- Permanyer, I. (2017) 'A new approach to length of life inequality measurement'. Presented in Session 144: Lifespan inequality. 28th International Population Conference of the International Union for the Scientific Study of Population-IUSSP, Cape Town (South Africa), 29 October - 4 November. https://iussp.confex.com/iussp/ipc2017/meetingapp.cgi/Paper/2358

Introduction

During the last decade or so the study of lifespan variability has attracted a great deal of attention from demographers and other social scientists (see, among others, Wilmoth and Horiuchi 1999, Edwards and Tuljapurkar 2005, Smits and Monden 2009, Engelman, Canudas-Romo and Agree 2010, Edwards 2011, Vaupel et al 2011, Nau and Firebaugh 2012, Van Raalte and Caswell 2013, Gillespie et al 2014, Van Raalte et al 2014, Seligman et al 2016). As these studies suggest, variation in length of life is one of the most fundamental inequalities in human populations: living long and healthy lives is among the most highly valued and universal human goals, so the existence of very unequal length of life distributions might go beyond purely natural causes and could be indicative of an unfair state of affairs in which some population groups might be disadvantaged or discriminated against. This paper identifies fundamental problems with currently existing approaches to the conceptualization and measurement of length of life inequality and proposes constructive alternatives overcoming such limitations.

Length of life is a bounded variable. While its upper bound (denoted as ω) is uncertain for the case of human populations that are still alive[1], our impermanence in this planet is, alas, certain. Indeed, prestigious recent studies suggest that the maximum lifespan of humans is fixed and is unlikely to increase over time (Dong, Milholland and Vijg 2016). The uncertainty around the limits of human lifespan has led researchers to implicitly treat that variable as unbounded. While such assumption is a practical way of circumventing the uncomfortable decision of fixing an upper bound to human lifespan, it is not realistic[2]. Instead, we suggest categorizing lifespan as a member of a new class of variables: the set of *uncertainly bounded variables*. This set, which is a proper subset of the class of bounded variables, includes those variables that cannot take arbitrarily high (or low) values but whose bounds are uncertain[3]. In this paper, we discuss the methodological and substantive implications that such categorization has for our understanding of lifespan variation.

The bounded nature of length of life generates two problems when measuring its variability that have not been identified in the literature so far: the so-called (i) ‘boundary effects’ and (ii) the ‘(in)consistency’ problems. By ‘boundary effects’ we refer to the clustering that takes place across observations when the mean of the distribution converges towards some of its bounds.

---

[1] Clearly, such uncertainty disappears when one analyzes lifespan distributions among extinct populations or among non-human populations with well-defined lifespans.

[2] Pervasive as it has been through human history, the quest for *true* immortality still belongs to the realms of religion or science-fiction literature.

[3] Examples belonging to that class would include variables measuring the time of occurrence of any event that is certain to happen in an uncertain future.

When this happens, the corresponding inequality levels mechanically go to zero – simply because there is no room for further variation – an issue that complicates comparisons of the levels of inequality for distributions with different means. Indeed, virtually all empirical studies report decreasing inequality in overall lifespan variation as life expectancy increases over time (e.g. Smits and Monden 2009, Engelman, Canudas-Romo and Agree 2010, Edwards 2011, Vaupel et al 2011). Given the bounded nature of the age at death variable, one is left pondering whether those convergence results are, so to say, purely mechanically driven by the fact that life expectancy might be approaching its upper limit. In these circumstances, it is not clear that studying lifespan inequality can provide new insights above what we already know from studying the values of life expectancy alone.

The second problem discussed in this paper arises because of the dual representation one can make of bounded variables. If length of life is bounded between 0 and ω, it is a priori equally plausible to focus on the number of years of 'lived life' (x) or on the number of years of 'unlived life' ($\omega - x$, i.e. length of time between age at death and maximal lifespan). Thus, length of life distributions can be represented as we typically describe them (i.e. as achievement variables X, measuring how long we have lived) or as shortfalls with respect to the upper bound (i.e. as shortfall variables $\omega - X$, measuring how many more years we could have potentially lived). While unorthodox, the latter representation mirrors the former, so it conveys exactly the same information as its achievement counterpart. Curiously, both approaches can lead to non-concordant results for many of the inequality measures traditionally used in the literature on lifespan variation[4] (that is: population A might exhibit higher or lower lifespan variability than population B depending on whether we compare achievement or shortfall distributions). This is the so-called '(in)consistency problem', which, far from being a mere academic curiosity, poses several practical challenges to the study of inequality of bounded variables (Clarke et al 2002, Erreygers 2009, Lambert and Zheng 2011) – indeed, its presence is particularly acute for the age at death distributions we will be working with in the empirical section of this paper.

While traditional absolute inequality indices[5] are very sensitive to the boundary effects but immune from the (in)consistency problem (see Erreygers 2009 and Lambert and Zheng 2011), relative inequality measures[6] are equally affected by both. The main goal of this paper is to propose new inequality measures that overcome the aforementioned problems simultaneously by explicitly taking into account the (uncertainly) bounded nature of lifespan in their formulation. The new measures will be referred to as 'benchmark inequality indices' for reasons that will become apparent shortly. Given the dependency of our benchmark inequality measures to the choice of the upper bound and the uncertainty surrounding that value, we develop simple tests to investigate the robustness of any statement one might want to make as regards the inequality of length of life distributions to alternative values of ω. The rest of the paper is organized as follows. After introducing some basic notations and reviewing traditional inequality measures in section 2, section 3 will introduce our new proposal. In that section we will also introduce tests to investigate the robustness of our measures to alternative choices of

---

[4] According to the work of Erreygers (2009) and Lambert and Zheng (2011), all relative inequality measures (i.e. those whose values remain unaffected when all its arguments are scaled by the same proportionality factor – see below) can give non-concordant results when comparing achievement and shortfall distributions.

[5] This includes the standard deviation, the variance or the absolute Gini index.

[6] This includes the generalized entropy measures (i.e. the Theil or the Mean-Log-Deviation), the relative Gini index, the coefficient of variation or the Atkinson index.

lifespan's upper bound. In section 4 we explore how these measures perform empirically using the period life tables from the Human Mortality Database (HMD) and comparing them with respect to other inequality indices currently used in literature of length of life inequality. Contrary to previous studies, our findings suggest that beyond a certain longevity threshold length of life inequality reductions stall and, in some particular cases, bounce upwards. We conclude in section 5 with some discussion and remarks.

Summary and concluding remarks

This paper makes two important contributions to the bourgeoning field of length of life inequality measurement. One of them is methodological and the other empirical. Methodologically, we have conceptualized length of life as an uncertainly bounded variable – rather than an unbounded one, as is implicitly the case in currently existing approaches – and explored the consequences of such decision when it comes to measure the inequality in its distribution across individuals. Being a bounded variable, the measurement of its variability can be affected by two important problems: (i) the 'boundary effects' (i.e. the inequality of the distribution is strongly related to the mean when the latter approaches the upper bound); and (ii) the '(in)consistency' problems (i.e. achievement and shortfall distributions can be inconsistently ranked by standard inequality measures). In this context we introduce the so-called benchmark inequality measures, which compare observed inequality levels with respect to the ones that would be observed under a hypothetical distribution with the same mean that maximized inequality. The new class of measures – which is unaffected by the aforementioned problems – takes into account the bounded nature of length of life explicitly. Interestingly, benchmark inequality indices can be seen as the natural counterpart of a relative inequality measure in the context of bounded variables (indeed, when the upper bound of lifespan is allowed to be arbitrarily large our benchmark inequality indices converge towards classical relative inequality indices, so the latter can be seen as a particular case of the former). Hopefully, this can be a useful addition to the practitioner's toolkit that can complement currently existing methods to analyze lifespan variability.

The new inequality indices introduced in this paper explicitly depend on the choice of the upper bound to human lifespan (ω). Since this upper bound is highly uncertain there might be concerns regarding the reliability of the findings that make use of such measures. In order to overcome such problem we have developed simple tests to investigate the extent to which certain comparisons between pairs of length of life distributions are robust to alternative specifications of ω.

On the empirical side, we have investigated how the new benchmark inequality indices behave using data from the Human Mortality Database. Contrary to previous findings, we observe that the variability in the full mortality distribution does appear to have a strictly positive floor when life expectancy reaches a certain threshold. That is: rather than converging to zero (as the extreme version of Fries' (1980) compression-rectangularization hypothesis would predict), the lifespan inequality declines we have been observing during the last years (e.g. Monden and Smits 2009, Vaupel et al 2011) seem to reach a previously unobserved plateau well above zero as life expectancy goes beyond 70 years. Indeed, for a certain group of countries we even observe trend *reversals*, that is: at high longevity levels further increases in life expectancy increase length of life inequality as well (this is the case of Japan, Taiwan and Luxembourg for the case of women and Ireland, Taiwan, the United States and the United Kingdom for the case of men). Even if our findings are no doubt influenced by the choice of the lifespan upper bound

(ω), they are robust enough to consistently point to the emergence of a length of life inequality plateau at higher longevity levels.

How do our substantive findings relate to other studies? The plateauing and reversal of lifespan variability as longevity increases is reminiscent of the findings reported by Engelman et al (2010), who find that length of life inequality is increasing *among the elderly* due to the growing heterogeneity in old-age mortality. These authors suggest that as health improvements delay mortality an increasingly heterogeneous population is reaching every age and health disparities in early life are delayed and manifest themselves in mortality variation at increasingly older ages. According to the benchmark inequality indices introduced here, it is lifespan variability *for the entire population* that might be on the rise at higher longevity levels. Indeed, very recent studies carried out in high-income countries indicate that the causes of death that contributed most to declines in the variance are different from those that contributed most to increase in life expectancy (Seligman et al 2016), an issue that might rise trade-offs between equality and overall efficiency. If health improvements contribute to increase longevity *and* overall lifespan variability simultaneously, health care systems would be facing a difficult ethical dilemma (particularly in high-income countries) upon which it will be necessary to reflect.

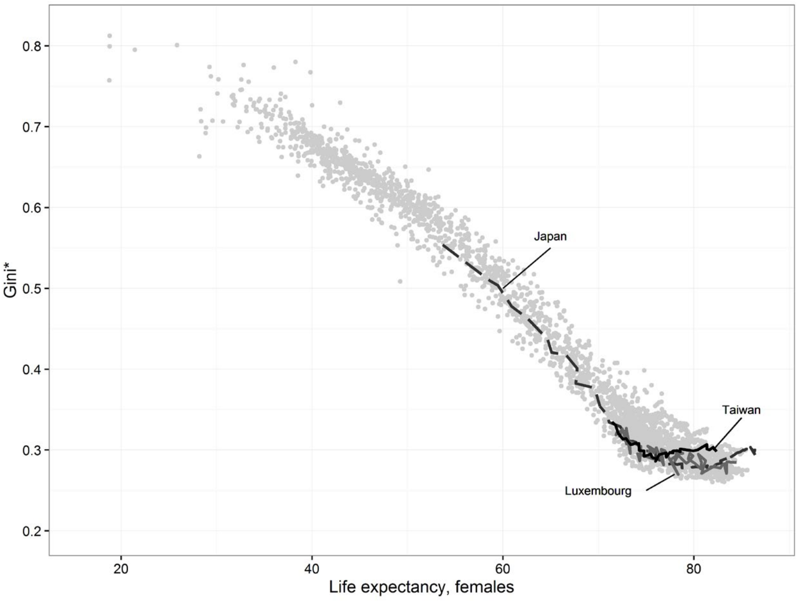

Figure 6. Length of life inequality as measured with $G_a^*(\boldsymbol{x}, 122)$ by life expectancy for women based on 3331 period life tables from the HMD.

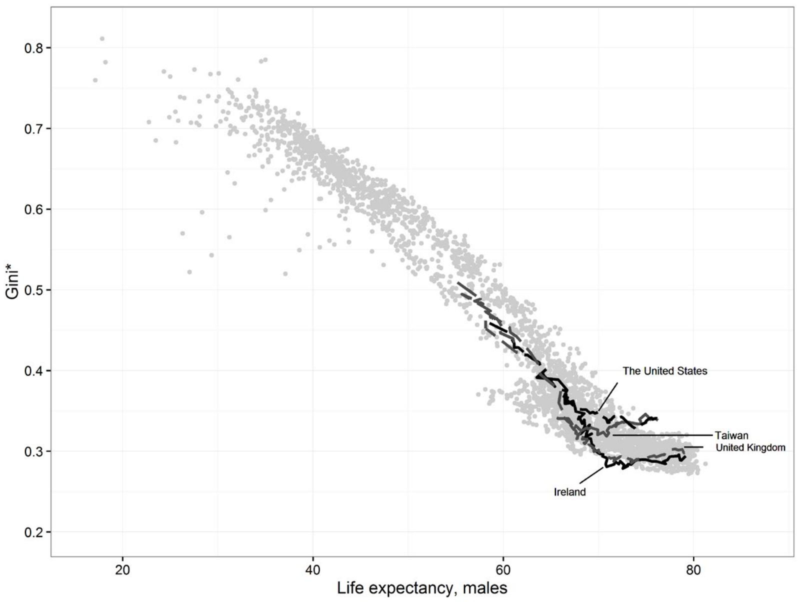

Figure 7. Length of life inequality as measured with $G_a^*(\boldsymbol{x}, 122)$ by life expectancy for men based on 3331 period life tables from the HMD.