


RESEARCH PAPER

The age-productivity profile: long-run evidence from Italian regions

Federico Barbiellini Amidei¹, Matteo Gomellini¹, Lorenzo Incoronato² and Paolo Piselli¹ 

¹Bank of Italy, Directorate General for Economics, Statistics and Research, Structural Economic Analysis Directorate – Economic History Division, Rome, Italy and ²Department of Economics and Statistics, CSEF-University of Naples Federico II, Naples, Italy, CESifo, CReAM, RFBerlin

Corresponding author: Paolo Piselli; Email: paolo.piselli@bancaditalia.it

(Received 19 June 2022; revised 11 October 2024; accepted 11 October 2024)

Abstract

This paper investigates the effects of demographic shifts on labor productivity by leveraging variation in the age structure of Italian regions. These effects are analyzed along a first channel – the direct relation between population age and productivity – and a second channel capturing the productivity implications of a more or less dispersed age distribution. We propose an estimation framework that relates regional productivity to the entire age distribution of the working-age population and use instrumental variable techniques to address endogeneity issues. The estimates yield a hump-shaped age-productivity profile peaking between 35 and 40 years. We also document non-linear effects of regional age dispersion on productivity.

Keywords: productivity; demography; age distribution; working-age population; long-run

JEL codes: J11; J21; N30

1. Introduction

European economies have witnessed unprecedented demographic shifts in the past decades. The decline in fertility and mortality rates has led to a progressively older population, with the old-age dependency ratio in Europe rising from 13% in 1955 to 31% in 2018 and set to reach 57% by 2100 (Eurostat, 2020).¹ Within Europe, Italy deserves most attention. As of 2018, Italy had the highest old-age dependency ratio (36%) and the highest share of people aged more than 80 (7%). Between 1950 and 2018, median age in Italy moved from below 29 to 45.5 years, while median age in Europe moved from just above 29 to 41.8 years. Importantly, these trends have been asymmetric across Italian regions: since 1950, median age has risen by about 20 years in the South and 15 years in the Center-North from starting levels of 24 and 30 years, respectively.

In parallel, Italy has experienced a prolonged productivity stagnation (Pellegrino & Zingales, 2017), which has also been heterogeneous across different areas of the

¹The old-age dependency ratio is defined as the share of people aged at least 65 to people aged 15 to 64.

country. Output per worker in the South of Italy is about 20% points lower than in the Center-North (Bugamelli *et al.*, 2018). These regional divides have been widely discussed in the economic literature (Federico *et al.*, 2019; Felice, 2019).

These phenomena are common to other countries and have been analyzed extensively by researchers. Demography has been identified as an important determinant of long-run economic performance across countries or regions (Bloom *et al.*, 2001). An established finding in the literature is a hump-shaped age-productivity profile, with larger productivity associated to cohorts in the middle of the age distribution.

A more neglected aspect is whether the overall shape of the age distribution plays any role in driving aggregate productivity. Productivity will be favored by, say, an increase in the pool of mid-age workers, if the individual-level age-productivity profile is hump-shaped. However, productivity may also depend on how dispersed the age distribution is, on top of the impact of specific age cohorts. There are two contrasting channels through which this effect operates. 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 and benefits productivity. On the other, a diverse workforce may suffer from communication challenges or conflicting values, with negative repercussions on productivity. These considerations suggest potentially non-linear effects of age dispersion (Zélity, 2023).

This paper aims to shed light on these issues and uncover the effects of age on productivity by exploiting the demographic shifts that took place across Italian regions during the second half of the twentieth century. We explore both the contribution of individual age cohorts to aggregate productivity, as well as the impact of a more or less dispersed age distribution.

To estimate the first effect, we relate regional labor productivity to age shares of the working-age population.² Identifying these effects causally is challenging due to potential unobserved confounders and reverse causality issues. We take a number of steps to minimize the bias in the estimated age-productivity profiles. First, we augment our specifications by including region and year effects, along with time-varying controls capturing cohort-specific human capital and the structure of regional economies. Because age-specific migration and mortality patterns may still bias our estimates, we resort to an instrumental variable (IV) strategy leveraging the strongly predetermined component of the population age structure. Namely, we instrument the number of people aged 15, 16, ..., 64 in a given year with the number of people aged 0, 1, ..., 49 fifteen years before (for a similar approach see Skans, 2008). Our estimates point to a hump-shaped age-productivity profile peaking in central age cohorts (35–44) and are robust across specifications. They also hold when using an alternative instrumental variable based on lagged births, constructed following Shimer (2001), with the caveat of a weaker first stage.

We then move to the second question and relate regional productivity to age dispersion, measured as the coefficient of variation of the age distribution of the working-age population. Following the insights in Zélity (2023), we specify a quadratic model to capture potential non-linearities in the effects of dispersion. We address endogeneity concerns by again instrumenting regional age dispersion with its 15-years lag and confirm results using the alternative IV based on lagged births. Our

²We are interested in how demographics affect productivity through labor market channels, hence our focus on the age distribution of the working-age population rather than total population.

estimates suggest that age dispersion benefits productivity. However, in line with Zélicity (2023), the coefficient on the quadratic term is negative and significant, hinting at a detrimental effect of age dispersion on productivity beyond a certain level of dispersion. Using our estimates, we compute the productivity-maximizing value of age dispersion and note that it lies above the levels of dispersion measured in our sample of Italian regions.

These results reflect stark regional heterogeneity. The hump-shaped age-productivity profile and the non-linear effects of age dispersion are driven by regions in Southern Italy, while our estimates are not significant in the North. An explanation we propose is that, as we document descriptively, demographic trends have been much more intense in the South over the last decades. In turn, Southern Italy offers a better setting to precisely estimate the effects of interest as demographic variables (our treatment) feature larger variation there compared to other areas of Italy. We also offer possible complementary explanations for this heterogeneity, related to differential structural change dynamics in the South vs. the North of Italy along the period under analysis. Our data do not allow us to test these economic channels more in depth.

In the last part of the paper, we estimate a “pointwise” age-productivity profile that pins down the coefficient specific to each 1-year age share in the working-age population without ex-ante constraining age cohorts to be grouped together (as commonly done in the literature). To overcome collinearity issues associated with the estimation of many (in our case, fifty) coefficients, we resort to the approach proposed in Fair and Dominguez (1991), which imposes that age coefficients lie on a low-order polynomial. If a second order polynomial is adopted, the number of parameters to be estimated collapses to two, attached to a first- and a second-order moment of the age distribution. The fifty “structural” age parameters are then backed out from the two “reduced-form” ones. This representation channels plenty of information about the population age structure, while allowing for a parsimonious parameterization.

Our preferred (IV) estimates point to a hump-shaped age-productivity profile peaking between 35 and 40 years, in line with the baseline findings and the existing literature. 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. We estimate future demographic trends to lead to productivity losses of about 1.5% per year until 2030, *coeteris paribus*.

Our work relates to the rich literature studying the interplay between demographics and productivity. To the extent that workers of different ages are differently productive, a change in the workforce age structure inevitably affects aggregate productivity (Ilmakunnas et al., 2010; Nagarajan et al., 2017; Prskawetz, 2005).³ The shape of the age-productivity profile, at different levels of aggregation, is the object of lively debate among economists. Theoretically, the individual-level age-productivity profile is expected to follow a hump-shaped pattern, peaking where the cognitive and physical abilities of the youth optimally combine with labor market experience (Skirbekk, 2004).⁴

³Changes in the age structure of the population might affect growth through other channels, which we do not investigate. For example, 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.

⁴Firm-level results are instead more controversial. Some studies confirm the hump-shaped profile observed in individual-level research (e.g., Aubert & Crépon 2003). Other studies find profiles that

At more aggregate levels, several studies in the empirical growth literature have explored the link between demographics and productivity. While results tend to vary somewhat across methodologies and samples, a negative impact of ageing on productivity seems to prevail (Aiyar *et al.*, 2016; Feyrer, 2007; Maestas *et al.*, 2023).⁵ We focus on Italy, which as noted earlier has witnessed dramatic demographic shifts in the last decades making the Italian case an interesting one to study.⁶ Our results, exploiting within-Italy variation, corroborate the existence of a hump-shaped age-productivity profile. Our innovation is also methodological: we estimate pointwise age-productivity profiles as in Fair and Dominguez (1991), explicitly accounting for the endogeneity of the age structure by means of an IV design.

This paper also speaks to the literature studying the economic impact of age diversity and, more broadly, population diversity (Alesina *et al.*, 2016; Docquier *et al.*, 2020; Orefice *et al.*, 2022). Evidence on the effects of workforce age diversity is still limited (OECD, 2020). If there is imperfect substitutability between age groups, age diversity can benefit productivity as the physical strength and the ease in using new technologies of young workers complements the managerial and soft skills of older workers (Guest & Stewart, 2011). In recent work, Zélity (2023) estimates the effects of age diversity on output, highlighting the complementarities between young and old cohorts and noticing the productivity costs of a sub-optimal combination of education and experience. These costs and benefits do indeed produce a non-linear relationship, both theoretically and empirically. These results resonate with those in micro-settings and extend to other outcomes such as innovation (Di Addario & Wu, 2024; Grund & Westergaard-Nielsen, 2008). Motivated by these insights, we estimate a quadratic specification and address endogeneity issues using IV techniques. While our aggregate-level analysis does not allow to identify the underlying microeconomic mechanisms, our new set-up and empirical approach confirm the evidence in Zélity (2023) and document how the shape of the age distribution has important economic consequences.

The paper is structured as follows: section 2 introduces the dataset and presents descriptive statistics; Section 3 shows our results separately for age cohorts and age dispersion; Section 4 presents the polynomial specification and results; Section 5 conducts a simple back-of-the-envelope analysis quantifying how future demographic trends may impact productivity. The last section concludes.

2. Data and descriptive analysis

This section shows a descriptive analysis using data on the age structure of Italian regions over the second half of the twentieth century, collected from the databases of the Italian National Institute of Statistics (Istat). Due to the lack of historical data on the age composition of regional workforces, the analysis will focus on the working-age population. We follow the OECD definition of working-age population

either flatten after the peak at around 40 years (e.g., Göbel & Zwick 2009), or even keep trending upwards (Mahlberg *et al.* 2013). Iparaguirre (2020) provides a survey of the literature.

⁵Some studies show that older cohorts might contribute positively to productivity, see for example Skans (2008). Acemoglu and Restrepo (2017) argue that an ageing-induced decline in the working-age population can stimulate the adoption of labor-replacing technologies, which in turn positively affect growth.

⁶For Italy, Barbiellini Amidei *et al.* (2018) document a positive demographic dividend – the share of economic growth accounted for by demographic shifts (Bloom *et al.* 2001) – until the 1990s, when the dividend turned negative following the surge in the dependency ratio. See also Ciccarelli *et al.* (2019).

and consider the total number of people aged 15 to 64. While workforce data would better suit our study, the age structure of the working-age population – the pool of potential workers available at a given point in time – is a good proxy for the age structure of the workforce. In addition, using population rather than workforce data circumvents possible bias coming from endogenous labor market participation.

The age structure of the Italian population has undergone large shifts over the past decades (Fig. 1). Between 1952 and 2011, mean age rose by 11.5 years for the total population and by 4.5 years for the working-age population. Holding age shares fixed at the 1952 level, this increase in 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 rising share of elderly people, with the old-age dependency ratio growing steadily especially in recent years (0.5% each year between 1990 and 2011).

The intensity of these demographic shifts has not been homogeneous across Italy – something we will notice also later in our empirical analysis. The working-age population in Southern regions was on average younger than in the rest of Italy in 1952 but has aged more rapidly in the ensuing 60 years (Fig. 2, left panel). For instance, the mean age of the working-age population in Basilicata (in the South) rose from 34 years in 1952 to 40 years in 2011, twice as much as in Piedmont (in the North) where the increase was of just 3 years (from 38.6 to 41.6) over the same period. In part, this has come from a more marked decline in the share of younger cohorts in Southern regions over this period.⁷ Despite these shifts, the working-age population in the South remains on average younger than in the rest of the country.

Another dimension of interest is the dispersion of the age distribution, which we capture here with the coefficient of variation of age within the working-age population (Simons et al., 1999). A glance at the data reveals a country-wide reduction in age dispersion over the sample period, as well as clear differences across macro-areas. Despite having become more homogeneous, the working-age population remains more (age) diverse in the South than in the rest of the country (Fig. 2, right panel).

The outcome variable of our analysis is regional labor productivity measured as the natural logarithm of real regional GDP per worker in Euros (chained values, 2000). The data are sourced from the economic research center Prometeia over a period spanning 1971 to 2011. Descriptive analysis shows that productivity correlates positively with mean age and negatively with age dispersion. Regions in the South, which as noted earlier are younger and have a more diverse age structure relative to other areas of the country, are also less productive. These patterns are showed in Fig. 3, which plots the correlation between demographic variables and labor productivity. The graphs show the natural logarithm of the real output per worker across Italian regions scattered against mean age (left) and age dispersion (right) in the working-age population. All variables are averaged between 1971 and 2011.

This preliminary exploration of the data shows that productivity is higher in regions with an older and relatively homogeneous working-age population, which are characteristics of regions in the Center-North. These correlations clearly do not provide evidence about the causal link between demographics and productivity. Unobserved factors can determine both the age structure and labor productivity of a

⁷The share of people aged 15 to 34 decreased on average by 16% points in the South vs. 12% points elsewhere in Italy between 1952 and 2011. This process has partly been driven by migration away from the South after World War II.

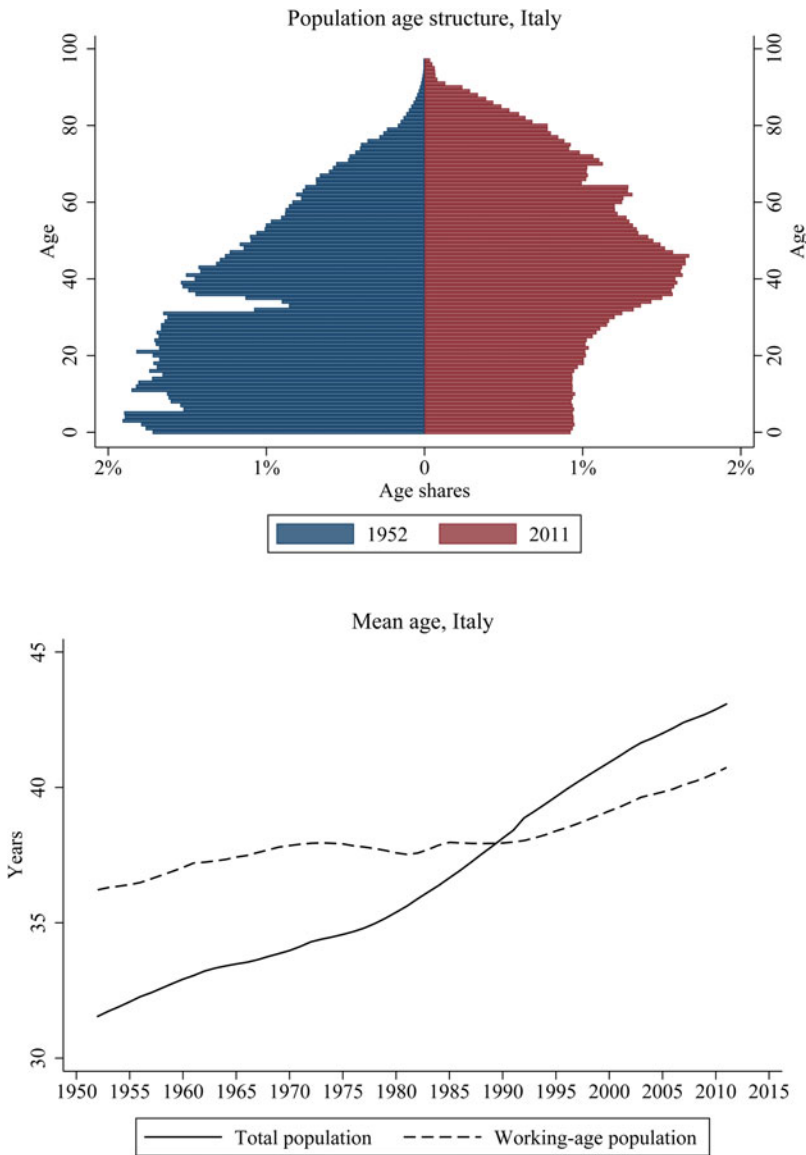


Figure 1. Demographic trends in Italy.
 Note: Authors' elaborations on Istat data.

region and thus generate spurious correlation between them. Even if accounting for all potential confounders, estimates might still suffer from reverse causality bias. An attempt at identifying causal effects is the objective of the next section.

Our regressions will include various controls, such as the sectoral composition of value added, sourced from Prometeia, and the share of young people with a college degree sourced from decennial censuses (<http://ottomilacensus.istat.it>) and linearly

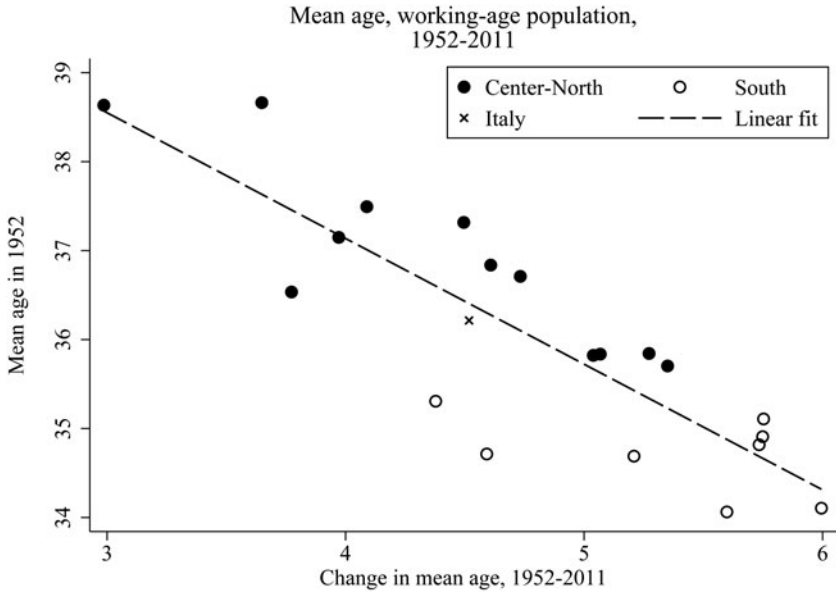


Figure 2. Regional heterogeneity.

Note: Authors' elaborations on Istat data. Italian regions include Piemonte, Valle d'Aosta, Liguria, Lombardia, Trentino Alto Adige, Veneto, Friuli Venezia Giulia, Emilia Romagna, Toscana, Marche, Umbria and Lazio (Center-North); Abruzzo, Basilicata, Calabria, Campania, Molise, Puglia, Sardegna, Sicilia in the Mezzogiorno (South).




Figure 3. Demographics and productivity: descriptive evidence.
 Note: All variables averaged between 1971 and 2011. See Fig. 2 and text for details.

interpolated across years. To construct our alternative instrument based on lagged births, we manually digitized data on births from the historical Yearbook of Demographic Statistics available in the Istat archives.

3. Empirical analysis

3.1. Age shares and productivity

In this section, we investigate the impact of the regional age distribution on labor productivity. Our dataset is a balanced panel of the twenty Italian regions (NUTS-2) between 1981 and 2011. We begin with regressions where the working-age population is broken down into five 10-year sub-groups (15–24 to 55–64), which enter our specification as independent variables, in natural logarithm. The estimated equation is as follows:

$$\ln(y_{it}) = \alpha + \gamma_i + \theta_t + \sum_{j=1}^5 \delta_j \cdot \ln(sh_{it}^j) + X'_{it} \cdot \beta + \varepsilon_{it} \quad (1)$$

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 five age shares sh_{it}^j are the demographic variables of interest.

Notice that, since age shares are expressed in logarithm, there is no need to exclude one to avoid perfect collinearity. This modelling choice allows to interpret coefficients as elasticities denoting 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 results are unchanged when we do not take the natural logarithm of age shares and exclude one share. We also caution that our strategy does not overcome the classic age-cohort-year identification problem (Bell & Jones, 2013) and, for instance, cannot 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 (1). The estimation period is 1981–2011 and standard errors are clustered regionally to allow for arbitrary correlations between observations within a region. Column (1) shows the estimated coefficients of a simple pooled OLS regression of labor productivity exclusively on the (log) age shares of the working-age population. A hump shaped pattern emerges, with a peak in the 25–34 cohort. In Columns (2) and (3) we test richer specifications including region and year fixed effects and time-varying controls. The within-country design significantly narrows down the list of possible confounders as many of them, such as the legal and (formal) institutional setting, do not vary within national borders. In addition, regional dummies account for region-specific unobserved traits linked, for example, to informal institutions. However, the substantial degree of time-varying heterogeneity across Italian regions justifies the inclusion of additional controls to corroborate the results. We control for the value-added sectoral composition (Maestas et al., 2023) and, as anticipated above, for the share of people aged 30 to 34 with a college degree to address concerns about cohort-specific education trends driving our results.⁸ Across the various specifications, the age-productivity profile peaks in the central age cohorts. According to the estimates in the richest specification in Column (3), a one percent

⁸Results are unchanged when using average years of schooling as a control for education. This variable is obtained by updating the reconstruction performed in Bronzini and Piselli (2009).

Table 1. Age effects on labor productivity

	Pooled OLS	Fixed effects		2SLS	
	(1)	(2)	(3)	(4)	(5)
Share of 15–24	0.311 (0.288)	0.444 (0.298)	0.581** (0.263)	0.398 (0.332)	0.548* (0.284)
Share of 25–34	1.485*** (0.297)	0.812** (0.330)	0.913** (0.334)	0.877* (0.471)	1.199*** (0.368)
Share of 35–44	1.399*** (0.373)	0.827* (0.442)	0.912** (0.387)	1.048** (0.467)	1.246*** (0.349)
Share of 45–54	0.852** (0.309)	0.432* (0.215)	0.341* (0.195)	0.555*** (0.215)	0.518*** (0.166)
Share of 55–64	0.659** (0.288)	0.281 (0.203)	0.341 (0.204)	0.272 (0.214)	0.338 (0.211)
Value added, agric. share			0.894 (0.781)		0.891 (0.777)
Value added, industry share			0.304 (0.357)		0.188 (0.360)
Share of young with higher education			0.215 (0.127)		0.239** (0.105)
Kleibergen-Paap <i>F</i> statistic	–	–	–	13.647	30.545
Region effects	No	Yes	Yes	Yes	Yes
Year effects	No	Yes	Yes	Yes	Yes
Observations	620	620	620	620	620
<i>R</i> ²	0.681	0.941	0.949	0.940	0.947

Notes: Estimates for equation (1) (except for Column (1), which excludes fixed effects). 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 share of young with higher education is the share of those aged 30–34 with a college degree. Columns (4)–(5) instrument age shares with their 15-years lag. Standard errors clustered at the regional level in parentheses. See text for details. **p* < 0.1, ***p* < 0.05, ****p* < 0.01.

inflow into the 35–44 years cohort from any of the other cohorts results in a 0.9% productivity rise.

As mentioned earlier, these estimates still fail to account for simultaneity. We control for time-invariant unobserved heterogeneity through region identifiers and eliminate spurious correlations with the business cycle using year effects. Still, reverse causality may bias the estimates, which could then pick up not only the impact of the age structure on productivity but also any effect moving in the opposite direction. One such effect could be driven by, say, migration (either domestic or foreign, regular or irregular) of younger workers toward more economically vibrant regions, which

would shift the age structure of both the receiving and the sending region and in turn bias the estimated coefficients.⁹

Causal identification requires isolating exogenous variation in the age structure of region i in year t . We seek such variation by means of an IV strategy that instruments 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 larger or smaller productivity.

Column (4) in Table 1 reports the Two-Stage Least Squares (2SLS) coefficient estimates when age shares 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, as evidenced by the large first-stage statistic. Relative to the 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).

We perform three additional checks. First, we test our results using a second instrument based on Shimer (2001) and constructed using lagged births (see also Lugauer, 2012). The share 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 share of those aged 16 in 1981 is instrumented with the number of those born 16 years before in 1965. And so on for the whole 15–64 cohort. While the instrument's first stage is not very high (F -statistic of 5.7), we still estimate a hump-shaped distribution as for the baseline IV results (Appendix Table A.1). Second, our inference might suffer from the small number of clusters (20). We address this concern by using the wild bootstrap procedure devised in Roodman et al. (2019). As showed in Appendix Table A.2, the estimates are less precise, but still overall significant in the central age cohorts. Third, we notice that a drawback of using population age shares as a proxy for workforce age shares (a choice determined by data availability), is that differential participation rate by age group and region may affect our results without being captured by population data. Using data on labor force participation by age group between 1993 and 2020, we find a high correlation (0.95) between population and labor force age shares, which mitigates these concerns somewhat.¹⁰

As anticipated, we also test a different specification where age shares (not their natural logarithms) enter the right-hand side. As these shares sum to one, we

⁹Note that a permanently high/low level of regional labor 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.

¹⁰Table A.3 shows that controlling for regional unemployment (source: Prometeia) does not alter the estimates.

exclude one of them to avoid perfect collinearity. We omit the central cohort (35–44) so that the coefficient attached to one of the included shares denotes the productivity impact of a population flow into that 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. The model is:

$$\ln(y_{it}) = \alpha + \gamma_i + \theta_t + \sum_{j=1}^4 \delta_j \cdot sh_{it}^j + X_{it}' \cdot \beta + \varepsilon_{it} \quad (2)$$

Table 2 reports 2SLS coefficient estimates of equation (2) under our preferred specification including the vector of controls, region and year effects (Table 1 Column (5)). A clear hump-shaped profile can still be detected, as outflows from the omitted cohort (35–44) into the remaining four lead to lower productivity.

3.2. Age dispersion and productivity

We now explore the role of age dispersion, measured as the coefficient of variation of the working-age population. We estimate the following equation:

$$\ln(y_{it}) = \alpha + \gamma_i + \theta_t + \varphi_1 CV_{it} + \varphi_2 CV_{it}^2 + X_{it}' \cdot \beta + \varepsilon_{it} \quad (3)$$

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 above; φ_1, φ_2 are the coefficients of interest, picking up any (potentially non-linear) effect of age dispersion – denoted by CV_{it} .

Table 3 reports fixed effects OLS (Columns (1)–(2)) and 2SLS (Columns (3)–(4)) estimates for equation (3). The IV estimates instrument the coefficient of variation with its 15-year lag, that is, the coefficient of variation of the regional age structure 15 years earlier. Columns (2) and (4) include the time-varying controls described earlier. In all specifications, we estimate a positive effect of age dispersion on regional productivity. However, our estimates imply that this positive impact exists up to a certain degree, beyond which dispersion becomes detrimental for productivity.¹¹ This result is consistent with the theoretical and empirical findings in Zélity (2023) and described in the introduction of this paper: if education and experience are complementary, age diversity may be beneficial. However, because only young cohorts have up-to-date education, too much age diversity might eventually harm productivity. Using the coefficients in Column (4), we estimate that the quadratic function peaks at a coefficient of variation equivalent to 4.56. We measure an average coefficient of variation of 4.31 for Italy in our data, suggesting a still positive contribution of age dispersion to productivity (Appendix Fig. A.1).

3.3. Regional differences

We now explore the regional heterogeneity of our results. The North-South economic divide is a long-lasting structural feature of Italian economy and its roots have been

¹¹Appendix Table A.4 confirms our estimates when age dispersion is instrumented using lagged births. Appendix Table A.5 shows results using wild bootstrap inference.

Table 2. Age effects on labor productivity, non-logarithmic shares

	(2SLS)
Share of 15–24	–3.884** (1.720)
Share of 25–34	–1.274 (0.883)
Share of 45–54	–2.245** (0.914)
Share of 55–64	–4.821*** (1.835)
Value added, agric. share	1.331 (1.011)
Value added, industry share	0.127 (0.505)
Share of young with higher education	0.170 (0.146)
Kleibergen-Paap <i>F</i> statistic	16.831
Region effects	Yes
Year effects	Yes
Observations	620
R^2	0.927

Notes: 2SLS coefficient estimates from equation (2). See text and Table 1 for details. Standard errors clustered at the regional level in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

investigated by many economists and historians. Here we just recall that after the brief parenthesis of the 1950s and the 1960s, characterized by “general convergence,” the South has fallen behind the rest of the country in the ensuing decades (Felice, 2019).

When running our analysis separately for the South and the Center-North (Table 4) we notice that the productivity effects of age shares are limited to the South. For the Center-North, the coefficients are not significant. A similar conclusion holds for the effects of age dispersion on labor productivity (Table A.6 in the Appendix). How can this be explained? A first possible cause is merely statistical. As described in section 2, there are substantial differences in the intensity of demographic shifts across different areas of Italy. In turn, larger variation in the independent variable allows us to better estimate the impact of age on productivity in the South, compared to the Center-North where demographic trends have been relatively more muted.

Another possible explanation stems from structural changes in the Italian economy in the period under analysis, when the gap between the two parts of Italy crystalized. While the South remained specialized in agriculture and traditional productions, the North completed structural transformation toward modern manufacturing. While our models control for broad sectoral shares (agriculture and industry), specialization

Table 3. Age dispersion effects on labor productivity

	Fixed effects		2SLS	
	(1)	(2)	(3)	(4)
Coeff. var. of age	1.689*	1.738**	1.508	1.577**
	(0.973)	(0.828)	(0.943)	(0.794)
(Coeff. var. of age) ²	-0.195*	-0.192*	-0.172*	-0.173*
	(0.109)	(0.094)	(0.104)	(0.090)
Value added, agric. share		0.575		0.631
		(0.991)		(0.938)
Value added, industry share		0.404		0.410
		(0.343)		(0.329)
Share of young with higher education		0.122		0.118
		(0.137)		(0.127)
Kleibergen-Paap <i>F</i> statistic	-	-	155.317	226.530
Region effects	Yes	Yes	Yes	Yes
Year effects	Yes	Yes	Yes	Yes
Observations	620	620	620	620
<i>R</i> ²	0.935	0.940	0.935	0.940

Notes: Estimates from equation (3). The dependent variable is the natural logarithm of real GDP per worker. The coefficient of variation of age is computed as the ratio between the standard deviation and the mean of age in the working-age population. Columns (3)–(4) instrument the coefficient of variation using its 15-years lag. Standard errors clustered at the regional level in parentheses. See text and Table 1 for details. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

within more or less skilled manufacturing might still drive some of our results. More generally, the marked technological transformation in the North and its effects on productivity may have obfuscated the effects of changing demographic structure, which are instead more easily identified (and estimated) for the South.

4. Polynomial approach

The empirical strategy of equation (1) implicitly assumes that the age effects 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 subgroups, without theoretical or empirical justifications. We now allow each 1-year cohort to independently affect regional productivity and obtain “pointwise” 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 , $j = 15, \dots, 64$. In the most general formulation, we also include a matrix of controls X_{it} , region and year effects:

$$\ln(y_{it}) = \alpha + \gamma_i + \theta_t + \sum_{j=15}^{64} \delta_j \cdot sh_{it}^j + X'_{it} \cdot \beta + \varepsilon_{it} \quad (4)$$

Table 4. Age effects on labor productivity: Center-North vs. South

	Fixed effects		2SLS	
	North	South	North	South
Share of 15–24	–0.452 (0.508)	1.161*** (0.212)	–0.600 (0.596)	1.113*** (0.197)
Share of 25–34	–0.597 (0.470)	1.263*** (0.318)	–1.050 (1.015)	1.566*** (0.348)
Share of 35–44	–0.916 (0.756)	1.941*** (0.382)	–1.061 (1.046)	2.028*** (0.353)
Share of 45–54	–0.518 (0.394)	0.972*** (0.266)	–0.440 (0.329)	1.071*** (0.227)
Share of 55–64	–0.710 (0.436)	0.846*** (0.162)	–0.752* (0.404)	0.835*** (0.146)
Value added, agric. share	–1.289 (1.397)	0.807 (0.452)	–2.115 (2.369)	0.935** (0.470)
Value added, industry share	0.764* (0.390)	0.442*** (0.116)	0.678** (0.338)	0.408*** (0.124)
Share of young with higher education	0.120 (0.126)	0.359*** (0.091)	0.025 (0.143)	0.332*** (0.079)
Kleibergen-Paap <i>F</i> statistic	–	–	1.960	13.140
Region effects	Yes	Yes	Yes	Yes
Year effects	Yes	Yes	Yes	Yes
Observations	372	248	372	248
<i>R</i> ²	0.950	0.985	0.947	0.985

Notes: Estimates for equation (1) separately for Center-North and South. The dependent variable is the natural logarithm of real GDP per worker. Standard errors clustered at the regional level in parentheses. See Table 1 and text for details. **p* < 0.1, ***p* < 0.05, ****p* < 0.01.

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 – complicates such a granular estimation of the age profiles. The finer the division between consecutive age shares, the more serious estimation issues become (Juselius & Takáts, 2018). Multicollinearity between regressors would severely affect coefficient estimates, whose precision would also decrease as the number of age cohorts rises relative to the number of periods. Moreover, leaving regression parameters unconstrained may lead to confusing age profiles, with estimated age effects possibly varying dramatically between adjacent 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-order polynomial to smooth out the estimated effect between consecutive cohorts.

$$\delta_j = \eta_0 + j \cdot \eta_1 + j^2 \cdot \eta_2, \quad j = 15, \dots, 64 \tag{5}$$

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:

$$\sum_{j=15}^{64} \delta_j = 0 \tag{6}$$

Combining (5) and (6) we express η_0 as a function of η_1 and η_2 :

$$\eta_0 = -\frac{1}{50} \cdot \left(\eta_1 \cdot \sum_{j=15}^{64} j + \eta_2 \cdot \sum_{j=15}^{64} j^2 \right) \tag{7}$$

Plugging (5) and (7) into (4) leads to the following expression:

$$\begin{aligned} \ln(y_{it}) = & \alpha + \gamma_i + \theta_t + \eta_1 \cdot \underbrace{\left(\sum_{j=15}^{64} j \cdot sh_{it}^j - \frac{1}{50} \cdot \sum_{j=15}^{64} j \right)}_{M_{it}^1} + \eta_2 \\ & \cdot \underbrace{\left(\sum_{j=15}^{64} j^2 \cdot sh_{it}^j - \frac{1}{50} \cdot \sum_{j=15}^{64} j^2 \right)}_{M_{it}^2} + X'_{it} \cdot \beta + \varepsilon_{it} \end{aligned} \tag{8}$$

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 moment of the age distribution (M_{it}^1, M_{it}^2).¹² The last reduced-form coefficient η_0 can be easily pinned down from the estimates of η_1 and η_2 using (7) and in turn the values of the structural age parameters $\delta_j, j = 15, \dots, 64$ are backed out from equation (5).¹³

Column (1) in Table 5 reports OLS coefficient estimates for equation (8). The age-specific effects on productivity resulting from the above estimates are depicted in the left panel of Fig. 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

¹²Given the attractiveness of tracing out pointwise age profiles while preserving a parsimonious model parameterization, several studies have adopted this methodology. See for instance Higgins (1998), Skans (2008), Juselius and Takáts (2018).

¹³The standard errors of the structural age coefficients are computed in a similar way. Plugging (7) in (5):

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

Table 5. Polynomial specification, coefficient estimates

	OLS (1)	2SLS (2)	2SLS – Lagged Births (3)
M_{it}^1	0.309 (0.240)	0.612** (0.285)	2.707** (1.369)
M_{it}^2	-0.004 (0.004)	-0.008** (0.004)	-0.035* (0.018)
Kleibergen-Paap <i>F</i> statistic		50.112	1.915
Region effects	Yes	Yes	Yes
Year effects	Yes	Yes	Yes
Additional controls	Yes	Yes	Yes
Observations	620	620	620
R^2	0.934	0.936	0.783

Notes: Estimates from equation (8). All are fixed-effect specifications and include controls for value added sectoral composition and the share of young people with college degree. Column (2) instruments age shares using 15-year lags; Column (3) instruments age shares using lagged births. See text for details. Standard errors clustered at the regional level are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

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

As previously noted, however, reverse causality concerns make demographic variables highly likely to remain endogenous even after including fixed effects and other controls. To address this issue, we augment the Fair and Dominguez (1991) polynomial specification by again resorting to an IV approach. We 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:

$$M_{it}^{k,IV} = \sum_{j=0}^{49} j^k \cdot sh_{it}^{j,IV} - \frac{1}{50} \cdot \sum_{j=0}^{49} j^k, \quad k = 1, 2 \tag{9}$$

where $sh_{it}^{j,IV} = pop_{it-15}^j / pop_{it-15}^{0-49}$ 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 5 reports coefficient estimates for equation (8) when estimated through a 2SLS procedure using $M_{it}^{k,IV}$ as instrument for M_{it}^k .¹⁴ The right panel of Fig. 4 shows the resulting age-productivity profile. The estimates again point to a hump-shaped relationship, with a peak between 35 and 40 years old.

¹⁴Column (3) reports the 2SLS results when instrumenting age groups using lagged births (Shimer 2001), which confirms our results although with a weak first stage.

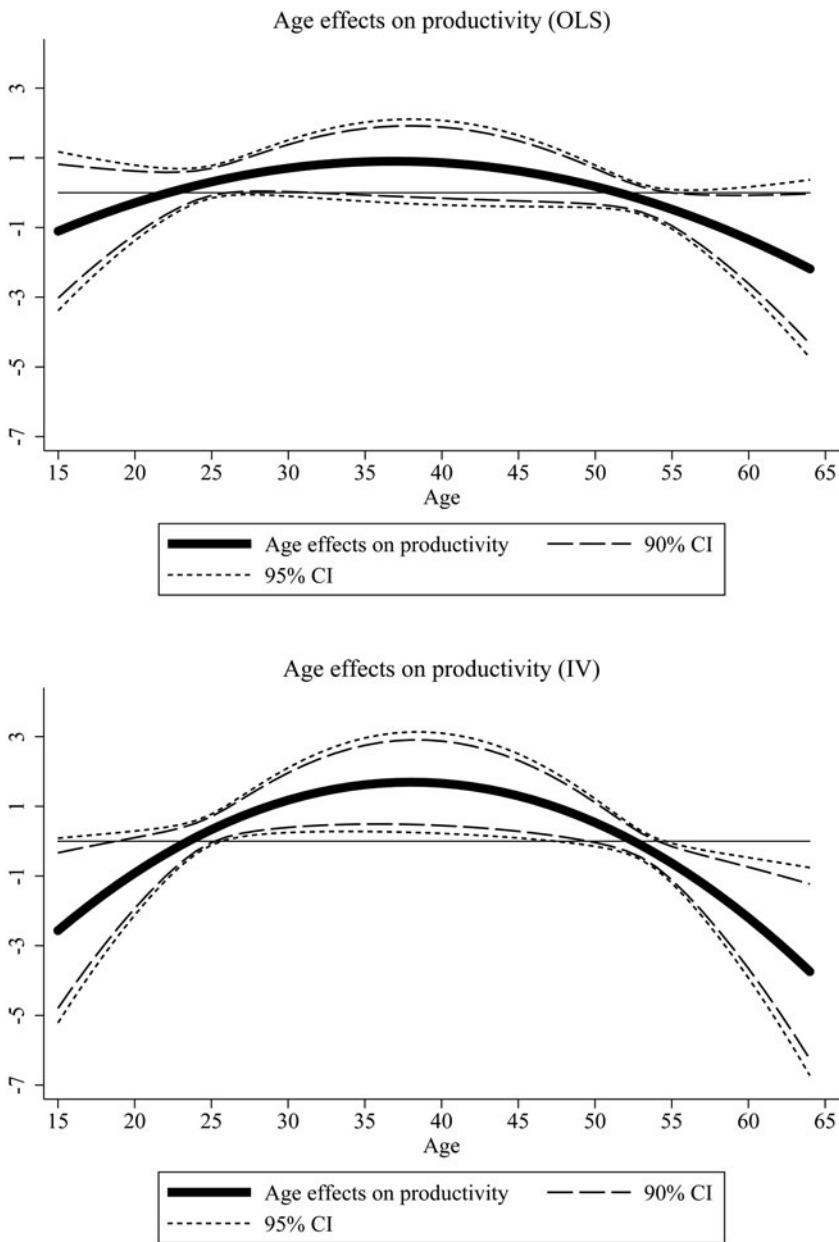


Figure 4. Age-productivity profiles for Italy, 1981–2011.

Notes: Estimates of structural age coefficients (the δ_j) from equation (4). These result from the estimation of equation (8) with fixed-effects OLS (left) and 2SLS (right). 2SLS estimates are obtained using 15-year lagged shares as instruments. Standard errors are clustered at the regional level. See text for details.

Incidentally, this approach allows to explore the effect of dispersion, too. 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 sample. The coefficient attached to M_{it}^2 thus conveys similar information as that associated with the coefficient of variation in Table 3, that is, the productivity implications of having a more or less dispersed age distribution. While the OLS estimates in Table 5 point to insignificant second-order effects of age on productivity, the IV approach restores their significance and confirms the positive relationship between age dispersion and labor productivity observed earlier. We show two extensions in the Appendix. Appendix Fig. A.2 shows the output of the polynomial IV estimation (the right panel in Fig. 4) separately for the Center-North and South of Italy. In line with the estimates showed in section 3, we find flat age-productivity profiles for Northern regions, and significant results only in the South. Appendix Fig. A.3 experiments with a cubic (rather than quadratic) polynomial, confirming a hump-shaped pattern.¹⁵

5. Future impact: back-of-the-envelope calculations

The main takeaway from this analysis is that the age-productivity relation is not monotonic. To gauge the economic consequences of demographic trends, one thus 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. Having estimated age-specific coefficients, we can quantify how past and future changes in the age structure of Italy's working-age population have affected, or are likely to affect, productivity. We thus perform a simple back-of-the-envelope exercise by first-differencing equation (8) and plugging our estimates for the δ_j 's, along with (i) changes in age distribution between 2000 and 2019 and (ii) projected changes between 2019 and 2030.¹⁶

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 Fig. 5, the age groups that grew (decreased) relatively to others between 2000 and 2019 are those providing a negative (positive) contribution to productivity, according to our estimates. These demographic shifts have thus likely been associated with a significant drop in productivity. We estimate such drop at around -0.7% per year on average (-12.5% over the whole period), *coeteris paribus*.¹⁷ This is a rather large impact, especially when compared with the 0.9% effective increase in GDP per worker over the same period (Bugamelli et al., 2018).

We observe negative demographic effects also for the near future (2019–2030). The age structure of the Italian working-age-population is expected to shift toward 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 labor productivity as

¹⁵In other words, we assume that Equation (5) reads:

$$\delta_j = \eta_0 + j \cdot \eta_1 + j^2 \cdot \eta_2 + j^3 \cdot \eta_3, \text{ for } j = 15, \dots, 64$$

¹⁶The 2019 age distribution is the actual distribution as of January 1 (see Istat, "Demografia in cifre," www.demo.istat.it). For the projected distribution we use Istat population forecasts, "Previsioni della Popolazione 2018–2065;" see <http://dati.istat.it>.

¹⁷In this calculation, coefficients not statistically different from zero (at the 95% confidence level) are set at zero. Our results are unchanged when including the whole set of coefficients, irrespective of their significance.




Figure 5. Demographic shifts and effects on productivity.

Notes: The left axis shows the percentage change in the working-age population share of each one-year age group between 2000 and 2019 (grey bars) and 2019 and 2030 (white bars) (source Istat). The right axis shows the estimates of structural age coefficients in equation (4), resulting from 2SLS estimation of equation (8) – the bottom panel chart in Fig. 4.

large as -1.5% per year by 2030 due to these projected shifts and abstracting from movements in the other control variables. Interestingly, longer-term projections paint a less pessimistic picture. Population statistics foresee a relative increase in the cohorts aged 30–40 by 2040 and 2050, partly compensating for the expected rise in older cohorts. As a result, the predicted productivity effects of demographic shifts hover around -0.4% per year until 2040 and -0.1% until 2050 (cumulatively -7.4 and -3.1% , respectively). We place less emphasis on these projections in light of their distance in the future, which implies higher uncertainty about other drivers of productivity different from population age.

6. Conclusions

This paper produces novel evidence about the relationship between population age and labor productivity. The empirical analysis exploits variation in the age structure of the working-age population across Italian regions between 1981 and 2011. The economic transformations witnessed by Italy over the past decades and the notable heterogeneity in how regional age structures have evolved offer an advantageous perspective to explore these phenomena.

Our instrumental variable estimates are in line with the existing literature and point to a hump-shaped age-productivity profile, 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% per year until 2030, abstracting from changes in other variables. We

also allow the dispersion of the age distribution to impact labor productivity, a channel which has received little emphasis in the literature. We estimate a quadratic relationship between age dispersion and productivity. Our estimates imply a positive relationship up to a certain level of dispersion, beyond which age dispersion becomes detrimental for productivity.

Further investigations are warranted to shed light on our results. Importantly, more theoretical work supporting our findings and micro-empirical research on possible channels through which demography affects productivity – such as innovation and entrepreneurship – are much needed.

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.1017/dem.2024.24>.

Acknowledgements. We thank Fabrizio Balassone, Marta De Philippis, Sabrina Di Addario, Dario Pellegrino, Paolo Sestito and seminar participants at the Bank of Italy for their suggestions. The views expressed in the article are those of the authors and may not reflect those of the Bank of Italy.

Competing interests. None.

References

- Acemoglu, D., & Restrepo, P. (2017). Secular stagnation? The effect of aging on economic growth in the age of automation. *American Economic Review*, 107(5), 174–179.
- Aiyar, S., Ebeke, C., & Shao, X. (2016). The Impact of Workforce Aging on European Productivity. *IMF Working Paper* 16/238.
- Alesina, A., Harnoss, J., & Rapoport, H. (2016). Birthplace diversity and economic prosperity. *Journal of Economic Growth*, 21(2), 101–138.
- Aubert, P., & Crépon, B. (2003). La productivité des salariés âgés: une tentative d'estimation. *Économie et Statistique*, 368(1), 95–119.
- Barbiellini Amidei, F., Gomellini, M., & Piselli, P. (2018). Il Contributo della Demografia alla Crescita Economica: Duecento Anni di Storia Italiana. *Bank of Italy Occasional Papers*, n. 431.
- Bell, A., & Jones, K. (2013). The impossibility of separating age, period and cohort effects. *Social Science & Medicine*, 93, 163–165.
- Bloom, D. E., Canning, D., & Sevilla, J. (2001). Economic Growth and the Demographic Transition. *NBER Working Papers*, N. 8685.
- Bönte, W., Falck, O., & Heblich, S. (2009). The impact of regional age structure on entrepreneurship. *Economic Geography*, 85(3), 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, 187–199.
- Bugamelli, M., Lotti, F., Ciapanna, E., D'Amuri, F., Linares, A., Manaresi, F., Palumbo, G., & Sette, E. (2018). La crescita della produttività in Italia: la storia di un cambiamento al rallentatore. *Bank of Italy Occasional Papers*, n. 422.
- Ciccarelli, C., Gomellini, M., & Sestito, P. (2019). Demography and Productivity in the Italian Manufacturing Industry: Yesterday and Today. *CEIS Research Paper*, 457, Tor Vergata University.
- Di Addario, S., & Wu, A. (2024). Employer learning and sorting in the labor market of inventors, mimeo. Bank of Italy.
- Docquier, F., Turati, R., Valette, J., & Vasilakis, C. (2020). Birthplace diversity and economic growth: Evidence from the US states in the Post-World War II period. *Journal of Economic Geography*, 20(2), 321–354.
- Eurostat (2020). *Population structure and ageing*. Eurostat.
- Fair, R., & Dominguez, K. (1991). Effects of the changing US age distribution on macroeconomic equations. *American Economic Review*, 81(5), 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), 63–98.
- Felice, E. (2019). The roots of a dual equilibrium: GDP, productivity and structural change in the Italian regions in the long-run (1871–2011). *European Review of Economic History*, 23(4), 499–528.
- Feyrer, J. (2007). Demography and productivity. *Review of Economics and Statistics*, 89(1), 100–109.
- Göbel, C., & Zwick, T. (2009). Age and Productivity – Evidence From Linked Employer Employee Data. *ZEW Discussion Paper*, 09-020, Centre for European Economic Research, Mannheim.
- Grund, C., & Westergaard-Nielsen, N. (2008). Age structure of the workforce and firm performance. *International Journal of Manpower*, 29(5), 410–422.
- Guest, R., & Stewart, H. (2011). The age dispersion of workers and firm productivity: A survey approach. *Australian Journal of Labor Economics*, 14(1), 59–75.
- Higgins, M. (1998). Demography, national savings, and international capital flows. *International Economic Review*, 39(2), 343–369.
- Ilmakunnas, P., van Ours, J., Skirbekk, V., & Weiss, M. M. (2010). Age and productivity. In P. Garibaldi, J. O. Martins, & J. van Ours (Eds.), *Ageing, health, and productivity: The economics of increased life expectancy* (pp. 133–240). Oxford University Press.
- Iparaguire, J. L. (2020). *Economics and ageing*. Springer Books.
- Juselius, M., & Takáts, E. (2018). The Enduring Link between Demography and Inflation. *Bank of Finland Research Discussion Paper No. 8/2018*.
- Lugauer, S. (2012). Estimating the effect of the age distribution on cyclical output volatility across the United States. *Review of Economics and Statistics*, 94(4), 896–902.
- Maestas, N., Mullen, K. J., & Powell, D. (2023). The effect of population aging on economic growth, the labor force and productivity. *American Economic Journal: Macroeconomics*, 15(2), 306–332.
- Mahlberg, B., Freund, I., & Fürnkranz-Prskawetz, A. (2013). Ageing, productivity and wages in Austria: Sector level evidence. *Empirica*, 40(4), 561–584.
- Nagarajan, R., Teixeira, A., & Silva, S. (2017). The impact of population ageing on economic growth: A bibliometric survey. *The Singapore Economic Review*, 62(2), 275–296.
- OECD (2020). *Promoting an age-inclusive workforce: Living, learning and earning longer*. OECD Publishing.
- Orefice, G., Rapoport, H., & Santoni, G. (2022). How Do Immigrants Promote Exports? Networks, Knowledge, Diversity. *IZA Discussion Paper No. 15722*.
- Pellegrino, B., & Zingales, L. (2017). Diagnosing the Italian disease. *NBER Working Paper no. 23964*.
- Prskawetz, A. (2005). Will population ageing decrease productivity? Summary of the debate. *Vienna Yearbook of Population Research*, 3(1), 1–3.
- Roodman, D., MacKinnon, J., Nielsen, M. O., & Webb, M. D. (2019). Fast and wild: Bootstrap inference in stata using boottest. *The Stata Journal*, 19(1), 4–60.
- Shimer, R. (2001). The impact of young workers on the aggregate labor market. *The Quarterly Journal of Economics*, 116(3), 969–1007.
- Simons, T., Pelled, L. H., & Smith, K. A. (1999). Making use of difference: Diversity, debate, and decision comprehensiveness in top management teams. *The Academy of Management Journal*, 42(6), 662–673.
- Skans, O. N. (2008). How does the age structure affect regional productivity? *Applied Economics Letters*, 15(10), 787–790.
- Skirbekk, V. (2004). Age and individual productivity: A literature survey. *Vienna Yearbook of Population Research*, 2(1), 133–154.
- Zélyty, B. (2023). Age diversity and aggregate productivity. *Journal of Population Economics*, 36(3), 1863–1899.