Swimming with Wealthy Sharks: Longevity, Volatility and the Value of Risk Pooling

Moshe A. Milevsky[1]

12 August 2018

[1]Milevsky is a professor of finance at the Schulich School of Business and can be reached at: milevsky@yorku.ca, or at Tel: 416-736-2100 x 66014. His address is: 4700 Keele Street, Toronto, Ontario, Canada, M3J 1P3. The author would like to acknowledge and thank Michael Stepner for access and help with the U.S. mortality data, research assistance from Victor Le, Dahlia Milevsky and Yossef Bisk, editorial assistance from Alexandra Macqueen, kind comments (and some corrections) from Natalia Gavrilova & Leonid Gavrilov, and many conversations with Tom Salisbury and Huaxiong Huang.

Abstract

Swimming with Wealthy Sharks:
Longevity, Volatility and the Value of Risk Pooling

Who *values* life annuities more? Is it the healthy retiree who expects to live long and might become a centenarian, or is the unhealthy retiree with a short life expectancy more likely to appreciate the pooling of longevity risk? Conventional economic wisdom assumes that higher mortality rates lead to an increase in the annuity equivalent wealth (AEW), which is the standard metric used to subjectively value annuities.

In this paper I offer a different perspective and argue that the value of pooling originates from the *individual volatility of longevity*, which is defined as the ratio of the standard deviation to the mean of life. Simply stated, consumers who face greater uncertainty in their longevity, subjectively value annuities the most and should own more. A healthy (and wealthy) retiree who expects to live 30 years but with a longevity volatility of 25%, doesn't receive as much utility from annuities as someone who expects to live only 10 years, but with a volatility of 50%.

All this assumes annuities are fairly priced, tailored to individual mortality or that retirees *swim within their species.* In contrast, when different risk-types are mixed together in a large pool – such as with mandatory social security – but all participants are forced to pay the same price, the situation is more nuanced and complex. I demonstrate how the so-called *Compensation Law of Mortality*, which implies that individuals with higher mortality (e.g. lower income) experience greater volatility of longevity, leads to pooling benefits for both high and low risk types.

In sum, this paper describes the conditions under which retirees benefit (or may not) from longevity risk pooling by linking AEWs to the biology of aging. The impetus for this research today, versus 50 years ago, is the growing evidence on the disparity in longevity expectations between rich and poor, especially in the U.S.

JEL: H55 (Pensions), G22 (Insurance)
Keywords: Annuities, Retirement, Utility, Social Security

There's nothing serious in mortality,
all is but toys, renown and grace is dead.
Macbeth, Act II, Scene III

1 Background and Motivation

The noted Princeton economist and Nobel laureate Angus Deaton recently wrote that: "The finding that income predicts mortality has a long history," having been noted as far back as the 19th century by Friedrich Engels in Manchester, England. According to Deaton (2016), commenting on similar findings by Chetty et al. (2016), "There is little surprising in yet another study that shows that those with higher income can expect to live as much as 15 years more than those with lower income." It simply isn't news. Indeed, the focus among those (economists) who study *mortality and its inequality*, using a phrase coined by Peltzman (2009), has shifted to the causes and consequences as opposed to proving its existence. The question en vogue is: *Why is the mortality gradient steepening?*

But what has received less attention from economists – and in fact may be surprising – is that not only is the life expectancy or longevity of those at the lowest income percentiles in the U.S. lower, the uncertainty or *risk* of their remaining lifetime is higher. It is the exact opposite of the well-known relationship in portfolio theory. If one thinks of the conditional random life: T_x in terms of return (its expected length) and variability, the mean is lower but the standard deviation in higher. In fact, the positive (cor)relation between mortality rates and the volatility of longevity follows from the (so called) *Compensation Law of Mortality* (CLM), a phenomena in the bio-gerontology field, introduced by Gavrilov and Gavrilova in (1991). I'll elaborate on CLM soon enough.

More than a statistical curiosity or something to idly puzzle over, nowhere is the natural link between life expectancy and it's volatility more pertinent than in the area of pensions, retirement planning and (subjective) annuity valuation.

1.1 Pension Subsidies

For the sake of a wider readership (but at the acknowledged risk of alienating an academic audience) I'll dispense with tradition and motivate this paper with a very simple example. Assume that Mrs. Heather White is about to retire at the age of 65 and is now entitled to a pension annuity, a.k.a. social security, or guaranteed lifetime income of exactly \$25,000 per year, paid monthly. For the record, this is the legislated maximum she can receive after having worked the requisite number of years. In other words, she has also contributed the maximum, perhaps explicitly by having a fraction of her paycheck withheld or implicitly via the income tax system. The pension annuity payments are adjusted for the cost of living or price inflation as measured by some national index, but the income will cease upon her death. The pension annuity contains no cash value or liquidity provisions, nor can she bequeath the income to her children, grandchildren or loved ones. I'm not describing any one specific country or government plan, but rather a generic, no-frills defined benefit (DB) pension scheme managed by any large sponsor.

Coincidentally, her next-door neighbor Mr. Simon Black was born in the same year, is also about to retire at 65 and is entitled to the same annuity of $25,000 per year. He, too, has contributed the maximum to the scheme. The exact details of how Simon and Heather paid for their pension annuity are unimportant at this juncture. The key is that over his working life – and in particular by the time he retires – both Simon and Heather have contributed fully to the pension system.

There is one important difference between the two, though. Regrettably, Simon has a medical life expectancy of 10 years and is rather sickly, whereas Heather is in perfect health with a corresponding life expectancy of 30 years – and they both know it. Heather is expected live to age 95 (from her current age of 65) and outlive Simon who will only make it to 75. Stated differently, although their chronological ages (CA) are both 65, Heather's biological age (BA) is much lower than Simon's BA. In actuarial terms, his mortality hazard rate is substantially higher than Heather's.

And yet, despite Simon's poor health and the financial fact he contributed the exact same amount to the retirement program or pension scheme, he isn't entitled to any more income from his pension scheme's annuity compared to Heather. In the language of insurance, retirement programs aren't underwritten or adjusted for individual health status. Instead, *all* government social security programs around the world are unisex or gender neutral. If you contribute (the same amount) into the system, you receive the same benefit regardless of your sex, health status or any other bio-marker for longevity.

Obviously, Simon's shorter life expectancy of ten years implies that he will be receiving (much) less money back compared to Heather. Moreover, if the retirement program or scheme is designed to neither make or lose money in the long run, a.k.a. it is actuarially balanced, then the (sick) Simons of the world are subsidizing the (healthy) Heathers. This isn't news to economists, who know this very well and is the nature of all government pension programs. In fact, in most of Europe today, insurance companies are prohibited from using gender to price *any* type of retail policy, whether it be life, health, home or even car insurance. (So, perhaps Heather has to pay a bit more for car insurance, in Europe, relative to Simon.)

My objective in introducing this very simple framework or toy story, is to quantify the magnitude of the financial subsidy from Simon to Heather, one that will form the basis and intuition for what follows later. Of course, despite the large dollar value of the transfer, the entire point of this paper is to argue why (and quantify how) Simon might still benefit from being a member of a pension scheme due to the fact that his longevity risk is greater. Yes, he isn't expected to live as long, but his *volatility of longevity is greater.*

To properly analyze the subsidy from a financial perspective, the natural next step is as follows: How much would Simon have to pay in the open (a.k.a. retail) market to acquire a pension annuity of $25,000 per year, and how much would Heather have to pay? That price or cost should give a rough sense of the magnitude of the transfer provided by Simon to Heather.

Now, this can get tricky. In practice, the market price will depend on many factors, such as commissions and profits and, more importantly, the magnitude of the uncertainty around Simon and Heather's life expectancy. But to keep things very simple (at this early intuitive stage) I'll assume that Simon will live for exactly 10 years and Heather will live for exactly 30 years. In other words, it's not an *average.* Remaining lifetimes are deterministic.

Based on these numbers, Simon would be charged \$212,750 at the chronological age of 65 and Heather would be charged \$487,225 for the same exact pension (actually, term-certain) annuity. These numbers are based on a 3% effective (real) annual interest rate but do not require much else in terms of assumptions or parameters.

Stated differently, the present value of \$25,000 per year for 10 years is exactly \$212,750, for 30 years of cash flows, the present value is \$487,225. Moreover, the market cost of their combined pension annuity entitlement is \$212,750 + \$487,225 = \$700,000, a very important number from a funding and pension solvency perspective. So, if – and this is a big if – the pension scheme is actuarially balanced, it should have exactly \$700,000 set aside in reserves to pay pension annuities when Heather and Simon both retire[1].

To be clear, real-world insurance companies will not charge \$212,750 and \$487,225 to Simon and Heather for these annuities. First and foremost, these companies have to make profits, so they would mark up or "load" the price, just like the retail vs. wholesale cost of coffee. More importantly, insurance companies have to budget and provision for uncertainty, including the risk of how long their annuitants will live. I'll get to more refined mortality models in section #3. For now I'll simply allude to Jensen's inequality and note that if life expectancy is $E[T]$ years, the pension annuity will cost more than the PV of a deterministic annuity to $E[T]$.

The relevant upshot is as follows: Simon is transferring \$137,225 to Heather – a subsidy amounting to 64.5% of the hypothetical value of Simon's pension pot. Where did this number come from? Again, the entire system should have \$700,000 set aside for both of them, of which \$487,225 is needed for Heather and only \$212,750 is required for Simon. And yet, by my construction, they both contributed the same amount of money to the pension scheme, which presumably is a total of \$350,000 x 2 = \$700,000 over the course of their lives. Not to hit readers over the head with this, but Simon contributed \$350,000 and is getting something worth \$137,225 *less* in the open market. Heather contributed \$350,000 and is getting something worth \$137,225 *more* in the open market.

Yes, these are all straw men and women. The numbers assume a pension system with no spousal or survivor benefits and an (extreme) 20-year gap in life expectancy between the two. More importantly the simple tale assumes that Simon and Heather die precisely at their life expectancy, which presumes the absence of any longevity risk or uncertainty. In fact, Simon might live beyond age 75 (or he may die even earlier) and Heather might not make it to age 95 (or she may live even longer). Under these *ex-post* outcomes the cross subsidy will be smaller (or perhaps even larger).

But, the *ex-ante* reality is that there is a large gap between the expected present value of the benefits they receive even though they have paid similar amounts into the retirement program. Whenever you mix (a.k.a. pool) heterogenous people with different longevity prospects into one scheme in which everyone gets a pension annuity for the rest of their life, there will be winners and losers *ex ante* as well as *ex post*. This outcome is well-known in the pension and insurance economics literature[2], but it is often surprising to non-specialists.

[1] Spoiler alert: Few schemes have anywhere near \$700,000 set aside to pay all the guaranteed pension annuities under reasonable discount rates. For the most prominent voice in this area, see Joshua Rauh.

[2] This is a concern with Notional Defined Contribution (NDC) schemes, as noted recently by Holzman et al. (2017)

1.2 Enter Longevity Risk

What happens if we incorporate longevity risk or horizon uncertainty into our straw neighbors? Well, as I noted earlier there is a small probability that Simon lives for more than 10 years beyond age 75 and/or that Heather dies before age 95. Nobody really knows exactly how long they are going to live *ex ante*. In that case the *ex-post* transfer of wealth from Simon to Heather was less than 60% of the value of his pension annuity. At the extreme edge of your imagination there is a (very) small probability that Simon actually outlives Heather and the *ex-post* transfer goes in reverse; she subsidizes him. We won't know until all the Heathers and Simons are dead.

Here is the main economic point. The pension annuity they are entitled to for the rest of their life provides them with more than just a periodic cash flow or income, it provides longevity insurance. Moreover, the value or benefit of any type of insurance can't be quantified in terms of what might happen *on average.* It must account for the so-called tails of the distribution, which is best measured via (some sort of) discounted expected utility. None of this should come as a surprise to insurance economists; the only debate is magnitude.

We now must parameterize the utility value of risk mitigation. Back to Heather and Simon. As mentioned, their income annuity entitles them to more than a term-certain annuity for 30 and 10 years respectively – they have acquired longevity insurance that protects them in the event they live longer than average. Sure, Simon would rather be *pooled* with people like him who share the same risk profile, as he would then expect a "more equal" distribution. But even Simon is willing to be pooled with Heathers – if the alternative is no pooling at all.

It's time to get technical. Who values this insurance more? Heather or Simon? Or is the insurance benefit symmetric? Stated differently, could Simon be gaining more (in utility) from pooling with Heather, even if he is losing on an expected present value basis? The answer is yes, Simon could be winning (economic utility) even if he appears to be losing (dollars and cents). Why? In a nutshell, his *volatility of longevity* is greater. Simon has a short life expectancy, but relatively speaking the range of how long he might actually live, expressed as a percentage, is actually greater than Heather's.

Think about it: If Simon lives 30 years instead of 10 years, that is equal to a 300% (of mean lifetime) shock. It's unlikely, but in the realm of possibility. In contrast, Heather who is expected to live 30 years will never experience a 300% shock. This would implies she lives 90 more years (from age 65) to the age of 155. It simply won't happen. The odds are zero. Ergo, Simon's *individual volatility of longevity*, that is in relative percentage terms, is higher than Heather's. That's biostatistics and displayed in Figure (#1), with much more explanation to come in Section (#4). From an insurance economics point of view, it implies that Simon values the risk pooling benefits of the pension more than Heather does. The volatility of which I speak (and model) is more subtle than the likelihood of living 300% longer than her current life expectancy. It will be defined precisely in section (#3.)

Most pension economists know that the transfer from Simon to Heather isn't as large as the expected dollars indicate, because (using my term) Simon's *volatility of longevity* is larger than Heather's. He places greater value on the insurance component. More importantly, he is willing to swim in a *pool* with Heather, rather than taking his longevity chances.

Of course, Simon is a mere euphemism for all the unhealthy males or individuals who retire and aren't expected to live very long, whereas Heather represents retirees with very long life expectancies. At this point I should make it clear that this isn't just a question of sex or gender. There is a growing body of evidence that we can identify the Simons of the world *ex ante* (i.e. in advance, not after they die young) based on the size of their wallets and magnitude of their income. Genetic testing is yet another pathway to gleaning this knowledge well before you reach your chronological 60s.

Economists have long-known most of this, or that forced risk pooling can (still) benefit everyone when measured in units of utility. What is new? Well, a recent article by Chetty et al. (2016) has documented a growing and increasing gap in life expectancy between U.S. taxpayers in the lowest income percentile versus the highest income percentile. For example, at the chronological age of 50, taxpayers in the lowest income percentile (have much higher mortality rates and) are expected to live 15 years less than taxpayers in the highest income percentile. And yet, they are all forced to participate in the same (mandatory) social security program. Using the Chetty et al. (2016) mortality data (which I'll explain) I am calibrate the extent of the transfer via the so-called *annuity equivalent wealth.* Bottom line: I show that even the lowest income percentiles in the U.S., do still benefit from pooling because their volatility of longevity is large enough to overcome the implicit loading that comes from pooling, *for the time being...*

1.3 Outline of the Paper

The remainder of this paper is organized (and presented in a more academic tradition) as follows. In the next section #2 I provide a proper literature review, linking the current paper to prior work (and interest) in the field. Then, in section #3, I provide analytic context to the *volatility of longevity*, by introducing the Gompertz law of mortality as well as the so-called *Compensation Law of Mortality.* Section #4 provides an expression for the *annuity equivalent wealth* (AEW), which is another way of presenting the *willingness to pay* (WtP) for longevity insurance. I illustrates exactly how it is a function of mortality characteristics. Section #5 displays the AEW as a function of income percentiles in the U.S. Section #6 concludes the paper with the main takeaways. All 5 tables and 7 figures are located at the very end of this documents, and all technical derivations are relegated to an appendix.

2 Scholarly Literature Review

This paper sits squarely within the so-called *annuity economics* literature, which – broadly speaking – attempts to model and explain the demand, or lack thereof, for insurance products that hedge personal longevity risk. Life annuities are an important form of retirement income insurance, very similar to Defined Benefit (DB) pensions, as explained and advocated by Bodie (1990) for example. This literature began close to 50 years ago with the 1965 article by Menachem Yaari, in which he extended the standard deterministic lifecycle model to include actuarial notes (a.k.a. life annuities.) That paper, which has been cited over 3000 times according to Google Scholar, is the economic "workhorse" for all lifecycle models of investment, consumption and retirement planning.

Yaari (1965) proved that for those consumers with no bequest motive, the optimal lifecycle strategy is to annuitize 100% of assets. Clearly, few people have 100% of their wealth annuitized (or "pensionized") and even fewer actively purchase annuities, as pointed out by Franco Modigiliani in his 1986 Nobel Prize address.

The restrictive conditions in the original Yaari (1965) model were relaxed by Davidoff, Brown and Diamond (2005), and still the important role of annuities prevail. In fact, to quote the recent paper by Reichling and Smetters (2015), "The case for 100% annuitization of wealth is even more robust than commonly appreciated" and it takes quite a bit of modeling effort to "break" the Yaari (1965) result. Of course, including bequest and altruistic motives will reduce the 100% annuitization result because annuity income dies with the annuitant and there is no legacy value. In a comprehensive review and modeling effort, Pashchenko (2013) pinpoints the extent to which bequest motives, pre-annuitized wealth and impediments to small annuity purchases can deter the full annuitization result. The (negative) impact of bequest motives is also echoed in work by Inkman, Lopes and Michaelidis (2010), who interestingly find a positive relationship between annuity market participation and financial education.

Reichling and Smetters (2015) succeed in "cracking" the 50-year-old model by introducing stochastic mortality rates, in which the present value of the annuity is correlated with medical costs. According to them, although the annuity helps in protecting against the impact of longevity risk, its economic value is reduced in states of nature that are most costly to the retiree – namely in the case of medical emergencies. In that state of nature annuities aren't as desirable; and as a result fewer people (than previously thought) should be acquiring any more annuities (a result which has been received with some controversy).

Other attempts to "break" the Yaari (1965) model revolve around the underlying (additive) lifecycle model and moving-away from the implied risk neutrality over the length of life, towards a model with recursive preferences. See Bommier (2006), who questions the Yaari (1965) framework as being synonymous with risk neutrality over the length of lifetimes.

To be clear, I don't aim for another crack in the Yaari (1965) model, or provide reasons for why consumers don't annuitize. The recent work in behavioral economics, specifically the article by Brown et al. (2008), provides a rather convincing explanation for why consumers dislike annuities; having to do with framing, anchoring, loss aversion and the usual culprits. The current paper stays well within the neoclassical paradigm, assuming that consumers are rational, risk-averse and maximizing an additive utility of instantaneous consumption over a stochastic life horizon. This is the approach taken by Levhari and Mirman (1977), Davies (1981), Sheshinski (2007), or more recently Hosseini (2015), to name just a few. Moreover, I assume the consumer values annuities using the *annuity equivalent wealth* (AEW) metric introduced by Kotlikoff and Spivak (1981), also used by Brown (2001) and others who have calibrated these numbers around the world. The AEW is just another (reciprocal) way of presenting the *willingness-to-pay* metric, which is widely used in economics and recently reviewed in Barseghyan et al. (2018). Brown (2001) showed that an increase in the individual's AEW leads to an increase in the propensity to annuitize. It partially predicts who is likely to buy an annuity,[3] which is yet another reason to dig deeper in the structure of AEW.

[3]See also Brown, Mitchell, Poterba and Warshawsky (2001) and the link to money's worth ratio.

That said, the main focus of attention in the current paper – as alluded to in the motivating introduction – has to do with mortality *heterogeneity* and the subjective or personal value of annuities when everyone is forced to pay the same price, i.e., they all must swim in the same pool.

Evidence of increasing mortality inequality continues to accumulate, and in particular the recent work by Chetty at al. (2016) indicates that the gap in expected longevity between the highest and lowest income percentiles in the U.S. can be as much as 15 years. These numbers are greater than (say) the 10 years reported by De Nardi, French and Jones (2009), or the 5-year gap noted in Poterba (2014, table 3) within the context of pensions and social security. Peltzman (2009) notes that in the year 2002 U.S. life expectancy (by county) at the highest decline was 79.83 years, and at the lowest decile was 73.17 years, a gap of less than seven years. It's large, but Chetty et al. (2016) indicate that the measurable gap can extend to 15 years. Echoing the same trend, within the context of social security, Goldman and Orszag (2014) discuss and confirm the "growing mortality gradient by income" and report a life expectancy gap of 13 years between those in the lowest versus highest average indexed monthly earning (AIME). It's worth noting that the correlation between (lower) income and (higher) mortality isn't only a U.S. phenomenon. It is discussed and reviewed in Andersson, Lundborg and Vikstrom (2015) within the context of Sweden, for example, where one wouldn't expect to observe such a mortality gradient. The question here is: how does this growing heterogeneity in mortality and longevity affect the value of annuitization?

Using numbers available in the late 20th century, Brown (2003) manufactures annuity prices and mortality tables based on race and education and concludes that "complete annuitization is welfare enhancing *even* for those with higher than average mortality, provided administrative costs are sufficiently low." This result was echoed (and cited by) Diamond (2004) in his presidential address to the *American Economics Association*. He starts by noting that "uniform annuitization would favor those with longer expected lives [such as] high earners relative to low earners" but concludes that Brown (2003) "shows much less diversity in the utility value of annuitization than previous comparisons."

And yet, the range of life expectancy between healthy and unhealthy in the Brown (2003) analysis (Table 1, to be specific) was a mere 3.4 years at the age of 67. He assumed a conditional life expectancy of 81.0 years (lowest) for male Blacks with less than a high school diploma vs 84.8 years (highest) for male Hispanics in the U.S. Contrast these differences with the more recent and granular numbers provided by Chetty et al. (2016), or even Goldman and Orszag (2014), where the gap in life expectancy between the highest and lowest income percentile (calibrated to the same age of 67) is more than 10 to 15 years. It's unclear whether the uniformly positive *willingness to pay* for insurance values, i.e. *annuity equivalent wealth* values greater than one, can survive such a large gap in longevity expectations. The *annuity equivalent wealth* might be lower than *endowed* wealth and the value of longevity risk pooling might be negative when retirees with low life expectancy are forcefully *pooled*, that is required to swim with individuals who are expected to live much longer. It's easy to construct such (mathematical) counterexamples.

Can one say unequivocally that no matter how high Simon's mortality rate is, relative to Heather's, that he is willing to be pooled and benefits from annuitization? Surely there must be a point at which the answer is no because the implicit loading is (too) high.

My point here isn't simply to argue for an update or revision of possibly stale AEW numbers to reflect the increasing heterogeneity of mortality (although that is part of the agenda of the paper.) Rather, my primary objective is to focus attention on the *individual volatility of longevity* as the main driver of the value of annuitization. I do this by presenting a simple (closed-form) analytic expression for computing the AEW and then using mortality rates by percentile, from the Chetty et al. (2016) data, as a calibration exercise.

Conceptually the main argument of this paper can be summarized as follows. Although recent data indicates the heterogeneity of mortality is increasing and the gap in life expectancy is increasing, the same data show that individuals with higher mortality experience a higher volatility of longevity. This then serves to increase the value of longevity insurance. In other words, it isn't the shorter expected longevity (or higher mortality rate) that makes annuities appealing. Rather, it's the volatility of longevity that drives its value. More on this will be provided in the body of the paper.

To be clear, there are a number of other authors and papers that have focused attention and made the link between the *standard deviation* (SD) of lifespans and optimal lifecycle behavior. Most prominent in this category would be Edwards (2013), building on the work of Tuljapurkar and Edwards (2011), who document a 15 year standard deviation at the age of $x = 10$. Edwards (2013) builds on the Yaari (1965) model and arrives at estimates for the increased longevity that a rational lifecycle consumer would demand in exchange for being exposed to a higher variance of life. Although much of Edwards (2011) is based on normally distributed lifespans (and I work within a Gompertz framework), he shows that one additional year of standard deviation (in years) is "worth" about six months of life. Needless to say, Mother Nature doesn't compensate individuals according to the Yaari (1965) model and retirees with lower life expectancies have (not only) a higher individual volatility of longevity, but a higher standard deviation of life as well. Nevertheless, as far as the literature review is concerned, Edwards' is one of the few papers to focus economists' attention on the second (versus the first) moment of life and show *how exactly* it affects optimal behavior.

3 Matters of Life and Death

3.1 Mortality by Gender and Income

Table (#1) displays mortality rates for males and females in the U.S. as a function of various ages (columns) and income percentile (rows). These numbers represent realized mortality rates per 1,000 people during the period 2001 to 2014 and are extracted from the data collected by Chetty et al. (2016). The phrase *realized mortality* means that the numbers provided are the actual ratio of observed deaths at a given age (say age 50) divided by the total number of people at that age (say 50). In other words, these rates aren't projections, forecasts or predictions and are based on over 1.4 billion (yes, with a b) person-year observations and close to 6.7 million deaths. The methodology is described in the article by Chetty et al. (2016), and the entire dataset (of mortality rates as a function of income percentile) is available online. Their (lagged 2-year income) numbers are for ages 40 to 63, and one must employ forecasting procedures to obtain values in later life.

These mortality rates in table (#1) contain various insights or takeaways, some immediately obvious and intuitive and some (much) less so. Focus first on the middle row, with the so-called median mortality rates. At the age of 40, a total of 1.2 per 1,000 (median income) males died, whereas for (median income) females the rate was only 0.8 per 1,000 individuals. Stated differently, the one-year mortality rate for (median income) males at the age of 40 is 50% higher than it is for females, which naturally leads to a lower life expectancy for (median income) males. Continuing along the same row, at the age of 50 the male mortality rate is now higher at 2.9 per 1,000 and for females it is 2.0, where I have dispensed with the phrases one-year and median income for the sake of brevity. At the age of 60, the rates are 7.3 (males) and 4.5 (females) respectively. This is simply the effect of aging, which affects both males and females. There is nothing surprising quite yet, but notice how the excess of male-to-female mortality shrinks from 150% (=1.2/0.8) at the age of 40, to 145% (=2.9/2.0) at the age of 50. This isn't quite a downward trend (at least in the table), since at the age of 60 the excess death is back to 162% (=6.3/4.5).

Moving on to the (more interesting) rows, we now have the opportunity to measure the impact of income percentile on mortality rates. Notice that for a U.S. male in the lowest income percentile (south of the median), the mortality rate at the age of 50 is over four times higher at 12.5 deaths per 1,000, versus 2.9 at the median income. In stark contrast, a 50-year-old male at the highest income percentile (north of the median) experiences a mortality rate of only 1.1 per 1,000. This is less than half the median (income) rate. Stated differently, the range in mortality rates between the 1st percentile and 100th percentile is (12.5/1.1) or over eleven to one. To those who haven't seen such numbers before they might seem extreme, but they are by no means original. As noted by Deaton (2016) and quoted in the first paragraph of this paper, the link between mortality and income is well-established in the economics literature. The Chetty et al. (2016) article upon which these numbers are based is simply one of the most recent and comprehensive documentations of the *mortality to income gradient.* In fact, Goldman and Orszag (2014) offer similar evidence and their data seem to indicate an even wider gap (i.e. greater than 15 years) in life expectancy based on income and wealth factors.

Digging a bit deeper and upon further inspection of these numbers, notice how the ratio of mortality rates between the lowest-income percentile (top of table #1) and the highest-income percentile (bottom of table #1) shrink or decline over time. For example, for males at the age of 40 the ratio of worst-to-best is 9.67 (=5.8/0.6), whereas at the age of 60 the ratio falls to 7.89 (=22.1/2.8). The same decline (in relative rates) is observed for females. At the age of 40 the ratio of worst to best is 14 (=4.2/0.3), but by the age of 60 it shrinks to a multiple of 5.8 (12.8/2.2). Stated differently, mortality rates appear to converge with age.

One might take some solace in the fact mortality inequality (or the gap) declines with age, but in fact it is more a matter of biology – with some help from the laws of conditional probability – as opposed to any improvement in their wealth or health fortune. After all, to misquote the old proverb: what doesn't kill (those with lower income), likely makes them stronger. Cute sayings aside, this brings me to the next topic on the analytic agenda which is the rate of change in mortality rates as a function of age.

3.2 Trends by Age: A Glimpse of Gompertz

Table (#2) is based on the numbers contained in Table (#1) and displays the annual rate at which mortality rates themselves increase with age, as a function of income percentile. For example, the increase in mortality rates at the median income level was 9.21% per year for males and 8.64% for females. The number next to it is the projected mortality rate at the age of 100, denoted by $\tilde{q}_{100}$.

To be precise, this growth number is expressed in continuous time. It is computed by solving for g in the relationship: $q_{63} = q_{50}e^{g13}$, or equivalently: $g = (\ln q_{63}/\ln q_{50})/13$ where q_{63} denotes the mortality rate (for either males of females) at age: $x = 63$ and q_{50} is corresponding number at age: $x = 50$. The upper age of 63 isn't arbitrary, but in fact is the highest age for which Chetty et al. (2016) report realized mortality rates as a function of (lagged 2-years) income.

The mortality *growth* rates of 9.21% for males and 8.64% for females (or approximately 9% on a unisex basis) at the median income level aren't restricted to the age range of 50 to 63 and are not an artifact of this particular dataset. The (approximate) 9% rate growth in mortality is observed in most human species from the age of 35 to the age of 95. It is known as the Gompertz (1825) law of mortality, named after Benjamin Gompertz; possibly the world's first actuary.

Back to the topic of mortality inequality, notice though how the growth (rate) of mortality (rates) for individuals at the lowest income percentile is only 5.63% for males and 4.81% for females, which is close to half of the corresponding rate at the median income level. Perhaps this can be interpreted as some modicum of good news for the less economically fortunate. Their mortality rates don't grow or increase as fast. Of course, they have started off (at the age of 50) from a much higher base. In contrast, those fortunate to live in the highest income percentile experience a 10% growth in mortality as they age. Stated differently, they *age faster* than the median person in the population and *much* faster than those at the lowest income percentile. In fact, the difference between males (10%) and females (9.97%) is almost negligible. Notice how by age $x = 100$ the mortality rates are much closer to each other. I'll get back to this.

The disparity in mortality *growth* rates between high and low-income individuals, that is between the wealthy (sharks) and the poor (fish) noted in the title, leads to a corresponding gap in the *dispersion* of the remaining lifetime random variable: T_x. I'll define this variable formally and precisely, shortly, but there is in fact an inverse mathematical relationship between the mortality growth rate and the standard deviation: $SD[T_x]$. The lower growth rate is synonymous with an increase in the *individual volatility of longevity*, which is defined as the ratio of the standard deviation of remaining lifetime $SD[T_x]$ to the mean remaining lifetime $E[T_x]$. In sum, the demand for longevity insurance is relatively higher – and *ceteris paribus* they are willing to pay more for insurance – compared to those at the lowest income percentile. The formal derivation of the AEW will be presented in Section (#4.)

3.3 Compensation Law of Mortality

Figure (#2) displays the empirical relationship between income percentile and mortality growth rate in one summary graphic. The x-axis represents age x between 40 and 90, and the y-axis is the mortality rate, q_x as well as $\ln q_x$. The individual values or points displayed in table #1 are highlighted, but additional points are included, all from the Chetty at al. (2016) dataset. In addition, for ages beyond 63, a regression line is fit up to the age of 100, which was used to generate the necessary value of q_x. In fact, a variant of this picture (with log mortality rates) is presented in their article.

The bottom curve is the highest income percentile. Initial mortality rates are lower, but the growth rate of mortality (and slope of the $\ln q_x$ line) is higher. The upper curve captures the lowest income percentile. Mortality rates are higher, but they increase at a slower rate.

The point at which the $\ln q_x$ lines intersects the y-axis (denoted by c_0) is the log mortality rate at the age of 40. The slope of the curve (denoted by c_1) for regression purposes, is the individual mortality growth rate g. The lower c_0 and higher c_1 for the wealthy induces a long expected life and lower volatility.

Indeed, the negative and statistically significant relationship between the expected remaining lifetime $E[T_x]$ and standard deviation of lifetime $SD[T_x]$, conditional on income percentile has a biological basis in models of aging. Hypothetically, the individual regression lines underlying Figure (#2) would meet at some advanced age (e.g. 100+) at which point mortality rates across all income percentiles would be the same. Needless to say, the probability a low-income retiree will ever reach this advanced age is much lower relative to the high-income retiree. But conditional on survival they eventually unite.

The convergence of these curves (or regression lines when plotted against $\ln q_x$) is known as the *Compensation Law of Mortality* in the biology and gerontology literature. To this author's knowledge it was first identified by L.A. Gavrilov and N.S. Gavrilova and fully explained in the book by Gavrilov and Gavrilova (1991, page #148) using what they call a reliability theory of aging[4]. I also follow Gavrilov and Gavrilova (1991, 2001) and assume Gompertz to (more) advanced ages than assumed by Chetty et al (2016), although whether or not mortality plateaus at age 95, 100 or 105 isn't quite pertinent.

Without drifting too far from the economically-motivated script of this particular article, although the *Compensation Law of Mortality* and the negative relationship between (i.) middle-age death rates and (ii.) growth rates in any sub-group of species is known to biologists, it remains controversial because it's difficult to reconcile with popular theories of aging and longevity. Nonetheless, a framework that can explain this compensation law is the so-called reliability theory of aging, which also explains why organisms prefer to die according to the Gompertz law of mortality. We refer the interested reader to Gavrilov and Gavrilova (2001) for more on the underlying *biology of lifespan* and return to (the safety of) economic risk and return.

[4]The negative relationship (between c_0 and c_1) is occasionally referred to as the Strehler-Mildvan correlation but isn't quite the nature of the link described above. See also the interesting and related paper by Marmot and Shipley (1996), which suggests that mortality rates may not converge as fast at advanced ages.

3.4 Law of Mortality and Volatility of Longevity

To analyze the impact of the *individual volatility of mortality* on the demand for insurance one requires a parsimonious model for mortality rates over the lifecycle, and not a few discrete age data points. Recall from the earlier subsection – and in particular figure (#1) – that mortality rates q_x, increase exponentially in age x, and $\ln q_x$ increases linearly in x. This suggests that in continuous time, a suitable model for the log *hazard* rate is:

$$\lambda_{x+t} \;=\; h_x e^{gt} \;=\; \frac{1}{b} e^{\frac{x+t-m}{b}} \;=\; \left(\frac{e^{(x-m)/b}}{b}\right) e^{t/b} \tag{1}$$

Enter the Gompertz law of mortality. The (m, b) parameterization is common on the actuarial finance literature. The simpler (h, g) formulation is more common in demographics, statistics and economics. I will use both of them interchangeably, depending on context and need.

Under the (m, b) formulation, two free (or governing) parameters have a more intuitive probabilistic interpretation. They are labeled the *modal* value m, and the *dispersion* parameter b, both in units of years. Different groups within a species or population have different (m, b) values, in particular those with lower incomes will be characterized by low values of m and high values of b. Or, stated in terms of (h, g), at any given age, the hazard rate h for low income groups is higher, but the growth rate g is lower.

The intuition behind the (m, b) perspective – and why they are labeled modal and dispersion – only becomes evident once the remaining lifetime random variable: T_x, is defined via the hazard rate with a corresponding density function. Formally:

$$\Pr[T_x \geq t] := p(x, t, m, b) = e^{-\int_0^t \lambda_{x+s} ds} = \exp\{e^{(x-m)/b}(1 - e^{t/b})\}, \tag{2}$$

where going from the third to the fourth term in equation (2) follows from the definition of the hazard rate in equation (1). The expected value of the remaining lifetime random variable: $E[T_x]$, as well as any higher moment can be computed via the cumulative distribution function (CDF) induced by: $F(x, t, m, b) := 1 - \Pr[T_x \geq t]$, or more commonly using the probability density function (PDF) defined by: $f(x, t, m, b) = -\frac{\partial}{\partial t} p(x, t, m, b)$. One can re-write the density function in terms of hazard rate and growth rate, or (h, g) as well.

Getting back to the main point of this entire exercise, the *individual volatility of longevity* at the age of x, which I denote by the symbol φ_x, indexed by current age x, is as follows. It is defined equal to the standard deviation of the ratio of: T_x to it's expectation: $E[T_x]$. Formally it is computed as follows:

$$\varphi_x \;= SD\left(\frac{T_x}{E[T_x]}\right) \;=\; \frac{SD[T_x]}{E[T_x]} \;=\; \frac{\sqrt{\int_0^\infty t^2 f(t)dt - (\int_0^\infty t\, f(t)dt)^2}}{\int_0^\infty t\, f(t)dt}, \tag{3}$$

where the abbreviated: $f(t)$ is shorthand notation for the full probability density function (PDF) under a Gompertz law of mortality. Figure (#1), which I cited and mentioned earlier in the introduction when comparing Heather to Simon, is essentially a plot of $f(t)$ under two sets of (m, b) values. I'll present a number of examples and explicitly compute values of $E[T_x]$, $SD[T_x]$ and φ_x, in the next sub-section.

Before I get to numbers, I would like to help readers develop some intuition for the φ_x function (and the individual volatility of longevity) by examining its properties under the simplest of (non-Gompertz) mortality laws, namely when: $\lambda_x = \lambda$ and is constant at all ages. Humans age over time and their hazard rate increase, but there are actually a few species (such as lobsters) that never age. In our language, their hazard rate remains constant regardless of how old they are. They die (obviously) but the rate at which this occurs never changes. In fact, a constant hazard rate is associated with an exponential remaining lifetime and $\Pr[T_x \geq t] = e^{-\lambda t}$. The probability of survival declines with t, but the rate doesn't change. Under this non-Gompertz law, the expected value: $E[T_x] = \lambda^{-1}$ and the standard deviation is (also) $SD[T_x] = \lambda^{-1}$; both are well-known properties of exponential lifetimes.

Back to the volatility; according to the general definition of *individual volatility of longevity* presented in equation (3), but applied to the exponentially distributed lifetimes; $\varphi = 1$, at all ages and under all parameter values. In other words, volatility of longevity is invariant to λ. Your iVoL (risk) is always 100%. Indeed, this is more than a nice coincidence and is in fact the reason I chose to define iVoL in this particular way. Under a Gompertz law of mortality the *individual volatility of longevity* is **always** less than 100%. Regardless of wealth or income, life is less risky for humans (versus lobsters). And now to some numbers.

3.5 Numbers and intuition: iVoL

Table (#3) offers a range of numerical values for the *individual volatility of longevity* function: φ_x, at the age of x, under an assortment of parameter values for m (the mode of the distribution) and b (the dispersion coefficient.) Note that by fixing (m, b) the hazard rate at any age x, is simply: $\lambda_x = (1/b)e^{(x-m)/b}$, as per the definition of the Gompertz hazard rate.

For example, the values in the first row are computed assuming that $m = 98$ and the mortality growth rate is $1/b = 11.5\%$ per year, which implies that the dispersion coefficient is: $b = 6.696$ years. For reasons that should by now be obvious, I refer to this combination of (m, b) as rich (i.e. highest income percentile), and the life expectancy at birth: $E[T_0] = 92.98$ years. The standard deviation of the lifetime random variable at birth is: $SD[T_0] = 11.15$ years, which is higher than the dispersion coefficient $b = 8.696$ by slightly more than two years. At birth, the standard deviation: $SD[T_0] = b\,\pi/\sqrt{6}$, which is 28% greater than b.

To be clear, computing the mean and standard deviation at birth under a Gompertz assumption for the evolution of the hazard rate $\lambda_x, x = 0..\omega$, is somewhat artificial (or disingenuous). After all, even in the most developed countries mortality rates are relatively high in the first few years of life, reaching a minimum around the age of 10 and only starting to "behave" in a Gompertz-like manner between the age of: $x = 30$ to $x = 95$. I'm not saying all mortality tables at all ages are Gompertz. Rather, the point here in table (#3) is to provide some intuition for the moments of the Gompertz random variable as opposed to accurately modeling the earliest ages in the lifecycle. Needless to say, infants and children aren't purchasing life annuities or pooling longevity risk – at least not in the current environment.

Moving along the first row of the table, the *individual volatility of longevity* is a mere: $\varphi_0(98, 8.686) = 12\%$ birth, then increases over the lifecycle to reach $\varphi_{65}(98, 8.686) = 33.7\%$ at the age of $x = 65$. Note – for the sake of replicability – that although φ_0 can be derived analytically the iVoL numbers at $x >> 0$ are computed numerically.

Although the denominator of φ_x, that is the mean of the remaining lifetime, is available in closed form (see Appendix A.3 with $r = 0$), the numerator requires a numerical procedure (or some calculus dexterity beyond this author's abilities). Bottom line: The rich woman's *iVoL* at the age of retirement age of $x = 65$ is 33.7%. Again, the $m = 98$ and $b = 8.696$ parameters are representative of the highest income percentiles in the Chetty et al. (2016) data. This 65-year-old has a hazard rate of $\lambda_{65} = 0.2586\%$ at the age of $x = 65$.

Now, as I move down the rows within the first panel and artificially reduce the mortality growth rate from 11.5% to 5.5%, while holding the modal value of life fixed at 98 years, the *iVoL* values increase. To be clear, holding m fixed and reducing $1/b$ induces a corresponding increase in λ_x as per the 6th column in the table.

Although these parameter combinations don't necessarily correspond to any particular income percentile (the rich have high modal values and mortality growth rates) the point here is to develop some intuition for the link between age x, Gompertz parameters (m, b), hazard rates λ_x and the critical iVoL. Notice how reducing the mortality growth rate $(1/b)$, that is increasing b, uniformly reduces the life expectancy at birth $E[T_0]$ and increases the standard deviation $SD[T_0]$ at birth. The numerator goes up, the denominator goes down and so the *individual volatility of longevity* increases with higher values of b.

The situation becomes less obvious or clear-cut at the age of 65. Notice how the mean $E[T_{65}]$ starts-off at 28.82 years (when $b = 8.696$) and then declines as one would expect at higher values of b, but then flattens-out at 28.59 years and starts to increase. It seems the increased dispersion b then serves to increase the mean value, as often happens with highly skewed distributions. This is also driven by the fact the initial hazard rate at age 65 is higher. All in all, the standard deviation $SD[T_{65}]$ does in fact increase monotonically with b and the net effect is that the *individual volatility of longevity* does in increases in b.

Moving from rich to poor, in the bottom panel of table (#3) in which the modal value of life is $m = 78$ years, the corresponding life expectancy values are (much) lower at ages. The hazard rates are higher as well. At the very bottom row, where $b = 18.182$ years (i.e. the slowest aging), the life expectancy at retirement age $x = 65$ is now 17 years (vs. 28 years). The poor man's *iVoL* is 61.9%, which is almost double the top row, compared to the rich woman's 33.7%.

The artificial numbers in table (#3) and the underlying intuition are quite important for understanding the comparative statics in the next section. Notice how the *individual volatility of longevity* is subtly affected by the interplay between age x, modal value m and dispersion value b. As $x \to m$ the *iVoL* increases purely as a result of the underlying Gompertz law, even when b is unchanged. In other words, aging increases the relative riskiness of your remaining lifetime. From an economic point of view – as I will show in the next section – the demand for longevity insurance and risk pooling increases in age, a.k.a. with higher mortality rates. Just as importantly, fixing both x and m, and (only) increasing the value of b also increases *iVoL*. At any given age x, your *iVoL* is higher with higher m and/or with higher b. Figure (#3) displays iVoL numbers over the entire lifecycle for two sets of (m, b) parameters. Notice how it plateaus at higher ages and never exceeds 100%. One can think of it as the expected value of the Gompertz remaining lifetime: $E[T_x]$ converging to $SD[T_x]$, as $x \to \omega$. In probabilistic terms, this is how to think about it. At very advanced ages the amount of time you *will actually* live is very close to what you expect to live.

To be crystal clear, the *volatility* to which I refer to in this paper is unrelated to notions of stochastic mortality, Lee and Carter (1992) models, or aggregate changes in population q_{65} values over time. I am not using volatility to forecast mortality in (say) 50 years and the models in this paper are entirely deterministic. That said, if one graduates from a deterministic model of hazard rates λ_x, and postulates a stochastic model: $\tilde{\lambda}_x$, the φ_x metric would still be defined as the ratio of moments and should converge to a value of one at very advanced ages.

4 Annuity Equivalent Wealth and Willingness to Pay

We now arrive at the main reason for introducing the *individual volatility of longevity*.

4.1 Pricing Annuities

I use the symbol $\mathtt{a}(.)$ to denote the market price of a life annuity at age x. In words, paying: $\mathtt{a}(.)$ to an insurance company or pension fund, obligates the issuer to return \$1 of income per year (or $\$1dt$ in continuous time) for the life of the buyer, a.k.a. retiree or annuitant. The items inside the bracket (.) are the conditional factors, which could be age, gender, health, etc. For example, a retiree might pay: $\mathtt{a}(65) = \$20$ to purchase a life annuity providing \$1 per year for life starting at the age of: $x = 65$, but the price might be: $\mathtt{a}(75) = \$14$ if purchased at age: $x = 70$, etc. Both are completely arbitrary numbers, although throughout the paper for the sake of this discussion I'll assume life annuities scale and the price of $\$100,000$ of annual income is exactly $\$100000\mathtt{a}(.)$. There are no bulk (economies of scale) discounts or (adverse selection-induced) penalties. Moreover, the only loading or frictions will come (implicitly) from being pooled with healthier and longer lived individuals, as opposed to profits or other institutional features.

Now, in the above example market prices are differentiated by age x. In practice (most) companies and issuers differentiate by gender, while some go even further and *underwrite* annuities, that is price based on the health status of the annuitant. Therefore when necessary I will augment notation to include biological characteristics, namely the two Gompertz parameters (h, g) introduced in the prior section. Thus, $\mathtt{a}(h, g)$ denotes the market price of a \$1 per year life annuity, purchased at age: x when the hazard rate is $h = h_x$, for an individual whose remaining lifetime random variable is modeled in Gompertz (h, g) space. These were all explained in prior Section #3. The point I make here is that these two bio-demographic characteristics are easily observable and can (legally) be used for underwriting – or at least for theoretical *valuation* purposes. More on this to come.

As far as finance and markets are concerned, interest rates (obviously) impact the price of pension and life annuities, so in the event I must draw attention to the underlying pricing rate: r, assuming it is constant, I will append a third parameter to the very beginning of the expression and write the annuity factor as: $\mathtt{a}(r, h, g)$ for completeness. Notice that absence of any age (x) in the expression, since this is already contained and included within the hazard rate (h). Occasionally the expression $\mathtt{a}(r, x, m, b)$ will make an appearance when I want to draw specific attention to the impact of a modal value (m) or global dispersion parameter (b), on the annuity valuation factor at an explicit age denoted by (x).

Notice that up to now I have only mentioned market *prices* as opposed to say theoretical model or economic *values*, which naturally might differ from each other. To link these two numbers, I refer to what actuaries might call the fundamental law of pricing mortality-contingent claims, or what financial economists might simply call *No Arbitrage* valuation. Either way, the market price: $\mathtt{a}(x)$ under a Gompertz law of mortality <u>should</u> be equal to the following value:

$$\mathtt{a}(r, h, g) \ := \int_0^{\omega - x} e^{-rt} \, p(t, h, g) dt, \tag{4}$$

where ω denotes the last possible age (perhaps 125) to which people are assumed to live. The underlying economic assumption is that if a large enough group of known (h, g)-types are pooled together they will – by the law of large numbers – eliminate any idiosyncratic mortality risk and the valuation is easily conducted by discounted cash-flow expectations. Another (minor) embedded assumption within equation (4) is that the term structure of interest rates is a constant r, although that really isn't critical. By assuming a constant rate r though, the integral in equation (4) can be explicitly solved and expressed in closed-form using the Gamma function, although annuity factors are usually computed numerically. See appendix A.3 for more on the various techniques to calculate: $\mathtt{a}(r, h, g)$. Note that equation (4) can be expressed in discrete time, and is quite commonly used in a variety of papers in the economics literature to *price* annuities, for example in Poterba, Venti and Wise (2011), footnote #3. Either way, from this point onward I will refrain from using the phrase market *price* or *value* and henceforward refer to: $\mathtt{a}(.)$ as the annuity *factor*.

Before I proceed to the economics of the matter, it's important to focus attention on the sign of the partial derivatives of: $\mathtt{a}(r, h, g)$, with respect to the three explicitly listed arguments. First, the annuity factor declines with increasing age and hazard rate h. Intuitively (and unlike a perpetuity) the cost of a constant \$1 of lifetime income declines as you get older (and closer to death). Likewise, the factor declines at higher valuation rates r, after all it's just a present value. Just as importantly, at any conditional age x and hazard rate h, the annuity factor declines as the growth rate g is increased, which is synonymous with individuals who have a lower remaining life horizon. These insights don't require much calculus and are discussed at greater length in Appendix A.3

Now back to the economics of the matter. Let $(\hat{h}, \hat{g})$ denote the Gompertz parameters that best fit population (group) mortality while (h^i, g^i) denotes the parameters that best fit individual (sub-group) mortality. In particular, using ideas introduced in section #3, I let: (h^1, g^1) denote the Gompertz parameters of individuals (i.e. sub-group) in the lowest income percentile, whereas (h^{100}, g^{100}) denotes the Gompertz parameters of individuals in the highest income percentile. Therefore, the population modal and dispersion parameters would be the median values: $\hat{h} = h^{50}$, and $\hat{g} = g^{50}$, respectively (albeit with a bit of hand waving).[5] As mentioned earlier, on the occasion that I want to draw attention to population averages for (m^i) and (b^i) parameters, I'll use the obvious: $(\hat{m})$ and $(\hat{b})$.

[5]This gets messy. First of all the weighted average of Gompertz variables isn't Gompertz. Second, even if I select the best fitting average line, the h and g values will *not* be linear averages. So, this is an approximation, but precisely what Chetty et al. (2016) assumed as well.

4.2 Measuring Utility

Let: $U^{**}(w)$ denote the value function (maximal utility) of the individual who annuitizes their wealth w, and $U^{*}(w)$ the value function of the individual who decides to self-insure (i.e. not own any annuities at all) and instead decides to fund retirement with a systematic withdrawal program, then:

$$U^{**}(w) \ \geq U^{*}(w). \tag{5}$$

This is the well-known result in annuity economics, noted and cited in section #2. There exists a constant $\delta \geq 0$ such that:

$$U^{**}(w) \ = U^{*}(w(1+\delta)). \tag{6}$$

A retiree who doesn't annuitize would require the δ percent increase in their wealth w to induce the same level of utility as someone who does annuitize. Given that we are operating with constant relative risk aversion (CRRA) utility and no pre-existing annuity income, I will set $w = 1$ and refer to the *annuity equivalent wealth* by: $(1+\delta)$, and the value of longevity risk pooling by δ. To close the loop on all these (utility based) definitions, note that if my subjective value of risk pooling is: $\delta = 25\%$, and the AEW is \$1.25, then someone with $w = \$125$ of initial wealth would be willing to pay \$25 or $\$100\delta/(1+\delta)$ to have access to the annuity. The *willingness to pay* is then $\delta/(1+\delta)$.

Whether it be AEW, WtP or simply *the value of pooling* δ, all essentially measure the same thing and will be used interchangeably in the paper, unless the numbers themselves are important. Either way, the δ is a function of the individual Gompertz parameters (h^i, g^i), the market pricing parameters $(\hat{h}, \hat{g})$, and the utility-based preferences involving risk γ and discounting $\rho = r$.

Main Claim Assume an individual (denoted by i) and their force of mortality are individually Gompertz parameters and the population is also (approximated as) Gompertz, upon which annuity factors are based. The individual's AEW can be expressed as:

$$1+\delta_i \ = \frac{\mathtt{a}(r,x,m^i,b^i)^{1/(1-\gamma)}\,\mathtt{a}(r,x,\hat{m},\hat{b})^{-1}}{\mathtt{a}(r,x-b^i\ln[\gamma],m^i,b^i)^{\gamma/(1-\gamma)}} \ = \ \frac{\mathtt{a}(r,h^i,g^i,)^{1/(1-\gamma)}\,\mathtt{a}(r,\hat{h},\hat{g})^{-1}}{\mathtt{a}(r,h^i/\gamma,g^i)^{\gamma/(1-\gamma)}} \tag{7}$$

where γ, denotes the coefficient of relative risk aversion within CRRA utility, and assuming the subjective discount rate is (also) equal to r. Note that I have expressed the annuity factor using both (m, b) and (h, g) formulations, mainly so that δ can be (easily) computed when either set of parameters are available.

Corollary The annuity equivalent wealth $(1+\delta)$ is an increasing function of hazard rate h (and age), but a declining function of g, and therefore an increasing function of $1/g$. In other words, δ is an increasing function of the *individual volatility of longevity.* Supporting proofs are presented and referenced in the technical (A4) appendix.

Equation (7) might appear unwieldy and counterintuitive (with annuity factors to the power of risk aversion), so I'll proceed by first discussing the homogenous case when everyone in the sub-group experiences the exact same Gompertz force of mortality. In that case the *annuity equivalent wealth* is:

$$1+\delta \ = \ \left(\frac{\mathtt{a}(r,x,m,b)}{\mathtt{a}(r,x-\hat{b}\ln\gamma,m,b)}\right)^{\frac{\gamma}{1-\gamma}} \ = \ \left(\frac{\mathtt{a}(r,h,g)}{\mathtt{a}(r,h/\gamma,g)}\right)^{\frac{\gamma}{1-\gamma}}. \tag{8}$$

The *annuity equivalent wealth* is a simple function of the ratio of two actuarial annuity factors. In particular, notice that the respective annuity factors in equation (8) are computed at two distinct ages (or hazard rates.) The numerator is computed at the biological age: x (or hazard rate h), and the denominator is computed at a modified (risk-adjusted) age: $x - \hat{b} \ln \gamma$ (or hazard rate h/γ.) The modified age in the denominator's factor is under (younger than) x whenever $\gamma > 1$. The lower modified age increases the annuity factor and the AEW. I'll now provide some examples and intuition for the *annuity equivalent wealth* using both equations (7) and (8).

4.3 An International Numerical Example

Based on mortality data for OECD countries (presented in the appendix) the remaining life expectancy at the age of 65 in *Latvia* for the 1940 cohort is: 11.95 years. Latvia was selected because its implied: $E[T_{65}]$ is the lowest across all OECD countries. The underlying mortality rates correspond to Gompertz parameters of: $m = 75.02$ and $b = 11.87$. The volatility of longevity at age 65 (for a Latvian born in 1940) is 60.4%. That, recall, is the standard deviation of remaining lifetime: $SD[T_{65}]$ scaled by mean remaining lifetime: $E[T_{65}]$.

In contrast, the remaining life expectancy at the age of 65 in *Japan*, which was highest across all OECD countries, was $E[T_{65}] = 23.64$ years. The associated best-fitting Gompertz parameters are: $m = 91.72$ and $b = 12.87$, for a volatility of longevity of (only): 47.4%. As this is only a numerical example, I'll assume that the (subjective discount rate and) valuation rate are: $r = 3\%$, and the coefficient of relative risk aversion is: $\gamma = 3$ in the CRRA utility. These values are not unreasonable for the valuation of subjective utility, AEW and WtP.

Under these parameter values a representative Latvian (65 year old, retiree) would value the longevity insurance at: $\delta = 89.32\%$ if they could pool this risk with other similar Latvians, by acquiring fairly priced life annuities based on their own population: $m = 75.02$ and $b = 11.87$ parameters. This 89.32% number is based on equation (8), first by computing $\mathtt{a}(0.03, 65, 75.02, 11.87) = 9.493$, dividing by $\mathtt{a}(0.03, 51.96, 75.02, 11.87) = 14.528$, raising to the power of $(-3/2)$ and then subtracting one to express as a percentage. In particular, notice the age setback in the denominator from: $x = 65$ to the modified age of $51.96 = 65 - (11.87) \ln 3$, which I labeled a risk-adjusted age. This AEW is in line with the (generally) large benefits from annuitization reported in (many) other studies over the last 30 years, noted in the literature review of section #2. For example, Brown (2003) computes δ (numerically) over a range of coefficient of relative risk aversion (CRRA) values, $\gamma = 1..5$. His numbers range from 36% to 90%, depending on demographic factors.

Now, let's examine the representative Japanese retiree, assuming the same $r = 3\%$ valuation rate and $\gamma = 3$, coefficient of risk aversion. If they are pooled with homogeneous risks (i.e. Japanese), their δ value at the age of $x = 65$ is lower than their Latvian counterparts. Intuitively, their volatility of longevity (VoL) is also lower at age 65. In particular, using the $m = 91.72, b = 12.87$ values, equation (8) leads to a value of $\delta = 48.39\%$ which is lower than what a Latvian would be willing to pay. This is despite the obvious (but not quite relevant) fact that the Japanese retiree is expected to live 24 years versus the 12 years for the Latvian. Indeed, what drives the δ for (fair) longevity insurance is the volatility of longevity (via the m and b) and not the demographic fact that Japanese retirees live longer.

Moving on to a larger heterogenous pool, imagine Latvians and Japanese (both at chronological age 65) are mixed together in equal amount in an OECD pool. To keep the system fair, the pension annuities must be priced based on the entire (Latvian plus Japanese) population mortality curve. With a bit of hand waving, assume the resulting Gompertz parameter upon which the group annuities are priced are: $m = 85.45$ and $b = 12.41$, which are OECD averages and explained in the appendix. Intuitively now, the Latvian is presented with a relatively worse (loaded) annuity price and the Japanese is getting a relatively better (subsidized) price. Again, this assumes both groups are forced to purchase annuities at the same price in a compulsory system. The *annuity equivalent wealth* is based on the group population (for market pricing) and individual (for lifecycle utility) mortality. I'm in the realm of equation (7). The group annuity factor – or price they both pay – is now $\mathsf{a}(r = 3\%, 65, \hat{m} = 85.45, \hat{b} = 12.41) = 13.583$, which isn't actuarially fair to either of them. It's advantageous to the Japanese (who would have had to pay 15.97) and disadvantageous to the Latvian (who would only have to pay 9.493). Bottom line, the Japanese now has a higher $\delta = 74.48$, which is higher than the prior (homogenous case of): $\delta = 48.39\%$.

In contrast, the Latvian is faced with loaded prices (13.583 vs. 9.493) and they are only valuing the annuity at: $\delta = 32.32\%$. In fact, if the Latvian had a life expectancy that was (even) lower, but still paid the same group price, it's conceivable the δ in equation (7) might actually be negative. Under those conditions Latvians would not be willing to swim in a pool with healthy Japanese. They would rather take their chances and self-insure longevity.

Figure (#4) illustrates this relationship graphically over a spectrum of possible m values. On the left are individuals (Latvians) with low life expectancy (proxied by m) values and a correspondingly higher volatility of longevity at retirement. To the right are individuals (Japanese) with higher life expectancy and lower volatility of longevity. The Gompertz dispersion of longevity – in contrast to the *volatility* of longevity – for the purposes of computing equation (7), is held constant at: $b = 12$ years. The upper range is based on $\gamma = 5$, that is higher levels of risk aversion, whereas the lower range is for $\gamma = 1$, albeit with slight modifications in equation (7) to account for logarithmic utility.

Notice that as m increases, the value of longevity risk pooling and willingness to pay declines – when the pooling is homogenous. A Japanese retiree (paying fair prices) experiences a δ that is lower than a Latvian (paying fair prices); but when they are mixed together and both pay the same group price, the curve is reversed. Intrinsically the Latvian values longevity insurance more, due to his/her higher volatility of longevity, but the positive loading reduces its appeal. In contrast, the Japanese who – shall we say loosely – was somewhat ambivalent about the benefit of longevity risk pooling, is now willing to pay more for a relatively cheap life annuity. To be clear, it might seem odd to talk about a change in the subjective willingness to pay for something based only on its market price. After all, my willingness to pay for apples or oranges should not depend on price. Indeed, that's the point of the metric, to be invariant to market prices. Rather, the proper way to think about this as the willingness to pay to have some oranges (i.e. Latvians) mixed in with your apples (i.e. Japanese) when annuitizing. The Japanese is willing to pay more for the mixed bag. The Latvian is willing to pay less. The Latvian values the mixed bag less than the Japanese. Needless to say, Latvians and Japanese aren't forced to swim together in longevity risk pools, and these labels are a euphemism for individuals with easily identifiable characteristics.

This brings me to the next juncture. What happens in the U.S. across different income percentiles? Does the lowest income percentiles (i.e. Simon) experience a large enough volatility of longevity to over-come the implicit loading induced by having to subsidize the annuities of the highest income (i.e. Heather)?

5 AEW by Income Percentile in the U.S.

I now use the main equation (7) with plausible values for the coefficient of relative risk aversion, $\gamma = 1..5$, and actual values for the Gompertz parameters (m_i, b_i) as a function of income percentile. Figure (#5) displays the relationship between (m) and (b) as a function of income percentile, based on the earlier-mentioned data contained in Chetty et al. (2016)[6].

Table (#4) displays results using the mid-point of $\gamma = 3$, with additional (tables of values) available upon request. First, I estimate the relevant Gompertz m^i and b^i (or h^i and g^i) values, then I compute the remaining lifetime expected value: $E[T_x]$ and standard deviation: $SD[T_x]$. With those numbers in hand, I can show the *individual volatility of longevity* and finally the two AEW values. Notice how both the Gompertz b value, which measures the random variable's intrinsic dispersion, and the volatility of longevity decline at higher income percentiles. Again, the iVoL is lower due to the decline in b and the intrinsic increase of iVoL as age x edges closer to the value m.

The last two columns in table (#4) are computed using equation (7) and (8), for individuals and groups respectively. To be clear, the column labeled AEW for Individual, assumes that at any percentile, these 65-year-olds are pooled with individuals who share their same m_i, b_i and are identical risk types. They pay fair actuarial prices for their annuity. Moving down the panels, higher-income percentiles are associated with (lower iVoL) and lower values of annuity equivalent wealth under **individual** pricing. The intuition for this result was presented in section (#4.1) and displayed in figure (#4). Notice how much more annuities are worth to the poor vs. rich, when they are fairly priced.

Moving on to mandatory annuities pricing, the right-most column, is based on loaded pricing which is obviously disadvantageous for some (poor) and beneficial for others (rich.) As noted many times in the paper, if you are healthier than some people in the annuity pricing pool – and you aren't paying the fair rate for your risk – group pricing offers a higher annuity equivalent worth. The group value of AEW *increases* with income percentile for males, because they get better relative discounts. This pattern (and one that was illustrated hypothetically in figure #4) is not observed for females. The reason for this discrepancy is likely that the unique structure of the implicit annuity loading which destroys the uniformity The important point here is that **all** values of $\delta > 0$, for the $\gamma = 3$ case. They all benefit (as far as utility is concerned) from swimming together.

[6] As far as measuring risk aversion is concerned, I'm aware of the ongoing debate and well known problems in calibrating γ and refer the interested reader to recent work by O'Donoghue and Somerville (2018), or Schildberg-Hörisch (2018) and rely on Andersen (2008), for example, for the justification in using such values. Likewise, while making comparison across different percentiles, I'll assume that γ remains constant and doesn't depend on the particular choice of (m_i, b_i) or (h^i, g^i) values. In fact, an argument could be made that individuals with higher mortality might actually have lower levels or risk aversion. See Cohen and Einav (2007) for a possible link between risk exposure and demand for insurance.

Again, the key (policy) takeaway from this table is that *even* at the lowest income percentile, where the Gompertz hazard rate is higher and life expectancy at retirement is a mere decade or so, the value of δ is positive *even* at the group pricing rate. Of course, this assumes that prices are purely based on group mortality, that is determined by the 50-percentile parameters. If there are additional loadings or costs added-on, the δ might become negative. I should note that in my extensive analysis (not all displayed) at low income percentiles, large amounts of pre-existing pension annuity income and lower levels of risk aversion γ, the δ value is barely positive.

6 Summary and Conclusion

In a mandatory pension system participants with shorter lifetimes *ex ante* subsidize those expected to live longer. Moreover, since individuals with higher incomes tend to have lower mortality, the poor end-up subsidizing the rich. To quote Brown (2003), "When measured on a financial basis, these transfers can be quite large, and often away from economically disadvantaged groups and towards groups that are better off financially." This uncomfortable fact is well established in the literature and occasionally touted as a (social) justification for transitioning to Defined Contribution (DC) schemes. And yet, Brown (2003) goes on to write that "the insurance value of annuitization is sufficiently large that relative to a world with no annuities, all groups can be made better off through mandatory annuitization."

The question motivating this paper is: At what point does the gap in longevity expectations become such that the value of annuitization is actually negative for the ones who are expected to live the least? This is an empirical question and quite relevant to the growing disparity in U.S. mortality as a function of income. Against this background, this paper focuses attention on the heterogeneity of the second moment of remaining lifetimes, something that has not received much interest in the economics literature. I demonstrate that at any given chronological age the *individual volatility of longevity* for low-income earners is larger relative to high-income earners. Life (financially) is riskier for the poor. This then implies that their *willingness-to-pay* (WtP) for longevity insurance and the *annuity equivalent wealth* is greater, relative to high-income earners. Ergo, in a mandatory DB pension system there are two competing or opposite effects. On the one hand there is a clear and expected transfer of wealth from poor (i.e. higher mortality) to rich (i.e. lower mortality). On the other hand, economically disadvantaged participants benefit more from risk pooling due to their higher *individual volatility of longevity.* This paper – and in particular the main equation for δ – helps locate the cut-off point.

If I can summarize with the characters I introduced at the very beginning, both Heather and Simon benefit from longevity risk pooling, that is owning life annuities, regardless of whether they are pooled (and swim) with people like themselves or forced to pool with others. When annuities are fairly priced, that is tailored to their own risk-type and we swim in small segregated pools, Simon derives relatively higher benefit from pooling longevity risk and holding annuities. Mother nature endows him with a higher mortality rate together with a slower mortality *growth* rate. This then is synonymous with a higher volatility of longevity and risk averse retirees are willing to pay (dearly) to mitigate the higher longevity risk.

In some sense, nature's *Compensation Law of Mortality* leads to a higher demand for longevity insurance from those with the highest mortality rate. In contrast, Heather's relative benefit from annuitization might be (much) smaller than Simon's because her life expectancy is (much) higher, which means her mortality rate is lower and her individual volatility of longevity is lower. Practically speaking (and in the real world), after adding a relatively small insurance profit loading – and perhaps some pre-existing annuity income to her portfolio – Heather might derive no value from additional annuitization. Again, this is when she swims alone. Once they are forced into the same pool, although Simon is paying a penalty (implicit insurance loading) and is subsiding Heather, she now is benefiting from cheaper annuities. This increases her willingness to pay. So, whether the value of pooling is positive is an empirical question and the data presented indicate that (for now) it is still the case.

References

[1] Andersen, S. , Harrison, G. W., Lau, M. I. and Rutstrom, E. E. (2008), Eliciting Risk and Time Preferences, *Econometrica*, Vol. 76, pg. 583-618.

[2] Andersson, E., P. Lundborg and J. Viksrom (2015), Income receipt and mortality: Evidence from Swedish public sector employees, *Journal of Public Economics*, Vol. 131, pg. 21-32.

[3] Barseghyan, L., F. Molinari, and T. O'Donoghue (2018), Estimating risk preferences in the field, *Journal of Economic Literature*, Vol. 56(2), 501-564.

[4] Bodie, Z. (1990), Pensions as retirement income insurance, *Journal of Economic Literature*, Vol. 28, pg. 28-49

[5] Bommier, A. (2006), Uncertain lifetime and inter-temporal choice: risk aversion as a rationale for time discounting, *International Economic Review*, Vol. 47, pp. 1223-1246.

[6] Brown, J.R. (2001), Private pensions, mortality risk and the decision to annuitize, *Journal of Public Economics*, Vol. 82, pg. 29-62.

[7] Brown, J.R. (2003), Redistribution and insurance: Mandatory annuitization with mortality heterogeneity, *Journal of Risk and Insurance*, Vol. 70(1), pg. 17-41.

[8] Brown, J. R., O.S. Mitchell, J.M. Poterba, and M.J. Warshawsky (2001), *The Role of Annuity Markets in Financing Retirement*, MIT Press, Cambridge.

[9] Brown, J.R., J.R. Kling, S. Mullainathan and M.V. Wrobel (2008), Why don't people insure late-life consumption? A framing explanation of the under-annuitization puzzle, *American Economic Review*, Vol. 98(2), 304-309.

[10] Bütler, M., S. Staubli and M.G. Zito (2013), How much does annuity demand react to a large price change? *Scandinavian Journal of Economics*, Vol. 115(3), pg. 808-824.

[11] Cannon, E. and I. Tonks (2008), *Annuity Markets*, Oxford University Press, New York.

[12] Chetty, R., M. Stepner, S. Abraham, S. Lin, B. Scuderi, N. Turner, A. Bergeron and D. Cutler (2016), The association between income and life expectancy in the United States, 2001-2014, *Journal of the American Medical Association*, Vol. 315(16), pg. 1750-1766.

[13] Cocco, J.F. and F.J Gomes (2012), Longevity risk, retirement savings and financial innovation, *Journal of Financial Economics*, Vol. 103(3), pp. 507-529.

[14] Cohen, A. and L. Einav (2007), Estimating risk preferences from deductible choice, *American Economic Review*, Vol. 97(3), pg. 745-788.

[15] Davidoff, T., J.R. Brown, and P.A. Diamond (2005), Annuities and individual welfare, *American Economic Review*, Vol. 95(5), pg. 1573-1590.

[16] Davies, J.B. (1981), Uncertain lifetime, consumption and dissaving in retirement, *Journal of Political Economy*, Vol. 89, pp. 561-577.

[17] Deaton, A. (2016), On death and money: History, facts and explanations, *Journal of the American Medical Association*, Vol. 315(16), pg. 1703-1705.

[18] De Nardi, M., E. French and J.B. Jones (2009), Life expectancy and old age savings, *American Economic Review*, Vol. 99(2), pg. 110-115.

[19] Diamond, P. (2004), Social security, *American Economic Review*, Vol. 94(1), 1-24.

[20] Edwards, R.D. (2013), The cost of uncertain life span, *Journal of Population Economics*, Vol. 26, pg. 1485-1522.

[21] Finkelstein, A. and J. Poterba (2004), Adverse selection in insurance markets: Policyholder evidence from the U.K. annuity market, *Journal of Political Economy*, Vol. 112(1), 183-208.

[22] Gan, L., G. Gong, M. Hurd and D. McFadden (2015), Subjective mortality risk and bequests, *Journal of Econometrics*, Vol. 188, pg. 514-525.

[23] Gavrilov, L.A. and N.S. Gavrilova (1991), *The Biology of Lifespan: A Quantitative Approach*, Harwood Academic Publishers, United Kingdom.

[24] Gavrilov, L.A. and N.S. Gavrilova (2001), The Reliability Theory of Aging and Longevity, *Journal of Theoretical Biology*, Vol. 213(4), pg. 527-545.

[25] Goldman, D.P. and P.R. Orszag (2014), The growing gap in life expectancy: Using the future elderly model to estimate implications for social security and medicare, *American Economic Review*, Vol. 104(5), pg. 230-233.

[26] Gompertz, B. (1825), On the nature of the function expressive of the law of human mortality and on a new mode of determining the value of life contingencies, *Philosophical Transactions of the Royal Society of London*, Vol. 115, pg. 513-583.

[27] Holzman, R., J. Alonso-Garcia, H. Labit-Hardy, and A.M. Villegas (2017), NDC Schemes and Heterogeneity in Longevity: Proposals for Redesign, *CEPAR*.

[28] Hosseini, R. (2015), Adverse selection in the annuity market and the role for social security, *Journal of Political Economy*, Vol. 123(4), pg. 941-984.

[29] Inkman, J., P. Lopes and A. Michaelides (2010), How deep is the annuity market participation puzzle, *The Review of Financial Studies*, Vol. 24(1), pg. 279-317.

[30] Kotlikof, L.J. and A. Spivak (1981), The family as an incomplete annuity market, *Journal of Political Economy*, Vol. 89(2), pg. 372-391.

[31] Levhari, D. and L.J. Mirman (1977), Savings and consumption with an uncertain horizon, *Journal of Political Economy*, Vol. 85(2), pp. 265-281.

[32] Milevsky, M.A. and H. Huang (2018), The utility value of longevity risk pooling: Analytic insights, *North American Actuarial Journal*, in press. Draft available at SSRN: *https://dx.doi.org/10.2139/ssrn.3138280*

[33] Marmot, M.G. and M. J. Shipley (1996), Do socioeconomic differences in mortality persist after retirement? 25 year follow up of civil servants from the first Whitehall study, *British Medical Journal*, Vol. 313, pg. 1177-80.

[34] O'Donoghue, T. and J. Somerville (2018), Modeling risk aversion in economics, *Journal of Economic Perspectives*, Vol. 32(2), pg. 91-111.

[35] Pashchenko, S. (2013), Accounting for non-annuitization, *Journal of Public Economics*, Vol. 98, pg. 53-67.

[36] Peltzman, S. (2009), Mortality inequality, *Journal of Economic Perspectives*, Vol. 23(4), pg. 175-190.

[37] Poterba, J.M. (2014), Retirement security in an aging population, *American Economic Review*, Vol. 104)5), pg. 1-30.

[38] Poterba, J.M., S. Venti and D. Wise (2011), The composition and drawdown of wealth in retirement, *Journal of Economic Perspectives*, Vol. 25(4), pg. 95-118.

[39] Reichling, F. and K. Smetters (2015) Optimal Annuitization with Stochastic Mortality and Correlated Mortality Cost, *American Economic Review*, Vol. 11, pg. 3273-3320.

[40] Schildberg-Hörisch, H. (2018), Are risk preferences stable? *Journal of Economic Perspectives*, Vol. 32(2), pg. 135-145.

[41] Sheshinski, E. (2007), *The Economic Theory of Annuities*, Princeton University Press, Princeton.

[42] Tuljapurkar, S. and R.D. Edwards (2011), Variance in death and its implications for modeling and forecasting mortality, *Demographic Research*, Vol. 24, pg. 497-526.

[43] Yaari, M.E. (1965), Uncertain lifetime, life insurance and the theory of the consumer, *The Review of Economic Studies*, Vol. 32(2), pp. 137-150.

A Technical Appendix

This technical appendix (A1.) describes how to calibrate the Benjamin Gompertz (1825) law of mortality to any set of discrete mortality rates via a linear regression methodology; then (A2.) uses the fitting methodology to estimated Gompertz parameters for OCED countries, (A3) derives the annuity factor $\mathtt{a}(x)$ under a Gompertz law of mortality; and finally (A4.) sketches the derivation for the analytic expression for the *annuity equivalent wealth* (AEW) as a function of the Gompertz parameters.

A.1 Calibrating the Gompertz Model

I start with a vector $\{\tilde{q}_x\}$ of one-year death rates, where the subscript x denotes chronological age. Note that $\{\tilde{q}_x\}$ can be based on a given *period* or *cohort* mortality table, or it can be the *expected* death rates for a group based on their medical condition and ailments. In fact, the individual elements within $\{\tilde{q}_x\}$ do not have to be increasing in chronological age, although one would expect (and like) them to, after the age of 10 or so. It is best to think of: $\{\tilde{q}_x\}$ as being observed with *error*, and the true one-year death rates are: $\{q_x\}$. This means:

$$\tilde{q}_x \;=\; q_x + \epsilon_x, \tag{9}$$

where ϵ_x is an age-dependent error (or noise) term. From a modeling perspective the (input, data) one-year death rate q_x at any given chronological age x is related to the continuous time mortality (a.k.a. hazard) rate denoted by λ_x, via the relationship:

$$q_x \;=\; 1 - e^{-\int_0^1 \lambda_{(x+s)} ds}. \tag{10}$$

Naturally, when the mortality rate $\lambda_x = \lambda$ is constant (and nobody ever ages), then $q_x = 1 - e^{-\lambda}$. In this (simplistic) case, λ is synonymous with a continuously compounded (interest) rate and q_x is the *effective* rate. Recall that I am operating under the Gompertz law of mortality, so the λ_x is assumed to satisfy the following relationship:

$$\lambda_{x+t} \;= h_x\, e^{gt} \;=\; \frac{1}{b} e^{(x+t-m)/b}, \tag{11}$$

See section (#3.4) for more information about the intuition. In this (self-contained) technical appendix I am interested in demonstrating how to fit or calibrate the Gompertz law, using the (m, b) formulation, to any given mortality table or set of mortality rates. Note that under the (m, b) formulation in equation (11), the current age can be expressed as:

$$x \;=\; b \ln[b\, \lambda_x] + m, \tag{12}$$

which might seem odd at this point, but can be used to recover age (x) from a given hazard rate h_x, assuming a (global) Gompertz law of mortality links hazard to age and vice versa. The point, once again, is to illustrate how: (h, g) and (m, b) are synonymous.

Combining equation (10) and (11), plus a bit of calculus, leads to the following relationship between the one-year death rate q_x and the two underlying parameters (m, b).

$$q_x \;= 1 - \exp\{-e^{(x-m)/b}(e^{1/b} - 1)\} \tag{13}$$

$$= 1 - \exp\{b\, \lambda_x\, (1 - e^{1/b})\} \tag{14}$$

I can now invert and express the continuous hazard rate as:

$$\lambda_x \;=\; \frac{\ln[1-q_x]}{b(1-e^{1/b})}. \tag{15}$$

For example, a one-year death rate of $q_x = 2\%$ is associated with a mortality rate of $\lambda_x = 1.937\%$ when $b = 12$, but leads to a (lower) mortality rate of $\lambda_x = 1.897\%$ when $b = 8$. (Think curvature.) Now, I can rewrite (a transformation of the) one-year death rate q_x as a linear function of the chronological age x by taking double logs of the one-year death rate from equation (13), as follows:

$$y_x \;:=\; \ln\left(\ln(\frac{1}{1-q_x})\right) \;=\; \left[\ln(e^{1/b}-1) - \frac{m}{b}\right] + \frac{1}{b}x. \tag{16}$$

To get a sense of the magnitude of the left-hand side of equation (16), for values of $q_x \in [10^{-6}, 1/4]$, the corresponding value of $y_x \in [-13.8, -1.25]$.

Again, this transformation might seem needlessly complicated, but the benefit of the (non) linear form in chronological age x, is that I can estimate the underlying (unknown) parameters (m, b), or (h, g), using basic regression techniques. I could do this nonlinearly (here) via Maximum Likelihood Estimates (MLE) or take one final step and linearize the problem, which has some appealing pedagogical features and is the methodology adopted in the paper. Chetty et al. (2016) compare both methodologies and conclude that in large populations they produce virtually identical estimates. Moving on, define two (new) constants:

$$\begin{aligned} c_0 &= \ln\left[e^{1/b}-1\right] - \frac{m}{b}, \\ c_1 &= 1/b. \end{aligned} \tag{17}$$

From here on, the y_x defined by equation (16) can be written neatly as: $y_x = c_0 + c_1 x$. So, back to our original estimation task, assuming that one-year death rates $\tilde{q}_x$ are observed and estimated with noise, the expression we need is:

$$\tilde{y}_x \;=\; c_0 + c_1 x + \epsilon, \tag{18}$$

where ϵ is the standard error term (with all its baggage and assumptions). We are now ready for some statistics. Once I locate the best-fitting values for the parameters (c_0, c_1), which recall are the values that maximize the R^2 in the linear regression specified by equation (18), I can recover estimates for the original (m, b) parameters in equation (17) using the inverted relationship:

$$b = \frac{1}{c_1}, \quad m = \frac{\ln\left[e^{c_1}-1\right] - c_0}{c_1} \tag{19}$$

Note that the confidence intervals and sampling properties of (c_0, c_1) will induce different (and slightly more complicated) ranges for (m, b), but the point estimates listed in equation (19) are indeed *consistent*.

A.2 Volatility of Longevity Around the World

Using the algorithm or methodology described above, Table (#5) and Figure (#6) display the (m^i, b^i) and iVoL numbers across 30 OECD countries for which survivor data (of the 1940 cohort) are available. To be very specific, I downloaded mortality rates for the 1940 cohort from the Human Mortality Database (HMD) for countries that are part of the Organization for Economic Cooperation and Development (OECD), as of July 2018. The countries for which these q_x values were available are listed in the first column of table (#5), ranked by the *individual volatility of longevity* which I displayed in the last column. Mortality rates between the age of $x = 30$ and $x = 90$, are averaged for both males and females to produce one unisex q_x vector of approximately 60 elements. This then became the basis of the country-specific y_x vector, that is the left-hand side of equation (15) noted in appendix A1. From these individual (country by country) regressions, I extracted the relevant intercept c_0 and slope c_0, and then converted them into (m, b) as well as (h, g) values using equation (18). They are displayed in columns #2 and #3 of the table. Note that although modal m values range from a high of 91.72 years using Japanese data, to a low of 75.02 using Latvian data, the dispersion b values (that is $1/g$) do not exhibit the negative pattern or relationship to m, found in the Chetty et al. (2016) data for income percentiles. The average b is 12.41 years, ranging from 10.66 (Greece) to 14.11 in Slovenia. Nevertheless, the (lack of) pattern for b is not enough to offset the inverse relationship between life expectancy at retirement and the volatility of longevity. In other words, even when b stays exactly the same, the closer x gets to the modal value m, the higher is the iVoL. In table #5 we see a clear pattern. A 65-year-old in Latvia is a mere: $(m - x) = 10$ years away from m. So, despite the fact his/her b value is 11.87 (which is lower than it is for the Japanese) the *individual volatility of longevity* is 60.4%, which is higher than Japan's 47.4%, because the Japanese retiree at the age of 65 is $(m - x) = 26.72$ years away from the modal value of life. In sum, partially as a result of the nature of the Gompertz law of mortality, countries with low $E[T_{65}]$ exhibit high iVoL, and vice versa.

A.3 Closed-form Annuity Factors and Moments

Recall that I define and express the continuous-time *hazard rate* function as:

$$\lambda_{x+t} \;=\; h_x \, e^{gt} \;=\; \frac{1}{b} e^{(x+t-m)/b}, \tag{20}$$

where $h_x = h$ is the hazard rate at (an arbitrary baseline) age x, and g is the hazard *growth* rate. The parameter g represents the slope in a regression of *log* hazard rates (as an independent variable) on age (the dependent variable.) Note that: $b = 1/g$, which measures the *dispersion* of the remaining lifetime random variable denoted by: T_x in the (m, b) formulation. For what follows I work with the: (h, g) formulation, which leads to a cleaner and more intuitive relationship with δ, although one can easily move from (m, b)-space to (h, g)-space. For example, fixing the hazard rate at age 65: $h_{65} = 0.5\%$ and assuming a growth rate of: $g = 10\%$, the modal value of the lifetime is: $m = 94.957 = 65 - \ln[0.005/0.1]/0.1$ years and the dispersion value is: $b = 1/g = 10$. Likewise, if: $h_{65} = 0.5\%$, but the growth rate is $g = 8\%$, the: $m = 98.761$ years and $b = 12.5$ years.

Moving on, under the (h, g) hazard rate formulation, the conditional survival probability, denoted by: $p(t, h, g)$ is equal to:

$$p(t,h,g) \;=\; \exp\{-\int_0^t h_{x+s}\,ds\} \;=\; \exp\{(h/g)(1-e^{gt})\}. \tag{21}$$

Recall that any immediate life annuity factor can be expressed as:

$$\begin{aligned} \mathtt{a}(r,h,g) \;&=\; \int_0^\infty e^{-rt}\,p(t,h,g)\,dt \\ &=\; \int_0^\infty e^{-rt}\exp\{(h/g)(1-e^{gt})\}\,dt \\ &=\; e^{h/g}\int_0^\infty e^{-rt}\exp\{-(h/g)e^{gt}\}\,dt \end{aligned} \tag{22}$$

I can simplify (or solve) the integral using change of variable techniques. In particular, define a new variable $s = (h/g)e^{gt}$ therefore:

$$s = \frac{h}{g}e^{gt} \quad\rightarrow\quad t = \ln[\frac{sg}{h}]/g \quad\rightarrow\quad dt = \frac{1}{sg}ds \tag{23}$$

Using the new variable s, instead of $(h/g)e^{gt}$, we can simplify the integrand:

$$e^{h/g}\int_0^\infty e^{-rt}\exp\{-(h/g)e^{gt}\}\,dt = \; e^{h/g}\int_0^\infty e^{-rt}e^{-s}dt \tag{24}$$

We can now replace the t with: $\ln[sg/h]/g$ to obtain:

$$e^{h/g}\int_0^\infty e^{-rt}e^{-s}dt = \; e^{h/g}\int_0^\infty e^{-r(\ln[\frac{sg}{h}]/g)}e^{-s}dt \;=\; e^{h/g}\int_0^\infty e^{-s}(\frac{sg}{h})^{(\frac{-r}{g})}dt \tag{25}$$

Replacing dt with $(1/sg)ds$ and moving all non s terms outside the integral:

$$\mathtt{a}(r,h,g) = \; e^{h/g}\int_0^\infty e^{-s}(\frac{sg}{h})^{(\frac{-r}{g})}\frac{1}{sg}ds \;=\; e^{h/g}(\frac{g}{h})^{(\frac{-r}{g})}\frac{1}{g}\int_{h/g}^\infty e^{-s}s^{(\frac{-r}{g})-1}ds \tag{26}$$

The terms outside and to the left the integral can be simplified to:

$$e^{h/g}(\frac{g}{h})^{(\frac{-r}{g})}\frac{1}{g} \;=\; \frac{1}{g}\,\exp\{\frac{1}{g}\,(h + r\,\ln[h/g])\} \;=\; \frac{1}{g\,\exp\{(-1/g)\,(h + r\ln[h/g]\,\}} \tag{27}$$

The annuity factor can now be re-written as:

$$\mathtt{a}(r,h,g) \;=\; \frac{1}{g\,\exp\{(-1/g)\,(h + r\ln[h/g]\,\}}\int_{h/g}^\infty e^{-s}s^{(\frac{-r}{g})-1}ds \tag{28}$$

From the messy looking equation (28) it might not *appear* as if I have improved matters, but the integral can actually be identified as the Incomplete Gamma (IG) function:

$$\Gamma(\alpha,\beta) \;=\; \int_\beta^\infty e^{-s}\,s^{\alpha-1}ds. \tag{29}$$

When the lower bound of integration $\beta = 0$ the IG function collapses to the basic Gamma function and when α is an integer then $\Gamma(\alpha, 0) = (\alpha-1)(\alpha-2)$... etc., a.k.a. $(\alpha-1)$ factorial, with the understanding that both $\Gamma(1, 0) = 1$ and $\Gamma(2, 0) = 1$.

For general values of α and β the IG function is readily available in most business and scientific software packages (as well as R, of course), similar to the *error function* or the normal distribution. For example, the value of $\Gamma(-0.5, 1) = 0.178148$ to five digits and the value of: $\Gamma(-0.5, 0.3678) = 0.89635$. I do caution that for non-positive values of α there are some numerical stability issues. Merging equation (29) and (28), I can write the annuity factor using a closed-form expression:

$$\mathtt{a}(r, h, g) \;=\; \frac{\Gamma(-r/g, h/g)}{g\,\exp\{(-1/g)\,(h + r\ln[h/g]\,\}}, \tag{30}$$

This is our basic annuity factor. Here are some numerical examples. Let's arbitrarily set the interest rate $r = 3\%$. If we keep g constant at 0.08, here is what happens when h_x takes values 0.1 , 0.2, and 0.3, respectively. $\mathtt{a}(0.03, 0.1, 0.08) = 5.552432$, $\mathtt{a}(0.03, 0.2, 0.08) = 3.464195$, and: $\mathtt{a}(0.03, 0.3, 0.08) = 2.543422$ Note that as h increases, the value of the annuity factor declines. This should be intuitive because the force of mortality (or hazard rate) kills you, so if it increases you will live a shorter life and thus receive less income, making the annuity factor cheaper. Likewise, if we fix: $h = 0.1$ and change g to take values 0.09, 0.12, 0.15, respectively. Then, $\mathtt{a}(0.03, 0.1, 0.09) = 5.392625$, and: $\mathtt{a}(0.03, 0.1, 0.12) = 4.981276$, and finally: $\mathtt{a}(0.03, 0.1, 0.15) = 4.646376$, all in dollars. So, as g increases, the value of the annuity factor declines.

Figure #7 shows this graphically. The x-axis represents the current (or initial) hazard rate h, ranging from $h = 0.005$ to $h = 0.5$ in equal increments. This happens to correspond to a range of: $x = 64$ to $x = 101$, assuming $m = 90$ and $b = 8$, in the alternate (m, b) specification. The x-axis is labeled with both a rate (top) and an age (below). And, while the former (top row) increases linearly, the latter (bottom row) does not, since $x := b\ln[hb] + m$ under the Gompertz law. The y-axis in figure #7a represents the annuity factor corresponding to that particular (x-axis) hazard rate and age, assuming: $g = 12.5\%$ in the above (h, g) formula. Notice how the annuity factor declines, as we move from left to right and the hazard rate (as well as the age) increases. Panel #7b plots the value of the annuity factor over the same range of $h = 0.005$ to $h = 0.5$, assuming a reduced: $g = 8\%$. The annuity factor is higher at every value (as one can see from the few points that are highlighted), although it also declines in hazard rate and/or age. Note that in this case the corresponding $b = 1/g = 12.5$ years, and the corresponding (second) x-axis values range from $x = 55$ to 113.

A.3.1 Gompertz (m,b)

Recall that under the Gompertz (m, b) formulation, common in actuarial finance, the hazard rate is expressed as: $\lambda_{x+t} = (1/b)e^{(x+t-m)/b}$. In that case, the immediate annuity factor $\mathtt{a}$ can be derived in terms of (x, m, b) by substituting: h with $(1/b)\exp\{\frac{x-m}{b}\}$ and replacing g with $1/b$. For completeness I present:

$$\mathtt{a}(r, x, m, b) \;=\; \frac{b\Gamma(-rb, \exp\{\frac{x-m}{b}\})}{\exp\{r(m - x) \;-\; \exp\{\frac{x-m}{b}\}\}}. \tag{31}$$

A.4 Deriving the Annuity Equivalent Wealth or δ

Here I sketch a quick proof of the expression for the annuity equivalent wealth (AEW): $1+\delta$, under the assumption mortality obeys a Gompertz law of mortality and the individual has **no** pre-existing annuity income. This is based on the derivation in Milevsky and Huang (2018), who provide various closed-form expressions for the AEW under alternate mortality assumptions and more general (non-zero) pension income. See also the excellent reference book by Cannon and Tonks (2008), and in particular the derivation in Chapter 7.6. Let $u(c)$ denote a constant relative risk aversion (CRRA) utility (a.k.a. felicity) function parameterized by risk aversion γ, and a subjective discount rate $\rho = r$. Formally, $u(c) = c^{1-\gamma}/(1-\gamma)$. The maximal utility *without* annuities is:

$$U^*(w) \;= \max_{c_t} \int_0^{\omega - x} e^{-rt}\, p(t, h^i, g^i)\, u(c_t)\, dt, \tag{32}$$

where (h^i, g^i) are the Gompertz parameters at the relevant income percentiles, and the budget constraint is:

$$dW_t \;=\; (rW_t \;-\; c_t)\, dt, \;\; W_0 = w. \tag{33}$$

The optimal consumption function is denoted by: c_t^*, and do not allow any borrowing so that wealth $W_t \geq 0$ at all times. The only reason to prefer early vs. late consumption, is due to mortality beliefs and the inter-temporal elasticity of substitution, $1/\gamma$. Now, since there is no pre-existing pension income, the relevant consumption rate must be sufficient to last all the way to ω, so that:

$$w = c_0^* \int_0^{\omega - x} e^{-rt} p(t, h^i, g^i)^{1/\gamma} dt, \tag{34}$$

which leads to the corresponding:

$$c_t^* \;=\; \left(\frac{w}{\int_0^{\infty} e^{-rt} p(t, h^i, g^i)^{1/\gamma} dt} \right) p(t, h^i, g^i)^{1/\gamma}, \tag{35}$$

The integral in the denominator of equation (35) is an annuity factor (of sorts) assuming the survival probability $p(t, h^i, g^i)$ is shifted by $1/\gamma$. For example, when $\gamma = 1$, the optimal consumption function c_t^* in equation (35) collapses to the hypothetical annuity: $w/\mathtt{a}(r, h^i, g^i)$ times the survival probability $p(t, h^i, g^i)$, which is less than what a life annuity would have provided. The individual who converts all liquid wealth w into the annuity would consume $w/\mathtt{a}$, but the non-annuitizer reduces consumption in proportion to survival probabilities as a precautionary measure. In contrast to: $U^*(w)$, let $U^{**}(w)$ denote discounted lifetime utility of wealth, assuming wealth w **is entirely** annuitized or pooled at age x. Discounted utility is:

$$U^{**}(w) \;=\; \int_0^{\infty} e^{-rt}\, p(t, h^i, g^i)\, u(w/\mathtt{a}(r, \hat{h}, \hat{g}))\, dt, \tag{36}$$

where the optimized consumption path is trivially $c_t^* = w/\mathtt{a}(r, \hat{h}, \hat{g})$, for all t. The next (and essentially final) step is to note that δ will satisfy the following equation:

$$U^*((1+\delta_i)w) \;= U^{**}(w), \tag{37}$$

as per the definition of AEW. We refer the interested reader to Milevsky and Huang (2018) for the algebra that extracts: δ_i from the above equation. **Q.E.D.**

Table #1:						
U.S. Death Rates per 1,000 individuals.						
	MALE			**FEMALE**		
Income Group	**Age=40**	**Age=50**	**Age=60**	**Age=40**	**Age=50**	**Age=60**
Lowest (1st pct.)	5.8	12.5	22.1	4.3	8.0	12.8
25th percentile	2.0	4.5	10.9	1.2	2.7	5.9
Median (50th pct.)	1.2	2.9	7.3	0.8	2.0	4.5
75th percentile	0.8	1.8	4.9	0.5	1.3	3.5
Highest (100th pct.)	0.6	1.1	2.8	0.3	0.8	2.2

Data source: Chetty et al. (2016). Mortality over the period: 2001 to 2014 with income lag of two years. Note how the ratio of *high-to-low* mortality rates decline at higher ages.

Table #2:				
Mortality Growth Rate and Centenarian Projection				
Income	**Male**		**Female**	
Group	$g := \frac{\ln q_{x+T}}{\ln q_x}/T$	$1000\tilde{q}_{100}$	$g := \frac{\ln q_{x+T}}{\ln q_x}/T$	$1000\tilde{q}_{100}$
Lowest (1st pct.)	5.63%	208.5	4.81%	89.2
25th Percentile	8.74%	355.2	8.32%	171.1
Median (50th pct.)	9.21%	285.7	8.68%	158.9
75th Percentile	10.22%	302.1	10.21%	211.4
Highest (100th pct.)	10.00%	163.2	9.97%	115.0

Data source: Chetty et al. (2016). Mortality over the period: 2001 to 2014 with income lag of: two years. Calculation (by author) of g was based on growth from age $x = 50$ to $x = 63$. Each percentile's g value was used to forecast: $\tilde{q}_{100}$. Note that by age $x = 100$, mortality rates are within a multiple of two of each other.

Table #3								
Computing the individual Volatility of Longevity (iVoL): φ_x								
...when modal value of life (m) is: **98 years**.								
Mortality	**Disp.**	**Mean**	**StDev.**	**iVoL**	**Hazard**	**Mean**	**StDev.**	**iVoL**
Growth	b	$E[T_0]$	$SD[T_0]$	φ_0	λ_{65}	$E[T_{65}]$	$SD[T_{65}]$	φ_{65}
11.5% (**rich**)	8.696	92.98	11.15	12.0%	0.2586%	28.82	9.70	**33.7%**
10.5%	9.524	92.51	12.20	13.2%	0.3284%	28.68	10.29	35.9%
9.5%	10.526	91.93	13.47	14.7%	0.4132%	28.59	10.95	38.3%
8.5%	11.765	91.23	15.00	16.4%	0.5143%	28.59	11.69	40.9%
7.5%	13.333	90.37	16.90	18.7%	0.6312%	28.72	12.55	43.7%
6.5%	15.385	89.30	19.26	21.6%	0.7609%	29.08	13.58	46.7%
5.5%	18.182	87.99	22.21	25.2%	0.8956%	29.83	14.89	49.9%
...when modal value of life (m) is: **78 years**.								
Mor.Gro	b	$E[T_0]$	$SD[T_0]$	φ_0	λ_{65}	$E[T_{65}]$	$SD[T_{65}]$	φ_{65}
11.5%	8.696	72.99	11.11	15.2%	2.5789%	12.30	6.55	53.3%
10.5%	9.524	72.53	12.14	16.7%	2.6815%	12.64	6.90	54.6%
9.5%	10.526	71.97	13.35	18.5%	2.7629%	13.08	7.33	56.0%
8.5%	11.765	71.32	14.79	20.7%	2.8153%	13.66	7.85	57.5%
7.5%	13.333	70.55	16.51	23.4%	2.8289%	14.43	8.51	59.0%
6.5%	15.385	69.65	18.57	26.7%	2.7921%	15.49	9.36	60.4%
5.5% (**poor**)	18.182	68.69	21.07	30.7%	2.6906%	17.00	10.53	**61.9%**

These numbers (and ranges) are for illustrative purposes and based on theoretical Gompertz values, motivated by the Chetty et al. (2016) data, and do not correspond to any specific income percentile or group.

Table #4						
Annuity Equivalent Wealth (AEW) and Value of Longevity Risk Pooling:						
Demographics					**Risk Aversion:** $\gamma = 3$	
Income Percentile	**Gompertz** (h_{65}, g)	$E[T_{65}]$ **Years**	φ_{65} **iVoL**	**Annuity Factor $**	$\delta :=$ **AER-1 Individual**	$\delta :=$ **AER-1 Group**
FEMALE						
Lowest	(1.64%, 5.29%)	22.75	56.73%	15.23	62.18%	46.52%
5th	(1.22%, 6.68%)	23.36	51.05%	15.72	53.59%	43.18%
10th	(1.15%, 8.08%)	21.52	48.63%	14.97	52.45%	35.39%
20th	(1.00%, 8.9%)	21.58	46.33%	15.08	49.20%	33.46%
30th	(0.86%, 8.61%)	23.41	45.24%	15.96	45.87%	38.06%
40th	(0.78%, 8.84%)	23.94	44.11%	16.23	43.93%	38.57%
50th	(0.69%, 8.73%)	25.31	43.10%	16.86	41.46%	41.46%
60th	(0.69%, 10.06%)	23.13	41.91%	15.94	41.89%	34.20%
70th	(0.55%, 9.08%)	26.68	40.95%	17.51	37.82%	43.18%
80th	(0.50%, 10.35%)	25.30	39.15%	17.00	36.81%	37.97%
90th	(0.45%, 10.49%)	26.05	38.15%	17.36	35.12%	39.12%
95th	(0.38%, 9.74%)	28.89	37.46%	18.53	32.44%	45.61%
Highest	(0.34%, 9.89%)	29.72	36.46%	18.89	30.89%	46.66%
MALE						
Lowest	(3.02%, 6.56%)	14.76	61.24%	11.15	84.26%	38.25%
5th	(2.10%, 6.63%)	17.93	56.97%	12.96	70.29%	48.52%
10th	(2.00%, 7.46%)	17.37	55.13%	12.72	68.20%	43.96%
20th	(1.75%, 8.31%)	17.55	52.56%	12.89	63.64%	41.90%
30th	(1.50%, 8.78%)	18.29	50.41%	13.33	59.15%	42.76%
40th	(1.18%, 8.43%)	20.82	48.41%	14.65	52.95%	50.76%
50th	(1.06%, 8.83%)	21.16	46.97%	14.86	50.53%	50.53%
60th	(0.89%, 8.68%)	23.01	45.45%	15.77	46.54%	55.47%
70th	(0.82%, 9.31%)	22.75	44.08%	15.70	45.01%	53.18%
80th	(0.70%, 9.49%)	23.81	42.58%	16.23	42.14%	55.20%
90th	(0.60%, 9.8%)	24.68	40.98%	16.67	39.45%	56.39%
95th	(0.51%, 9.68%)	26.26	39.83%	17.38	36.87%	60.09%
Highest	(0.42%, 8.74%)	30.13	39.01%	18.93	33.24%	69.77%
Assumes $r = 3\%$. No explicit insurance loading, other than when paying group rates.						

Table # 5							
Global 1940 (unisex) Cohort Mortality							
The individual volatility of longevity (iVoL) at age 65							
OECD	**Gompertz Parameters**				**Mean**	**St.Dev**	**iVoL**
COUNTRY	**m**	**b**	$\mathbf{h_{65}}$	**g**	$E[T_{65}]$	$SD[T_{65}]$	φ_{65}
Latvia	75.02	11.87	3.62%	8.42%	11.95	7.22	60.4%
Estonia	77.26	12.77	3.00%	7.83%	13.71	8.10	59.1%
Hungary	76.34	11.21	3.24%	8.92%	12.38	7.25	58.6%
Slovakia	78.92	11.58	2.60%	8.64%	14.14	7.98	56.4%
Poland	79.69	11.75	2.44%	8.51%	14.71	8.22	55.9%
Portugal	82.40	13.28	2.03%	7.53%	17.14	9.47	55.3%
Czechia	80.88	11.18	2.16%	8.94%	15.24	8.24	54.1%
Slovenia	85.62	14.11	1.64%	7.09%	19.63	10.52	53.6%
USA	83.81	12.22	1.76%	8.18%	17.65	9.31	52.7%
Finland	85.86	13.32	1.57%	7.51%	19.48	10.23	52.5%
Spain	85.64	12.97	1.57%	7.71%	19.20	10.03	52.2%
France	88.20	14.36	1.38%	6.96%	21.51	11.18	52.0%
Belgium	85.07	12.39	1.60%	8.07%	18.59	9.65	51.9%
Denmark	83.47	11.39	1.73%	8.78%	17.11	8.88	51.9%
Luxembourg	88.57	13.77	1.31%	7.26%	21.56	10.99	51.0%
Austria	88.75	13.68	1.29%	7.31%	21.66	10.99	50.7%
UK	84.89	11.30	1.52%	8.85%	18.09	9.13	50.5%
Canada	87.83	12.74	1.31%	7.85%	20.69	10.38	50.2%
Italy	87.72	12.61	1.31%	7.93%	20.57	10.30	50.1%
Ireland	87.07	12.15	1.34%	8.23%	19.95	9.96	49.9%
New Zealand	88.76	13.05	1.24%	7.66%	21.46	10.70	49.9%
Norway	87.89	11.99	1.24%	8.34%	20.51	10.05	49.0%
Iceland	88.43	12.20	1.20%	8.20%	20.97	10.25	48.9%
Australia	90.78	13.40	1.09%	7.46%	23.07	11.27	48.9%
Sweden	88.52	12.13	1.19%	8.24%	21.02	10.24	48.7%
Netherlands	86.41	11.02	1.30%	9.07%	19.13	9.31	48.7%
Switzerland	90.39	12.77	1.07%	7.83%	22.60	10.90	48.2%
Japan	91.72	12.87	0.97%	7.77%	23.64	11.20	47.4%
Israel	89.04	11.51	1.08%	8.69%	21.25	10.04	47.2%
Greece	88.49	10.66	1.04%	9.38%	20.64	9.53	46.2%
AVERAGE	**85.45**	**12.41**	**1.55%**	**8.06%**	**18.98**	**9.72**	**52%**

Figure 1: Remaining lifetime, dispersion and individual volatility of longevity (iVoL).

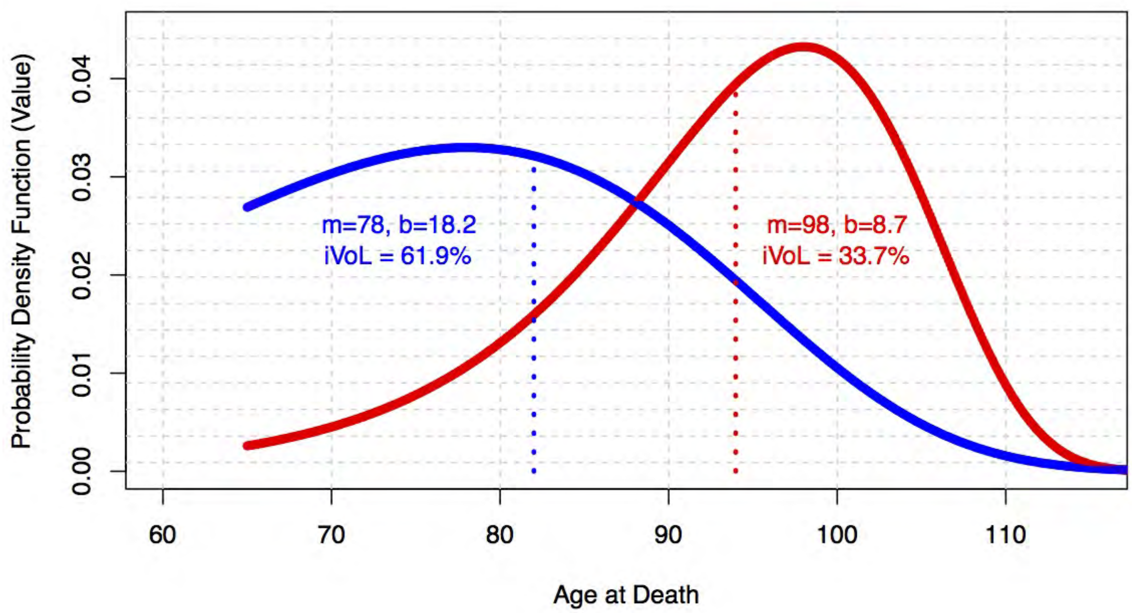

Note: Both curves are based on the Gompertz law of mortality, conditional on age $x = 65$. The right curve represents a retiree with a higher life expectancy ($m = 98$) and lower dispersion parameter ($b = 8.7$), leading to an *individual volatility of longevity* of 33.7%. The left curve represents a retiree with lower life expectancy ($m = 78$) and higher dispersion ($b = 18.2$), whose *individual volatility of longevity* is double, at: 61.9%.

Figure 2: Mortality Rates and Logarithms: Income Percentiles

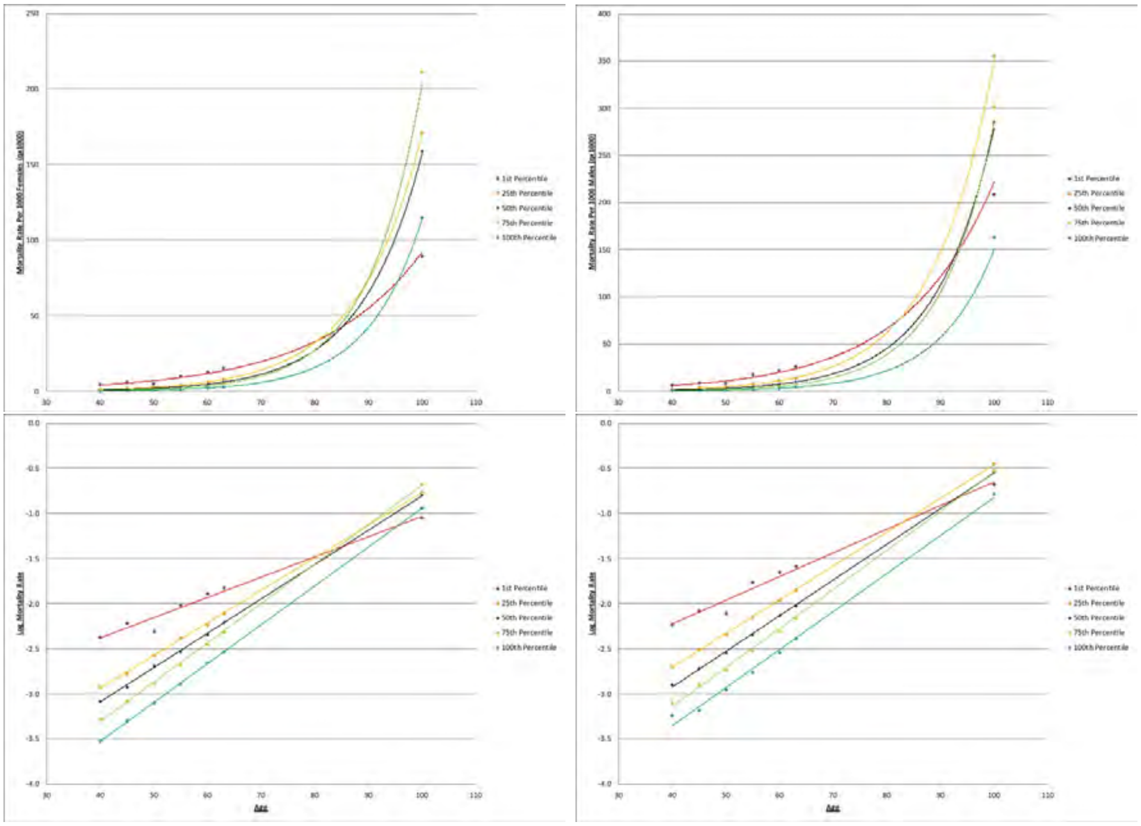

Note: Left panels are female. Right panels are male. The highlighted points are from the Chetty et al. (2016) data for mortality by income percentile. The $\ln q_x$ lines (bottom) and q_x curves (top) are extended to age 100 based on the regression coefficients explained in the paper. The convergence in mortality rates (which is clearer when viewed in: $\ln q_x$ rates) is a manifestation of the so-called *Compensation Law of Mortality.*

Figure 3: The Individual Volatility of Longevity (iVoL) over the Lifecycle.

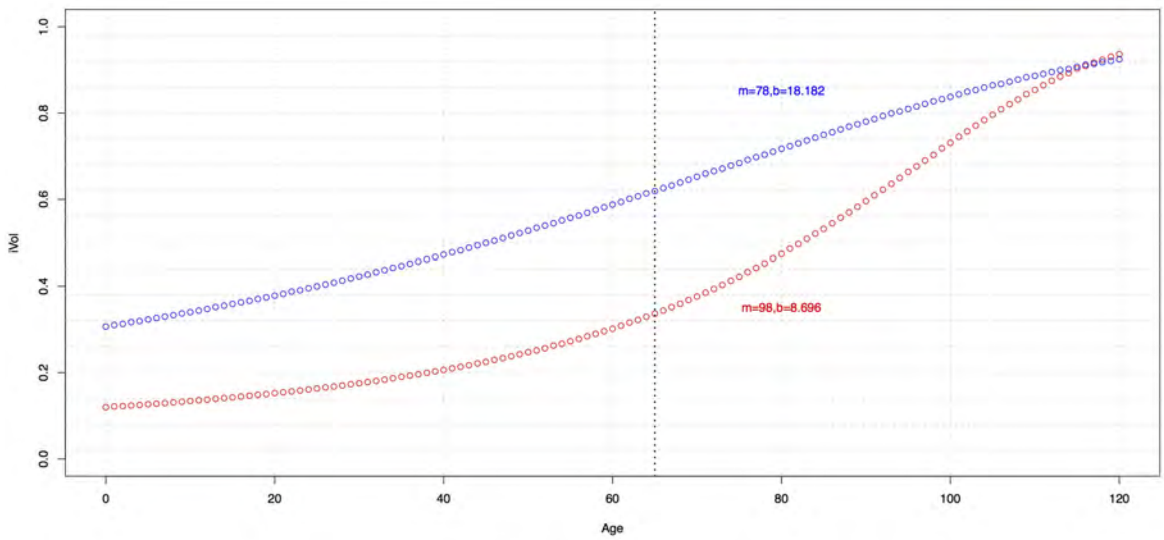

Note: The y-xis is defined as the ratio of the standard deviation of expected lifetime $SD[T_x]$ to mean lifetime $E[T_x]$ at age x. The vertical line is at age $x = 65$, with exact values noted in table (#3). The Gompertz parameters are: $m = 98, b = 8.696$ (rich) and $m = 78, b = 18.182$ (poor.) Notice how both curves and the iVoL converge to a value of $\varphi = 1$ at advanced ages.

Figure 4: Annuity Equivalent Wealth (AEW) under a range of values for m.

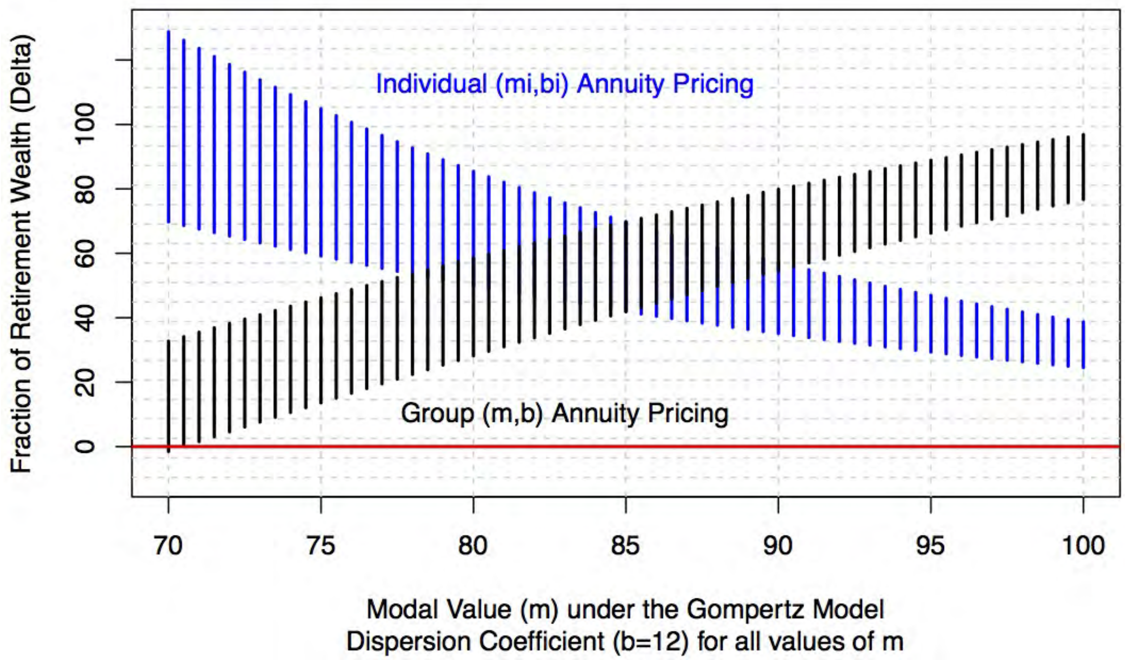

Note: Figure is based on a valuation rate $r = 3\%$ and coefficient of relative risk aversion from $\gamma = 5$ (top) to $\gamma = 1$ (bottom). When annuities are priced based on individual mortality $(m_i, b = 12)$, the value of AEW declines in m_i. But when annuities are priced based on group mortality $(\hat{m}, b = 12)$, the value of AEW increases in m_i, due to the implicit loading.

Figure 5: Estimated Gompertz Parameter Values versus Income Percentiles

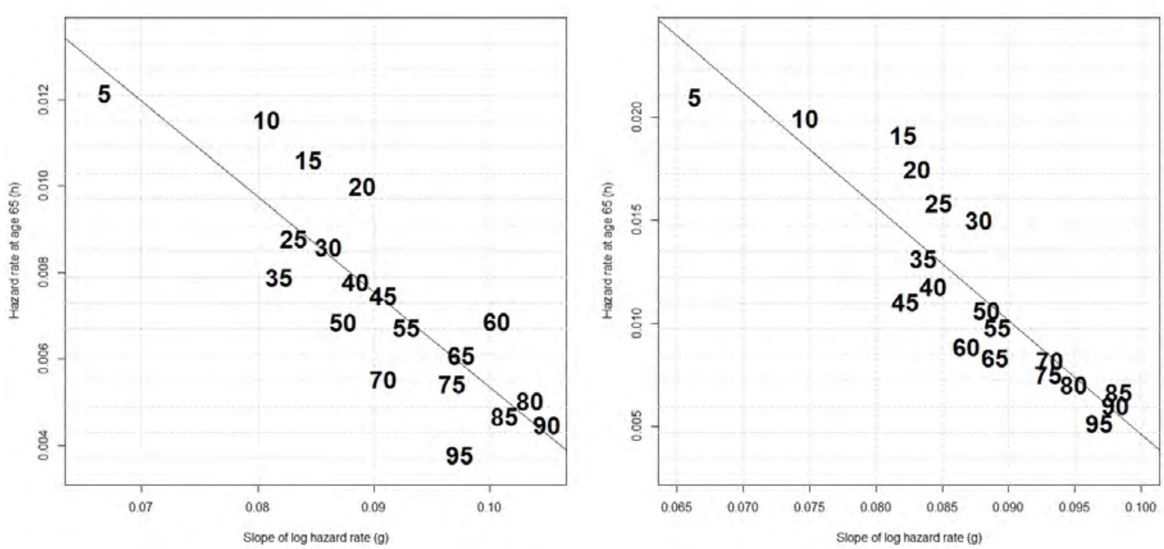

Note: Using data from Chetty et al. (2016), average mortality rates during the 2001 to 2014 period are converted into Gompertz (h, g) parameters for various income percentiles. Each percentile point noted in the figure (left panel for females and right for males) is placed at the co-ordinate for the relevant value. Generally speaking higher (wealthier) income percentiles (numbers) are located at the bottom right and lower percentiles are at the top left. Wealthier retirees have lower mortality rates, but age faster.

Figure 6: Individual Volatility of Longevity at Age 65: A Global Comparison

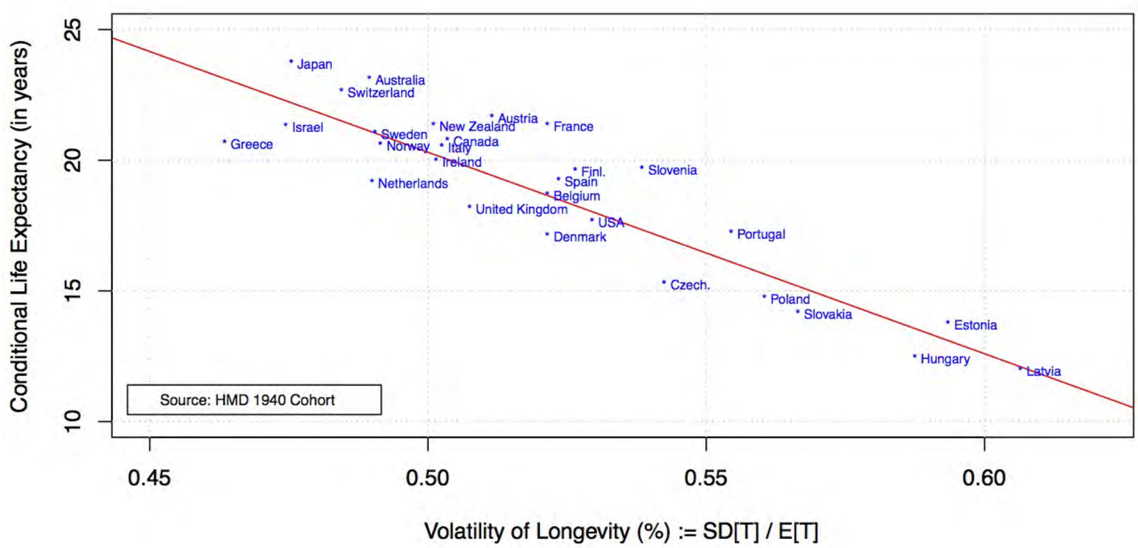

Note: Based on HMD data from OECD countries. The entire mortality rate curve for the 1940 cohort is converted into Gompertz parameters using the procedure described in appendix A1, and then mapped into the *individual volatility of longevity* at age 65, which is the ratio of the standard deviation to the mean and denoted by φ_{65}.

Figure 7: The Annuity Factor in (h, g)-space

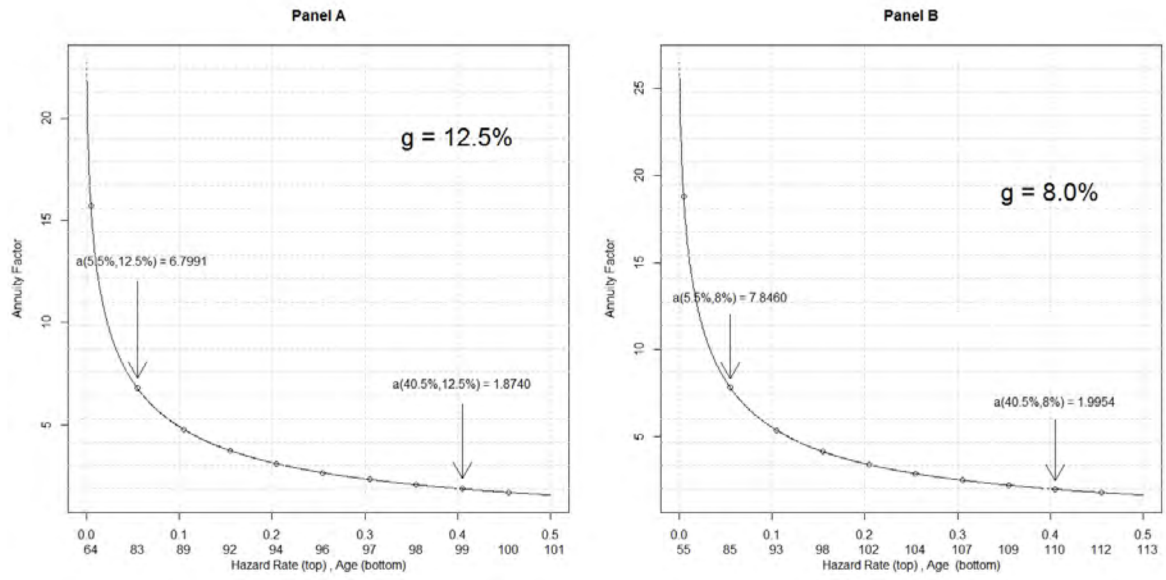

Note: The cost of \$1 lifetime income is cheaper (and the annuity valuation factor is lower) when the current hazard rate ($h_x = h$) is higher and/or the hazard growth rate (g) is higher. Both are equivalent to increasing the discount rate (r) in the present value factor. Technically; $\partial \mathsf{a}(r, h, g)/\partial r \leq 0$, $\partial \mathsf{a}(r, h, g)/\partial h \leq 0$, and: $\partial \mathsf{a}(r, h, g)/\partial g \leq 0$.