The US Productivity Slowdown, the Baby Boom, and Management Quality *

James Feyrer†
Dartmouth College

June 13, 2008

Abstract

This paper examines whether management changes caused by the entry of the baby boom into the workforce explain the US productivity slowdown in the 1970s and resurgence in the 1990s. Lucas (1978) suggests that the quality of managers plays a significant role in determining output. If there is heterogeneity across workers and management skill improves with experience, an influx of young workers will lower the overall quality of management and lower total factor productivity. Census data shows that the entry of the baby boom resulted in more managers being hired from the smaller, pre baby boom cohorts. These marginal managers were necessarily of lower quality. As the boomers aged and gained experience, this effect was reversed, increasing managerial quality and raising total factor productivity. Using the Lucas model as a framework, a calibrated model of managers, workers, and firms suggests that the management effects of the baby boom may explain roughly 20 percent of the observed productivity slowdown and resurgence.

*I am grateful to Peter Klenow, Jon Skinner, Doug Staiger, Jay Shambaugh and participants at the Population Aging and Economic Growth conference at the Harvard School of Public Health 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.

Introduction

Annual productivity growth fell by approximately one percentage point in the 1970s relative to the preceding decade. It was not until the mid 1990s that productivity growth returned to the level of the 1960s. The causes of the slowdown and recovery have been much debated. Interestingly, discussions of the resurgence have often failed to talk about what caused the slowdown in the first place. This paper will investigate one potential cause of both the slowdown and subsequent recovery, demographic change. The channel through which demographic change will impact productivity is management quality.

The slowdown in the US roughly corresponds to the entry of the baby boom into the workforce. The recovery occurred as the mass of the baby boom entered their prime working years. Feyrer (2007) finds that there is a strong and robust correlation between relative cohort sizes and total factor productivity. In particular, the proportion of workers aged forty and older is positively correlated with total factor productivity. The magnitude of the effect is an order of magnitude larger than one would expect from the returns to experience measured at the micro level. In other words, the social return to experience appears to be much larger than the private return.

This paper looks to management quality as one possible source of externalities to experience. Lucas (1978) proposes a model where the quality of the manager of a firm maps directly into firm output differences. Given the same quantity of labor and capital, firms with more talented managers will produce more output. The effectiveness of management is also affected by the scope of the enterprize. Given a firm with a particular manager, there are decreasing returns to the scale of the firm.

If we assume some distribution of managerial talent those with the most managerial talent will take managerial positions and those with less talent will become workers. The heterogeneity in management talent combined with decreasing returns to scale results in tension between firm size and managerial quality. In order to reduce the size of firms, additional managers are needed. The additional managers will be lower quality than the existing managers.

Given this structure, it should be clear that an influx of workers that have low levels of managerial talent will require changes on both margins. The number of workers per manager will rise for existing managers – reducing their effectiveness – and some workers will be drawn into the ranks of management. These marginal managers will obviously lower the overall talent of management.

The entry of the baby boom into the US workforce presents us with precisely this situation. When the baby boomers entered the workforce, they were not immediately useful as managers, either through a lack of experience or institutional restrictions on younger workers managing older workers. The managers for these new workers necessarily came from the smaller and older cohorts of workers. This dynamic is evident in the data.

An examination of census data shows that the entrance of the baby boom into the US workforce caused significant changes in the age structure of management. First, workers in cohorts born before the boom were drawn into management in larger numbers. Second, workers in the baby boom cohort were called upon to manage at earlier ages than in previous generations. This implies that from 1960 until 1980, firms had to go deeper into the management distribution causing a fall in average management quality. This trend reverses by 2000. This implies falling management quality from 1960 until 1980 and rising quality thereafter.

The timing of this change in management structure roughly parallels the slowdown in US productivity growth. This paper examines how important the demographic change in management quality was compared to the overall US productivity slowdown. Using the Lucas model as a framework, we calibrate a model of managers, workers and firms and examine how changes in the observed age distribution of workers in the US affects productivity. The results suggest that the management effects of the baby boom may have caused productivity growth in the 1990's to be 0.15-0.25 percentage points higher than in the 1970s. This represents roughly 20% of the observed slowdown.

1 Background

Productivity growth slowed in the US and other industrialized nations during the 1970's and remained below the earlier trend until the mid 1990's. Figure 1 shows yearly growth in output per hour in the nonfarm business sector for the US.[1] The yearly data are overlaid with a five year moving average. The fall in average growth rates from 1975 until 1995 is obvious.

There has been enormous debate about the causes of the US productivity slowdown and the recent resurgence in productivity growth. Interestingly, these two events are treated somewhat separably in the literature. In a 1988 symposium on the productivity slowdown in the *Journal of Economic Perspectives*, Fischer (1988) suggests that the main causes were oil price increases, as argued by Jorgenson (1988) in the same issue, and a slowdown in the rate of new technology production. Griliches (1988) expresses skepticism of the latter view. Interestingly, the productivity resurgence debate is a bit disconnected from the slowdown literature. Debate has largely revolved around the role of computer technology.[2]

[1] (Lazear, Baicker and Slaughter 2007)

[2] See Nordhaus (2002), Gordon and Sichel (2002), Oliner and Sichel (2000), and Fernald, Thipphavong and Trehan (2007) among many others.

Figure 1: Growth in US Labor Productivity 1960-2005

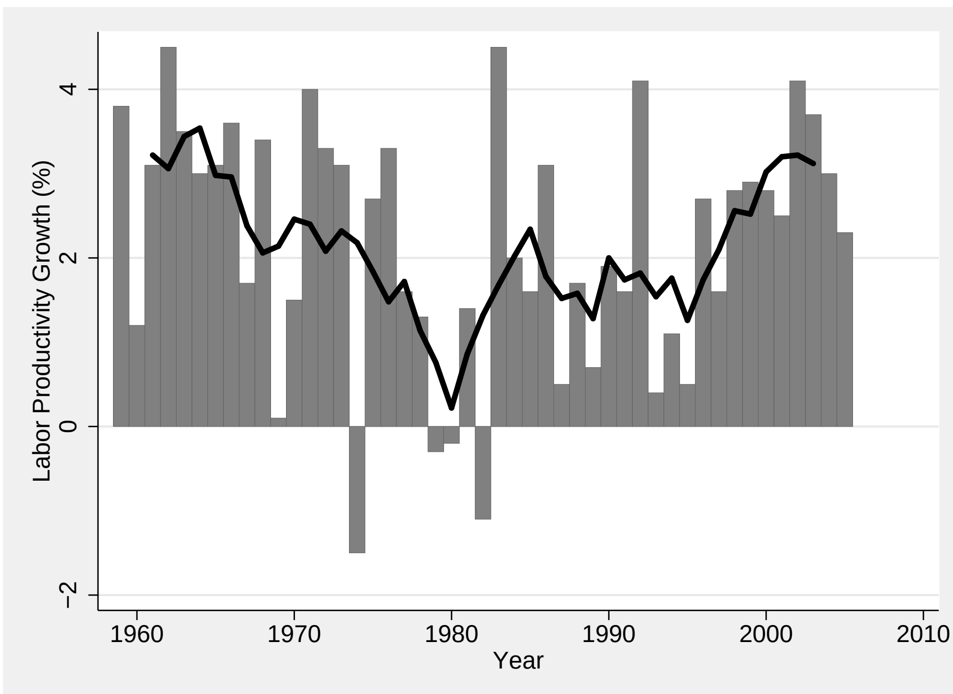

Gray bars are yearly observations. Five year moving average in bold.
source: The Economic Report of the President 2007.
Table B50 Growth in output per hour, nonfarm business sector.
http://www.gpoaccess.gov/eop/2007/B50.xls

This paper will examine one potentially overlooked determinant of both the productivity slowdown and resurgence, demographic change. To suggest that the entry of the baby boomers into the workforce lowered labor productivity is hardly a new idea. Standard Mincer regressions suggest that increases in experience result in higher wages. The entry of the baby boomers into the workforce in the 1970s certainly lowered average experience in the US. Insofar as simple labor productivity measures do not account for the experience of the workforce, this may be causing part of what we see in Figure 1. However, the impact of experience is relatively modest and even accounting for it does not change the basic picture very much. In addition, the baby boomers were more educated than their elders, offsetting the change in experience. Baily, Gordon and Solow (1981) summarize the adjustments one can make for demographics and finds that they generate relatively small effects.

It should be noted, however, that the standard Mincer regressions are based on private returns to experience and will not be capable of detecting externalities to a more experienced workforce. Feyrer (2007) and Feyrer (2006) show that the effect of demographic change on aggregate output may be much larger than suggested by the private returns to experience.[3] The impact of an increase in the proportion of experienced workers (aged forty plus) is found to have an effect an order of magnitude larger than expected from the private return to experience. This paper is an attempt to suggest a mechanism through which changes in experience may generate large externalities.

One way in which the externalities to experience may matter is through manage-

[3]Other work has also found a relationship between demographic change and output. Focusing on the dependency ratio, Bloom, Canning and Sevilla (2001) find that increases in the size of the working age population can produce a "demographic dividend" to economic growth. Kogel (2005) finds a relationship between total factor productivity and the dependency ratio. Persson (2002) finds that the age structure of US states affects output. Sarel (1995) finds a significant effect of the age structure of the population on output in a cross section of countries. Bloom, Freeman and Korenman (1988) find that being a member of a large cohort leads to lower lifetime earnings.

ment quality. It is not unreasonable to think that lowering management quality will have effects on productivity that go beyond the change in private returns. Good management may also be important for the implementation of new technologies. There is microeconomic evidence that age matters in the adoption of technology. Weinberg (2002) finds that both experience and age matters for technology adoption. Since schooling tends to be concentrated early in life, young managers have the advantage of more recent human capital.[4]

Much of the current debate on the productivity resurgence revolves around whether it is due to the falling cost of information technology capital (essentially a capital deepening story) or through the ability of information technology to improve productivity in all sectors. If management quality matters, it may be that demographic shifts have had important effects on the ability of firms to implement computer technology. If so, the proximate cause of the productivity resurgence may be the implementation of IT, but the ultimate cause may be the enhanced ability of management to take advantage of the new opportunities.

1.1 Changes in the US Workforce Age Structure

The baby boom in the United States generated significant changes in the age distribution of the population. Figure 2 shows the number of live births in the US from 1920 until about 2000. Peaking in the late 1950's, the baby boom cohorts were significantly larger than those preceding and following. Figure 3 shows the resulting changes in the age distribution of the US workforce over time. The proportion of twenty year olds in the workforce rose from 20 percent in 1960 to 30 percent in 1980 as the baby boom enters the workforce. In each ten year interval, this effect repeats

[4] Chari and Hoenhayn (1991) find that technologies diffuse slowly due to vintage human capital effects.

Figure 2: Live Births in the US

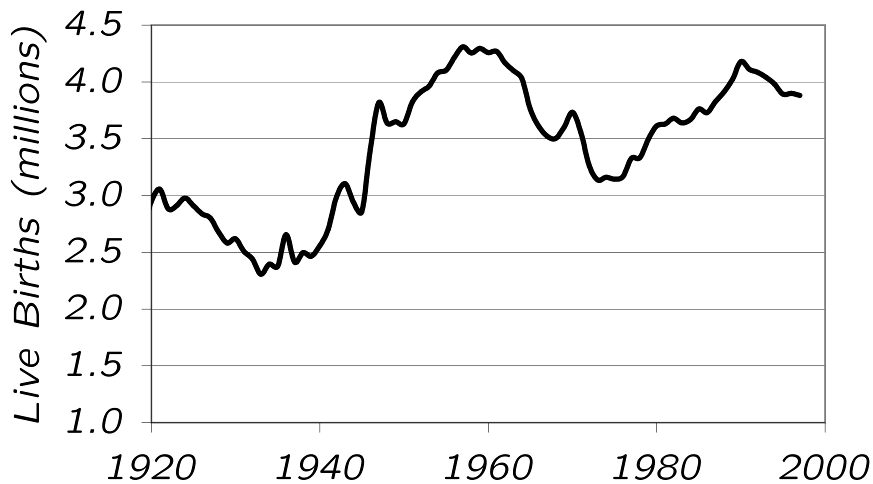

source: National Center for Health Statistics
http://www.cdc.gov/nchs/fastats/births.htm

Figure 3: The Baby Boom and the Age structure of the US workforce

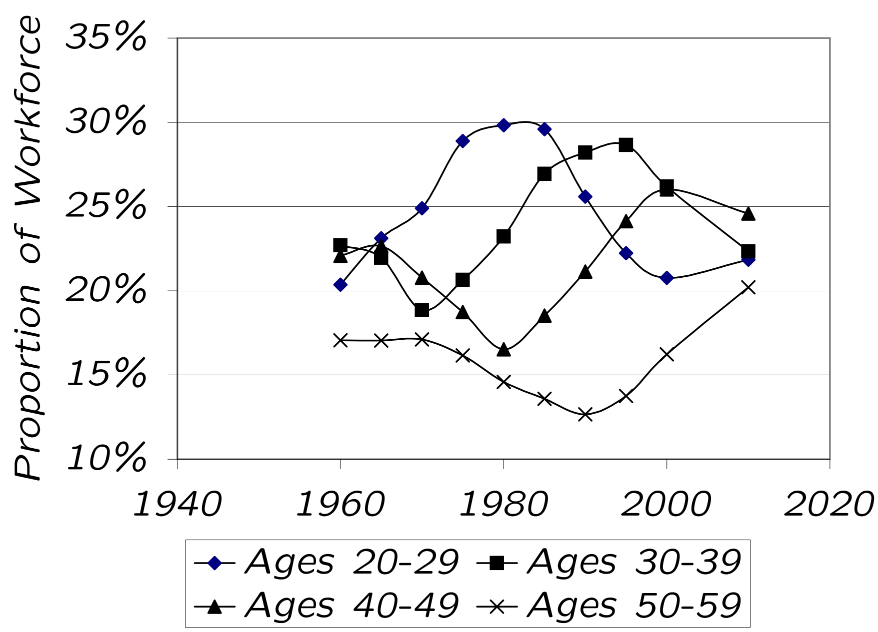

source: IPUMS, author's calculations

in the next age category. The number of thirty year olds peaks in 1990-1995 and the number of forty year olds peaks in 2000.

Figure 4 shows the age distribution of full time workers in the US overlaid with the age distribution of workers classified as managers for each census between 1960 and 2000. The underlying data for these histograms are from the Integrated Public Use Microdata Series of the US census.[5] The entire workforce distribution is the solid line while the managerial workforce distribution is dashed. At the left hand side of each distribution the dashed line falls below the solid line, illustrating that workers in their twenties are managers at lower rates than older cohorts. In 1970, the entry of the baby boomers into the workforce is very apparent in the worker distribution, but barely discernable in the manager distribution. By 1980, however, the presence of the baby boom in the managerial workforce is obvious. Figure 5 shows the evolution of the mean and median ages of US managers over time. Over the decade of the 1970s the median age of managerial workers falls by five years. Between 1980 and 2000 the baby boom begins to have enough workers in the over thirty distribution to manage itself. By 2000 the two distributions have returned to a pattern very similar to 1960.

If the baby boomers were entering the non-managerial workforce in large numbers by 1970, who was managing them? The short answer is the older, smaller cohorts. Figure 6 shows the proportion of workers classified as managers by age group over time. This data has been detrended by time to remove secular movements in the proportion of workers classified as managers. From 1960 until 1980 the probability that any given person enters management rises for all age groups. These are the workers called into management ranks by the entry of the baby boom into the

[5] http://usa.ipums.org/usa/. For 1980, 1990 and 2000, the data are a 5% sample. For the earlier years a 1% sample is used. The sample is comprised of full time workers. Workers categorized as "Managers, Officials, and Proprietors" under the 1950 occupational coding are coded as managers.

Figure 4: Age Distribution of US Workforce categorized as managers, by year

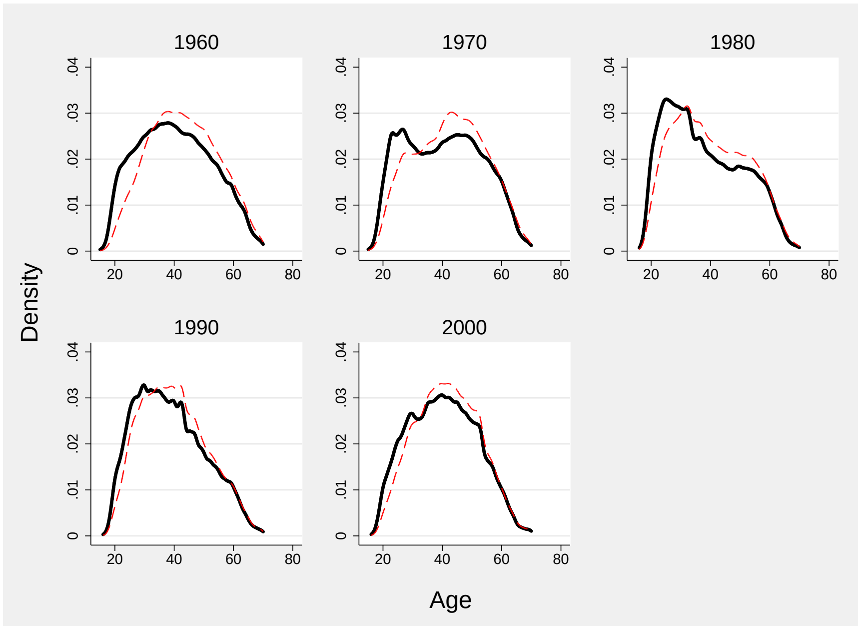

Solid Line - Entire Workforce, Dashed Line - Managers

source: IPUMS, author's calculations

Figure 5: The Mean and Median Ages of US Managers, 1960-2000

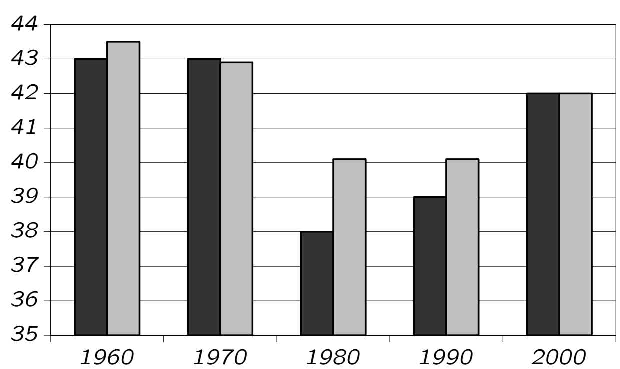

	1960	1970	1980	1990	2000
Median Age	43	43	38	39	42
Mean Age	43.5	42.9	40.1	40.1	42

source: IPUMS, author's calculations

workforce. From 1980 to 2000 this trend reverses as the baby boom itself provides the managers.

It should be noted that this is not due to a secular increase in workers classified as management – though this is also happening. As long as young cohorts manage at a lower rate than older cohorts, it is possible for all cohorts to see an increase in the proportion managing despite the management to worker ratio staying constant in the aggregate. Imagine that there are two groups of workers, the old and the young. The overall proportion of workers to managers is constant at 10 percent. In year one, there are an equal number of young and old workers, but only the old workers take on management roles, so 20 percent of the old cohort manages. In year two a large cohort enters such that the young cohort is now twice the size of the old cohort. If no one from the young cohort is allowed to manage, the proportion of the old cohort managing will need to rise to 30 percent in order to keep the aggregate proportions constant. The overall proportion of managers will remain constant despite the fact that the proportion of managers went up in the older group.

Figure 7 shows the change in management proportions over time. The overall numbers rise dramatically from 1960 until 1990 and subsequently fall. Nothing discussed so far gives an answer to whether the entry of the baby boom should *ceteris paribus* raise or lower the proportion of managers in the economy. It is almost certainly true that the large rise is due to secular changes in the US economy or changes in the way forms classify jobs. However, we will return to this question with the calibrated model. In particular, we will look for suggestions that the fall between 1990 and 2000 is related to the aging of the boomers.

Intuitively, the baby boomers were entering the workforce in large numbers and needed to be managed by *someone*. Since young workers tend to be managers in lower numbers than older workers, the additional managers needed to come from the

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

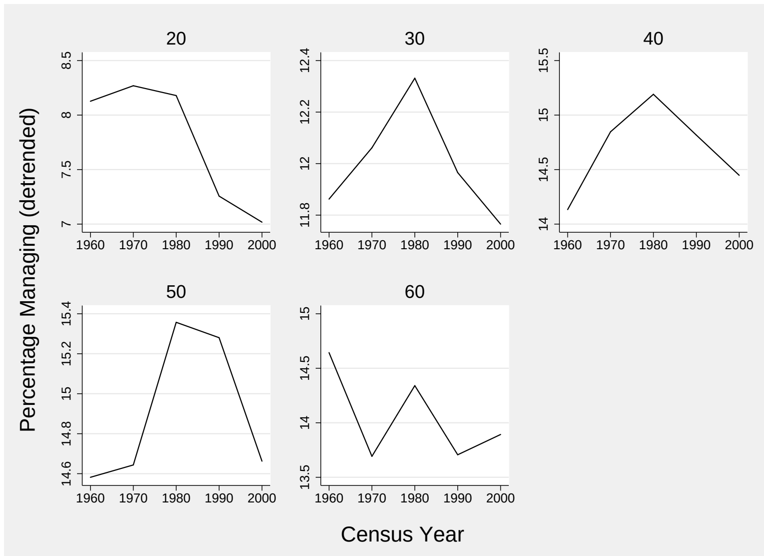

source: IPUMS, author's calculations

Figure 7: Proportion of US Workforce categorized as managers over time

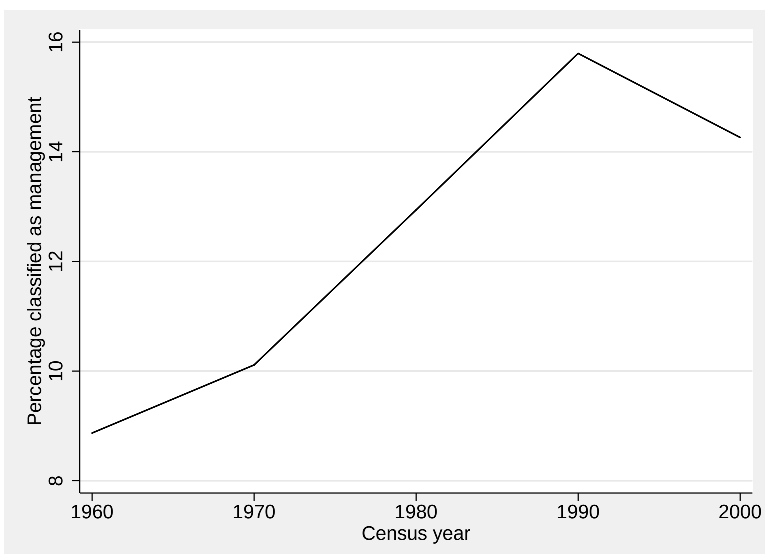

source: IPUMS, author's calculations

older and smaller cohorts. The additional managers are necessarily of lower talent than the existing managers, otherwise they would have been managing already. This leads to an overall drop in management quality and a fall in TFP. After 1980, when the baby boomers enter the years when workers normally enter management roles, the situation reverses. At this point the baby boom can manage itself and has a relatively small group of young workers to manage. This results in a fall in the proportion of 40 year old workers managing over time and an increase in the overall quality of management. The following section will make this argument more formally in an attempt to estimate how large these effects could be.

2 The Lucas Span of Control Model

This section presents a simplified version of the model from Lucas (1978). Each firm consists of a single manager and a number of homogeneous workers. There is heterogeneity in management talent and management talent has a multiplicative effect on firm output. There are decreasing returns in the number of workers employed by a firm.

More formally, a firm with a manager of quality x managing L workers and K units of capital will produce the following amount of output,

$$Y(x) = xg[f(K, L)] \tag{1}$$

where $f()$ is a standard neoclassical production function, and $g[]$ has decreasing returns. The decreasing returns to $g[]$ imply that larger firms will have lower per worker output given a fixed level of management talent, x. As the 'span of control' of a manager gets larger, their effectiveness diminishes.

Assume a heterogeneous distribution of managerial talent, $\Gamma(x)$. All individuals will be either managers or workers. Individuals with the highest level of management talent will become managers while workers of lower talent will be employed by firms. Smaller firms are more productive, but since each firm needs a manager, the additional managers needed to reduce firm size will be of lower quality. This tradeoff drives the equilibrium level of firm size and management quality. There will be a cutoff level of managerial talent, z, below which individuals will not be managers. If the total quantity of labor is normalized to one, the resource constraint on labor is

$$1 - \Gamma(z) = \int_z^\infty L(x) d\Gamma(x) \leq 1. \tag{2}$$

Total output is

$$Y = \int_z^\infty Y(x) d\Gamma(x). \tag{3}$$

The profit of each firm is equal to the output minus the cost of inputs.

$$\pi = Y - Kr - Nw \tag{4}$$

where r is the real return to capital and w is the real wage. Firms maximize profits subject to the first order conditions

$$\frac{dY}{dK} = r \tag{5}$$

and

$$\frac{dY}{dL} = w. \tag{6}$$

The production function $F(K, L)$ from (1) is taken to be Cobb-Douglas with capital share of $\alpha \in (0, 1)$. Lucas shows that if firm size is independent of firm growth rates (Gibrat's Law), the function g[] must take on the following simple parameterization,

$$Y = x(K^{\alpha}L^{1-\alpha})^{\beta}, \tag{7}$$

where $\beta \in (0,1)$ defines the degree of diminishing returns to firm size.[6] The first order conditions imply the following optimal capital labor ratio, which is identical for all firms, regardless of management talent, x.

$$\frac{K}{L} = \frac{\alpha}{1-\alpha}\frac{w}{r} \tag{8}$$

Combining (8) with either of the first order conditions results in the following optimal level of workers for any given level of management talent, x, wage rate, w, and rate of return to capital, r.

$$L^* = L(r,w,x) = \left[(1-\alpha)\beta\left(\frac{\alpha}{1-\alpha}\right)^{\alpha\beta} r^{-\alpha\beta}w^{\alpha\beta-1}x\right]^{\frac{1}{1-\beta}} \tag{9}$$

The number of workers for any given firm is increasing in the talent of management and decreasing in the rental rate of capital and the real wage. Combining the profit function with (8) and (9) you can solve for the profit level of a firm with management talent x in terms of the real wage and real return on capital. For simplicity we can assume that the real rental rate of capital is exogenously determined (US firms can borrow in world markets). The wage, however is going to be endogenously determined subject to the resource constraint (2). Consider the choice of a individual between working for wages at an existing firm and managing their own firm. They will choose to manage only when the profit of the firm will exceed their wages as a worker. There will be some cutoff level of management talent where a worker is

[6] Lucas (1978), p. 514-515

indifferent between managing and working.

$$\pi(z) = w \tag{10}$$

Where this cutoff falls will be a function of the distribution of management talent, $\Gamma(x)$, and the resource constraint (2). The following section will outline the parameterization of this basic model and the algorithm for a numerical solution.

3 A Numerical Solution

In this section we will take the basic structure outlined above and describe the steps needed to generate numerical estimates of the behavior of the US economy based on the model. To simulate the model we first must specify the distribution of managerial talent. The basic method is to generate a large number of discrete agents and aggregate. These agents are distributed into groups by age, with the number of agents from each age group reflecting the observed workforce proportions in the US (see Figure 3). Each age group has a different managerial talent distribution, $\Gamma_i(x)$. The overall distribution of management talent, $\Gamma(x)$ is generated through summing the distributions for each age group.

Each of these agents has a managerial talent, x, taken from an age specific distribution of talent.[7] Each agent takes the rental rate of capital, r, and the wage rate, w, as given. The output for a firm headed by a manager of talent level x can be calculated from (7) subject to the profit maximizing conditions (5) and (6). Equations (9) and (8) give the profit maximizing levels of labor and capital.

[7]These talent levels are deterministic and not stochastic. The solution technique should therefore be seen as a numeric integration rather than a Monte Carlo exercise. To be more concrete, if there were 99 agents, the first would be assigned the first percentile cutoff of the distribution, the second the second percentile and so on to the 99 percentile.

Plugging these values into (4) gives the profit level of a firm headed by a manager of talent x. If the profits are greater than the wage, the firm is viable and will operate with $L(x)$ employees and one manager with the manager receiving the profits. If the profits are less than the wage, the agent will work for one of the viable firms.

Aggregate demand for workers at any given wage w can be determined by summing $L(x) + 1$ over the agents where $\pi(x) \geq w$. The solution of the model requires finding the wage which equalizes supply and demand for workers. As the wage increases, the labor demand will fall along two margins. First, a higher wage results in a lower number of firms, since the profit threshold is higher. Second, any firm which remains above the profit threshold will employ fewer workers. There is therefore a downward sloping labor demand function with a fixed supply.

The equilibrium wage is the wage for which the total desired employment is equal to the total number of agents in the economy. This can easily be found numerically. Once the equilibrium wage, w^*, has been determined we can calculate total output by summing up individual firm output. Other metrics are also easily calculated including the proportion of managers from each age cohort and the profit distribution of managers.

3.1 Parameterization of the Model

To complete the parameterization of the model we need to specify the coefficient on capital in the production function, α, the degree of decreasing returns at the firm level, β, the exogenous rate of return to capital, r, and the age specific distributions of management talent.

The exogenous rate of return to capital is set to the risk free rate of return, 5%. We can rely on relative income shares to determine the production function parameters, α and β. It can easily be shown that the following shares of total income hold.

$$\text{CAPITAL'S SHARE} = \alpha \times \beta \tag{11}$$

$$\text{NONMANAGERIAL LABOR'S SHARE} = (1-\alpha) * \beta \tag{12}$$

$$\text{MANAGERIAL LABOR'S SHARE} = 1-\beta \tag{13}$$

We choose α and β to produce income shares that roughly match the actual economy. We set capital's share of income to match the standard figure from national income and product accounts, 40 percent. Piketty and Saez (2003) show that the wage share of the top 10 percent of wage earners ranges from 25-35 percent over the period of interest. Since the proportion of managers in the dataset is 10-15 percent, this implies that the share of labor income going to managers should be roughly in the range of 35-45 percent. The share of labor income going to managers is

$$\text{MANAGER'S SHARE OF LABOR INCOME} = \frac{1-\beta}{1-\beta\alpha}. \tag{14}$$

Given a share of capital in total income and a managerial share of labor income appropriate values of α and β can be calculated from (11) and (14).

To generate the distribution of management talent, agents in the economy are split across five age groups, 15-29, 30-39, 40-49, 50-59 and 60 and older. The number of agents from each group is generated such that the proportions match the data. The variation in these proportions over time will generate our results. Within each age group the distribution of managerial talent is Pareto, a distribution commonly associated with the distribution of income. Since firm profits will be proportional to management talent, a Pareto distribution of management talent will tend to generate similarly distributed management income.

The Pareto distribution has pdf

$$pdf = \frac{kx_m^k}{x^{k+1}}. \tag{15}$$

The shape parameter of the distribution, k, is common across all distributions and is chosen so that the proportion of managers in equilibrium roughly matches the observed proportions in the data overall given the values for α and β derived from income shares. The lower support of the distribution, x_m, will vary by age group. Differences in x_m result in proportional differences in the means of the resulting distributions, so a 10% difference between the x_m between two age groups shifts the average managerial talent for that group by the same amount. These shifts in the talent distribution will generate different proportions of managers from the various age groups. Since the forty year old cohort has the highest proportion of managers in the data, it is used as a benchmark with $x_m = 1$. The values of x_m for each of the other age groups are chosen so that the equilibrium proportion of workers in management from each age group roughly matches the data.

With these parameters in place the observed age proportions can be run through the model from 1960 until 2000. The only variables that change over time are these age proportions. The calibration parameters k, α, β, and x_m are all fixed as the model is marched forward through time. For each year the model is solved numerically and equilibrium wages, the proportion of managers in each group and total output can be determined. The size of the changes in output generated by these demographic shifts in the model should be informative about how large the effects may have been in the real economy.

Figure 8: Simulated Response of output per worker to Demographically Induced Management Changes (1960==100)

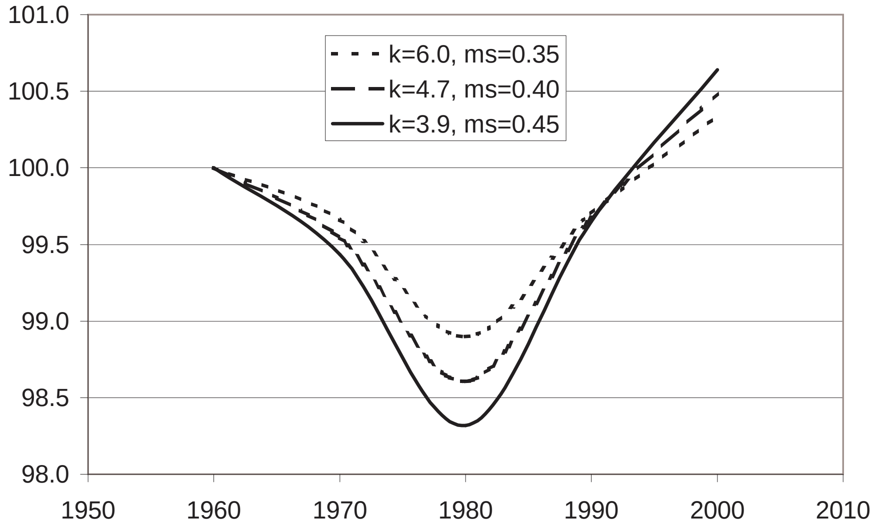

source: IPUMS, author's calculations
k is the shape parameter for the Pareto distribution.
ms is managements share of labor income.
Given ms, k is chosen to match the data illustrated in Figure 6.

4 Results

Figure 8 shows the progression of income from 1960 to 2000 in the simulated economy under three different parameter assumptions. The parameter sets were chosen by first varying management's share of labor income between 35 and 45 percent. The value of the Pareto distribution shape parameter, k, was then adjusted until the overall percentage of workers in management roles matched the data. The values for the lower supports of each distribution, x_m were then adjusted to get the relative managerial proportions of the age groups correct. Output per worker is falling from 1960 until 1980 with the more dramatic fall happening between 1970 and 1980. Since the simulated model has no underlying productivity growth this fall in the growth rate should be interpreted as a deviation from some constant underlying trend growth rate. Output per worker drops by 1 to 2 percentage points over this period. After 1980 this reverses with output per worker in 2000 higher than in 1960 under all parameterizations. Table 1 shows the yearly effect on growth rates simulated by the model under the three different parameters sets. The model

Table 1: Percentage Change in Growth Rate Relative to Trend

Year	ms=0.35 k=6.0	ms=0.40 k=4.7	ms=0.45 k=3.9
1960-1970	-0.034	-0.046	-0.057
1970-1980	-0.077	-0.095	-0.113
1980-1990	0.081	0.108	0.135
1990-2000	0.064	0.081	0.099

source: Author's calculations

generates a fall in yearly growth rates of 0.05-0.15 percent relative to trend from 1970 until 1980. From 1980 until 2000 the growth rates reverse and are 0.05-0.15 percentage points above trend. The simulated difference between growth rates in the 1970's and the 1990s is between 0.14 and 0.21 percentage points depending on the specific parameterization. This compares to a swing in the actual data of about one percentage point. The model is therefore generating effects that are 15 to 20 percent of the effect measured in the actual economy.

The most significant discrepancy with the model relative to reality is the timing. The simulated economy predicts an increase in labor productivity growth in both the 1980's and 1990's while the actual productivity slowdown continued into the 1980's (though the 1970's were certainly the low point). Part of the difficulty is that the census years match poorly with turns in the productivity series and the productivity series is quite noisy.

The model matches other moments of the economy reasonably well. Each of the parameter sets were chosen such that the proportion of managers from each group matches the detrended data from Figure 6. Figure 10 shows the proportions of each group in management generated by the model for the central scenario overlayed with the actual data. The model does a good job of matching the overall proportion of managers from each group and also generates about a one percent increase in the proportions managing from each group between 1960 and 1980. This compares quite well with Figure 6.

The average proportion of managers from each group is a direct consequence of the parameterizations and was intentionally designed to match the data. The changes over time, however, are entirely the result of changing cohort sizes, The fact that the changes over time match the data indicates that the model and parameterization are capturing some of the dynamics in a reasonable way.

The differences in average management proportions across the age groups are the result of adjusting the lower support, x_m of each group's distribution. These values are constant over time for each of the parameterizations. As mentioned earlier, the lower support is proportional to the mean for the group, so a 10% lower value for x_m lowers the group mean by a similar amount. Table 2 shows the values of x_m for each of the parameterizations. The parameters necessary to generate the proper

Table 2: Relative Managerial Talent, by age group

Age Group	ms=0.35 k=6.0	ms=0.40 k=4.7	ms=0.45 k=3.9
15-29	0.927	0.907	0.889
30-39	0.986	0.982	0.978
40-49	1.000	1.000	1.000
50-59	0.995	0.994	0.992
60 plus	0.994	0.993	0.992

source: Author's calculations

management proportions suggest that workers under the age of 30 have 7-11 percent

Figure 9: Simulated Manager Proportions, by age group

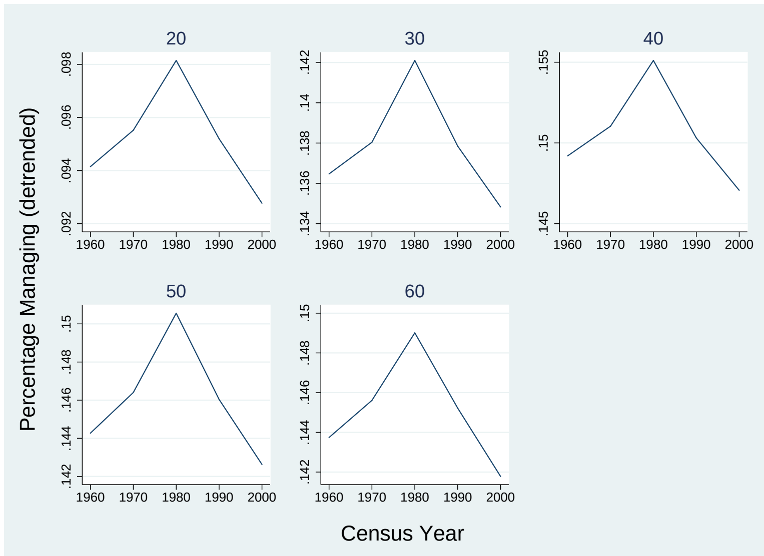

source: Author's calculations

Figure 10: Simulated Manager Proportions, by age group

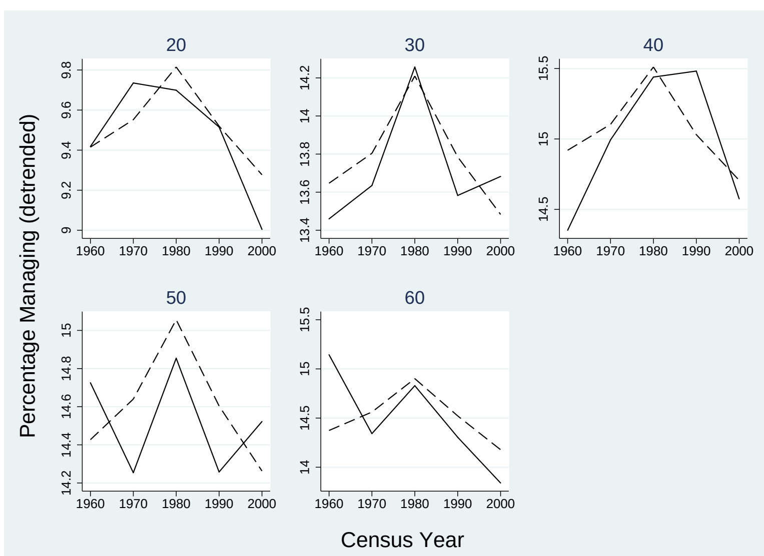

solid − actual data, dashed − simulated

source: Author's calculations

lower management talent than their elders. The 30-39 age group is about 2 percent lower in talent than the older groups. There is no significant difference between the over 40 groups.

5 Conclusions

This exercise obviously does not capture the US productivity slowdown perfectly, though this is likely not a reasonable expectation. We can show that the Lucas model is capable of generating movements in labor productivity that are large. This suggests that there are external effects of changes in experience which are larger than we should expect from an examination of micro returns to experience. Changes in management talent due to demographic changes are potential factor contributing to US productivity movements. The magnitudes suggest that these movements are insufficient to explain the entire US slowdown, but may be causing one fifth of the observed effect.

It is obviously true that the management story offered here is incomplete. The composition of the US labor force changed in a number of dramatic ways during the 1970's. In addition to the entry of the baby boom, participation rates for women rose dramatically. If we assume that women were systematically barred from some management roles during the 1970s is seems possible that there was an additional drain on management quality in the 1970s which reversed as women took on more management responsibility. If true, it may be that the effects in our simulated model are understated.

Additionally, the experience level of managers was necessarily changed by these demographic movements. In the 1970s, the marginal forty and fifty year old managers had no managerial experience since they were only managing due to the entry of the baby boom. Conversely, when the boomers hit these ages, the average manager had more experience in management roles than in previous generations since they had typically been called on to manage at earlier ages.

Understanding the role of demographic change in productivity movements is of interest for more than historical reasons. Many developing countries have large demographic bubbles in the wake of fertility declines. As these cohorts age, their impact on productivity will be significant.

References

Baily, Martin Neil, Robert J. Gordon, and Robert M. Solow, "Productivity and the Services of Capital and Labor," *Brookings Papers on Economic Activity*, 1981, *1981* (1), 1–65.

Bloom, David E., David Canning, and Jaypee Sevilla, "Economic Growth and the Demographic Transition," *NBER Working Paper*, December 2001, *8685.*

___ **, Richard B. Freeman, and Sanders D. Korenman**, "The labour-market consequences of generational crowding," *European Journal of Population*, 1988, *3* (2), 131–76.

Chari, V. V. and Hugo Hoenhayn, "Vintage Human Capital, Growth, and the Diffusion of New Technology," *Journal of Political Economy*, 1991, *99* (6), 1142–1165.

Fernald, John, David Thipphavong, and Bharat Trehan, "Will Fast Productivity Growth Persist?," *FRBSF Economic Letter*, April 6 2007, *2007* (9).

Feyrer, James, "Aggregate Evidence on the Link Between Demographics and Productivity," *Population and Development Review*, December 2006. forthcoming.

___ , "Demographics and Productivity," *Review of Economics and Statistics*, February 2007, *89* (1), 100–109.

Fischer, Stanley, "Symposium on the Slowdown in Productivity Growth," *The Journal of Economic Perspectives*, 1988, *2* (4), 3–7.

Gordon, Robert J. and Daniel E. Sichel, "[Productivity Growth and the New Economy]. Comments and Discussion," *Brookings Papers on Economic Activity*, 2002, *2002* (2), 245–265.

Griliches, Zvi, "Productivity Puzzles and R&D: Another Nonexplanation," *The Journal of Economic Perspectives*, 1988, *2* (4), 9–21.

Jorgenson, Dale W., "Productivity and Postwar U.S. Economic Growth," *The Journal of Economic Perspectives*, 1988, *2* (4), 23–41.

Kogel, Tomas, "Youth dependency and total factor productivity," *Journal of Development Economics*, February 2005, *76* (1), 147–173.

Lazear, Edward P., Katherine Baicker, and Matthew J. Slaughter, *Economic Report of the President*, Washington, D.C.: Government Printing Office, February 2007.

Lucas, Robert E., "On the Size Distribution of Business Firms," *The Bell Journal of Economics*, Autumn 1978, *9* (2), 508–523.

Nordhaus, William D., "Productivity Growth and the New Economy," *Brookings Papers on Economic Activity*, 2002, *2002* (2), 211–244.

Oliner, Stephen D. and Daniel E. Sichel, "The Resurgence of Growth in the Late 1990s: Is Information Technology the Story?," *The Journal of Economic Perspectives*, 2000, *14* (4), 3–22.

Persson, Joakim, "Demographics, Human Capital, and Economic Growth: A Study of US States 1930-2000," *FIEF working paper*, February 2002.

Piketty, Thomas and Emmanuel Saez, "Income Inequality In The United States, 19131998," *Quarterly Journal Of Economics*, February 2003, *118* (1), 1–39.

Ruggles, Steven, Matthew Sobek, Trent Alexander, Catherine A. Fitch, Ronald Goeken, Patricia Kelly Hall, Miriam King, and Chad Ronnander, *Integrated Public Use Microdata Series: Version 3.0 [Machine-readable database]* Minnesota Population Center 2004.

Sarel, Michael, "Demographic Dynamics and the Empirics of Economic Growth," *IMF Staff Papers*, 1995, *42* (2), 398–410.

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