

Aggregate Evidence on the Link Between Demographics and Productivity*

James Feyrer[†]
Dartmouth College

This Draft: September 30, 2005

Abstract

This paper examines the relationship between workforce demographics and aggregate productivity. Cross country regressions show changes in the age structure of the workforce to be significantly correlated with changes in aggregate productivity. Different demographic structures may be related to almost one quarter of the persistent productivity gap between the OECD and low income nations. The magnitude of these correlations is significantly larger than one would expect from the private return to experience, suggesting externalities. Causal mechanisms through inventive activity and management are suggested and evidence from the US census is shown to be consistent with these mechanisms. Results using US state and metropolitan area data suggest that the scope of the externality is largely national, not regional

Keywords: Productivity, human capital, demographics.
JEL Classification: E23,O30,O47

Introduction

This paper will examine the relationship between workforce demographics and aggregate productivity. It is well known from the labor literature that there is a robust relationship between years of experience and income. If workers are paid their marginal product than this suggests a relationship between worker productivity and age. In the aggregate, therefore, we should expect to see changes in the age structure of a population to be correlated with changes in productivity.

*I am grateful to Peter Klenow, Doug Staiger, Bruce Sacerdote, and the participants in the Symposium on Population Aging and Economic Productivity at the Vienna Institute of Demography for their helpful comments and advice. All errors are my own.

[†]james.feyrer@dartmouth.edu, Dartmouth College, Department of Economics, 6106 Rockefeller, Hanover, NH 03755-3514. fax:(603)646-2122.

Feyrer (2004) finds that there is a strong and robust correlation between workforce demographics and total factor productivity. The impact of demographics on productivity is much larger than is estimated by microeconomic evidence on the relationship between wages and experience. The magnitude of the result does not appear to be driven by obvious reverse causality from productivity to work force demographics through immigration or participation rates. This suggests that the social return to a workforce with a particular experience profile is higher than the private return to experience. This paper is interested in examining the nature of this externality in several ways.

First, I will explore changes in the age distributions of the US workforce for several subcategories of worker. One mechanism through which the age distribution could produce large externalities is through innovative activity. I will examine the evolution of the age distribution of patent holders in the US over time. Another mechanism through which the age distribution of the workforce may be important is through changes in management. Lucas (1978) suggests that the quality distribution of managers may play a large role in determining output. Using census data, I will explore how the entry of the baby boom into the workforce changed the composition of the managerial workforce over time.

Second, it may be useful to identify the scope of the externality. The initial results show that demography has an impact on country level productivity. Does this extend to smaller geographic units like US states and localities? This is an important question because different scopes may suggest different mechanisms at work. Does a large proportion of prime age workers in a city raise productivity in that city alone or do effects spill over to the state and national level?

The paper will proceed as follows. Section 1 will examine the relationship between demographics and productivity at the country level. Section 2 will discuss some of the implications of the cross country results for cross country economic performance and look at the relative performance of the US and Japan in light of the results. Section 3 will suggest channels through which demographics may be affecting output and review some evidence from the US census. Section 4 will examine US state and metropolitan area data

to see if the cross country effects are evident at lower levels of aggregation.

1 Cross Country Evidence

For use in a cross country regression, demographic measures have several useful characteristics which make identification more straightforward than with many variables typically used in this literature. First, demographic measures are strongly predetermined. The current age structure of the workforce was determined roughly twenty years ago and should be predetermined with respect to current output movements. Second, demographics have significant time series variation. This time series variation allows for exploiting the panel nature of the data.

The following results largely follow Feyrer (2004) and focus on total factor productivity. The results would not be substantively different if the dependent variable were to be changed to per worker output. Total factor productivity in country i at time t , $y_{i,t}$, is assumed to be a function of a time invariant country fixed effect, f_i , a time trend common to all countries, μ_t , and a vector of explanatory variables $x_{i,t}$,

$$y_{i,t} = f_i + \mu_t + \beta x_{i,t} + u_{i,t}. \quad (1)$$

The regressors are the proportion of the workforce by age group, with W10 indicating workers between 10 and 19, W20 workers between 20 and 29, etc. W60 indicates workers older than 60 years of age. Since these variables are proportions, the sum of all the age groups is one for each country year pair. For this reason, one group is excluded.¹ For all the reported regressions in this section, first differencing was used to eliminate the country specific effect.²

Productivity is calculated as a Solow residual. I assume a Cobb-Douglas production

¹I choose to exclude W40 because the forty year old age group generally has the highest coefficient when included. By excluding W40, significant coefficients on the other age groups indicate that they are significantly different from the implied zero coefficient on W40.

²Standard errors were clustered at the country level to deal with serial correlation. The results do not change if estimated in levels.

function taking physical capital, human capital from schooling, and productivity as inputs.

$$y_{i,t} = k_{i,t}^\alpha (A_{i,t} h_{i,t})^{1-\alpha} \quad (2)$$

where $y_{i,t}$ is output, $k_{i,t}$ is capital per worker, $h_{i,t}$ is human capital per worker, and $A_{i,t}$ represents productivity. Capital's share of output, α is assumed to be $1/3$.³ The human capital production function is assumed to have a Mincer form

$$h_{i,t} = e^{\phi(s_{i,t})} \quad (3)$$

where $s_{i,t}$ is the average years of schooling in country i at time t and $\phi(s)$ is an increasing function that is assumed to be piecewise linear with decreasing returns to scale. The coefficients are taken from Psacharopoulos (1994), which surveys the literature on returns to schooling.⁴ The production function can be solved for log total factor productivity.

$$\log(A_{i,t}) = \log(y_{i,t}) - \frac{\alpha}{1-\alpha} \log\left(\frac{K}{Y}\right)_{i,t} - \log(h_{i,t}). \quad (4)$$

Data for output are from the Penn World Tables version 6.0. Following Hall and Jones (1999), output data are adjusted to exclude income from mining and oil.⁵ Data for capital per worker are from Easterly and Levine (2001).⁶ The schooling data used to calculate human capital stocks are from Barro and Lee.

The data on workforce composition are from two sources. The International Labor Organization (ILO) has compiled cross country data on the number of workers by five year age groups spanning age ten to age sixty five. These are available at ten year intervals starting

³Gollin (2002) shows that capital's share is roughly equal across counties.

⁴The choice of coefficients follows Hall and Jones (1999). For the first four years of schooling the return to schooling in sub-Saharan Africa, 13.4 percent, is used. For schooling from four to eight years the world average return to schooling, 10.1 percent, is used. For schooling beyond 8 years the OECD return to schooling, 6.8 percent, is used. The precise method of calculating human capital from schooling turns out to be unimportant for the following results.

⁵This correction is taken from UN national accounts data, as collected in Aiyar and Feyrer (2002). Because the regressions in this paper exclude oil exporting countries, the corrections are quite minor and have very little impact on the results

⁶Their calculations, in turn, are based on the Penn World Tables 5.6. Both are available from the World Bank website (<http://www.worldbank.org/research/growth>).

in 1960. Population by five year age groups is available from the UN. The population data is used to impute the intermediate values for the workforce data.⁷

The availability of both workforce and demographic data allows for the use of instrumental variables to address several issues. First, if participation rates are systematically related to productivity, results may reflect causality from productivity to participation rates. By instrumenting workforce values on population proportions, this channel is eliminated. Second, we may worry that immigration is leading people to migrate to high productivity area. By instrumenting on lagged population, this channel can be eliminated.

Table 1 present the results of a series of cross country regressions. Column 1 is the basic results of total factor productivity versus the age structure of the workforce. All point estimates are negative, indicating that an increase in the size of the excluded group, aged forty to fifty years, is associated with higher productivity. The coefficients on W10, W20, and W30 are significant at the 1% level. The coefficients on W50 and W60 are significant in all the regressions.

The differences between the age groups are extremely large. According to the column one estimates, a five percentage point shift from the thirty year age group to the forty year age group is associated with over a 16% increase in per worker output.⁸ Supposing this shift occurred over a 10 year period, this would add approximately 1.6 percentage points to output growth in each year. Column two restricts the data to the OECD and the results are similar.

Columns 3, 4, and 5 are three robustness checks, which focus on the potential endogeneity problems identified above. Column (3) uses only unimputed values of the demographic measures as regressors. This columns test to see if the imputation procedure used to allow 5 year data is biasing the results. Columns (4) and (5) report the results of IV estimations where workforce measures are instrumented on population measures. For column (4) contemporaneous population measures are used and are limited to the working age population.

⁷The population demographic data used in the imputation is limited to the working age population in order to avoid contaminating the imputed data with information about dependency ratios.

⁸Demographic shifts of this size, while not the norm, are present in the data. Between 1980 and 1990, the proportion of workers in the US aged between 40 and 49 rose by 4.6%.

Table 1: The Effect of Changes in Workforce Composition on TFP - Cross Country Evidence

Sample imputed W	(1)	(2)	(3)	(4)	(5)
	OLS nonoil yes	OLS OECD yes	OLS nonoil no	IV nonoil no	IV 10 yr lag nonoil no
	$\Delta\log(A)$	$\Delta\log(A)$	$\Delta\log(A)$	$\Delta\log(A)$	$\Delta\log(A)$
ΔW_{10}	-3.832 (1.086)**	-3.959 (0.729)**	-4.447 (1.338)**	-5.848 (1.477)**	-6.362 (1.602)**
ΔW_{20}	-3.22 (1.047)**	-3.112 (0.720)**	-2.858 (1.154)*	-3.093 (1.011)**	-3.092 (1.076)**
ΔW_{30}	-3.339 (1.062)**	-2.358 (0.576)**	-3.336 (1.111)**	-3.509 (1.047)**	-3.362 (1.159)**
ΔW_{50}	-2.73 (0.978)**	-2.093 (0.795)*	-1.978 (1.270)	-1.464 (1.640)	-1.002 (1.797)
ΔW_{60}	-3.143 (1.081)**	-2.75 (0.889)**	-4.484 (1.675)**	0.941 (2.698)	1.332 (3.127)
year=1965	0.183 (0.024)**	0.228 (0.020)**			
year=1970	0.04 (0.030)	0.051 (0.043)	0.218 (0.036)**	0.239 (0.038)**	0.243 (0.040)**
year=1975	0.087 (0.038)*	0.124 (0.029)**			
year=1980	-0.053 (0.051)	-0.123 (0.053)*	0.001 (0.043)	0.043 (0.052)	0.037 (0.058)
year=1985	0.007 (0.045)	0.137 (0.039)**			
year=1990	-0.113 (0.060)+	-0.121 (0.048)*	-0.124 (0.043)**	-0.131 (0.049)**	-0.144 (0.053)**
Observations	499	126	246	246	246
Countries	87	21	87	87	87
R-squared	0.15	0.55	0.24		

Standard errors in parentheses

+ significant at 10%; * significant at 5%; ** significant at 1%

This column test whether endogeneity of participation rates is biasing the base results. For column (5) lagged values of the population are used as instruments.⁹ This column tests whether cross country migration is significantly biasing the results.

The results for the robustness tests are similar to the base result for each estimation though less precisely estimated due to reductions in the sample size. The 95% confidence intervals overlap with the base case for all regressors in all three regressions. For all but the W60 group, all point estimates are negative, indicating that movements into the 40 year old group from these groups is associated with higher productivity. For the younger groups, the coefficients are significant in all but one case. For the IV results, W60 has positive point estimates though the standard errors are sufficiently large that the error bands overlap with the base case.

Additional robustness checks were performed which are not presented here. One concern might be that the output data used is in terms of output per worker and does not take into account differences in hours worked which may be age specific. In general, the productivity calculations are quite crude and do not take into account many factors which would be appropriate for a careful analysis of total factor productivity. This is largely due to data limitations. However, some estimates can be made on the subsample of the data for which more detailed information is available. Regressions were run using hours worked data from the OECD. Also, more detailed productivity numbers from Jorgenson (2003) are available for the G7 countries. The results from these subsamples do not contradict the base results.

These results suggest that the age structure of the workforce has a significant correlation with total factor productivity. The regressions using lagged demographics indicate that movements in productivity are not causing contemporaneous changes in demographics. The endogeneity of participation rates and migration are not driving the results.

Though the evidence in this section does not make a conclusive case for a causal link between demographic change and productivity growth, the results certainly suggest that

⁹The instrument set in these cases is restricted to the population structure of people who will be in the workforce at the end of the lag period. For example, the instruments for the 10 year lagged IV regressions are constructed from the population aged 5 to 55 since they will comprise the working age population 10 years hence.

it is possible. Many alternative explanations have been eliminated by the IV results. Any non-causal explanation would require some omitted factor which had an impact on the demographic structure in the past but which affects productivity with long lags. Given this, looking for further evidence of contemporaneous causal links seems sensible.

2 Implications

2.1 Cross Country Productivity Differences

The results of the previous section can be used to provide insight into cross country productivity patterns. The demographic characteristics of the workforce differ greatly across countries with different income levels. Figure 1 illustrates the proportion of the workforce between the ages of 40 and 49 by income level.

Two facts are immediately apparent. The poorer nations have a lower proportion of forty year old workers than the richer nations in every year. The second aspect of the graph is the trend. The wealthy nations saw a relatively static 40 year old cohort until about 1980. From 1980 until 2000 the proportion of 40 year olds increases dramatically. This is not true of the poor nations.

The results of the previous section lead to two obvious conclusions. First, some proportion of the income gap between rich and poor nations can be attributed to persistent differences in the demographic structures. Poor nations typically have younger workforces which the results suggest lead to lower productivity. Feyrer (2004) suggests that one quarter to one third of the rich-poor productivity gap can be explained by steady state demographics. Second, over the second half of the sample the demographically induced productivity gap has gotten worse.

2.2 The US and Japan

Relative demographic movements can also inform us about relative growth rates between different rich countries. The demographic composition of the Japanese workforce has dif-

ferred greatly from the US in the postwar period. Figure 2 shows the number of live births in Japan and the US in the post war period.¹⁰ The most remarkable feature of this graph is the degree to which US and Japanese birthrates move in opposite directions.¹¹ During the peak of the baby boom (around 1960), Japan was experiencing a local minimum in births. Japan had an upsurge in births during the mid seventies as the US was experiencing a significant slowdown in births. Consequently the Japanese workforce has very different demographic movements than the US. Japan has a steeply rising cohort of workers in their forties from 1960 to 1980, a period when the US saw this cohort fall in size. From 1990 to 2000 the situation reverses.

The demographic effect maps roughly to the observed growth pattern between the US and Japan. Between 1960 and 1980, the US was experiencing worsening demographics and low productivity growth. Figure 3 shows the demographic effect on productivity implied by the results presented earlier in the paper.

The model suggests that 2-3% of the differences between US and Japanese growth in the seventies is correlated with demographic shifts. In the nineties, this situation reverses. The US saw higher productivity growth due to demographic change while Japan experienced declining productivity growth. Demographics are associated with a 2% differential between the US and Japan during the nineties. The model predicts that relative growth rates are set to reverse once again in the coming decade. The US is about to enter a period of slower productivity growth while Japan should see a significant improvement in productivity growth relative to the nineties.

¹⁰This graph is the raw number of births and is not scaled to the size of the population. The important features of the graph are the locations of the peaks and troughs, which are more easily seen in the unscaled graph.

¹¹In 1966 there was a dramatic one year downturn of almost one half million births in Japan. Apparently, 1966 was the most recent “Year of the Fire Horse.” According to Japanese superstition, girls born in the year of the Fire Horse will have very unhappy lives and most likely will kill their husbands.

Figure 1: The Proportion of Workers Aged 40-49 By Income Group




Figure 2: Live Births in the US and Japan 1950-2000





Figure 3: The Demographic Effect on US and Japanese Productivity Growth



3 Externalities

The cross country results show that demographics have a significant correlation with output and productivity. This should come as no surprise since labor economists have long identified experience effects in the wages of workers. The canonical Mincer wage regression takes the following form.

$$\log(wage) = \alpha + \beta_1 * school + \beta_2 * experience + \beta_3 * experience^2 + \epsilon. \quad (5)$$

Bils and Klenow (2000) collect a sample of these coefficients estimated for 52 countries. Using the average coefficients from their sample produces

$$\log(wage) = \alpha + 0.096 * school + 0.051 * experience - 0.00071 * experience^2. \quad (6)$$

According to these estimates, an additional year of schooling increases the wage by 9.6%.¹² Experience has diminishing returns, with each addition year of experience increasing the wage by some amount less than 5.1%

Worker productivity rises with age up to about age fifty then falls somewhat. The Mincer evidence implies that there is about a 60% difference between the productivity of twenty years old workers and fifty year old workers. For the aggregate data, this implies that an economy with a large cohort of young workers will have lower productivity than an economy with an large cohort of older workers.¹³

The Mincer evidence is therefore relatively similar to the results presented in this paper. However, there are enormous differences in the magnitude of the effects. The Mincer evidence suggests that moving 5% of the population from the 20-30 year old age category to the 40-50 year old age category will increase wages (and output) by 1-2%. The evidence presented in this paper suggests that this same demographic shift is associated with a 10-15% increase in output, an effect an order of magnitude larger than predicted by the Mincer evidence.

However, the Mincer evidence may not tell the entire story at the aggregate level. The micro evidence, based on wage data, only captures the private return to experience and education. It may be that the social returns are higher than the private returns. Externalities to experience (or age) may mean that the Mincer coefficients are understating the aggregate productivity effects.¹⁴ The results of this paper suggest that there are externalities to workforce demographic composition that go beyond the private return to experience.

The next two sections will suggest two possible channels through which social returns to

¹²While there is cross country variation in the coefficient estimates, the range of variation is relatively small. Bils and Klenow (2000) find coefficients on schooling as high as 0.28 (Jamaica) and as low as 0.024 (Poland). The majority of coefficients, however, fall between 0.05 and 0.15.

¹³Assuming, of course that productivity is measured in a way that does not account for differences in human capital through experience.

¹⁴The importance of externalities to education has long been emphasized beginning with the theoretical work of Nelson and Phelps (1966). It has, however, been difficult to show that these externalities exist empirically. Panel growth regressions such as Caselli, Esquivel and Lefort (1996) take into account country specific productivity effects and try to deal with the endogeneity of schooling relative to output. These regressions fail to find coefficients on schooling consistent with large externalities. In a study of US states, Acemoglu and Angrist (2000) also fail to find evidence of large externalities to education. Aiyar and Feyrer (2002) find evidence of dynamic externalities to human capital that act over long time periods.

the age structure might be realized. First, idea creation through inventive activity. Second, idea adoption through managerial talent and entrepreneurial activity. The first is important because the non-rival aspect of ideas increases the potential for large externalities.

3.1 Innovation

Suppose that productivity changes are driven by individuals engaged in innovative activity. The private returns to experience are unlikely to capture the full societal gains from innovation because of the inability of firms to capture the full surplus created by innovation. Many types of innovation are, by their nature, non-rival. Non-rivalry may make it particularly difficult to capture more than a small fraction of the gains of innovative activity. In many innovative industries a large proportion of productivity increases may benefit consumers far beyond the price that they pay for the product. Take as an extreme example the Google search engine. Google has almost certainly increased the productivity of academic researchers as well as anyone else who relies on the internet for productivity enhancing information.¹⁵ Yet, most people have never paid any money to Google. While the creators of Google have benefitted from their creation, it seems likely that their revenues represent only a small fraction of the aggregate gains in output that their invention has made possible.

Suppose that the age structure of the workforce affects the probability that an invention like Google will be created. If a country has an age structure which increases the likelihood of Google being invented, productivity will be higher for all workers. Only a very small fraction of these productivity gains will be captured by the original inventors. If this hypothesis is true, then we should not be surprised that the aggregate productivity effects are much larger than the micro Mincer effects.

There is evidence that generating and implementing new ideas varies by age. Lehman (1953) finds that creative output in science and invention varies substantially by age. There is some variation among disciplines, but Lehman finds that peak productivity tends to be

¹⁵Even if you are not using the Google search engine, Google's competitors have almost certainly improved as a result of the competition.

in the interval between ages 30 and 40. If there is indeed an age effect in idea generation, having a larger cohort of workers in the peak idea generating ages should result in more rapid production of new ideas and new technologies. As an extreme example, consider the world of academic mathematics, where a significant portion of the innovative ideas are produced by people between the ages of twenty five and thirty five.¹⁶ If the world were like a mathematics department, we would expect to see more new ideas being produced in countries with a large cohort of young workers.


More recent work by economists have also found a link between age and creative performance. Galenson and Weinberg (2000) find that artistic output is related to the age of the artist. Galenson and Weinberg (2004) find that the peak years for Nobel Prize winning economists tends to be in their forties. Jones (2004) collects the birth dates for a sample of inventors granted patents in the NBER patent database. Figure 4 presents kernel density estimates of the age distribution of these patent grantees by year from 1975 until 1995. While the age profile of inventors changes somewhat in response to the large underlying changes in the age distribution of the workforce as a whole, the median age of 48 does not vary by more than one year during the sample period. This is in stark comparison to the age profile of managers which will be presented in the next section. The relatively stable distribution of the patent holders suggests that creativity profile of inventors may be quite stable with a peak somewhere in the mid forties. When demographic change results in a low number of workers (and therefore inventors) in this age group it seems likely that there will be a reduction in the level of inventive activity.

3.2 Idea Adoption

While creative output is one potential channel through which age may impact productivity, it may not be the most relevant for cross country comparisons. For most of the countries in the world, it is not idea creation that matters so much as idea adoption. Organizations (or countries) that increase productivity by producing new ideas are different than

¹⁶Lehman (1953), p. 8

Figure 4: Age Distribution of US Inventors Granted Patents, by year



organizations that adopt ideas generated elsewhere.

Idea creators operate at the technological frontier at all times because they define the frontier. The rate of new idea creation determines the rate of expansion of the frontier. For technology adopters, the technological frontier is a given. Nothing an adopter does impacts the rate of expansion of the frontier and adopters are always operating below the frontier. The relevant question is how far below the frontier they are operating. If age structure affects the rate of technology adoption, then favorable demographic shifts may make a country more effective at implementing ideas generated elsewhere. This allows the country to get closer to the frontier and in the short run this means more rapid productivity growth. However, in the long run growth will be determined by the movement of the frontier which is exogenous from the point of view of the adopter. It seems apparent that most countries in the world are technology adopters.

There is microeconomic evidence that age matters in the adoption of technology. Weinberg (2002) finds that both experience and age matters for technology adoption. Among

high school graduates, technology adoption complements experience while among college graduates, technology adoption complements youth. This evidence points toward a tension between youth and experience. Since schooling tends to be concentrated early in life, young workers have the advantage of more recent human capital.¹⁷ It may also be that younger workers are less bound by tradition and more likely to take risks. Young workers, on the other hand, lack human capital in the form of experience.

Lucas (1978) suggests that the quality distribution of managers may play a large role in determining output. In the Lucas model, a firm with a manager of quality x managing n workers and k units of capital will produce the following amount of output,

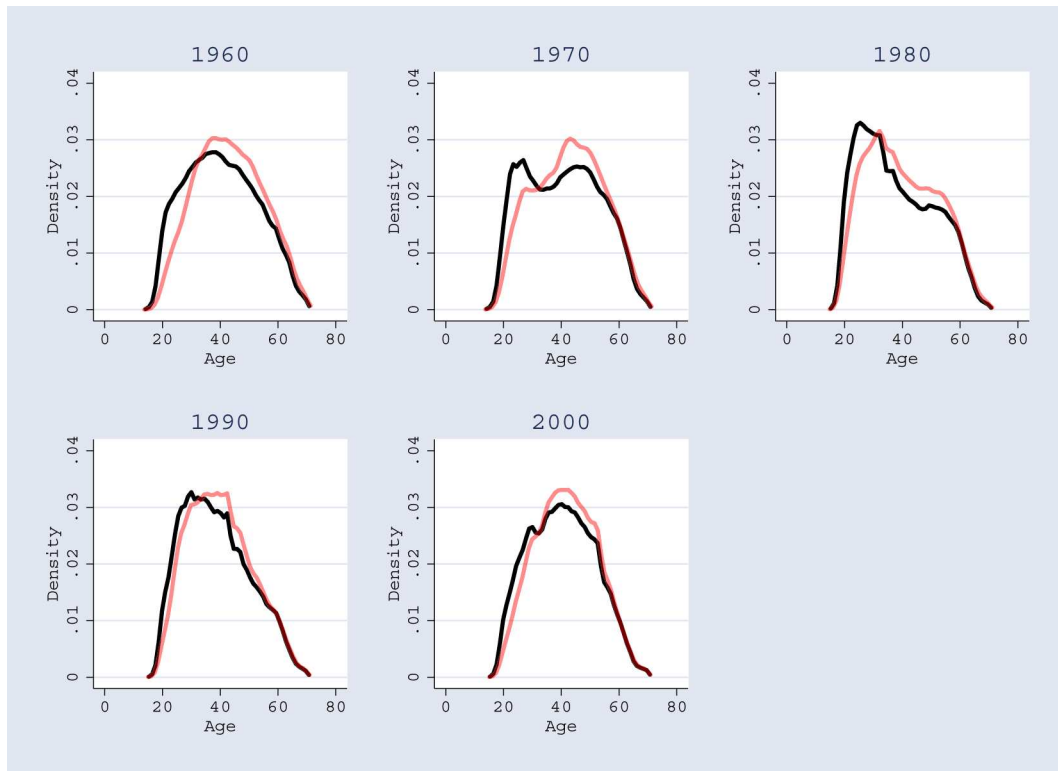
$$y = xg[f(n, k)] \tag{7}$$

where $f()$ is a standard neoclassical production function, and $g[]$ has decreasing returns. The decreasing returns to $g[]$ imply that increasing the size of any given firm will reduce per worker output. This indicates that there are advantages to having smaller firms, on average. However, each firm needs to have a manager. In order to have smaller firms, you must have a larger group of managers. Assuming heterogeneity in management talent, an efficient allocation of workers into management positions will result in a talent cutoff, v . Workers with managerial talent $x > v$ will be managers and all other workers will be normal workers in firms. In order to reduce average firm size, this threshold will need to be reduced, causing a fall in overall management quality. These two competing factors result in an equilibrium number of managers.

This model would seem to apply to managerial age insofar as age affects managerial talent. We observe that young workers are much less likely to take management positions than older workers. This is likely because some amount of experience is important in managing other workers. It may also be that social constraints prevent young workers from managing older workers even if they are particularly talented. Up and out promotion systems of the sort used in the military tend to produce a structure where people are always

¹⁷Chari and Hoenhain (1991) find that technologies diffuse slowly due to vintage human capital effects.

Figure 5: Age Distribution of US Workforce categorized as managers, by year



Entire Workforce - Bold line, Managers - Grey Line

managed by someone older than themselves.

In either case, a large influx of young workers will increase the probability that a worker in one of the smaller and older cohorts will be called upon to take a management role. This suggests that the marginal manager will be less talented since we are dipping farther into the talent pool of the older cohorts. The Lucas model suggests that less talented managers will make all workers less productive.

An examination of census data suggests that the entrance of the baby boom into the US workforce caused significant changes in the age structure of the the management of US firms. Figure 5 shows the evolution of the age distribution of managers in the US over time against the evolution of the workforce as a whole. The dark line is a kernel density estimate of the age distribution of the US workforce over time. The grey line is the age distribution of US workers categorized as managers.

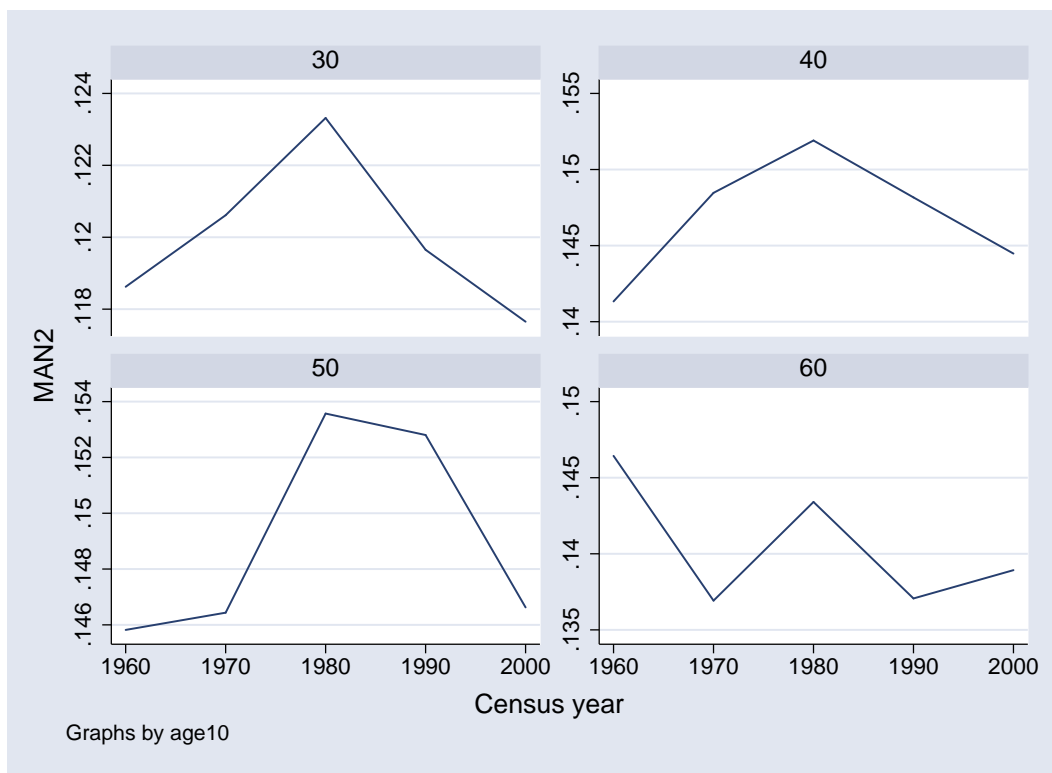
The baby boom first enters into the workforce in large numbers in the 1970 census, but

they are not well represented in the management workforce in this year. This is consistent with the idea that young workers are not chosen to be managers due to lack of experience. This implies that a worker with the necessary experience to manage was more likely to be a manager in 1970 than in 1960. The marginal manager was therefore likely to be less talented as the baby boomers entered in large numbers. By 1980, the baby boom has fully entered the workforce, but is still quite young and is proportionally under represented in the managerial workforce. However, the overall size of the boom was such that the average age of managers fell by 4 years from 1970 to 1980 from 43 years old to 39 years old. As the mass of the boomers enters their thirties in the 1990 census, the managerial workforce begins to return to its earlier shape. By 2000, when the boomers are of an age when people typically are in management, the distribution looks almost identical to the 1960 distribution. Indeed, the median age of managers in 2000 is nearly the same as in 1960.

Figure 6 shows the proportion of each age category in management job classifications over time. Over this time period there was a secular increase in the proportion of workers classified as managers, so the data has been detrended to emphasize the within group effects. The most striking feature of this graph is the increase in the proportion of workers classified as management (relative to trend) from 1960 followed by a decline from 1980 until 2000. As argued earlier, when the baby boomers were young, they were under represented in the management workforce. This necessitated that a larger percentage of the older cohorts enter management roles. As the boom aged, they began taking over the management burden generated by the size of their cohort. Between 1980 and 2000, this results in a lowering of the proportion of managers in each age cohort (relative to the time trend).

The data suggests that from 1960 until 1980 the entry of the baby boomers resulted in a lowering of marginal manager quality, while from 1980 until 2000 the baby boom's aging resulted in higher manager quality. This effect is likely magnified by the fact that the workers in the baby boom cohort were called on to manage earlier in their careers, so that by 2000 they not only had an appropriate experience level to manage other workers, but also had more specific experience as managers than other cohorts at the same age. The

Figure 6: Proportion of US Workforce categorized as managers (detrended), by age group



Lucas model suggests that higher manager quality will cause higher overall productivity, potentially contributing to the aggregate results.

4 Tests on US state and Metropolitan area data

The results presented thus far suggest that the composition of the workforce at the country level matters. If this result also holds at the US state and local level this may suggest a different set of mechanisms than a result which is confined to the national level. This section will attempt to replicate the cross country results using US state and metropolitan area data.

State and metropolitan area data on income and the demographic composition of the workforce is taken from the Integrated Public Use Microdata Series of the US census.¹⁸ The base data is a 5% sample of the US census from 1960 until 2000. Population proportions

¹⁸<http://www.ipums.org>

are from the entire sample while workforce proportions are from a subsample of full time workers. Income is measured as the average hourly income of full time workers.

Estimation is identical to the cross country sample, but the use of US data presents several challenges. One problem is that of endogenous participation rates. For example, in the US high participation rates among teenagers may be correlated with unobservable area characteristics that drive wages downward. By instrumenting the workforce proportions on population proportions, this source of bias can be addressed.

Within the US, the workforce is highly mobile. If migration is driven by wage differentials and people of different age groups migrate differentially, the US estimations will likely reflect reverse causality. In the cross country data, this was much less of an issue because cross country mobility is much smaller than US cross state and metro area mobility. It is possible to eliminate the impact of mobility by instrumenting on lagged population data. In the cross country data, this showed that mobility was not driving the results. The same test will be applied for the US data.

Table 2 shows the results of regressions at the state and metropolitan area level of log hourly wage on demographic proportions. Estimation is done in differences to eliminate fixed effects. Standard errors are clustered at the regional area to deal with serial correlation. Three different estimations were performed on each sample. Columns (1) and (4) were estimated using OLS. Columns (2) and (5) were estimated using IV with current population demographics as instruments. This estimation deals with the problem of endogenous participation rates. Columns (3) and (6) estimate using IV with ten year lagged population demographics as instruments.

These results are not nearly as clear as the cross country results and are somewhat inconclusive. Point estimates and almost all significant coefficients are negative.¹⁹ In general, the point estimates suggest that the effect of demographics is smaller at the state

¹⁹The exception to this is the positive and significant coefficient on twenty year old workers in regression (5). Recall that one of the problems in using US state and MSA data is migration. If twenty year olds are differentially likely to migrate to high income area, we should expect a larger coefficient on the 20 year old age category. Instrumental variables regressions using lagged population data as instruments should eliminate migration as a source of bias. Doing so eliminates the positive coefficient on twenty year old workers (albeit with a large standard error).

Table 2: The Effect of Changes in Workforce Composition on US wages - State and metro Area

Sample	1 State OLS	2 State IV	3 State IV-lag	4 MSA OLS	5 MSA IV	6 MSA IV-lag
	$\Delta\log(\text{wage})$	$\Delta\log(\text{wage})$	$\Delta\log(\text{wage})$	$\Delta\log(\text{wage})$	$\Delta\log(\text{wage})$	$\Delta\log(\text{wage})$
$\Delta W15$	-3.497 [1.855]+	-3.206 [11.395]	-7.253 [5.806]	-1.967 [0.620]**	-5.402 [2.132]*	1.608 [7.551]
$\Delta W20$	0.032 [0.406]	0.19 [0.826]	-0.827 [1.523]	0.273 [0.184]	0.529 [0.254]*	0.388 [0.586]
$\Delta W30$	-0.743 [0.546]	-0.943 [1.281]	-1.536 [1.194]	-0.135 [0.184]	-0.315 [0.308]	-0.521 [0.823]
$\Delta W50$	-1.357 [0.780]+	-4.371 [1.229]**	-4.773 [1.417]**	-0.007 [0.226]	-0.759 [0.443]+	-0.04 [0.809]
$\Delta W60$	-0.354 [1.243]	12.76 [11.843]	1.195 [1.947]	0.157 [0.327]	-0.451 [2.480]	2.88 [1.795]
Observations	190	190	188	441	441	335
R-squared	0.79	0.46	0.74	0.72	0.67	0.73

Standard errors in parentheses

+ significant at 10%; * significant at 5%; ** significant at 1%

and MSA level than at the cross-country level. In the case of the state level results, we cannot reject that the effect of demographics is different than at the cross country level, at least in the IV regressions. However, the confidence intervals for the OLS regression at the state level suggest that this set of coefficients is significantly smaller than the point estimates for the cross country regressions. Moving to a lower level of aggregation, all three regressions at the MSA level have coefficients which are significantly closer to zero than the cross country point estimates.

These results suggest that the demographic effects diminish at lower levels of aggregations. This should not come as a large surprise. The US has highly integrated labor and product markets. Suppose that the increased availability of prime age workers makes an individual plant more efficient. These efficiency gains will not necessarily be captured completely in the wages of the workers of that plant. Shareholders of the firm, who do not necessarily live in the immediate area, may reap some of the gains. If product markets are competitive, consumers will gain due to lower marginal costs of production. Obviously,

these consumers are not necessarily located in the same geographic area as the workers and plants.

5 Conclusion

The results presented in this paper show that workforce demographics are strongly correlated with productivity and output. A significant portion of the productivity gap between rich and poor countries may be related to different demographic structures. The results also appear to capture some of the productivity divergence between the poor and rich countries since 1980. Looking at the experience of Japan and the US over the last 40 years, the relative demographic movements are consistent with the cross country results.

Given the importance of productivity in explaining cross country income differences, this is a useful result. Demographics have substantial predictable time series variation that is largely exogenous to contemporaneous events, at least at the country level. Also, the regressions using lagged demographics indicate that movements in productivity are not causing contemporaneous changes in demographics. The magnitudes of the results are much larger than one would expect from the standard labor results, suggesting that externalities play a large role.

Two possible hypotheses are suggested as mechanism through which the age distribution might effect aggregate output. First, the productivity of innovative activity is undoubtedly related to age. However, US patent data shows that the age distribution of innovators did not change substantially in the US as a result of the entry of the baby boom to the workforce. This suggests that changes in the supply of workers that are at the prime age to innovate may have an impact on the rate of innovation.

By contrast there were substantial changes in the age distribution of management in the US. Initially the baby boom were inexperienced and could not provide their own management talent, necessitating the use of less talented managers from older cohorts. As the boom aged, they entered management ranks earlier than previous cohorts. This had the

net effect on increasing the proportion of managers drawn from all age cohorts from 1960 to 1980, almost certainly lowering management quality. This trend reverses from 1980 until 2000.

Results at the US state and MSA level, while less conclusive than the cross country results, suggest that effect of demographics is smaller at the lower levels of aggregation. This suggests that the externalities at work are stronger at higher levels of aggregation. Given the integrated product and labor markets in the US, this is not surprising. The non-rivalry of ideas makes it likely that an age-innovation link may not be evident at the state or MSA level because the gains of inventive activity are spread out quickly with little regard to geography. For management, however, it is not hard to imagine that gains will be more (though not completely) local.

Understanding the relationship between demographics and productivity is important because of the useful and predictable characteristics of demographics and because the significance of the relationship is strong. Almost every region in the world is experiencing significant demographic change. The rich nations are rapidly becoming older and most have birthrates below replacement level. Some poor countries are experiencing dramatically reduced birthrates in the wake of population explosions. Understanding how these changes will affect productivity over the coming decades is of crucial importance. While this paper shows that there is a relationship between productivity and demographics, more research is needed to understand the mechanisms behind this relationship.

References

- Acemoglu, Daron and Josh Angrist**, “How Large Are Human Capital Externalities? Evidence from Compulsory Schooling Laws,” *NBER Macro Annual*, 2000, pp. 9–59.
- Aiyar, Shekhar and James Feyrer**, “A Contribution to the Empirics of Total Factor Productivity,” *Dartmouth College Working paper*, August 2002.
- Bils, Mark and Peter Klenow**, “Does Schooling Cause Growth?,” *American Economics Review*, December 2000, *90* (5), 1160–1183. Also NBER 6393.
- Caselli, F., G. Esquivel, and F. Lefort**, “Reopening the Convergence Debate: A New Look at Cross-Country Growth Empirics,” *Journal of Economic Growth*, September 1996, *1* (3), 363–389.
- Chari, V. V. and Hugo Hoenhain**, “Vintage Human Capital, Growth, and the Diffusion of New Technology,” *Journal of Political Economy*, 1991, *99* (6), 1142–1165.
- Easterly, William and Ross Levine**, “It’s Not Factor Accumulation: Stylized Facts and Growth Models,” *World Bank Economic Review*, 2001, *15* (2), 177–219.
- Feyrer, James**, “Demographics and Productivity,” *Dartmouth College Working Paper*, 2004.
- Galenson, David W. and Bruce Weinberg**, “Age and the Quality of Work: The Case of Modern American Painters,” *Journal of Political Economy*, August 2000, *108* (4), 761–777.
- and — , “Creative Careers: Age and Creativity among Nobel Laureate Economists,” *Working paper*, June 2004.
- Gollin, Douglas**, “Getting Income Shares Right,” *Journal of Political Economy*, April 2002, *110*, 458–475.
- Hall, Robert and Charles I. Jones**, “Why Do Some Countries Produce So Much More Output per Worker Than Others?,” *Quarterly Journal Of Economics*, February 1999.
- Jones, Ben**, “Age and Great Invention,” *Working Paper*, 2004.
- Jorgenson, Dale**, “Information Technology and the G7 Economies,” *World Economics*, October-December 2003, *4* (4), 139–169.
- Lehman, Harvey C.**, *Age and Achievement*, Princeton University Press, 1953.
- Lucas, Robert E.**, “On the Size Distribution of Business Firms,” *The Bell Journal of Economics*, Autumn 1978, *9* (2), 508–523.
- Nelson, Richard and Edmund Phelps**, “Investment in Humans, Technological Diffusion, and Economic Growth,” *American Economics Review*, 1966, *56*, 69–75.
- Psacharopoulos, George**, “Returns to Investment in Education: A Global Update,” *World Development*, 1994, pp. 1325–1343.

Weinberg, Bruce A., "Experience and Technology Adoption," *Ohio State University working paper*, May 2002.