Wages, Productivity, and Worker Characteristics: Evidence from Plant-Level Production Functions and Wage Equations

Judith K. Hellerstein, *University of Maryland and NBER*

David Neumark, *Michigan State University and NBER*

Kenneth R. Troske, *University of Missouri—Columbia*

We use a unique new data set that combines data on individual workers and their employers to estimate marginal productivity differentials among different types of workers. We then compare these to estimated relative wages, leading to new evidence on productivity-based and nonproductivity-based explanations of the determination of wages. Among our findings are (1) the higher pay of prime-aged workers (aged 35–54) and older workers (aged 55+) is reflected in higher point estimates of their relative marginal products, and (2) for

We are grateful to Joe Altonji, Zvi Griliches, Bruce Meyer, Dek Terrell, and seminar participants at the University of Chicago, Hebrew University, Michigan, Michigan State, Northwestern, New York University, the University of California, Los Angeles, the University of California, San Diego, the Federal Trade Commission, the Census Bureau, and the NBER for helpful comments, and to Kim Bayard and Daniel Hansen for research assistance. The opinions expressed herein are solely those of the authors and do not in any way reflect the views of the U.S. Census Bureau. This research was supported by NSF grant SBR95-10876. Neumark's research was also supported by NIA grant K01-AG00589.

[*Journal of Labor Economics*, 1999, vol. 17, no. 3]

0734-306X/99/1703-0006$02.50

the most part, the lower relative earnings of women are not reflected in lower relative marginal products.

I. Introduction

The existence of wage differentials across workers in different demographic groups has been documented in many empirical studies. Three types of differences across workers and the reasons behind them have received a great deal of attention and are the focus of this article. First, estimates of wage differentials associated with age or experience are used to examine implications of human capital models of wage growth. Second, estimates of wage differentials associated with sex or race are used to test for wage discrimination. Third, estimates of wage differentials associated with marriage have been interpreted as reflecting productivity effects. Additional areas of inquiry include wage differentials by union status, education, and industry.

The problem with the traditional approach of estimating wage regressions to test theories of wage determination is that, without independent measures of worker productivity, it is difficult to determine whether wage differentials associated with worker characteristics reflect productivity differentials or some other factor, such as discrimination. For example, with data only on wages and worker characteristics over the life cycle, it is difficult to distinguish human capital models of wage growth (such as Ben-Porath 1967; Mincer 1974; Becker 1975) from incentive-compatible models of wage growth (Lazear 1979) or forced-saving models of life-cycle wage profiles (Loewenstein and Sicherman 1991; Frank and Hutchens 1993). Typical wage regression results report positive coefficients on age, conditional on a variety of covariates, but these positive coefficients neither imply that older workers are more productive than younger ones, nor that wages rise faster than productivity. Similarly, without direct measures of the relative productivity of workers, discrimination by sex, race, or marital status cannot be established based on significant coefficients on sex, race, or marital status dummy variables in standard wage regressions, since the usual individual-level wage regression controls may not fully capture productivity differences (e.g., Becker 1985).

The major contribution of this article is to use a unique new data set that combines data on individual workers with data on their employers to estimate relative marginal products for various groups of workers, which we then compare with relative wages. This employer-employee data set, the Worker Establishment Characteristics Database (WECD), matches long-form respondents to the 1990 Decennial Census of Population to data on their employers from the Longitudinal Research Database (LRD). These data are a major improvement over previously available data sources because they combine detailed demographic information on workers in a sample of plants with information on plant-level inputs and outputs.[1] We use these data to estimate production functions in which workers with different demographic characteristics have potentially different marginal products, thereby obtaining estimates of these relative marginal productivities. In addition, we explore numerous issues regarding the estimation of these production functions in an attempt to obtain reliable estimates of these productivity differentials. For the most part, we find that our estimates of marginal productivity appear relatively robust and reasonable, although, not surprisingly, they do change somewhat as we vary our specification and sample.

Because we have information on plant labor costs, we also specify and estimate plant-level earnings equations. These plant-level earnings equations represent the aggregation of individual-level earnings equations over workers employed in a plant, and hence are the plant-level counterparts to the individual-level wage regressions that motivate this research. By simultaneously estimating the production functions and earnings equations at the plant level, we can compare the relative marginal products and relative wages of workers distinguished by various demographic characteristics.[2] Thus, the data and empirical framework we develop supply the independent productivity measures needed to draw more decisive conclusions on numerous topics regarding the determination of wages, in-

[1] However, they are somewhat limited in that they are only cross sectional, only cover the manufacturing sector, and are weighted toward large plants.

[2] The WECD is a very rich and useful data set, and has so far been utilized only in a few other studies (Carrington and Troske 1998; Troske 1999). There are clearly many important issues that these data may be able to address; we limit this article solely to the analysis of the relationship between the productivity and wage differentials among workers with different demographic characteristics. This article builds on the framework used in Hellerstein and Neumark (1995, 1999) to analyze Israeli manufacturing data (although the WECD offers numerous advantages over the Israeli data), and it represents a departure from most of the existing empirical literature on wage determination. As discussed in Hellerstein and Neumark (1995, 1999) there is little existing research comparing productivity and wage data, and even less using firm-level data. Brown and Medoff (1978) estimate a production function using state-by-industry level data to test whether the union wage premium is associated with higher productivity of union labor. Leonard (1984) uses similar data over time to examine the impact of affirmative action laws on productivity in the United States. One firm-level productivity and wage study examines evidence of sex discrimination using data from the nineteenth-century French textile industry (Cox and Nye 1989). Studies applied to more narrowly-defined industries have been pursued in the union literature (Clark 1980; Allen 1984). Other research has used proxies for productivity, including using piece-rate pay to measure productivity in time-rate work (Foster and Rosenzweig 1993) and performance ratings (Medoff and Abraham 1980; Holzer 1990; Korenman and Neumark 1991).

cluding race and sex discrimination in wages, the causes of rising wages over the life cycle, and the returns to marriage.

II. The Relationship between Wages and Productivity

In order to motivate the approach we take in this article, we first present the simplest model illustrating the relationship between wages and productivity under perfect competition. Consider an economy consisting of plants that produce output Y with a technology that utilizes two different types of perfectly substitutable labor inputs, L_1 and L_2. The production function of these plants is

$$Y = F(L_1 + \phi L_2), \tag{1}$$

where ϕ is the marginal productivity of L_2 relative to L_1. These plants are assumed to operate in perfectly competitive spot labor markets, and labor supply is assumed to be completely inelastic. The price of the output Y is normalized to equal one. Wages of workers of types L_1 and L_2 are w_1 and w_2, respectively. Define the relative wage rate (w_2/w_1) to be λ. Given this setup, the proportional mix of the two types of labor in each plant will be determined by the relationship between ϕ and λ. If $\phi = \lambda$, then under profit maximization or cost minimization plants will be indifferent to the proportional mix of the two types of labor in the plant. If there is a wedge between the relative marginal product and relative wage so that $\phi \neq \lambda$, then profit-maximizing or cost-minimizing plants will be at a corner solution, hiring either only workers of type L_1 (if $\phi < \lambda$) or only workers of type L_2 (if $\phi > \lambda$). The only equilibrium in this model is when wages adjust so that $\phi = \lambda$, and plants are indifferent between the two types of labor.

Evidence that $\phi \neq \lambda$ is inconsistent with the assumption that we are observing profit-maximizing or cost-minimizing plants in a competitive spot labor market.[3] This article can be interpreted as providing empirical tests of this characterization of labor markets. We estimate variants of the plant-level production function in equation (1) simultaneously with plant-level wage equations in order to obtain estimates of parameters corresponding to ϕ and λ for various types of workers. We interpret cases where we cannot reject the equality of ϕ and λ as evidence consistent with competitive spot labor markets. Cases in which we reject the equality of ϕ and λ indicate some deviation from this characterization of labor markets, such as long-term incentive contracts or discrimination.

[3] Labor supply could be less than completely inelastic; as long as market wages remain above reservation wages, the conclusions are unchanged.

III. The Data

The WECD, constructed at the U.S. Census Bureau, links information for a subset of individuals responding to the long form of the 1990 Decennial Census of Population with information about their employers in the 1989 LRD. Long-form Census of Population respondents report the location of their employer in the prior week and the type of business or industry in which they work. The Census Bureau then assigns a code for the location of the employer, corresponding to a unique city block for densely populated areas, or corresponding to a unique place for sparsely populated areas. The Census Bureau also classifies workers into industries using census industry codes so that respondents can be assigned to a unique industry-location cell. In addition, the Census Bureau maintains a complete list of all manufacturing establishments operating in the United States in a given year, along with location and industry information for these establishments that is similar to the data available for workers. Thus, it is possible to assign all plants in the United States to an industry-location cell. The WECD is constructed by first selecting all manufacturing establishments in operation in 1990 that are unique in an industry-location cell. Then all workers who are located in the same industry-location cell as a unique establishment are matched to that establishment. This results in a data set consisting of 199,558 workers matched to 16,144 plants.

To obtain data on a worker's employer, these data must be matched to the plant-level data in the LRD. The LRD is a compilation of plant responses to the Annual Survey of Manufacturers (ASM) and Census of Manufacturers (CM). The CM is conducted in years ending in a two or a seven, while the ASM is conducted in all other years for a sample of plants. The LRD contains plant data from every CM since 1963 and every ASM since 1971. Data in the LRD are of the sort typically used in production function estimation, such as output, capital stock, materials expenditures, and number of workers. In addition, the LRD contains information on total salaries and wages and total nonsalary compensation paid by the plant in a given year (McGuckin and Pascoe 1988).

Since worker earnings and labor force information in the Decennial Census of Population refer to 1989, we match the worker data to the 1989 plant data in the LRD. Since 1989 is an ASM year, data are only available for a sample of plants. Furthermore, since plant-level capital stock information is only available in Census of Manufacturers years, we require all plants to be in the LRD in both 1989 and 1987.[4] Finally, to increase the representativeness of the sample of workers in each plant, we require plants in our data set to have at least 20 employees in 1989 (as reported in the LRD), and at least 5% of their workforce contained in the WECD. Our final sample contains data

[4] Total capital in the plant is measured as the sum of the end-of-year book value of buildings and machinery in 1987.

Table 1
Descriptive Statistics for Matched Establishments

LRD Data	Mean (1)	Standard Deviation (2)
Log output ($1,000)	10.19	1.33
Log value added ($1,000)	9.34	1.39
Log capital ($1,000)	8.82	1.53
Log cost of materials ($1,000)	9.43	1.51
Log wages and salaries ($1,000)	8.40	1.17
Log compensation costs ($1,000)	8.62	1.18
Employment	353.0	846.8
Establishment size:		
1–75 employees	.22	
76–150 employees	.25	
151–350 employees	.29	
351+ employees	.25	
Industry:		
Food products and tobacco products	.16	
Textile mill products, apparel, and leather and leather products	.07	
Lumber and wood products and furniture fixtures	.04	
Paper and allied products and printing and publishing	.16	
Chemicals and petroleum refining	.10	
Rubber and plastics	.05	
Stone/clay/glass/concrete	.04	
Primary metals	.08	
Fabricated metal products	.08	
Machinery/computer equipment	.06	
Electrical/electronic equipment	.06	
Transportation equipment	.06	
Instruments/clocks/optical goods and miscellaneous manufacturing	.03	
Region:		
Northeast	.29	
Midwest	.44	
South	.23	
West	.05	
Multiple-establishment unit	.82	

Census of Population Data	Mean (3)	Standard Deviation (4)
Log estimated wages and salaries ($1,000)	8.38	1.17
Proportion of LRD employment matched	.12	.08
Proportion with 2–5 workers matched	.14	
Proportion with 6–10 workers matched	.18	
Proportion with 11–20 workers matched	.24	
Proportion with 21–40 workers matched	.21	
Proportion with 41+ workers matched	.24	
Proportions:		
Female	.30	.23
Black	.07	.12
Aged 34 or less	.39	.20
Aged 35–54	.48	.18
Aged 55 or more	.13	.12
Some college	.36	.21
Ever married	.84	.14
Managerial/professional workers	.15	.15
Technical, sales, administrative, and service workers	.20	.15
Precision production, craft, and repair workers	.20	.15
Operators, fabricators, and laborers	.45	.22

NOTE.—There are 3,102 establishment-level observations, and 128,460 matched individuals from the Census of Population. The sample is restricted to those establishments with total employment of 20 or more, for which at least 5% of employees are matched.

on 3,102 plants and 128,460 workers. Summary statistics for plant-level data are given in table 1. The average plant has 353 employees, and on average 12% of a plant's workforce is matched to the plant.[5]

[5] We have no fewer than two workers per plant. Table 1 also reports the distribution of plants based on number of workers matched.

Troske (1998) concludes that workers are matched to their correct plants— based on the match rate and on high correlations between variables available in the two data sets—with approximately 5% of manufacturing workers from the Census of Population long-form represented in the WECD. The matching process does not, however, yield a representative sample of workers, as nonblack, male, married workers are overrepresented in the WECD. Below we discuss some of the implications of this for our empirical results.

IV. Estimating Marginal Productivity Differentials

A. Basic Approach

To estimate parameters corresponding to ϕ—the relative marginal productivities of various types of labor—we estimate a translog production function in which the value of output Y is a function of capital K, materials M, and a quality of labor aggregate QL.[6] In logs, this is

$$\begin{aligned}\ln(Y) = \ln(A) + \alpha \ln(K) + \beta \ln(M) \\ + \gamma \ln(QL) + g(K, M, QL) + \mu,\end{aligned} \tag{2}$$

where $g(K, M, QL)$ represents the second-order terms in the production function (Jorgenson et al. 1973), and μ is an error term.

For each plant in our data set, we have demographic information on a sample of the workforce from the WECD. We assume that in the quality of labor aggregate QL, workers with different demographic characteristics are perfectly substitutable inputs with potentially different marginal products.[7] For example, assume that workers are distinguished only by sex. Then QL would be defined as

$$QL = L\left(1 + (\phi_F - 1)\frac{F}{L}\right), \tag{3}$$

where L is the total number of workers in the plant, F is the number of women in the plant, and ϕ_F is the marginal productivity of women relative to men. Substituting equation (3) into equation (2), we obtain a production function with which we can estimate ϕ_F, using plant-level data on output, capital and materials inputs, and the number of workers and sex composition of the workforce.

We actually define QL to assume that workers are distinguished not only

[6] The results reported in the article were very similar when a Cobb-Douglas production function was used.

[7] Issues relating to this specification of the labor input are discussed in Rosen (1983). Below, we report some estimates dropping the perfect substitutes assumption.

by sex but also by race (black and nonblack), marital status (ever married), age (divided into three broad categories—aged under 35, 35–54, 55 and over), education (defined as having attended at least some college), and occupation (divided into four groups—[1] operators, fabricators, and laborers [unskilled production workers]; [2] managers and professionals; [3] technical, sales, administrative, and service; and [4] precision production, craft, and repair). A firm's workforce can then be fully described by the proportions of workers in each of 192 possible combinations of these demographic characteristics.

To reduce the dimensionality of the problem, for much of our work we impose two restrictions on the form of QL. First, we restrict the relative marginal products of two types of workers within one demographic group to be equal to the relative marginal products of those same two types of workers within another demographic group. For example, the relative productivity of black women to black men is restricted to equal the relative marginal productivity of nonblack women to nonblack men. Similarly, the race difference in marginal productivity is restricted to be the same across the sexes. Second, we restrict the proportion of workers in an establishment defined by a demographic group to be constant across all other groups; for example, we restrict blacks to be equally represented in all occupations, education levels, marital status groups, and so forth. We impose these restrictions due to data limitations. For each establishment, we do not have data on the actual number of workers in each of the 192 possible combinations of demographic characteristics, but instead estimate that number using our sample of workers matched to the plant. It is likely, therefore, that we cannot obtain accurate estimates of the representation of workers in narrowly defined sets of demographic groups. Our restrictions on QL reduce the number of sample estimates based on small numbers of workers, as well as the number of parameters.

With these assumptions, the log of the quality of labor term in the production function becomes

$$\begin{aligned}\ln(QL) = \gamma \ln\Bigg\{ &[L + (\phi_F - 1)F]\left[1 + (\phi_B - 1)\frac{B}{L}\right]\left[1 + (\phi_R - 1)\frac{R}{L}\right] \\ &\times \left[1 + (\phi_G - 1)\frac{G}{L}\right]\left[1 + (\phi_P - 1)\frac{P}{L} + (\phi_O - 1)\frac{O}{L}\right] \\ &\times \left[1 + (\phi_N - 1)\frac{N}{L} + (\phi_S - 1)\frac{S}{L} + (\phi_C - 1)\frac{C}{L}\right]\Bigg\},\end{aligned} \tag{4}$$

where B is the number of black workers, R is the number of workers ever married, G is the number of workers who have some college education, P is the number of workers in the plant between the ages of 35 and 54, O is the number of workers who are aged 55 or older, and N, S, and C are the numbers of workers in the second through fourth occupational categories defined above.[8] Note that the way QL is defined, productivity differentials between groups are indicated when the estimate of the relevant ϕ is significantly different from one (rather than zero). For example, a finding of $\phi_R = 1.3$ would imply that ever-married workers are 30% more productive than never-married workers.[9]

We also allow productivity to vary by size of plant (see Lucas 1978; Baily et al. 1992), industry, region, and whether or not the plant is part of a multiplant firm, by adding controls for these plant-level characteristics to the production function.[10]

B. Assessing the Robustness of the Relative Marginal Productivity Estimates

To this point, we have described the basic approach to the estimation of productivity differentials across workers in different demographic groups. Because the estimation of relative marginal productivities is the central contribution of this article, we carry out a number of additional analyses relaxing various assumptions imposed on the production function estimation, to assess the robustness of the estimates.

Among the endogeneity biases that might be most troubling is the potential endogeneity of materials. We first address this issue by

[8] For example, suppose workers are distinguished by race and sex. Then the unrestricted quality of labor term is $QL = L + (\phi_F - 1)WF + (\phi_B - 1)BM + (\phi_F \cdot \phi_B \cdot \phi_{FXB} - 1)BF$, where WF is the number of white females, BM the number of black males, and BF the number of black females. The restriction of equal relative marginal productivities implies $\phi_{FXB} = 1$. The equiproportionate distribution restriction implies $BF = B \cdot (F/L)$, $BM = B(1 - (F/L))$, and $WF = F(1 - (B/L))$. Substituting, we obtain $QL = L + (\phi_F - 1)F[1 - (B/L)] + (\phi_B - 1)B[1 - (F/L)] + (\phi_F \cdot \phi_B \cdot \phi_{FXB} - 1)B(F/L)$, which reduces to $QL = [L + (\phi_F - 1)F][1 + (\phi_B - 1)(B/L)]$, paralleling equation (4).

[9] In the text of the article, we sometimes report the estimate of ϕ, and whether it is significantly different from one, and sometimes refer to the implied percentage differential ($\phi - 1$), and whether it is statistically significant (i.e., significantly different from zero). The tables report estimates of the ϕ's.

[10] As Griliches and Ringstad (1971) point out, estimates of the first-order terms in the translog production function are not invariant to the units of the data. We therefore de-mean the (log of) capital, materials, and labor quality inputs prior to estimating the production function, so that the coefficients on the productive inputs in the production function are estimated at the mean of the sample. Following Crepon and Mairesse (1993), we de-mean the log quality of labor term, $\ln(QL)$, by first estimating the translog production function without de-meaning, constructing plant-level estimates of $\ln(QL)$, and then taking the mean over the sample of the estimated values of $\ln(QL)$. This allows us to measure the returns to scale parameter by adding up the coefficients on the linear terms. We initially also entered controls for age of plant in the production function (and wage equation). However, the estimated coefficients on these variables were individually and jointly insignificant, and their exclusion had little effect on the other estimated coefficients.

estimating a value-added version of the production function, using $\ln(Y - M)$ as the dependent variable, since the value-added specification finesses the endogeneity issue by avoiding estimation of a coefficient on materials. There are also other potential virtues of a value-added specification, as Griliches and Ringstad (1971) discuss. The value-added specification enhances comparability of data across industries and across establishments within industries, when industries or establishments differ in their degree of vertical integration. In addition, the value-added specification can be derived from quite polar production function specifications: one in which the elasticity of substitution between materials and value added is infinite (i.e., $Y = f(K, QL) + M$); and one in which this elasticity of substitution is zero (so that materials have to be used in a fixed proportion to output). The second approach to the endogeneity problem is instrumental variables estimation of the output specification of the production function, treating materials as endogenous and using data on materials usage in the 1987 CM to form an instrument for current materials use. The specifics of the approach are discussed in the empirical section.

As explained above, the way that labor enters the production function is restrictive in two senses.[11] First, the relative marginal productivities of two types of workers within one demographic group are restricted to equal the relative marginal productivities of those types of workers within another demographic group. Second, the proportion of workers defined by one demographic group is restricted to be constant across all other groups. We do, however, present estimates relaxing some of these assumptions to see if they have a substantive effect on the estimates. To explain how we relax these assumptions, consider the approach used for one specific set of estimates in which we relax both types of restrictions with regard to marriage, race, and sex.[12] In the production function, this yields a quality of labor term of the form

$$\begin{aligned} QL = [&L + (\phi_F - 1)WFS + (\phi_R - 1)WMR + (\phi_B - 1)BMS \\ &+ (\phi_R \cdot \phi_B - 1)BMR + (\phi_R \cdot \phi_F \cdot \phi_{FXR} - 1)WFR \\ &+ (\phi_F \cdot \phi_B \cdot \phi_{FXB} - 1)BFS + (\phi_F \cdot \phi_B \cdot \phi_R \cdot \phi_{FXR} \cdot \phi_{FXB} - 1)BFR] \end{aligned} \tag{5}$$

[11] The same restrictions are imposed in the plant-level wage equation estimation described below, and in the empirical section the restrictions are relaxed in the same manner described here.

[12] One motivation for this is evidence from wage equations that the marriage wage premium for men does not carry over to women and that the race differential is larger for men than for women (e.g., Corcoran and Duncan 1979).

$$\times\left[1+(\phi_G-1)\frac{G}{L}\right]\left[1+(\phi_P-1)\frac{P}{L}+(\phi_O-1)\frac{O}{L}\right]$$
$$\times\left[1+(\phi_N-1)\frac{N}{L}+(\phi_S-1)\frac{S}{L}+(\phi_C-1)\frac{C}{L}\right],$$

where WFS denotes the number of nonblack, never-married females in the plant, WMR the number of nonblack married males, BMS the number of black, never-married males, and so forth. Introducing these variables relaxes the equiproportionate restriction regarding the distribution of workers. The term ϕ_{FXR} is the coefficient on the interaction for being female and married. A finding that $\phi_{FXR} = 1$ would indicate that the marriage productivity premium for women is no different than that for men. The term ϕ_{FXB} is the interaction coefficient for black females, where $\phi_{FXB} = 1$ would imply that the productivity differential between men and women does not vary by race. Introducing these parameters relaxes the equal relative productivity restriction.[13]

Empirically, when we estimate this augmented specification we first estimate the unrestricted model, using the expression for QL in equation (5). We then reestimate the equations retaining only the "interaction" coefficients (such as ϕ_{FXR}) that are significantly different from one. We report the latter (restricted) set of estimates. However, we always use the disaggregated estimates of the distribution of workers (*WFS*, *WMR*, etc.).

In addition to the specific case considered here, we carry out estimations relaxing both types of restrictions for the following combinations of variables: sex and occupation, sex and age, and age and education. The sex by occupation restrictions are natural to question because the occupational distribution differs markedly by sex. The sex by age restrictions are natural to question because younger cohorts of women are likely to be quite different (relative to men) from older cohorts. Because of the limitations of the Census of Population we do not have data on experience (or tenure) but, instead, use age as a proxy. Younger cohorts of women are likely to have more continuous experience than older women, so that for them age should be a better proxy; by comparing sex differences by age cohort, we can assess the sensitivity of the estimates to this measurement issue. Finally, relaxing the age by education restrictions is appealing because education levels are higher for younger cohorts.

To this point, the production function has been specified so that

[13] One way to see that the formulation in equation (5) is correct is to impose these restrictions on the parameters, impose the equiproportionate assumption on the data (e.g., $WMR = (R/L)\cdot(1 - \{F/L\})\cdot(1 - \{B/L\})$), and note that the original quality of labor term in equation (4) results.

workers of different types have different marginal products but are perfectly substitutable. Because this specification may be too restrictive, we also consider evidence from estimates of a production function in which workers are imperfect rather than perfect substitutes. It seems to us most natural to separate labor inputs along occupational lines. We therefore estimate a production function of the form

$$\begin{aligned}\ln(Y) = \ln(A) + \alpha \ln(K) + \beta \ln(M) + \gamma_P \ln(QL_P) \\ + \gamma_{NP} \ln(QL_{NP}) + g(K, M, QL_P, QL_{NP}) + \mu,\end{aligned} \tag{6}$$

where the subscripts P and NP denote production and nonproduction workers, respectively, and $g(K, M, QL_P, QL_{NP})$ represents the higher-order terms in the translog production function.[14] The QL terms in equation (6) are of the same form as equation (4) but defined for the two subsets of workers.[15] We assume that the relative marginal productivities (and wages) are the same for workers in different demographic groups within each of these broad occupational groups, which permits us to focus on the consequences of relaxing the perfect substitutes assumption.

To preview the production function results, we find that the methods described in this subsection indicate that the productivity estimates are relatively robust. This suggests that we have obtained reliable estimates of marginal productivity differentials across different types of workers. Because we can supply robust estimates of the marginal productivity differentials, we can accomplish the central goal of this article, which is to compare estimated marginal productivity differentials with estimated wage differentials.

Finally, there is one fundamental identification issue regarding the estimation of plant-level production functions. As the preceding equations make clear, identification of productivity differentials associated with demographic characteristics of workers comes from covariation across plants in the demographic composition of the workforce and output. If we find evidence suggesting that, for example, women are less productive than men, the plant-level data do not enable us to determine whether the estimated lower productivity of women comes from the

[14] Production workers include the two blue-collar occupations, and nonproduction workers include the other two occupations. With this form of the production function, output is zero for any plants without workers in an occupation category. We had to drop 219 plants with either no production workers or no nonproduction workers in the matched sample of workers. Had we entered all four occupations as imperfectly substitutable labor inputs, we would have had to drop many more observations.

[15] In the wage equation that is estimated jointly with this production function, described below, we also break out production and nonproduction workers.

segregation of women into low-productivity plants,[16] with the productivities of women and men within plants roughly the same, or instead from the lower productivity of women relative to men within plants. However, when it comes to wages or earnings (for which precisely the same question arises), we can assess evidence on this question, since we have data on earnings at both the individual and plant level. We return to this point at the end of the next section when we describe our estimation of earnings differentials among workers.

V. Earnings Differentials among Workers

The goal of this article is to estimate the relative marginal products of different types of workers and then to compare these estimates with estimated relative wage differentials. There are a number of ways one could obtain the relative wage differentials. One possibility would be to use the estimates of wage differentials from some of the many papers filled with wage regressions. We prefer instead to use wage regression estimates computed from the sample of workers matched to plants in the WECD in order to obtain the most comparable estimates of relative wages and relative productivity. To do this, we estimate wage differentials using plant-level (rather than individual-level) earnings equations. We have chosen to focus on plant-level earnings equations (although we verify that these estimates are not at odds with individual-level wage equation estimates) for three reasons. First, we are ultimately interested in testing the equality of wage and marginal productivity differentials. Focusing on plant-level wage equations allows us to jointly estimate the production function and wage equations, yielding straightforward statistical tests of the equality of relative wages and relative marginal products. Second, the labor cost measure from the Census of Population (the only one available at the individual level) refers to all jobs worked in the year and, therefore, may not reflect earnings worked at the plant to which the worker is matched. Thus, the plant-level labor cost measure from the LRD is a more reliable earnings measure to use for estimating wage differentials to compare to our plant-level productivity differentials. Third, there may be some unobservables in the production function and wage equation. However, as long as we estimate both of these at the plant level, any biases from these unobservables ought to affect the estimated productivity and wage differentials similarly, at least under the null hypothesis, thus minimizing their impact on tests of the equality of relative marginal products and relative wages.

To establish a baseline for comparison with other data sets, table 2

[16] This could parallel evidence that women are crowded into low-wage employers (e.g., Blau 1977).

Table 2
Individual-Level Census of Population Log Earnings Regressions

	Specifications with Usual Individual-Level Controls (1)	Specifications with Variables Used in Plant-Level Analysis (2)	Fixed Plant Effects (3)
Individual-level variables:			
Female	−.35	−.38	−.32
	(.003)	(.003)	(.003)
Black	−.05	−.03	−.08
	(.01)	(.01)	(.01)
Age	.08	. . .	. . .
	(.001)		
$\text{Age}^2 \times 10^{-2}$	−.08	. . .	. . .
	(.001)		
Age 35–54	. . .	.24	.19
		(.003)	(.003)
Age 55+	. . .	.23	.18
		(.005)	(.004)
Ever married	.14	.28	.25
	(.004)	(.004)	(.004)
Highest degree attained:			
High-school diploma	.14	. . .	. . .
	(.004)		
Some college/no degree	.19	. . .	. . .
	(.004)		
A.A. degree	.22	. . .	. . .
	(.01)		
B.A. or B.S. degree	.37	. . .	. . .
	(.01)		
Advanced degree	.47	. . .	. . .
	(.01)		
Some college or higher	. . .	.15	.11
		(.003)	(.003)
MSA	.13	. . .	. . .
	(.003)		
Log establishment employment	.08	. . .	. . .
	(.001)		
Dummy variables included for:			
Region (4)	Yes	Yes	. . .
Occupation (one-digit)	Yes	No	No
Occupation (4)	No	Yes	Yes
Industry (two-digit)	Yes	No	. . .
Industry (13)	No	Yes	. . .
Establishment size (4)	No	Yes	. . .
R^2	.48	.39	. . .

NOTE.—The dependent variable is log earnings. Standard errors of the estimates are reported in parentheses. The sample size is 128,460. The sample includes all individuals matched to the establishments used in the analysis in the following tables. Less than high-school diploma is the omitted education category. MSA = metropolitan statistical area.

reports individual-level earnings regression estimates using the individual workers in the WECD. These regressions obviously use the Census of Population earnings measure, which is earnings on all jobs in the year. The first column reports estimates from a standard earnings regression. The estimates display results common to numerous other data sets. There is a significant wage gap between men and women, and a smaller but still significant wage gap between blacks and nonblacks.[17] The estimated life-cycle wage profile has the usual quadratic shape and the positive marriage premium (of 14%) parallels that found elsewhere.

As explained above, to get more reliable estimates of the demographic composition of plants' workforces, in the plant-level estimation we define workers' demographic characteristics more broadly than is typical for individual-level wage equations. In order to provide direct comparability between individual-level wage equations and the plant-level equations we discuss below, column 2 of table 2 reports the results of the individual-level regression using the more aggregated forms of these characteristics, including using cells for age ranges and establishment size, and using more limited education, occupation, and industry controls. The only major qualitative difference is that the magnitude of the estimated marriage premium increases. Other than that, the estimated coefficients of race and sex scarcely change, and the estimated age coefficients broadly reflect the quadratic shape from column 1.

To obtain a plant-level wage equation, for most of the analysis we retain the equiproportionate distribution restriction made in defining QL in the production function. We also (again paralleling the production function) restrict the relative wages of workers within a demographic group to be constant across all other demographic groups. Furthermore, we assume that all workers within each unique set of demographic groupings are paid the same amount, up to a plant-specific multiplicative random error. Under these assumptions, total log wages in a plant can be written as

$$\begin{aligned}\ln(w) = a' + \ln\Big\{ & [L + (\lambda_F - 1)F]\left[1 + (\lambda_B - 1)\frac{B}{L}\right] \\ & \times\left[1 + (\lambda_R - 1)\frac{R}{L}\right]\left[1 + (\lambda_G - 1)\frac{G}{L}\right] \\ & \times\left[1 + (\lambda_P - 1)\frac{P}{L} + (\lambda_O - 1)\frac{O}{L}\right] \\ & \times\left[1 + (\lambda_N - 1)\frac{N}{L} + (\lambda_S - 1)\frac{S}{L} + (\lambda_C - 1)\frac{C}{L}\right]\Big\} + \varepsilon,\end{aligned} \tag{7}$$

where a' is the log wage of the reference group (nonblack, never married,

[17] A race-wage gap of this magnitude (5%) is standard in manufacturing and suggests that we may be unable to detect significant differences between blacks and nonblacks in plant-level estimates of wage equations (and production functions).

male, no college, young, unskilled production worker) and the λ terms represent the relative wage differentials associated with each characteristic. To see that this plant-level equation can be interpreted as the aggregation over workers in the plant of the individual-level wage equation, consider a simpler version of the wage equation involving only men and women. The total wage bill in levels implied by equation (7) is

$$w = w_M(L - F) + w_F F, \tag{8}$$

where w_M and w_F are the average wages of men and women. This can be rewritten as

$$w = w_M(L - F) + \lambda_F w_M F = w_M[L + (\lambda_F - 1) F], \tag{8a}$$

which in logs is

$$\ln w = a' + \ln[L + (\lambda_F - 1) F], \tag{8b}$$

as in equation (7), where $a' = \ln(w_M)$.

Next, consider the individual-level wage equation in levels

$$w_i = w_M M_i + w_F F_i, \tag{9}$$

where M_i and F_i are dummy variables for men and women, respectively. Clearly, the aggregation of this equation over all workers in the plant yields equation (8), from which, as we have shown, equation (7) can be derived.

We interpret equation (7) not as a behavioral equation but simply a definitional one. It assumes that all plants are wage takers in a competitive labor market so that wages do not vary systematically across plants.[18] In order to relax this assumption somewhat, in the empirical analysis we allow wages to vary systematically with industry and plant size.[19]

We actually have three compensation measures available in our data set: the plant's total annual wage and salary bill as reported in the LRD; the plant's total annual wage and salary bill plus expenditures on nonwage compensation as reported in the LRD; and an estimate of the plant's total annual wage and salary bill derived from Census of Population data on the sample of workers matched to the establishment, which, as noted above, refers to all jobs worked in the year. For simplicity, in the following discussion we refer to each of these measures as the plant's total

[18] As discussed in Section II, this is the correct assumption to make given that we are testing the null hypothesis of competitive spot labor markets.

[19] We also report some results estimating the wage equation and production function for various subsets of the data, in which case wage differentials across workers are not constrained to be equal in all plants.

wages. We examined results with each of the compensation measures, although our preferred measure is the plant's wages and salaries from the LRD, as this measure avoids the problems with using Census of Population earnings and is closer to the measures used in the vast literature on individual-level wage or earnings regressions than the total labor cost measure.

As we noted at the end of the previous section, we cannot examine within-plant productivity differentials across groups of workers. However, we can use the Census of Population earnings to look at within-plant earnings differentials. We do this in column 3 of table 2 by adding plant fixed effects to the individual-level earnings equation. The estimates indicate that most of the estimated wage differentials (with the exception of those associated with race) are largely within plants.[20] Given this, it seems valid to interpret the plant-level wage equation as the plant-level aggregation of the equation for individual wages. In contrast, if the wage differentials were largely between plants, we could not interpret confidently our estimates as measuring differences between demographic groups. In the absence of measures of productivity for individual workers, we cannot test whether estimated productivity differentials also reflect primarily within-plant differentials. However, given the evidence from the within-plant wage regressions, it is reasonable to assume this to be the case.

VI. Measurement Issues

Before turning to the results, we note that two additional measurement issues arise with respect to the demographic composition of the workforce. First, we measure the percentages of labor input from each demographic group as the percentages of workers in each demographic group in the sample of workers matched to each plant. However, if workers in different demographic groups work different numbers of hours, then we will mismeasure the proportion of labor supplied by workers in different groups. For example, if women on average work fewer hours per week than men, then we will overstate the female labor input and will underestimate the relative productivity of female labor (ϕ_F).[21] An alternative is to calculate the percentage of labor supplied by each type of worker using data on weeks worked and usual hours worked per week, as reported in the Census of Population, to construct annual hours estimates. In results not reported in the tables, when we measure the percentages of labor input for each demographic group using annual hours we obtain the

[20] Groshen (1991) finds a larger role for between-plant wage variation in the male-female wage gap. However, her results are not very comparable to ours. First, she has much finer controls for occupation, and second, she studies only five detailed industries, three of which are not in manufacturing.

[21] The same problem will arise in the plant-level wage equation.

expected result: women have higher estimated relative marginal products and relative earnings than when we do not use hours data to measure labor inputs. However, the Census of Population data measure weeks and hours on all jobs worked in the year. Because men may be more likely than women to hold multiple jobs, the annual hours estimate we construct from the Census of Population may go too far in adjusting hours supplied to the plant. Indeed, there is some evidence that this is occurring. Plant-level estimates of the male-female wage difference that we obtain when we use earnings data from the Census of Population are larger than those we obtain when we use the plant-level earnings measure from the LRD. This suggests that men's earnings in the Census of Population come partly from hours worked in jobs other than those in the plants to which we match these men. As a result, we have chosen to maintain the measurement of demographic composition using simply the number of workers of each type employed at the plant. We note, however, that with respect to the questions of substantive interest—such as whether there is a difference between the relative marginal product and relative earnings of women—the hours corrections did not affect our conclusions.

The second measurement problem arises because we treat the percentages of workers in each demographic group as known for the purposes of estimation, although they are in fact estimated from a sample. This measurement error may bias the estimated productivity and wage differentials. The measurement error in the estimates of the percentage of workers in each demographic group is likely to affect both the productivity and wage equations, however, and it is the comparison between corresponding coefficients in the two equations that is of primary interest.[22] Nonetheless, to the extent that productivity differentials across workers may be of independent interest, and to the extent that measurement error may under some circumstances bias coefficients differently in the wage and productivity equations, it is an issue that merits consideration. To explore the consequences of measurement error, we carry out a Monte Carlo simulation that essentially mimics the sampling we do in constructing the WECD. The precise methods are described below. To preview our results, we find what we regard as relatively small biases in the estimated productivity (and wage) differentials, which do not affect the qualitative conclusions.

[22] In the textbook linear model, the measurement error bias is proportional to the signal to signal-plus-noise ratio, and, therefore, under the null of equality of relative marginal products and wages, the bias would be the same in the two equations.

VII. Plant-Level Estimates of Marginal Productivity and Wage Differentials

Up to this point we have described our methods of estimating relative marginal productivities in detail and have also explained how we estimate corresponding relative wages. With these estimates in hand, it is straightforward to test whether relative wages across workers in different demographic groups reflect differences in marginal productivity. Because the production functions and wage equations are estimated jointly and because we test for differences between them, we present the production function estimates and wage equation estimates together in each of the next four tables.

Our main results are presented in tables 3–6. Each table presents joint estimates of equation (4) and equation (7), for alternative specifications of the production function and, when appropriate, corresponding differences in the specification of the earnings equation. The parameters are estimated using nonlinear least squares. In each case, we present p-values for Wald tests of the equality of the estimates of the ϕ's (the productivity differentials) with the corresponding estimates of the λ's (the wage differentials).

Table 3 begins by reporting basic results of the joint estimation of the wage and productivity equations. In this table, we use LRD wages and salaries to measure earnings. We first discuss the production function estimates. The coefficient for females indicates that women are somewhat less productive than men, with an estimate of ϕ_F of 0.84, which is significantly less than one. The estimates indicate that blacks are if anything slightly more productive than nonblack workers, as the estimate of ϕ_B is 1.18, although this estimate is not significantly different from one. The estimated age profile suggests that productivity increases somewhat with age, although the differences are not statistically significant. Finally, the estimates indicate that married workers are considerably more productive than unmarried workers, with an estimated relative marginal productivity of 1.45, which is significantly different from one.

The other estimates are of some interest. The estimated coefficients on the capital, materials, and labor inputs provide evidence against a Cobb-Douglas specification of the production function, since many of the estimated coefficients on the higher-order terms are significantly different from zero. The estimated coefficients of the relative marginal productivity of more educated workers and workers in higher-skilled occupations (the omitted occupation is operators, fabricators, and laborers, or unskilled blue-collar workers) are in line with expectations, as all of these exceed one, some significantly.

The estimates of relative wage differentials are reported in column 2. The estimates indicate that women's wages are 45% lower than men's, a difference that is strongly statistically significant and is consistent with

Table 3
Joint Production Function and Wage Equation Estimates: Translog Output Production Function, Using LRD Wages and Salaries

	Log(Output) (1)	Log(Wages and Salaries) (2)	p-Value, Column 1 = Column 2 (3)
Demographic characteristics:			
Female	.84	.55	.00
	(.06)	(.02)	
Black	1.18	1.12	.63
	(.14)	(.05)	
Aged 35–54	1.15	1.19	.71
	(.11)	(.04)	
Aged 55+	1.19	1.18	.95
	(.15)	(.05)	
Ever married	1.45	1.37	.68
	(.21)	(.07)	
Productive inputs:			
Log capital	.05	...	
	(.01)		
Log capital × log capital	.021	...	
	(.008)		
Log capital × log materials	−.03	...	
	(.01)		
Log capital × log labor quality	.014	...	
	(.009)		
Log materials	.59	...	
	(.02)		
Log materials × log materials	.15	...	
	(.01)		
Log materials × log labor quality	−.12	...	
	(.01)		
Log labor quality	.34	...	
	(.02)		
Log labor quality × log labor quality	.11	...	
	(.02)		
Other worker controls:			
Some college	1.67	1.43	
	(.16)	(.04)	
Managerial/professional	1.13	1.00	
	(.14)	(.04)	
Technical, sales, administrative, and service	1.27	1.11	
	(.12)	(.04)	
Precision production, craft, and repair	1.06	1.02	
	(.12)	(.04)	

NOTE.—Standard errors of the estimates are reported in parentheses. The sample size is 3,102. Estimates of the intercept are not reported. Test statistics are from Wald tests. The excluded occupation is operators, fabricators, and laborers. Other control variables included in the production function are industries (13), size (4 categories), region (4), and establishment part of multiplant firm. Other control variables included in the wage equation are industries (13), size (4 categories), and region (4). These control variables were selected by estimating the production function and wage equation jointly without the demographic controls and retaining those sets of control variables that were jointly significant. The model is estimated with the data transformed so that output is homogeneous of degree S in the inputs, where S is the sum of the coefficients of the linear terms for the production function inputs. For variables that enter linearly, we use deviations from the means. For variables that enter nonlinearly, we first estimate the model using the data in levels and then take deviations from the means of the nonlinear terms. This two-step procedure is valid because the estimated coefficients of all of the nonlinear terms are invariant to the deviations from the mean transformation.

Table 4
Joint Production Function and Wage Equation Estimates: Translog Value-Added Production Function, and Instrumental Variables Estimates of Translog Output Production Function, Using LRD Wages and Salaries

	Translog Value-Added Production Function, NLLS			Translog Output Production Function, IV		
	Log (Value Added) (1)	Log (Wages and Salaries) (2)	p-Value, Column 1 = Column 2 (3)	Log (Output) (4)	Log (Wages and Salaries) (5)	p-Value, Column 4 = Column 5 (6)
Demographic characteristics:						
Female	.83	.56	.00	1.01	.55	.00
	(.07)	(.02)		(.10)	(.02)	
Black	1.13	1.12	.92	1.19	1.12	.66
	(.15)	(.05)		(.18)	(.05)	
Aged 35–54	1.25	1.20	.65	1.23	1.20	.87
	(.12)	(.04)		(.21)	(.04)	
Aged 55+	.93	1.18	.07	1.17	1.18	.94
	(.14)	(.05)		(.20)	(.05)	
Ever married	1.72	1.37	.22	1.48	1.39	.73
	(.29)	(.07)		(.29)	(.07)	
Productive inputs:						
Log capital	.16	. . .		.07	. . .	
	(.01)			(.02)		
Log capital × log capital	.04	. . .		.06	. . .	
	(.01)			(.01)		
Log capital × log materials	. . .	. . .		−.030	. . .	
				(.014)		
Log capital × log labor quality	−.027	. . .		−.04	. . .	
	(.019)			(.03)		
Log materials	. . .	. . .		.60	. . .	
				(.05)		
Log materials × log materials	. . .	. . .		.17	. . .	
				(.01)		
Log materials × log labor quality	. . .	. . .		−.14	. . .	
				(.02)		
Log labor quality	.82	. . .		.32		
	(.04)			(.13)	. . .	
Log labor quality × log labor quality	.03	. . .		.19	. . .	
	(.04)			(.08)		
Other worker controls:						
Some college	1.73	1.44		1.61	1.44	
	(.17)	(.04)		(.29)	(.04)	
Managerial/professional	.86	.99		1.19	.99	
	(.12)	(.04)		(.17)	(.04)	
Technical, sales, administrative, and service	1.13	1.11		1.28	1.11	
	(.12)	(.04)		(.17)	(.04)	
Precision production, craft, and repair	1.00	1.02		1.02	1.02	
	(.12)	(.04)		(.15)	(.04)	

NOTE.—See the note to table 3 for details. NLLS = nonlinear least squares. For the IV estimation, since the production function and wage equation are nonlinear, there is no clear choice of instruments because identification can be achieved with both linear and nonlinear combinations of the exogenous variables in the equations and the excluded variables (in this case, the log of lagged materials). We used as our instrument list the most obvious choice given the production function and wage equation: all demographic characteristics and dummy variables, and linear and second-order combinations (squares and interactions) of log(capital), log(lagged materials), and log(total employment).

Table 5
Joint Production Function and Wage Equation Estimates: Translog Output Production Function, Using LRD Wages and Salaries, Dropping Equiproportionate Restrictions on Means and Parameters, and Relaxing Perfect Substitutes Assumption, Estimated Coefficients of Demographic Characteristics

	A. Drop Restrictions on Marriage, Sex, and Race			B. Drop Restrictions on Sex and Occupation		
	Log (Output) (1)	Log (Wages and Salaries) (2)	p-Value, (1) = (2) (3)	Log (Output) (4)	Log (Wages and Salaries) (5)	p-Value, (4) = (5) (6)
Female	.84	.55	.00	.71	.49	.00
	(.06)	(.02)		(.07)	(.02)	
Black	1.18	1.13	.70	1.22	1.13	.51
	(.14)	(.05)		(.14)	(.05)	
Aged 35–54	1.15	1.19	.73	1.15	1.20	.65
	(.11)	(.04)		(.11)	(.04)	
Aged 55+	1.19	1.18	.93	1.17	1.18	.93
	(.15)	(.05)		(.14)	(.05)	
Ever married	1.46	1.38	.68	1.47	1.38	.64
	(.21)	(.06)		(.21)	(.07)	
Female × managerial/professional	. . .	. . .		1.26	1.42	
				(.32)	(.15)	
Female × technical, sales, administrative and service	. . .	. . .		1.81	1.42	
				(.38)	(.12)	
Female × aged 35–54	. . .	. . .		. . .	. . .	
Female × aged 55+	. . .	. . .		. . .	. . .	
Aged 35–54 × some college	. . .	. . .		. . .	. . .	
Aged 55+ × some college	. . .	. . .		. . .	. . .	
Relative marginal product and wage of women within occupations:						
Managerial/professional				.89	.70	.33
Technical, etc.				1.29	.70	.01
Relative marginal product and wage of women within age groups:						
Aged 35–54						
Aged 55+						
Relative marginal product and wage of workers aged 35–54 with some college						
Relative marginal product and wage of workers aged 55+ with some college						

NOTE.—See the note to table 3 for details. In panels A–D, we dropped the equiproportionate restriction in estimating the proportion of workers in each demographic group and the corresponding restriction on the parameters and estimated the unrestricted model. We then reestimated the model, imposing those parameter restrictions that were not rejected at the 5% significance level; if either the productivity or wage interaction was significant, we retained both. These latter estimates are reported. The bottom rows report estimated productivity and wage differentials for pairings of demographic groups for which parameter restrictions were not rejected. In panel E, the sample size is 2,883. It is slightly smaller than in the previous tables because plants with either no production workers or no nonproduction workers in the matched sample of workers had to be dropped. In these specifications, the relative wages and marginal productivity of production and nonproduction workers are allowed to differ, but separate occupational differentials within the production and nonproduction groups are not included.

C. Drop Restrictions on Sex and Age			D. Drop Restrictions on Age and Education			E. Production and Nonproduction Workers as Imperfect Substitutes		
Log (Output) (7)	Log (Wages and Salaries) (8)	p-Value, (7) = (8) (9)	Log (Output) (10)	Log (Wages and Salaries) (11)	p-Value, (10) = (11) (12)	Log (Output) (13)	Log (Wages and Salaries) (14)	p-Value, (13) = (14) (15)
.99	.65	.02	.84	.55	.00	.90	.56	.00
(.15)	(.04)		(.06)	(.02)		(.07)	(.02)	
1.18	1.11	.62	1.19	1.12	.59	1.43	1.16	.10
(.14)	(.05)		(.14)	(.05)		(.17)	(.05)	
1.25	1.26	.96	1.22	1.26	.80	1.22	1.17	.70
(.14)	(.04)		(.16)	(.05)		(.12)	(.04)	
1.23	1.22	.97	1.16	1.31	.45	1.19	1.16	.84
(.18)	(.06)		(.20)	(.07)		(.15)	(.05)	
1.45	1.36	.66	1.48	1.39	.64	1.48	1.51	.88
(.21)	(.06)		(.21)	(.07)		(.22)	(.09)	
...	...		...	...		...	...	
...	...		...	...		...	...	
.75	.77		...	...		...	...	
(.17)	(.07)							
.88	.83		...	...		...	...	
(.26)	(.11)							
...	...		.86	.87		...	...	
			(.16)	(.05)				
...	...		.96	.72		...	...	
			(.26)	(.07)				
.74	.50	.01						
.87	.54	.10						
			1.05	1.10	.71			
			1.11	.94	.38			

Table 6
Joint Production Function and Wage Equation Estimates: Translog Output Production Function, Using LRD Wages and Salaries, Estimated Coefficients of Demographic Characteristics, Subsamples of the Data Set

A. High and Low Percent Female

	Plants above Median (25%) Percent Female (N = 1,508)			Plants at or below Median Percent Female (N = 1,594)		
	Log(Output) (1)	Log(Wages and Salaries) (2)	p-Value, Column 1 = Column 2 (3)	Log(Output) (4)	Log(Wages and Salaries) (5)	p-Value, Column 4 = Column 5 (6)
Female	1.01	.56	.00	1.13	.74	.10
	(.15)	(.03)		(.24)	(.07)	
Black	1.41	1.12	.17	.83	1.10	.11
	(.22)	(.07)		(.18)	(.06)	
Aged 35–54	1.05	1.13	.63	1.27	1.18	.50
	(.16)	(.06)		(.14)	(.05)	
Aged 55+	1.29	1.15	.53	1.04	1.17	.44
	(.24)	(.08)		(.17)	(.06)	
Ever married	1.56	1.42	.67	1.34	1.26	.72
	(.34)	(.11)		(.23)	(.07)	

B. High and Low Employment

	Above Median (166) Employment (N = 1,551)			Below Median Employment (N = 1,551)		
Female	1.11	.53	.00	.73	.61	.05
	(.14)	(.02)		(.07)	(.03)	
Black	1.04	1.15	.57	1.30	1.10	.24
	(.21)	(.06)		(.17)	(.07)	
Aged 35–54	1.76	1.51	.48	1.05	1.05	.92
	(.36)	(.08)		(.10)	(.04)	
Aged 55+	1.03	1.22	.58	1.20	1.13	.60
	(.35)	(.10)		(.15)	(.06)	
Ever married	1.65	1.90	.68	1.34	1.27	.69
	(.63)	(.24)		(.18)	(.07)	

NOTE.—See the note to table 3 for details.

the individual-level regression results in table 2. The estimates also indicate significantly higher wages for workers aged 35–54, aged 55+, and for married workers. In addition, although we do not focus as much on these results, paralleling the productivity results the estimated wage differential for workers with more education is strongly positive, and the wage differentials for the three excluded occupations are positive or nonnegative.

Comparing the estimated wage and marginal productivity differentials shows that for most groups of workers, the estimated differentials are statistically indistinguishable. For the age categories, marital status, and race, the p-values for the tests of the equality of the estimated wage and productivity differentials range from 0.63 to 0.95. Moreover, the point estimates of the marginal productivity and wage differentials are very close, particularly for the age estimates.[23] Thus we fail to reject the hypothesis that the wage differentials reflect differences in marginal products for these workers. This is also true for our occupational groupings. Our estimates do suggest, however, that for educated workers the difference between the estimated productivity and wage differential is marginally statistically significant (with a p-value of 0.11).[24]

The sharp departure from these results is the evidence for sex differences. The production function estimates indicate that women's marginal productivity is somewhat lower than men's, as the estimate of ϕ_F is 0.84. However, the wage equation estimates point to a larger wage gap, with an estimate of λ_F of 0.55. The p-value of 0.00 shows that we strongly reject the hypothesis of equality of productivity and wage gaps. The results imply that, on average, women's wages fall short of men's by considerably more than can be explained by their lower marginal productivity. This is consistent with the standard wage discrimination hypothesis.

The plant-level wage equation results for blacks showing that blacks earn slightly higher wages than nonblacks conflicts with the individual-level wage regressions reported in table 2 (and the commonplace finding) showing a small negative, but statistically significant, wage differential between blacks and nonblacks. The reason for this difference is probably due to upward bias in our estimates of ϕ_B and λ_B, which arises for two reasons. First, as discussed in Section III, our sample of workers underestimates the number of blacks working in manufacturing. The underrepresentation of blacks would cause the estimates of both λ_B and ϕ_B to be biased away from one. However, since we are testing for the difference between λ_B and ϕ_B, our test of the equality of the wage and productivity differentials is still valid under the null hypothesis that $\phi_B = \lambda_B$.[25] Second, the fixed-plant-effects estimate of the wage differential between blacks and nonblacks in the WECD is a bit larger (−.08) than the cross-section differential (−.05), indicating that within plants, blacks earn less than nonblacks, but that blacks work in slightly higher-paying plants (see also Carrington and Troske 1998). This suggests that the production function and wage equation estimates—which use cross-plant variation—

[23] This parallels results for Israeli manufacturing reported in Hellerstein and Neumark (1995).

[24] The estimated productivity and wage differentials by education are, unfortunately, of no help in distinguishing among the two major theories of the positive association between wages and education—the human capital and signalling models. In either model, workers with more education are more productive.

[25] Because white, male, and married workers are overrepresented in the WECD, the same argument may apply to the estimated ϕ's and λ's for women and marital status, although the problem will be most severe with respect to blacks since they make up such a small percentage of the workforce in the WECD.

mask lower relative wages paid to blacks. Nonetheless, our results should be biased toward finding no evidence of discrimination only if blacks tend to work in plants that pay relatively higher wages, but in which productivity is not relatively higher. Given that we cannot measure within-plant productivity differentials between workers, all we can conclude is that our results from the between-plant estimates are not consistent with discrimination against black workers.

The finding of equal changes in relative marginal productivity and relative wages with age is most consistent with the general human capital model of investment, in which wages rise in lockstep with productivity (Mincer 1974). In contrast, these results are less consistent with models in which wages rise faster than marginal product over the life cycle (Lazear 1979) or, as in some models of specific human capital investment, more slowly. The equality of relative marginal productivity and wages of married workers shows that the marriage wage premium reflects an underlying productivity differential and is not attributable to discrimination in favor of married workers. However, the result does not help sort out whether marriage increases the productivity of men or whether high productivity men are selected into marriage (see Korenman and Neumark 1991).[26]

We now turn to numerous analyses of the robustness of the results reported to this point, focusing in particular on alternative estimates of

[26] We carried out two other robustness checks that are largely unrelated to the estimation of marginal productivity differentials. First, we estimated the equations using LRD total compensation costs and also using Census of Population earnings. In the former case, the results were very similar to those using LRD wages and salaries. Using LRD compensation costs led to slightly higher estimates of relative earnings than with LRD wages and salaries for workers aged 35–54 (1.24) and 55+ (1.26); however, these are only about 0.07–0.09 greater than the estimated relative marginal productivities of these workers, and the p-values for the tests of equality are high (.36 for ages 35–54, and .61 for ages 55+). The finding that total compensation costs are relatively higher for older workers than wage and salary costs is not surprising, since these workers are more likely to receive costly benefits such as health insurance. Using Census of Population wages and salaries, estimated wage growth over the life cycle is considerably higher than we found using LRD earnings measures, with relative earnings of 1.43 for 35–54-year olds, and 1.32 for those aged 55 and over. In addition, for the 35–54 age group, this estimate was significantly higher than the estimate of relative marginal productivity. However, relative earnings of prime-age workers in the Census of Population (relative to the LRD) is likely to be overstated if prime-age workers have a greater tendency to report earnings from more than one job. Second, we estimated the specifications in table 3 including controls for capital and materials usage of the plant in the wage equation. These variables may help capture unobservable worker quality. Including these variables did not affect the qualitative conclusions. Some of these results are reported more fully in the working paper version of this study (Hellerstein et al. 1996).

the production function from which the estimated marginal productivity differentials are obtained. Table 4 examines the issue of the potential endogeneity of materials in the production function, reporting estimates from the value-added specification, and instrumental variables (IV) estimates of the output production function. In the value-added estimates, reported in columns 1–3, most of the estimated productivity differentials are very similar to the results in table 3. The only differences are that the estimated relative productivity of married workers rises, and that of workers aged 55 and over falls (to a bit below one, although not significantly). Instrumental variables estimates are reported in columns 4–6. Since the production function and wage equations are nonlinear, there is no clear choice of instruments because identification can be achieved with both linear and nonlinear combinations of the exogenous variables in the equations and the excluded variables (in this case, the log of lagged materials). We used as our instrument list the most obvious choice given the production function and wage equation: all demographic characteristics and dummy variables, and linear and second-order combinations (squares and interactions) of log(capital), log(lagged materials), and log(total employment). We do not emphasize the IV estimation elsewhere in the article because if there are omitted plant fixed effects that are correlated with materials, then lagged materials is not a valid instrument (the instrument is valid, however, if the output differences over time are due to serially uncorrelated period-specific effects). However, the results reported in columns 4–6 in table 4 are similar to those in table 3, with two exceptions. First, the estimated marginal productivity of women relative to men rises to 1.01, indicating no difference in productivity. Since the estimate of relative earnings is unchanged, this only strengthens the evidence that women's relative earnings are less than their marginal products. Second, the estimated relative marginal productivity of more-educated workers falls somewhat to 1.61 and becomes less precise, so that the deviation between relative marginal products and relative earnings of these workers shrinks and is no longer significant (with a p-value of 0.55). Overall, although the value-added and IV estimation results vary somewhat from the results found in table 3, these changes are generally not large and do not affect the substantive conclusions. In particular, regardless of our specification we find evidence consistent with sex discrimination.

In panels A–D of table 5, we report results based on estimations for which we drop the restrictions of equiproportionate distributions of workers across demographic groups and equiproportionate productivity restrictions. In the estimation in panel A we define the proportions of female, married workers and female, unmarried workers directly, rather than assuming that the proportion female is the same among both married and unmarried workers. Similarly, we relax the restriction that the marginal productivity of female married workers relative to male married workers equals that of female unmarried workers to male unmarried workers. We also relax the corresponding restrictions by race. In this panel, for neither wages nor productivity were any of the estimated interaction coefficients significantly different from one. We therefore report the fully restricted model but without imposing the equiproportionate restriction on the data.[27] The estimates and test results closely parallel the corresponding estimates in table 3. Thus, the imposition in table 3 of the equiproportionate assumption on the data—at least for the demographic categories that we have considered here—has little or no effect on the estimates.

In panel B we carry out a similar exercise but relax the restrictions with regard to sex and occupation, allowing the proportion of the workforce in each occupation to vary by sex, and wage and productivity differentials to vary by sex, across occupations. Our primary interest is in the sensitivity of estimated wage and productivity differentials by sex to these restrictions—especially the restriction that the occupational distribution by sex is the same. In the unrestricted model, the interactions for managerial/professional workers and technical, sales, administrative, and service workers were significant for either the production function or the wage equation (or both), so the specification retaining these is reported in panel B. In this case, the coefficients for female in the top panel refer to unskilled workers (operators, fabricators, and laborers) and precision production, craft, and repair workers. Among these workers, the relative wage of women is less than their relative marginal productivity (0.49 vs. 0.71), and the estimates are significantly different from each other (p-value = 0.00). The last two rows of the panel report sex differences in wages and productivity for the other two occupational categories.[28] The sex gap in wages exceeds that in productivity for both occupations, but only for technical, sales, administrative, and service workers, is the difference significant. For female managerial/professional workers, although we find that relative earnings fall short of relative marginal products (0.70 vs. 0.89), the difference between productivity and wages is insignificant. The extra information that we obtain from panel B of table 5, then, is that the evidence consistent with sex discrimination comes from the nonmanagerial and nonprofessional occupations, in which 86% of the women in the sample work.

[27] Manipulation of the equations in n. 8 above shows that this leads to a different specification from the fully restricted version.

[28] We obtain these by multiplying the parameter estimates for the reference group by the estimated interaction parameters. For example, the marginal productivity of managerial/professional women (relative to managerial/professional men) is $0.71 \times 1.26 = 0.89$.

Panels C and D, in turn, relax these restrictions for sex differentials by age and age differentials by education. In both cases, the interactions on the wage differentials are significantly different from one, so we report these richer specifications in the table. Looking first at panel C, we find that the relative marginal productivities and relative wages of older cohorts of women are lower than for the youngest cohort. This could occur because age overstates experience more for older cohorts, so that conditional on age older women are less productive (and paid less). For the youngest and prime-aged cohorts, the evidence that the sex gap in wages exceeds the productivity gap remains very strong, with p-values of 0.01 or less. For the oldest workers the estimates point in the same direction, but because the marginal productivity differential is imprecisely estimated, the p-value rises to 0.10. Panel D introduces interactions of age and education to allow for educational differences across age cohorts. For the oldest group the point estimates suggest wage profiles rising faster than marginal productivity for the less-educated group (with relative earnings of 1.31 vs. relative marginal productivity of 1.16) and the opposite for the more-educated group (with relative earnings of 0.94 vs. relative marginal productivity of 1.11). However, the estimates of this richer model are sufficiently imprecise that the statistical conclusion is unchanged, as we fail to reject the equality of changes in relative marginal products and wages with age for either education level.[29]

The estimates in panel E are for the specification of the production function in which production and nonproduction workers are imperfect substitutes. This has relatively little impact on the estimated relative marginal products of different types of workers. The estimated marginal productivity of women rises slightly to 0.90. The corresponding estimate rises more for blacks, but this estimate is relatively imprecise. The estimated productivity differentials for older and married workers are similar to those in table 3 in which perfect substitutability is imposed. Although not reported in the table, it is of interest that for this specification the estimated relative marginal product of more-educated workers falls to 1.43, virtually the same as the estimated relative earnings of these workers (1.44). Thus, the differential between relative marginal products and earnings for more-educated workers appears due to the perfect substitutes production function specification. On the other hand, this panel shows that imposing the perfect substitutes assumption has little effect on the

[29] Interestingly, the marginal productivity and wage differentials by education are quite different by age group. For the young and prime-aged group, the differences between these differentials fall to 0.20 and 0.15, respectively, compared with 0.24 in table 3; neither of these is statistically significant, with p-values around 0.4. But for the oldest group the estimated relative marginal product is 1.70, and the estimate of relative earnings is 1.13, with the p-value falling to 0.14.

estimated relative marginal products and earnings for the groups of workers in which we are most interested.

The WECD only contains information on a cross section of plants and workers. Because of this, we are unable to account formally (say, through a fixed-effects analysis) for differences across plants in unobservables that may be correlated with the demographic characteristics of a plant. If these unobservable plant-level characteristics affect the productivity of workers in the plant (and also perhaps wages), then we would expect the omission of these plant-level characteristics to bias our estimates of productivity (and wage differences) across different types of workers.[30] While we obviously cannot account for all unobservable plant-level differences, we can try to get a sense of the magnitude of this problem by breaking up the sample along dimensions in which we think plants may differ. This should account for at least some differences across plants that may be related to their demographic composition.

In panel A of table 6, we divide the sample into plants with above- and below-median percentages female in the workforce, for two reasons. First, the nature or extent of sex discrimination may differ in plants with varying proportions of female workers. Second, women may disproportionately work in plants with different technologies than in plants that employ mostly men. Some of the estimates vary quite a bit across these two subsamples—in particular, the estimated relative marginal productivity of blacks, and of workers aged 35–54 and 55+. However, the estimates are not sufficiently precise that the statistical conclusions regarding these groups are affected. The estimated relative marginal productivity of women is approximately one in both subsamples, but relative earnings are estimated to be a bit higher in plants with a lower percentage of females in the workforce. This fact, coupled with the imprecision of the relative marginal productivity estimate, means that we cannot strongly reject the hypothesis of the equality of relative marginal products and relative wages (p-value = 0.10) for women in these plants.

One particular hypothesis regarding technological differences and the percent female is that the percent female is highest in plants that have adopted technology that conserves on (generally male) production labor.[31] This would bias upward our estimates of ϕ_F because high percent female plants would also tend to be highly productive, technically sophisticated plants. However, as long as these plants take wages as given, relative wages will not be affected, hence potentially leading to spurious

[30] If these unobservables affect productivity and wages similarly, however, they should not affect the tests of equality of relative marginal products and relative wages.

[31] For evidence that technological change has reduced the proportion of production worker employment, see Berman et al. (1994) and Dunne et al. (1997).

evidence of discrimination in the tests we report in the earlier tables. If our previous estimates of ϕ_F were biased upward because of this cross-plant variation in labor-saving technology coupled with a positive co-variation between such technology and the percent female, then when we split the sample into plants with relatively higher or lower percentages of women, the estimate of ϕ_F should fall, as we effectively condition on this technology.[32] In fact, however, the estimate of ϕ_F rises in both sub-samples.[33]

In panel B we disaggregate the plants into those with employment levels above and below the median. These results may provide some indication of possible differences in the extent of discrimination between large and small plants. The estimated degree of discrimination against women (measured by the estimate of $\phi_F - \lambda_F$) is much smaller in the smaller plants. These results suggest that smaller firms are less able to indulge in sex discrimination, which may be in part because they have less market power (Becker 1971). Some of the other estimates vary quite a bit between the two subsamples, but only in cases in which the estimates become imprecise; the qualitative conclusions are not affected.

Overall, the disaggregated estimates indicate that the estimates of relative marginal products of workers in different demographic groups are somewhat sensitive to the sample composition. However, the qualitative nature of the evidence from the separate subsamples is consistent, indicating that the full-sample estimates of the relative marginal products of different types of workers are relatively robust. In our view, the stability of these estimates across the various specifications we present provides

[32] This bias will arise if such technological change is not fully accounted for in the book value of capital (and we would not expect it to be). This problem would not arise if technological change is capital augmenting and is correctly captured in capital price deflators. While these desirable circumstances seem unlikely to hold exactly, Baily et al. (1992) find that results for total factor productivity regressions are the same using book value of capital and a more carefully constructed capital series based on initial capital stocks and annual investment data.

[33] In Hellerstein et al. (1996) we examine additional evidence on this question. For a subset of industries we have independent information on technological innovation from the Census Bureau's 1988 Survey of Manufacturing Technology (see Doms et al. 1997). This survey covered over 10,000 establishments in Standard Industrial Classification (SIC) industries 34–38 (which are high-technology industries). Over 350 establishments in the WECD can be matched to establishments in this survey. The results indicate that there is no evidence that these advanced technologies are associated with fewer production workers, and that the percent female tends to be lower, rather than higher, in plants using the advanced technologies. Thus, for this subset of industries at least, there is no evidence suggesting that the estimated relative marginal productivity of women is biased upward because the percent female tends to be higher in plants that have installed technology that saves on male production labor, or on male labor generally.

strong evidence that we are obtaining parameter estimates of marginal productivity differentials that can be meaningfully compared with estimated earnings differentials, to test alternative theories of wage differences between these types of workers. We would not argue, however, that additional research utilizing the approach we pursue in this article could not generate more reliable estimates of variations in marginal productivity across different types of workers.

VIII. The Role of Measurement Error from Matched Samples of Workers

As mentioned earlier, although we only have estimates of the percentage of workers with each set of demographic characteristics in each plant, until now we have treated these percentages as known for the purposes of estimation. In this section, we explore more fully the potential effects of measurement error that arises from estimating these percentages. Specifically, we quantify the magnitudes of measurement error biases with a Monte Carlo simulation.[34]

Consider the production function and wage equations given by equations (2), (4), and (7).[35] From the data, we know the true values of Y, K, M, and L for each plant. All of the remaining variables (F/L, R/L, etc.) are estimated from the sample of T workers within each plant. We simulate the effects of measurement error by creating a synthetic workforce of L workers for each plant. We do this by creating L/T (rounded to the nearest integer) synthetic workers for each of the T workers in the sample. With this new synthetic workforce of L workers, we sample randomly without replacement T workers and use this simulated sample to estimate the proportions of workers in each demographic group.[36] Finally, we use these simulated estimates of these proportions to jointly

[34] Although we know the number of workers sampled in each plant, we do not implement a formal correction for the measurement error bias that results from sampling error. This correction would require a consistent estimate of the variance of the measurement error, which varies by plant depending on the true proportion of workers in any particular category. (For example, at one extreme, in a plant with no female workers the variance of the measurement error in the proportion female is zero.) In other contexts, measurement-error corrections of this type (with nonhomogeneous error variances across observations) result in near-singular covariance matrices, because of a high ratio of error variance to total variance (Cockburn and Griliches 1987).

[35] In this section we do not estimate a separate coefficient for blacks versus nonblacks. The simulation requires repeated sampling of workers within a plant, and there are too few blacks in the sample to successfully estimate a race coefficient for many of the simulations.

[36] For example, suppose in the WECD we have four men and six women matched to a plant ($T = 10$), and we know from the LRD that there are 100 workers in the plant ($L = 100$). We then create a simulated sample of 40 men (4

Table 7
Measurement Error Simulation Results

	Estimated Productivity Differential (1)	Estimated Wage Differential (2)	Simulated Mean Productivity Differential (3)	Simulated Mean Wage Differential (4)
Female	.84	.55	.87	.63
	(.06)	(.02)	(.03)	(.01)
Aged 35–54	1.15	1.19	1.12	1.15
	(.11)	(.04)	(.06)	(.02)
Aged 55+	1.19	1.18	1.14	1.14
	(.15)	(.05)	(.07)	(.03)
Ever married	1.45	1.37	1.23	1.21
	(.21)	(.07)	(.08)	(.03)

NOTE.—The estimated productivity and wage differentials in columns 1 and 2 are from table 3, columns 1 and 2. The standard errors of the estimates and the standard deviations of the simulated values are reported in parentheses.

estimate the production function and wage specifications in columns (1) and (2) of table 3, obtaining new estimates of the productivity and wage differentials (the ϕ's and λ's) across demographic groups. We repeat this process 1,000 times, yielding 1,000 different values for each of the ϕ's and λ's. This procedure enables us to assess the impact on our results of measurement error in the estimated proportions of workers in each demographic group, by comparing model estimates based on the simulated data to model estimates based on the WECD data, which we treat as true. In other words, we assess the impact of measurement error on the estimated ϕ's and λ's by adding sampling error to the estimated proportions of workers in each demographic group and reestimating these parameters.[37]

Summary results of the 1,000 simulations are reported in columns (3) and (4) of table 7, and can be compared to the results from table 3 (repeated in cols. 1 and 2) to assess the magnitude of the biases caused by estimating the demographic proportions. The results indicate measurement error biases that, as expected, pull the estimated coefficients toward one, and are greater in magnitude the farther from one is the true value. For example, the mean estimate of ϕ_F in the simulations is 0.87, and the estimate using the actual data is 0.84; the mean estimate of λ_F in the

$\times$ 100/10) and 60 women (6 $\times$ 100/10). Finally, we sample 10 members of this sample and obtain a new estimate of the proportion female.

[37] Because the estimated proportions of workers in each category that we treat as known for our simulation are overdispersed relative to the true population distribution, the extent of measurement error bias indicated by our simulation method is a lower bound for the magnitude of measurement error in the real data.

simulations is 0.63, and the estimate using the actual data is 0.55. These results show that the effect of measurement error is to bias us toward finding no discernable productivity or wage differentials across workers and toward finding no differences between the relative productivity and wage estimates for a given type of worker (since estimates of parameters that are further from one have larger absolute biases toward one). Thus, the power of the tests of the equality of wage and productivity differentials is somewhat reduced because of measurement error.

We therefore conclude from the simulation that our earlier estimates indicating significant estimated gaps between relative wages and productivity for women are robust to the measurement error problem. In contrast, we consistently found evidence of the equality of relative wages and relative marginal products for married workers and for workers of different ages, using the actual data. However, the closeness of the point estimates of these premia, the relatively large estimated standard errors of the productivity differentials, and the similarity between the distributions of the simulated estimates for the wage and productivity premia in table 7, lead us to believe that even in the absence of measurement error, we would still find no significant difference between the relative wages and relative marginal products for married workers or for workers in different age categories.[38]

IX. Conclusions

Using evidence on wage differentials among different types of workers to test theories of wage determination is one of the most common avenues of research in labor economics. Often, the alternative theories being considered are "productivity-based" versus "nonproductivity-based," such as in the discrimination literature. What is almost invariably missing from these studies, however, is an independent measure of productivity. Most studies instead use observable individual-level characteristics that are presumed to be proxies for productivity.

This article uses plant-level data on inputs and outputs matched with individual-level data on workers to estimate relative marginal products of

[38] For example, for the estimates for married workers the distribution of the productivity premium bounds that of the wage premium. Below the median, the productivity premium is slightly below the wage premium; above the median, the productivity premium is greater than the wage premium. The magnitude of the estimated standard error on the marriage productivity premium in table 3 alone suggests that measurement error would have to be reducing the gap between the wage and productivity premia by a factor of five before we would reject equality of the wage and productivity marriage premia. For the age coefficients, this factor would have to be about the same for the 35–54 age group, and even larger for the 55+ age group, before we would reject equality.

workers with different demographic characteristics. Although production function estimation is a complicated task, and even more so in our case where we are adding labor quality terms that distinguish among many types of workers, we obtain relatively robust (and seemingly reasonable) estimates of these relative marginal products. We then compare these estimates of relative marginal products to estimates of relative earnings and address many of the same questions that have previously been addressed without the advantage of an independent productivity measure.

With one major exception, our basic results indicate that for most groups of workers wage differentials do, in fact, match productivity differentials. Workers who have ever been married are paid more than never-married workers, and the wage premium they receive reflects a corresponding productivity premium. This suggests that the marriage premium does not simply reflect discrimination against unmarried workers but reflects actual productivity differences. However, our data do not allow us to distinguish between the hypothesis that marriage reflects an unobservable variable associated with higher wages and the hypothesis that marriage makes workers more productive.

We find that for prime-aged workers (aged 35–54) and older workers (aged 55 and over) productivity and earnings rise at the same rate over the life cycle. Although the estimated productivity differentials by age are not very precise, they are very close to the estimated earnings differentials. This evidence is most consistent with models in which wages rise in lockstep with productivity, such as the general human capital model.

We find no evidence consistent with discrimination against blacks in manufacturing. In addition, there does not appear to be any productivity differential between blacks and nonblacks that might be attributable to premarket discrimination or other unobserved characteristics, although we are less confident in our separate estimates of the race gap in wages and the race gap in productivity than in our estimates of the difference between them.

Finally, in contrast, in nearly all of our specifications and samples, we find that women's marginal product is perhaps somewhat below that of men. But we find that women are paid significantly less than men, with the wage differential between men and women generally much larger than the productivity differential. These results are strongest for women who are not managers (most of the sample), women who work in plants that employ a lot of women, and larger plants. The statistical evidence is strong enough to reject the null hypothesis that relative marginal products and relative wages are equal, which would be implied by a spot labor market with no sex discrimination. Although there is probably no single decisive test for discrimination in labor markets, we regard this evidence — based on independent estimates of relative marginal products to compare with relative earnings — to be considerably more convincing than the evidence produced by the typical wage regression discrimination study.

References

Allen, Steven G. "Unionized Construction Workers Are More Productive." *Quarterly Journal of Economics* 91 (May 1984): 543–69.

Baily, Martin Neil; Hulten, Charles; and Campbell, David. 1992. "Productivity Dynamics in Manufacturing Plants." *Brookings Papers on Economic Activity: Microeconomics* (1992): 187–249.

Becker, Gary S. *The Economics of Discrimination.* 2d ed. Chicago: University of Chicago Press, 1971.

———. *Human Capital.* 2d ed. Chicago: University of Chicago Press, 1975.

———. "Human Capital, Effort, and the Sexual Division of Labor." *Journal of Labor Economics* 3, suppl. (January 1985): S33–S58.

Ben-Porath, Yoram. "The Production of Human Capital and the Life Cycle of Earnings." *Journal of Political Economy* 75 (August 1967): 352–65.

Berman, Eli; Bound, John; and Griliches, Zvi. "Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufacturing." *Quarterly Journal of Economics* 109 (May 1994): 367–98.

Blau, Francine D. *Equal Pay in the Office.* Lexington, MA: Lexington Books, 1977.

Brown, Charles, and Medoff, James L. "Trade Unions in the Production Process." *Journal of Political Economy* 84 (June 1978): 355–78.

Carrington, William, and Troske, Kenneth. "Interfirm Racial Segregation." *Journal of Labor Economics* 16 (April 1998): 231–60.

Clark, Kim. "The Impact of Unionization on Productivity: A Case Study." *Industrial and Labor Relations Review* 33 (July 1980): 451–69.

Cockburn, Ian, and Griliches, Zvi. "Industry Effects and Appropriability Measures in the Stock Market's Valuation of R & D and Patents." Working Paper no. 2465. Cambridge, MA: NBER, 1987.

Corcoran, Mary, and Duncan, Greg J. "Work History, Labor Force Attachment, and Earnings Differences between Races and Sexes." *Journal of Human Resources* 14 (Winter 1979): 3–20.

Cox, Donald, and Nye, John V. "Male-Female Wage Discrimination in Nineteenth-Century France." *Journal of Economic History* 49 (December 1989): 903–20.

Crepon, Bruno, and Mairesse, Jacques. "Productivite, recherche-developpement et qualifications." *INSEE-Methodes* 37–38 (1993): 181–200.

Doms, Mark; Dunne, Timothy; and Troske, Kenneth R. "Workers, Wages, and Technology." *Quarterly Journal of Economics* 112 (February 1997): 253–90.

Dunne, Timothy; Haltiwanger, John; and Troske, Kenneth. "Technology and Jobs: Secular Change and Cyclical Dynamics." *Carnegie-Rochester Conference Series on Public Policy* 46 (June 1997): 107–78.

Foster, Andrew D., and Rosenzweig, Mark R. "Information, Learning, and Wage Rates in Low-Income Rural Areas." *Journal of Human Resources* 28 (Fall 1993): 759–90.

Frank, Robert H., and Hutchens, Robert M. "Wages, Seniority, and the Demand for Rising Consumption Profiles." *Journal of Economic Behavior and Organization* 21 (August 1993): 251–76.

Griliches, Zvi, and Ringstad, Vidar. *Economies of Scale in the Form of the Production Function.* Amsterdam: North-Holland, 1971.

Groshen, Erica. "The Structure of the Female/Male Wage Differential: Is It Who You Are, What You Do, or Where You Work?" *Journal of Human Resources* 26 (Summer 1991): 457–72.

Hellerstein, Judith K., and Neumark, David. "Are Earnings Profiles Steeper than Productivity Profiles? Evidence from Israeli Firm-Level Data." *Journal of Human Resources* 30 (Winter 1995): 89–112.

———. "Sex, Wages, and Productivity: An Empirical Analysis of Israeli Firm-Level Data." *International Economic Review* 40 (February 1999): 95–123.

Hellerstein, Judith K.; Neumark, David; and Troske, Kenneth R. "Wages, Productivity, and Worker Characteristics: Evidence from Plant-Level Production Functions and Wage Equations." Working Paper no. 5626. Cambridge, MA: NBER, 1996.

Holzer, Harry J. "The Determinants of Employee Productivity and Earnings." *Industrial Relations* 29 (Fall 1990): 403–22.

Jorgenson, Dale W.; Christensen, Laurits R.; and Lau, Lawrence J. "Transcendental Logarithmic Production Frontiers." *Review of Economics and Statistics* 55 (February 1973): 28–45.

Korenman, Sanders, and Neumark, David. "Does Marriage Really Make Men More Productive?" *Journal of Human Resources* 26 (Spring 1991): 282–307.

Lazear, Edward P. "Why Is There Mandatory Retirement?" *Journal of Political Economy* 87 (December 1979): 1261–84.

Leonard, Jonathan S. "Anti-discrimination or Reverse Discrimination: The Impact of Title VII, Affirmative Action, and Changing Demographics on Productivity." *Journal of Human Resources* 19 (Spring 1984): 145–74.

Loewenstein, George, and Sicherman, Nachum. "Do Workers Prefer Increasing Wage Profiles?" *Journal of Labor Economics* 9 (January 1991): 67–84.

Lucas, Robert E., Jr. "On the Size Distribution of Business Firms." *Bell Journal of Economics* 9 (Autumn 1978): 508–23.

McGuckin, Robert, and Pascoe, George. "The Longitudinal Research Database (LRD): Status and Research Possibilities." *Survey of Current Business* (November 1988): 30–37.

Medoff, James L., and Abraham, Katherine G. "Experience, Performance, and Earnings." *Quarterly Journal of Economics* 95 (December 1980): 703–36.

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

Rosen, Sherwin. "A Note on Aggregation of Skills and Labor Quality." *Journal of Human Resources* 18 (Summer 1983): 425–31.

Troske, Kenneth R. "The Worker-Establishment Characteristics Database." In *Labor Statistics Measurement Issues,* edited by John Haltiwanger, Marilyn Manser, and Robert Topel, pp. 371–404. Chicago: University of Chicago Press, 1998.

———. "Evidence on the Employer Size-Wage Premium for Worker-Establishment Matched Data." *Review of Economics and Statistics* 81 (February 1999): 15–26.