Earnings Functions When Wages and Prices Vary by Location

EARNINGS FUNCTIONS WHEN WAGES
AND PRICES VARY BY LOCATION

DAN BLACK, NATALIA KOLESNIKOVA, AND LOWELL TAYLOR

ABSTRACT. Economists generally assume, implicitly, that "the return to schooling" is invariant across local labor markets. We demonstrate that this outcome pertains if and only if preferences are homothetic—a special case that seems unlikely. Our theory predicts that returns to education will instead be relatively low in expensive high-amenity locations. Our analysis of U.S. data provides support for this contention; returns to college are especially low in such cities as San Francisco and Seattle. Our findings call into question standard empirical exercises in labor economics which treat the returns to education as a single parameter.

JEL: J24, J31, R23.

Keywords: earnings functions, return to education, local labor markets.

I. INTRODUCTION

The development of human capital theory, and the application of this theory to the estimation of earnings functions, is a landmark contribution in applied economics. The central logic of human capital theory, as set out in the classic works of Becker (1964 and 1967) and Mincer (1974) and developed in such subsequent treatments as Willis (1986) and Card (1999 and 2001), is straightforward: Education is understood to entail an investment—in tuition, foregone earnings, and possibly loss in utility—which has a return in the form of increased earnings in the labor market. As in the theoretical treatment of most financial investments,

the only relevant prices are those that affect the cost of investment (e.g., tuition) and those that affect the return (e.g., the education-wage locus); prices of other goods and services are ignored, either explicitly or implicitly.

In this paper we revisit the seemingly innocuous practice of ignoring the general price vector in applied human capital theory. Our concern about the potential role of "other prices" stems from the observation that, unlike returns on an equity stake in General Electric, returns to an investment in human capital are usually realized in a local labor market. We demonstrate that in an equilibrium that has variation in local prices, not only do wage levels differ across locations, but so too do education-wage gradients. We show that in the U.S. such variation does exist across local labor markets. And we argue that this complication can lead to serious misunderstandings in standard empirical applications of human capital earnings functions, e.g., the identification of the causal return to education using instrumental variables, or the estimation of racial wage disparities.

Our paper proceeds in three additional sections. In Section II, we demonstrate that in a model of local labor markets the education-earnings gradient—known in labor economics as the "return to schooling"—is a constant across locations if and only if preferences are homothetic. We also show that even if preferences are homothetic, common implementation practices with Mincer earnings regressions may be problematic. When we consider the more likely case of non-homothetic preferences, theoretical reasoning leads us to believe that the observed returns to education will be particularly low in expensive cities (e.g., San Francisco, Seattle, and New York) and relatively high in inexpensive cities (e.g., Houston and Pittsburgh).

In Section III, we examine the return to college education, relative to high school education, for large cities in the U.S. in 1980, 1990, and 2000. We find substantial cross-city heterogeneity in the return to the college degree, and we find that this cross-city variation is generally persistent across decades. This heterogeneity does not appear to be the consequence of sampling variation, nor is it the result of differences across cities in the labor force age distribution, nor differences in the industry or occupation mix. We find support for the prediction that local return to schooling is inversely related to housing prices.

Section IV provides a discussion of the implications of our findings for empirical work in labor economics.

II. EDUCATION AND EARNINGS IN A MULTIPLE-LOCATION MODEL

Perhaps the most familiar analytical expression in labor economics is the Mincer earnings function,

$$(1) \qquad\qquad\qquad \ln(w_i) = \alpha + \beta E_i + \epsilon_i;$$

the expected log earnings (or wage) of individual i is a linear function of that individual's level of education (and possibly other covariates that we suppress here). The issue we treat in this section concerns the properties of the estimates of regression (1) when individuals live in locations that have differing prices. In our exploration, we set up a model of price variation across location and then ask when the education-wage gradient (β in (1)) is indeed a single parameter, i.e., we ask when the returns to education will be the same in all locations.

A. The Basic Model

We consider a model in which locations differ in attractiveness or in worker productivity. The consequence is that wages and prices differ across locations, e.g., some locations have

relatively high housing prices by virtue of their high production or consumption amenities. Many such models exist in urban economics, including the pioneering papers by Haurin (1980) and Roback (1982). Among the issues considered in these papers is the questions of where firms choose to locate (e.g., firms that are land-intensive will not want to locate where land is expensive). As will be clear below, we abstract from the question about where firms will choose to locate (by implicitly assuming that land is not an important component to production), and instead focus on an issue that is generally ignored in previous models—the implication of heterogeneity in workers' human capital.[1]

To streamline the initial presentation of our key idea, we restrict attention to the case in which local price variation is due to underlying productivity differences across location. There has been much work on possible causes of such productivity differences (e.g., see Acemoglu, 1996, Glaeser and Mare, 2001, and other work on agglomeration) and we remain agnostic as to the source of the variation. Whatever the source of productivity variation across locations, the location with higher productivity will have greater labor demand, which in turn will generally result in higher wages and housing prices. Our question is whether the return to education is likely to be the same across locations in such a model.[2]

In our model, individuals choose one of two levels of human capital, they make consumption decisions over two goods, and they choose to live in one of many cities, $j = 1, \ldots, n$. The price of one of the consumption goods is set by a national market and is thus the same

---

[1]The papers that are perhaps closest in spirit to ours are Lee (2007) and Beeson (1991). We discuss Lee's paper below. Beeson's innovative work provides documentation of large variation in local returns to schooling. For instance, using data from the 1980 CPS she finds that the return to a year of schooling is only 0.024 in Seattle, but is 0.050 in Tampa. She also finds that these rates of return are correlated with various measures of location-specific amenities. It seems that very little other work examines regional variation in the return to education. In their study of black-white education and wage disparity, Card and Krueger (1992) notice that in the U.S. the return to education is lower in the South than elsewhere. They credit Chiswick (1974) for first making this observation.

[2]Below we also study a variant of our model in which local price differences are due to a consumption amenity, i.e., we let utility depend on the consumption of some location-specific amenity.

in all cities. The price of the second consumption good, housing, varies across cities. Because productivity and housing prices differs across cities, wages of course also vary across cities. Throughout our analysis we assume, as is standard in models of human capital investment, that labor supply is fixed (at one unit) and assume further that there is no non-labor income.

Individuals have identical preferences except along one dimension—the extent to which they suffer disutility from acquiring education. In particular, preferences are characterized by a utility function $U = u(X, H) - c(E; \alpha)$, where X is the good that has a common price across the cities, H is housing, and $c(E; \alpha)$ is the utility cost of acquiring schooling level E. We let education be one of two levels, $E \in \{0, 1\}$, normalize $c(0; \alpha)$ to be 0, and suppose $\dfrac{\partial c(1; \alpha)}{\partial \alpha} > 0$ (so that the parameter α scales the cost of acquiring education). Utility maximization entails choosing the optimal level of education (0 or 1), the preferred location, and the best consumption bundle (X^*, H^*) given the education level and location.

In our analysis below we examine an equilibrium in which some, but not all, individuals optimally choose the higher level of education, and in which people of both education levels live in each city (so that we can study the nature of cross-city differences in the returns to education). In the equilibrium we describe shortly, each individual is indifferent over which city to live in, i.e., the utility (net of the education cost) must be the same in each city. Given that the education cost component of utility $c(E; \alpha)$ is independent of the city of residence, the optimal education choice in such an equilibrium is trivial to characterize: There will be some critical value of α, say α^*, such that people with $\alpha < \alpha^*$ acquire education $E = 1$ while people with $\alpha > \alpha^*$ do not. Essentially we can then treat the two levels of human capital, 0 and 1, as being predetermined. To simplify notation, we henceforth omit the education cost component from utility.

Let the non-housing good X be the numeraire, and let p_j be the rental price per unit of housing in city j. Individuals with human capital 0 earn wage w_j^0 in city j, while those with human capital 1 earn $w_j^1 > w_j^0$. Define the expenditure function for workers with human capital k ($k = 0$ or 1) living in city j: $e_j^k = e(p_j, u_j^k)$. The key equilibrium condition is that workers of both education levels must be indifferent over their city of residence; utility u_j^k, is the same in each city. We thus drop the subscript j on utility, and note that equilibrium entails, for $j = 1, \ldots, n$,

$$e(p_j, u^0) = w_j^0 \quad \text{and} \quad e(p_j, u^1) = w_j^1.$$

The gross return to education in location j—the wage of the well-educated individual relative to the poorly-educated individual—is $R_j = \dfrac{e(p_j, u^1)}{e(p_j, u^0)}$. In general, this ratio depends on the housing price p_j. Obviously, the return to education generally differs across locations, in which case we cannot ignore local prices in the empirical implementation of the human capital returns function.

When are the returns to education independent of location-specific price variation? First, note that if preferences are such that the expenditure function takes the form $e(p, u) = f(u)\psi(p)$, the return to education in location j is $R_j = \dfrac{f(u^1)\psi(p_j)}{f(u^0)\psi(p_j)} = \dfrac{f(u^1)}{f(u^0)}$, which does not depend on local prices. Second, and more importantly, note that the converse is true. The proof is simple: Let $R_j = g(u^0, u^1)$, so that the return in location j does not depend on that location's prices. Without loss of generality we can take $u^0 = 1$, $u^1 = u$. Then $R_j = \dfrac{e(p_j, u)}{e(p_j, 1)} = g(u, 1)$ and $e(p_j, u) = g(u, 1) \cdot e(p_j, 1)$. Setting $f(u) \equiv g(u, 1)$ and $\psi(p) \equiv e(p, 1)$ we obtain an expenditure function of the form $e(p, u) = f(u)\psi(p)$.

A standard result from price theory is that the expenditure function takes the form $e(p, u) = f(u)\psi(p)$ if and only if preferences are homothetic. We thus have a key proposition: *The returns to education are the same across locations if and only if preferences are homothetic.*

We can easily summarize the economic logic of our model: Utility-maximizing individuals (i) choose their level of education, 0 or 1, (ii) choose their location from among several cities, and then (iii) given their earnings in the chosen location, consume a locally-priced good H and another good X. By assumption, the utility cost of education is independent of choices (ii) and (iii), so decision (i) is trivial; individuals with a sufficiently low utility cost acquire the higher level of education. Individuals who make optimal consumption decisions (iii) are indifferent over location choices (ii).

In thinking about the equilibrium characterized by the previous paragraph, the "return to education function" in city j, $R_j = \dfrac{e(p_j, u^1)}{e(p_j, u^0)}$, is a simple welfare measure that answers this question: By what proportion do we need to increase the wage of an individual with education level 0 to make her as well off as an individual with education level 1?[3] Now equilibrium wages and housing prices differ across cities, so it is not obvious when this return function will yield the same answer for all cities. We have shown that if (and only if) preferences are homothetic, the return to education is the same in all cities.

Figure 1 illustrates our case with homothetic preferences. Suppose that workers with human capital 1 (and utility u^1), say financial analysts, sell their services in a national market, i.e., buyers who do not care where the analysts live. Suppose further that these analysts are more productive in city a than in city b, so that $w_a^1 > w_b^1$. If in equilibrium

---

[3]Of course, the answer given by the returns function does not factor in any costs of acquiring the higher level of education; it is a measure of *ex post* monetary value of having that human capital.

financial analysts live in both cities, obviously the price of housing must be higher in city a. Next, consider workers with human capital 0 (and utility u^0) who sell their labor in local markets, say janitors. These workers will locate in either city only when there is equality in the ratios $\dfrac{w_a^1}{w_a^0} = \dfrac{w_b^1}{w_b^0}$, i.e., when the (proportional) return to education is the same in the two cities.

If, in contrast, preferences are non-homothetic, returns to education differ by location; the education-wage gradients drawn in Figure 1 will not be parallel.

We could just as easily have explored a model in which location-specific differences in wages and prices are driven by differences in consumption amenities, and indeed we return to such a model shortly. The conclusion is the same as in our model with production amenities: a single return to education generally pertains across all cities only if preferences are homothetic.

As long as preferences are homothetic, it proves quite easy to relate our theory back to the familiar Mincer wage regression (1). Under homotheticity, the form of the expenditure function is $w = f(u)\psi(p)$. Using the logarithmic form of this latter equation, if person i living in city j has education k, and therefore utility level k, we have

$$\ln(w_{ij}) = \ln(\psi(p_j)) + h(u^k),$$

where $h(u^k)$ is an index of utility ($k = 0, 1$). As the $\psi(p_j)$ is independent of utility, wage levels vary with prices, but the log difference of wages (by education) remains a constant across cities. If we let E_i be an indicator variable equal to person i's education level, this leads to an earnings function of the form

$$(2) \qquad\qquad\qquad\qquad \ln(w_{ij}) = \alpha_j + \beta E_i,$$

where β, the return to the higher education level, is a constant across locations. If we add a random error term to equation (2) we have the familiar Mincer regression (1), with one exception: if housing prices and wage levels differ by location, we must include location-specific fixed effects.

In fact, in the U.S. there is large variability in housing prices, e.g., in the 1990 Census, the median housing price in New York is over three times that of the median housing price in Cleveland.[4] Not surprisingly, the wage level is higher in New York than Cleveland, e.g., men with a college degree earned about 22 percent more in New York in 1990.[5] It is clear that if researchers fail to include city fixed effects, the error term in regression (2) contains the city-fixed effects, and the OLS estimate of β will be inconsistent, except in the special case in which the distribution of education is the same across cities.[6]

In short, our theory leads us back to the traditional specification of the earnings functions *if preferences are homothetic,* and with the important caveat that we must include city fixed effects unless there is little price variation between cities (or unless the distribution of education is identical across locations).

Unfortunately, homotheticity is a strong restriction, implying that for all goods the income elasticity is equal to one. In fact, a large literature suggests that for many goods the income

---

[4]Gabriel and Rosenthal (2004) and Chen and Rosenthal (2005) show that housing prices differ widely across cities even after careful adjustment for quality.

[5]This calculation, from the 1990 PUMS, is for non-Hispanic white men aged 25 to 55 and holds constant the age distribution.

[6]In fact, in the U.S. there are large differences in the education distribution across cities (e.g., in New York there are 0.68 men with a high school education in our sample for every man with a college education, while the corresponding ratio in Cleveland is 1.43). A simple example illustrates the problem for estimating the returns to education. Suppose we have two equally-sized cities, a and b, but the ratio of high school graduates to college graduates is 3 in city a and 1/3 in city b. Suppose that in b, high school graduates earn \$30,000, while college graduates earn \$42,000, which gives a log wage difference of 0.336. In a the corresponding wages are \$50,000 and \$70,000, also a 0.336 log wage difference. Using log-linear OLS regression (without location controls), the estimated return to college is 0.592. For consistent estimation one needs to include location fixed effects.

elasticity is very different than one. Hausman, Newey, and Powell (1995), for example, estimate income elasticities of demand of 0.7 for food, 1.4 for clothing, and 1.3 for recreation. More importantly for our analysis, a large literature suggests that the income elasticity of housing differs from one.[7] If preferences are *not* homothetic, returns to education will vary across cities, i.e., there will be no single parameter β in equation (2), but rather a parameter β_j for each city. We next ask if these city-specific returns are likely to conform to a predictable pattern. We carry out our investigation using two variants of our model: First, we consider the case in which wages and housing prices vary across cities because of variation in city-specific productivity. Second, we examine the case in which local price variation stems from differences in a consumption amenity across locations.

B. The Model with Location-Specific Productivity Differences

In this variant of the model, wages will be higher in a high-productivity city than in a low-productivity city, and the housing price will of course be higher in the high-productivity city. We are interested in comparing the returns to education in the two locations.

Let u^1 and u^0 be utility levels, respectively, of individuals with high and low education (so that $u^1 > u^0$). The return to education in a city with a housing price p is $R = \dfrac{e(p, u^1)}{e(p, u^0)}$. To streamline notation we denote the expenditure function $e^k = e(p, u^k)$ for an individual with education k.

We want to know how the return in a low-price, low-productivity city compares to the return in a higher-price city. We conduct this thought experiment by evaluating the derivative

$$R_p = \frac{\partial R}{\partial p} = \frac{e_p^1 e^0 - e^1 e_p^0}{(e^0)^2},$$

---

[7]The task of measuring an income elasticity of demand on housing is complicated (see Olsen, 1987, for a discussion). One widely cited study, Rosen (1985), reports an income elasticity of demand on housing of 0.76. Harmon (1988) reviews a large number of studies and concludes that an estimate of the income elasticity of demand of 0.7 may be appropriate for most applications.

which has the same sign as the numerator

$$e_p^1 e^0 - e^1 e_p^0,$$

or, after we divide by a positive quantity $e^1 e^0/p$, as

$$e_p^1 \frac{p}{e^1} - e_p^0 \frac{p}{e^0}.$$

The last expression can be written in terms of relevant budget shares:[8]

$$s_H^1 - s_H^0,$$

which is negative if the share of income allocated to housing decline as income increases.

We conclude that if the income elasticity of housing is less than one, as suggested in the literature, then $R_p < 0$. The return to education is lower in cities that are more expensive, i.e., in the higher-productivity cities.

C. The Model with Location-Specific Consumption Amenities

The case with a location-specific consumption amenity is only slightly more complicated. In this case the price of housing is a function of the amenity, say A, and the amenity level is also an argument in the expenditure function. So the return to education is written $R = \frac{e(p(A), u^1, A)}{e(p(A), u^0, A)}$. Now our interest is comparing the return in a given city to a comparable city with a higher level of the amenity. Thus we evaluate the derivative $R_A \equiv \frac{\partial R}{\partial A}$. The return to education is lower in the higher-amenity city (which is also the more expensive city) when this derivative is negative. We are interested, therefore, in the conditions on preferences that relate to this latter inequality.

---

[8]To convert to budget shares, we use Shephard's lemma: the derivative of the expenditure function with respect to p, e_p, gives housing demand.

The derivative of interest is

$$R_A = \frac{\left(e_p^1 \frac{dp}{dA} + e_A^1\right) e^0 - \left(e_p^0 \frac{dp}{dA} + e_A^0\right) e^1}{(e^0)^2},$$

which is negative when

$$(3) \qquad e_p^1 \frac{dp}{dA} + e_A^1 - R\left(e_p^0 \frac{dp}{dA} + e_A^0\right) < 0.$$

Rearranging equation (3), we obtain

$$\frac{dp}{dA} \frac{1}{p} \left(e_p^1 \frac{p}{e^1} - e_p^0 \frac{p}{e^0}\right) + \frac{1}{A} \left(e_A^1 \frac{A}{e^1} - e_A^0 \frac{A}{e^0}\right) < 0,$$

which can be written in terms of relevant elasticities or budget shares:[9]

$$(4) \qquad \eta_A(\varepsilon_p^1 - \varepsilon_p^0) + (\varepsilon_A^1 - \varepsilon_A^0) = \eta_A(s_H^1 - s_H^0) + (\varepsilon_A^1 - \varepsilon_A^0) < 0,$$

where for utility levels $k = 0$ and 1, ε_p^k are the elasticities of the expenditure function with respect to the price of housing (which in turn equal housing budget shares s_H^k), ε_A^k are elasticities of the expenditure function with respect to the amenity level, and η_A is the elasticity of the equilibrium price of housing with respect to the amenity level.

Theoretical considerations as well as empirical observation suggest that amenities are at least partially capitalized into housing values; we expect $\eta_A > 0$. Thus, for equation (4) to hold it is sufficient that $0 < s_H^1 \le s_H^0$ and $\varepsilon_A^1 \le \varepsilon_A^0 < 0$, with strict inequality holding for one condition. The condition $0 < s_H^1 < s_H^0$ simply requires that the share of income allocated to housing decline as income increases; housing is a *necessity*. The condition $\varepsilon_A^1 < \varepsilon_A^0 < 0$ can be written $|\varepsilon_A^1| > |\varepsilon_A^0| > 0$, which requires that the value placed on the amenity be higher at the higher utility level; the amenity is a *luxury*.

---

[9]Again we use Shephard's lemma to convert to budget shares.

We can summarize the logic, starting at the beginning. In our model, people are essentially endowed with one of two "real wealth levels" in the form of innate ability to acquire education. The relatively fortunate acquire education level 1 and subsequently have higher utility than those who have education level 0 (i.e., $u^1 > u^0$). Then our analysis shows the following: If the amenity is a luxury good—in the sense that individuals with the higher wealth level also have a higher "marginal willingness to pay for the amenity"—then high-amenity locations will also have low returns to education. It is easy to think of location-specific amenities, like an ocean view, that are almost certainly luxury goods; wealthy people are willing to sacrifice a higher fraction of their wealth to purchase these amenities than are poor people. Our theory suggests, then, that places with high levels of such amenities—places that in turn have high housing prices—will have relatively low returns to education.

To build intuition for this finding we refer to Figure 2, which depicts the theoretical relationship between an individual's earnings and utility in each of two cities, a and b, with $A_a > A_b$. An individual with the high education level 1 has utility level u^1. We illustrate an example in which this utility is achieved by locating in either city and receiving w^1 in either city. (For example, this individual might be a computer programmer, whose skills are sold in a national labor market. Because in this example her productivity is the same in each city, so too are her wages; the assumption that w^1 is the same in both cities, however, is just for expositional simplicity.) A person with education level 1 enjoys the relatively high amenity level if she locates in the high-amenity city a, but "pays" for the amenity by facing high housing prices in that same city. An individual with education level 0 has utility level u^0 in either city. (For example, this individual might be the janitor; wages w_a^0 and w_b^0 can differ.) Because the amenity is a luxury good relative to housing, this poorly-educated individual

has a lower "willingness to pay" for the amenity than does the well-educated individual. He is indifferent between living in the high- or low-amenity city only if his wage is *higher* in the high-amenity city. In equilibrium, therefore, the education-wage gradient must be flatter in the high-amenity city.[10]

As Haurin (1980) and Roback (1982) note, the degree to which the amenity is capitalized into wages or housing prices may differ by local demand conditions. Obviously, at a fixed level of amenities, an increase in housing prices requires that wages too increase in order to keep the worker at a fixed level of utility. It is straightforward to show that, holding constant the amenity level and the utility levels u^0 and u^1, an increase in housing prices reduces the returns to schooling, or $\frac{\partial R}{\partial p} = p \, R \, (s_H^1 - s_H^1) < 0$. Thus, the greater the capitalization of the amenity into housing prices, the lower the financial returns to schooling.

III. An Empirical Examination of City-Specific Returns to Education

We turn next to an empirical exploration of cross-city heterogeneity in the education-wage gradient, examining specifically the return to a college education (relative to high school education). Our focus on the return to college stems in part from the fact that roughly 90 percent of young people now graduate high school, so that most of the meaningful variation in schooling is at the post-high school level. We take a nonparametric approach to estimation, focusing on those with college and high school levels of education provides us with relatively large samples for carrying out this exercise.

---

[10]The preferences illustrated in Figure 2 are quasi-homothetic. For the case in Figure 2 it is easy to show that our elasticities condition (4) boils down to this: The return to education is lower in the city where the "minimum consumption bundle" —the bundle consumed at the lowest defined utility level—is more expensive.

Our empirical explorations exploit the 1990 public use micro-samples (IPUMS) of the U.S. Census, and the 1990 Census complete long form data.[11] We are interested to see if variation in city-specific returns persist over time, so we also present estimates using 1980 IPUMS and 2000 IPUMS data (see Ruggles, *et al.*, 2004).

We consider only respondents with a bachelors degree or high school degree as the highest level of reported schooling.[12] We limit our analysis to men whose main job is a wage and salary position and who have no imputed values for variables used in this analysis. Due to concerns that arise with selection into the labor force, we restrict our sample to prime-age men, aged 25 to 55. We also restrict our sample to non-Hispanic white men, which allows us to abstract from any cross-city variation owing to race and ethnicity in labor market outcomes.[13]

We use a simple matching estimator to calculate, for each city j, the rate of return to college. We assume, as in the traditional Mincer set-up, that productivity is a function of education and experience x.[14] For an individual with experience $x = X$ in city j we would like to estimate the causal effect of college education ($BA = 1$):

$$(5) \qquad \Delta(X, j) = E(y_1 | x = X, BA = 1, j) - E(y_0 | x = X, BA = 1, j),$$

where y_1 is the logarithm of the worker's wage if the individual is college educated, and y_0 is the logarithm of the worker's wage if the individual stops his education at high school. Of

---

[11]These latter data provide extremely large samples, representing an approximately a one-in-six sample of the US population. Use of the confidential version of the data is helpful for our purposes because recorded earnings are top coded at $1,000,000$, rather than the $150,000$ mark used in the public-use version of Census data.

[12]Unfortunately, the data do not allow us to distinguish between the high school degree and GED.

[13]An additional consideration is that Hispanics and non-whites have much higher error rates than non-Hispanic whites in Census responses to education. See Black, Sanders, and Taylor (2003).

[14]Given concerns raised by Heckman, Lochner, and Todd (2003), though, we do not adopt parametric assumptions often used in estimating Mincer wage regressions, e.g., entering experience as a quadratic in the wage equation.

course we cannot directly observe the second term in equation (5); we never observe what a person with a bachelor's degree would have earned if he had only a high school education.

If we are willing to assume away selection problems, though (including the issue of ability bias that has received close attention in the literature), we have

$$E(y_0|x = X, BA = 1, j) = E(y_0|x = X, BA = 0, j).$$

In implementing our estimation strategy we follow standard practice of using "potential experience", age minus schooling minus six.[15] Then the mean return in a particular city j, say $\Delta(j)$, is

$$\Delta(j) = \int \Delta(x|j)dF(x|j),$$

where $dF(x|j)$ is the distribution of x in the city.

In principle, $\Delta(j)$ might vary across cities owing simply to differences in the age distributions in these cities. Those differences would be of little interest to us, so we "standardize" our estimates using the national cumulative distribution function of x, i.e., calculate

$$\Delta_n(j) = \int \Delta(x|j)dF_n(x),$$

where $F_n(x)$ is derived from the national data.

A. MSA-Specific Returns to College Education

Table 1 presents our initial evidence about the mean return to college education in 21 large U.S. urban locations. Using PUMS data, we include in our analysis all metropolitan

---

[15]We use potential experience rather than age because it is the variable implied by conventional human capital theory. There may be considerable measurement error in this variable, but our data do not allow us to improve on the measure. Because men with a high school degree only have more potential work experience than college-educated men, matching on potential experience implies that we typically match men with a bachelor's degree to men with a high school degree who are four years younger. Thus, we match men with a bachelor's degree aged 29 to 55 to men with a high school degree aged 25 to 51.

statistical areas (MSAs) with a sample of at least 1500 men with a bachelor's degree.[16] In

the table, MSAs are ordered from low to high estimated returns in 1990. There is substantial

heterogeneity in the return to a college education: In 1990, the year for which we conduct

our more detailed analysis, the log wage return varies from 0.33 in Seattle to over 0.54 in

Houston. Thus, the returns are 64 percent higher in Houston than in Seattle. Our estimates

are quite precise, with standard errors typically in the 0.01 to 0.02 range; this variation is not

likely due to sampling variation. Using test procedures described in Horrace and Schmidt

(2000), we conclude that in 1990 the lowest returns to college are in San Francisco or Seattle,

and the highest returns are in the subset: Houston, Pittsburgh, Dallas, Phoenix, Atlanta

and Tampa (based on a one-sided test, with a 0.95 critical value).

In two other columns of Table 1 we also list measured returns to college in each of these

MSAs in 1980 and 2000. There is a fair amount of persistence in the variation in the returns

to college. The correlation is 0.55 for returns 1980 to 1990, 0.67 for returns 1990 to 2000, and

0.55 for returns 1980 to 2000. Two other features of the data are readily apparent. First, over

the past two decades the return to college has generally increased. This fact about higher

education is well known and has been widely studied. Second, although there is clearly

some persistence in the returns to education, there is also a fair amount of idiosyncratic

heterogeneity. For example, San Francisco has a relatively low return in both 1980 and

1990, but a quite high return in 2000 (at the peak of the "dot com" boom).

The estimates we report in Table 1 are based on means. One potential source of bias

stems from the fact that income data in the PUMS are top-coded at \$150,000. We can

base our estimates on median regression, though, and when we do so results are unchanged;

---

[16]For cities that were a part of a large consolidated metropolitan area, we retain only the principle city
in the CMSA.

the correlation between the means-based estimates and the median-based estimates is 0.97. Similarly, it makes little difference whether we use the wage as our object of study "weekly earnings" or "annual earnings."[17] For interest sake, we also examined how our estimates would differ had we used the usual OLS approach to estimating returns to education. In fact, OLS estimates are nearly identical to our nonparametric estimates if we allow experience to be entered in a flexible way (with a dummy for each level of the 31 experience levels).

One issue concerns the extent to which cross-MSA variation in the returns to education is due to differences across MSA in industry and occupation. Fortunately, the Census provides reasonably detailed industry and occupation description (over 240 industries and over 480 occupations). The Appendix describes a simple semi-parametric approach to standardizing across MSAs on industry and occupation.

We find that standardizing for industry and occupation makes little change to the results in Table 1. More importantly, we conduct these detailed analyses using the complete long-form data of the 1990 Census, which allows us to substantially increase the number of MSAs we can study. In particular, we calculate the returns to college in 286 metropolitan statistical areas. Table 2 provides our key findings. Column (1) shows that when we include smaller MSAs in our analysis, there are very large differences in the returns, ranging now from 0.17 to 0.70. Adjusting for the age structure makes virtually no difference in our estimates (column 2). Conditioning on the industry (column 3) or occupation (column 4) distribution decreases the variation in our estimated returns, but range in cross-MSA estimates is still substantial: from 0.26 to 0.64. The variation in the measured returns is not being driven by

---

[17]The correlation between estimates based on wage and annual earnings is 0.96, and the correlation between estimates based on wage and weekly earnings is 0.97.

a few outliers. Even after conditioning on occupation, for example, the mean return is less than 0.35 in 10 percent of MSAs and it exceeds 0.51 in 10 percent of MSAs.

B. The Relationship Between the Return to Education and Local Housing Prices

Our empirical investigation indicates that there are large persistent differences in cross-MSA returns to education, and indicates that these differences are not due to differences in the industry or occupation mix across these MSAs. This is just as theory would lead us to expect (given non-homothetic preferences). The theory outlined above also provides reasons to expect that education-wage gradients will be flatter in high-amenity cities. We do not have a measure of local amenities, either in consumption or production.[18] But if desired amenities are capitalized in the housing price, our reasoning leads us to anticipate that there will be an inverse relationship between the observed return to education and the price of housing.

Returning to Table 1 we notice that indeed the MSAs with particularly low returns, e.g., Seattle, San Francisco, and New York, are among the most expensive urban locations in the country. The highest-return locations, e.g., Dallas, Pittsburgh, and Houston, are MSAs with relatively low housing prices. More generally, if we take a measure of quality-adjusted housing prices for these MSAs in 1990 (from Chen and Rosenthal, 2005), we find a Spearman rank-order correlation with the return to education of -0.54 (p-value < 0.01). If we estimate a regression with the return to college as a dependent variable and the price index (normalized between 0 and 1) as an independent variable, we get a slope estimate of -0.121 with a standard error of 0.048.

---

[18]Indeed the problem of measuring location-specific amenities is a difficult one. See, for example, Gyourko, Kahn and Tracy (1999) for a discussion of consumption amenities, and Glaser and Mare (2001) for some evidence about the presence of production amenities in expensive cities.

We turn next to a more systematic evaluation of this regularity. In this analysis we consider a linear model that specifies an MSA's "return to college" as a function of the MSA-level local housing price index and also the ratio of college-educated to high-school educated individuals in the MSA:

$$(6) \qquad\qquad R_j = \gamma_0 + \gamma_1 P_j + \gamma_2 (BA/HS)_j + \epsilon_j.$$

Our theory provides the rationale for including the local housing price in the regression; in equilibrium we expect the return to education to be lower in high-priced cities. Berry and Glaeser (2005) provide solid economic reasoning for conditioning also on the MSA's education mix. In particular, following the well-known work of Rauch (1993), Berry and Gleaser show that there is a strong positive correlation between the general educational level in a metropolitan area and average local wage. They then present theoretical arguments and supporting evidence indicating that these urban agglomeration effects may be growing stronger over the 1980–2000 period. Importantly, for our purposes, the agglomeration benefits created for people who live in metropolitan areas with large concentrations of skilled people are larger for skilled workers than for unskilled workers.[19] If Berry and Glaeser's arguments are correct, we would expect that in 1990 the observed return to college would be higher in metropolitan areas with higher concentrations of college-educated individuals (i.e., we would expect γ_2 to be positive). In any event, we want to control for this possibility.

Table 3 provides estimates using complete long-form 1990 Census data. For each metropolitan area we calculate the return to a bachelor's degree (relative to a high school degree),

---

[19]Over the last two decades metropolitan areas that initially had relatively high concentrations of college graduates generally saw those concentrations increase. To explain this phenomenon Berry and Glaeser posit a model of an agglomeration economy in which entrepreneurial innovations are disproportionately made by skilled people, and in turn these innovations result in an increase in local labor demand that disproportionately benefits other skilled individuals. (Lower-skilled individuals might also benefit from being around skilled people, but these benefits are smaller.)

standardizing on experience, as described above. We use as our housing measure a quality-adjusted housing price index developed by Chen and Rosenthal (2005), which we normalize between 0 and 1.[20] We report Huber-White (robust) standard errors unless we note otherwise.

In the first column we presents our baseline results. The estimated coefficient on the housing price index, $\hat{\gamma}_1$, is indeed negative; in metropolitan areas with higher prices, the return to education is lower. Also, the estimated coefficient $\hat{\gamma}_2$ is positive. In our regression analysis observations are unweighted by the metropolitan area size. The rationale for this approach is that our theory makes predictions about how local economies price education, and each city represents a distinct realization. One may worry, however, that our results are being driven by small MSAs. Furthermore, one could argue that in large metropolitan areas the returns to education are being estimated more precisely, and that these observations are therefore more informative. In any event, in the second column of Table 3 we present results for a specification that weights by metropolitan size. Our key result—that the return to education is lower in expensive metropolitan locations—is if anything strengthened.

Analysis presented in the remaining columns provides robustness checks, investigating whether our findings are due to "outliers." In particular, we tried three additional approaches: we use median regression, calculating standard error based on 999 bootstrapped replications; we used Stata's robust regressions, which reduces the weight on "outliers" using Huber and bi-weights; and we deleted 22 influential observations using a critical value suggested by Belsley, Kuh, and Welsch (1980) ($2\sqrt{k/n}$ or approximately 0.178 with our two regressors and 253 observations). Results are quite similar with each of these procedures.

---

[20]We exclude from analysis the 33 MSAs for which the index is unavailable.

One might be concerned that the our education composition variable $(BA/HS)_j$ is itself endogenous, i.e., influenced by the same factors that drive the return to education. In turn, this might bias the estimated coefficient on housing prices, $\hat{\gamma_1}$. Although we have no entirely satisfactory way to deal with this issue, we did estimate our model using a substantially lagged value of city j's educational composition as an instrument for our contemporaneous measure of education composition.[21] In particular, we used as our instrument the 1940 ratio of individuals with 16 or more years of education to individuals with a high school degree in city j. This instrument is positively correlated with 1990 values of our $(BA/HA)_j$ variable (and the F statistic for the first stage is over 30). Importantly, the 2SLS estimate of γ_1 is virtually the same as the OLS estimate; the point estimate is -0.100 with a standard error of 0.040.

Recent innovative work by Lee (2007) provides theoretical reasoning, complementary to ours, for understanding differences in observed returns to education across cities. In Lee's model people value consumption variety, and in equilibrium land values are higher in large cities than in small cities because large cities provide more consumption variety. High-skill (thus high-income) people place a higher value on variety than low-skill people. In consequence, low-skill individuals will require a wage *premium* to live in large expensive cities, while high-skill individuals might well accept an urban wage *discount.* Lee proceeds to show that such a pattern exists among medical professionals in the U.S. For example, doctors in large cities are paid less than their peers in small cities, while the converse is true of nurses.

Notice that if housing in larger cities is relatively expensive, as it must be if city size *per se* is the key to the valued amenity (consumption variety), then the prediction of Lee's

---

[21]See, e.g., Berry and Glaeser (2005).

model is the same as ours: the observed return to education will be lower in expensive cities. Equivalently, in Lee's set-up, the return to education will be lower in large cities. As a rough way of providing empirical evidence about the hypothesized key role of city size for our context—evaluating differences in the return to college across metropolitan areas—we substitute a measure of MSA size for the housing price index in our regression model (6). When we do so, using either *MSA Population* or *Log of MSA Population* as an explanatory variable, we find that the coefficient on the metropolitan size variable is near zero (when we use *MSA Population* as the explanatory variable) or unexpectedly positive and marginally significant.[22] Housing prices appear to be more helpful than metropolitan size in explaining cross-MSA variation in the return to college. Having said this, though, we view all of the results we present here as a mere starting point for further analysis. There is persuasive evidence of substantial systematic differences in the return to education across metropolitan areas in the U.S., but these differences are part of a complex set of interactions in urban labor and housing markets that merit further exploration.[23]

IV. Concluding Remarks

We have described a simple model in which prices vary across location. In our theory, the (proportional) returns to education are same across locations if and only if preferences are homothetic. We show that even for this case, the proper empirical approach to estimating Mincer earnings functions is to include a fixed effect for each location.

---

[22]When we use *Log of MSA Population* as an explanatory variable, the point estimate of the associated coefficient is 0.010 with a standard error of 0.055. It seems that MSA size doesn't work as expected here because some smaller MSAs in our sample are quite pricey (e.g., Napa, California) while some larger MSAs are rather inexpensive (e.g., Cleveland). It is worth noting that our analysis excludes consideration of smaller urban areas, with population under 100,000, and rural areas.

[23]Lee's (2007) interesting empirical analysis of local wage variation among medical professionals is a good example of such work.

Our theory shows, more importantly, that in the more general case of non-homotheticity in preferences, returns to education will vary by location. Moreover, empirical evidence suggests that the monetary reward to a college education does indeed vary widely across U.S. cities. This variation is persistent over decades, and is clearly systematic. In particular, it appears that the return to education is relatively low in expensive high-amenity cities.

On the empirical front, our work has a number of implications that bear further investigation. We would argue that considerable care must be exercised in estimating earnings regressions using data that span locations. We provide two examples in which attention to the issues raised above may affect inferences in empirical work.

Our first example concerns the well-documented increase in the return to college seen in the U.S. during the 1980s. The underlying causes of this increase are widely studied, and there is considerable concern on the policy front about the effects of this continuing trend for societal inequality.[24]

Now even if one is willing to assume homotheticity in preferences—so that the return to education is a single parameter—OLS estimates of this parameter may be systematically biased if location is ignored in the estimation procedure. In particular, because there are differences in wage levels over locations, one needs to include labor-market fixed effects. Consider the return to college in 1980 and in 1990. In the Census PUMS for non-Hispanic white men, regressing log wage on a vector of dummies for potential experience and education—restricting attention to individuals with exactly 12 and 16 years of schooling—we estimate an annual return of 0.088 in 1980 and 0.118 in 1990, giving an approximately 3.0 percentage

---

[24]Work on this topic includes include important contributions of Murphy and Welch (1992), Katz and Murphy (1992), Bound and Johnson (1992), and Juhn, Murphy, and Pierce (1993). Policy issues are discussed in, e.g., the Council of Economic Advisers (1997).

point increase. Conducting that same exercise but in a regression that includes also location indicators for city of residence (or state of residence for those not residing in large and medium-size cities), we estimate returns of 0.082 and 0.104 respectively for 1980 and 1990, giving a 2.2 percentage point increase. Omitting location dummies causes us to overestimate the change in the mean return to college by 36 percent.[25]

Also note that some of the increase in nominal earnings inequality might stem from the reallocation of workers between regions with differing education-earnings gradients. Consider, for instance, Figure 2. In this example, if the population grows more rapidly in the low-priced city than the high-priced city, measured inequality increases, even though there are no welfare changes for either the poorly-educated (u^0) or well-educated (u^1).

One could undertake a location-based decomposition of changes in the returns to education in the U.S. Here we simply note that despite the sharp increase in the returns to college from 1980 to 1990, in the 21 large cities listed in Table 1 there was virtually no increase in the average return to college over this decade.[26] Apparently the increase can be attributed primarily to changes in returns in other cities and rural areas, and in migration between areas.

Our second example concerns the many papers that rely on natural experiments that exogenously influence individuals' schooling decisions (e.g., variation in institutional features in the provision of education) to identify the causal effect of schooling on earnings. Our work argues that the monetary return to education will vary across locations; the "causal return to education," measured as the relationship between schooling and wages, is not a single

---

[25]Our regressions have over 300 location indicators. Standard errors in the estimated returns were very small (0.0005 or less) owing to sample sizes in excess of 500,000. Of course, given our discussion above, we are not particularly enamored of these estimates. Including location dummies does not deal with the more fundamental issue that returns to education vary across locations. A more palatable alternative might be to evaluate changes in the *distribution* of returns to education.

[26]In our 21 large cities, the average total return to four years of college was 0.43 in 1980 and 0.44 in 1990.

parameter. With this in mind, one might wonder how instrumental variable (IV) approaches will be affected if they fail to account for localized differences in returns to education.

A typical OLS estimate of "the return to education" will give an average of education-earnings gradients for many locations. An IV estimator re-weights observations—placing increased weight on those observations where the exogenous variation influences schooling outcomes. To take a specific instance, consider work by Angrist and Krueger (1991). In that paper, the authors noticed that because of the specifics of compulsory school attendance laws, children born in the fourth quarter of the calendar year receive more education than children born earlier in the year. In principle, their IV estimator identifies the causal effect of additional schooling for increments of education around the compulsory schooling level. They find that in various specifications IV estimates are similar to or higher than corresponding OLS estimates, an outcome might occur if OLS estimates are downward biased. This outcome would also occur, though, if compulsory schooling laws have a disproportionate impact on completed education in locations that have high location-specific returns, for example if individuals affected by the laws tend to live in low-amenity locations.[27]

Our work surely raises concerns for other empirical work in labor economics.[28] As for future theoretical development, there are a number of potentially interesting paths that might be worth following. In our model the equilibrium "utility return" to education is the same in each location, while the "monetary return" is lower in high-amenity cities. Put another

---

[27]Similar concerns might arise with many IV approaches. Institutions that generate variation in schooling outcomes are invariably location-specific. One might be particularly concerned about research strategies explicitly based on location, such as studies that exploit variation in individuals' proximity to a college or university.

[28]For instance we might also worry about the literature on earnings differentials between races and ethnic groups. The geographic patterns of residence of minority groups are very different than for non-Hispanic whites. For example, 53 percent of blacks but only 33 percent of whites now live in the South, while 20 percent of whites but only 9 percent of blacks live in the West. Given substantial regional differences in levels and slopes of schooling-earnings gradients, racial and ethnic differences in earnings and returns to education will appear, even if in each location all workers with a given schooling level have the same earnings.

way, as is readily apparent in Figure 2 (when one flips the axes), the marginal utility of money is lower for individuals in low-amenity cities than for individuals in the high-amenity cities. This feature of a location-based economy might well have interesting implications for a variety of behaviors, including investment in human capital, migration, fertility decisions, labor supply, and interpretation of evidence concerning agglomeration economies.

Appendix. Description of the Semi-Parametric Estimation

We are interested in estimating the return to college for an individual with characteristics $x = X$ in city j:

$$\Delta(X, j) = E(y_1|x = X, BA = 1, j) - E(y_0|x = X, BA = 1, j),$$

where y_1 is the logarithm of the worker's wage if the individual receives a bachelor's degree, y_0 is the logarithm of the worker's wage if the individual stops his education at high school, and BA is an indicator variable equal to 1 if the respondent has a college education. As noted in the text, we of course observe at most either y_1 or y_0, never both, so we adopt as our identification condition,

$$E(y_0|x = X, BA = 1, j) = E(y_0|x = X, BA = 0, j).$$

We assume that the data generating process for the wage y_1 is

$$(7) \qquad\qquad\qquad\qquad y_1 = g_1(x, j) + \varepsilon_1,$$

where $g_1(x, j)$ is an unknown function of the covariates and ε_1 is a mean zero error term that is independent of x and the vector of cities. Similarly, we assume that the data generating process for the wage y_0 is

$$(8) \qquad\qquad\qquad\qquad y_0 = g_0(x, j) + \varepsilon_0,$$

where again $g_0(x, j)$ is an unknown function of the covariates and ε_0 is a mean zero error term that is independent of X and the vector of cities.

In our initial results, reported in Table 1 and the first column of Table 2, the only covariate x that we use is the worker's potential experience. In this case, we use a completely non-parametric specification of g_0 and g_1; we directly match workers on their exact potential experience.

The Census provides detailed occupation and industry codes. We could, of course, match workers on their industry and/or occupation as well. Our goal, however, is *not* to compare, say, accountants in Bakersfield with a four-year college degree and 14 years of potential experience to the accountants of Bakersfield with a high school degree and 14 years of potential experience. (Indeed, we believe that a portion of the returns to a college education is captured by entry into higher paying occupation.) Rather, we wish to "control for" the occupation distribution across cities. Moreover, given the extremely large number of cities and occupations (or industries), we would undoubtedly suffer from severe support problems. For instance, we doubt that there are any employed coal miners residing in New York City or Los Angeles.

We pursue the semiparametric approach, but we wish to use relatively flexible functional forms so that we do not put too much structure on the data.[29] We therefore re-specify our equations (7) and (8) as

$$y_1 = \tilde{g}_1(x, j, z) + \varepsilon_1 = g_1(x, j) + z\gamma_1 + \varepsilon_1,$$

$$y_0 = \tilde{g}_0(x, j, z) + \varepsilon_0 = g_0(x, j) + z\gamma_0 + \varepsilon_0,$$

where functions $g_0(x, j)$ and $g_1(x, j)$ are left non-parametric (i.e., we match on potential experience), and z is a vector of occupation (or industry) indicators.

---

[29]See Horowitz (1998) for an excellent discussion of the relative merits of semiparametric estimation.

Then in estimating the return to college we replace y_1 and y_0 with $\tilde{y}_1$ and $\tilde{y}_1$, where

$$(9) \qquad \tilde{y}_1 \;=\; g_1(x, j) - z\gamma_1 + \overline{z_1}\gamma_1 + \varepsilon_1,$$

$$(10) \qquad \tilde{y}_0 \;=\; g_0(x, j) - z\gamma_0 + \overline{z_0}\gamma_0 + \varepsilon_0,$$

and $\overline{z_j}$ is the mean of the industry or occupation controls for the jth group. While much less restrictive than most wage equations, our specification of equations (9) and (10) requires an occupation (or industry) to shift the wage profile by an equal amount for each potential experience level and city. Conceptually we are asking what would be the observed return to education in each city had the same distribution of workers' expected experience and the same industry or occupation mix.

References

Acemoglu, Daron. 1996. A microfoundation for social increasing returns in human capital accumulation. *Quarterly Journal of Economics* 111, no. 3:779-804.

Angrist, Joshua and Alan Krueger. 1991. Does compulsory school attendance affect schooling and earnings? *Quarterly Journal of Economics* 106, no. 4:979-1014.

Becker, Gary S. 1964. *Human capital: a theoretical and empirical analysis with special reference to education.* National Bureau of Economic Research, New York.

——. 1967. *Human capital and the personal distribution of income: an analytical approach.* Ann Arbor, Michigan: University of Michigan Press.

Beeson, Patricia E. 1991. Amenities and regional differences in returns to worker characteristics. *Journal of Urban Economics* 30, no. 2:224-241.

Belsley, David, Edwin Kuh, and R. Welsch. 1980. *Regression diagnostics: identifying influential data and sources of collinearity.* New York: John Wiley and Sons.

Berry, Christopher R., and Edward L. Glaeser. 2005. The divergence of human capital levels across cities. *Papers in Regional Science* 84, no. 3:407-444.

Black, Dan, Seth Sanders, and Lowell Taylor. 2003. Measurement of higher education in the census and Current Population Survey. *Journal of the American Statistical Association* 98, no. 463:545-554.

Bound, John and George Johnson. 1992. Changes in the structure of wages in the 1980's: an evaluation of alternative explanations. *American Economic Review* 82, no.3:371-92.

Card, David. 1999. The causal effect of education on earnings. In *Handbook of labor economics* 3A., eds. Orely C. Ashenfelter and David Card, 1801-1863. Amsterdam: Elsevier Science.

——. 2001. Estimating the returns to schooling: progress on some persistent econometric problems. *Econometrica* 69, no.5:1127-60.

Card, David and Alan Krueger. 1992. School quality and black-white relative earnings: a direct assessment. *Quarterly Journal of Economics* 107, no. 1:151-200.

Chen, Yong and Stuart Rosenthal. Forthcoming. Local amenities and life cycle migration: do people move for jobs or fun? *Journal of Urban Economics.*

Chiswick, B. 1974. *Income inequality: regional analyses within a human capital framework.* New York: Columbia University Press.

Council of Economic Advisers. 1997. *Economic report of the President*, Washington D.C.

Gabriel, Stuart and Stuart Rosenthal. 2004. Quality of the business environment versus quality of life: do firms and households like the same cities? *Review of Economics and Statistics* 86 no. 1:438-444.

Glaeser, Edward L. and David C. Mare. 2001. Cities and skills. *Journal of Labor Economics* 19, no. 2:316-342.

Gyourko, Joseph, Matthew Kahn, and Joseph Tracy. 1999. Quality of life and environmental comparisons. In *Handbook of regional and urban economics*, ed. Paul Cheshire and Edwin S. Mills, 1413-1454. Elsevier.

Harmon, Oskar. 1988. The income elasticity of demand for single-family owner-occupied housing: an empirical reconciliation. *Journal of Urban Economics* 24, no. 2:173-185.

Haurin, Donald R. 1980. The Regional distribution of population, migration, and climate. *Quarterly Journal of Economics* 95, no. 2:293-308.

Hausman, J. A., Newey, W. K., and J. L. Powell. 1995. Nonlinear errors in variables: estimation of some Engel curves. *Journal Of Econometrics* 65, no. 1:205-233.

Heckman, James, Lochner, Lance, and Petra Todd. 2003. Fifty Years of Mincer Earnings Regressions. Working Paper no. 9732, National Bureau of Economic Research, New York.

Horowitz, Joel L. 1998. *Semiparametric methods in econometrics*. Berlin: Springer-Verlag.

Horrace, William C. and Peter Schmidt. 2000. Multiple comparisons with the best, with economic applications. *Journal of Applied Econometrics* 15, no. 1:1-26.

Juhn, Chinhui, Kevin M. Murphy, and Brooks Pierce. 1993. Wage inequality and the rise in returns to skill. *Journal of Political Economy* 101, no. 3:410-442.

Katz, Lawrence F. and Kevin M. Murphy. 1992. Changes in relative wages, 1963-1987: supply and demand factors. *Quarterly Journal of Economics* 107, no. 1:35-78.

Lee, Sanghoon C. 2007. Ability sorting and consumer city. draft, University of British Columbia.

Mincer, Jacob, 1974. *Schooling, experience, and earning*. New York: National Bureau of Economic Research.

Murphy, Kevin M. and Finis Welch, 1992. The structure of wages. *Quarterly Journal of Economics* 107, no. 1:285-326.

Olsen, Edgar. 1987. The demand and supply of housing service: a critical survey of the empirical literature. In *Handbook of regional and urban economics*, ed. Edwin S. Mills. Elsevier Science. 989-1022.

Rauch, James E. 1993. Productivity gains from geographic concentration of human capital: evidence from the cities. *Journal of Urban Economics* 34, no. 3:380-400.

Roback, Jennifer. 1982. Wages, rents, and the quality of life. *The Journal of Political Economy* 90, no. 6:1257-1278.

Rosen, Harvey S. 1985. Housing subsidies: effects on housing decisions, efficiency, and equity. In *Handbook of public economics*, ed. A. J. Auerback and M. Feldstein, 375-420. North-Holland: Elsevier.

Ruggles, Steven, Matthew Sobek, Trent Alexander, Catherine Fitch, Ronald Goeken, Patricia Hall, Miriam King, and Chad Ronnander. 2004. Integrated Public Use Microdata Series: Version 3.0 [Machine-readable database]. Minneapolis, MN: Minnesota Population Center. http://usa.ipums.org/usa/

Stata. 2007. Reference, Q-Z Version 10. College Station, Texas: Stata Press.

Willis, Robert J. 1986. Wage determinants: a survey and reinterpretation of human capital earnings functions. In *Handbook of labor economics* 1, ed. Orley C. Ashenfelter and Richard Layard, 525-602. Amsterdam: Elsevier Science.

FIGURE 1. Earnings-Education Profiles when Preferences are Homothetic
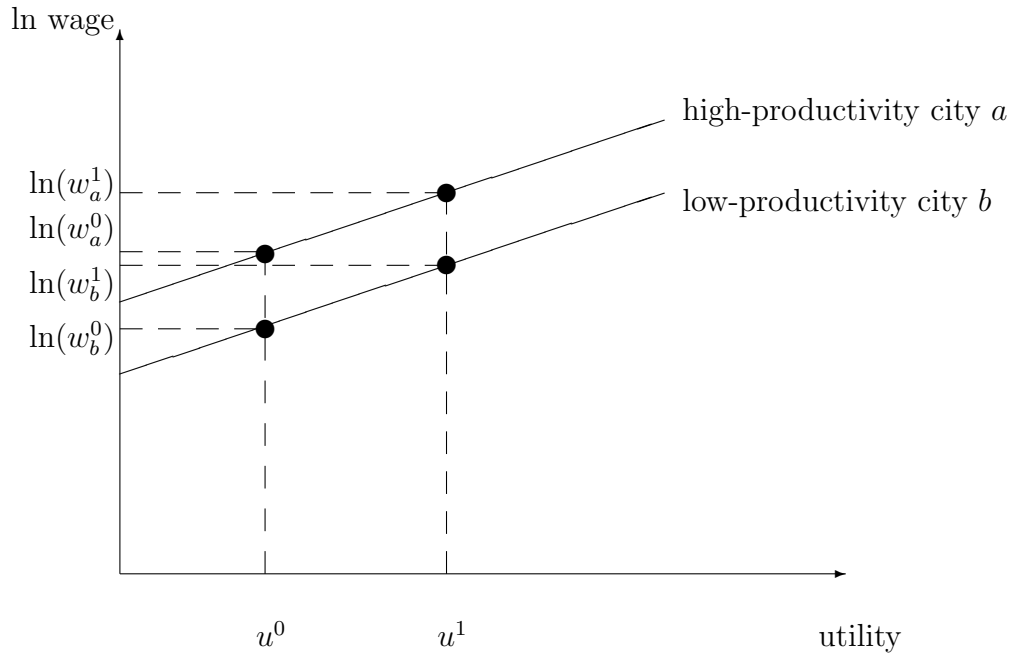
FIGURE 2. Earnings-Education Gradients when Cities Have Different Consumption Amenities (and Preferences are Not Homothetic)
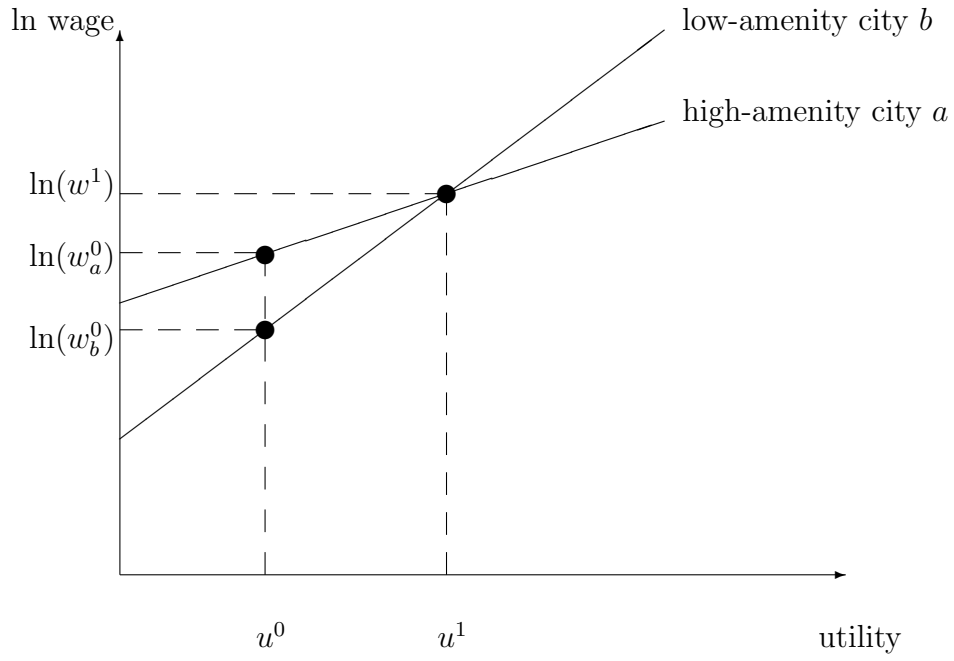
TABLE 1. Local Variation in the Returns to a Bachelor's Degree, PUMS

City	2000	1990	1980
Seattle	0.405 (0.0182)	0.331 (0.0147)	0.335 (0.0171)
San Francisco	0.573 (0.0360)	0.378 (0.0296)	0.341 (0.0146)
Minneapolis	0.457 (0.0154)	0.386 (0.0146)	0.395 (0.0131)
New York	0.497 (0.0186)	0.388 (0.0142)	0.450 (0.0098)
Chicago	0.434 (0.0116)	0.389 (0.0107)	0.390 (0.0087)
St. Louis	0.436 (0.0176)	0.404 (0.0149)	0.438 (0.0155)
Detroit	0.472 (0.0135)	0.421 (0.0120)	0.396 (0.0117)
Los Angeles	0.497 (0.0184)	0.429 (0.0118)	0.438 (0.0110)
Washington	0.551 (0.0136)	0.430 (0.0119)	0.459 (0.0120)
Boston	0.536 (0.0140)	0.437 (0.0124)	0.420 (0.0147)
Philadelphia	0.482 (0.0121)	0.438 (0.0100)	0.444 (0.0099)
Baltimore	0.473 (0.0162)	0.439 (0.0132)	0.453 (0.0160)
Cleveland	0.474 (0.0168)	0.442 (0.0151)	0.389 (0.0153)
Denver	0.471 (0.0194)	0.447 (0.0191)	0.408 (0.0161)
San Diego	0.577 (0.0210)	0.458 (0.0169)	0.489 (0.0195)
Tampa	0.547 (0.0180)	0.491 (0.0176)	0.436 (0.0214)
Atlanta	0.556 (0.0152)	0.497 (0.0140)	0.483 (0.0158)
Phoenix	0.506 (0.0163)	0.518 (0.0176)	0.396 (0.0223)
Dallas	0.617 (0.0181)	0.524 (0.0155)	0.487 (0.0128)
Pittsburgh	0.529 (0.0177)	0.530 (0.0159)	0.399 (0.0147)
Houston	0.609 (0.0179)	0.542 (0.0142)	0.472 (0.0124)

Note. Authors' calculations, Five-Percent PUMS of the 2000, 1990, and 1980 Census. Workers are non-Hispanic white males who have either a high school degree or bachelor's degree with between seven and 33 years potential experience (aged 25 to 55 years). Workers are matched within cities to workers with exactly the same number of years of potential experience. For each city, the distribution of potential experience is standardized to the national average of bachelor's degree holders. Bootstrapped standard errors using 499 replications are reported in parentheses.

TABLE 2. Estimated Returns to a Bachelor's Degree, 1990 Census Long Form

	Mean Return	Standardizing on Experience	Standardizing on Experience and Industry	Standardizing on Experience and Occupation
Mean	0.413	0.414	0.422	0.425
Std. Deviation	0.077	0.076	0.063	0.062
Lowest MSA	0.168	0.197	0.261	0.263
10th Percentile	0.318	0.321	0.341	0.351
25th Percentile	0.362	0.363	0.384	0.385
Median MSA	0.413	0.414	0.420	0.422
75th Percentile	0.461	0.457	0.463	0.466
90th Percentile	0.505	0.507	0.498	0.506
Highest MSA	0.701	0.703	0.642	0.636

Note. Results are from semiparametric estimation. The data are from the 1990 Census complete long form. Data are weighted to account for sample stratification. Our sample consists of 1,032,629 non-Hispanic white men aged 25 to 55 years with high-school or college degrees, reporting positive earnings for the year, with non-imputed data on earnings, weeks worked, and usual hours of work per week. The unit of observation is the MSA. The 286 MSAs used had population greater than 100,000 and had all the age groups present in the sample.

TABLE 3. Estimated Regression Coefficients: Dependent Variable is the Return to College

Independent Variable	Unweighted OLS	Weighted OLS	Median Regression	Robust Regression	Trimmed, Based on DFIT Stat.
Housing Price Index	−0.083	-0.108	−0.088	−0.098	−0.106
	(0.0248)	(0.0262)	(0.0263)	(0.0222)	(0.0173)
Ratio of College Graduates to HS Graduates	0.135	0.139	0.146	0.139	0.155
	(0.0181)	(0.0252)	(0.0249)	(0.0187)	(0.0141)
R^2 or pseudo R^2	0.163	0.197	0.099	—	0.272
N (number of metro areas)	253	253	253	253	231

Note. Huber-White robust standard errors are in parentheses, except for the median regression, which are bootstrapped. The data are from the 1990 Census complete long form and 5% PUMS of the 1990 Census. The unit of observation is the MSA. Housing price index is Chen and Rosenthal's (2005) quality-adjusted housing index normalized between 0 and 1.