CReAM
Centre for Research & Analysis of Migration

Discussion Paper Series

CDP 19/20

▶ **The Age-Productivity Profile: Long-Run Evidence from Italian Regions**

▶ Federico Barbiellini Amidei, Matteo Gomellini, Lorenzo Incoronato and Paolo Piselli

Centre for Research and Analysis of Migration
Department of Economics, University College London
Drayton House, 30 Gordon Street, London WC1H 0AX

www.cream-migration.org

The Age-Productivity Profile: Long-Run Evidence from Italian Regions[1]

Federico Barbiellini Amidei[*], Matteo Gomellini[*], Lorenzo Incoronato[†] and Paolo Piselli[*]

This version: September 2020

Abstract

This paper leverages spatial and time-series variation in the population age structure of Italian regions to uncover the causal effect of demographic shifts on labour productivity. Such effect is analysed along a 'first-order' channel stemming from the direct relation between an individual's age and productivity, and a 'second-order' channel that captures the productivity implications of a more or less dispersed age distribution. We propose an estimation framework that relates labour productivity to the entire age distribution of the working-age population and employs instrumental variable techniques to address endogeneity issues. The estimates return a hump-shaped age-productivity profile, with a peak between 35 and 40 years, as well as a positive productivity effect associated with a more dispersed age distribution.

Keywords: productivity, demography, age distribution, working-age population, long-run.

JEL codes: J11, J21, N30.

[1] This note is part of a broader research project that investigates the links between demography, productivity, innovation and entrepreneurship. We thank Fabrizio Balassone, Marta De Philippis, Sabrina Di Addario, Dario Pellegrino and Paolo Sestito for suggestions and comments on previous versions of the work. The views expressed in the article are those of the authors and do not necessarily reflect those of the Bank of Italy; all errors are our own.

[*]Bank of Italy, Directorate General for Economics, Statistics and Research, Structural Economic Analysis Directorate – Division of Economic History and Historical Archives.

[†]University College London (UCL) and Centre for Research and Analysis of Migration (CReAM).

1. Introduction

Developed economies have witnessed unprecedented demographic shifts over the past decades. These trends have been particularly pronounced in Europe, where the parallel decline in fertility and mortality rates led to a progressively older population. The old-age dependency ratio[2] rose from around 13 percent in 1955 to 31 percent today. According to the latest projections, the old-age dependency ratio in the EU-27 is set to reach 57 percent by 2100 (Eurostat, 2020). Within the European borders, the Italian picture deserves particular attention. The peninsula features in fact the highest old-age dependency ratio (almost 36 percent) and the highest share of the population aged more than 80 (roughly 7 percent). The median age of the Italian population stood at slightly below 29 years in 1950 and reached 45.5 years in 2018; by contrast, the median age in Europe moved from just above 29 to 41.8 over the same period. Importantly, the ageing trend has affected Italian regions asymmetrically: since 1952, the median age has risen by about 20 years in the South, 17 years in the Centre and 15 years in the North, from starting levels of 24, 32 and 30 years respectively.

In parallel, Italy is facing a prolonged period of stagnating productivity (e.g. Pellegrino and Zingales, 2017; Bugamelli et al., 2018) with notable heterogeneity in economic performance across different areas of the country. The North-South divide has been a widely discussed theme in the literature[3], with regional disparities detected across a broad range of economic development indicators. Output per worker in the South is currently about 20 percent lower than in the Centre-North (Banca d'Italia, 2018).

These phenomena are in part common to other countries and have been jointly analysed by an extensive body of research. Demography has been identified as an important determinant of long-run economic performance (e.g. Bloom et al., 2001; for Italy, see Barbiellini Amidei et al., 2018), and different demographic trends across countries or regions are now thought to contribute, at least in part, to diverging growth paths.

The proper quantification of the causal nexus between demographics and productivity, at different levels of aggregation, has thus become an important goal in empirical research. The sign of this relationship is a-priori undetermined, as a number of factors interplay and should be considered when attempting to identify such nexus (Skirbekk, 2003). First, the age-productivity relation is likely to be heterogeneous across sectors and occupations. Changes in age induce changes in individual skills (both physical and cognitive) that might matter differently for productivity depending on the worker's task[4]. Moreover, individual skills and the demand for such skills evolve over time, possibly making the age-productivity profile differ across generations. The age-productivity profile is also highly endogenous with respect to labour market institutions (e.g., retirement age, employment protection policies, human

[2] Defined as the ratio between the number of persons aged 65 and over and the number of persons aged between 15 and 64. The value is expressed per 100 persons of working age (15-64).
[3] For recent contributions see Federico et al. (2017), Felice (2017), Istat (2018).
[4] van Ours (2010): "*while older people do not run as fast, there is no evidence of a mental productivity decline and little evidence of an increasing pay-productivity gap. The negative effects of ageing on productivity should not be exaggerated*".

resource management), which complicates its estimation whenever such institutions cannot properly be accounted for[5].

The aggregate impact of demographics on productivity has been estimated to be stronger than the corresponding direct, individual-level relation predicted by classic Mincer equations linking wages to experience. This hints at the presence of indirect effects making social returns to experience/age larger than the sum of the private returns (Feyrer 2008). Inspired by these observations, we posit that the age composition of the population affects productivity through both a 'first-order' effect – the individual-level relation between age and productivity - as well as a 'second-order' effect associated with the overall shape of the age distribution. Regional productivity will be favoured by, say, an increase in the pool of young workers, if the true individual-level age-productivity profile is downward sloping. Yet aggregate productivity may also depend on how dispersed the distribution of age is, on top of the impact of specific age cohorts.

What are the theoretical channels through which the second-order effect is likely to operate? Intuitively, age diversity is a proxy for overall population/workforce diversity, which in turn might shape productivity in different ways. On the one hand, a more (age) heterogeneous population brings with it a diverse set of skills and expertise, which spurs cross-fertilization of ideas and creativity. On the other hand, a diverse workforce might also be characterized by communication challenges or conflicting values and inclinations, with negative repercussions on regional productivity. These intuitions, borrowed from firm-level studies, can easily be extended to higher levels of aggregation such as the regional one.

This work embeds these theoretical considerations in a unified estimation framework, and explores the relationship between demographics and labour productivity by examining a panel of Italian regions between 1981 and 2011. The within-country design exploits both time-series and cross-sectional variation in the variables of interest, while eliminating confounding factors - e.g., labour market institutions differing across countries - that may undermine the validity of cross-country studies[6]. We are going to focus on the working-age population (those aged 15-64) rather than the total population throughout the analysis, as our reasoning relates to labour market-specific dynamics.

Our goal is to trace out a "pointwise" age-productivity profile. In other words, we aim to pin down the coefficient specific to each 1-year age share without ex-ante constraining particular age cohorts to be grouped together as commonly done in the literature. However,

[5] Economists have thus started to investigate other channels through which demographic forces might shape economic performances, partly to avoid some of the complications associated with the use of productivity as outcome variable (e.g. its multifaceted nature). One such channel is innovation; another one is entrepreneurship. Factors that determine an individual's propensity to innovate or start a firm – creativity, human capital, risk aversion, time discounting – are in fact strongly linked to age. One would thus expect shifts in the population age distribution to exert some influence on innovation output and firm creation.

[6] Admittedly, labour market and other economic institutions may differ and have asymmetric impacts also across regions within a country. However, this unobserved heterogeneity is likely to be less relevant if compared to the between-country case.

this is not straightforward. For the case of the working age population, one would in fact need to estimate fifty age coefficients (one for each group from 15 to 64), together with those attached to the fixed effects and other control variables. Multicollinearity issues would make coefficient estimates in such a granular specification highly imprecise and potentially very different between each other even for adjacent shares.

A convenient approach is that proposed in Fair and Dominguez (1991), which restricts age coefficients to sum to zero and imposes each of them to lie on a low-order polynomial. If a second order polynomial is adopted, the number of age-related parameters to be estimated collapses to just two, attached to a first- and a second-order moment of the age distribution. The fifty "structural" age parameters are then easily backed out from the two "reduced-form" ones. Moreover, the reduced-form coefficient associated with the second-order moment enables the researcher to identify the productivity implications of a more or less dispersed age distribution – what we referred to as second-order effect. Such a representation therefore channels plenty of information about the population age structure, while allowing for a parsimonious parameterization of the model.

We aim to estimate age-productivity profiles that are robust to potential confounders and to other sources of unobserved heterogeneity. We thus augment our specification by including region- and year-specific effects, along with time-varying controls capturing cohort-specific human capital and the structure of the regional economies. However, age-specific migration and mortality patterns still undermine the validity of the above strategy, as long as OLS techniques are adopted. For example, young workers might have incentives to migrate towards more productive regions, which would induce simultaneity bias in the estimates. We therefore resort to a number of instrumental variable (IV) strategies to address the endogeneity of demographic indicators. The most robust strategy leverages the strongly predetermined component of the population age structure and instruments current demographic indicators with their lagged values. More precisely, the number of people aged 15, 16…64 in a specific year is instrumented with the number of those aged 0, 1…49 fifteen years before (for a similar approach see Skans, 2008).

Our preferred specification points to a hump-shaped age-productivity profile peaking between 35 and 40 years. We exploit these estimates to perform simple back-of-the-envelope calculations that quantify the long-run productivity implications of the projected shifts in the age structure of Italian regions. Moreover, we estimate a positive contribution of regional age dispersion to labour productivity. The note is structured as follows: Section 2 briefly reviews the literature examining the relation between demographics and productivity; Section 3 introduces the dataset and presents a descriptive exercise; Section 4 presents the econometric analysis and Section 5 concludes.

2. Review of the literature

A rich literature has developed investigating the economic consequences of demographic shifts. Changes in the age structure of the population are commonly thought to affect growth through two main channels. First, a higher old-age dependency ratio implies a contraction in the workforce and a rise in retirees, thus affecting production inputs, government expenditure and consumption/savings patterns. Second, to the extent that workers of different ages are differently productive, a shift in the workforce age structure inevitably affects aggregate productivity (Prskawetz 2005, Ilmakunnas et al. 2010, Nagarajan et al. 2013). In broad terms, the present paper aims to contribute to the literature that focuses on the latter channel.

The true shape of the age-productivity profile, at different levels of aggregation, is subject to a lively debate among economists (Ilmakunnas et al. 2010). From a theoretical standpoint, the individual-level age-productivity profile is expected to follow a hump-shaped pattern with a peak around 40, an age where the cognitive and physical abilities typical of the youth could optimally combine with the gains from the experience of older workers. In this regard, Skirbekk (2004) outlines how several factors that change with age, such as physical abilities, mental abilities, education and job experience, are also thought to contribute to an individual's productivity potential. The empirical studies surveyed by the author tend to confirm, at the individual level, the hump-shaped pattern predicted by the theory.

At the firm level, empirical results are instead more controversial. Some studies confirm the hump-shaped profile observed in individual-level research (e.g., Aubert and Crépon 2003). Other studies, however, find profiles that either flatten after the peak at around 40 years (e.g. Göbel and Zwick 2009), or even keep trending upwards (Mahlberg et al. 2013). Iparaguirre (2020) provides an excellent and up-to-date survey.

At higher levels of aggregation, a number of contributions to the empirical growth literature have explored the connection between the demographic structure of the workforce and aggregate productivity. While results tend to differ somewhat depending on the adopted methodology and the sample under scrutiny, an overall negative impact of ageing on productivity seems to prevail. Feyrer (2007) studies a cross-country panel between 1960 and 1990 and identifies a hump-shaped pattern whereby a relative increase in the cohort aged 40-49 raises aggregate productivity growth. Aiyar et al. (2016) report a negative and significant effect of workforce ageing - a higher fraction of workers older than 55 - on the real growth of output per worker using a panel of Euro area economies between 1950 and 2014. Maestas et al. (2016) conduct a similar analysis across US states and estimate that a 10 percent increase in the share of the population aged at least 60 would result in a 5.5 percent reduction in the growth of GDP per capita, largely due to a slowdown in labour productivity.

On Italy, Barbiellini Amidei et al. (2018) perform an accounting exercise to quantify the long-run contribution of demographic forces to the growth in GDP per capita in Italy and other industrialized economies. They document a positive demographic dividend[7] for most of the 20th century until the 1990s, when the dividend turned negative following the surge in the dependency ratio. Using population forecasts they show that, all else equal (in particular,

[7] The portion of economic growth accounted for by the growth in the working-age population (Bloom et al., 2001).

holding productivity at its 2016 level), the demographic dynamics expected over the next 45 years would determine a 24.4 percent drop in GDP compared to the 2016 value, 16.2 percent in per capita terms (around -0.4 percent on average each year). Ciccarelli et al. (2018) estimate the relationship between labour productivity in manufacturing and the availability of young people, and compare the results obtained focusing on the beginning and the end of the twentieth century, respectively. Using province-level data, they observe a positive relationship in both periods, with no sign of weakening over time.

Other authors argue instead that the relation between age and productivity might even be positive. In this vein, Skans (2008) examines a panel of Swedish local labour markets and estimates a positive contribution to regional productivity coming from an increase in the workforce share of individuals aged 50-60. Acemoglu and Restrepo (2017) argue that an ageing-induced reduction in the working-age population can stimulate the adoption of labour-replacing technologies, which in turn positively affect the growth of GDP per capita.

Lastly, a related field of research explores the productivity implications stemming from workforce age dispersion, at the firm level. Again, the net effect arises as the combination of opposing forces. On the one hand, workforce age homogeneity might proxy similarity of values and inclinations across workers, thereby easing communication and knowledge transmission. On the other hand, a different age implies diversity of skills and expertise, which may generate fruitful complementarities. The overall net effect of age dispersion hence varies across sectors and over time, and its sign is a priori undetermined. Unsurprisingly, findings in the literature range from a negative effect of age diversity, to a hump-shape relation, to no relationship at all (see Østergaard et al. 2011, Frosch 2011, Hammermann et al. 2019). A recent contribution (Zelity, 2020) establishes a hump-shaped pattern between age diversity and aggregate productivity using regional data from Europe.

3. Data and descriptive analysis

In this section we use data about the age structure of Italian regions collected from the databases of the Italian National Institute of Statistics (Istat), over a period spanning the entire second half of the 20th century. Due to the limited availability of historical information about the detailed age composition of regional workforces, the core of the analysis is going to focus on the working-age population (the total number of people aged 15 to 64). If in principle workforce data would more neatly suit the purpose of this study, the age structure of the working-age population – the pool of potential workers available at a given point in time – serves as a valid proxy for the age structure of the workforce. In addition, the use of population rather than workforce data circumvents a possible source of bias linked to the labour market participation rate, which may be endogenously determined.

The age structure of the Italian population has undergone relevant transformations over the past decades (Figure 1). Between 1952 and 2011, the average age has risen by 11.5 years for the total population and by 4.5 for the working-age population. Holding age shares fixed at the 1952 level, such an increase in the mean age has been equivalent, for the case of the working-age population, to the entire cohort aged 15-22 hypothetically disappearing. Progressively, the Italian workforce has been bearing the pressure of a mounting share of elderly people, with the old-age dependency ratio rising steadily between 1990 and 2011 by almost half percent a year. The country's total population became (on average) older than the working-age population in the early 1990s.

Figure 1. Demographic Trends in Italy

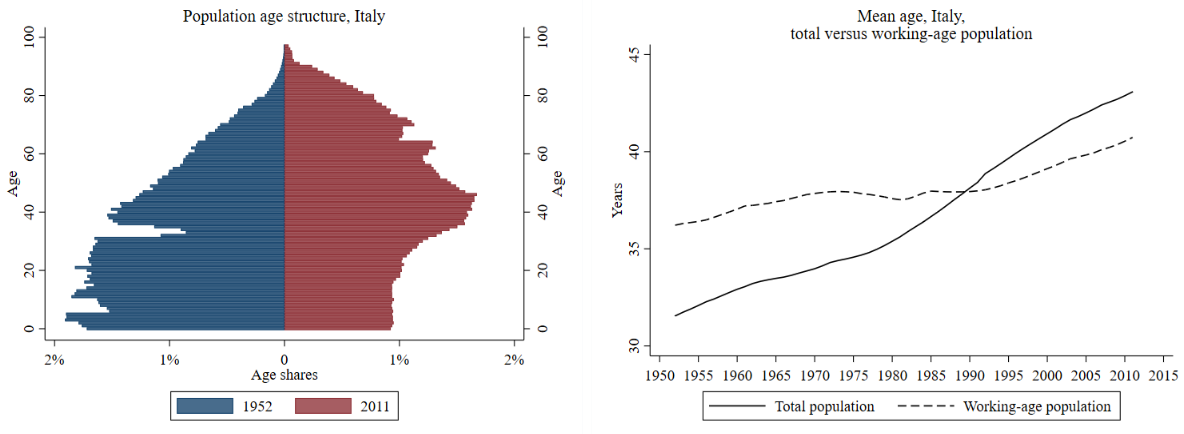

Note: Authors' elaborations on Istat data.

The intensity of these demographic shifts has not been homogeneous across Italy. The working-age population in regions belonging to Southern Italy was on average younger than in other regions in 1952 but has been going through a more rapid ageing process over the ensuing 60 years[8] (Figure 2, left panel). For instance, the mean age of the working-age population in Basilicata rose from 34 years in 1952 to 40 years in 2011, twice as much as in Piemonte where the increase was of just 3 years (from 38.6 to 41.6) over the same period. This has largely been the result of a more marked decline in the share of younger cohorts for Southern regions over this period[9]. Despite these shifts, the working-age population in the South remains on average younger than in the rest of the country. A further dimension of interest is the dispersion of the age distribution, which we capture using the coefficient of variation of age within the working-age population. A first glance at the data reveals a country-wide reduction in age dispersion over the sample period, as well as clear differences in terms of macro-areas. Despite having become more homogeneous, the working-age population remains in fact more (age) diverse in the South than in the rest of the country (Figure 2, right panel).

[8] This process has been largely driven by South-to-North migration flows following World War II.

[9] For instance, the share of people aged 15-34 decreased on average by 16 percent in the South and 12 percent in Northern regions between 1952 and 2011.

Figure 2. Regional Heterogeneity

Note: Authors' elaborations on Istat data. Mezzogiorno refers to the following Southern regions: Abruzzo, Basilicata, Calabria, Campania, Molise, Puglia, Sardegna, Sicilia.

The outcome variable of our analysis is regional labour productivity measured as the natural logarithm of real regional GDP per worker in Euros (chained values, 2000). The data is sourced from Prometeia over a period spanning 1971 to 2011. As anticipated earlier, a neat divide characterizes the country: labour productivity has been on average about 30 percent lower in the South than in the rest of Italy between 1971 and 2011.

At the descriptive level, productivity positively correlates with mean age and negatively with age dispersion. This is obviously a quite preliminary glance at the first- and second-order age effects introduced above: as already showed, regions in the South are younger and feature a more diverse age structure within the working-age population relative to other areas of the country. These patterns are showed in Figure 3, which portrays the cross-sectional variation between these demographic traits and labour productivity. The graphs show the natural logarithm of the real output per worker across Italian regions scattered against mean age (left) and age dispersion measured as the coefficient of variation of age (right) in the working-age population. All variables are averaged between 1971 and 2011.

Figure 3. Demographics and Productivity: Descriptive Evidence

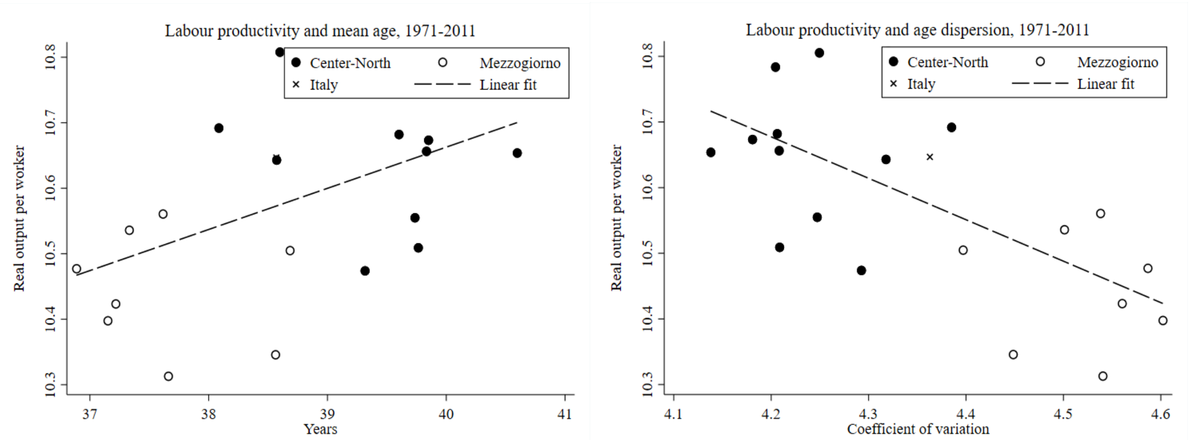

Note: Authors' elaborations on Istat data. Mezzogiorno refers to the following Southern regions: Abruzzo, Basilicata, Calabria, Campania, Molise, Puglia, Sardegna, Sicilia.

This preliminary exploration of the data shows how productivity reaches higher levels in regions featuring an older and relatively homogeneous working-age population - incidentally, these are all characteristics of regions in the Centre-North. A closer look at the above Figures reveals that the within macro-area correlations tend to differ from the between macro-areas correlations. In turn, this hints at the fact that the patterns of correlation just examined do not provide evidence about the direction of the nexus between demographics and output per worker, let alone about the existence of any causal link. Unobserved factors can in fact determine both the age structure and labour productivity of a region and thus generate spurious correlation between them. Furthermore, even accounting for any potential source of unobserved heterogeneity, estimates would still suffer from reverse causality bias. An attempt at identifying causal effects is the objective of the next section.

4. Empirical analysis

In this section, we perform a more thorough investigation of the impact that the regional age distribution has on labour productivity. We begin with preliminary regressions where the working-age population is broken down into five 10-year sub-groups (15-24 to 55-64), which enter our regressions as independent variables, in natural logarithm. In the attempt to capture possible second-order age effects, we allow the regional age distribution to affect labour productivity through its dispersion, too. In its most general formulation, the estimated equation can be written as follows:

$$\ln(y_{it}) = \alpha + \gamma_i + \theta_t + \sum_{j=1}^{5} \delta_j \cdot \ln(sh_{it}^{j}) + \rho \cdot \ln(CV_{it}) + X_{it}' \cdot \beta + \varepsilon_{it} \tag{1a}$$

where y_{it} denotes real output per worker in region i and year t, γ_i and θ_t are region and time fixed effects and X_{it} is a matrix of control variables specified below. The demographic indicators of interest are the five age shares sh_{it}^{j} and the coefficient of variation CV_{it}, a standardized measure of dispersion of the regional age distribution capturing second-order age effects[10].

Notice that, since age shares enter the regression in natural logarithm, there is no need to exclude one to avoid perfect collinearity issues, as would be the case had we not computed their logarithms. This modelling choice has important consequences on the interpretation of coefficients. On the one hand, it allows to view them somewhat more intuitively as elasticities. On the other, these elasticities denote the productivity effect of an inflow into the age group of interest from any of the other four, so that the effect of a change in a share also depends on the elasticity of the share that shrinks. We show below that our qualitative results remain unchanged when we do not take the natural logarithm of age shares and exclude a share as, for instance, in Feyrer (2007). At this stage, we would also caution the reader against interpreting

[10] We have also tested a richer specification including a higher-order moment of the age distribution such as its skewness. The coefficient associated with the latter is not significant and the magnitude and significance of the other coefficients remains unchanged.

our results as simple ageing effects. Our empirical strategy does not in fact overcome the classic ageing-cohorts-year effects identification problem and, more specifically, is unable to disentangle cohort-specific traits that change over time and are related to productivity. One such trait might be human capital. For example, the share of people aged 30-34 with a college degree rose more than threefold over our sample period. We show below that results are robust to including cohort-specific schooling levels as a control.

Table 1 reports the estimation output for a battery of variants of Equation (1a). The estimation sample period is 1981-2011 and standard errors are clustered regionally to allow for arbitrary time correlations among observations within a region. Column (1) shows the estimated coefficients of a simple pooled OLS regression of labour productivity exclusively on the (log) age shares of the working-age population. A hump shaped pattern begins to emerge, with a peak in the 25-34 cohort. In Column (2) we test a richer specification including region and year fixed effects along with the coefficient of variation of age within the working-age population. A hump-shaped profile can still be detected, yet coefficients lose most of their statistical significance. The only age group that still appears to matter is the aforementioned 25-34 cohort: a one percent inflow to that cohort from any of the others would result, according to these estimates, in a 0.8 percent rise in labour productivity. In line with the right panel of Figure 3, a negative correlation exists between age dispersion and productivity, albeit not statistically significant.

These findings are broadly confirmed in the most complete specification of Column (3), where also other potential confounders are controlled for. The within-country design significantly narrows down the list of possible confounders as many of them, such as the institutional setting, do not vary within the Italian borders. However, the substantial degree of (time-varying) heterogeneity across Italian regions makes it reasonable to include additional controls to corroborate the results. Specifically, in the same vein of Maestas et al. (2016), we choose to control for the value-added sectoral composition[11]. As anticipated below we also control for the share of young people (30-34) with a college degree, to partially address concerns about cohort-specific trends driving our results[12]. If anything, including these controls restores some of the explanatory power of demographic indicators.

A remaining concern with these estimates is that they fail to account for the simultaneity bias mentioned earlier. The panel structure of the dataset allows to control for time-invariant, region-specific unobserved heterogeneity through the inclusion of region identifiers. Moreover, year effects eliminate any spurious correlation with the business cycle. Yet reverse causality between demographics and productivity may still bias the estimates, which as a result could pick up not only the impact of the regional age structure on the outcome variable but also

[11] In order to identify sector-specific age effects, we have also interacted demographic variables with the shares of workers in different sectors of the economy. The coefficients associated with interaction terms are not significant and the other coefficients do not substantially change relative to the baseline specification.
[12] We use decennial census data (http://ottomilacensus.istat.it), linearly interpolated on an annual basis. Results are unchanged when using average years of schooling as a control for education, by updating the reconstruction performed in Bronzini and Piselli (2009).

any effect moving in the opposite direction. One such effect could be driven by, say, the migration of younger workers towards more economically vibrant regions, which would shift the age structure of both the receiving and the sending region and in turn induce a bias in the estimated coefficients.

Note that a permanently high/low level of regional labour productivity would be captured by the fixed effects and thus not pose a threat to identification. The concern is instead that any temporary shock to productivity in one region may induce a change in the age structure of its population. Therefore, proper identification of the causal effect of interest requires isolating a source of exogenous variation in the age structure of region i in year t that does not correlate with regional output per worker. We seek to isolate such variation by means of an IV strategy.

We instrument the regional age structure with its 15-years lag. For example, the 25-34 share of the working-age population (15-64) in 1995 is instrumented with the 10-19 share of the population aged 0-49 in 1980 (for a similar approach see Skans, 2008 and Bönte et al., 2009). Instrument validity requires that, conditional on the included controls, region effects and year effects, the regional age structure in year $t-15$ and its determinants have no effect on productivity in year t if not through the year t regional age structure. In the above example, the exclusion restriction would fail if a shock that affected, say, the number of people aged 10-19 in 1980, continues to weigh on productivity in 1995 not via the number of those aged 25-34 in 1995. In practice, such restriction would be violated if the number of people aged 10-19 in 1980 in a given region was to vary in anticipation of a larger - not permanently, as region dummies would capture it - productivity output.

Column (4) in Table 1 reports the Two-Stage Least Squares (2SLS) coefficient estimates when age shares and the coefficient of variation are instrumented with their lags. The same is performed in Column (5), where we also include the same set of controls as in Column (3). Unsurprisingly, the instrument does well at predicting the endogenous regressors. Relatively to OLS estimates, a more symmetric hump-shape peaking in the middle cohort (35-44) is observed. The explanatory power and statistical significance of demographic forces rises overall, especially in the most complete specification of Column (5). Remarkably, while estimates in Columns (1) to (3) fail to detect any significant second-order effect, the adoption an IV approach uncovers a positive impact of age diversity on productivity across regions, in stark contrast with the descriptive evidence above[13].

We also resort to a second instrument based on Shimer (2001) and constructed using lagged births. The number of, say, those aged 15 in 1981 in a given region is instrumented with the number of those born in the same region 15 years before in 1966. The number of those aged 16 in 1981 is instrumented with the number of those born 16 years before, in 1965. And so on

[13] All our results hold when a broader definition of working-age population that encompasses also the 65-74 age cohort is adopted. In particular, the hump-shape is still peaking in the 35-44 age group and a negative productivity effect is associated with the 65-74 cohort.

for the whole 15-64 cohort. The instrument's first-stage is however below standard threshold levels, and we report the related estimates in the Appendix[14].

A remaining issue with our estimates is that the low number of clusters (20) might make inference misleading. We address this concern by computing standard errors using the wild bootstrap procedure devised in Roodman et al. (2019). While this partially reduces statistical significance of demographic indicators, our overall results remain unchanged as is clear from Table A2 in the Appendix.

Table 1: Age Effects on Labour Productivity

	Pooled OLS	*Fixed effects*		*2SLS*	
	(1)	(2)	(3)	(4)	(5)
Share of 15-24	0.311	0.557	0.707**	-0.501	-0.262
	(0.288)	(0.324)	(0.303)	(0.495)	(0.474)
Share of 25-34	1.485***	0.814**	0.914**	0.847*	1.185***
	(0.297)	(0.328)	(0.329)	(0.476)	(0.376)
Share of 35-44	1.399***	0.769	0.845*	1.498***	1.671***
	(0.373)	(0.484)	(0.441)	(0.499)	(0.379)
Share of 45-54	0.852**	0.374	0.274	1.003***	0.937***
	(0.309)	(0.253)	(0.256)	(0.259)	(0.222)
Share of 55-64	0.659**	0.293	0.354*	0.165	0.246
	(0.288)	(0.199)	(0.199)	(0.234)	(0.231)
Age Variation Coefficient	-	-0.574	-0.650	4.481**	4.078**
		(1.345)	(1.251)	(1.781)	(1.812)
Share of Young with Higher Education	-	-	0.213	-	0.251**
			(0.127)		(0.108)
Value Added, Agriculture Share	-	-	0.940	-	0.615
			(0.790)		(0.848)
Value Added, Industry Share	-	-	0.304	-	0.186
			(0.354)		(0.387)
Kleibergen-Paap rk Wald F statistic (H0: Weak IV)	-	-	-	12.00	24.74
Stock-Wright LM S statistic (H0: Orthogonality)	-	-	-	7.07	11.92
Region effects	No	Yes	Yes	Yes	Yes
Year effects	No	Yes	Yes	Yes	Yes
Observations	620	620	620	620	620
R-squared	0.681	0.941	0.950	0.937	0.944

Notes: The dependent variable is the natural logarithm of real GDP per worker. Population shares computed as fraction of the working-age-population are in natural logarithm. The coefficient of variation is computed as the natural logarithm of mean age divided by age standard deviation. The share of young with higher education is the share of those aged 30-34 with a college degree. Standard errors cluster-corrected at regional level in parentheses. *p<0.1, **p<0.05, ***p<0.01.

[14] We also constructed a Bartik-style shift-share instrument that imposes beginning-of-period regional age shares to grow at the same rate as the correspondent age share at the national level. However, the instrument's first stage is extremely weak and the related results are omitted.

As anticipated above, we also test a slightly different specification where simple age shares and not their natural logarithms enter the right-hand side. As these shares sum to one, we need to exclude one of them to avoid multicollinearity issues. We omit the middle cohort (35-44) so that the coefficient attached to a particular share among the included ones (15-24, 25-34, 45-54 and 55-64) denotes the productivity impact of a (working-age) population flow into that specific share from the 35-44 group. Statistical significance of a coefficient would thus imply that it is significantly different from the implied zero coefficient on the 35-44 age group.

$$\ln(y_{it}) = \alpha + \gamma_i + \theta_t + \sum_{j=1}^{4} \delta_j \cdot sh_{it}^{j} + \rho \cdot \ln(CV_{it}) + X'_{it} \cdot \beta + \varepsilon_{it} \tag{1b}$$

Table 2 below reports 2SLS coefficient estimates of Equation (1b) under our preferred specification including the vector of controls, region and year effects. A clear hump-shaped profile can still be detected, as outflows from the omitted cohort (35-44) into the remaining four always result into significantly lower productivity. Again, a positive coefficient is associated with age dispersion, hinting at the presence of positive second-order age effects.

Table 2: Age Effects on Labour Productivity, Non-Logarithmic Shares

	(1)
Share of 15-24	-10.06**
	(4.122)
Share of 25-34	-2.141**
	(1.047)
Share of 45-54	-3.199***
	(1.210)
Share of 55-64	-6.779***
	(2.506)
Age Variation Coefficient	5.053**
	(2.399)
Share of Young with Higher Education	0.255**
	(0.117)
Value Added, Agriculture Share	0.883
	(0.912)
Value Added, Industry Share	0.183
	(0.431)
Kleibergen-Paap rk Wald F statistic (H0: Weak IV)	13.77
Stock-Wright LM S statistic (H0: Orthogonality)	10.51
Region effects	Yes
Year effects	Yes
Observations	620
R-squared	0.937

Notes: 2SLS coefficient estimates for Equation (1b). See text and Table 1 for details. Standard errors cluster-corrected at regional level in parentheses. *p<0.1, **p<0.05, ***p<0.01.

Polynomial approach

The empirical strategy adopted so far implicitly assumes that the causal effect of interest might change when moving from one of the five age groups to another. Yet these groups are formed ex-ante by grouping 1-year age cohorts together into five equally-sized subgroups, without any theoretical nor empirical justification. We thus allow each 1-year cohort to independently affect (log) regional productivity and obtain 'point-wise' age effects by relating the outcome variable for region i and year t, y_{it} to each of the fifty cohorts composing the working-age population, $sh_{it}^{j}, j = 15, \dots, 64$[15]. In the most general formulation, we also include a matrix of controls X_{it}, region and year effects:

$$\ln(y_{it}) = \alpha + \gamma_i + \theta_t + \sum_{j=15}^{64} \delta_j \cdot sh_{it}^{j} + X_{it}' \cdot \beta + \varepsilon_{it} \tag{1c}$$

Where $sh_{it}^{j} = pop_{it}^{j} / pop_{it}^{15-64}, j = 15, \dots, 64$ and pop_{it}^{15-64} is the total working-age population in region i and year t.

The large number of explanatory variables – fifty age coefficients along with fixed effects and other controls – makes such a granular estimation of the age profiles not immune to threats. The finer the division between consecutive age shares, the more serious estimation issues become (Juselius and Takats 2018). Multicollinearity between regressors would in fact severely affect coefficient estimates, whose precision would also worsen as the number of age cohorts rises compared to the number of periods. Moreover, leaving regression parameters unconstrained may lead to confusing age profiles, with estimated age effects possibly varying dramatically between consecutive cohorts.

Inspired by Fair and Dominguez (1991), we overcome these issues by imposing each of the fifty "structural" age coefficients to lie along a second-degree polynomial to smoothen out the estimated effect between adjacent cohorts[16]:

$$\delta_j = \eta_0 + j \cdot \eta_1 + j^2 \cdot \eta_2 , \quad j = 15, \dots, 64 \tag{2}$$

[15] The approach we use has the main advantage of being less arbitrary than grouping 1-year cohorts together before estimating age effects, in that each 1-year share is now allowed to influence the outcome variable independently of the others.

[16] Higher degree polynomials were also tested, but this led to no noticeable change in results. As will be clear shortly, this is consistent with previous estimates reporting the non-significance of higher-order moments of the age distribution, such as its skewness. At least in this case, first and second moments thus appear to be "sufficient statistics" to describe the age distribution in the working-age population.

We also constrain these coefficients to sum to zero in order to remove perfect collinearity between age shares (which sum to one) and the constant term:

$$\sum_{j=15}^{64} \delta_j = 0 \tag{3}$$

Combining (2) and (3) we express η_0 as a function of η_1 and η_2:

$$\eta_0 = -\frac{1}{50} \cdot \left(\eta_1 \cdot \sum_{j=15}^{64} j + \eta_2 \cdot \sum_{j=15}^{64} j^2\right) \tag{4}$$

Plugging (2) and (4) into (1c) leads to the following expression:

$$\ln(y_{it}) = \alpha + \gamma_i + \theta_t + \eta_1 \cdot \underbrace{\left(\sum_{j=15}^{64} j \cdot sh_{it}^j - \frac{1}{50} \cdot \sum_{j=15}^{64} j\right)}_{M_{it}^1} + \eta_2 \cdot \underbrace{\left(\sum_{j=15}^{64} j^2 \cdot sh_{it}^j - \frac{1}{50} \cdot \sum_{j=15}^{64} j^2\right)}_{M_{it}^2} + X_{it}' \cdot \beta + \varepsilon_{it} \tag{5}$$

These restrictions dramatically reduce the number of age-related parameters to be estimated to just two "reduced-form" ones (η_1, η_2) attached to a first- and a second-order moment of the age distribution (M_{it}^1, M_{it}^2)[17]. The last reduced-form coefficient η_0 can be easily pinned down from the estimates of η_1 and η_2 using (4) and in turn the values of the structural age parameters $\delta_j, j = 15, \ldots, 64$ are backed out from equation (2)[18].

Column (1) in Table 3 below reports OLS coefficient estimates for equation (5). The age-specific effects on productivity resulting from the above estimates are depicted in the left panel of Figure 4. Estimates for the δ_j's can be thought of as the (relative) productivity contribution associated with each 1-year age share – a granular breakdown of the first-order age effect

[17] Given the attractiveness of tracing out 'pointwise' age profiles while preserving a parsimonious model parameterization, several studies in the literature on demographics have adopted this methodology. See for instance Higgins (1998), Skans (2008), Juselius and Takats (2018).

[18] The standard errors of the structural age coefficients are computed in a similar way. Plugging (4) in (2):

$$\delta_j = \eta_1 \cdot \underbrace{\left(j - \frac{1}{50} \cdot \sum_{j=15}^{64} j\right)}_{c_1(j)} + \eta_2 \cdot \underbrace{\left(j^2 - \frac{1}{50} \cdot \sum_{j=15}^{64} j^2\right)}_{c_2(j)}, \qquad j = 15, \ldots, 64 \tag{7}$$

So that:

$$var\left(\widehat{\delta_j}\right) = c_1^2(j) \cdot var\left(\widehat{\eta_1}\right) + c_2^2(j) \cdot var(\widehat{\eta_2}) + 2c_1(j)c_2(j) \cdot cov(\widehat{\eta_1}, \widehat{\eta_2}), \qquad j = 15, \ldots, 64 \tag{8}$$

introduced above and first estimated in Table 1. More specifically, each point on the curve shows the age-specific contribution relative to the mean contribution, which is normalised to zero. A hump-shaped pattern emerges in line with preliminary estimates in Table 1, although the estimated age effects do not significantly differ from zero.

This approach allows to explore second-order effects, too. Specifically, we observe an almost one-to-one negative correlation between the second-order moment M_{it}^2 and the coefficient of variation of age, within each Italian region over our three-decades sample. The coefficient attached to M_{it}^2 thus conveys the same information as that associated with the variation coefficient in Table 1, that is, the productivity implications of having a more or less dispersed age distribution. OLS estimates in Table 3 point to insignificant second-order effects of age on productivity, similarly to those in Columns (2) and (3) of Table 1.

As previously noted, however, reverse causality concerns make demographic variables highly likely to remain endogenous even after including fixed effects and other potential confounders. The estimation bias affecting reduced form coefficients in equation (5) would in turn deliver incorrect estimates for the structural age coefficients and their standard errors. To address this issue, we augment the Fair and Dominguez (1991) polynomial specification by again resorting to an IV approach. We choose to employ lagged population shares as an instrument for current shares. Specifically, current values of demographic indicators $M_{it}^k, k = 1, 2$ are instrumented as follows:

$$M_{it}^{k,IV} = \sum_{j=0}^{49} j^k \cdot sh_{it}^{j,IV} - \frac{1}{50} \cdot \sum_{j=0}^{49} j^k, \;\; k = 1, 2 \tag{6}$$

Where $sh_{it}^{j,IV} = pop_{i,t-15}^{j} / pop_{i,t-15}^{0-49}$ is used as instrument for sh_{it}^{j}. In other words, each 1-year age share is instrumented with its 15-years lag – for instance, the working-age population share of people aged 50 in 1995 is instrumented with the share of people aged 35 in 1980 relative to the total population aged 0-49 in 1980.

Column (2) in Table 3 reports coefficient estimates for equation (5) when estimated through a 2SLS procedure using $M_{it}^{k,IV}$ as instrument for M_{it}^{k}. The right panel of Figure 4 portrays the resulting age-productivity profile.

The estimates again point to a hump-shaped age-productivity profile with a peak between 35 and 40 years old. Differently from OLS estimates, and again in line with results displayed in Table 1 (Columns (4) and (5)), the IV approach is able to restore statistical significance of demographic variables. Finally, the negative and significant coefficient attached to M_{it}^2 in Column (2) confirms what already observed in the previous section, that is, a positive relationship between age dispersion and labour productivity[19].

[19] These results are confirmed when standard errors are computed via wild bootstrap (Roodman et al. 2019). A hump-shaped profile still appears when we extend our definition of the working-age population to include the 65-74 age cohort. This is in line with the evidence referred to in Footnote 13.

Table 3. Estimation Output, Equation (5)

	OLS (1)	2SLS (2)
M_{it}^{1}	0.309	0.612**
	(0.240)	(0.285)
M_{it}^{2}	-0.00420	-0.00804**
	(0.00308)	(0.00365)
Kleibergen-Paap rk Wald F statistic (H0: Weak IV)	-	50,112
Stock-Wright LM S statistic	Chi (2):	8.69
(H0: Orthogonality)	p-value:	0.013
Region effects	Yes	Yes
Year effects	Yes	Yes
Additional controls	Yes	Yes
Observations	620	620
R-squared	0.936	0.934

Notes: Estimation output for Equation (5). Both specifications include controls for value added sectoral composition and the share of young people with college degree. See text for details. Standard errors cluster-corrected at regional level in parentheses. *p<0.1, **p<0.05, ***p<0.01.

Figure 4: Age-Productivity Profiles for Italy, 1981-2011

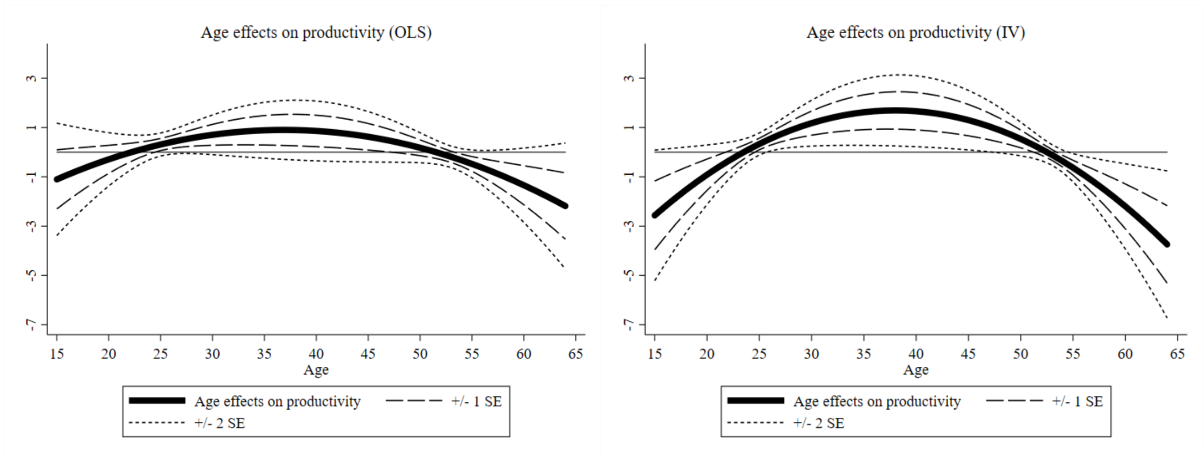

Notes: Estimates of structural age coefficients in equation (1c), resulting from estimation of equation (5) with fixed-effects OLS (left) and 2SLS (right). Standard errors clustered at the regional level.

Discussion

The main takeaway from this analysis is that the age-productivity relation is not monotonic. Therefore, to correctly gauge the economic consequences of demographic trends, one needs to consider how the age distribution evolves in its entirety without limiting the focus to, say, mean age, median age or particular age groups. In this regard, the availability of age-specific coefficients enables us to precisely quantify how past and future changes in the overall age structure of Italy's working-age population have affected, or are likely to affect, the country's productivity. We thus perform a simple back-of-the-envelope exercise by first-differencing Equation (1c) and plugging our estimates for the δ_j's, along with i) changes in age distribution between 2000 and 2019[20] and ii) projected changes between 2019 and 2030[21].

Between 2000 and 2019, the Italian working-age population witnessed a major increase in the cohorts aged 45 to 60, at the expenses of those aged 25 to 40. As is clear from Figure 5 below, the age groups that grew (decreased) relatively to others between 2000 and 2019 are those providing a relative negative (positive) contribution to productivity, according to our estimates. This demographic shift has thus likely been associated with a significant drop in productivity. We estimate such drop at around -0.7 percent per year on average (-12.5 percent over the whole period)[22], *coeteris paribus*. Quite a large impact, especially when compared with the 0.9 percent effective increase in GDP per worker over the same period (Bugamelli et al., 2018).

We observe negative demographic effects also when looking at the near future (2019-2030). The age structure of the Italian working-age-population is expected to shift towards the oldest cohorts (55-64) over the next decade, with most of the loss concentrated in the 40-52 age group. Indeed, our estimates point to a loss in labour productivity as large as -1.5 percent per year by 2030 due to the projected shifts in the working-age population and abstracting from movements in the other variables included in our specification. Interestingly, longer-term population projections paint a less pessimistic picture. Istat foresees a relative increase in the cohorts aged 30-40 by 2040 and 2050, which partly compensates for the expected rise in older cohorts. As a result, the predicted productivity effects of demographic shifts hover around -0.4 percent per year for 2040 and -0.1 percent for 2050 (cumulatively -7.4 and -3.1 per cent, respectively). We place less emphasis on these projections in light of their distance in the future, which implies higher uncertainty about other possible drivers of productivity[23].

[20] The 2019 age distribution is the actual distribution at January 1st; see Istat, "Demografia in cifre", www.demo.istat.it.
[21] We use Istat population forecasts, "Previsioni della Popolazione 2018-2065"; see http://dati.istat.it.
[22] In this calculation, coefficients not statistically different from zero (at the 95% confidence level) are set at zero. The result is practically unchanged when including the whole set of coefficients, irrespective of their significance.
[23] By and large, our estimates remain unchanged when including the 65-74 age group in the analysis.

Figure 5: Demographic Shifts and Effects on Productivity

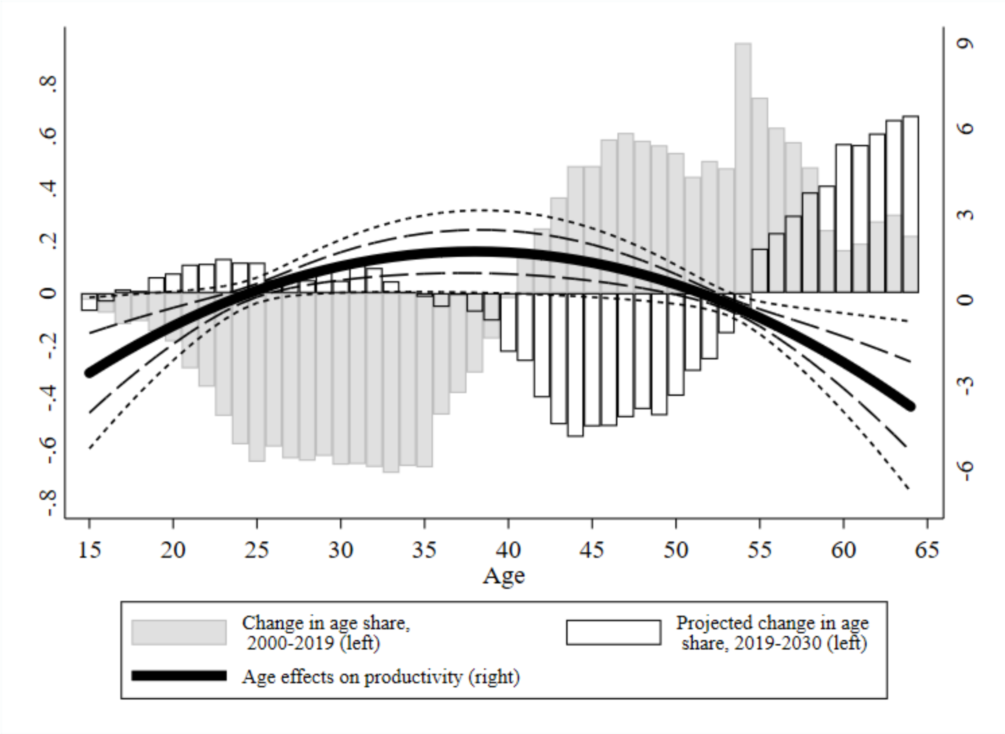

Notes: The percentage change in the working-age population share of each one-year age group between 2000 and 2019 and 2019 and 2030 (expected) is on the left axis (source Istat). Estimates of structural age coefficients in equation (1c), resulting from estimation of equation (5) with fixed-effects 2SLS, are on the right axis.

5. Conclusions

This paper produces novel evidence about the relationship between population age and aggregate labour productivity. The empirical analysis exploits time-series and cross-sectional variation in the age structure of the working-age population across Italian regions between 1981 and 2011. The still large productivity divide between the stagnating South and the rest of Italy, the relevant transformations witnessed by the country's population over the past decades as well as the notable heterogeneity in how regional age structures have evolved offer an advantageous perspective to explore the interplay between these phenomena.

Our instrumental-variable estimates are in line with the previous literature and point to a hump-shaped profile for the age-productivity relationship, with a peak between 35 and 40 years. Based on current population projections, these results imply a potential productivity loss as large as -1.5 percent per year until 2030, abstracting from changes in other variables.

The empirical framework explicitly allows the dispersion of the age distribution to impact labour productivity. This 'second-order' channel has received little emphasis in the literature. We estimate a positive and significant relationship between regional age dispersion, proxied by the coefficient of variation of the age distribution, and productivity. At the basis of this positive 'second-order' age effect might be that age heterogeneity entails diversity of skills and expertise and spurs cross-fertilization of ideas within the workforce.

Further investigations are warranted to shed light on our results. More importantly, a suitable theoretical framework that rationalizes our findings and additional empirical research on possible channels through which demography affects productivity – such as innovation and entrepreneurship – are much needed.

Appendix

Table A1

Replication of Table 1, Columns (4) and (5) Using Lagged Births Instruments (Shimer, 2001)

	(1)	(2)
Share of 15-24	10.48	4.554
	(15.35)	(5.596)
Share of 25-34	5.231*	3.538***
	(3.107)	(1.342)
Share of 35-44	-1.649	0.656
	(7.782)	(2.947)
Share of 45-54	-3.118	-0.709
	(7.207)	(2.727)
Share of 55-64	0.791	0.742
	(1.031)	(0.602)
Age Variation Coefficient	-52.99	-20.20
	(90.08)	(33.45)
Share of Young with Higher Education	-	0.302
		(0.230)
Value Added, Agriculture Share	-	2.997
		(3.306)
Value Added, Industry Share	-	-0.439
		(0.936)
Kleibergen-Paap rk Wald F statistic (H0: Weak IV)	0.1	0.2
Stock-Wright LM S statistic (H0: Orthogonality)	10.55	11.3
Region effects	Yes	Yes
Year effects	Yes	Yes
Observations	620	620
R-squared	0.202	0.799

Notes: The dependent variable is the natural logarithm of real GDP per worker. Population shares computed as fraction of the working-age-population are in natural logarithm. The coefficient of variation is computed as the natural logarithm of mean age divided by age standard deviation. The share of young with higher education is the share of those aged 30-34 with a college degree. Standard errors cluster-corrected at regional level in parentheses. *p<0.1, **p<0.05, ***p<0.01.

Table A2

Age Effects on Labour Productivity, Wild Bootstrap SEs (Roodman et al., 2019)

	Pooled OLS	*Fixed effects*		*2SLS*	
	(1)	(2)	(3)	(4)	(5)
Share of 15-24	0.311	0.557	0.707*	-0.501	-0.262
	(0.308)	(0.300)	(0.280)	(0.460)	(0.438)
Share of 25-34	1.485**	0.814**	0.914*	0.847*	1.185**
	(0.348)	(0.319)	(0.318)	(0.459)	(0.361)
Share of 35-44	1.399**	0.769	0.845	1.498***	1.671***
	(0.433)	(0.478)	(0.434)	(0.491)	(0.376)
Share of 45-54	0.852**	0.374	0.274	1.003***	0.937**
	(0.335)	(0.247)	(0.251)	(0.249)	(0.218)
Share of 55-64	0.659	0.293	0.354	0.165	0.246
	(0.311)	(0.193)	(0.193)	(0.222)	(0.218)
Age Variation Coefficient		-0.574	-0.650	4.481**	4.078**
		(1.250)	(1,150)	(1.607)	(1.639)
Share of Young with Higher Education			0.213		0.251*
			(0.125)		(0.104)
Value Added, Agriculture Share			0.940		0.615
			(0.790)		(0.826)
Value Added, Industry Share			0.304		0.186
			(0.350)		(0.379)
Kleibergen-Paap rk Wald F statistic (H0: Weak IV)	-	-	-	12.0	24.74
Stock-Wright LM S statistic (H0: Orthogonality)	-	-	-	7.07	11.92
Region effects	No	Yes	Yes	Yes	Yes
Year effects	No	Yes	Yes	Yes	Yes
Observations	620	620	620	620	620
R-squared	0.681	0.941	0.950	0.937	0.944

Notes: The dependent variable is the natural logarithm of real GDP per worker. Population shares computed as fraction of the working-age-population are in natural logarithm. The coefficient of variation is computed as the natural logarithm of mean age divided by age standard deviation. The share of young with higher education is the share of those aged 30-34 with a college degree. Standard errors computed with wild cluster bootstrapping procedure (Roodman et al., 2019) are in parentheses. *p<0.1, **p<0.05, ***p<0.01.

References

Acemoglu, D., Akcigit, U. and Celik, M. (2017), Young, Restless and Creative: Openness to Disruption and Creative Innovations, NBER Working Papers 19894 (updated), National Bureau of Economic Research, Inc.

Acemoglu, D. and Restrepo, P. (2017), Secular Stagnation? The Effect of Aging on Economic Growth in the Age of Automation, American Economic Review, American Economic Association, vol. 107(5), pages 174-179, May.

Aiyar, S., Ebeke, C. and X. Shao (2016), The Impact of Workforce Aging on European Productivity, IMF Working Paper, WP/16/238.

Aubert, P. and Crépon, B. (2003), La productivité des salariés âgés: une tentative d'estimation, Économie et Statistique, 368, issue 1, p. 95-119.

Banca d’Italia (2018), The economy of the Italian regions, n. 24.

Barbiellini Amidei, F., Gomellini, M. and Piselli, P. (2018), Il Contributo della Demografia alla Crescita Economica: Duecento Anni di Storia Italiana, Bank of Italy Occasional Papers, n. 431.

Bartik, Timothy. 1991. Who Benefits from State and Local Economic Development Policies? W.E. Upjohn Institute.

Bloom, D., Canning, D. and Sevilla, J.P., (2001), Economic Growth and the Demographic Transition, N. 8685, NBER Working Papers, National Bureau of Economic Research, Inc.

Bönte, W., Falck, O. and Heblich, S., (2009), The Impact of Regional Age Structure on Entrepreneurship, Economic Geography, 85, issue 3, p. 269-287.

Bronzini R., Piselli P. (2009), Determinants of long-run regional productivity with geographical spillovers: The role of R&D, human capital and public infrastructure, Regional Science and Urban Economics 39, pp. 187–199

Bugamelli, M. F. Lotti, M. Amici, E. Ciapanna, F. Colonna, F. D'Amuri, S. Giacomelli, A. Linarello, F. Manaresi, G. Palumbo, F. Scoccianti, E. Sette (2018), La crescita della produttività in Italia: la storia di un cambiamento al rallentatore, Bank of Italy Occasional Papers, n. 422, Roma.

Ciccarelli, C., Gomellini, M. and Sestito, P. (2019), Demography and Productivity in the Italian Manufacturing Industry: Yesterday and Today, CEIS Research Paper 457, Tor Vergata University, CEIS.

Eurostat (2020). Population Structure and Ageing.

Fair, R., and K. Dominguez (1991), Effects of the Changing US Age Distribution on Macroeconomic Equations, The American Economic Review, pp. 1276–1294.

Federico, G., Nuvolari, A., Vasta, M. (2019), The Origins of the Italian Regional Divide: Evidence from Real Wages, 1861-1913 The Journal of Economic History, 79, 1, pp. 63-98.

Felice, E. (2017), The Roots of a Dual Equilibrium: GDP, Productivity and Structural Change in the Italian Regions in the Long-Run (1871-2011). Bank of Italy Economic History Working Paper No. 40.

Feyrer, J. (2007), Demography and Productivity, Review of Economics and Statistics, 89, 1, pp. 100-109.

Feyrer, J. (2008), Aggregate Evidence on the Link between Age Structure and Productivity, Population and Development Review, pp. 78–99. Thuenen-Series of Applied Economic Theory, University of Rostock, Institute of Economics.

Frosch, K. H. (2011), Workforce Age and Innovation: A Literature Survey, International Journal of Management Reviews, 13(4), 414–430.

Göbel, C. and Zwick, T. (2009), Age and Productivity – Evidence From Linked Employer Employee Data, ZEW Discussion Paper No. 09-020, Centre for European Economic Research, Mannheim.

Hammermann, A., Niendorf, M. and Schmidt, J. (2019), Age Diversity and Innovation: Do Mixed Teams of Old and Experienced and Young and Restless Employees Foster Companies Innovativeness?, No 201904, IAB Discussion Paper

Higgins, M. (1998), Demography, National Savings, and International Capital Flows, International Economic Review, 39(2), pp. 343–369.

Ilmakunnas, P., van Ours, J., Skirbekk, V. and Weiss, M. (2010), Age and Productivity, in Garibaldi, P., Martins, J.O. and van Ours, J. (eds.), Ageing, Health, and Productivity: The Economics of Increased Life Expectancy, Oxford University Press, pp. 133-240.

Iparaguirre, J. (2020), Economics and Ageing, Springer Books.

ISTAT (2018). Il Benessere Equo e Sostenibile in Italia.

Juselius, Mikael and Takáts, Elod, The Enduring Link between Demography and Inflation (April 6, 2018). Bank of Finland Research Discussion Paper No. 8/2018.

Kopecky, J. (2017), An Aging Dynamo: Demographic Change and the Decline of Entrepreneurial Activity in the United States, Working paper.

Liang, J., Wang, H. and Lazear, E. P. (2018), Demographics and Entrepreneurship, Journal of Political Economy 126, 140–196.

Maestas, N., Mullen, K. and Powell, D. (2016), The Effect of Population Aging on Economic Growth, the Labour Force and Productivity, NBER Working Papers 22452, National Bureau of Economic Research, Inc.

Mahlberg, B., Freund, I. and Fürnkranz-Prskawetz, A. (2013), Ageing, productivity and wages in Austria: sector level evidence, Empirica, 40, issue 4, p. 561-584.

Nagarajan, R., Teixeira, A., Silva, S. (2013), The Impact of an Ageing Population on Economic Growth: an Exploratory Review of the Main Mechanisms, FEP Working Papers 504, Universidade do Porto, Faculdade de Economia do Porto.

Østergaard, C. R., Timmermans, B., Kristinsson, K. (2011), Does a Different View Create Something New? The Effect of Employee Diversity on Innovation. Research policy 40, 500–509.

Pellegrino, B. and Zingales, L. (2017), Diagnosing the Italian disease, NBER Working Paper no. w23964.

Prskawetz, A. (2005), Will Population Ageing Decrease Productivity?, Vienna Yearbook of Population Research, vol. 3, 2005, pp. 1–3.

Roodman, D., MacKinnon, J., Webb, M. and Nielsen, M. (2018), Fast And Wild: Bootstrap Inference In Stata Using Boottest, No 1406, Working Paper, Economics Department, Queen's University.

Shimer, R. (2001), The Impact of Young Workers on the Aggregate Labour Market, The Quarterly Journal of Economics, 116, pp. 969-1007.

Skans, O.N. (2008), How Does the Age Structure Affect Regional Productivity?, Applied Economics Letters, 15, pp. 787-790.

Skirbekk, V. (2004), Age and Individual Productivity: a Literature Survey, Vienna Yearbook of Population Research, 2 (1), pp. 133–154.

van Ours, J. C. and Stoeldraijer, L. (2010), Age, Wage and Productivity, CESifo Working Paper Series No. 2965.

Zelity, B. (2020), Age Diversity and Aggregate Productivity, No 2020-004, Wesleyan Economics Working Papers, Wesleyan University, Department of Economics.