Aggregate Evidence on the Link Between Age Structure and Productivity

James Feyrer

This chapter examines the relationship between the age structure of the workforce 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 then this suggests a relationship between worker productivity and age. In the aggregate, therefore, we should expect changes in the age structure of a population to be correlated with changes in productivity.

Feyrer (2007) finds a strong and robust correlation between workforce age structure and total factor productivity. The impact of age structure 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 reverse causality from productivity to workforce age structure 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 chapter examines the nature of this externality in several ways.

To begin, I 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 examine the evolution of the age distribution of patent holders in the United States 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 explore how the entry of the baby boom cohorts into the workforce changed the composition of the managerial workforce over time.

For this reason, it may be useful to identify the scope of the externality. Does it 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?

First, I examine the relationship between age structure and productivity at the country level. Next I discuss the implications of the cross-country results for cross-country economic performance and look at the relative performance of the United States and Japan in light of the results. I suggest channels through which age structure may be affecting output and review evidence from the US census. I then examine US state and metropolitan area data to see if the cross-country effects are evident at lower levels of aggregation.

Cross-country evidence

For use in a cross-country regression, demographic measures have several characteristics that make identification more straightforward than with many variables typically used in the literature on cross-country growth. First, demographic measures are strongly predetermined. The current age structure of the workforce was determined roughly 20 years ago and should be predetermined with respect to current output movements. Second, demographic structure has significant time series variation. This time series variation allows for exploiting the panel nature of the data.

The following results largely follow Feyrer (2007) and focus on total factor productivity. Feyrer (2007) shows that the impact of demographic shifts on physical and human capital accumulation is relatively small and uncorrelated with the productivity effects. For purposes of examining output, productivity is the key variable, and 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 $x_{i,t}$,

$$y_{i,t} = f_i + \mu_t + \beta_{x_{i,t}} + \mu_{i,t}. \tag{1}$$

The regressors are the proportion of the workforce by age group, with W10 indicating workers between ages 10 and 19, W20 workers between 20 and 29, and so on. W60 indicates workers age 60 years and older. Since these variables are proportions, the sum of all the age groups is 1.0 for each country year pair. For this reason, one group is excluded.[1] For most of the reported regressions in this section, first differencing was used to eliminate the country-specific effect.

Productivity is calculated as a residual. I assume a Cobb-Douglas production function taking physical capital, human capital from schooling, and productivity as inputs.

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

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.[2] The human capital production function is assumed to have a Mincer form

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

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.[3] The production function can be solved for log total factor productivity.

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

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.[4] Data for capital per worker are from Easterly and Levine (2001).[5] The schooling data used to calculate human capital stocks are from Barro and Lee (2001). All variables used in calculating total factor productivity (TFP) are levels from the individual year in question.

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 10 to age 65. These are available at ten-year intervals starting in 1960. Population by five-year age groups is available from the United Nations. The population data are used to impute the intermediate values for the workforce data.[6]

The availability of both workforce and demographic data allows for the use of instrumental variables (IV) to address several issues. First, if participation rates are systematically related to productivity, results may reflect causality from productivity to participation rates. Instrumenting workforce values on population proportions eliminates this channel. Second, we may worry that immigration is leading people to migrate to high-productivity areas. Instrumenting on lagged population eliminates this channel. Another potential area of concern is that dependency ratios are correlated with workforce age structure. The dependency ratio is added as an additional control and does not change the basic results.

The basic sample contains 87 countries at five-year intervals between 1960 and 1990. The years are limited by the workforce proportion data. The base regressions are also run on an OECD subsample of 21 countries. For the five-year frequency regressions, the workforce data are imputed from the ten-year data and population data. Since the IV strategy uses population data to instrument for the workforce, these imputations cannot be used for

IV. The IV results therefore are at ten-year intervals. The dependent variable is log TFP in all regressions.

Tables 1 and 2 present the results of a series of cross-country regressions. The two tables differ only in their estimation method. Table 1 is estimated in differences and Table 2 is in levels with unreported country-level dummies. The coefficients are directly comparable. In both cases, standard errors are clustered to account for serial correlation. Since the differenced estimator is robust to a unit root in TFP, these estimates are generally preferred. The results are nearly identical, although the differenced estimates are more precisely estimated.

Column (1) is the basic result 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 40–50 years, is associated with higher productivity. The coefficients on W10, W20, and W30 are significant at the 1 percent level. The coefficients on W50 and W60 are significant in all the regressions.

The differences between the age groups are very large. According to the column (1) estimates, a 5 percentage point shift from the 30-year age group to the 40-year age group is associated with over a 16 percent increase in per worker output.[7] Supposing this shift occurred over a 10-year period, this would add approximately 1.6 percentage points to output growth in each year. Column (2) adds in the dependency ratio as an additional control. It does not affect the results. Columns (3) and (4) replicate columns (1) and (2) for the OECD sample with similar results.

Columns (5), (6), and (7) are three robustness checks, which focus on the potential endogeneity problems identified above. Column (5) uses only unimputed values of the demographic measures as regressors. This column tests whether the imputation procedure used to allow five-year data is biasing the results. Columns (6) and (7) report the results of IV estimations where workforce measures are instrumented on population measures. For column (6) contemporaneous population measures are used and are limited to the working-age population. This column tests whether endogeneity of participation rates is biasing the base results. For column (7) lagged values of the population are used as instruments.[8] 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 although less precisely estimated because of reductions in the sample size. The 95 percent 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, although the standard errors are sufficiently large that the error bands overlap with the base case.

TABLE 1 Effect of changes in workforce composition on changes in total factor productivity: Cross-country regressions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	OLS	OLS	OLS	OLS	OLS	IVL0	IVL10
Sample:	nonoil	nonoil	OECD	OECD	nonoil	nonoil	nonoil
Imputed W:	yes	yes	yes	yes	no	no	no
ΔW10	–3.774	–3.797	–3.996	–4.063	–4.42	–5.753	–6.254
	(1.085)**	(1.109)**	(0.739)**	(0.778)**	(1.342)**	(1.470)**	(1.594)**
ΔW20	–3.152	–3.704	–3.095	–3.233	–2.766	–3.022	–3.02
	(1.044)**	(1.028)**	(0.723)**	(0.844)**	(1.148)*	(1.004)**	(1.068)**
ΔW30	–3.312	–3.661	–2.323	–2.395	–3.296	–3.468	–3.317
	(1.059)**	(1.029)**	(0.580)**	(0.610)**	(1.110)**	(1.046)**	(1.157)**
ΔW50	–2.661	–2.731	–2.04	–2.122	–1.877	–1.392	–0.927
	(0.972)**	(0.985)**	(0.800)*	(0.801)*	(1.265)	(1.627)	(1.784)
ΔW60	–3.046	–3.309	–2.709	–2.81	–4.305	1.038	1.444
	(1.079)**	(1.052)**	(0.899)**	(0.919)**	(1.666)*	(2.689)	(3.114)
ΔDependency ratio		–1.812		–0.4			
		(0.723)*		(0.868)			
Year							
1965	0.183	0.185	0.231	0.229			
	(0.024)**	(0.024)**	(0.020)**	(0.019)**			
1970	0.044	0.049	0.051	0.052	0.223	0.243	0.247
	(0.030)	(0.031)	(0.043)	(0.044)	(0.035)**	(0.037)**	(0.040)**
1975	0.088	0.071	0.126	0.125			
	(0.038)*	(0.041)+	(0.029)**	(0.029)**			
1980	–0.048	–0.048	–0.127	–0.132	0.007	0.05	0.044
	(0.051)	(0.051)	(0.053)*	(0.056)*	(0.043)	(0.052)	(0.058)
1985	0.009	–0.016	0.139	0.134			
	(0.045)	(0.049)	(0.040)**	(0.037)**			
1990	–0.107	–0.109	–0.124	–0.13	–0.117	–0.122	–0.135
	(0.060)+	(0.060)+	(0.048)*	(0.049)*	(0.043)**	(0.049)*	(0.052)*
Observations	499	499	126	126	246	246	246
Countries	87	87	21	21	87	87	87
R-squared	0.15	0.16	0.55	0.55	0.24		

Standard errors clustered by state in parentheses.
+ significant at 10 percent; * significant at 5 percent; ** significant at 1 percent
NOTE: ΔW10 is the change in the proportion of workers aged 15–19. ΔW20, W30, W40, W50 are the changes in the proportion of workers ages 20–29, 30–39, 40–49, and 50–59. ΔW60 is the change in the proportion of workers ages 60 and older. The dependency ratio is the proportion of the population younger than 15 and older than 64. The nonoil sample excludes the set of countries that the World Bank classifies as oil exporters.

Additional robustness checks were performed but are not presented here. One concern might be that the output data used are measured in terms of output per worker and do 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 that would be appropriate

TABLE 2 Effect of workforce composition on total factor productivity: Cross-country regressions

Sample: Imputed W:	(1) nonoil yes	(2) nonoil yes	(3) OECD yes	(4) OECD yes	(5) nonoil no	(6) nonoil no	(7) nonoil no
W10	–4.913	–3.404	–5.85	–5.637	–4.697	–6.744	–6.149
	(1.507)**	(1.633)*	(1.074)**	(1.255)**	(1.822)*	(2.111)**	(2.292)**
W20	–0.928	–2.257	–2.165	–1.607	–1.109	–1.773	–2.099
	(1.391)	(1.157)+	(0.858)*	(1.084)	(1.796)	(1.493)	(1.549)
W30	–2.307	–2.811	–3.931	–3.625	–1.926	–1.8	–1.376
	(1.155)*	(1.121)*	(0.943)**	(0.930)**	(1.666)	(1.667)	(1.852)
W50	–0.987	–1.124	–1.447	–1.113	–1.251	–1.682	–0.367
	(1.565)	(1.554)	(0.930)	(0.921)	(1.972)	(2.249)	(2.471)
W60	–4.905	–5.46	–2.974	–2.476	–4.817	0.874	2.578
	(1.911)*	(1.704)**	(2.069)	(1.999)	(2.368)*	(4.475)	(5.158)
Dependency ratio		–2.795		1.311			
		(1.002)**		(1.209)			
Year							
1965	0.256	0.222	0.255	0.261	11.018	11.102	10.711
	(0.037)**	(0.034)**	(0.030)**	(0.034)**	(1.278)**	(1.301)**	(1.378)**
1970	0.212	0.224	0.216	0.216	0.215	0.233	0.247
	(0.040)**	(0.039)**	(0.053)**	(0.057)**	(0.043)**	(0.049)**	(0.050)**
1975	0.38	0.329	0.389	0.391			
	(0.054)**	(0.057)**	(0.058)**	(0.060)**			
1980	0.154	0.175	0.226	0.244	0.163	0.223	0.269
	(0.056)**	(0.053)**	(0.065)**	(0.069)**	(0.063)*	(0.082)**	(0.094)**
1985	0.295	0.224	0.432	0.468			
	(0.053)**	(0.059)**	(0.054)**	(0.054)**			
1990	0.004	0.022	0.234	0.291	0.007	0.033	0.098
	(0.083)	(0.077)	(0.087)*	(0.099)**	(0.095)	(0.121)	(0.135)
Observations	586	586	147	147	333	333	333
Countries	87	87	21	21	87	87	87
R-squared	1	1	1	1	0.89	0.88	0.87

Standard errors clustered by state in parentheses.
+ significant at 10 percent; * significant at 5 percent; ** significant at 1 percent
NOTE: W10 is the proportion of workers aged 15–19. W20, W30, W40, W50 are the proportion of workers ages 20–29, 30–39, 40–49, and 50–59. W60 is the proportion of workers ages 60 and older. For definitions of dependency ratio and nonoil sample, see note to Table 1.

for a careful analysis of total factor productivity. This is largely attributable 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 data on hours worked 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 age structure indicate that movements in productivity are not causing contemporaneous changes in age structure. Possible endogeneity of participation rates and migration is not driving the results.

Although 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 such a link is likely. Many alternative explanations have been eliminated by the IV results. Any noncausal explanation would require some omitted factor that had an impact on the demographic structure in the past but that affects productivity with long lags. Given this, looking for further evidence of contemporaneous causal links seems sensible.

Implications

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 groups of countries of differing income levels.

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

The results of the previous section lead to two obvious conclusions. First, some proportion of the income gap between rich and poor countries can be attributed to persistent differences in the age structure. Poor countries typically have younger workforces, which the results suggest lead to lower productivity. Feyrer (2007) suggests that one-quarter to one-third of the rich–poor productivity gap can be explained by steady-state demographic structure. Second, over the second half of the sample the demographically induced productivity gap has further widened.

The United States and Japan

Relative demographic movements can also inform us about relative growth rates between rich countries. The demographic composition of the Japanese workforce has differed greatly from that of the United States in the period since World War II. Figure 2 shows the number of live births in Japan and the United States in the postwar period.[9]

FIGURE 1 Workers aged 40–49 as percent of all workers by groups of countries differing in income levels

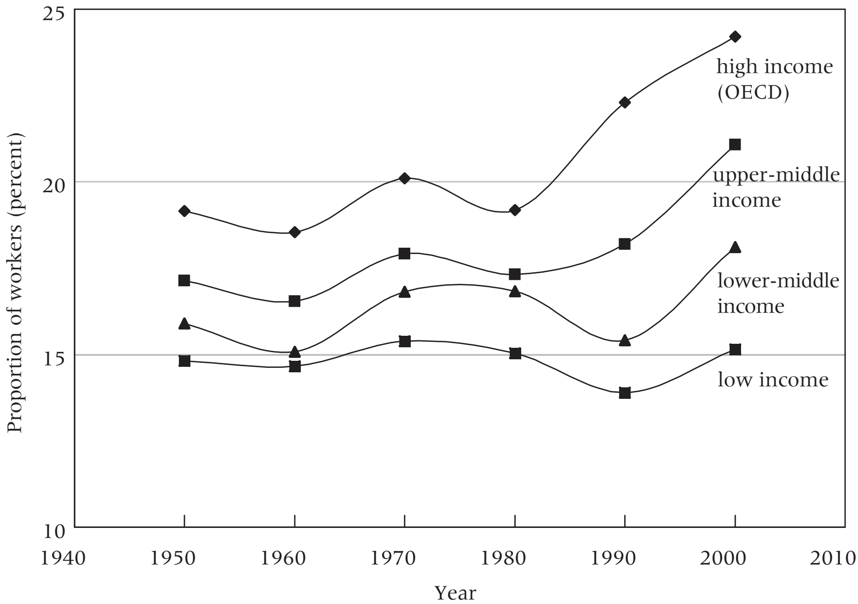

NOTE: Data are for 1995. The classifications are standard World Bank income classification codes. High income: 25 OECD countries, median GDP per worker $39,000. Upper-middle income: 14 countries, median GDP per worker $17,000. Lower-middle income: 21 countries, median GDP per worker $12,000. Low income: 26 countries, median GDP per worker $3,000.
SOURCE: ILO «http://laborsta.ilo.org».

The most remarkable feature of this graph is the degree to which US and Japanese birth rates move in opposite directions.[10] During the peak of the US baby boom (around 1960), Japan was experiencing a local minimum in births. Japan had an upsurge in births during the mid-1970s as the United States was experiencing a significant slowdown in births. Consequently the Japanese workforce has very different demographic movements than that of the United States. Japan has a steeply rising cohort of workers in their 40s from 1960 to 1980, a period when the United States saw this cohort fall in size. From 1990 to 2000 the situation reverses.

The demographic effect roughly maps to the observed growth pattern between the United States and Japan. Between 1960 and 1980, the United States was experiencing worsening demographic structure—in the sense of a shrinking proportion of workers in their 40s—and low productivity growth. Figure 3 shows the demographic effect on productivity implied by the results presented earlier in the chapter.

The model suggests that 2–3 percent of the difference between US and Japanese growth in the 1970s is correlated with demographic shifts. In the 1990s, this situation reverses. The United States saw higher productivity growth attributable to demographic change while Japan experienced declining productivity growth. Age structure is associated with a 2 percent dif-

FIGURE 2 Annual number of live births in the United States and Japan, 1920–2000

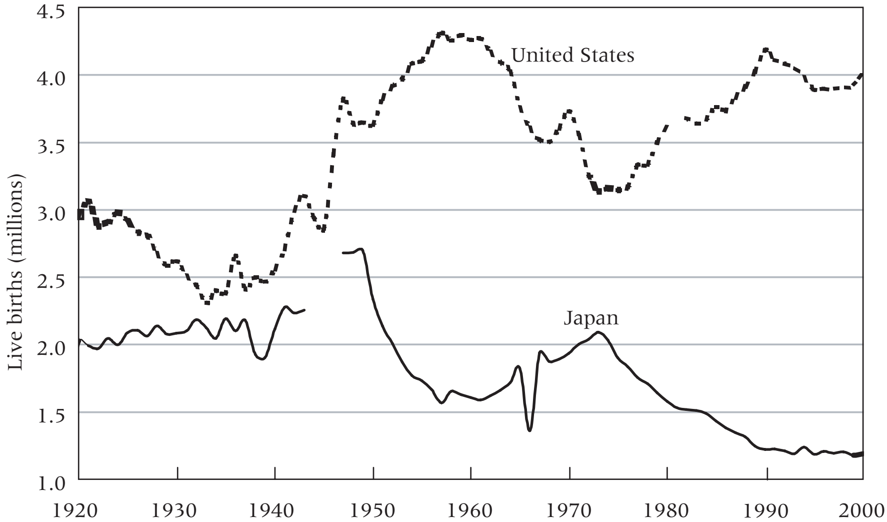

SOURCES: Vital Statistics of the United States; Japanese Ministry of Health, Labor and Welfare.

FIGURE 3 Estimated demographic effect on US and Japanese productivity growth, from the 1950s to the early 2000s

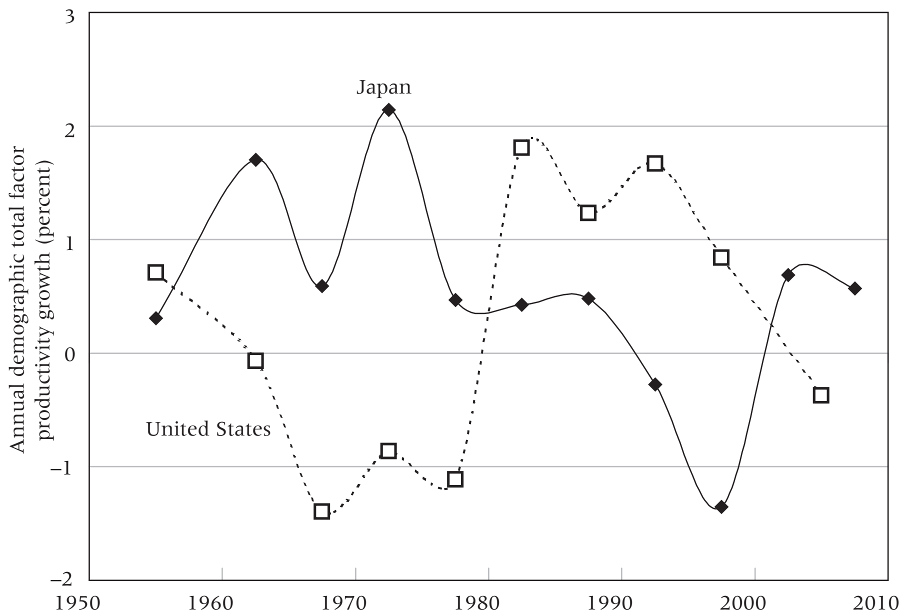

SOURCE: Author's calculations.

ferential between the United States and Japan during the 1990s. The model predicts that relative growth rates will reverse once again in the coming decade. The United States is about to enter a period of slower productivity growth while Japan should see a significant improvement in productivity growth relative to the 1990s.

Externalities

The cross-country results show that demographic structure has 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.

$$\log(wage) = \alpha + \beta_1 * school + \beta_2 * experience + \beta_3 * experience^2 + \varepsilon. \tag{5}$$

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

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

According to these estimates, an additional year of schooling increases the wage by 9.6 percent.[11] Experience has diminishing returns, with each additional year of experience increasing the wage by some amount less than 5.1 percent.

Worker productivity rises with age up to about age 50 then falls somewhat. The Mincer evidence implies that there is about a 60 percent difference between the productivity of 20-year-old workers and 50-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 a large cohort of older workers.[12]

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

The Mincer evidence may not tell the entire story at the aggregate level, however. The micro evidence, based on wage data, only captures the private return to experience and education. The social returns may be higher than the private returns. Externalities to experience (or age) may mean that the Mincer coefficients understate the aggregate productivity effects.

The 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 empirically that these externalities exist. 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. Some recent work has had more success in finding external effects of education. Aiyar and Feyrer (2002) find evidence of dynamic externalities to human capital that act over long time periods. Moretti (2004) and Bloom, Hartley, and Rosovsky (2006) find evidence of externalities to higher education at the state and city level.

The results of this chapter thus far suggest that externalities to workforce demographic composition go beyond the private return to experience. The next two sections suggest two possible channels through which social returns to 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 nonrival aspect of ideas increases the potential for large externalities.

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, nonrival. Nonrivalry 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 the productivity of anyone else who relies on the internet for productivity-enhancing information.[13] Yet, most people have never paid any money to Google. While the creators of Google have benefited 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 that increases the likelihood of Google being invented, productivity will be higher for all workers. Only a 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 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 25 and 35.[14] 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 has 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 (2005) find that the peak years for Nobel Prize–winning economists tend to be in their 40s. Jones (2005) collects the birth dates for a sample of inventors granted patents in the NBER patent database. Figure 4 presents 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 the creativity profile of inventors may be quite stable with a peak somewhere in the mid-40s. 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.

This argument is essentially one of scale. More potential inventors equals more invention. While scale effects in the production of nonrival goods seem undeniable on theoretical grounds, the evidence of them in the aggregate data is harder to find. Jones (1995) argues that productivity growth has not increased despite the increase in the number of researchers. His argument relies on the decreasing returns in the search for new knowledge. Kremer (1993) suggests that over longer time horizons there is strong evidence of scale effects in knowledge production.

Over the shorter time horizons used in this study, it is not clear that raw knowledge creation explains the aggregate results, especially outside the OECD where knowledge adoption is more important. The following section examines how idea adoption might be affected by demographic change.

FIGURE 4 Distribution by single years of age of US inventors granted patents, 1975–95

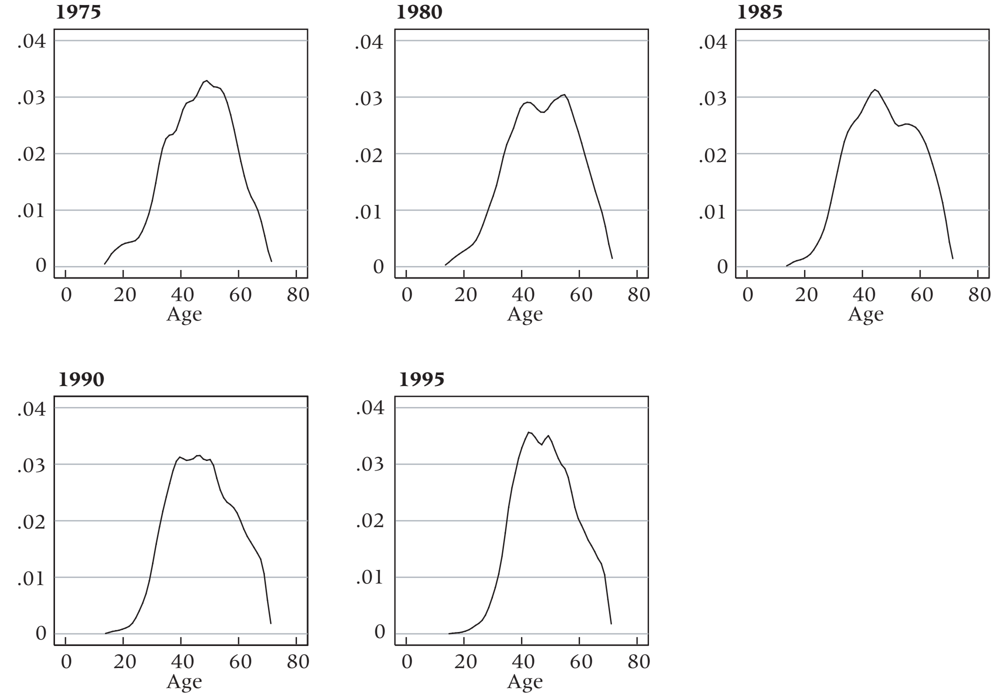

SOURCE: Jones (2005).

Idea adoption

While creative output is one potential channel through which age may affect productivity, it may not be the most relevant for cross-country comparisons. For most of the countries in the world, idea creation matters less than idea adoption. Organizations (or countries) that increase productivity by producing new ideas are different from 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 affects 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 matter 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.[15] 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.

Large demographic shifts may also matter through the effect on the quality of management. 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,

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

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, there must be 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 probably 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 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 there is a need to dip farther into the talent pool of the older cohorts. The Lucas model suggests that less talented managers will make all workers less productive. Indeed, one need not rely on this specific model to accept the basic argument. Any model where the quality of management has spillovers for all workers will produce similar results.

An examination of census data suggests that the entrance of the baby boom cohorts into the US workforce caused significant changes in the age structure of the management of US firms. Figure 5 shows the evolution of the age distribution of managers in the United States over time against the evolution of the workforce as a whole. The latter is shown by the dark line of the age distribution. The dashed line is the age distribution of US workers categorized as managers.

The baby boom cohorts first entered the workforce in large numbers in the 1970 census, but they were not well represented in the management workforce in that year. This is consistent with the idea that young workers are not chosen to be managers, because of their 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 the workforce in large numbers. By 1980 the baby boom cohorts had fully entered the workforce

FIGURE 5 Distribution by single years of age of the entire US workforce and of those categorized as managers, 1960–2000

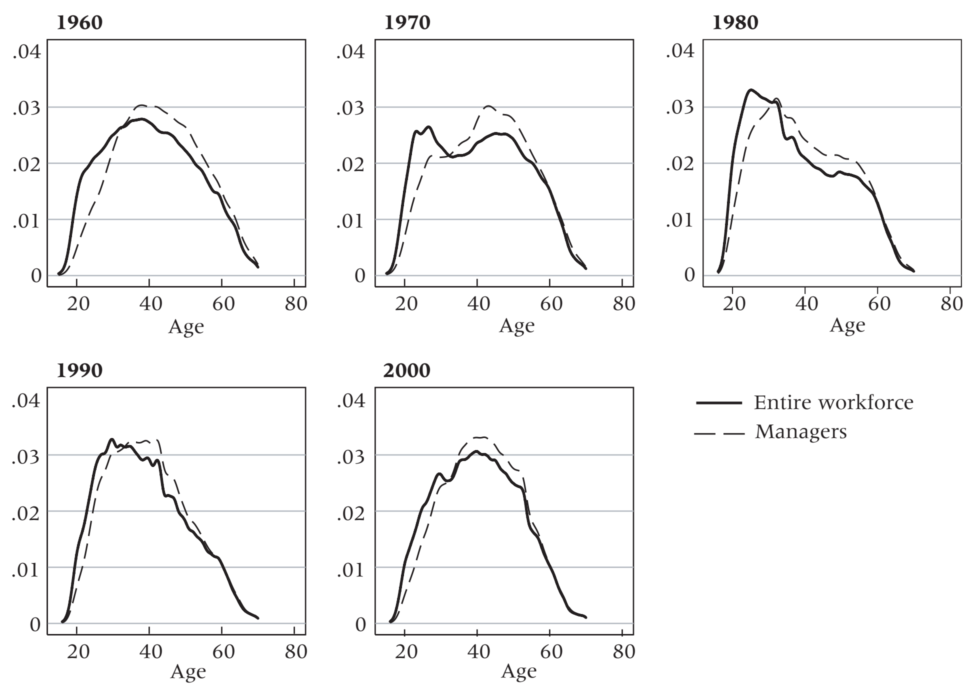

SOURCES: IPUMS; author's calculations.

but were still quite young and were 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 between 1970 and 1980 from 43 years old to 39 years old. As the mass of the boomers entered their 30s in the 1990 census, the managerial workforce began to return to its earlier shape. By 2000, when the boomers were 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 have 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 managers (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 boomers aged, they began taking over the management burden generated by the size of their cohort. Between 1980 and 2000, this resulted

FIGURE 6 Proportion of US workforce categorized as managers (detrended), by age group

SOURCE: IPUMS, author's calculations.

in a lowering of the proportion of managers in each age cohort (relative to the time trend).

The data suggest 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. Higher manager quality will cause higher overall productivity, potentially contributing to the aggregate results.

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 that is confined to the country level. This section attempts 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 are taken from the Integrated Public Use Microdata Series of the US census.[16] The base data is a 5 percent sample of the US census from 1960 until 2000. Population proportions 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.

While hourly income is not a perfect measure of total factor productivity, it is the best available proxy for this purpose. If we assume that workers earn their marginal product and that capital is mobile within the United States, the regional differences in wages should reflect productivity differences. The use of wages in studies of schooling externalities is common (see, e.g., Moretti 2004).

Estimation is identical to that in 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 United States high participation rates among teenagers may be correlated with unobservable area characteristics that drive wages downward. This source of bias can be addressed by instrumenting the workforce proportions on population proportions.

Within the United States, 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 3 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 age structure as instruments. This estimation deals with the problem of endogenous participation rates. Columns (3) and (6) estimate using IV with ten-year lagged age structures as instruments.

These results are not nearly as clear as the cross-country results and are inconclusive. Point estimates and almost all significant coefficients are negative.[17] In general, the point estimates suggest that the effect of changes in the age structure is smaller at the state and MSA level than at the cross-country level. In the case of the state-level results, we cannot reject that the effect of changes in the age structure 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. All three

TABLE 3 Effect of changes in workforce composition on US wages: State and metro area, 1960–2000

	Δlog(wage)					
Sample:	**(1) State OLS**	**(2) State IV**	**(3) State IV-lag**	**(4) MSA OLS**	**(5) MSA IV**	**(6) MSA IV-lag**
ΔW15	–3.497	–3.206	–7.253	–1.967	–5.402	1.608
	(1.855)+	(11.395)	(5.806)	(0.620)**	(2.132)*	(7.551)
ΔW20	0.032	0.19	–0.827	0.273	0.529	0.388
	(0.406)	(0.826)	(1.523)	(0.184)	(0.254)*	(0.586)
ΔW30	–0.743	–0.943	–1.536	–0.135	–0.315	–0.521
	(0.546)	(1.281)	(1.194)	(0.184)	(0.308)	(0.823)
ΔW50	–1.357	–4.371	–4.773	–0.007	–0.759	–0.04
	(0.780)+	(1.229)**	(1.417)**	(0.226)	(0.443)+	(0.809)
ΔW60	–0.354	12.76	1.195	0.157	–0.451	2.88
	(1.243)	(11.843)	(1.947)	(0.327)	(2.480)	(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 percent; * significant at 5 percent; ** significant at 1 percent
NOTE: ΔW15 is the change in the proportion of workers aged 15–19. ΔW20, W30, W40, W50 are the changes in the proportion of workers ages 20–29, 30–39, 40–49, and 50–59. ΔW60 is the change in the proportion of workers ages 60 and older. MSA = metropolitan statistical area.

regressions at a lower level of aggregation, the MSA level, have coefficients that are significantly closer to zero than the cross-country point estimates.

These results suggest that the demographic effects diminish at lower levels of aggregation. This should not come as a surprise. The United States 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 by 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 as a result of lower marginal costs of production. Obviously, these consumers are not necessarily located in the same geographic area as the workers and plants.

Conclusion

The results presented in this chapter show that changes in workforce age structure are strongly correlated with productivity and output. A significant portion of the productivity gap between rich and poor countries may be related to different age structures. The results also appear to capture some of the productivity divergence between poor and rich countries since 1980. In Japan and the United States over the last 40 years, the relative demographic movements are consistent with the cross-country results in productivity changes.

Given the importance of productivity in explaining cross-country income differences, this is a useful result. Demographic changes have substantial predictable time series variation that is largely exogenous to contemporaneous events, at least at the country level. Also, the regressions using lagged demographic data indicate that movements in productivity are not causing contemporaneous changes in age structure. 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 mechanisms through which the age distribution might affect aggregate output. First, the productivity of innovative activity is undoubtedly related to age. However, US patent data show that the age distribution of innovators did not change substantially in the United States as a result of the entry of the baby boom cohorts into the workforce. This suggests that changes in the supply of workers who 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 managers in the United States. Initially the baby boomers were inexperienced and could not provide their own management talent, necessitating the use of less talented managers from older cohorts. As the boomers aged, they entered management ranks earlier than previous cohorts. This had the net effect of increasing the proportion of managers drawn from all age cohorts from 1960 to 1980, almost certainly lowering management quality. This trend reversed from 1980 through 2000.

Results at the US state and MSA level, while less conclusive than the cross-country results, suggest that the effect of age structure is smaller at 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 United States, this is not surprising. The nonrivalry 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.

These explanations for the aggregate results are hardly exhaustive. This chapter has focused on the direct impact of the age profile on production. It is also possible that certain market demand effects matter. For example, a particular age profile might result in consumption patterns that have aggregate effects. Taxation patterns may also differ across age profiles.

Understanding the relationship between age structure and productivity is important because of the useful and predictable characteristics of age structure and because the significance of the relationship is strong. Almost every region in the world is experiencing significant demographic change. Rich countries are rapidly becoming older and most have birth rates below replacement level. Some poor countries are experiencing dramatically reduced birth rates in the wake of rapid population growth. Understanding how these changes will affect productivity over the coming decades is of crucial importance. While this study indicates a relationship between productivity and age structure, more research is needed to understand the mechanisms behind this relationship and their strengths.

Notes

I am grateful to David Bloom, Peter Klenow, Doug Staiger, and Bruce Sacerdote for their helpful comments and advice.

1 I choose to exclude W40 because the 40-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.

2 Gollin (2002) shows that capital's share is roughly equal across countries.

3 The choice of coefficients follows Hall and Jones (1999), who in turn use returns to schooling data compiled in Psacharopoulos (1994). The present chapter differs in using data from a recent update in Psacharopoulos and Patrinos (2004). The differences are minor. For the first four years of schooling the return to schooling in sub-Saharan Africa, 11.7 percent, is used. For schooling from four to eight years the world average return to schooling, 9.7 percent, is used. For schooling beyond eight years the OECD return to schooling, 7.5 percent, is used. The results are not sensitive to the precise method of calculating human capital from schooling.

4 This correction is taken from UN national accounts data, as collected in Aiyar and Feyrer (2002). Because the regressions in this chapter exclude oil-exporting countries, the

corrections are quite minor and have very little impact on the results.

5 Their calculations, in turn, are based on the Penn World Tables 5 and 6. Both are available from the World Bank website «http://www.worldbank.org/research/growth».

6 The population 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.

7 Demographic shifts of this size, while not the norm, are present in the data. Between 1980 and 1990, the proportion of workers in the United States aged 40–49 rose by 4.6 percentage points.

8 The 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 5–54 since they will comprise the working-age population 10 years later.

9 Figure 2 shows the raw number of births and is not scaled to the size of the population. The important features of the figure are the locations of the peaks and troughs, which are more easily seen in the unscaled graph.

10 In 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.

11 While 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.

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

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

14 Lehman (1953), p. 8.

15 Chari and Hoenhayn (1991) find that technologies diffuse slowly because technology is embedded in long-lived capital.

16 «http://www.ipums.org».

17 The exception to this is the positive and significant coefficient on 20–29-year-old workers in regression (5). Recall that one of the problems in using US state and MSA data is migration. If 20–29-year-olds are differentially likely to migrate to high-income areas, we should expect a larger coefficient on the 20–29-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 20–29-year-old workers (albeit with a large standard error).

References

Acemoglu, Daron and Josh Angrist. 2000. "How large are human capital externalities? Evidence from compulsory schooling laws," *NBER Macro Annual* 2000: 9–59.

Aiyar, Shekhar and James Feyrer. 2002. "A contribution to the empirics of total factor productivity," Dartmouth College Working Paper, August.

Barro, Robert J. and Jong-Wha Lee. 2001. "International data on educational attainment: Updates and implications," *Oxford Economic Papers* 53(3): 541–563.

Bils, Mark and Peter Klenow. 2000. "Does schooling cause growth?," *American Economic Review* 90 (5): 1160–1183.

Bloom, David E., David Canning, and Jaypee Sevilla. 2004. "The effect of health on economic growth: A production function approach," *World Development* 32 (1): 1–13.

Bloom, David E., Matthew Hartley, and Henry Rosovsky. 2006. "Beyond private gain: The public benefits of higher education," in James J. F. Forest and Philip G. Altbach (eds.), *International Handbook of Higher Education*. Springer, pp. 293–308.

Caselli, F., G. Esquivel, and F. Lefort. 1996. “Reopening the convergence debate: A new look at cross-country growth empirics,” *Journal of Economic Growth* 1(3): 363–389.

Chari, V. V. and Hugo Hoenhayn. 1991. “Vintage human capital, growth, and the diffusion of new technology,” *Journal of Political Economy* 99(6): 1142–1165.

Easterly, William and Ross Levine. 2001. “It’s not factor accumulation: Stylized facts and growth models,” *World Bank Economic Review* 15(2): 177–219.

Feyrer, James. 2007. “Demographics and productivity,” *Review of Economics and Statistics* 89(1): 100–109.

Galenson, David W. and Bruce Weinberg. 2000. “Age and the quality of work: The case of modern American painters,” *Journal of Political Economy* 108(4): 761–777.

———. 2005. “Creative careers: The life cycles of Nobel Laureates in economics,” NBER Working Paper series no. w11799. Cambridge, MA: National Bureau of Economic Research.

Gollin, Douglas. 2002. “Getting income shares right,” *Journal of Political Economy* 110: 458–475.

Hall, Robert and Charles I. Jones. 1999. “Why do some countries produce so much more output per worker than others?,” *Quarterly Journal of Economics* 114(1): 83–116.

Jones, Benjamin F. 2005. “Age and great invention,” NBER Working Paper 11359, May.

Jones, Charles I. 1995. “R&D-based models of economic growth,” *Journal of Political Economy* 103 (4): 759–784.

Jorgenson, Dale. 2003. “Information technology and the G7 economies,” *World Economics* 4(4): 139–169.

Kremer, Michael. 1993. “Population growth and technological change: One million B.C. to 1990,” *Quarterly Journal of Economics* 108(3): 681–716.

Lehman, Harvey C. 1953. *Age and Achievement*. Princeton University Press.

Lucas, Robert E. 1978. “On the size distribution of business firms,” *The Bell Journal of Economics* 9(2): 508–523.

Mincer, Jacob. 1974. *Schooling, Experience, and Earnings*. New York: Columbia University Press.

Moretti, Enrico. 2004. “Estimating the social return to higher education: evidence from longitudinal and repeated cross-sectional data,” *Journal of Econometrics* 121: 175–212.

Nelson, Richard and Edmund Phelps. 1966. “Investment in humans, technological diffusion, and economic growth,” *American Economic Review* 56: 69–75.

Psacharopoulos, George. 1994. “Returns to investment in education: A global update,” *World Development* 22(9): 1325–1343.

Psacharopoulos, George and Harry Anthony Patrinos. 2004. “Returns to investment in education: A further update,” *Education Economics* 12(2): 111–134.

Weinberg, Bruce A. 2002. “Experience and technology adoption,” Ohio State University working paper, May.