



Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

Population age structure and secular stagnation: Evidence from long run data

Joseph Kopecky

TEP Working Paper No. 0523

February 2023

Trinity Economics Papers
Department of Economics



Full length article

Population age structure and secular stagnation: Evidence from long run data

Joseph Kopecky

Trinity College Dublin, Ireland



ARTICLE INFO

Keywords:

Secular stagnation
Demographic transition
Population aging
Economic growth

ABSTRACT

A large literature has reopened the secular stagnation hypothesis, first proposed near the end of the great depression as a warning for anemic growth resulting from long run trends in population aging. In this paper, I explore the relationship between population age structure and growth in: investment, consumption and output, in a long run panel of advanced economies. The evidence is largely consistent with proposed channels for secular stagnation. Investment growth, in its level and as a fraction of GDP, appears much stronger in young populations, while facing demographic headwinds in older economies. Consumption and output growth are positively associated with late career workers, with a negative relationship coming from both young and old dependents. Consistent with the recent secular stagnation literature, interest rate channels appear to have strong interactions with population age structures. I find that for investment and output growth, estimated impacts of age-structure are more pronounced in low interest rate environments, with high rates mitigating some of their effect.

Introduction

How does population aging affect growth of investment, consumption and output? A large literature has reopened the *secular stagnation* hypothesis, first proposed by Hansen (1939) who predicted that downward trends in population growth would lead to declines in investment, and through this act as a drag on growth. The central arguments of his 1938 address to the American Economic Association are that growing populations require increasing investment expenditures, such as housing, while stagnant or declining populations would see a dearth of investment opportunity. He describes a future more heavily reliant on technology driven growth, while cautioning that such growth is likely to be “discontinuous, lumpy, and jerky” (Hansen, 1939). While these prognostications never materialized in the 20th century, it was not possible to foresee the incredible influence that the post-war baby boomer cohort would exert on the populations of developed economies. Recently, Summers (2014a,b), and others¹ have revived the secular stagnation hypothesis in light of the anemic growth of advanced economies following the global financial crisis. Here, among other factors, population aging is cited as reducing investment demand, pushing down natural rates, slowing growth, and making zero lower bound (ZLB) episodes more likely. Many of the mechanisms proposed in this new literature appeal to similar logic as those in Hansen (1939).

Is secular stagnation consistent with the long run empirical data? Certainly the views put forth in Hansen (1939) look incorrect given the post-war track record of advanced economies. However, the unprecedented reversal of long run trends in fertility during the baby boom reflect a dramatic change in the underlying population dynamics on which the hypothesis was founded. Were these predictions an incorrect reading of the empirical relationship between age structure and growth rates or just badly timed? This paper tests the relationship between age structure and these macroeconomic variables in the long run, using historical macroeconomic data going back to 1870. In line with a large literature on the importance of demographics, I find that the age structures of advanced economies do appear to be strongly related to growth rates in investment, consumption, and output. Aging economies are estimated to have a drag on both the growth rate of investment and of investment-to-output ratios. Demographic headwinds to consumption and output look poised to push growth rates down as the baby boomer cohort moves fully into retirement. In the case of real output and investment growth, these estimates appear to be strongly correlated to the interest rate environment, suggesting that low interest rates may magnify this relationship.

A key element of the (Hansen, 1939) address is growth in investment. He is clear that this can come from capital *deepening* through

E-mail address: jkopec@tcd.ie.

¹ This literature has since grown substantially, see Krugman (2014) for another early call to revitalize the secular stagnation hypothesis in relation to current weakness in growth.

² Acemoglu and Restrepo (2017) suggest more rapid adoption of automation technology in countries with older populations.

³ See, for example: Eggertsson et al. (2019b), Carvalho et al. (2016), Kopecky and Taylor (2020), and Gagnon et al. (2021), though notably in some of these models capital-to-output/labor ratios are fixed in equilibrium.

increases in investment relative to output, or capital *widening*, the expansion of capital *pari passu* with output. Much of the modern secular stagnation literature is not clear on which of these should matter. Speaking directly to the question of secular stagnation, Eggertsson et al. (2019a) explore the empirical relationship between aging and output growth, reconciling the positive estimates found in Acemoglu and Restrepo (2017).² Eggertsson et al. (2019a) show for recent decades that positive growth associated with aging, measured as the ratio of old-age population to working-age population, comes during non-ZLB periods when aging works to push interest rates down. This price mechanism leads to capital deepening, which should have positive growth impacts, *ceteris paribus*. They show in a simple model that during ZLB periods, when rates cannot accommodate these demographic forces, that investment as a ratio of output falls. This is consistent with their empirical evidence that old-age dependency ratios are negatively correlated with growth at times when the economy faces this effective lower bound. This interest rate effect has now been widely studied in life-cycle macroeconomic models,³ with empirical evidence of a long run relationship between them in Lunsford and West (2019) and Kopecky and Taylor (2020). I will not test the ZLB relationship directly as such episodes are fairly rare in the historical data. However, I show that the strength of estimates of age structure do vary with interest rates, consistent with general equilibrium price mechanisms. These work to offset the impact of population age when rates are higher, reinforcing them when rates are low.

Another large literature explores the impact of fertility and aging on long run macroeconomic development. This includes the unified growth theories suggesting demographic channels as important long run drivers of growth in transitions from a Malthusian economy to a modern growth path as in the seminal works of Galor and Weil (2000), Galor and Moav (2004). In such a model, endogenous fertility choices due to *quantity-quality* tradeoffs, from the work of Becker (1960), lead to lower population growth and increases in human capital. Cervellati and Sunde (2005) complement this literature suggesting a positive feedback between life expectancy and human capital formation, which itself drives technological progress. Cervellati and Sunde (2011) show that post-demographic transition economies experience positive relationships between life expectancy and growth, supporting this channel as a mechanism. The economies I study will be, in general, post-transition, with some still experiencing the tail end of their falls in fertility at the start of the 1870 sample.⁴ While my focus will be on age structure rather than these long run development trends in mortality and fertility it is important to keep these mechanisms in mind as they ultimately determine the long run shape of age distributions.

Empirically there have been a number of papers establishing the relationship between age structure and macroeconomic variables. Work such as Bloom and Williamson (1998) and Bloom and Finlay (2009) suggest that demographics can account for large fractions of the *economic miracle* of rapid growth experienced in East Asia during the latter half of the 20th century. Lee and Mason (2010) and Mason et al. (2016) describe two dividends coming from demographic transitions. The first dividend comes due to baby boom-bust dynamics that create a relatively large cohort in the population. As they move into working age, increases in labor per-capita provide a boost to growth that in principle should persist as long as this group remains economically active. A second potential demographic dividend comes from the excess savings of this group as they near retirement, with savings accumulation promoting capital deepening. Cooley and Henriksen (2018) use growth accounting and a quantitative life-cycle model to quantify the various channels through which demographics will contribute to an upcoming demographic *deficit* as the United States and Japan age structures move past this dividend. They find pronounced demographic headwinds for

future growth due to both intensive and extensive margins of labor supply, as well as through lower measured TFP growth. These are offset only partially in their model by capital accumulation due to larger shares of households in these high savings groups, suggesting that this second dividend may slow some of the labor market effects of aging, but not all. Similar evidence is found in Cooley et al. (2019) for European countries. Bloom et al. (2010) estimate that the negative implications of moving to an old population will be relatively small for OECD countries, citing changes in behavior as a result. Both this second dividend, through life-cycle savings mechanisms and behavioral differences in young and old populations suggest that analysis should consider not just population growth rates, or working age population, but also the full age structure.

To bring a flexible approach to studying age structure, I use a methodology first proposed by Fair and Domínguez (1991), who fit parameters for the entire age distribution to understand how various macroeconomic equations (consumption, housing investment, and money demand) correlate with age. They find evidence in favor of Ando and Modigliani (1963) style effects with households consuming a lower share of their income in prime ages, and for money demand as might be expected from a (Baumol, 1952; Tobin, 1956) style model with higher opportunity costs of prime age workers. This methodology is adopted for the question of investment and international capital flows in Higgins (1998), Higgins and Williamson (1997). My estimates on investment-to-gdp growth are extensions of some of the results from Higgins (1998), adding the nearly thirty years of new data available from after 1990, while also extending much further back in time. Unlike Higgins (1998), I focus on a subset of advanced economies for which this long historical panel is available. Given the slow moving nature of demographics, this trade-off of increasing the time series dimension at the expense of cross sectional coverage is instructive. This likely limits my ability to understand the impact that aging has had on cross border flows,⁵ which may arise as older countries seek investment opportunity abroad from younger countries. My longer time series should better clarify the empirical relationships between aging and domestic investment in well developed economies, a key question in the secular stagnation literature and likely more quantitatively important for the long run growth in this subset of advanced economies. In recent work, Aksoy et al. (2019) use an unconditional VAR to show that changes in the distribution of population across age groups from 1970–2014 leads to about a 2% decline in investment-to-gdp and about a 1.25% decline in output growth in OECD countries. They show that these results are consistent with a model where productivity of innovative sectors depends on the stock of young workers. My estimates will show a population age relationship consistent with those found in Higgins (1998), with quantitative implications on a similar scale to those found in Aksoy et al. (2019).

I contribute to existing literature in a number of ways. First, I take advantage of a longer panel of data to provide evidence in line with existing estimates from Higgins (1998) who studies the period from 1950 to 1989, and Aksoy et al. (2019) who use a panel of OECD countries from 1970–2014. My estimates will not rely as heavily on the periods with the large baby boom cohort in the labor force, and will allow for better historical context for this inherently long run question. If the baby boomer group was a temporary reversal of the trends driving secular stagnation in growth it will be important to study demographic movements that substantially predate their entry into the labor force. Second, I contribute to the extensive literature on correlations between age structure and growth, again using a longer

⁴ Most of the countries where my data extends furthest back in time are those that transitioned relatively early.

⁵ In addition to Higgins (1998), for more on cross border impacts of aging see Taylor and Williamson (1994), who study the role that capital flows act as an intergenerational transfer, as well as Backus et al. (2014) and Lane and Milesi-Ferretti (2001) who find them to be an important component of low frequency capital movements.

run panel of OECD countries than many existing estimates, but also applying the more flexible methodology of Fair and Domínguez (1991) to test correlations of age structure with fewer a priori assumptions regarding the important cutoffs in the age distribution.

In Section “Data and methodology”, I describe my data and estimating equations. In Section “Empirical results”, I present the main results for the relationship between age structure on a number of economic outcomes. First I show that both investment and investment-to-gdp ratios have similar correlations with age structure, then I explore the relationship between age structures and consumption and output growth. In Section “Implications for aging and secular stagnation: A quantitative exercise”, I use the results from Section “Empirical results” to show the quantitative implications of my primary estimates. I illustrate the estimated impact of age structure on growth rates in the past, as well as projected future demographic influences out to the year 2050. Section “Conclusions” provides further discussion and concludes.

Data and methodology

I primarily use two sources of data. The first is the long run macroeconomic panel from Jordà et al. (2017) (JST). This gives macroeconomic data from 18 countries going back as far as 1870, with financial variables included since the update from Jordà et al. (2019). I merge this with population data from the Human Mortality Database (2020), (HMD). I construct population shares for five-year groups using their population-by-age data, which in turn will be used to construct my preferred demographic controls. For the United States I augment this for years before 1933 with census data to ensure a sample going back to 1900. I also use data from the United Nations World Population Prospects 2019 (UN, 2019), which provides estimates of forward projections of population age structures to 2100. Not all variables are complete for the full 1870–2018 sample. In my analysis, estimates will be made on data with an average panel length of 102.2 years. The shortest is 61 years for Ireland, whose data only go back to 1960, with the longest, Sweden, covering the full 147 years of data available. All countries begin before the large boomer cohort enters the workforce, with most having substantial data from before they were born.

I report estimates using population growth and old-age dependency ratios to provide context relative to the literature and some easy to interpret baselines. However, my preferred estimates seek to understand the underlying correlations between age structure and outcomes in a way that does not provide a priori assumptions on their relationships. As such my preferred specifications will use a modified version of the (Fair and Domínguez, 1991) method for controlling for population age structure. Consider the naive equation:

$$g_{i,t} = \sum_j \alpha_j \Delta p_{j,i,t} + X_{i,t} \beta + \mu_i + \mu_t + \nu + \epsilon_{i,t} \quad (1)$$

Where $g_{i,t}$ is some outcome of interest for a particular country i at time t , the growth rate of output for example. The age distribution is divided into J groups, and the term $\Delta p_{j,i,t}$ is the annual change in the share of population in age group j in country i at time t . Including controls ($X_{i,t}$), fixed effects (μ_i, μ_t), and a regression constant (ν), the primary object of interest is α_j , which measures the linear relationship between change in the share of a particular age group and the outcome. This might be preferred to using simple age structure measures such as old-age dependency ratios, or working age populations, because it allows the data to determine where important age cutoffs are. Further there is concern that correlation between movements in age shares might attribute a relationship to one variable that may belong to another. Population age shares are strongly correlated with each other since growth in the share of one must lead to a fall in others. Further, as cohorts move through the age distribution, growth in one group today will lead to future growth in subsequent ages. While this provides motivation to use complete and relatively fine age bins, as in Eq. (1), it also creates challenges for estimation. For one, it is not possible to jointly estimate

coefficients on all perfectly colinear terms⁶ at once (while keeping a regression constant). Second, and likely more important, is that due to the strong colinearity between age groups, increasing in the J number of groups studied, can lead to parameter instability in α_j . Finally there could be, in smaller samples, a power issue from increasing the size of J . Two assumptions allow for: joint estimation of all α_j terms, stability of adjacent parameter estimates of α_j , and tractability.

$$\sum_j \alpha_j = 0 \quad (2)$$

$$\alpha_j = \theta_0 + \theta_1 j + \theta_2 j^2 + \theta_3 j^3 \quad (3)$$

Eq. (2) facilitates joint estimation of all age coefficients with a regression constant. This is necessary due to the perfect colinearity of all population shares and without it one share must always be dropped from the estimation. In principle this is enough to econometrically estimate all of the α_j coefficients of interest for growth rates on each population age share. However, because of the issue of highly colinear shares described above, using only Eq. (2) generally leads to estimates of α_j that vary widely (often switching sign) from one age group to the next. Eq. (3) assumes that age specific effects are fit by a third-order polynomial. This forces some degree of smooth transition between the coefficients of adjacent age groups, while also working to reduce the number of parameters to estimate. Substituting the assumptions in Eqs. (2) and (3) into Eq. (1) and rearranging implies the following three demographic terms.

$$\begin{aligned} D_{1,it} &= \sum_{j=1}^J j \Delta p_{j,it}, \\ D_{2,it} &= \sum_{j=1}^J j^2 \Delta p_{j,it}, \\ D_{3,it} &= \sum_{j=1}^J j^3 \Delta p_{j,it}. \end{aligned} \quad (4)$$


When these are included in my regression, their estimated coefficients are parameter estimates of the θ_k coefficients from Eq. (3), with θ_0 determined implicitly due to Eq. (2). I provide a detailed explanation of how Eq. (4) are constructed in Appendix “Construction of demographic controls”. There is a trade-off between the imposition of smoothness made by Eq. (3), and the gains from using these (Fair and Domínguez, 1991) controls. The benefits is that the researcher need not choose relevant age cutoffs a priori, allowing that coefficients can move from positive to negative anywhere along the distribution. Further, jointly estimating the full distribution avoids potentially misleading correlations among age groups. The share of population in my sample aged 50–64 has a -0.56 correlation with the share of population aged 20–34. In a regression including the 50–64 age group without jointly controlling for the younger share may be picking up its own correlation with the outcome, or a relationship that works through its correlation with the younger group. The strong assumption of the polynomial shape can be relaxed by allowing more coefficients (eg a fourth or fifth order polynomial in Eq. (3)). I show in Appendix C “Varying the degree of the population polynomial” that there is little change in the shape of my estimated age coefficients when increasing up to a sixth order. The Fair and Domínguez (1991) allows for all groups to be included while dealing with some of the practical difficulties of doing so. I proceed by estimating the following equation:

$$g_{i,t} = \theta D_{1,it} + \beta r_{i,t} + \omega D_{i,t} \times r_{i,t} + \phi g_{i,t}^{pop} + \rho X_{i,t} + \mu_i + \mu_t + \nu + \epsilon_{i,t} \quad (5)$$

where $g_{i,t}$ is an outcome of interest, which for my analysis will be growth rates of macroeconomic aggregates, in particular growth of: investment, the investment-to-GDP ratio, consumption, and real output.

⁶ Age shares sum to one, my changes in age shares sum to zero.

(a) Trends in Young, Old, and Working Populations



(b) Changes in Age Structure Relative to 1938




Fig. 1. Variation in population data: United States.

Any vector of controls for the population age structure are included with $D_{i,t}$. My preferred estimates will include the three variables in Eq. (4), but I will also include estimates using the old-age dependency ratio as a single control. To construct these I split the population into $J = 14$ age groups consisting of: the share less than age 15, greater than age 75, and shares of the twelve five-year age groups between them. In addition to my (Fair and Domínguez, 1991) style controls, I estimate specifications without an age structure term (population growth alone) and with the old-age dependency ratio. These provide more intuitive results, though will be unable to detect variation occurring within working life, or across young and old dependents. All specifications control for population growth, $g_{i,t}^{pop}$, to ensure that other demographic terms pick up variation due to population distribution rather than broad population growth rates. In addition, I control for the short term bill rate, $r_{i,t}$, defined from the Jordà et al. (2017) data as a weighted sum of bill rates and longer term government bonds. I include interactions with age structure variables (either old-age dependency ratios or D1–D3 defined above) with these bill rates to determine the sensitivity of age structure estimates to changing interest rate conditions. I can include any number of macroeconomic controls available in the JST data in the vector $X_{i,t}$. All estimations contain country (μ_i) and time (μ_t) fixed effects to control for country specific differences in outcomes as well as long run trends unrelated to aging. I also include controls for both world wars and financial crises as defined in the JST data.

Estimations reported below do not use additional macroeconomic control variables available in the data other than the contemporaneous movement of other growth rates of interest. This is in part to ensure that as much of the long time-series variation as possible is leveraged. The estimates presented appear to be quite robust to inclusion of other variables. Including these in the vector $X_{i,t}$ to account for short run dynamics could improve the fit of the empirical model. The JST dataset has information on price levels, unemployment rates, and government debt. Including these in my estimations does not appear to have large impact on point estimates, but reduces the sample substantially from 1840 observations to 1574, as many of these variables are incomplete. I include these estimates for robustness in Appendix B “Robustness to macroeconomic controls” to show that my outcomes of interest are largely unchanged.

I note that while many of my estimates are fairly robust to specification, they are largely descriptive in nature. It is not possible in this context to causally identify movements in underlying age structure that are plausibly exogenous from the macroeconomic outcomes that I study. Nor is it possible to properly deal with potential endogeneity of

interest rates to the movements in demographics that I wish to capture. My estimations seek to provide guidance as to what the descriptive evidence says in regard to the secular stagnation literature. This is a problem that exists for many of the existing empirical estimates of population age structure.

Data: Population trends

It is useful before moving to analysis to visualize some of these population changes and understand the variation in demographics that will drive any results. Fig. 1 shows two visualizations of population data for the United States. The first, Fig. 1(a), shows the time series variation in young dependents (aged under 15), economically active individuals (aged 20–64), and old age retirees (greater than 75) beginning in 1900 and extending to 2100 using UN population projections. The two broad trends in long run demographics are evident with young dependents falling dramatically, while the oldest age groups go from a trivially small part of the population, to one that is projected to be just shy of one-fifth of the total. Also evident is the large baby boom cohort, who shrink the 20–64 group while boosting young dependents during the period from roughly 1946–1964, their presence in the workforce has supported working age populations, but they will soon drive fairly dramatic falls in this group as they retire. This trend will continue at a slower pace as retirees make up ever larger shares of the population due to projected increases in longevity. Further, this baby boom represented only a brief reversal from the long run decline in fertility. These declines in fertility are projected to continue for advanced economies, though the pace of decline has slowed relative to what was seen early in this sample.

The right hand panel, Fig. 1(b), gives a picture of how the population in finer, five-year, age groups has evolved relative to the time when Hansen (1939) made his secular stagnation prediction just before World War II. Each bar shows the percentage point difference between a population age shares in a given year relative to their value in 1938. Young dependents were a much larger share of the population in 1900, with high fertility in prior decades generally supporting stronger shares of early career workers, but with the much higher mortality of the 19th century implying smaller shares for age groups over 30. The distribution in 1980 shows that the boomer cohort, who are in early career at this stage (aged 16–34), keep early career shares elevated relative to the late 1930s to nearly their 1900 levels. This is because fertility rates during the baby boom were similar those from the start of the 20th century. By 1980, fertility, and therefore young dependents,

Table 1
Population age structure and investment.

	Investment growth								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
g_{pop}	0.74 (1.35)	0.60 (1.37)	0.40 (1.40)	0.37 (1.46)	−1.38 (1.34)	0.76 (1.24)	0.64 (1.28)	0.58 (1.29)	−1.55 (1.17)
r (demeaned)	0.47* (0.26)		0.40* (0.23)	0.22 (0.25)	0.24 (0.23)		0.25 (0.23)	0.38 (0.48)	0.48 (0.46)
OADR		−1.10** (0.40)	−1.02*** (0.35)	−0.93** (0.36)	−0.89* (0.49)				
OADR $\times r$				0.21*** (0.06)	0.22*** (0.07)				
D1						8.30*** (2.22)	7.87*** (2.25)	8.55*** (2.22)	9.01*** (2.56)
D2						−10.59** (3.91)	−9.89** (4.15)	−11.61*** (3.84)	−14.49*** (4.14)
D3						3.47† (2.13)	3.18 (2.25)	4.24* (2.15)	6.02** (2.13)
D1 $\times r$								−0.31 (0.56)	−0.77 (0.58)
D2 $\times r$								−0.34 (1.00)	0.49 (1.09)
D3 $\times r$								0.46 (0.50)	0.06 (0.56)
Year	✓	✓	✓	✓	✓	✓	✓	✓	✓
Country	✓	✓	✓	✓	✓	✓	✓	✓	✓
γ_c & γ_y					✓				✓
F: D1–D3 (p -val)						6.49 0.00	5.93 0.01	5.68 0.01	5.14 0.01
F: D1–D3 $\times r$ (p -val)								3.77 0.03	4.88 0.01
R2	0.42	0.42	0.42	0.42	0.46	0.43	0.43	0.43	0.47
N	1840	1840	1840	1840	1840	1840	1840	1840	1840

All regressions include controls for WW1, WW2, and Financial Crises. $\gamma_{i/y/c}$ refer to inclusion of controls for the growth rate of investment/output/consumption respectively. Standard errors clustered by country, † $p < 0.15$ * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

has fallen back onto the longer run trend seen in Fig. 1(a), while the large improvements in health and mortality have started increasing the share of late career workers and retirees. By 2020 the population pyramid starts to look like the one described in Hansen (1939), with all age groups under 45 having a diminished share, while the boomers (aged 56–74) are a large cohort enjoying better longevity than those who came before and are poised to fully leave their economically productive stage of life.

Empirical results

I now present estimations of Eq. (5). For each outcome of interest I specify three equations. First, a naive regression using just the population growth rate. Next I include the growth of the old-age dependency ratio, defined as the ratio of individuals past retirement age to the working age population. This is a common measure of age structure used in other secular stagnation research. My preferred specification controls for growth rates of population shares and includes information on the full age distribution captured by my polynomial (Fair and Domínguez, 1991) method described above. I first investigate the impact that population age structure has on investment and investment-to-GDP ratios. This is a question at the heart of the original (Hansen, 1939) theory. I then turn to how consumption and real output growth respond to these demographic forces.

Age structure and investment

In Table 1 and Table 2, I report estimates for Eq. (5) for two outcomes: growth rate of investment and growth of the investment-to-gdp ratio. For each, I include nine specifications of Eq. (5) with two different sets of demographic controls. My preferred specification

includes the full polynomial controls for population age, as well as their interactions with short term interest rates. While the individual significance of these terms is reported, the more relevant statistic is the *joint* significance of these variables. For this reason I include an F-test for joint significance of D1–D3 as well as for the three interaction terms. For convenience, I also include the associated p -value for this test. All estimates include country and year fixed effects as well as controls for both world wars and financial crises. In addition to the full specification with short term interest rates (demeaned) and their interactions with my population age-structure controls, I also include specifications that remove these interactions and interest rate controls. This is to show that while there are many potential exogeneity concerns with how population age structure might influence the interest rate, inclusion of the interaction has little effect on the coefficients on the demographics directly. Finally I include an additional specification that includes the contemporaneous growth rates of consumption and output, the outcomes I investigate in the next section.

In Table 1 I report these results for investment growth. Notably population growth rate is not significant in any specification, and is sensitive in terms of sign to the inclusion of contemporaneous controls for consumption and output growth. There are of course many other potential competing mechanisms at play using broad population growth alone, which is why I am more interested in specifications that also account for age structure and include it to avoid any confounding effects. The safe rates are somewhat puzzlingly positive, but not statistically distinguishable from zero when interactions with age variables are included. I note that the value of the short rate coefficient in column one falls to nearly zero in my robustness specifications shown in Table 5 Appendix B “Robustness to macroeconomic controls”. Those estimates include controls for short term business cycle variables and should

Table 2
Population age structure and investment-to-GDP.

	Investment-to-GDP growth								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
g_{pop}	-0.25 (0.64)	-0.63 (0.62)	-0.45 (0.63)	-0.47 (0.63)	-0.94 (0.68)	-0.72 (0.61)	-0.51 (0.58)	-0.52 (0.58)	-1.10† (0.65)
r (demeaned)	-0.31** (0.12)		-0.35*** (0.11)	-0.48*** (0.13)	-0.47*** (0.12)		-0.45*** (0.10)	-0.55*** (0.19)	-0.52*** (0.18)
OADR		-0.55** (0.25)	-0.62** (0.26)	-0.55** (0.24)	-0.52* (0.27)				
OADR $\times r$				0.16** (0.06)	0.16** (0.06)				
D1						4.21*** (1.10)	4.97*** (1.07)	4.99*** (0.82)	5.02*** (0.86)
D2						-6.25*** (1.89)	-7.50*** (1.95)	-7.94*** (1.50)	-8.53*** (1.49)
D3						2.34** (0.85)	2.85*** (0.93)	3.22*** (0.75)	3.62*** (0.73)
D1 $\times r$								-0.07 (0.34)	-0.16 (0.34)
D2 $\times r$								-0.13 (0.64)	0.04 (0.63)
D3 $\times r$								0.17 (0.31)	0.09 (0.30)
Year	✓	✓	✓	✓	✓	✓	✓	✓	✓
Country	✓	✓	✓	✓	✓	✓	✓	✓	✓
$\gamma_c \& \gamma_y$					✓				✓
F: D1-D3						9.04	13.01	20.25	17.49
(p -val)						0.00	0.00	0.00	0.00
F: D1-D3 $\times r$								2.65	2.29
(p -val)								0.08	0.11
R2	0.31	0.31	0.31	0.31	0.32	0.31	0.31	0.31	0.32
N	1840	1840	1840	1840	1840	1840	1840	1840	1840

All regressions include controls for WW1, WW2, and Financial Crises. $\gamma_{i/y/c}$ refer to inclusion of controls for the growth rate of investment/output/consumption respectively. Standard errors clustered by country, † $p < 0.15$ * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

improve estimates of this term which is likely biased due to reverse causality.⁷

All of the population age structure variables show a strong correlation with investment growth. A negative old-age dependency ratio suggests that increasing shares of retirees relative to working age population will act as a drag on investment growth. This is quite close to -1 , which means a one percentage point increase in the growth rate of old-age dependency ratios shrinks investment growth by close to one percentage point. Such a change would be less than a standard deviation shift in this growth rate and would reflect a quantitatively large impact on investment growth, which has a mean in the estimation sample of 8.59%. The old-age dependency ratio actually shrunk from 1996 to 2004 in the United States, with the boomers still in the workforce and their children (the relatively large *echo boom*) also beginning their working lives. In recent years this ratio has grown substantially, and is projected to continue doing so, reversing this long standing demographic dividend. Interaction of old-age dependency with demeaned short rates in both specifications has the opposite sign. This implies that when rates are above their mean, the impact of the relative share of retirees is diminished, while it is strengthened when rates are relatively low. The standard deviation of the short rate is 3.56, so when rates are one standard deviation above their means the implied old-age dependency impact on these outcomes fall in absolute value to close to zero (-0.11), with symmetric increases in magnitude when rates are low. In a theme that will be common for many of my estimates, low interest rate regimes appear to strengthen the implied relationship between age structure and investment growth.

⁷ ie: Interest rates are high because they endogenously rise with the business cycle as capital becomes scarce.

Turning to my preferred specification, I first note that these three demographic variables constructed from changes in five-year population shares are hard to interpret directly, a point I will address shortly. Their coefficients are the estimates of the $\hat{\theta}_{1-3}$ terms in Eq. (3) and therefore represent the curvature of the polynomial relationship that defines the coefficients on each individual age-share. Many of the individual estimates of the age structure on outcomes directly are significant. However, more relevant is the joint significance of these estimates, for which I provide F-tests and their corresponding p -values at the bottom of Table 1. They are quite strongly significant for the direct correlations, and their interactions with short term interest rates are jointly significant.

I report the same estimates using the growth rate of investment-to-GDP as an outcome in Table 2. The broad results are quite similar, with a few notable differences. Here, the coefficient on the short rate has the expected negative sign, and in the full specification in column 9, the growth rate of population has a weak significant impact. While demographics have the same sign and are strongly significant, their interaction with interest rates in my polynomial specification are only weakly jointly significant at the 10% and 15% level, with no single term independently significant. The broad implication of these results is that not only might investment growth be affected by population age structure, but also capital deepening.⁸ Since these polynomial controls are difficult to interpret directly, I back out the age specific coefficients, $\hat{\alpha}_j$, by using the parameters $\hat{\theta}_{1-3}$ in Eq. (3), calculating 95% confidence intervals using the delta method. These are reported in Fig. 2.

⁸ Capital deepening can be defined differently, and I might prefer investment-to-labor but given that the closest proxy I have in this long run data is rather working age population I prefer to use the investment-output ratio.




Fig. 2. Implied coefficients on five-year population share growth rates: Investment growth.

Implied five-year population share growth rate coefficients from Column 9 in Table 1 and Table 2 are calculated from coefficients on (D1–D3), as $\hat{\theta}$ in Eq. (3), with 95% standard errors calculated via the delta method.




Fig. 3. Conditional interest rate impacts on age structure coefficients.

Average coefficients on age share growth rates calculated as in Fig. 2, which represent estimates when interest rates are at long run mean. Dashed maroon line and dot-dashed Fuschia line represent implied coefficients under high and low rate environments respectively. These are defined as the short term interest rate as one standard deviation (3.56 ppts) above/below its in-sample mean (4.98).

Each point along the lines in Fig. 2 is the estimated impact of increasing the growth rate of the corresponding population age share by one percentage point according to the D1–D3 variables in column 9 in Table 1 and Table 2. The impact of young workers is strongly positive, shifting into negative territory for population age groups over 55 for investment growth, and by age 50 for investment-to-GDP. One of the key benefits of using these polynomial controls is that the researcher need not specify ex-ante the important age groups or cutoffs. While old-age dependency ratio looks like it might do a reasonably good job, it likely misses some late-career dynamics for this 50–64 age group. Such a distinction is non-trivial when assessing the past two decades (as well as the next) as the large boomer cohort has moved through this group. It is important to keep in mind is that since these are changes in shares, which must sum to zero, any increase in one bin must be met with a subsequent decrease elsewhere. In some cases, such as increasing population in the > 60 age bins while decreasing population in the < 20 groups this could compound the estimated impact. These are *direct* correlations between age structure and these growth rates from Table 1 and Table 2, which can be interpreted as age estimates when the interest rate is in a somewhat neutral position (its long run

in-sample mean). The results for Fig. 2(b) look quite similar to those reported in Higgins (1998), who performs a similar estimation on a larger cross section of countries from 1950–1989

To unpack the meaning of the interaction terms between population age structure and short rates I plot the same $\hat{\alpha}_j$ coefficient estimates in Fig. 3, along with the equivalent estimates when short term rates are high, and low. These simply use both the $\hat{\theta}_{1-3}$ estimates, along with the $\hat{\omega}_{1-3} \times r^{h/l}$ terms, to back out new $\hat{\alpha}_j^{h/l}$ terms. Since these new terms depend on a particular level of the short term rate I take high and low rates to be one standard deviation changes above and below their long run mean. The mean short rate in my estimation sample is 4.98, with a standard deviation of 3.56. Thus the high rate environment for these and future graphs is when the short rate is 8.54%, while the low rate environment is 1.42%. Broadly speaking these interactions imply that low rate environments exaggerate the existing life-cycle estimates putting stronger negative weight on retirees and higher positive weight on the rest of the population. Generally the high rate environment weakens these. An important note is that interaction coefficients need not sum to zero, which is true by assumption in the average estimates. Given that any net impact of changes in the growth rates of population

Table 3
Population age structure and consumption.

	Consumption growth								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
g_{pop}	1.57** (0.74)	1.51* (0.79)	1.52* (0.79)	1.52* (0.79)	0.58 (0.55)	1.73* (0.85)	1.74* (0.84)	1.74* (0.85)	0.62 (0.65)
r (demeaned)	−0.01 (0.05)		−0.02 (0.05)	−0.03 (0.04)	−0.03 (0.03)		−0.03 (0.04)	−0.06 (0.06)	−0.01 (0.06)
OADR		−0.16 (0.18)	−0.17 (0.18)	−0.16 (0.18)	−0.15† (0.09)				
OADR $\times r$				0.01 (0.02)	0.01 (0.02)				
D1						0.61† (0.40)	0.66† (0.42)	0.67 (0.48)	0.94*** (0.28)
D2						0.44 (0.76)	0.37 (0.79)	0.31 (0.88)	−1.28** (0.52)
D3						−0.57 (0.43)	−0.54 (0.44)	−0.51 (0.47)	0.47* (0.26)
D1 $\times r$								0.07 (0.12)	−0.20** (0.08)
D2 $\times r$								−0.11 (0.19)	0.38** (0.13)
D3 $\times r$								0.05 (0.10)	−0.18** (0.07)
Year	✓	✓	✓	✓	✓	✓	✓	✓	✓
Country	✓	✓	✓	✓	✓	✓	✓	✓	✓
γ_y & γ_i					✓				✓
F: D1-D3						14.06	14.27	12.91	4.31
(p -val)						0.00	0.00	0.00	0.02
F: D1-D3 $\times r$								0.22	2.80
(p -val)								0.88	0.07
R2	0.28	0.28	0.28	0.28	0.52	0.29	0.29	0.29	0.52
N	1840	1840	1840	1840	1840	1840	1840	1840	1840

All regressions include controls for WW1, WW2, and Financial Crises. $\gamma_{i/y/c}$ refer to inclusion of controls for the growth rate of investment/output/consumption respectively. Standard errors clustered by country, † $p < 0.15$ * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

age share must be a weighted average of the changes in individual shares and the coefficients in Fig. 3, what matters is not necessarily the sign of an age's specific coefficient, but rather the relative magnitude, which will determine that age group weight. In both cases the low rate line has more extreme swings, implying that shifting population weight from mid-career groups to late career and retirees will have a *stronger* impact than the average rate, with a muted effect in a high rate environment, where the relative change from one age group to the next is smaller. These interest rate interactions are consistent with those found for old-age dependency ratios above. Notably the statistical significance of the interactions for growth rate of investment-to-GDP are somewhat marginal and their quantitative magnitude looks small here compared those on investment growth. In the latter, the high rate curve looks nearly flat, suggesting little change in investment growth rate as population weights are moved around this age distribution. Meanwhile, the low rate estimates in Fig. 3(a) have roughly double the range of the average coefficient estimates implying large differences associated with shifting population weights. In Fig. 3(b) the difference between the average lines and these high/low rate lines is much less pronounced. My estimates thus imply that the impact of age structure on investment growth appears to be highly sensitive to the level of the interest rate, while estimates of age relationships to growth of investment-to-GDP will be less affected.

It should be kept in mind that these interaction coefficients are only significant jointly, and while they are strongly jointly significant for growth in investment in Fig. 3(a), they are only weakly so for growth in investment-to-GDP in Fig. 3(b). I include them here primarily for consistency with significant interactions on output below, but also to show a potential relationship between these variables that has been documented extensively in the literature.

Age structure, consumption, and output

I now repeat the above exercises for consumption and output. While much of the original secular stagnation hypothesis suggested investment and capital deepening as a *mechanism* the ultimate concern is the role that population age variables will have on output growth. Further, some have suggested that sluggish output growth may not be a concern if consumption growth remains robust, a key thesis behind theories of slow growth reflecting a maturing of economies, where output growth itself is not necessarily fully reflective of overall welfare.⁹ Growth in output may not be important for overall wellbeing if it is not met with an equivalent fall in consumption growth. I report the estimated coefficients for the same population age estimations first on consumption growth in Table 3 and on output growth in Table 4. The full specifications (column 9 in each table) include the growth rates of investment and output in Table 3 and the growth rates of investment and consumption in Table 4.

Here population growth is now a significant positive variable. There is a large literature estimating the effect of population growth on output, with mixed signs. Notably when controlling for the growth of economically active population in a panel of rapidly developing Asian economies, Bloom and Williamson (1998) find a negative impact of population growth (and a positive effect for the economically active population). Such an estimate would be more in line with the long run macroeconomic development literature, which uses Galor and Weil (2000) style models of fertility transition suggesting a movement from high-fertility low-growth economies to low-fertility high-growth. As I

⁹ For a full treatment see: Vollrath (2020).

Table 4
Population age structure and output.

	Output growth								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
g_{pop}	1.38** (0.61)	1.37** (0.64)	1.38** (0.63)	1.38** (0.63)	0.66** (0.31)	1.67** (0.66)	1.66** (0.66)	1.65** (0.64)	0.83** (0.32)
r (demeaned)	-0.02 (0.07)		-0.02 (0.06)	-0.01 (0.06)	-0.01 (0.05)		0.01 (0.05)	-0.08 (0.07)	-0.07 (0.06)
OADR		-0.00 (0.20)	-0.01 (0.19)	-0.01 (0.19)	0.10 (0.13)				
OADR $\times r$				-0.00 (0.03)	-0.02 (0.02)				
D1						-0.75 (0.74)	-0.78 (0.75)	-0.47 (0.69)	-1.09* (0.53)
D2						2.99** (1.33)	3.02** (1.36)	2.44* (1.34)	2.73** (1.04)
D3						-1.71** (0.62)	-1.72** (0.64)	-1.48** (0.66)	-1.40** (0.51)
D1 $\times r$								0.39*** (0.08)	0.37*** (0.06)
D2 $\times r$								-0.71*** (0.17)	-0.64*** (0.14)
D3 $\times r$								0.34*** (0.10)	0.30*** (0.08)
Year	✓	✓	✓	✓	✓	✓	✓	✓	✓
Country	✓	✓	✓	✓	✓	✓	✓	✓	✓
γ_c & γ_i					✓				✓
F: D1-D3						9.96	10.05	9.15	5.63
(p -val)						0.00	0.00	0.00	0.01
F: D1-D3 $\times r$								10.20	14.82
(p -val)								0.00	0.00
R2	0.32	0.32	0.32	0.32	0.56	0.34	0.34	0.35	0.58
N	1840	1840	1840	1840	1840	1840	1840	1840	1840

All regressions include controls for WW1, WW2, and Financial Crises. $\gamma_{i/y/c}$ refer to inclusion of controls for the growth rate of investment/output/consumption respectively. Standard errors clustered by country, $^{\dagger}p < 0.15$ $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.




Fig. 4. Implied coefficients on five-year population share growth rates: consumption and output growth.

Implied five-year population share growth rate coefficients from Columns 9 in Tables 3 and 4 are calculated from coefficients on (D1–D3), as $\hat{\theta}$ in Eq. (3), with 95% standard errors calculated via the delta method.

will show in my age-coefficient graphs below, my preferred specifications in columns 9 suggest very strong *negative* impact of young age groups. In this specification this may absorb some of the negative correlations of population growth found in the literature, and is directly in line with the quantity-quality style trade-off, which theoretically motivates such a channel for negative associations between population and growth in that literature. Of course existing estimates may have

population growth measures picking up other correlations with age groups that are picked up in my D1–D3 controls, leading to differences in estimates on the broad population growth term. Another explanation is that my sample, though stretching back further in history than many, still only observes developed economies *after* some of the more dramatic falls in fertility and population growth have taken place. Finally, unlike the developing context, some rich economies (especially the




Fig. 5. Conditional interest rate impacts on age structure coefficients.

Average coefficients on age share growth rates calculated as in Fig. 4, which represent estimates when interest rates are at long run mean. Dashed maroon line and dot-dashed fuchsia line represent implied coefficients under high and low rate environments respectively. These are defined as the short term interest rate as one standard deviation (3.56 ppts) above/below its in-sample mean (4.98).

United States early in my sample) have a large migration component of population growth which may have a different estimated coefficient than that of internal growth.¹⁰

My primary interest is in the age structure themselves, so while the strength of these population growth estimates are interesting in their own right, I include them primarily to avoid picking up age-structure impacts that come from long run changes in population growth rather than impacts of how that population is distributed across ages. Unlike Tables 1 and 2, consumption and output show no strong relationship with old-age dependency ratios. The polynomial controls on the other hand are strongly jointly significant for both outcomes directly, with consumption growth quite strongly jointly significant, despite none of the individual coefficients passing the 10% statistical significance threshold outside of the regression that includes investment and output growth. Consumption interactions with short rates are only weakly significant and only when including these contemporaneous controls for investment and output growth. With output while there are varying degrees of statistical significance for any given demographic term, these polynomial controls are broadly significantly estimated both in their direct impact and their interaction with interest rates. It is worth noting that consumption is the only outcome where the signs of these estimations are not robust across all specifications. The age specific coefficients under average interest rates implied by D1–D3 are shown in Fig. 4.

Output growth displays a strong negative relationships with young dependents and early career workers, with consumption sharing this for young dependents but turning positive for early career workers. Interestingly though demographics are strongly jointly estimated no single age coefficient is significant when backing out their point estimates via the delta method using the polynomial variables in Table 3. This along with the changing signs of coefficient estimates on consumption outcomes should perhaps caution the stability of these point estimates, whereas growth in investment and output are robust to all specifications estimated. Output has much more precisely estimated positive

age coefficients from age 40 up until age 70. Both have a negative but statistically weak coefficient for the oldest age groups. The implication for output is broadly that growth of middle-to late career workers has positive growth effects, while growth of young and old dependents, and early career workers negatively correlate with growth.

Both consumption growth and output growth have strong interactions with interest rates, with output's coefficients robust to the two specifications reported. Fig. 5 plots the same exercise as Fig. 3 to include the age coefficients for high and low interest rate environments. Both interactions, like those in Fig. 3 seem to imply that low rate environments accentuate the curvature of these coefficient estimates. Notably for consumption old-age individuals who have near zero average demographic effect have a relative effect that changes from negative to positive among older age groups. The output age-specific coefficients in Fig. 5(b) suggest, similar to those from investment growth, that low rate regimes reinforce the average estimates for age structure, with a larger change between the negative young and old coefficients and those of mid-late career workers. This appears to be consistent with empirical evidence found in Eggertsson et al. (2019a), who suggest that when rates are near zero that general equilibrium mechanisms, which would normally offset at least some of these demographic impacts on output, are shut down. Their results show that demographic forces exert strong negative influence on output growth during these low interest (ZLB) periods. Eggertsson et al. (2019a) propose a mechanism where capital deepening occurs when interest rate adjustments, which come due to aging, encourage investment. Something to this effect can be seen in Fig. 5(b), where the high rate line is much flatter from late career workers to old retirees, suggesting that output growth should not be greatly altered by moving population weight from these older workers into retirement. Contrast this to the low rate line where there is a pronounced drop in the coefficient estimates from their peak at age 55–59 and low rates imply larger reduction in output growth.

Notably while I show the strong correlation between interest rates and the transmission of population age to growth, I do not try to model the age-structure interest rate relationship directly. However, the implications are similar to that of Eggertsson et al. (2019a) and provide further evidence that the mechanisms linking aging, interest rates, and macroeconomic variables are worthy of investigation. Eggertsson et al. (2019a) use a kink in the relationship around the ZLB to show this connection, while mine imply a smooth linear relation. It seems likely

¹⁰ Indeed such migration likely suffers from reverse causality with growth and as such, estimates of population growth, which are not my primary object of interest, should be taken with a grain of salt here. It should be possible to get some measure from the (Human Mortality Database, 2020) birth and mortality data of internal and external contributions to population growth over my sample, but is not at present included in this analysis.

that there are some non-ZLB forces at play given the rarity of ZLB episodes in my sample¹¹ and significance of these results.

The age specific coefficients presented in this section give some useful intuition on how population age structure correlates with macroeconomic variables of interest. However, while it is easy to visualize the average impact of changing just one group, with fourteen coefficients interacting with fourteen population changes it is incredibly difficult to visualize from the information in Figs. 3 and 5 what the net impact of a changing population. In Section “Implications for aging and secular stagnation: A quantitative exercise”, I perform a quantitative exercise that will do exactly that, and help unpack what these coefficients imply in terms of this net impact of population aging.

Robustness: Results on various subsets of the data

In Appendix D “Robustness: Changing the estimation sample”, I briefly explore the implications for changing the estimation sample used in my main results by estimating the full specifications of my demographic models above on three variations of the data. The first two exercises test whether the estimated impact of aging is different before and after the baby boomer cohort by splitting the sample in to the period before 1945 and after 1950. The period from 1870 to 1945 is marked by long run declines in fertility with rapid health improvements that increased survival across all ages. As such many countries saw relatively flat working age shares, though the contribution of later career workers in those shares was growing. The post-war baby boom created a cohort that triggered a massive fall in the working age share upon their birth, with a subsequent rise as they entered the labor force in the 1970s and 1980s. As a result, the period from 1950–2018 is an incredibly volatile time from the perspective of historical population age distributions. While much of the existing research on the impacts of population age structure on macroeconomic outcomes has been conducted in this period, the future is likely to consist of more stable movements as this boomer cohort enters retirement age, with long run trends continuing to increase the relative size of old-age retirees relative to workers. The third specification I test is a replication of the results in Sections “Age structure and investment” and “Age structure, consumption, and output” using ten year averages. This is to show that the shorter run estimates that I find above are hold when looking at lower frequency movements in macroeconomic variables.

Broadly speaking the direct effects of demographics appear to be robust in terms of signs and significance, but may display important heterogeneity over time. Comparing the age specific coefficients reported in Fig. 10 for each of these three models to their counterparts in Figs. 2 and 4 paints a broadly similar portrait with much more volatile estimates for the pre-1945 period relative to those estimated on the post-1950 sample. This may come about for a number of reasons. First, there may be some true structural change across these periods either due to the nature of the demographic movements themselves.¹² If this is the case then it would mean that the smaller, though significant and qualitatively similar, estimates may be more relevant for the current situation. It may also be the case that more stable trends with large potentially exogenous shocks such as the Spanish flu and World War I may better identify these shocks relative to the movements in the post-war period that are ultimately going to be correlated with many other trends taking place at the same time: for example, post-war reconstruction and globalization. In such a case one might consider the earlier

estimates as better indicators of the true demographic relationship. Finally there are differences in data coverage between the two samples. While the pre-boom sample nominally covers more years (75) than the post-1950 (68), limitations in demographic and macroeconomic data make the pre-1945 sample much smaller, with only 14 out of the 18 countries in the sample having any observations in this window and the average panel length only 45 years relative to the nearly 64 in the post-1950 period. On the one hand this exercises the value of considering longer run data as there is clearly insight to be gained from the longer historical sample. On the other it shows the need for more work to study the potential asymmetries in the impact of population age structure across both time and space. This is a valuable area of future research.

Implications for aging and secular stagnation: A quantitative exercise

The above analysis suggests that age structure is important for these macroeconomic variables in the long run for this sample of developed countries. However, it is not easy to understand from the estimated coefficients in Fig. 3 or Fig. 5 what the net impact will be for various changes in population age distributions. To make the implications of these estimates clearer, I now conduct a quantitative exercise to show the impact of observed changes in population age shares (through D1–D3) on these macroeconomic outcomes. To do this, I set any controls equal to their sample mean (and set year fixed effects and other dummy variables to zero), and plot movements in the predicted value of each outcome using the estimated models in Column 9 of Table 1, 2, 3, and 4 for the United States as its age structure, D1–D3, varies over time. These are reported as changes relative to the mean prediction of each variable so that a positive and negative values represent demographic head and tailwinds respectively. I note that I also hold population growth fixed. I do this because it is not significantly estimated across all specifications, and because I am more interested in the relationship of age structure independent of any population growth effects. Including population growth time series would have little impact on predictions for investment. For consumption and output growth it would introduce, or in places add to, a downward trend. In general the purpose of this exercise is to be able to understand how the changes in age structure estimates from the estimates in Section “Empirical results” impact outcomes, so including it would ultimately muddy that analysis. This exercise is not intended to forecast these variables, but rather to provide context as to the quantitative magnitude implied by these estimated coefficients.

The results of this exercise are reported in Fig. 6. The solid lines in each of these figures are the contribution of demographics to the model predicted growth rate of each outcome when interest rates are at their long run mean. This experiment amounts to taking the observed (and projected) growth rates of each five year age group, multiplying them with their coefficients in Fig. 2 or Fig. 4, and summing each age group’s contribution to get a net impact and adding that to the United States long run average. Movement of the solid lines in this figure are thus the variation in an outcome expected purely due to changes in the underlying age structure of the United States in that year. I then repeat the exercise, setting interest rates at their high and low values as defined above (+/– one standard deviation), which is the same as repeating the exercise of weighting contributions from each age group, but rather using the high/low rate age coefficients from Figs. 3 and 5. I emphasize that the interest rate interactions are very strongly jointly significant for output growth and investment growth, but the three interactions have *p*-values of their F-statistic of joint significance of 0.07 and 0.11 for consumption and investment-to-GDP ratio, respectively. While I include them all for completeness the latter two should be taken with some additional caution. In addition to this, it is worth

¹¹ There are 36 observations of this short rate at or below zero in my 1840 observation estimation sample, all of these occur between 2012 and 2018. Allowing for rates up to 0.5 expands this slightly to 119 cases, but again the vast majority are from the last decade.

¹² For example, the effects of short term reversals of longer run trends, as in the baby boom may not have the same impact as deeper structural changes in fertility and longevity.




Fig. 6. Predicted impact of age structure: United States.

Linear prediction for the United States with all variables set to long run mean and all year fixed effects set to zero, predicted values of growth rates plotted relative to their in-sample mean. Changes in growth rates in each variable reflect movement in model predicted value driven by changes in D1–D3.

remembering that while both investment growth rates and output are strongly robust to specification, consumption varies quite a bit.¹³

¹³ Repeating this exercise using other specifications of consumption paint broadly similar qualitative pictures, with larger reductions in projected future stagnation.

There are a number of important trends to notice in these estimates. To help in understanding it is worth noting the position of the outsized baby boom cohort in the United States, who are born between 1946 and 1964. Ignoring the large movements in age shares during World War 2, the pictures of investment, consumption, and output look consistent with the future that Hansen (1939) cautioned in his 1938 address.

These are in part driven by the relatively small *silent* generation entering the workforce, in line with the hypothesis from Hansen (1939). However the large drops, particularly in consumption and output, are exaggerated by the large share of young dependent boomers, who have strong negative correlations with consumption and output growth. I note however that dependents have a zero (small negative) point estimate for investment-to-GDP growth, which also falls post-war fairly dramatically.

For all outcomes, these estimates suggest that the entry of this boomer cohort into the workforce in the late 1960s and early 1970s dramatically reversed these demographic headwinds, creating sustained high levels of predicted investment growth, capital deepening, and consumption growth well above their pre-war level. Growth in output sees positive trends coming from demographics only later in the late 1980s. This picture suggests a demographic dividend similar to that described in Mason et al. (2016), though it appears to take some time for the investment boom to translate into output growth. From here there are three important trends. First, the boomer cohort gradually moves through the workforce, and by the 1990s the oldest groups start to provide headwinds to investment and consumption growth. Second, increasing longevity continually drives growth in the oldest population age groups, which provide additional headwinds for all of these outcomes. Third, falling fertility after the boomer cohort implies smaller shares of young dependents throughout the 1970–80s, contributing to positive output estimates (small impact on investment). This latter point is consistent with the kind of mechanisms in the long run development literature using a quantity-quality model from Becker (1960). These smaller cohorts eventually reverse the demographic dividend as they provide little support to the working age population, though this has limited impact on output until the outsized boomer cohort leaves the workforce. Broadly speaking the model characterizes the past fifty years as seeing a boom in investment and consumption, followed by a boom in output.

With projected increases in longevity slowing from the rapid improvements over the 20th century, a return to lower fertility rates, and the boomer cohort fully entering retirement: both investment growth rates and consumption growth have already passed their modern low point and are set to return to a level that is more or less demographically neutral. For both investment variables the point estimates suggest that this should largely have already taken place by 2020. Notably this neutral level is far below peaks from a demographic dividend in the late 1970s, so while population age structures are not predicted to provide any further negative impact on growth rate of these variables in this estimated model, research that looks only to the recent past for a reference may be misleading. Consumption follows a similar path but is estimated to have persistent, but small, negative demographic forces out to 2050 of roughly a tenth of a percentage point. Output is estimated to have only recently had positive impact from age structure due to the positive coefficients on the large late-career boomers who remain in the workforce. At present their departure pushes the age effect on output growth to negative, hitting a low point at roughly three-quarters of a percentage point in 2030.

The recent literature on secular stagnation focuses largely on the demographic trends occurring from 1970 to present, with dramatic acceleration of stagnation from about 1990 onward. I note that while my age structure estimates are generally consistent with these trends, the interest rate interactions during this period look particularly important for growth rates in investment and output. Demographics reduce estimates of growth rates for these two outcomes much more aggressively when rates are low, with high interest rates mitigating or even reversing any stagnation. Indeed if demographics themselves are pushing interest rates down, through a channel of demand for savings by older age groups,¹⁴ then these partial equilibrium results may undersell

the overall impact of aging. Demographics may endogenously move the economy from the high rate line in the 1970–80s, slightly reducing the demographic dividend, to the low rate lines in the 2010s and onward, exaggerating the demographic deficit. Until very recently economies have been stuck in a low interest rate environment, and many predict that such pressures will remain for some time.

Comparing to recent literature, Aksoy et al. (2019) show using an unconditional VAR estimate from 1970–2014 that age structure variables (similarly specified to those used here), account for a 2 percentage point decline in investment-to-GDP in OECD countries over the period from 1970 to 2014. While the roughly 5 percentage point fall that I find over this period for the United States is larger, this is driven mostly by the finding that population aging was exhibiting 3 percentage points of positive pressure in 1970. They find a 1.25 percentage point fall in output over the same period, which is quite close to the fall I find from 1990, though again it is hard to directly compare as I find demographic impact on output rising and then falling from that period. While it is not possible to have a perfect apples-to-apples comparison of our results I see mine as largely in support of their findings.

What are the predictions for secular stagnation in the United States? It would appear that investment growth has already seen its stagnation play out, both in terms of growth rate and declines in capital deepening. Both of these saw an estimated demographic contribution of roughly –2 percentage points around 2010, with very little predicted movement away from a zero net demographic effect going forward. Consumption growth, though the least robust of these estimates, is predicted to have faced demographic resistance and will continue to have minor reductions in growth due to demographics in the far future. Output is predicted to have prolonged period of below trend growth, all things equal. This is more than a percentage point swing from the estimated demographic dividend still enjoyed when the boomer cohorts were in their late career in the late 1990s and early 2000s. While these are large estimated impacts of demographics, they need not be overtly pessimistic. There are some cases to be made for mature economies growing at a slower pace, as outlined in Vollrath (2020). Of course these demographic forces need not be the only factor determining long run growth trends, and themselves may drive some innovation as economies adapt to living in a world with high old-age dependencies. The goal of this exercise is not to forecast future growth trends, but to demonstrate that the historic macroeconomic data appears consistent with a story of population aging and secular stagnation.

Secular stagnation: A global phenomenon?

While much of the discussion has been framed around the experience of the United States, population aging is a global phenomenon that is poised to impact all advanced economies. In fact, the United States is on the younger side of this group. To illustrate the implications of changing age structures more broadly I now repeat this quantitative exercise for four other economies in my sample: France, Italy, the United Kingdom, and Japan. To contextualize some of their broad demographic trends I plot both the working age population and share of population over age 75 for each of these countries in Fig. 7. Data from the HMD for France go back to the start of the JST macroeconomic data, and so it is possible to extend this series to 1870. While France was the first advanced economy to begin its demographic transition in the eighteenth century, it remains on the young end of the spectrum relative to its European peers. The United Kingdom faces very similar evolution of young dependents to the trend seen in France due to similar fertility rates, but has fewer in the oldest age group. On the other end of the spectrum are Japan, often viewed as the canary in the coal mine with respect to aging populations, and Italy whose old-age dependency ratio is the highest in Europe. These two countries have seen a complete collapse in fertility and explosion of old age shares, trends that are projected to continue.

¹⁴ As is common in the theoretical models such as Gagnon et al. (2021), Eggertsson et al. (2019b), and others.




Fig. 7. Changing demographics: France, Italy, UK, and Japan.

I repeat the quantitative exercise conducted for the United States with each of these countries, here considering only the mean demographic effect. These are reported in Fig. 8. The results for France and the UK in Fig. 8 paint a similar picture to those for the United States in Fig. 6. Both see an investment boom coinciding with the baby boomer cohort entering the labor force, with a subsequent boom in consumption and output. Future projections show a fairly immediate return to neutral demographics for both investment growth variables, with a prolonged journey back toward zero for both consumption and output growth which will have negative demographic headwinds for some time. The idea of secular stagnation as a prolonged deviation from trend growth is a much clearer implication of my results when taken through the lens of these and many similar European countries.

The picture for Italy and Japan looks considerably more bleak. Age structure estimates for Japan, whose macroeconomic data do not extend earlier than the 1950s, seem to be in line with a large post-war boom, with population forces driving decreases in investment, consumption, and output in the 1990s. While Italy's postwar path looks similar to that of France and the UK it faces much more rapid increases in the share of old-age retirees in the coming years. In this sense it looks much more like Japan than its European peers. It too has predicted continual stagnation if these model estimates are to be taken seriously. While future out-of-sample projections for investment growth is estimated to rebound substantially, the net demographic contribution for consumption and output remain much lower. Consumption trends towards zero, but remains at an estimated loss of 0.2 percentage points in 2050 for Japan. Output completely collapses with Japan's estimated demographic contribution at -2 percentage points at the end of the projection.

Secular stagnation may be even more pronounced in my estimates when considering the older economies of Japan and Europe. While I can only speak to these estimates that have been made on 18 advanced economies, there are similarly challenging underlying demographic trends for much of east Asia, with China and Korea set to rapidly age in the coming years. Countries like India, Brazil and Indonesia represent large populations still on the positive side of their demographic dividend but are also starting to transition toward a lower fertility rate while Africa and much of central Asia continue to see rapid growth. Whether or not secular stagnation spreads from an advanced economy problem to a global one may depend on the degree to which these areas of population growth are poised to take the advantage of the potential demographic dividends ahead of them.

Conclusions

A retrospective on the secular stagnation hypothesis suggests that population age structure indeed appears poised to be a drag on growth as advanced economies age. The boomer cohort reversed this trend for much of the late 20th and early 21st centuries, driving a great deal of movement with their entry and exit from the workforce. While the above estimates are descriptive in nature, and should not be taken as causal, they establish important evidence in the long run that these

channels are meaningful. These forces are being studied widely in macro models, and this historical context should be useful to researchers wishing to study these mechanisms in such theoretical work. In particular, it would be useful to turn the lens of quantitative macro models of secular stagnation to a historical context, incorporating not only these findings, but also the vast existing empirical literature on population age structures and growth. Current aging may pose new challenges, but the implications of these models would be easier to contextualize if such research is conducted.

Not modeled here is a causal link between demography and interest rates, nor am I able to causally identify movements in age structure relative to growth rates in investment, consumption, and output. What the above estimations show is a strong empirical connection. My findings are broadly in line with others that the link between aging and these financial markets appear to be important for the transmission of population age estimates to output growth. These results support ongoing work to understand mechanisms whereby aging may drive secular stagnation, with a particular focus on the specific mechanisms that connect them to interest rates. However more empirical work exploring this transmission mechanisms will be critical in sorting whether a linear relationship like those shown above is appropriate, or whether non-linearities are required. The modern empirical evidence of Eggertsson et al. (2019a) around recent zero lower bound estimates suggests such non-linearities, but my results may imply some degree of correlation between demographic impact for a broader range of interest rates above the ZLB. If indeed demographic determination of interest rates through models such as Gagnon et al. (2021) or Eggertsson et al. (2019b) are true, then the demographic headwinds implied by my estimates would be stronger. However, I find little evidence that this correlation matters much for consumption growth or capital deepening, and so while demographics themselves seem important, this financial channel may ultimately be less important for welfare.

The failure of the predictions (Hansen, 1939) appear to be in part related to the dramatic, but temporary, changes to population trends rather than failure of macroeconomic aggregates to move in expected ways in response to them. In this paper, I provide evidence that largely supports the idea that population aging may exhibit persistent impacts on growth rates of investment, consumption, and output. These longer run patterns in the future may be less pronounced than the demographic impact of the large baby boomer cohort moving through the labor force, but my estimates still project quantitatively meaningful stagnation in these growth rates going forward.

CRedit authorship contribution statement

Joseph Kopecky: Research, Analysis, Writing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.




Fig. 8. Predicted impact of age structure: France, Italy, UK, and Japan.

Linear prediction with all variables set to long run mean and all year fixed effects set to zero, changes in growth rates in each variable thus reflect only influence of variation in five-year population share growth rates captured by D1–D3.

Appendix A. construction of demographic controls

Consider the naive specification from Eq. (1) in the paper:

$$g_{i,t} = \sum_j \alpha_j \Delta p_{j,i,t} + X_{i,t} \beta + \mu_i + \mu_t + \nu + \epsilon_{i,t} \quad (6)$$

Where $g_{i,t}$ an outcome, the growth rate of a macroeconomic aggregate, $X_{i,t}$ is an arbitrary vector of controls, $\epsilon_{i,t}$ is error, μ_i country fixed effects, μ_t time fixed effects, and ν a constant. The variable $\Delta p_{j,i,t}$ are changes in population age shares, where the population is divided into J bins. This equation is inestimable as currently specified due to the




Fig. 9. Age coefficients: Changing polynomial order.

perfect colinearity of the population shares $p_{j,i,t}$. In principle one might like a very granular approach, estimating a large number of J groups. As described in the paper these shares are highly colinear with one another, more so as J increases. As in the paper I make the following two assumptions:

1. Letting α_j be the coefficient on population share $\Delta p_{j,i,t}$ of age group j in country i and time t . Assume that all of the effects of these coefficients across the age distribution sum to zero. In other words:

$$\sum_j \alpha_j = 0$$

2. Assume that the age coefficients α_j can be fitted with a K order polynomial. In other words:

$$\alpha_j = \sum_k \theta_k j^k \quad (7)$$

The problem is then transformed from one of estimating J coefficients into estimating the K , θ coefficients of the polynomial. The

assumption that all age effects, α_j , sum to zero makes them jointly estimable without having to drop the regression constant, implicitly defining the constant θ_0 of the age-coefficient polynomial. Forcing the age coefficient estimates to lie on this fitted polynomial requires that there be relatively smooth transitions from the effect of one age group to another. Taking only the first assumption makes estimation of Eq. (6) possible. Resulting estimations with large number of colinear age groups tend to be highly unstable, often switching signs from group to the next. To further show how this methodology works, substitute α_j as defined in Eq. (7) into Eq. (6). As in the paper, I set $K = 3$, using a third order polynomial. This yields:

$$g_{i,t} = \sum_j [(\theta_0 + \theta_1 j + \theta_2 j^2 + \theta_3 j^3) \Delta p_{j,i,t}] + \beta X_{i,t} + \mu_i + \mu_t + \nu + \epsilon_{i,t} \quad (8)$$

Summing over both sides of Eq. (7), and using the assumption that the sum of α_j must be zero, it can easily be shown that θ_0 is equal to:

$$\theta_0 = -\frac{1}{J} \left[\theta_1 \sum_j j + \theta_2 \sum_j j^2 + \theta_3 \sum_j j^3 \right] \quad (9)$$

Table 5
Robustness: including macroeconomic controls.

	Investment growth			Investment-to-GDP growth			Consumption growth			Output growth		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Population Growth	−0.65 (1.26)	−1.35 (1.31)	−1.47 (1.42)	−0.53 (0.92)	−1.16 (1.01)	−1.23 (1.16)	0.91*** (0.31)	0.89*** (0.30)	1.04*** (0.30)	0.73** (0.29)	0.78** (0.28)	0.95*** (0.29)
Demeaned Safe Rate, (r)	0.05 (0.19)	−0.22 (0.17)	−0.20 (0.21)	−0.05 (0.13)	−0.26** (0.12)	−0.27† (0.16)	0.06 (0.06)	0.04 (0.05)	0.00 (0.07)	0.06 (0.07)	0.06 (0.06)	0.02 (0.06)
Old-age Dependency Ratio		−0.98** (0.46)			−0.96*** (0.32)			−0.01 (0.13)			0.12 (0.16)	
Old-age Dependency Ratio × r		0.21** (0.08)			0.15** (0.06)			0.02 (0.02)			0.01 (0.02)	
D1			6.84*** (1.66)			4.63*** (1.06)			0.47 (0.47)			−0.25 (0.56)
D2			−10.06*** (2.98)			−7.37*** (1.89)			0.58 (0.95)			1.80† (1.13)
D3			3.85** (1.45)			2.92*** (0.92)			−0.62 (0.48)			−1.13* (0.60)
D1 × r			−0.25 (0.51)			−0.17 (0.38)			0.08 (0.10)			0.26** (0.10)
D2 × r			−0.40 (0.94)			0.00 (0.70)			−0.16 (0.16)			−0.52*** (0.15)
D3 × r			0.50 (0.47)			0.12 (0.35)			0.08 (0.09)			0.26*** (0.08)
Year	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Country	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
F-test D1-D3			6.69			9.68			8.07			4.56
(p-val)			0.00			0.00			0.00			0.02
F-test D1-D3 × r			3.92			1.87			0.39			4.03
(p-val)			0.03			0.17			0.76			0.02
R2	0.44	0.44	0.45	0.33	0.34	0.35	0.41	0.41	0.43	0.44	0.44	0.46
N	1548	1548	1548	1548	1548	1548	1548	1548	1548	1548	1548	1548

All regressions include controls for WW1, WW2, and Financial Crises as well as current inflation, unemployment, and government debt to GDP ratios. Standard errors clustered by country, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Distributing the summation term and recognizing that the first term, $\theta_0 \sum_j \Delta p_{j,i,t} = 0$ this can then be expressed as:

$$g_{i,t} = \theta_1 \sum_j (\Delta p_{j,i,t} j) + \theta_2 \sum_j (\Delta p_{j,i,t} j^2) + \theta_3 \sum_j (\Delta p_{j,i,t} j^3) + X_{i,t} \beta + \mu_i + \mu_t + \nu + \epsilon_{i,t} \quad (10)$$

The terms in parenthesis are the demographic variables I describe in Eq. (4) in the paper, which when used in this regression yield estimates of the polynomial fitting of the age specific coefficients. Notably in the original (Fair and Domínguez, 1991) paper population shares are used, rather than changes in these shares. This creates a small difference as these sum to one, instead of zero, which means that the θ_0 term does not drop out as it did here. The first term in this case is rather, $\theta_0 \sum_j \Delta p_{j,i,t} = \theta_0$, which simply requires one additional step of substituting the expression for θ_0 in Eq. (9) into Eq. (8) and combining like terms. The resulting demographic variables are each simply adjusted by a constant so ultimately this is not a critical distinction.

Appendix B. robustness to macroeconomic controls

To ensure the largest possible sample I exclude additional controls from my estimates. To show that their inclusion does not substantially alter my primary results I now report two tables, replicating Table 1, 2, 3, and 4 but now estimating:

$$g_{i,t} = \theta D_{i,t} + \beta r_{i,t} + \omega D_{i,t} \times r_{i,t} + \phi X_{i,t} + g_{i,t}^{pop} + \mu_i + \mu_t + \nu + \epsilon_{i,t} \quad (11)$$

To include a vector of, largely short run, variables that may be important. Notably this reduces the sample size substantially from 1840 to 1548, so not only do the results hold up to a change in control set, but also a reduction in power and change in sample used to estimate them. I choose three controls in particular to add to these estimates: inflation, unemployment rates, and government debt-to-gdp levels. The first two are meant to control for short term fluctuations, which in principle should be independent of longer term demographic forces. There is some evidence from work such as Juselius and Takáts (2021) that inflation may be linked to population age structure, an argument recently popularized by Goodhart and Pradhan (2020). Additionally unemployment might be affected by labor market tightness, which may be a function of population age structure. Finally government debt-to-gdp may likely be a function of population age structure as public pensions put a strain on government finances. Such cross correlations should be useful if they drive shorter term fluctuations in growth rates in a way that is related to age structure, but of course may also act as “bad controls” if they either remove the “good” variation in growth rates or provide a link for confounding biases.

These estimations are reported in Table 5. Broadly speaking there are not many reasons to be concerned. Despite inclusion of these controls, and a substantial change in sample, all of the demographic variables of interest, both in the old-age dependency ratio specifications and my baselines with full population age controls, have parameters with identical signs and similar magnitude. Their significance is weakened, but other than the interaction terms between polynomial controls and short term interest rates all remain significant and with lower power due to increases parameters and smaller sample some reduction in coefficient significance is to be expected.

Table 6
Additional tests of age structure.

	Investment growth			Investment-to-GDP growth			Consumption growth			GDP growth		
	Pre-1945	Post-1950	Ten year means	Pre-1945	Post-1950	Ten year means	Pre-1945	Post-1950	Ten year means	Pre-1945	Post-1950	Ten year means
g_{pop}	-0.99 (2.13)	0.57 (0.58)	-3.86** (1.71)	-0.56 (1.28)	-0.20 (0.49)	-1.76** (0.69)	0.76 (1.23)	0.34** (0.15)	0.84** (0.34)	0.95 (0.64)	0.35 (0.35)	0.32 (0.37)
stir	-0.67 (1.58)	0.25† (0.16)		-2.60* (1.37)	-0.34*** (0.09)		-0.23 (0.41)	0.01 (0.05)		-0.05 (0.32)	-0.03 (0.04)	
D1	19.08*** (4.39)	1.38† (0.89)	5.91† (3.79)	8.24 (5.58)	1.47** (0.56)	2.94*** (0.91)	2.79* (1.45)	0.57* (0.27)	0.79† (0.50)	-2.90*** (0.93)	-0.09 (0.25)	-0.94** (0.36)
D2	-33.47*** (10.56)	-5.12*** (1.67)	-12.04* (5.73)	-14.47 (11.15)	-3.70** (1.30)	-7.38*** (1.72)	-4.66* (2.32)	-0.44 (0.48)	-1.00 (1.05)	6.86*** (1.80)	0.78 (0.60)	2.58*** (0.85)
D3	16.04** (6.20)	3.09*** (0.86)	5.67** (2.42)	6.77 (6.30)	2.07*** (0.69)	3.88*** (0.84)	2.05* (1.14)	0.01 (0.24)	0.36 (0.53)	-3.71*** (1.03)	-0.51† (0.33)	-1.42*** (0.43)
D1 × r	-0.64 (3.52)	0.24 (0.30)	-0.25 (0.82)	1.78 (2.07)	0.19 (0.18)	0.05 (0.26)	-1.74* (0.86)	-0.11† (0.07)	-0.24* (0.12)	0.91 (0.73)	0.09 (0.09)	0.23† (0.14)
D2 × r	0.57 (5.53)	-0.73 (0.59)	-1.04 (1.43)	-3.25 (3.70)	-0.39 (0.40)	-0.52 (0.60)	4.07** (1.73)	0.20* (0.12)	0.37† (0.22)	-1.94 (1.61)	-0.17 (0.14)	-0.23 (0.21)
D3 × r	0.02 (2.70)	0.44 (0.30)	0.98 (0.73)	1.93 (1.94)	0.21 (0.22)	0.41 (0.35)	-2.31** (0.99)	-0.10† (0.06)	-0.16 (0.11)	0.97 (0.95)	0.08 (0.06)	0.05 (0.09)
Year	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Country	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
$\gamma_c \& \gamma_y$	✓	✓	✓	✓	✓	✓						
$\gamma_i \& \gamma_y$							✓	✓	✓			
$\gamma_i \& \gamma_c$										✓	✓	✓
F: D1-D3	8.10	5.60	1.90	3.16	3.95	7.74	1.37	3.09	2.74	5.04	1.91	4.46
(p-val)	0.00	0.01	0.17	0.06	0.03	0.00	0.30	0.05	0.08	0.02	0.17	0.02
F: D1-D3 × r	0.32	1.29	2.53	1.24	0.38	2.33	2.23	1.11	1.54	3.57	1.63	2.85
(p-val)	0.81	0.31	0.09	0.34	0.77	0.11	0.13	0.37	0.24	0.04	0.22	0.07
R2	0.39	0.55	0.66	0.28	0.32	0.57	0.51	0.68	0.85	0.55	0.73	0.86
N	633	1148	157	633	1148	157	633	1148	157	633	1148	157

All regressions include controls for WW1, WW2, and Financial Crises. $\gamma_{i|y|c}$ refers to inclusion of controls for the growth rate of investment/output/consumption respectively. Standard errors clustered by country, † $p < 0.15$, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix C. varying the degree of the population polynomial

An inherent weakness of my preferred choice of demographic control is the structure it imposes on the shape of coefficients across the age distribution. By imposing a third order fit I limit the number of turning points to two, and it is difficult for coefficients to rapidly change sign from one age to the next. This of course would be remedied by using distinct age shares, but these come with their own set of issues outlined in the paper. I can state the polynomial assumption more generally as:

$$\alpha_j = \sum_{k=0}^K \theta_k j^k \quad (12)$$

where I can choose any polynomial order, K . Increasing this order allows for more turning points in the age specific coefficients. As it stands I use $k = 3$, which allows that the age relationship can change direction twice. Using a $k = 6$ polynomial should allow fairly rapid changes in sign between age groups if the data fits better to do so. As a check that the demographic data are not overly constrained by my choice I re-estimate the main specifications in Section “Empirical results” using polynomials of degree 1–6. The corresponding age coefficients are plotted in Fig. 9. While the shape changes, the broad pattern across age remains similar to that from those reported in the paper. Indeed using a fourth order effect might strengthen my results in some ways, but since the literature studying Fair and Domínguez (1991) generally cautions against *over-fitting* these age shares, I take what I interpret as a more conservative estimate.

Appendix D. robustness: Changing the estimation sample

In this section I report estimates subject to changes in the sample. While I will include demeaned interest rates and their interaction with

demographics in this specification, their interpretation is less clear while conducting this exercise. My motivation for including interactions between demographics and interest rates above are to capture the relationship that has been well documented papers such as Gagnon et al. (2021), to allow that there may be direct correlation between age structure and my outcomes of interest, as well as one that works jointly with interest rates, due to shifting supply of savings and demand for investment conditional on the underlying age structure.

I estimate these equations for each of my outcomes of interest in Table 6. I begin by noting that, although there are some notable differences across these samples, the sign of all of the direct demographic controls are consistent across all specifications, and consistent with the signs reported in Tables 1–4. The interactions with demographics now tell a mixed picture. For the long run, ten year estimate this is likely in part due to the fact that there is little variation in this demeaned value when taken at this low frequency. For the pre-1945 and post-1950 estimates this may in part be a function of large differences in the level and volatility of rates across these periods.

As discussed in the text there are many intriguing aspects to splitting these estimates that may further be explored. For the context of my main demographic results I note that the shape of the implied age coefficients for these estimations in Fig. 10 are in line with those seen in the paper, with the notable difference being the extreme values estimated for the pre-1945, and the muted (though for many, still significant) for those post-1950. In particular, the coefficients at the extremes (aged over 75 below 15) have very large point estimates. A sign in the literature using Fair and Domínguez (1991) that the model may be overfitting these parameters. Future work that may investigate these asymmetries may wish to explore whether a different model is more appropriate, changing either the degree of the polynomial fit, or using more traditional age shares.




Fig. 10. Age coefficients: Various samples.

References

- Acemoglu, D., Restrepo, P., 2017. Secular stagnation? The effect of aging on economic growth in the age of automation. *Amer. Econ. Rev.* 107 (5), 174–179.
- Aksoy, Y., Basso, H.S., Smith, R.P., Grasl, T., 2019. Demographic structure and macroeconomic trends. *Am. Econ. J.: Macroecon.* 11 (1), 193–222.
- Ando, A., Modigliani, F., 1963. The "life cycle" hypothesis of saving: Aggregate implications and tests. *Am. Econ. Rev.* 53 (1), 55–84.
- Backus, D., Cooley, T., Henriksen, E., 2014. Demography and low-frequency capital flows. *J. Int. Econ.* 92, S94–S102.
- Baumol, W.J., 1952. The transactions demand for cash: An inventory theoretic approach. *Q. J. Econ.* 545–556.
- Becker, G.S., 1960. An economic analysis of fertility. *Demographic and economic change in developed countries*. In: NBER Conference Series, Vol. 11. pp. 209–231.
- Bloom, D.E., Canning, D., Fink, G., 2010. Implications of population ageing for economic growth. *Oxf. rev. econ. policy* 26 (4), 583–612.
- Bloom, D.E., Finlay, J.E., 2009. Demographic change and economic growth in Asia. *Asian Econ. Policy Rev.* 4 (1), 45–64.
- Bloom, D.E., Williamson, J.G., 1998. Demographic transitions and economic miracles in emerging Asia. *World Bank Econ. Rev.* 12 (3), 419–455.
- Carvalho, C., Ferrero, A., Nechio, F., 2016. Demographics and real interest rates: Inspecting the mechanism. *Eur. Econ. Rev.* 88, 208–226.
- Cervellati, M., Sunde, U., 2005. Human capital formation, life expectancy, and the process of development. *Amer. Econ. Rev.* 95 (5), 1653–1672.
- Cervellati, M., Sunde, U., 2011. Life expectancy and economic growth: The role of the demographic transition. *J. Econ. Growth* 16 (2), 99–133.
- Cooley, T., Henriksen, E., 2018. The demographic deficit. *J. Monetary Econ.* 93, 45–62.
- Cooley, T.F., Henriksen, E., Nusbaum, C., 2019. The growth-cost of demographic change in Europe. *Rev. Econ. Dyn.* 1352, 1–30.
- Eggertsson, G.B., Lancastr, M., Summers, L.H., 2019a. Aging, output per capita, and secular stagnation. *Am. Econ. Rev. Insights* 1 (3), 325–342.
- Eggertsson, G.B., Mehrotra, N.R., Robbins, J.A., 2019b. A model of secular stagnation: Theory and quantitative evaluation. *Am. Econ. J.: Macroecon.* 11 (1), 1–48.
- Fair, R.C., Domínguez, K.M., 1991. Effects of the changing US age distribution on macroeconomic equations. *Am. Econ. Rev.* 81 (5), 1276–1294.
- Gagnon, E., Johannsen, B.K., Lopez-Salido, D., 2021. Understanding the new normal: the role of demographics. *IMF Economic Review* 69, 357–390.
- Galor, O., Moav, O., 2004. From physical to human capital accumulation: Inequality and the process of development. *Rev. Econom. Stud.* 71 (4), 1001–1026.
- Galor, O., Weil, D.N., 2000. Population, technology, and growth: From malthusian stagnation to the demographic transition and beyond. *Amer. Econ. Rev.* 90 (4), 806–828.
- Goodhart, C., Pradhan, M., 2020. *The Great Demographic Reversal: Ageing Societies, Waning Inequality, and an Inflation Revival*. Springer Nature.
- Hansen, A.H., 1939. Economic progress and declining population growth. *Am. Econ. Rev.* 29 (1), 1–15.
- Higgins, M., 1998. Demography, national savings, and international capital flows. *Internat. Econom. Rev.* 343–369.
- Higgins, M., Williamson, J.G., 1997. Age structure dynamics in Asia and dependence on foreign capital. *Popul. Dev. Rev.* 261–293.
- Human Mortality Database, 2020. University of california berkeley, and max planck institute for demographic research. <https://www.mortality.org>.
- Jordà, Ò., Knoll, K., Kuvshinov, D., Schularick, M., Taylor, A.M., 2019. The rate of return on everything, 1870–2015. *The Quarterly Journal of Economics* 134 (3), 1225–1298.
- Jordà, Ò., Schularick, M., Taylor, A.M., 2017. Macrofinancial history and the new business cycle facts. *NBER Macroecon. Ann.* 31 (1), 213–263.

- Juselius, M., Takáts, E., 2021. Inflation and demography through time. *J. Econom. Dynam. Control* 104136.
- Kopecky, J., Taylor, A.M., 2020. The Murder-Suicide of the Rentier: Population Aging and the Risk Premium. Technical report, National Bureau of Economic Research.
- Krugman, P., 2014. Four observations on secular stagnation. *Secul. Stagnation Facts Causes Cures* 61–68.
- Lane, P.R., Milesi-Ferretti, G.M., 2001. Long-term capital movements. *NBER Macroecon. Ann.* 16, 73–116.
- Lee, R., Mason, A., 2010. Fertility, human capital, and economic growth over the demographic transition. *Eur. J. Popul./Revue Européenne de DÉMographie* 26 (2), 159–182.
- Lunsford, K.G., West, K.D., 2019. Some evidence on secular drivers of US safe real rates. *Am. Econ. J.: Macroecon.* 11 (4), 113–139.
- Mason, A., Lee, R., Jiang, J.X., 2016. Demographic dividends, human capital, and saving. *J. Econ. Ageing* 7, 106–122.
- Summers, L.H., 2014a. Reflections on the ‘new secular stagnation hypothesis’. *Secul. Stagnation Facts Causes Cures* (27–38).
- Summers, L.H., 2014b. US economic prospects: Secular stagnation, hysteresis, and the zero lower bound. *Bus. Econ.* 49 (2), 65–73.
- Taylor, A.M., Williamson, J.G., 1994. Capital flows to the new world as an intergenerational transfer. *J. Polit. Econ.* 102 (2), 348–371.
- Tobin, J., 1956. The interest-elasticity of transactions demand for cash. *Rev. Econ. Stat.* 241–247.
- UN, 2019. World Population Prospects. Department of Economic and Social Affairs, Population Division (2019), United Nations, Online Edition. Rev. 1..
- Vollrath, D., 2020. Fully grown. In: *Fully Grown*. University of Chicago Press.