

NBER WORKING PAPER SERIES

WAGES, PRODUCTIVITY, AND THE DYNAMIC
INTERACTION OF BUSINESSES AND WORKERS

John C. Haltiwanger
Julia I. Lane
James R. Spletzer

Working Paper 7994
<http://www.nber.org/papers/w7994>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2000

We would like to thank Victor Aguirregabiria, David Card, Daniel Hamermesh, Andrew Hildreth, David Levine, Jonathan Leonard, Joe Tracy, and numerous conference and seminar participants for their thoughtful comments, as well as Javier Miranda and Bryce Stephens for their research assistance. We would like to thank the ASA/NSF/Census Fellowship program as well as grants from the NSF for financial support and the Center for Economic Studies and the Longitudinal Employer-Household Dynamics (LEHD) program at the Bureau of the Census where this research was conducted. The microdata used in this paper are confidential US Code Title 13. They were assembled as part of a precursor project to the Census Bureau's LEHD program, which is supported by NSF Grant SES-9978093. The Census Bureau will make a fully-documented version of these data available to external researchers as a part of the LEHD program, where John Haltiwanger and Julia Lane are also Research Fellows. The views expressed herein are those of the authors and do not necessarily reflect the views of the Bureau of the Census, the Bureau of Labor Statistics or the National Bureau of Economic Research.

© 2000 by John C. Haltiwanger, Julia I. Lane, and James R. Spletzer. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Wages, Productivity, and the Dynamic Interaction of Businesses and Workers
John C. Haltiwanger, Julia I. Lane, and James R. Spletzer
NBER Working Paper No. 7994
November 2000
JEL No. J21, J23

ABSTRACT

This paper exploits a new matched universal and longitudinal employer-employee database at the US Census Bureau to empirically investigate the link between firms' choice of worker mix and the implied relationships between productivity and wages. We particularly focus on the decision making process of new firms and examine the role of both learning and selection.

Our key empirical results are:

- (i) We find substantial and persistent differences in earnings per worker, output per worker, and worker mix across businesses within narrowly defined industries, which remain even after controlling for other observable characteristics.
- (ii) We find that new businesses exhibit even greater heterogeneity in earnings and productivity than do mature businesses, but that they adjust to the mature business pattern as they age. The adjustment process, while different for earnings and productivity, is consistent both with firms learning as they age and with the exit of "mistake" prone firms.
- (iii) The dynamics of the reduction in productivity heterogeneity of new firms as they age is both complex and very different from the dynamic reduction of earnings heterogeneity.

John Haltiwanger
Department of Economics
University of Maryland
College Park, MD 20742
and NBER
halt@cyclops.umd.edu

Julia Lane
Urban Institute
2100 M ST. NW
Washington, D.C. 20037

James Spletzer
Bureau of Labor Statistics
2 Massachusetts Ave, NE
Washington, D.C. 20212

1. Introduction

Research exploiting new longitudinal business datasets has shown that firms in narrowly defined industries exhibit very different outcomes in terms of productivity, wages, growth, and survival. These firms also exhibit very different choices in terms of technology, organizational structure, size, and factor mix. It is apparent that there is much heterogeneity and persistence in the ways firms produce and do business, and that learning and selection effects play large roles in these processes. In this paper we use a unique new database with longitudinal information on both firms and their workforce to find that this heterogeneity and persistence in productivity both extends to and is related to firm workforce composition. We propose a framework where the underlying source of these different outcomes is differences in firms' endowment of key factors such as technology, capital, organizational structure, and the ability of managers. New firms that make mistakes in recognizing their initial endowment either adjust their workforce as they learn about their capabilities or exit through a competitive selection mechanism.

This study follows in the footsteps of recent studies using longitudinal business data and the much larger and more established literature using worker level data. While the former contributes to understanding the nature and sources of firm heterogeneity and the latter to understanding worker heterogeneity, this paper contributes to the gap in knowledge about the connections between the two – and hence a key piece of the story underlying the interactions of workers and firms. Casual observation suggests that these interactions are important: similar firms, even when faced with similar markets, make very different choices of workers. For example, visits to various individual businesses will reveal different types of workers at each, and repeat visits will confirm that although the workers may have changed, the age, gender, and education level of the replacement workers are similar. Recent empirical evidence using longitudinal matched employer-employee data confirms this anecdotal impression: in their sample of long-lived large firms, Haltiwanger, Lane, and Spletzer

(1999) found that employers choose very different types of workforces and productivity levels, even within relatively narrow industries, and these choices are quite persistent over time.

Figure 1 illustrates this heterogeneity and persistence in workforce mix, productivity, and earnings per worker. Using a sample of long-lived single establishments in Maryland, the upper left panel of this figure plots the proportion of the workplace that is female for each firm in 1986 against the same proportion for the same firm in 1996, the upper right panel plots the proportion of the workplace that is highly educated for each firm in 1986 against the same proportion for the same firm in 1996, the lower left panel of this figure plots the productivity in 1986 against the productivity for the same firm in 1996, and the lower right panel plots the analogous relationship for earnings per worker.¹ The data show quite dramatic heterogeneity and persistence in firm personnel practices and firm outcomes. That is, firms not only choose widely different workforces (the percent female and percent highly educated range from 0% to 100%) but the result that workforce composition, productivity, and earnings per worker are very similar a decade later suggests that these firms are choosing their workforces quite deliberately.²

Such a deliberate choice of worker mix is likely to be complementary to other choices of the firm, such as technology and organizational practice, and to other intrinsic characteristics of the business, such as managerial or entrepreneurial ability. Consideration of these relationships leads one to think about the complex matching and sorting processes that must underlie the connection between firm and worker outcomes. In particular, it raises questions about the relationship between wages, productivity, and worker mix, and the way in which firms – particularly new firms - make such widely differing

¹ A description of the data underlying Figure 1 is provided in section 6. Note that 4-digit effects have been removed from the observations underlying Figure 1. Note that we find similar patterns of heterogeneity and persistence for other measures of worker characteristics including other categories of education, age and foreign born.

² For all cases, we present a 10 percent random sample of our data. For percent highly educated, we only show range from 0 to 0.4 since the mean is relatively low. In all cases, the persistence is measured as the first order AR coefficient from a regression of the 96 value on the 86 value -- for percent female, the AR1 coefficient is

choices. This raises other issues about the role of learning and selection in this process. In this paper we both propose a framework where the underlying source of different outcomes in both workforce composition and productivity is differences in firms' endowment of key factors such as technology, capital, organizational structure, and the ability of managers, and investigate the implications of the model using a new longitudinal matched employer-employee dataset.

Our key empirical results are:

- i) We find substantial and persistent differences in earnings per worker, output per worker, and worker mix across businesses within narrowly defined industries, which remain even after controlling for other observable characteristics.
- ii) We find that new businesses exhibit even greater heterogeneity in earnings and productivity than do mature businesses, but that they adjust to the mature business pattern as they age. The adjustment process, while different for earnings and productivity, is consistent both with firms learning as they age and with the exit of "mistake" prone firms.
- iii) The dynamics of the reduction in productivity heterogeneity of new firms as they age is both complex and very different from the dynamic reduction of earnings heterogeneity.

The paper is organized as follows. We first set the stage by describing how our approach draws on the previous literature and explore some of the implications of Figure 1 a bit more. We then develop a descriptive model that synthesizes a number of these ideas. We follow this by describing the construction of a linked employer employee dataset, which we use to empirically test our key hypotheses. We conclude by summarizing our contribution.

2. Background

The framework presented here views the choice of worker mix as part of the experimentation of how best to create and run a business. It is clear from the emerging literature using linked

0.82; for highly educated the coefficient is 0.45; for productivity the coefficient is 0.47, for earnings the coefficient is 0.45.

employer-employee data that worker heterogeneity and firm heterogeneity are closely linked. Our earlier work (as well as that of others) suggest firms ultimately locate along a productivity/earnings/skill locus with some firms being high productivity, high wage, and high skill firms while others are low productivity, low wage, and low skill firms. In the current paper, we provide further analysis of the connection between productivity, earnings, and worker mix and, in particular, focus on the learning and selection dynamics of new businesses and the role of worker mix in these dynamics. Prior to beginning our own analysis, it is useful to review briefly key strands of the theoretical and empirical literature on firm dynamics that provide background and additional motivation for our analysis.

First, firm performance and behavior, even within quite narrowly defined industries, is quite heterogeneous. Davis, Haltiwanger, and Schuh (1996) document the large magnitude of job creation and destruction and the dominance of idiosyncratic factors in accounting for the observed large pace of job reallocation. Spletzer (2000) reports that forty percent of new businesses die within three years of their birth, and more than half of all jobs destroyed in a three-year period are due to the death of establishments. In addition to growth and survival, firm heterogeneity in wages and productivity has been documented by Abowd, Kramarz, and Margolis (1999), Foster, Haltiwanger, and Krizan (1998), and Dunne, Foster, Haltiwanger, and Troske (2000), amongst others.

Second, complementarities in the firm's production process can lead to a choice of worker skill mix that varies across firms. Kremer's (1993) model of production as a series of tasks where quantity can not be substituted for quality provides insights into how workers of similar skill will match together in firms. Empirical evidence that worker skill is positively correlated across occupations within establishments is provided by Lane, Salmon, and Spletzer (1999). Milgrom and Roberts (1990) discuss the business strategy of exploiting complementarities and the resulting clustering of marketing, production, and organization. Ichniowski, Shaw, and Prennushi (1997) find that innovative employment practices tend to be complements, and these human resource policies are

important determinants of productivity. And Bresnahan, Brynjolfsson, and Hitt (1999) describe how information technology is complementary with the organization of work, including the demand for workers of various skill levels.

Third, there is considerable persistence in outcomes such as wages and productivity across mature businesses. See, for example, Baily, Hulten, and Campbell (1992) with regard to productivity, and Doms, Dunne, and Troske (1997) with regard to wages. In addition, there is considerable persistence in measures of choice of technology, factor mix, and worker mix across mature businesses -- see Doms, Dunne, and Troske (1997) and Haltiwanger, Lane, and Spletzer (1999).

Fourth, the idea that learning and selection dynamics are important for firm dynamics and heterogeneity across firms is most notably attributable to Jovanovic (1982), but also by Ericson and Pakes (1995). In this vein, Jovanovic develops the notion that efficient firms survive and others die. Consistent with the learning and selection models, Evans (1987), Audretsch and Mahmood (1995), and Dunne, Roberts, and Samuelson (1989) find that establishment specific characteristics such as size and age are important contributors to the success or failure of manufacturing establishments. Burgess, Lane, and Stevens (2000) find that firms both at the beginning and end of their lifecycle have much higher turnover rates than mature and successful firms do, and Lane, Isaac, and Stevens (1996) suggest that this contributes to firm exit. Foster, Haltiwanger, and Krizan (1998) note that a common finding in the emerging literature is that low productivity is an excellent predictor of exit.

The evidence of tremendous heterogeneity in outcomes such as productivity, employment, and wages (both in levels and growth rates) across businesses in the same narrowly defined industries is at the core of this burgeoning literature. In the next section, we provide some further analysis of the nature of this heterogeneity. We think this analysis is interesting in its own right but also helps in thinking about the connections between productivity, earnings, and worker mix.

3. The Sources of Firm Heterogeneity

The motivating force for this paper was presented in Figure 1, which showed that businesses in the same year exhibit widely different productivity, wages, and worker mixes – and that these differences are persistent. What might drive this? The most obvious argument is that this persistence and heterogeneity reflects different production techniques: different industries, quite naturally, require different levels of capital intensity, or organizational structure, and that this is reflected in the workforce composition. In this view, if we hone in closely enough to industry characteristics, the observed persistence and heterogeneity should be reduced: firms within very narrowly defined industries should choose roughly similar methods of production and consequently workforces. Figure 1 goes a long way in this direction by controlling for 4-digit effects but it may be that even narrower classification is required for this purpose.

In fact, we have a unique opportunity to examine this, because we can link our analytical dataset (described in detail in section 5 below) to internal Census files with detailed firm level responses to the quinquennial economic Censuses in 1987 and 1992. In Figures 2 and 3, we recreate the relationship described in Figure 1 for two detailed five-digit industries: restaurants and plumbing businesses.³ For each of these industries, we examine the heterogeneity and persistence for all businesses, and then businesses that fit into a very narrow set of what one might think of as observable dimensions of differences across firms.

For the restaurant industry we present eight panels of firm level scatter plots on the relationship between variables on worker mix, earnings and productivity for the firm in 1987 against the corresponding variable in 1992. The left column of panels present these relationships for the entire restaurant industry in the state of Maryland; the right-hand column for the subset of restaurants that have table service, that are in metropolitan area in the state of Maryland, and where the average customer spends the same amount (within a \$5 range). By restricting our attention in the right panels

³ We chose these industries both because of an abundance of data points and because the questions asked of firms in these industries were particularly well suited to our analytical purpose.

to the narrow set of restaurants on several dimensions, we think we have come fairly close to finding firms with observationally equivalent sets of characteristics.

Since we are looking within a relatively narrow five digit industry, it is “surprising” (although this mimics the results that controlled for 4-digit industry) to find considerable heterogeneity and persistence in the measures of worker mix, earnings, and productivity in the left panels for the entire restaurant industry. As noted, this “surprising” result may simply reflect observable differences in products and services provided – that is for the restaurant industry, some datapoints are for the fast food restaurants while others are expensive, fancy restaurants. The right panels of Figure 2 depict the same scatter plots focusing on the businesses with seemingly very similar observable characteristics. The striking finding is that there remains substantial heterogeneity and persistence in earnings, productivity, and worker mix even within as narrow a group as can be defined using detailed survey data.⁴

We repeat this same exercise for plumbing businesses in Figure 3. On the left hand column we recreate Figure 1 for all firms in the industry. On the right column we subset the firms to include those in the top decile of capital expenditures and capital assets. Again, despite controlling for these observable differences, we still find similar patterns – substantial heterogeneity and persistence across businesses not only in the same narrowly defined industry but also controlling for various detailed characteristics of the business.

Our interpretation of these findings of large residual heterogeneity and persistence is that they are prima facie evidence that there are some inherently unobservable (or difficult to measure) dimensions of firm heterogeneity. Indeed, this would mirror the prior research surveyed in Section 2, which documented substantial differences in capital intensities as well as propensities to adopt advanced technologies, organizational and human resource practices across business even in the same 4-digit

⁴ The autocorrelation coefficients from simple linear regression fit through these points is positive and significant for all of the plots in Figure 2 and 3.

industries. In our view, these differences are likely to reflect, at least in part, inherently difficult to measure characteristics such as managerial ability or related organizational practices.⁵

4. Theoretical Underpinnings – Learning and Selection

The preceding section described differences in productivity levels, workforce composition, and earnings for long-lived, mature businesses. In order to understand these different outcomes, we argue in this paper that the heterogeneity in each of these is a manifestation of an underlying quasi-fixed heterogeneity in firm type. This also explains the persistence in outcomes: since firm type is quasi-fixed, the manifestation of type should not change quickly. In the model we sketch in this section, we make three key assumptions about differences across firms. First, we assume that there are some underlying differences in firm type, labeled “k” for now. Second, we assume that firms of different types have incentives to use different mixes of workers through the nature of the interaction between firm type and the productivity of different types of workers. Third, since some aspects of firm type are arguably unobservable even to the firm itself initially, newly born firms learn and evolve towards their “long run” type (or potentially exit as discussed below). In what follows, we first focus on the nature of the connection between firm type and worker mix. We then consider the role of uncertainty, the effect of this on the learning of surviving firms, the associated belief distribution in this learning, and the consequences for the choice of worker mix. After discussing the optimal choice of worker mix for surviving firms, we discuss the factors that determine selection.

4A. The link between firm type and worker mix

We only sketch a model here, drawing heavily on the literature, in order to motivate and help interpret the empirical exercises we consider in section 5. The simple descriptive model focuses on the choice of worker mix and abstracts from scale effects. As such, we take the size of businesses as

⁵ It is important to emphasize that in all of the results to be presented in this paper control for 4-digit industry, which is a basic control for observable differences across businesses. There may well be, of course, other important observable dimensions of firm heterogeneity. As such, we believe that it would be a fruitful path for

given and develop the model focusing on the determinants of output per worker and the mix of workers by worker type within a business, although many of the same factors that we emphasize that influence mix also influence scale.⁶

We begin by assuming that the observed firm heterogeneity fundamentally reflects some variation in business type – the idiosyncratic ability of managers, the firm's organizational practices or policies, or the choice of productive inputs such as technology, physical capital, or organizational capital. In what follows, we refer to this type as “k” with high “k” firms inherently more productive. For now, it is useful to think of k as fixed but in practice it is better to think of k as changing either slowly or, even better, infrequently (as when a firm adopts a technology through an investment spike - see, for example, Cooper, Haltiwanger, and Power (1999)). Following the learning literature, we also argue that firms are not likely to be fully aware of their type at birth, and initially “guess” their type to be K. Since k is not directly observable to the firm, but output is, firms have to revise their priors each period based upon observations of their output. However, a firm's ability to extract information on its type is clouded to some extent by unobservable profit (i.e., productivity, demand, and cost) shocks.

Furthermore, assume that output per worker is given by $\theta f(k,s)$ (suppressing time subscripts), where θ is initially assumed to be a random (i.i.d.) profit shock, s is skill mix, f is increasing in its arguments and the cross partial between k and s is positive (i.e., there is k-skill complementarity).⁷

future research to explore on an industry-by-industry basis other observable dimensions of differences across businesses.

⁶ For example, Lucas (1977) develops a model of the size distribution of employment that depends upon heterogeneity in the managerial ability across businesses. The latter is part of what we have mind as “k” in the analysis that follows. In the Lucas (1977) model, the most productive entrepreneurs employ more workers but do not capture the entire market because of diseconomies of control – even the most productive individual can only effectively control a limited size operation. While this is outside the scope of our model, this argument is implicitly important in our context as well in that it provides one potential explanation for why the high and low “k” firms can operate simultaneously.

⁷ The assumption that θ is i.i.d. is not critical and likely unrealistic – we assume this for now to make some of the simple illustrative analysis easier to present.

Firms observe total output each period after decisions for the period are made, but do not observe k or θ directly. The observation of output implies there is learning over time about true k .

We specify wages as having two components: one a function of skill and one reflecting the ability of workers in a firm to extract rents (and or more generally internal labor market considerations that might generate firm specific components to wages – in what follows we use “rent sharing” as a catch all label for such effects). Specifically, wages are given by $w(s) = \omega(s) + \delta E\Pi$, where $\omega(s)$ is the reservation wage for workers of skill s (with $\omega' > 0$), and $\delta E\Pi$ is the rent sharing premium such that $\delta < 1$ and Π is the surplus in the period (see below). In principle, we permit δ to vary over time (e.g., over the life cycle of the firm) and to be idiosyncratic and stochastic. Put differently, one source of heterogeneity across businesses may be the nature of the internal labor market and, accordingly, the earnings profiles of businesses may change as the firm ages and changes its internal paysetting structure. Our descriptive model includes such firm-specific effects in a crude fashion to recognize their potential empirical relevance – modeling the nature of the firm-specific components of wage determination in this context should be a high priority for future research. In this regard, we recognize that our descriptive model is missing potentially interesting aspects of the firm-specific components of wages. It could be, for example, that δ not only varies with time but with s – many interesting internal labor market models may yield incentives for varying the wage differential based upon worker characteristics.

In this setting, expected profits in a given period are given by (normalizing the output price to one for now):

$$E[(1-\delta)\{\theta f(k,s) - \omega(s) - F\} | K], \quad (1)$$

where F represents fixed costs of operating in each period and K represents the belief about type in each period.⁸ In our simple model, firms will learn about their type and update K each period based upon prior observations. Then they will either decide to shut down (selection) or decide to continue to operate. If the latter, they will choose s , and generate their output.

This generates several concrete predictions. First, in the absence of uncertainty and rent-sharing, the choice of s each period will satisfy $\theta \partial f(k,s) / \partial s = \omega'(s)$ and k -skill complementarity implies that firms with higher k will have higher s and higher wages. Second, the slope of the wage and productivity profiles should be the same *for given k*. In this simple case, systematic firm productivity differentials, wage differentials, and workforce composition differences are entirely driven by differences in k .

The intuition of this result can be seen in figure 4, where we have graphed the productivity profiles for firms with three different values of k ($k_3 > k_2 > k_1$). For any given skill level s , firms that are type k_3 are the most productive, firms that are type k_1 are the least productive, and firms that are type k_2 are in between. Expanding this simple three firm example to a continuum of k , the points of tangency where marginal productivity equals marginal wages sketch out the equilibrium line of observed productivity/skill combinations (for simplicity in drawing the figure, we have assumed that wages are linear in skills). If we interpret skill on the horizontal axis as an index of observable worker characteristics, this line can be thought of a possible underlying theoretical justification for the statistically significant productivity/skill relationship found for long-lived large firms by Haltiwanger, Lane, and Spletzer (1999).⁹

⁸ We assume conditions are such that expected profits are a strictly concave function of s . In many of our examples, we assume f is strictly concave in s and ω is linear in s .

⁹ It is worth noting that there are a number of potential reasons why high “ k ” and low “ k ” firms can both exist in the same market. If k is managerial ability, economies of scope as in Lucas (1977) can limit the size of the business. If k is type of technology adopted, then costs of adopting and learning new skills may mean that there is heterogeneity in technologies being used as in Cooper, Haltiwanger and Power (1999) or Caselli (1999). It might also be that we want to think about k as the persistent component of idiosyncratic technology or cost shocks as in Hopenhayn (1992) or Hopenhayn and Rogerson (1993). Such idiosyncratic technology and cost

4B. The role of uncertainty about type and the implications for learning

In the more general case, with uncertainty about firm type, the skill mix s will be chosen so that $E[\theta \partial f(k,s)/\partial s \mid K] = \omega'(s)$. We still obtain the result that firms with higher K choose higher s and have higher wages. However, as is obvious from the first order condition, uncertainty about k implies that the slopes of the wage and productivity profiles need not be the same, even for firms with the same k . Figure 5 provides a visual interpretation of this. Suppose we have three firms in an entering cohort, each of which has the same k , but each of which initially believes it has a different type $K_3 > K_2 > K_1$. At birth, then, they choose different skill levels based upon their belief about K . However, because productivity is determined by actual type as well as the skill mix of workers, the optimistic firm's realization of productivity is less than anticipated. Similarly, the pessimistic firm's realization of productivity is greater than what it expected. Only firm 2 is at a point on the long-run (no uncertainty) productivity/skill locus. Firms 1 and 3 then adjust their skill level along the productivity curve, so that eventually they end up with a workforce composition appropriate to their type (and in this case the three firms with the same k would end up with the same workforce compositions and the same productivity level).

While the specific quantitative implications for productivity and worker mix dynamics depend upon the precise assumptions about the initial distribution of beliefs, there are a few key aspects of the implications illustrated that do not depend upon the precise assumptions regarding this distribution.

shocks might relate to the location of production or the specific technology adopted (and thus be related to some of these other factors as well). It is the case that the equilibrating forces act somewhat differently across these different possibilities. As noted above, in Lucas (1977) the differences in managerial ability and the economies of scope yield differences in the size of businesses (and thus determine the equilibrium size distribution). In Caselli (1999), there is some general capital that is perfectly mobile and thus sites that have adopted the latest technology have not only better capital but more capital. In Hopenhayn and Rogerson (1992), labor is the only input considered, labor is subject to adjustment costs and is subject to decreasing returns and there are entry costs (therefore the mass of entrants does not take over the market nor does any individual business). These different models and associated equilibrating forces have implications for the dispersion of productivity across businesses. It is beyond the scope of this paper to specify fully the nature of the factors generating the productivity dispersion and the associated equilibrating forces. Mapping the observed patterns and evolution of dispersion in productivity (and earnings) across businesses to the specific implications of these alternative models is an important area for future research.

For one, Figure 5 helps illustrate the more general prediction that businesses will learn that they have made mistakes about their type and will adjust their mix of workers in a systematic fashion: overestimates of type result in decreased s ; underestimates result in increased s .¹⁰ A second key related prediction that does not depend upon the initial distribution of beliefs is that the deviations from the long run profile should be reduced over time. In the limit, only the idiosyncratic productivity shocks should generate deviations in the long run. Wages, productivity, and worker mix are all still related through their common relationship to k , but the tightness of that relationship, as measured by a reduction in the variance of the deviations, increases as a cohort ages. The following example helps illustrate the working of the simple model, as well as the role of initial beliefs.

4C. A simple example

To more fully understand the implications of this model for learning dynamics and the link between wages, productivity, and skill mix, it is useful to consider a simple example. Suppose we have the following production function (for $0 < \alpha < 1$) and wage function (where we have abstracted from rent sharing – a point we return to below):

$$y = \theta k s^\alpha, \quad \omega(s) = \bar{\omega} s \quad (2).$$

The optimal s is then:¹¹

$$s^* = \left(\frac{\alpha K}{\bar{\omega}} \right)^{\left(\frac{1}{1-\alpha} \right)} \quad (3).$$

Equations (2) and (3) imply that as firms get older and learn about their type, the dispersion of skills, s , and the dispersion of productivity, y , depend directly on the dispersion of beliefs about type, K .

¹⁰ Note that it should be clear that even if firms had the same initial belief, firms with negative errors (overly optimistic) will decrease their skill mix while firms with positive errors (overly pessimistic) will increase their skill mix.

¹¹ The assumption that θ is i.i.d. makes this expression simple. Observations of past output yields information about k but no information about future productivity shocks. For this example, we have also normalized the mean of θ to one.

Accordingly, the changes in the dispersion of mix, productivity, and wages depend upon the evolution of $\text{var}(\mathbf{K})$.

To consider the empirical implications, it is helpful to take logs of the production function, defining $\tilde{\bullet} \equiv \ln(\bullet)$, which yields $\tilde{y} = \alpha\tilde{s} + \eta$. The “error” term in this expression has information on both type and productivity shocks: $\eta = \tilde{k} + \tilde{\theta}$. Clearly, estimating a simple regression of y on s does not yield an estimate of α , since s and k are positively correlated and k is unobserved at least to the econometrician, but rather an estimated coefficient, β where

$$E(\beta) = \alpha + \left(\frac{\text{cov}(\tilde{k}, \tilde{s})}{\text{var}(\tilde{s})} \right) \quad (4)$$

The “omitted variable” bias from regressing \tilde{y} on \tilde{s} differs for mature and young businesses, and hence help us understand the role of learning about type on the productivity/skill relationship. For mature businesses, where $\mathbf{K} = k$, estimating this regression yields the slope of the long run profile (which we refer to below as $\bar{\beta}$) depicted in Figures 4 and 5. For young businesses that have uncertainty about their type, the relationship between skills and productivity depends on the structure of their initial beliefs. If the initial distribution is similar to a type of classical measurement error ($\mathbf{K} = k + \kappa$, where κ is a random measurement error with mean zero), then learning effects as firms converge to the long run profile will yield a reduction in overall dispersion in productivity, a reduction in the dispersion of the worker mix, but an increase in the slope of the estimated productivity/skill relationship as firms age. If, in contrast, newly-born firms have identical ex ante beliefs, then we would expect to see rising dispersion in both productivity and worker mix, but a flattening of the slope of the estimated productivity/skill relationship as firms age. Note that regardless of the assumption about initial beliefs, the convergence to the long run profile implies that residual deviations from the long run profile will diminish over time.

In this simple example, the implied changes for wage dispersion are similar to those for productivity on some dimensions but different on others. The similarities arise because the dispersion of wages depends upon the dispersion of s . In the absence of rent sharing, the only effect of the learning about k on wages is the direct effect through s . In this case, the model predicts changes in the dispersion in wages only due to systematic changes in the distribution of s , but does not predict systematic changes in the estimated slope of the wage/skill profile or systematic changes in the distribution of the unobservables (i.e., residuals). With rent sharing, the qualitative predictions for the distribution of wages begin, in principle, to mimic those for productivity. Even here, the quantitative predictions may be different given the nature of rent sharing.

On the learning side of the story, then, as firms' age, declining uncertainty about firm type leads to systematic changes in the skill mix and decreases in the residual dispersion in the relationship between productivity and skills (with similar predictions potentially for wages). Changes in overall wage, productivity, and worker mix dispersion (and the estimated relationship between these) depend upon the initial distribution of beliefs and the associated resolution of uncertainty. We return to these predictions when we discuss the empirical specification in section 5.

4D. Selection

On the selection side, we draw heavily on the implications from the existing literature on selection as we have little to add (theoretically) except for some modest refinements. Following Jovanovic (1982) and Dunne, Roberts, and Samuelson (1989), we suppose that entering firms face the uncertainty about type that we have already outlined, entry costs, and a fixed cost of operating in every period (denoted F in the prior section).¹² In this environment, only firms with sufficiently high expected k will enter. The marginal entering firm will have zero expected discounted profits taking

¹² In a formal model of entry and exit, we might also want or need to guarantee some exit of mature businesses. The latter could be generated from there being some probability of any given production site becoming sufficiently unproductive (e.g., as in Davis and Haltiwanger (1990) or Ericson and Pakes (1995)) that it should

into account the costs of entry. However, given the learning as well as the fixed costs of operating each period, some firms will learn that they have negative expected profits and exit. The firms that are more likely to exit are those with low k since they will be closer to the marginal firm.

In terms of observable implications, low k firms will tend to have both lower productivity and lower “ s ”. Putting these pieces together, we get some modest refinements of the typical predictions of the selection literature. The standard prediction is that low productivity firms are more likely to exit and this reflects the fact that low productivity firms are more likely (on average) to have learned through experience of producing that they are expected negative profit firms. Our refinement is that we should see that such marginal firms have also selected a low “ s ” and thus we should see that low “ s ” predicts exit. Errors should also provide information: firms that deviate from the long run profile in the negative direction should also be more likely to exit since this yields information about the learning about type. That is, a negative error indicates that the firm is more likely to have overestimated rather than underestimated its type and such firms will revise their expected type downwards and be more likely to exit.

To sum up, like the standard selection model, firms that learn they are low productivity firms are more likely to exit. Our focus on the choice of worker mix suggests the following refinements. First, the choice of mix will be correlated with k and productivity and thus should help predict exit. Second, the component of productivity that is not predicted by worker mix may also have some additional power to predict exit as this indicates the direction of learning.

5. Empirical Specifications and Tests

The above theoretical discussion generates testable hypotheses about selection and the evolution of productivity, wages, and worker mix as a cohort of businesses enters and ages. We have

exit. In terms of the model sketched above, this could be captured by assuming that there is a component of the profit shock θ that has this feature.

specific predictions about the evolution of productivity and earnings as firms' age and the role of changes in the worker mix in this evolution.

The evolution of productivity and earnings can be traced empirically by using the following specification following directly from the discussion in the previous section:

$$y_{it} = X_{it} \beta_t + \varepsilon_{it}, \quad (5)$$

Here y_{it} is a measure of an outcome such as output per worker or payroll per worker, X_{it} is a vector of worker characteristics, β_t is the slope of the profile in the current year (and thus may not be the slope of the long run profile) and ε_{it} is the residual (reflecting a deviation from the current "profile").

On the learning side, the theoretical predictions of changes in dispersion can be tested directly using a Juhn, Murphy, and Pierce (1993) (hereafter JMP) full accounting approach to decompose changes in variance of outcomes into changes in β 's, changes in X 's, and changes in residuals. That is, by exploiting this specification for a cohort of entering business, we can examine the changes in the dispersion of productivity and wages and the extent to which changes in the dispersion is accounted for by changes in the distribution of observable and unobservable factors.

Our JMP full accounting distribution exercises are based upon exploiting the following two closely related decompositions:

$$y_{it} = X_{it} \bar{\beta} + v_{it} = X_{it} \bar{\beta} + X_{it} (\beta_t - \bar{\beta}) + \varepsilon_{it}, \quad (6)$$

where $\bar{\beta}$ represents a set of time invariant coefficient estimates and β_t represents the estimated coefficients for each age level of an entering cohort. In this methodology, the distribution of the outcome variable, y , is ultimately decomposed into three terms based upon the expressions on the right of the second equal sign: the first term captures the influence of observable characteristics for given β 's, incorporation of the second term captures the influence of changing differentials associated with the given characteristics, and incorporation of the third term captures the influence of unobservable

factors. However, as an intermediate step, it is useful to consider the simpler decomposition of y into two terms given by the expressions to the right of the first equal sign: the first term captures as before the influence of observable characteristics for given β 's and the second term everything else. This simpler intermediate decomposition is useful in this context because the "error" term in this simpler decomposition, v_{it} , can be interpreted as representing deviations from the long run profile. Note that by construction in what follows that performing the full 3-term decomposition automatically yields the intermediate decomposition so we can characterize how the distribution of deviations from the long term profile evolves over time.

To see how to use this decomposition for characterizing changing dispersion, consider using this decomposition to decompose changes in the differential between the 90th and 10th percentile of y for the surviving members of an entering cohort of businesses that we follow over successive years. Using the first term in each of the above decompositions and constructing the implied 90-10 differential based upon this term alone for each year and then considering the change in this differential over time provides a measure of the contribution of changing characteristics to the change in overall dispersion. Inclusion of the second term in the 3-way decomposition in a measure of the predicted 90-10 differential incorporates the contribution of the changes in the β 's. The marginal contribution of the changing β 's is measured as the difference between the change in dispersion generated by the first two terms from the change in dispersion generated from the first term. In a similar manner, the contribution of unobservables is generated by considering the change in dispersion implied by all three terms less the change in dispersion implied by the first two terms alone.¹³

¹³ One issue that is relevant is that the ordering of the decomposition may matter here. Covariance-like effects on the components of the decomposition of the 90-10 can make the ordering important. As such, appropriate caution about the interpretation of the contribution of the terms is required. Given the ordering we have chosen, there is a clear interpretation of the contribution of the first term but the second term and its interpretation is potentially impacted by covariance-like effects. Note that regardless of the ordering the contribution of the unobservables will be the same. The ordering we have chosen is the natural one for the exercises we are interested in – that is, we are especially interested in the contribution of changes in the distribution of worker mix across businesses to changes in dispersion of productivity and earnings across businesses. We have

As noted above, we use the estimated coefficients from a balanced panel of mature businesses to generate the “fixed” β 's (the $\bar{\beta}$ in the above equation). The motivation for this -- influenced largely by Figures 4 and 5 and our initial empirical work in Haltiwanger, Lane, and Spletzer (1999) -- is that mature businesses have presumably learned their type and thus the estimated coefficients will not reflect errors in type. The time varying β 's are generated from estimation of the specification for samples of entering cohort of businesses for each of the years that we follow over time.

The theory generates predictions about the changes in skill mix that should be observed for an entering cohort of firms. Firms that overestimate their type will decrease their skill mix while firms that underestimate their type will increase their skill mix. In addition, the magnitude of the skill change will be greater the greater the initial mistake made by the firm. These predictions can be tested for entering cohorts of businesses by examining the relationship between changes in the skill mix and the sign and magnitude of their errors. To be specific, we construct a measure of predicted productivity using the long run β 's and the actual skill mix chosen by the business. The difference between actual and predicted productivity yields a measure of the “error” made by the entering business (this measure of the error is captured as v_{it} in equation (6)). A similar measure of the wage “error” is also constructed.

Finally, on the selection side, the theory predicts that firms with observed initial low productivity are more likely to fail. It could be that the firm had expectations that it was a marginal firm with a low k and therefore had also chosen a low s and therefore worker mix should help predict exit. Alternatively, the low observed productivity could be a surprise (and thus show up in the error term) that indicates the firm overestimated its type and through learning the firm is more likely to exit. We can test these predictions by examining the probability of exit as a function of initial productivity

considered the alternative ordering and briefly discuss the results of this alternative ordering in the results that we present.

as well as the magnitude and sign of the errors made in k . The measure of the errors is constructed in the manner discussed above.

6. Longitudinal Linked Employer-Employee Data

In order to address these questions, we need a dataset that links workers to their employers, provides information on the characteristics of both, and follows both over time. We use three sources of data to create our longitudinal linked employer-employee dataset.¹⁴ Firm characteristics are obtained from the Standard Statistical Establishment List (SSEL), which is the U.S. Census Bureau's sampling frame for businesses in the United States. Demographic characteristics of individuals are obtained from administrative records at the U.S. Census Bureau. The bridge between the workers and the firms is the Unemployment Insurance (UI) wage records from the state of Maryland.

Every quarter, all employers subject to state Unemployment Insurance (UI) laws are required to submit quarterly contribution reports detailing their monthly employment and quarterly wages to the State Employment Security Agencies. The employer also provides quarterly wages for every employee. These data on individual employees are known as the wage records, and are used by the states to manage their unemployment insurance program. The UI wage records have been used by many previous authors for research into worker flows, job flows, and the earnings effects of labor reallocation -- see, for example, Jacobson, LaLonde, and Sullivan (1993), Anderson and Meyer (1994), Lane, Miranda, Spletzer, and Burgess (1999), and Burgess, Lane, and Stevens (2000). In this paper, we use wage records from the state of Maryland from the second quarter of 1985 through the third quarter of 1997. This dataset consists of 108,254,142 observations on 5,006,622 workers and 262,062 employers.

¹⁴ We should note that our linked employer-employee dataset has been created under strict protocols within the confidential firewalls at the U.S. Census Bureau. Following Census Bureau policy, the resulting micro analytic database can be used only for statistical purposes, and the output from such statistical analysis is carefully reviewed to ensure the confidentiality of individual households and businesses.

The main advantage of the UI wage records is their coverage of essentially a universe of workers and employers.¹⁵ Another advantage of the wage records data is the ability to construct longitudinal linkages for both individuals and employers. We have linked the UI wage records across quarters, using the employer as the unit of analysis. To ensure that we do not falsely define businesses that are involved in ownership changes as births and deaths, we use information on worker flows to distinguish administrative changes in employer identification numbers from true births and deaths.

Perhaps the main drawback to the UI wage records data is the lack of even the most basic demographic information on workers. We overcome this by linking the UI microdata to administrative data residing at the Census Bureau which contains information such as date of birth, place of birth, and gender. We are also particularly interested in the educational composition of the workforce. We were able to directly obtain education information on a sample of the workforce via other administrative data links to the long form of the decennial census.¹⁶ Since education can change for workers aged 18-24, we subset the data to only include workers over age 25. We then follow Angrist and Krueger (1991) in using detailed place of birth, month of birth, year of birth, race, and gender to impute education for the balance of the workforce.

The information in the UI wage records is also quite limited with regard to characteristics of the employer. We link the UI data to the Standard Statistical Establishment List (SSEL), which is the Census Bureau's business register. All SSEL information about individual businesses is confidential under Title 13, Unites States Code. The SSEL is the source of basic employment and payroll measures summarized by industry and geographic area in the Annual County Business Patterns series, and the SSEL also serves as a resource for research into topics such as longitudinal business demographics. Walker (1997) provides an excellent description of the SSEL. We use the natural log

¹⁵ The state based UI system provides a virtual census (98 percent) of employees on nonfarm payrolls. The major exclusions from private sector UI coverage are agricultural employees, self employed workers, private household workers, unpaid family workers, and employees in small nonprofit organizations.

¹⁶ The long form is only sent to 1/6 of households.

of payroll per worker as our measure of earnings in the empirical work below, and we use the natural log of sales per worker as our measure of productivity.¹⁷ Because of concerns about price variation or variation in capital usage that vary across years and industries, most of the empirical work below uses data measured as deviations from year and four-digit industry means.

The analytical dataset that we construct from these merged files has the employer as the unit of analysis. We aggregate the individual worker data into summary measures for each firm. For example, we construct firm level proportions of workers under age 30, between 30 and 55, and over 55. We construct the proportion of the firm's workforce that is female and the proportion that is foreign born. We also construct firm level proportions of workers with low, medium, and high amounts of education (where these categories roughly correspond to less than high school, high school graduates and those with some college, and college graduates). The resulting dataset has annual observations on essentially every firm in the State of Maryland between 1985 and 1997.

We impose a series of restrictions before turning to the analysis. We restrict the data to those firms in the UI with a valid match to the SSEL, and we keep single establishment firms with nonmissing sales data in the SSEL.¹⁸ While restricting attention to single establishment firms has limitations for some purposes, it is not a major concern in the current context where we want to focus on the dynamics in the first few years after entry. As Spletzer (2000) has shown, almost 90 percent of

¹⁷ As described by Haltiwanger, Lane, and Spletzer (1999), the sales to employment ratio should be regarded as a proxy for labor productivity since revenue is divided by employment rather than hours, and the GDP deflator is used rather than the appropriate firm specific price deflator. Variation in hours per worker may vary systematically across types of businesses because of, for example, the use of part-time workers – by controlling for 4-digit effects we have abstracted from differences in hours per worker across 4-digit industries. Still, one could easily imagine that within the same narrowly defined industry that some businesses have chosen to use part-time workers more intensively than others – thus, one of the sources of between business variation we may be capturing is this type of variation in the choice of the role of part-time workers.

¹⁸ Each of these restrictions is somewhat severe in terms of losing observations. The Maryland UI data identifies firms with a state-specific identifier, whereas the SSEL identifies firms with a federal identifier. We have the crosswalk between these two identifiers for 1988 Q1 and 1992 Q1. This implies that all births after 1992 Q1 and all deaths before 1988 Q1 will not be in the merged dataset. Because the sales data in the SSEL is reported at the EIN level rather than the establishment level, and because most births are single establishment births, we have chosen to delete multi-establishment companies from our analysis.

entering establishments are single unit firms and thus we have the set of firms that dominate entry dynamics. Because of concerns about partial year data, we delete observations from 1985 and 1997 and we delete observations in the year of birth and the year of death. We further restrict our data to firms not in the agriculture and public administration industries, and we delete observations in the top and bottom 1% of the productivity and earnings distributions.

Since the focus of the paper is on the evolution and selection of newly-born firms, we create several analytical datasets. The first dataset consists of a balanced panel of firms that exist every year 1985-1997. There are 22,420 of these firms, with full information available in 116,461 employer-year observations (sales data are missing in the SSEL in 1988, 1989, 1990 and 1993; sales data are, on average, missing for 32 percent of single establishments in the other years) The average firm in the balanced panel has 11.9 employees. This balanced panel is the sample of mature firms from which we estimate the $\bar{\beta}$ coefficients in the JMP decompositions. The data in Figure 1 are from this balanced panel. The second dataset consists of firms born within the 1985-1990 period; we restrict on births in these years so we can observe their growth and survival during their initial six years of life. There are 20,338 such births, contributing 48,664 employer-year observations. Of these births, 14,722 survive for at least six years, and these successful births contribute 39,981 employer-year observations.

7. Basic Patterns in the Relationship between Productivity, Earnings, and Worker Mix

Table 2 presents very basic regressions relating productivity, earnings, and worker mix. The regressions in table 2 use the simple specification described in section 5 to relate firm-level outcomes to the characteristics of the workers employed by the firm. We find that firms in our data have higher measures of labor productivity and earnings if they have workforces with a higher fraction of foreign born workers, a lower fraction of female workers, a higher fraction of prime age workers, and a higher

fraction of more educated workers.¹⁹ Many of these coefficients are statistically significant. While the estimates vary somewhat across the age of businesses (which our theoretical arguments suggest should happen), the slopes of the productivity and wage profiles with respect to worker characteristics appear to be qualitatively similar across businesses of different ages.

It is tempting of course to compare the magnitudes of the coefficients in the productivity and earnings regressions and draw inferences regarding the extent to which different workers are being paid their marginal products.²⁰ For example, it appears that increasing the percentage of highly educated workers raises earnings per worker more than output per worker. Our discussion of the simple theoretical model in section 4 cautions against interpreting such interesting findings as implying that, in this case, highly educated workers are being paid more than their marginal products. Our theory and evidence presented above suggests that, amongst other things, there are inherently unobserved factors “k” that may influence productivity and wages differentially. These factors are likely to be correlated with our measures of worker characteristics and as such, it is important to interpret the coefficients in Table 2 as reflecting the empirical covariances between productivity, earnings, and worker characteristics. Along related lines, the resolution of uncertainty may have

¹⁹ Haltiwanger, Lane, and Spletzer (1999) note that the female coefficient in these regressions requires special interpretation. Because we are using employment rather than hours in the denominator of the productivity and the earnings variables, these coefficients might be biased downward if the propensity to work part-time is greater for females. More generally of course, there may be many factors that influence the productivity and earnings effects that are outside the scope of the type of theoretical considerations we have discussed. Rather than take on those issues directly here, we take the approach here that we are interested in the connection between productivity, earnings, and a variety of different indicators of worker mix. It would be of interest to consider the respective contribution (and the associated interpretation) of each of the indicators we have separately.

²⁰ For example, Hellerstein, Neumark and Troske (1999) present analysis that compares slopes of similarly specified earnings and productivity profiles in order to examine questions of discrimination. As they carefully show there are important issues of functional forms that can influence tests of the productivity versus earnings differentials. Since such tests are not the focus of our analysis we do not attempt to investigate the sensitivity to functional form issues in the same manner.

differential impacts on wages and productivity so that the slopes of the estimated profiles for young businesses have additional sources of possible differences between the profiles.²¹

While some caution is called for, it is also the case that it is interesting to compare the coefficients – especially since we are interested in using the patterns in these coefficients in empirical exercises that follow. The last columns on the right of Table 2 report differences in the coefficients for each of the respective worker characteristic variables and a test of statistical significance. For the balanced panel, we find that the coefficient in the productivity equation is less than that in the wage equation for foreign born, older and more educated workers. Again, for the balanced panel, we find that the coefficient in the productivity equation is greater than that in the wage equation for female, younger and less educated workers. All of these differences are statistically significant. For the younger businesses, the pattern of the difference in coefficients is roughly similar but somewhat more erratic and less statistically significant. An interesting aspect of these findings is that whatever factors are driving the differences in the productivity and earnings coefficients are apparently more important for well-established businesses. Since uncertainty would presumably be less important for mature businesses, this suggests that the wedge between productivity and wages for mature businesses must be driven by other factors.

While the inherently difficult to measure or observe components of “k” make some inferences difficult (like interpreting slope comparisons), recall that our theory offers predictions about how the

²¹ It is true that we have some reasonable conjectures about the nature of the biases induced by unobservable “k” for particular specifications. For example, consider the example sketched out in section 4 but abstracting from uncertainty (Figure 4). In that example, we abstracted from rent sharing and internal labor market considerations and thus k only influences wages indirectly via s while k influences productivity both directly and indirectly via s. In addition, in that example the production process exhibited k-skill complementarity so higher k yielded higher s. Under this scenario, the “bias” in the coefficient on s in the productivity equation abstracting from uncertainty should be positive – that is, the coefficient should reflect both the direct impact of s on productivity and the fact that s should have a positive covariance with k and k is unobservable and thus not included in the empirical specification. The implication is that under this strict specification (abstracting from both uncertainty and rent sharing considerations), the estimated coefficient on s in the productivity equation should exceed the analogous coefficients in the wage equation. The results in Table 2 tend to reject this strict specification (abstracting from uncertainty with no rent sharing and internal labor market considerations, etc.) since the coefficients on “s” for wages tend to exceed those for productivity.

nature of these covariances should evolve over the life cycle of a business and moreover how the worker characteristics should evolve over the life cycle. It is the latter predictions that we can analyze even without measures of “k” and it is these predictions that we investigate in the analysis that follows.²²

8. Empirical Analysis

8A. *Learning: Changes in Productivity and Earnings*

We now investigate the evidence for and the consequences of firms learning about their type. Our interest here is to characterize the overall nature of dispersion of productivity and earnings, its evolution over time, and decompose the sources of its change.

Using our sample of successful births, the productivity and earnings dynamics are given in table 3.²³ An examination of Table 3 demonstrates that successful births exhibit systematic changes in the mean and variance of productivity and earnings during their first six years: as the firm ages, the mean productivity falls, the mean earnings rises, and the standard deviation of both productivity and earnings falls.²⁴

²² There are some other possible strategies to deal with “k”. For example, in section 3, we provide some exploratory analysis of potential observable indicators of “k”. However, as noted in the analysis in that section, it appears that there are important unobservable components of “k”. Note that in addition to trying to measure “k” directly, it would be possible to exploit the panel nature of the data for continuing businesses to attempt to abstract from the influence of “k” using fixed effect techniques. We plan on pursuing these ideas in future research.

²³ In all of the forthcoming analysis in tables 3-6, all data have year and four digit industry means removed.

²⁴ One hypothesis of the falling dispersion of earnings and productivity is that it simply reflects a statistical artifact due to successful businesses growing under the accompanying hypothesis of a random assignment of heterogeneous workers across businesses. That is, suppose that workers are heterogeneous in terms of productivity and are paid their marginal products but workers are just randomly assigned to businesses. Moreover, suppose that new businesses grow. Then, for a sample of new businesses that grow, the businesses will exhibit falling dispersion in earnings and productivity as businesses will inherently become more statistically similar as they grow. While we think the evidence is overwhelmingly against the hypothesis of random assignment of workers, we nevertheless conducted some further robustness checks since some elements of this statistical artifact hypothesis might operate even without complete random assignment as it is true that new successful businesses exhibit positive growth (on average). In particular, we examined the patterns of change in productivity and earnings dispersion for alternative groups of successful businesses. One check we made was based upon initial size – we classified businesses into quartiles based upon initial size. We found that the pattern of falling dispersion holds even more strongly for businesses in the top quartile as for the bottom quartile of initial size. This evidence does not support the statistical artifact hypothesis since this hypothesis

The JMP decompositions of the sources of changes in the declining productivity and earnings dispersion are presented in the top two panels of table 4. We find that changes in the distribution of worker characteristics and the differentials associated with those characteristics account for a modest fraction of the falling dispersion in productivity and a substantial fraction of the falling dispersion in wages.²⁵ Specifically, changes in the distribution of worker characteristics account for 3% of the decline in the 90-10 ratio of productivity during the first six years, and changes in the differentials (i.e., estimated regression coefficients) account for 5% of this decline. Changes in the distribution of worker characteristics account for 34% of the decline in the 90-10 ratio of earnings during the first six years, and changes in the differentials account for 43% of this decline. This implies that 92% of the declining productivity dispersion is due to unobservables, whereas 22% of the declining earnings dispersion is due to unobservables. In terms of the simpler intermediate decomposition in equation (6), these results suggest that only 3% of the decline in productivity dispersion can be interpreted as arising from reduced dispersion in worker mix and only 34% of the decline in earnings dispersion can be interpreted to reduced dispersion in worker mix. The vast majority of the decline is associated with changing deviations from the long run earnings/worker mix profile (i.e., residuals).

should be of particular relevance for the smallest businesses. Second, we classified new businesses into growth rate quartiles based upon growth over the first six years of existence. We found that the pattern of falling dispersion holds for all growth rate quartiles – moreover, the lowest quartile has strong negative growth while the top quartile has strong positive growth. It is true that the quantitative decrease in dispersion is somewhat greater for the growth rate group in the top quartile. However, under pure random assignment, we should actually see dispersion rising for the lowest quartile group (as they are getting smaller) and dispersion falling for the highest quartile group. Our interpretation of these robustness checks is that the falling dispersion in productivity and earnings is not being driven by this statistical artifact hypothesis. Moreover, the other evidence we have presented does not support random assignment. Such evidence includes the finding of persistent heterogeneity amongst mature businesses in earnings, productivity and worker mix and the finding of rising mean earnings and falling mean productivity amongst new businesses.

²⁵ Recall that the ordering in the JMP decomposition does not impact the contribution of the unobservables so the combined contribution of changes in worker characteristics and differentials is robust to ordering as well. When we change the ordering placing the contribution of the changing differentials prior to changes in the characteristics, we largely get the same patterns as reported in Table 4. There is a modest increase in the contribution of changes in the distribution of characteristics but we still find that changes in the distribution of characteristics matter a lot for earnings but not for productivity. We also find that changes in the differentials still contribute substantially for earnings but not for productivity (actually they modestly go the “wrong” way).

How do these decomposition results fit with the predictions of the learning model? There are several facets of the results that we discuss in turn and, in so doing, it is important to recall that some of the predictions of the model depend on the assumptions regarding the initial distribution of beliefs. One prediction of the model that does not depend upon the latter is that we should observe a reduction in the overall dispersion of residuals from the long run profile. We observe that this prediction holds for both productivity and wages. With regard to other predictions, the declining overall dispersion of productivity and earnings is, at first glance, consistent with the assumption that the variance of initial beliefs about type exceeds the variance of true types so that learning yields falling dispersion in productivity, earnings, and worker mix as firms age. However, the weak contribution of changes in the distribution of worker characteristics raises doubts that the falling overall dispersion in productivity is due to businesses having an initial dispersion of beliefs about type that exceeded the true dispersion in types. It is also the case that the changes in dispersion resulting from worker mix effects are not supportive of the hypothesis that all firms had common (identical) initial beliefs at entry.

It is interesting that changes in the distribution of worker characteristics account for a greater fraction of the changes in dispersion of earnings than productivity. This finding is consistent with the prediction that deviations from the long run profile are less likely to be important in accounting for changes in earnings dispersion. The reason is that mistakes about “ k ” may only (or primarily) have an impact on changes in earnings dispersion through changes in the workforce mix while these same mistakes imply a falling contribution of deviations from the long run profile for productivity. We also know, however, from Table 2 that, if anything, the slope of the earnings/skill mix profile is steeper than the productivity/skill mix profile. Thus, any change in the distribution of worker characteristics will have a larger impact on earnings dispersion than on productivity dispersion.

We focus on the ordering in Table 4 in the text because this ordering provides the most straightforward evidence on the contribution of changes in the distribution of worker characteristics.

We also find that changes in the slopes of the earnings and productivity profiles tend to decrease rather than increase dispersion – although the contribution of changing slopes for productivity is very small. The implied flattening of the schedules is consistent with the hypothesis that firms had similar or common initial beliefs that led to a steeper initial profile than the long run profile. However, since the changes in slope matter mostly for earnings, it may be (as we will discuss further below) that there are other factors contributing to the earnings dynamics beyond the learning effects emphasized in our model. If, for example, the propensity for rent sharing increases over the lifetime of a business, this could potentially yield a weaker relationship between earnings and characteristics as a business ages. Moreover, for earnings, taken together the contribution of changes in the distribution of worker characteristics and changes in the differentials associated with those characteristics work in the “wrong” way for the changes in earnings dispersion to be entirely driven by learning effects. That is, the changes due to distribution of worker characteristics suggest that initially $\text{var}(K) > \text{var}(k)$ while the changes due to differentials suggest that initially $\text{var}(K) < \text{var}(k)$. In short, there are some interesting dynamics for young businesses involving the interaction of changing worker characteristics and earnings differentials associated with those characteristics that appear to involve factors other than learning.

The results in Table 3 also show that mean productivity falls and mean earnings rises during the first six years of the successful birth's lifecycle. This difference in trends is, at first glance, somewhat puzzling since our model suggests that average earnings also captures some portion of average productivity, and thus mean earnings and mean productivity should move in the same direction. To further investigate this divergence, we present results from a JMP decomposition on the means in the bottom two panels of table 4. We find that changes in worker mix alone during the first six years would lead to falling mean earnings and falling mean productivity, and changes in the regression coefficients would account for a relatively small increase in both means. Therefore, the opposite trends we observe are primarily accounted for by changes in unobservable characteristics.

There could be a number of explanations for why unobservable factors explain the different trends in mean productivity and earnings. The decline in mean productivity might reflect mean-reverting tendencies for successful firms to have high initial draws of productivity (that is, high θ shocks). The increase in mean earnings might reflect an increasing share of rents as a business ages – in terms of the simple descriptive model, it may be that the rent sharing fraction (δ) increases with the age of the business. We do not have a formal explanation of why rent sharing would vary in this systematic way – however; it seems plausible that survival probabilities are related to the ability to enter with low initial wages and implicit commitments of higher future wages upon success (i.e., survival). A related possibility is that the productivity and earnings measures may capture different dimensions of entrepreneurial effort that are important in the first few years of a business's existence. It is possible that entrepreneurs put in long hours at the birth of a business and pay themselves very little salary, reversing this pattern as the business ages. This would be consistent with a falling mean sales-per-worker ratio and an increasing mean earnings-per-worker ratio. We believe that investigating the robustness of and the explanations for the opposite trends of mean earnings and mean productivity is an interesting area for future research.

To sum up, there are striking and systematic changes in the mean and dispersion of productivity and earnings per worker in the first several years after the birth of a cohort of firms. Some of these patterns (in particular, the reduced dispersion from deviations from long run profiles) are clearly consistent with the hypothesized effects of learning. There are, however, significant differences in the patterns of the findings across productivity and earnings. These differences are not inconsistent with learning effects per se but do suggest, especially for earnings, that other factors are playing a role in the evolution of the relationship between earnings per worker and worker characteristics. Perhaps this is not surprising as firm-specific components of earnings determination

are likely to evolve over the life cycle of new businesses and to vary with worker characteristics as well.

8B. Learning: Changes in Worker Mix

In this part of the empirical analysis, we are interested in describing how firms adjust their workforce in response to their initial errors in knowing their type. The theory predicts that firms with positive surprises in initial productivity are more likely to have underestimated k and will exhibit increases in skill, and firms with negative surprises in initial productivity are more likely to have overestimated k and will exhibit decreases in skill. The dependent variable we use to measure changes in skill is a change in a skill index from the second year to the sixth year for a successful entrant (measured formally as $(X_6 - X_2) \bar{\beta}$, where X_t is the vector of workforce composition variables in the t^{th} year after birth, and $\bar{\beta}$ is the vector of regression coefficients estimated from the balanced panel from the productivity regression on worker characteristics).²⁶ We measure the change in skill between years 2 and 6 to avoid problems with measures of initial year productivity as explanatory variables in the regression.

In column 1 of table 5, we regress the change in the productivity skill index on the firm's initial productivity error. We find that those firms with positive errors (such as firm 1 in figure 5 that underestimated its k) tend to increase their skill mix, and those firms with negative errors (such as firm 3 in figure 3 that overestimated its k) tend to decrease their skill mix. To see whether this result might be driven by regression to the mean effects, we control for the initial worker skill mix by including the initial predicted productivity in column 2.²⁷ The coefficient on the initial error is reduced somewhat

²⁶ We have also replicated the analysis in Table 4 using the earnings specification. We find qualitatively similar results although they are somewhat less strong and estimated less precisely. Part of the reason for this is that the implications for responses to “errors” may not carry over to earnings. In the absence of rent sharing, there is no sense in which the gap between actual and predicted earnings reflects an error.

²⁷ Of course, learning is a potential source of regression to the mean effects but there may be other factors generating regression to the mean effects in the worker mix. By controlling for initial worker mix, we abstract from such factors and focus on the impact of the “error” made in the choice of worker mix.

but remains positive and statistically significant. Finally, in column 3, using indicator variables for whether the firm's initial error is in the top or bottom third of the error distribution, we find that the largest changes in skill are from firms in the bottom third of the error distribution decreasing their skill levels.

We interpret these results as saying that the initial error in choosing skill level has important empirical content. This is consistent with the view that new firms systematically experiment with their mix of workers as they learn about their type.

8C. Selection

As discussed earlier, our theory has specific predictions about both selection and learning. With regard to selection, we present simple OLS death regressions in Table 6. For each cohort of entering businesses, we define an indicator variable that is equal to one if the business does not survive within the first six years and zero otherwise. In the first column, we regress this indicator of death within the first six years on the firm's actual productivity in its first year. We find that the firms with initially low productivity are more likely to exit. This is in keeping with the theoretical predictions, where we expect that firms with lower levels of initial productivity are more likely to be close to the marginal firm (i.e., close to the fixed cost threshold at which exit is optimal).²⁸

In keeping with the theoretical analysis, and with the visual depiction in Figures 4 and 5, we decompose actual productivity into two parts: the component that is on the long run productivity/skills locus, and the residual. The former is where the firm expects to be given its chosen skill mix, and the latter is what we term the firm's error in identifying its type. In addition to the predicted component

²⁸ This result is mirrored in a death regression that uses earnings instead of productivity: we find that firms with low initial earnings are likely to exit. At first glance, this result might seem counter-intuitive: one would expect that all else equal, firms with lower earnings would have higher expected profits and thus be more likely to survive. One explanation for this finding is that low earnings reflects low skill, and skill is chosen based upon the firm's expectation of type. Indeed, in the model, wages $w(s)$ are equal to $[(1-\delta)\omega(s)+\delta\{\theta f(k,s)-F\}]$, and in the presence of rent sharing ($\delta>0$), wages are a function of k and thus become an alternative measure of productivity. When we estimate a death regression with actual productivity and actual earnings as explanatory variables, both

signaling the firm's belief about its type k , we expect firms that make optimistic errors about k will exhibit negative productivity errors, which leads to firms learning that they are lower k and therefore a greater likelihood of exit. These hypotheses are empirically valid: the coefficients in column 2 are negative and statistically significant. Interestingly, the coefficient on the predicted component of productivity is substantially larger in magnitude than the coefficient on the error component. One interpretation for this result is that the observed worker mix is a good proxy for type and thus a good predictor for productivity, whereas the error is a noisy predictor of type (as it should be) but nevertheless has predictive power in terms of predicting survival.

We now extend the analysis to examine whether there are asymmetries in how various parts of the distributions of the predicted and error components of initial productivity influence exit. In column 3 of table 6, we find that firms that choose worker characteristics that map into the lower third of the predicted productivity distribution are more likely to exit. Consistent with our prediction that overly optimistic mistakes about type matter for survival, we find that firms with errors in the lower third of that distribution are more likely to exit. However, we find that firms with errors in the upper third of that distribution are not less likely to exit. Thus, there is an interesting asymmetry in the response to errors.

This interesting asymmetry is outside the scope of the simple selection model we have sketched. The latter predicts that firms with negative errors should be more likely to exit and firms with positive errors should be less likely to exit. We find that it is only negative errors that matter here. A potential explanation for this finding (outside of the model we have sketched) is that negative errors have a direct impact on productivity and profitability. That is, the mismatch between type and worker mix may have direct adverse consequences, particularly for those that are overly optimistic. It will also be the case that overly optimistic firms will be paying the higher wages that go along with the

come in with statistically significant negative coefficients. We interpret this to mean that both earnings and productivity are noisy measures of underlying firm type.

higher skill mix. It seems plausible that firms with such a mismatch problem might find themselves in difficulty and be more likely to exit.

9. Concluding Remarks

We began this study by exploiting a rich new employer-employee dataset to find substantial and persistent differences in earnings per worker, output per worker, and worker mix across businesses within narrowly defined industries. We found that these remained even after controlling for other observable characteristics. We argued that these were all manifestations of the same underlying dynamic: the type of firm that generated these differences. We sketched a model in which businesses deliberately choose different worker mixes because of differences in their types and productivity interactions (e.g., k-skill complementarity) between the type of business and the mix of workers. Following the industrial organization literature, we posited that businesses learn about their types over time so there is churning and evolution of the worker mix at businesses and this is related to the observed earnings and productivity dynamics. In order to understand the dynamics underlying the choice of workforce composition, our empirical analysis focused on the productivity and earnings dynamics of young businesses.

We found that, at entry, successful (i.e., surviving) new businesses exhibit substantially higher than industry-average productivity and substantially lower than industry-average earnings per worker – interestingly, however, as the business ages there is convergence between productivity and earnings.

We also found that new businesses exhibit even greater heterogeneity in earnings and productivity than do mature businesses, but that two factors contribute to the lower earnings dispersion of mature businesses. First, new firms learn to adjust their workforce composition as they age. Second, new firms with low initial productivity, or that make initial mistakes about their type (particularly those that are overly optimistic), are more likely to exit.

These results on earnings and productivity dynamics are of interest in their own right but our focus has been to examine the contribution of the choice and evolution of worker mix to these

dynamics. When we did this, we found that unobservable factors dominate the observed changes in the distribution of productivity - consistent with the learning model. However, the failure of worker characteristics to account for much of the considerable decline in productivity dispersion is somewhat puzzling.

The results for changes in earnings dispersion are quite different on some dimensions. As with productivity, overall dispersion fell and changes in the distribution of unobservable factors contributes substantially to this decline. Here, however, changes in the distribution of worker characteristics and changes in the differentials associated with those characteristics play a substantial role. The importance of worker characteristics is potentially consistent with the theory, given the uncertainty about firm type. However learning does not explain why both workforce mix adjustment and adjustment in differentials act to reduce earnings dispersion. In addition, the substantial differences in the adjustment process for earnings and productivity suggest that some other factors are playing a role in the interaction between earnings and worker characteristics for young businesses.

In sum, this paper extends a very basic result from the empirical industrial organization literature – namely that learning and selection effects matter for firm dynamics – to the labor economics literature in showing that this experimentation process involves the choice of worker mix in both the selection and the learning process. Put differently, new firms struggle to find their way in deciding how to undertake production and run a business – some do not succeed and fail while others learn and prosper. In this sense, our results, although derived from a new dataset, are very consistent with both the rich empirical literature documenting marked differences in firm behavior within industries, and the explicit theoretical understanding of the impact of different initial endowments – like managerial skill, physical and organizational capital – on firm outcomes.

There are clearly aspects of the results that require further research. A particularly interesting result is investigating the contribution of rent sharing and internal labor market considerations to the different dynamics between average earnings and productivity as firms age. In addition, it would be

useful to investigate the interaction between different types of k and different sets of worker characteristics on earnings and productivity. A question that we do not directly address but which is equally interesting is the impact of these firm dynamics on workers. In particular, we know that the types of firms which are more likely to exit are those which are low productivity, with a low level of workforce skill. We also know from the recent literature that this entry and exit plays an important role in productivity growth, yet that the effects of involuntary job loss can be large and persistent, at least for long tenure mature workers. Quantifying and understanding the potential economy wide gains in productivity from this process of creative destruction and adjustment via the costs borne by workers should be an important area for future research.

References

- Abowd, John M., Francis Kramarz, and David Margolis (1999). "High Wage Workers and High Wage Firms." *Econometrica*, pp. 251-334.
- Anderson, Patricia M. and Bruce D Meyer (1994). "The Extent And Consequences Of Job Turnover." *Brookings Papers On Economic Activity*, pp. 177-248.
- Angrist, Joshua D. and Alan B. Krueger (1991). "Does Compulsory School Attendance Affect Schooling and Earnings?" *The Quarterly Journal of Economics*, pp. 979-1014.
- Audretsch, David B. and Talat Mahmood (1995). "New Firm Survival: New Results Using a Hazard Function." *Review of Economics and Statistics*, pp. 97-103.
- Baily, Martin Neil, Charles Hulten, and David Campbell (1992). "Productivity Dynamics in Manufacturing Plants." *Brookings Papers on Economic Activity*, pp. 187-249.
- Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt (1999). "Information Technology, Workplace Organization and the Demand for Skilled Labor: Firm-Level Evidence." NBER Working Paper No. 7136.
- Burgess, Simon, Julia Lane, and David Stevens (2000). "Job Flows, Worker Flows and Churning." *Journal of Labor Economics*, pp. 473-502.
- Caselli, Francesco (1999). "Technological Revolutions." *American Economic Review*, pp. 78-102.
- Cooper, Russell, John Haltiwanger and Laura Power (1999). "Machine Replacement and the Business Cycle: Lumps and Bumps," *American Economic Review*, pp. 921-946.
- Davis, Steven J. and John C. Haltiwanger (1990). "Gross Job Creation and Destruction: Microeconomic Evidence and Macroeconomic Implications," *NBER Macroeconomics Annual*, pp. 123-168.
- Davis, Steven J., John C. Haltiwanger, and Scott Schuh (1996). *Job Creation and Destruction*. MIT Press, Cambridge, Massachusetts.
- Doms, Mark, Timothy Dunne, and Kenneth R. Troske (1997). "Workers, Wages, and Technology." *The Quarterly Journal of Economics*, pp. 253-290.
- Dunne, Timothy, Lucia Foster, John Haltiwanger, and Kenneth Troske (2000). "Wage and Productivity Dispersion in U.S. Manufacturing: The Role of Computer Investment." NBER Working Paper No. 7465.
- Dunne, Timothy, Mark J. Roberts, and Larry Samuelson (1989). "The Growth and Failure of U.S. Manufacturing Plants." *The Quarterly Journal of Economics*, pp. 671-698.

- Ericson, Richard and Ariel Pakes (1995). "Markov-Perfect Industry Dynamics: A Framework for Empirical Work." *Review of Economic Studies*, pp. 53-82.
- Evans, David S. (1987). "Tests of Alternative Theories of Firm Growth." *Journal of Political Economy*, pp. 657-674.
- Foster, Lucia, John Haltiwanger, and C.J. Krizan (1998). "Aggregate Productivity Growth: Lessons from Microeconomic Evidence." NBER Working Paper No. 6803.
- Haltiwanger, John C., Julia I. Lane and James R. Spletzer (1999). "Productivity Differences Across Employers: The Role of Employer Size, Age, and Human Capital." *American Economic Review Papers and Proceedings*, pp. 94-98.
- Hellerstein, Judith K., David Neumark, and Kenneth R. Troske (1999). "Wages, Productivity, and Worker Characteristics: Evidence from Plant-Level Production Functions and Wage Equations." *Journal of Labor Economics*, pp. 409-446.
- Hopenhayn, Hugo (1992). "Entry, Exit, and Firm Dynamics in the Long Run," *Econometrica*, pp. 1127-1150.
- Hopenhayn, Hugo and Richard Rogerson (1993). "Job Turnover and Policy Evaluation: A General Equilibrium Analysis," *Journal of Political Economy*, pp. 915-938
- Ichniowski, Casey, Kathryn Shaw, and Giovanna Prennushi (1997). "The Effects of Human Resource Management Practices on Productivity: A Study of Steel Finishing Lines." *The American Economic Review*, pp. 291-313.
- Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan (1993). "Earnings Losses of Displaced Workers." *The American Economic Review*, pp. 685-709.
- Jovanovic, Boyan (1982). "Selection and the Evolution of Industry." *Econometrica*, pp. 649-670.
- Juhn, Chinhui, Kevin M. Murphy, and Brooks Pierce (1993). "Wage Inequality and the Rise in the Return to Skill." *Journal of Political Economy*, pp. 35-78.
- Kremer, Michael (1993). "The O-Ring Theory of Economic Development." *The Quarterly Journal of Economics*, pp. 551-575.
- Lane, Julia, Alan Isaac, and David Stevens (1996). "Firm Heterogeneity and Worker Turnover." *Review of Industrial Organization*, pp. 275-291.
- Lane, Julia, Javier Miranda, James Spletzer, and Simon Burgess (1999). "The Effect of Worker Reallocation on the Earnings Distribution: Longitudinal Evidence from Linked Data." In *The Creation and Analysis of Employer-Employee Matched Data*, edited by John C. Haltiwanger, Julia I. Lane, James R. Spletzer, Jules J.M. Theeuwes, and Kenneth R. Troske, North-Holland Press, pp. 345-374.

Lane, Julia, Laurie Salmon, and James Spletzer (1999). "Establishment Wage Differentials: Evidence from a New BLS Survey." Unpublished Working Paper, Bureau of Labor Statistics.

Lucas, Robert (1977). "On the Size Distribution of Firms." *Bell Journal of Economics*, pp. 508-523.


Milgrom, Paul and John Roberts (1990). "The Economics of Modern Manufacturing: Technology, Strategy, and Organization." *The American Economic Review*, pp. 511-530.

Spletzer, James R. (2000). "The Contribution of Establishment Births and Deaths to Employment Growth." *Journal of Business and Economic Statistics*, pp. 113-126.


Walker, Ed (1997). "The Census Bureau's Business Register: Basic Features and Quality Issues." Unpublished paper presented at the 1997 American Statistical Association Meetings.

Figure 1: Heterogeneity & Persistence
Firms Alive in 1986 and 1996 with >10 Employees Each Year


Proportion Female: 1986 vs. 1996, by firm



Proportion Highly Educated: 1986 vs. 1996, by firm



Productivity: 1986 vs. 1996, by firm



Actual Earnings: 1986 vs. 1996, by firm




Figure 2: Productivity, Earnings and Worker-Mix in Restaurant Industry

All Restaurants







Table Service Restaurants


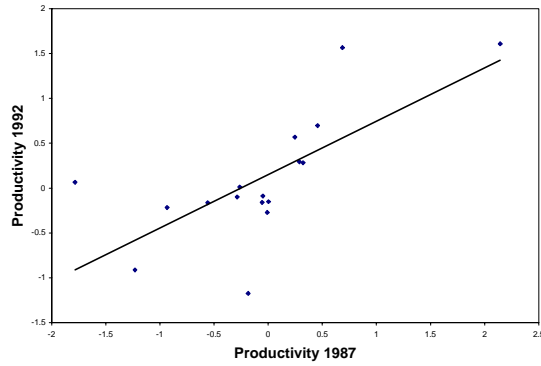
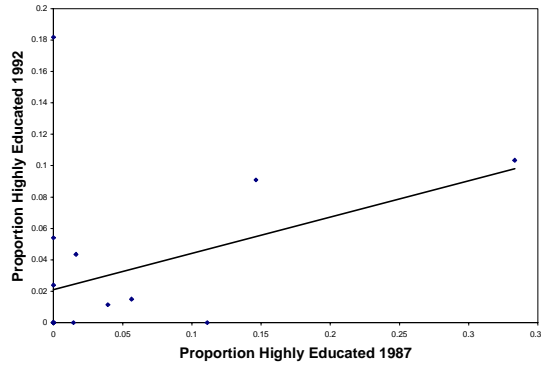




Figure 3: Productivity, Earnings and Worker-Mix in Plumbing Industry

All Plumbing Businesses



Capital-Intensive Plumbing Businesses










Figure 4: No uncertainty, No rent sharing, Linear Wages
Three different values of k




Figure 5: The link between k , skill, and productivity
Uncertainty: 3 different values of K , one value of k




Table 1: Descriptive Statistics

Mean (Std. Dev.)	6-year Successful Births: 1 year after birth	6-year Successful Births: 6 years after birth	Balanced Panel
Productivity	4.32 (.851)	4.26 (.827)	4.32 (.792)
Earnings	8.20 (.708)	8.25 (.699)	8.39 (.675)
% Foreign Born	.096 (.230)	.101 (.231)	.075 (.173)
% Female	.436 (.382)	.449 (.372)	.463 (.346)
% Age <30	.348 (.332)	.245 (.271)	.249 (.245)
% Age >55	.067 (.173)	.102 (.214)	.140 (.211)
% Low Educated	.216 (.276)	.284 (.286)	.257 (.239)
% High Educated	.054 (.158)	.048 (.141)	.038 (.106)
Sample Size	4,168	9,389	116,461

Descriptive Statistics are presented before removing Year & 4-Digit SIC means.

Productivity = $\text{Ln}\{ [\text{Annual Sales } (\$1000) / \text{GDP Deflator}] / \text{March 12th Employment} \}$

Earnings = $\text{Ln}\{ [\text{1st Quarter Payroll} / \text{GDP Deflator}] / \text{March 12th Employment} \}$.

Table 2: Regression Analysis

	Productivity Regressions			Earnings Regressions			Differences in Coefficients between Productivity and Earnings Regressions		
	6-year Successful Births: 1 year after birth	6-year Successful Births: 6 years after birth	Balanced Panel	6-year Successful Births: 1 year after birth	6-year Successful Births: 6 years after birth	Balanced Panel	6-year Successful Births: 1 year after birth	6-year Successful Births: 6 years after birth	Balanced Panel
% Foreign Born	.085 (.053)	.080 * (.032)	.068 * (.011)	.031 (.043)	.042 (.027)	.097 * (.010)	.053 (.055)	.038 (.032)	-.030* (.011)
% Female	-.181 * (.041)	-.134 * (.027)	-.181 * (.008)	-.398 * (.033)	-.431 * (.022)	-.562 * (.007)	.217* (.042)	.297* (.026)	.382* (.008)
% Age <30	-.030 (.041)	.046 (.042)	-.018 (.010)	-.213 * (.033)	-.141 * (.035)	-.145 * (.009)	.183* (.042)	.187* (.041)	.127* (.010)
% Age >55	-.189 * (.071)	-.230 * (.036)	-.206 * (.010)	-.170 * (.058)	-.217 * (.030)	-.160 * (.009)	-.019 (.074)	-.013 (.035)	-.046* (.010)
% Low Educated	-.150 * (.049)	-.215 * (.039)	-.181 * (.010)	-.280 * (.040)	-.188 * (.033)	-.249 * (.008)	.130* (.051)	-.027 (.038)	.069* (.010)
% High Educated	.142 (.075)	.070 (.052)	.128 * (.018)	.093 (.061)	.145 * (.043)	.209 * (.016)	.048 (.078)	-.076 (.051)	-.082* (.018)
Sample Size	4,168	9,389	116,461	4,168	9,389	116,461			
R-Squared	.017	.016	.017	.078	.069	.081			

All regressions include an intercept and an indicator for whether initial employment exceeds 100. The successful birth regressions include cohort dummies. The balanced paned regressions include controls for age.

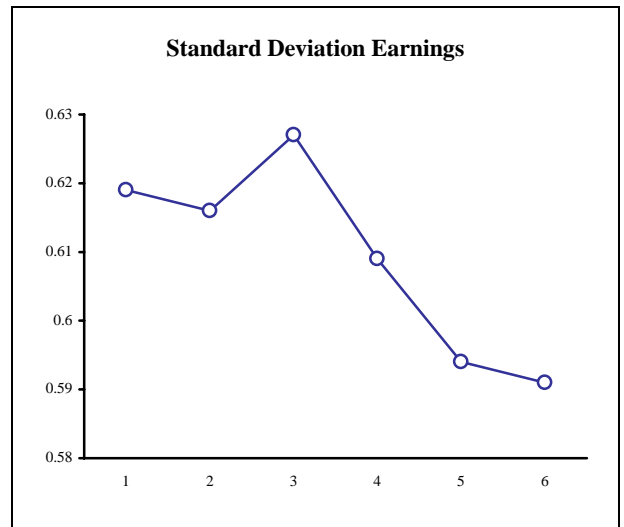

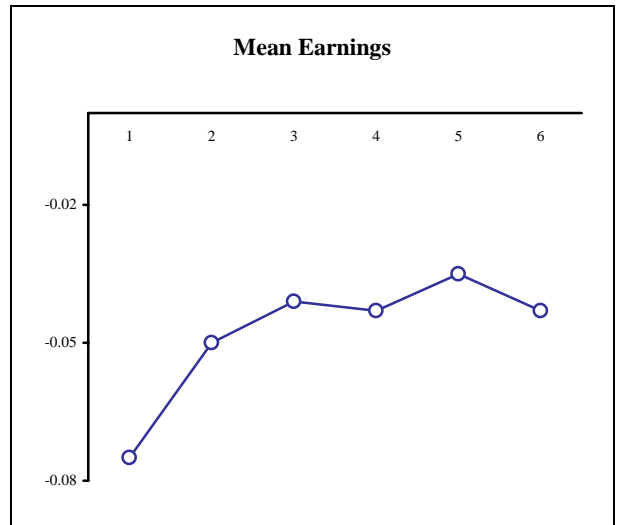

All data (except intercept and cohort dummies) have Year & 4-Digit SIC means removed. Productivity and earnings per worker are in logs.

Table 3: Productivity and Earnings Dynamics, 6-year Successful Births

Productivity	Mean	Standard Deviation
1 Year After Birth	.076	.737
2 Years After Birth	.063	.704
3 Years After Birth	.049	.700
4 Years After Birth	.038	.704
5 Years After Birth	.038	.696
6 Years After Birth	.017	.686

Earnings	Mean	Standard Deviation
1 Year After Birth	-.075	.619
2 Years After Birth	-.050	.616
3 Years After Birth	-.041	.627
4 Years After Birth	-.043	.609
5 Years After Birth	-.035	.594
6 Years After Birth	-.043	.591

All data have Year & 4-Digit SIC means removed.
Productivity and earnings per worker are in logs.



**Table 4: Productivity & Earnings Dynamics, 6-year Successful Births
Juhn-Murphy-Pierce Decompositions**

	90-10 Productivity	90-10 from Observable Worker Characteristics	90-10 from Differentials Associated with Worker Characteristics	90-10 from Unobservables
1 Year After Birth	1.830	.215	-.010	1.626
6 Years After Birth	1.699	.211	-.016	1.505
Difference	-.131	-.004 (3%)	-.006 (5%)	-.121 (92%)

	90-10 Earnings	90-10 from Observable Worker Characteristics	90-10 from Differentials Associated with Worker Characteristics	90-10 from Unobservables
1 Year After Birth	1.555	.486	-.062	1.131
6 Years After Birth	1.479	.460	-.095	1.114
Difference	-.076	-.026 (34%)	-.033 (43%)	-.017 (22%)

	Mean Productivity	Mean from Observable Worker Characteristics	Mean from Differentials Associated with Worker Characteristics	Mean from Unobservables
1 Year After Birth	.076	.019	-.002	.059
6 Years After Birth	.017	.006	.002	.010
Difference	-.059	-.013 (22%)	.004 (-7%)	-.049 (83%)

	Mean Earnings	Mean from Observable Worker Characteristics	Mean from Differentials Associated with Worker Characteristics	Mean from Unobservables
1 Year After Birth	-.075	.022	-.006	-.091
6 Years After Birth	-.043	.002	.001	-.046
Difference	.032	-.020 (-63%)	.007 (22%)	.045 (141%)

X = [% Foreign Born, % Female, % Age <30, % Age >55, % Low Educated, % High Educated].
Coefficient estimates are from table 2. All data have Year & 4-Digit SIC means removed. Productivity and earnings per worker are in logs.

Table 5: Skill Index Change Regressions, 1990 Successful Births

	(1)	(2)	(3)
$X * \bar{\beta}$: Predicted		-.244 *	
Productivity		(.024)	
1 if Bottom 1/3 of			.026 *
Distribution			(.006)
1 if Top 1/3 of			-.023 *
Distribution			(.006)
Error: Actual -	.010 *	.008 *	
Predicted	(.003)	(.003)	
1 if Bottom 1/3 of			-.015 *
Distribution			(.005)
1 if Top 1/3 of			-.003
Distribution			(.006)
R-Squared	.008	.097	.076

All regressions include an intercept. Sample size=1051.

Dependent Variable is $[X(\text{year 6 after birth}) - X(\text{year 2 after birth})] * \bar{\beta}$ (Productivity); mean=-.0030.

All explanatory variables are measured as of the first year after birth. $\bar{\beta}$ is estimated from the balanced panel.

All data (except intercept) have Year & 4-Digit SIC means removed. Productivity per worker is in logs.

Table 6: Probability of Death Regressions, 1986 & 1990 Births

	(1)	(2)	(3)
Actual Productivity	-.039 *		
	(.007)		
X* $\bar{\beta}$: Predicted		-.271 *	
Productivity		(.057)	
1 if Bottom 1/3 of			.053 *
Distribution			(.014)
1 if Top 1/3 of			.011
Distribution			(.014)
Error: Actual -		-.035 *	
Predicted		(.007)	
1 if Bottom 1/3 of			.076 *
Distribution			(.014)
1 if Top 1/3 of			.011
Distribution			(.014)
R-Squared	.031	.033	.035

All regressions include an intercept and cohort dummies. Sample size=6631.

Dependent Variable (before removing year and 4-digit SIC means) = 1 if establishment does not survive 6 years, 0 otherwise; mean (before removing year and 4-digit SIC means) = .3714.

All explanatory variables are measured as of the first year after birth. $\bar{\beta}$ is estimated from the balanced panel.

All data (except intercept and cohort dummies) have Year & 4-Digit SIC means removed. Productivity per worker is in logs.