely skewed top income distributions observed empirically. This conforms to received wisdom; the salaries of the highest earners reflects the skewed cross-sectional distribution of earning potentials rather than life-cycle factors.

¹⁴Given the finite-sample properties of the Gini Coefficient the choice of whether to group cohorts together rather than treating them individually will affect our results. Appendix E shows that treating them individually will minimise this bias which will be very small and negative.

¹⁵The pre-redistribution data are from the OECD (OECD, 2014) and the post-redistribution data from the Luxembourg Income Survey (LIS, 2013). Pre-redistribution data are unavailable for Taiwan.

Analysis of the wealth data in the lower panel of Table 1 suggests that, as should be expected, wealth is more unequally distributed than income. Gini coefficients are about 0.1 higher than they are for income, and are more varied. The concentration indices are now several times larger. The smallest increase is for the United States where they have, approximately, doubled.¹⁶ In Taiwan the wealth share of the top 1% is nearly ten times larger than its income share. Similarly, the 5% concentration index of wealth ranges from twice as large of that for income in the United States, to six times in Taiwan. It is typical in the data for the richest 5% of cohorts to own nearly half a nation’s wealth. The analysis of Cagetti and De Nardi (2006) suggests this additional, much more pronounced, variation may potentially be attributed to cross-country differences in bequest motives or borrowing constraints, particularly so between Taiwan and the United States.

The lower panel of Table 2 reports empirical wealth data for the countries in our sample, where available.¹⁷ As with income, the level of concentration predicted by the simulations is smaller than observed empirically. While the available data suggest that the 5% concentration index is qualitatively similar to the simulation results, the 1% index for the empirical data is substantially larger than for the simulation data. This suggests that, as for income, that within-cohort variation is necessary to understand the top end of the wealth distribution. This is unsurprising given the importance of dynastic wealth (see, Atkinson, 2013, Piketty and Zucman, 2015) and the distribution of returns to entrepreneurship (see, Cagetti and De Nardi, 2006, Quadrini, 1999, Hurst and Lusardi, 2004, Castañeda et al., 2003, Cagetti and De Nardi, 2009) in determining the right-tail of the wealth distribution.

Table 1: Natural Inequality Around the World

	Austria	Finland	Germany	Japan	Slovenia	Spain	Sweden	Taiwan	United States
Income									
Gini	0.46	0.48	0.48	0.47	0.50	0.47	0.47	0.47	0.46
1%	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03
5%	0.12	0.11	0.12	0.15	0.13	0.12	0.11	0.09	0.11
Wealth									
Gini	0.63	0.50	0.50	0.62	0.52	0.58	0.62	0.65	0.59
1%	0.14	0.14	0.10	0.08	0.16	0.11	0.09	0.29	0.05
5%	0.44	0.45	0.34	0.43	0.34	0.23	0.44	0.56	0.21

These figures are calculated using cohort income shares and cohort sizes from the NTA ((NTA, 2013)). We calculate inequality as if there is no within-cohort variation in income or wealth.

¹⁶This is not an artefact of the NTA, the Gini coefficient calculated using the Survey of Consumer Finances is 0.61, available on request.

¹⁷These data are taken from Davies et al. (2011).

Table 2: Actual Inequality Around the World

	Austria	Finland	Germany	Japan	Slovenia	Spain	Sweden	Taiwan	United States
Income									
Before	0.47	0.46	0.47	0.43	0.44	0.47	0.43	N/A	0.46
After	0.26	0.24	0.27	0.30	0.23	0.32	0.23	0.33	0.35
Wealth									
Gini	0.65	0.62	0.67	0.55	0.63	0.57	0.74	0.66	0.80

The income data are microdata based estimates from the OECD (2014) and the LIS (2013). Before refers to inequality before government transfers and taxes. After is inequality after all government redistribution. No such distinction is available for wealth, and the data reported include the effects of any government redistribution or taxes. The wealth data are taken from Davies et al. (2011).

To summarize – natural inequalities can account for the bulk of income inequality observed in a broad selection of countries around the world. Importantly, this holds much less for wealth inequality and we note that life-cycle factors cannot account for the extreme concentration of income and wealth in the tails of the distribution. Viewing the interplay between the life-cycle and inequality suggests that a.) life-cycle inequality is potentially a good benchmark to which actual inequality can be compared and b.) that a changing demographic structure has a direct impact on inequality measures. In what follows we flesh out these two suggestions.

4 Life-Cycle Adjusted Inequality

Table 3 contains calculated values of θ^{LA} for the countries in our sample for income and wealth. Consider first the results for income, calculated using pre-redistribution Gini indices. We see that the deviations from the natural rate are relatively small, and except for Austria they are either negative or zero. The difference is largest in Slovenia ($\theta^{LA} = -0.12$), Sweden, and Japan (both $\theta^{LA} = -0.08$). This suggests that, perhaps due to labour market institutions, pre-redistribution inequality in every country we study is lower than predicted by life-cycle factors alone. Turning now to the post-redistribution figures, we see that those for income are uniformly negative and larger in magnitude than those for pre-redistribution income. That they are larger is unsurprising, we expect redistribution in democracies (see, Meltzer and Richard, 1981), but that they are substantially negative suggests that both government and other institutions are reducing inequality such that it is substantially lower than would otherwise be expected solely due to life-cycle factors. Put differently, the equalising effect of these institutions exceeds the inequality owing to variation in earnings potentials.

The opposite is true for wealth inequality. These are mostly positive, and often substantial, suggesting much more wealth inequality than the natural-rate. The largest value is for the United States, where wealth inequality is 0.51 higher than would be expected solely due to wealth accumulation over the life-cycle, this is in stark contrast to Japan where it is 0.18 lower than expected. The combined difference of nearly 0.7 highlights the importance of taking into account differences in demographic structure when comparing Gini coefficients. Comparison of the unadjusted data suggests a difference of only 0.25. Taking these results together with those for income, we can see that excess inequality tends to be highest in the United States over all metrics, and lowest in Japan, but without a clear pattern in between. Considering Germany, Sweden, and the United States together, one interpretation is that these countries are more similar than is often suggested.

Considering the concentration indices reported in Table 1 and 4 together with the excess inequality reported in Table 3 suggests some implications for the possible additional extent of redistribution. Given that the natural rate Gini coefficients, particularly for income, are close to the observed level, as noted already this implies that additional redistribution would involve redistribution from the older to the younger. However, the shares of the top 1% are further removed from their natural rate implying that there is more room for redistribution from the very rich. This concentration of excess inequality at the rightmost extreme of the distribution is even more pronounced for wealth.

Table 3: Life-Cycle Adjusted Gini Indices

	Austria	Finland	Germany	Japan	Slovenia	Spain	Sweden	Taiwan	United States
Pre-Redistribution									
Income	0.02	-0.04	-0.02	-0.08	-0.12	0	-0.08	N/A	0
Post-Redistribution									
Income	-0.37	-0.46	-0.40	-0.32	-0.54	-0.28	-0.45	-0.26	-0.20
Wealth	0.05	0.24	0.34	-0.18	0.23	-0.02	0.32	0.03	0.51

The coefficients are the Life-cycle Adjusted Gini defined in 10 and described in Figure 2. A negative value indicates that observed inequality is lower than the level of inequality expected due to life-cycle factors alone. Pre-Redistribution data are unavailable for Taiwan. Data sources are as for Table 2.

5 Inequality and the Baby Boom

We have now seen that the life-cycle should be more prominent in understanding inequality. One further implication of this is that demographic dynamics will lead to changes in the distributions of income and wealth. Since the early work of Kuznets and Jenks (1953) and Atkinson and Harrison (1978), economists have paid considerable attention recently to long-run trends in inequality, prominent recent studies include Piketty (2003), Piketty and Saez (2003), Piketty (2011), Piketty and Saez


(2014) and Roine and Waldenström (2015). This section speaks to that literature and shows that the return after the baby-boom to the demographic steady-state will lead, *ceteris paribus*, to a substantial increase in inequality. This finding complements that of Almås et al. (2011) who study Norwegian micro-data and find that the baby-boom generation is responsible for a (temporary) reduction in inequality.

The baby-boom generation, for the US commonly considered those born between 1946 and 1964, represented a temporary upwards deviation from developed countries' otherwise stable demographic trajectories. Figure 4 reports long-run fertility data for a selection of countries. A first observation is that the Baby Boom was a common feature across many developed countries.¹⁸ Although, there are variations in timing and magnitude these fail to mask the overall scale of the boom - nearly an extra child per woman for 18 years. Also, notable is the rapidity with which it begins and ends. This large, sudden, and in demographic terms brief, rise in fertility has led to a one generation distortion in the demographic structure of the affected societies. Panel D of Figure 3 displays the consequences of the Baby Boom for the United States. There we can see that compared to a stable population distribution (the solid line) the actual population distribution (dashed line) features a bulge for the age groups most affected by the Baby Boom. Note that the post-baby-boom fertility bust can be seen by the indent in the actual population pyramid around ages 20 to 35. This difference between the actual and the stable demographic pyramid provides an interesting natural experiment for us to study. Our analysis suggests we should expect large increases in inequality as societies return to the demographic steady-state. Indeed, the estimated magnitudes are larger than the cross-country variation in current inequality levels.

Our approach is to calculate the cohort shares that would be consistent with the demographic steady-state given current mortality patterns for each country. We then repeat the analysis in Section 3 using these instead of the empirical shares. To do so we use the detailed demographic data in the Human Fertility Database (2013) and Human Mortality Database (2013), and apply the method of Mierau and Turnovsky (2014). Details of the calculations may be found in Appendix D. Our comparison is then between the natural rates of inequality in societies with current and steady-state demographics. That is we are asking how different the natural rate level of inequality would be if there had been no Baby Boom.

¹⁸All data are from the Human Fertility Database Human Fertility Database (2013). Germany refers to West Germany only, France excludes the overseas territories. The 'Average' series is the annual arithmetic mean of available observations.

Figure 4: The Baby Boom



The y-axis reports the number of children born per woman in a given year. The blue line is the (unweighted) mean fertility rate across the six countries reported. The red line highlights the USA for clarity but is otherwise identical in construction to those for other countries. The dotted vertical lines indicate the beginning and end of the baby-boom. All data are from the Human Fertility Database (2013).

Table 4 reports both empirical and steady state levels of natural inequality (expressed as Ginis) for each country, for both income and wealth.¹⁹ We also report the Life-Cycle Adjusted Gini which in this context reports the additional inequality due to the baby boomers.

While in Section 2 we introduced θ^{LA} as an adjustment of actual inequality rates, the same approach may be employed using alternative inequality rates as the bases for the adjustment. For our current purpose, an adjustment of natural rates depending on whether or not the demographic structure is in its long-run equilibrium will be useful. In that case the definition of θ^{LA} becomes:

$$\theta^{LA} = \frac{\theta^{NR} - \tilde{\theta}^{NR}}{1 - \tilde{\theta}^{NR}}, \quad (11)$$

where $\tilde{\theta}^{NR}$ is the natural rate that would prevail if income and wealth data were unchanged but cohort shares would reflect their equilibrium rates (see Panel D of Figure 3 for the difference between actual and equilibrium cohort shares). Viewing, θ^{LA} as in 11 means that here, if excess inequality is negative, then current natural rates are lower than their long-run counterparts.

¹⁹Data for Taiwan are not available.

We rely solely on the Gini coefficient for our comparison with the steady state as varying cohort sizes make concentration indices hard to interpret. Analysing first the income data it is clear that in each country a return to the steady state will lead to a marked increase in the natural level of inequality. This increase is largest in Spain, where the Gini index will increase by 0.11. A seven point increase is typical. To put this in context, inspection of Table 2 shows that this is larger than difference in *ex post* inequality between the United States and Sweden, or indeed any two countries in our data. The Life-Cycle Adjusted Gini tells a similar story, the effects of the Baby Boom are most pronounced in Spain and least so in Sweden. Most other countries have $\theta^{LA} \approx 0.15$, again suggesting that substantial increases in inequality are likely to be the norm.

The effects on the distribution of wealth are more muted yet still significant. As noted before, our calculation of cohort wealth shares assumes that savings behaviour is unchanged. This is a stronger assumption in the case of changing demographics but the qualitative conclusion is nonetheless clear – the baby boom means wealth inequality is considerably lower now than it will be in the steady state. There is on average an increase of around 0.02 points, and a 0.1 point increase in Spain. This is, however, relatively small compared to the cross-sectional variation in the simulation data. The Life-Cycle Adjusted Gini makes this discrepancy plainer. θ^{LA} is -0.31 for Spain and between 0 and -0.07 elsewhere. This suggests that while as we have seen life-cycle factors are much better able to explain the concentration in income than that of wealth, that in Spain life-cycle factors seem to drive both. The reasons for this are unclear. One explanation is that this partly reflects that less of wealth is driven by demographics than income, and partly reflects that life-cycle effects mean that wealth will remain relatively closely concentrated on those close to retirement age, as described by Panel C of Figure 3. A further explanation is that the changes are highly concentrated and thus not well captured by the Gini index, which is well-known to be more sensitive to changes in the middle of the distribution. The larger change in Spain, may well reflect differences in the nature of the demographic changes implied by convergence to the steady-state. As discussed in the introduction these results are particularly interesting in terms of the debate about the *fertility bust*. They suggest that concomitant with the problems of an ageing society discussed by Lee and Mason (2014) and Smeeding (2014) will be far higher levels of inequality, particularly in income. One implication of the natural rate, therefore, is that analysts of demographic change need to be mindful of distributional effects.

Table 4: Inequality and the Baby Boom

	Austria	Finland	Germany	Japan	Slovenia	Spain	Sweden	United States
Income (Gini)								
Current	0.46	0.47	0.49	0.47	0.50	0.47	0.47	0.46
Steady State	0.53	0.54	0.55	0.54	0.58	0.58	0.52	0.53
θ^{LA}	-0.15	-0.15	-0.13	-0.15	-0.19	-0.26	-0.1	-0.15
Wealth (Gini)								
Current	0.63	0.50	0.50	0.62	0.52	0.58	0.62	0.59
Steady State	0.65	0.50	0.52	0.63	0.55	0.68	0.63	0.60
θ^{LA}	-0.06	0.00	-0.04	-0.03	-0.07	-0.31	-0.03	-0.03

Current refers to the level of natural inequality given the current actual demographic structure. Steady State refers to the level of natural inequality that we would observe if each country were in its demographic steady state. That is the cohort size distribution that we would observe if the birth rate had been such for all cohorts that given current mortality patterns population equalled its current level there was no population growth. θ^{LA} is the Life-Cycle Adjusted Gini introduced in Section 2. The calculations of these birth rates are described in Appendix C. The Gini coefficient is then calculated as for Table 1.

6 Conclusion

Even a society in which everybody is the same at the same stage of the life-cycle will exhibit a substantial degree of income and wealth inequality. In this paper we take this notion to the data in order to quantify the share of observed income and wealth inequality that is attributable to life-cycle profiles of income and wealth. To this end we set out by formalizing the concept of life-cycle inequality using a simple life-cycle model. That model highlights that a stylized economy in which everybody is the same at the same stage of the life-cycle will generate Gini coefficients in the order of 0.1 and 0.5 for income and wealth, respectively. We then take this notion to the data – relying on the recently released National Transfer Accounts as a source for internationally comparable data on life-cycle income and wealth profiles. These data reveal that – when compared to actual inequality data – life-cycle factors alone seem to explain a lion’s share of income inequality but much less of wealth inequality. Indeed, deriving a Life-cycle Adjusted Gini coefficient suggests that while there is a substantial amount of wealth inequality in excess of what is to be expected on life-cycle factors alone, actual income inequality tends to be less than what would be expected based solely on life-cycle factors. To home in on the role of the demographic structure for inequality we close our analysis by focusing on the impact of the bulge on the demographic pyramid generated by the baby-boom generation. This shows that once cohort shares settle back into their long-run equilibrium levels, natural inequalities of income and wealth will increase. With the former increasing more so than the latter.

Our analysis speaks mainly to two literatures. First of all, by building on the work of Atkinson (1971) and Paglin (1975) we re-emphasize the central role that demographic structure has for the determination of inequality. Importantly, we add to this literature by quantifying the magnitude of life-cycle inequality for a collection of countries around the world. Moreover, we contribute to this literature by developing a Life-cycle Adjusted Gini coefficient. This coefficient takes into account the various caveats raised against Paglin's measure but also improves upon them by showing that it is both theoretically and empirically possible for there to be life-cycle inequality in excess of observed levels of inequality. Second, we advance on the literature dealing with long-run trends in inequality initiated by Atkinson and Harrison (1978) and reinvigorated by Piketty (2003), Piketty and Saez (2014), Roine and Waldenström (2015), and Milanovic (2016). In this regard we show that an additional factor contributing to any future rise in income and wealth inequality is the transition of the population pyramid back into its long-run shape.

References

- ALMÅS, I., A. W. CAPPELEN, J. T. LIND, E. SØRENSEN, AND B. TUNGODDEN (2011): "Measuring unfair (in)equality," *Journal of Public Economics*, 95, 488–499.
- ATKINSON, A. B. (1971): "The distribution of wealth and the individual life-cycle," *Oxford Economic Papers*, 23, 239–254.
- (2013): "Wealth and Inheritance in Britain from 1896 to the Present," *CASE Papers*.
- ATKINSON, A. B. AND A. J. HARRISON (1978): *Distribution of personal wealth in Britain*, CUP Archive.
- ATKINSON, A. B. AND T. PIKETTY, eds. (2010): *Top Incomes: A Global Perspective*, no. 9780199286898 in OUP Catalogue, Oxford University Press.
- ATKINSON, A. B., T. PIKETTY, AND E. SAEZ (2011): "Top Incomes in the Long Run of History," *Journal of Economic Literature*, 49, 3–71.
- ATTANASIO, O. P. AND G. WEBER (2010): "Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy," *Journal of Economic Literature*, 48, 693–751.
- BLANCHARD, O.-J. (1985): "Debts, Deficits, and Finite Horizons," *Journal of Political Economy*, 93, 223–247.
- BOUCEKKINE, R., D. DE LA CROIX, AND O. LICANDRO (2002): "Vintage human capital, demographic trends, and endogenous growth," *Journal of Economic Theory*, 104, 340–375.

- BROWNING, M. AND T. F. CROSSLEY (2001): "The life-cycle model of consumption and saving," *Journal of Economic Perspectives*, 3–22.
- CAGETTI, M. AND M. DE NARDI (2006): "Entrepreneurship, Frictions, and Wealth," *Journal of Political Economy*, 114, 835–870.
- (2009): "Estate taxation, entrepreneurship, and wealth," *American Economic Review*, 99, 85–111.
- CASTAÑEDA, A., J. DÍAZ-GIMÉNEZ, AND J.-V. RÍOS-RULL (2003): "Accounting for the U.S. Earnings and Wealth Inequality," *Journal of Political Economy*, 111, 818–857.
- DAVIES, J. B., S. SANDSTRÖM, A. SHORROCKS, AND E. N. WOLFF (2011): "The Level and Distribution of Global Household Wealth," *Economic Journal*, 121, 223–254.
- DELTAS, G. (2003): "The small-sample bias of the Gini coefficient: results and implications for empirical research," *Review of Economics and Statistics*, 85, 226–234.
- FEDERAL RESERVE BOARD (2013): "Survey of Consumer Finances," .
- FORMBY, J. P. AND T. G. SEAKS (1980): "Paglin's Gini Measure of Inequality: A Modification," *American Economic Review*, 70, 479–482.
- GOURINCHAS, P.-O. AND J. A. PARKER (2002): "Consumption over the Life Cycle," *Econometrica*, 70, 47–89.
- GUVENEN, F. (2006): "Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective," *Journal of Monetary Economics*, 53, 1451–1472.
- HANSEN, G. D. (1993): "The cyclical and secular behaviour of the labour input: Comparing efficiency units and hours worked," *Journal of Applied Econometrics*, 8, 71–80.
- HANSEN, G. D. AND S. İMROHOROĞLU (2008): "Consumption over the life cycle: The role of annuities," *Review of Economic Dynamics*, 11, 566–583.
- HEIJDRA, B. J. AND J. O. MIERAU (2012): "The individual life-cycle, annuity market imperfections and economic growth," *Journal of Economic Dynamics and Control*, 36, 876–890.
- HUGGETT, M. (1996): "Wealth distribution in life-cycle economies," *Journal of Monetary Economics*, 38, 469–494.

- HUGGETT, M., G. VENTURA, AND A. YARON (2006): "Human capital and earnings distribution dynamics," *Journal of Monetary Economics*, 53, 265–290.
- (2011): "Sources of lifetime inequality," *American Economic Review*, 101, 2923–2954.
- HUMAN FERTILITY DATABASE (2013): "Human Fertility Database," Available at www.humanfertility.org (data downloaded July 2013). Max Planck Institute for Demographic Research (Germany) and Vienna Institute of Demography (Austria).
- HUMAN MORTALITY DATABASE (2013): "Human Mortality Database." University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany) www.mortality.org or www.humanmortality.de (data downloaded July 2013).
- HURST, E. AND A. LUSARDI (2004): "Liquidity Constraints, Household Wealth, and Entrepreneurship," *Journal of Political Economy*, 112, 319–347.
- KUZNETS, S. AND E. JENKS (1953): *Shares of Upper Income Groups in Income and Savings*, NBER Books, National Bureau of Economic Research, Inc.
- LAMBERT, P. J. AND J. R. ARONSON (1993): "Inequality decomposition analysis and the Gini coefficient revisited," *Economic Journal*, 1221–1227.
- LEE, R. AND A. MASON (2014): "Is low fertility really a problem? Population aging, dependency, and consumption." *Science*, 346, 229–34.
- LIS (2013): LIS Inequality and Poverty Key Figures, <http://www.lisdatacenter.org> (data downloaded July 2013). Luxembourg: LIS.
- LOTKA, A. J. (1998): *Analytical theory of biological populations*, Springer Science & Business Media.
- MELTZER, A. H. AND S. F. RICHARD (1981): "A Rational Theory of the Size of Government," *Journal of Political Economy*, 89, pp. 914–927.
- MIERAU, J. O. AND S. J. TURNOVSKY (2014): "Demography, growth, and inequality," *Economic Theory*, 55, 29–68.
- MILANOVIC, B. (1997): "A simple way to calculate the Gini coefficient, and some implications," *Economics Letters*, 56, 45–49.
- (2016): "Income inequality is cyclical," *Nature*, 537, 479–482.

- MITCHELL, O. S., J. M. POTERBA, M. J. WARSHAWSKY, AND J. R. BROWN (1999): “New evidence on the money’s worth of individual annuities,” *American Economic Review*, 89, 1299–1318.
- MOOKHERJEE, D. AND A. SHORROCKS (1982): “A decomposition analysis of the trend in UK income inequality,” *Economic Journal*, 886–902.
- NEAL, D. AND S. ROSEN (2000): “Theories of the distribution of earnings,” *Handbook of Income Distribution*, 1, 379–427.
- NTA (2013): “National Transfer Accounts Website,” www.ntaccounts.org. Data downloaded July 2013.
- OECD (2014): “OECD.StatExtracts,” Dataset, downloaded from stats.oecd.org June 2014.
- PAGLIN, M. (1975): “The measurement and trend of inequality: a basic revision,” *American Economic Review*, 598–609.
- PIERRE-OLIVIER GOURINCHAS, J. A. P. (2002): “Consumption over the Life Cycle,” *Econometrica*, 70, 47–89.
- PIKETTY, T. (2003): “Income Inequality in France, 1901–1998,” *Journal of Political Economy*, 111, 1004–1042.
- (2011): “On the long-run evolution of inheritance France 1820-2050,” *The Quarterly Journal of Economics*, 126, 1071–1131.
- PIKETTY, T. AND E. SAEZ (2003): “Income Inequality in the United States, 1913–1998,” *The Quarterly Journal of Economics*, 118, 1–39.
- (2014): “Inequality in the long run,” *Science*, 344.
- PIKETTY, T. AND G. ZUCMAN (2015): “Wealth and Inheritance in the Long Run,” in *Handbook of Income Distribution*, vol. 2, 1303–1368.
- QUADRINI, V. (1999): “The importance of entrepreneurship for wealth concentration and mobility,” *Review of Income and Wealth*, 45, 1–19.
- ROINE, J. AND D. WALDENSTRÖM (2015): “Long-Run Trends in the Distribution of Income and Wealth,” in *Handbook of Income Distribution*, vol. 2, 469–592.
- SAEZ, E. AND G. ZUCMAN (2016): “Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data,” *The Quarterly Journal of Economics*, 131, 519–578.

SMEEDING, T. M. (2014): "Adjusting to the fertility bust," *Science*, 346, 163–164.

UN (2013): "National Transfer Accounts Manual: Measuring and analysing the generational economy," .

YAARI, M. E. (1965): "Uncertain lifetime, life insurance, and the theory of the consumer," *The Review of Economic Studies*, 137–150.

Appendix

A Proofs

Proof of Proposition 1. Focusing on income inequality and following Milanovic (1997) we can write the Gini Coefficient of Income as:

$$\theta(W) = \frac{1}{\sqrt{3}} \frac{\sigma_W}{\bar{W}} \rho(W, r_W) \frac{\sqrt{N^2 - 1}}{N} \approx \frac{1}{\sqrt{3}} \frac{\sigma_W}{\bar{W}} \rho(W, r_W),$$

where \bar{W} , σ_W are the mean and standard deviation of individual income W , r_W is the rank of a specific income level W and $\rho(W, r_W)$ is the correlation of W with its rank r_W . To proceed, observe that $\rho(W, r_W) \in [0, 1]$ and that $\rho(W, r_W) = 0$ if and only if $W = \bar{W} \forall W$, otherwise $\rho(W, r_W) \in (0, 1]$. In combination with the fact that $\sigma_W \geq 0$ but also $\sigma_W = 0$ if and only if $W = \bar{W} \forall W$, implies that as long as the set $W \neq \bar{W}$ is non-empty $\theta(W) > 0$. Results for the Gini Coefficient of Wealth can be established with the same arguments. \square

B Detailed Model Derivation

Individuals maximize lifetime utility:

$$\Lambda(v) = \int_v^{v+D} \frac{C(v, t)^{1-1/\sigma} - 1}{1 - 1/\sigma} e^{-\rho(t-v) - M(t-v)} dt \quad (\text{B.1})$$

subject to the budget constraint:

$$A_t(v, t) = (r + \mu(t-v))A(v, t) - C(v, t) + E(t-v)w, \quad (\text{B.2})$$

where we have imposed that the economy is in its steady-state throughout the life-cycle of the individual, as is indicated by the omission of the time indices on r and w . Defining the present value Hamiltonian for an agent born at time v :

$$\mathcal{H} \equiv e^{-\rho(t-v) - M(t-v)} \left\{ \frac{C(v, t)^{1-1/\sigma} - 1}{1 - 1/\sigma} + \lambda(v, t) [(r + \mu(t-v))A(v, t) - C(v, t) + E(t-v)w] \right\}, \quad (\text{B.3})$$

and optimizing with respect to $C(v, t)$ and $A(v, t)$ we obtain:

$$C(v, t)^{-1/\sigma} = \lambda(v, t) \quad (\text{B.4a})$$

$$\rho - \frac{\dot{\lambda}(v, t)}{\lambda(v, t)} = r. \quad (\text{B.4b})$$

Combining the equations in B.4a gives the consumption Euler equation:

$$\frac{\dot{C}(v,t)}{C(v,t)} = \sigma(r - \rho). \quad (\text{B.5})$$

Solving B.5 forward in time from v onward allows us to write consumption at time t as:

$$C(v,t) = C(v,v)e^{\sigma(r-\rho)(t-v)}. \quad (\text{B.6})$$

Solving B.2 forward from v provides:

$$w \int_v^{v+D} E(t-v)e^{-r(\tau-v)+M(\tau-v)} d\tau = \int_v^{v+D} C(v,\tau)e^{-r(\tau-v)+M(\tau-v)} d\tau, \quad (\text{B.7})$$

where we have used the initial and terminal assets (i.e., transversality) condition $A(v,v) = A(v,v + D) = 0$. Substituting B.5 into B.7 provides:

$$C(v,v) = \frac{w \int_v^{v+D} E(t-v)e^{-r(\tau-v)+M(\tau-v)} d\tau}{\int_v^{v+D} e^{-((1-\sigma)r+\sigma\rho)(\tau-v)+M(\tau-v)} d\tau}, \quad (\text{B.8})$$

which, in combination with B.6, allows us to draw out a path for consumption. Alternatively, solving B.2 forward from any other time t provides:

$$A(v,t) = e^{r(t-v)+M(t-v)} w \left(\int_v^t E(t-v)e^{-r(\tau-v)+M(\tau-v)} d\tau - \frac{\int_v^{v+D} E(t-v)e^{-r(\tau-v)+M(\tau-v)} d\tau}{\int_v^{v+D} e^{-((1-\sigma)r+\sigma\rho)(\tau-v)+M(\tau-v)} d\tau} \int_v^t e^{-((1-\sigma)r+\sigma\rho)(\tau-v)+M(\tau-v)} d\tau \right), \quad (\text{B.9})$$

which is the asset path traced out in the figures of the main text.

C Simulations

To simulate the model we need to associate values to the parameters and choose functional forms for the life-cycle income profile and the survival functions. For the parameters we follow Guvenen (2006) and set $\sigma = 0.5$. Not much is known about the values of ρ so we set it equal to 0.01, which, in combination $r = 0.025$ assures that individuals are patient in the sense that they opt for an upward sloping consumption profile. We normalize $w = 1$.

As regards the demographic structure, we use the survival function suggested by Boucekine et al. (2002):

$$S(t-v) \equiv e^{-M(t-v)} = \frac{\mu_0 - e^{\mu_1(t-v)}}{\mu_0 - 1}, \quad (\text{C.1})$$

with $\mu_0 > 1$, $\mu_1 > 0$ and $D = \ln \mu_0 / \mu_1$. To estimate the parameters in C.1 we follow Mierau and Turnovsky (2014) and employ non-linear least squares in combination with survival data from Human Mortality Database (2013). To this end we rewrite C.1 as:

$$S(u) = I(u \leq D) \frac{\mu_0 - e^{\mu_1 u}}{\mu_0 - 1} + \epsilon, \quad (\text{C.2})$$

where $\epsilon \sim i.i.d(0, \sigma^2)$ is the error term, $S(u)$ is the fraction of individuals surviving to age u and $I(u \leq D)$ is an indicator function which takes the term between brackets as logical input. Performing this estimation procedure for the United States 2006 cohort provides $\mu_0 = 78.3618$, $\mu_1 = 0.0566$ and a tight fit ($R^2 = 0.9961$).

Assuming that the demographic structure is in its steady state (see, Lotka (1998)) cohort shares $p(t-v)$ are given by:

$$p(t-v) = \beta e^{\pi(t-v) - M(t-v)}, \quad (\text{C.3})$$

where β is the crude birth rate and π is the population growth rate. For the figures and calculations in the main text we draw population growth rates from the World Bank, which for the US equals $\pi = 1\%$. Through the demographic steady state this implies a crude birth rate of 2.24%. For all other countries we follow the same procedure and the various estimates of the demographic parameters are available on request.

For the life-cycle income profile we follow Blanchard (1985) and employ the sum of two exponential functions:

$$E(t-v) = \alpha_0 e^{-\gamma_0 u(t-v)} - \alpha_1 e^{-\gamma_1 u(t-v)} \quad (\text{C.4})$$

which, under the assumptions $\alpha_0 > \alpha_1 > 0$, $\gamma_0 > \gamma_1 > 0$ and $\alpha_1 \gamma_1 > \alpha_0 \gamma_0$, leads a hump-shaped income profile. We estimate the underlying parameters of C.4 using non-linear least squares using data from Hansen (1993). This provides $\alpha_0 = 4.494$, $\alpha_1 = 4.010$, $\gamma_0 = 0.0231$, $\gamma_1 = 0.050$ and $R^2 = 0.80$.²⁰ Extrapolating the income profile up to D then provides the life-cycle income trajectory displayed in the main text.

D Counter Factual Analysis

To perform the counter factual analysis of Section 5 we use the following procedure. First, for each of the countries in our sample we estimate the Boucekkine et al. (2002) survival function using the method outlined above and employing data from the Human Mortality Database (2013). Second, using the estimates in combination with population growth rate data from the World Bank we generate

²⁰See, Heijdra and Mierau (2012) for details on the estimation procedure.

the cohort-shares using C.3. Finally, we combine the cohort shares with the income and wealth profiles from the NTAs to calculate the ensuing Gini coefficients and concentration measures as reported in the various tables of the main text.

E Properties of the Demographically Adjusted Gini Coefficient


As mentioned in the text, the demographically adjusted Gini coefficient, or excess inequality Gini, may be both positive or negative i.e., $\theta^{DA} \in (-\infty, 1]$. To see this, consider the following two extreme cases. Firstly, that there is no life-cycle inequality and thus $\theta^{NR} = 0$ and perfect overall inequality θ^U , then $\theta^{DA} = \frac{\theta^U - \theta^{NR}}{1 - \theta^{NR}} = \frac{1 - 0}{1 - 0} = 1$. The alternative extreme occurs when overall inequality is 0 and the natural rate of inequality approaches 1. Then, $\theta^{DA} = \lim_{\theta^{NR} \rightarrow 1} \frac{0 - \theta^{NR}}{1 - \theta^{NR}} = -\infty$.

It is possible that the demographic and the unadjusted Lorenz curves cross. That is, Lorenz dominance may be violated. This is not problematic for the calculation of θ^{DA} but it is worth noting that in the same way that the Gini coefficient almost always does not uniquely identify a particular Lorenz curve, θ^{DA} is uninformative about both Lorenz curves. Thus, almost always, there are many pairs of Lorenz curves that exhibit Lorenz dominance for each θ^{DA} and many that do not. This means that in the same way the Gini coefficient is only informative about the overall level of inequality, θ^{DA} only allows us to make statements about the overall level of excess inequality.

In calculating θ^{DA} we must take stock of the fact that θ^U and θ^{NR} may be calculated using different sample sizes. In particular, while θ^U will generally be generated from a larger amount of data, θ^{NR} is calculated using much less data because it is based on age data that is often aggregate at the 5 or 10-year cohort level. As the Gini coefficient is known to be underestimated in small samples (see, Deltas (2003)) we may be led to conclude that θ^{DA} is large simply because θ^{NR} has been underestimated due to small sample problems. To consider the magnitude of this problem consider Figure 5 in which we display the value of the natural Gini of wealth for the US calculated using different sample sizes (the solid line).²¹ As predicted by theory, for small sample sizes the Gini coefficient is substantially smaller than its true value. However, for our current context the bias generated from the underestimated Gini is likely to be small as convergence to the true Gini has already been reached for a sample size of 80. In the figure we also display the Gini coefficient taking into account the small sample adjustment suggested by Deltas (2003) (the dashed line). Clearly, this adjustment works well in the sense for sample sizes just above of 10 the small sample Gini approaches the true Gini very well. This is particularly encouraging for potential future analyses based on cohort data using age groups of 5 or 10 years.

²¹Technically we do this by reducing the length of the vector used to calculate the various simulations results reported in Section 2 of the main text.


Figure 5: Small Sample Properties of the Gini Coefficient



F Life-cycle Consumption Profile

The seeming discrepancy between the consumption profile arising from the NTA data, as in Fig. 3, and the typical, hump-shaped, consumption profile identified in the empirical literature by Pierre-Olivier Gourinchas (2002) amongst others, reflects that we use a simple definition of consumption including education and health expenditure. Whereas, the standard definition is excludes these. The broad definition of consumption fits well with our parsimonious model – in which there is no government, healthcare sector or education – but it is the narrower definition of consumption that reveals the typical hump-shaped profile. The NTA data also contain a measure that corresponds to the narrower definition (i.e., private consumption corrected for government expenditures on health and education) and we display it in Fig. 6. Alongside these data we also display the consumption profile arising from our model if, as discussed in footnote 4 in the main text, we relax our assumption of perfect annuity markets. We impose an annuity market imperfection of the type suggested by Heijdra and Mierau (2012) and Hansen and İmrohoroğlu (2008). Although the NTA profile has a more pronounced hump, both the data profile and the simulated profile suggest a similar – hump-shaped – pattern over the life-cycle, consistent with the standard finding in the literature.

Figure 6: Life-cycle Consumption Profile



G National Transfer Accounts versus the Survey of Consumer Finances

In order to calculate the life-cycle profile we have relied on the assumption that savings behaviour is constant over cohorts so that a flow variable – savings – can be transformed into a stock variable – wealth. This assumption is important and to vitiate it we compare the wealth profile implied by the NTA data to that generated by the Survey of Consumer Finances (Federal Reserve Board (2013)). To this end, we take the raw wealth data from the SCF and then smooth it using the Lowess smoother – one of the alternatives suggested by UN (2013). The results are displayed in Figure 7, where we plot the resulting wealth profile alongside the profile suggested by the NTA data. We see that the life-cycle patterns implied by both data sources are very similar and thus support our assumption.

Figure 7: Comparison of National Transfer Accounts and Survey of Consumer Finance Wealth Profiles

