Local Labor Markets and the Evolution of Inequality

Dan A. Black,[1,4,5] Natalia Kolesnikova,[2] and Lowell J. Taylor[3,4,5,6]

[1]Harris School of Public Policy, University of Chicago, Chicago, Illinois 60637; email: danblack@uchicago.edu

[2]Department of Economics, University of Mississippi, University, Mississippi 38677

[3]Heinz College, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213

[4]NORC, Chicago, Illinois 60637

[5]Institute for the Study of Labor (IZA), 53113 Bonn, Germany

[6]National Bureau of Economic Research, Cambridge, Massachusetts 02138

Annu. Rev. Econ. 2014. 6:605–28

First published online as a Review in Advance on April 10, 2014

The *Annual Review of Economics* is online at economics.annualreviews.org

This article's doi:
10.1146/annurev-economics-080213-040816

JEL codes: D31, D63, J31, R23

Keywords

wage regressions, cross-location variation in inequality, changes in wage inequality

Abstract

US labor markets have experienced rising inequality over the past 30 years—as evidenced by an increased gap in wages earned by high-skill workers (e.g., college graduates) and low-skill workers (e.g., high school graduates). Empirical evidence documenting this evolution of inequality comes from studies that assess wage-education gradients at the national level. But of course people work in local labor markets that differ in important ways. We provide a theoretical framework for evaluating inequality changes when individuals work in local labor markets, and we give an empirical reassessment of inequality changes in light of the insights that emerge from our framework.

1. INTRODUCTION

Over the past 30 years, US labor markets have been characterized by a striking rise in inequality, as evidenced by increasing gaps between wages at low-percentile and high-percentile benchmarks, or as indicated by rising returns to higher education. This phenomenon is among the most intensely studied topics in labor economics.

The prevailing wisdom assigns an outsized role for the "race between education and technology" as an explanation for increasing wage dispersion. The basic story is simple. Throughout the twentieth century, increasing educational attainment resulted in an upskilling of the US labor force. Over these same decades, however, there have been dramatic technological innovations in the production of goods and services, which have evolved in ways that generally favored highly skilled workers, especially those with a college education. From the early part of the twentieth century through approximately 1950, the increasing relative supply of highly skilled workers outpaced the increasing relative demand for these same workers, and in consequence, there was a drop in the wage gap between high-skill and low-skill workers. From 1950 through approximately 1980, the relative wage gap remained quite low. Thereafter, the trend reversed; since approximately 1980, skill-biased technological change has led to an increase in the demand for high-skill workers relative to low-skill workers, outpacing the corresponding shifts to relative supply, hence increasing wage dispersion in labor markets.[1]

Our article focuses on an underappreciated and potentially illuminating feature of the recent trend in the US wage structure. Specifically, the empirical documentation of trends in wage inequality comes almost exclusively from an examination of national data—thereby obscuring evolving differences in local labor markets in the United States. Our concern, in particular, is that many interactions between workers and firms occur in markets that have markedly different housing prices and nominal wages. To get a sense of the magnitude of the cross-city price variation under discussion, a real-earnings calculator available on the CNNMoney website indicates that a person presently earning \$100,000 in Austin, Texas, would require a salary of \$143,257 to live in Los Angeles, \$175,512 to live in San Francisco, and \$235,030 to live in Manhattan (**http://money.cnn.com/calculator/pf/cost-of-living/**). The calculations may be rough, but they likely give a good sense of the dangers of ignoring spatial variation in prices in local labor markets as disparate as Austin and Manhattan. In fact, the value of the Consumer Price Index (CPI) as of 2013 is approximately 233, relative to a value of 100 in 1983. It seems that cost-of-living differences between present-day Austin and Manhattan are on the same order as national cost-of-living differences from the early 1980s to the present day!

Complicating matters further is that the past three decades have been characterized by dramatic shocks to local markets, for example, the collapse in demand for labor in heavy manufacturing in the Rust Belt and the rise in labor demand in high-tech sectors in such locations as San Francisco and Boston, which were accompanied by the stagnation of housing prices in many Midwestern cities and sharp rise in housing prices in many East and West Coast localities. In short, there is large

[1]The notion of the "race between education and technology" appears in the pioneering work of Tinbergen (1974) and is set out in detail in Goldin & Katz's (2008) seminal monograph that adopts that phrase as a title. These authors trace the "race" in the US context, over the course of the twentieth century and the beginning of the twenty-first century. Their work includes a careful discussion of the role of institutions in shaping the wage structure. As for other literature, Acemoglu & Autor (2011) provide a definitive overview of the role of technological change for the evolution of inequality in labor markets and provide links to the vast literature on the topic. Card & DiNardo (2002) highlight difficulties in the identification of models of skill-biased technological change (i.e., problems that arise given that, in much of the literature, the term technological change is simply the name assigned to wage-change components that are not explained by the evolution of the relative supply of low- and high-skill workers). An important challenge is to understand why the evolution of wage inequality in labor markets differs across countries [see, e.g., Card et al. (2013), who study increasing wage inequality in German labor markets].

variation in the cost of living across locations and that variation has been shifting in important ways over the past three decades.

Our goal here is to provide a conceptual framework for thinking about the evolution of inequality in an environment such as the one described in the preceding paragraphs. The starting point is the canonical model from urban economics (Haurin 1980, Roback 1982), adapted to accommodate workers with differing skill levels. Specifically, as is standard in urban models, we suppose that locations are heterogeneous in productivity or consumption amenities, and individuals can costlessly move from one location to another. Individuals are alike in their tastes, but they differ in terms of skill levels. In such models, equilibrium nominal wages and prices differ across locations, but the models are constructed so that in equilibrium, individuals are indifferent about where to live. Importantly, therefore, real inequality must be the same in all locations. Our insight is that if we are willing to adopt the further assumption that consumer preferences are homothetic, then log nominal wage differences between, for example, college and high school graduates (i.e., inequality as measured by wage gaps) will also be the same across all cities.

In our model, if skill-biased technological change increases the relative demand for high-skill workers in some locations, our inequality measure must rise in all locations by the same amount—at least once markets reach the new equilibrium. This motivates an empirical approach to the evaluation of changing inequality that differs from standard practice. In particular, if our model is a reasonably accurate depiction of the real world, we expect that the within-location evolution of inequality will be similar in all locations, but the national trend need not mirror local trends. For instance, if migratory patterns entail a relocation of workers from low-wage, low-price cities to high-wage, high-price cities, nominal inequality might easily rise more rapidly than real inequality. The theory shows how, in principle, an empirical approach that relies solely on national averages might result in serious misunderstandings.

Our article proceeds in three additional sections. In Section 2, we set out the theory described in the preceding paragraphs. In Section 3, we reconsider empirical evidence about the evolution of inequality in the United States, 1980–2010, in light of the theoretical insights that we develop. Our work entails a detailed evaluation of inequality across 21 large cities—as measured by the wage of college graduates relative to similarly aged high school graduates.[2] We find very large variation across cities in the evolution of wages. For instance, if we focus on non-Hispanic white men, we find that, in Detroit, the log point change in real wage from 1980 through 2010, as adjusted using the national CPI, was -0.01 for college-educated workers and -0.31 for high school–educated workers. In contrast, the corresponding changes in New York were $+0.31$ for college-educated workers and $+0.04$ for those with a high school education. More generally, in line with our theory, we find that the increase in the log wage gap between college and high school graduates was reasonably similar across cities (e.g., 0.30 in Detroit and 0.27 in New York).[3]

In Section 4, we discuss implications of our theoretical approach for a variety of important questions. For instance, although our analysis focuses primarily on non-Hispanic white men, in Section 4 we ask how our insights pertain to the analysis of black-white wage differentials, and women in the labor market. We also ask how an approach that emphasizes the role of local labor markets can shed light on the causes of increasing wage inequality and the consequent impacts on welfare.

[2]More specifically, we use weekly wages of full-time workers. Thus, our empirical results can be interpreted as describing earnings inequality.

[3]Having said that, there certainly is evidence to suggest that our model is a serious oversimplification—definitely inadequate to the task of fully characterizing the evolution of cross-city wage inequality.

2. THEORY

Our model of location-specific wage inequality is based on the canonical model of local markets due to Haurin (1980) and Roback (1982). The setup allows for differences across locations in worker productivity and/or location-specific amenities. In turn, there is variation in local wages and prices. To begin our discussion, we divide all consumption goods into traded goods that have a common price across locations and local goods, which have prices that can differ across locations. Local goods include locally provided services, such as restaurant meals, and goods that must be consumed locally, such as housing. Traded goods are goods that are easily shipped to various locations. To simplify matters, we take the case with one traded good, which has a price of 1 everywhere, and one local good, housing.[4] We let p_j be the price of a unit of the local good in location $j = 1, \ldots n$.

Suppose people live in one of our n locations and have one of two levels of education. To set the stage for the empirical work that follows, we designate the two levels of education to be high school and college.[5] Wages vary across locations; individuals with a high school education earn wage w_j^{HS} in city j, and those with a college degree earn wage w_j^{C}. We assume that all workers supply one unit of labor and have the same preferences.

As is typical in these models, we also assume that the migration between locations is costless. Among the many insights of Haurin and Roback, one is extremely relevant for our problem: Competitive forces require utility levels to be equilibrated between all locations, so $u_j^{\mathrm{HS}} = u^{\mathrm{HS}}$ and $u_j^{\mathrm{C}} = u^{\mathrm{C}}$. Thus, in equilibrium, the expenditure function of workers with a high school diploma living in city j is $w_j^{\mathrm{HS}} = e_j^{\mathrm{HS}} = e(p_j, u^{\mathrm{HS}})$, and the expenditure function of workers with a college degree living in city j is $w_j^{\mathrm{C}} = e_j^{\mathrm{C}} = e(p_j, u^{\mathrm{C}})$.

To facilitate the study of inequality in such a model, assume that each location has workers of both education levels, who have utilities satisfying $u_j^{\mathrm{C}} > u_j^{\mathrm{HS}}$. Define the inequality index

$$I_j\left(p_j, u^{\mathrm{HS}}, u^{\mathrm{C}}\right) = \frac{w_j^{\mathrm{C}}}{w_j^{\mathrm{HS}}} = \frac{e\left(p_j, u^{\mathrm{C}}\right)}{e\left(p_j, u^{\mathrm{HS}}\right)}. \tag{1}$$

Of course, having distinct values for every city would render such an index essentially useless for summarizing inequality. An important question then is, When is the inequality index independent of location? Surprisingly, this question has an elegant answer.[6]

> **Proposition 1:** In an equilibrium model of local labor markets, the inequality index is independent of location if and only if preferences are homothetic.
>
> **Proof:** First, a familiar result from price theory tells us that the expenditure function takes the form $e(p, u) = \psi(p)f(u)$ if and only if preferences are homothetic. Note then that if preferences are homothetic, the equality index in location j is $I_j = \left(\psi(p_j)\, f(u^{\mathrm{C}})\right)/\left(\psi(p_j) f(u^{\mathrm{HS}})\right) = \left(f(u^{\mathrm{C}})\right)/\left(f(u^{\mathrm{HS}})\right)$, which does not depend on local prices. Second, and more importantly, note that the converse is true: Let $I_j = g(u^{\mathrm{HS}}, u^{\mathrm{C}})$, so the index in location j does not depend on that location's

[4]Results generalize easily to the case with a vector of local and traded goods.

[5]Again, results generalize easily to accommodate many levels of education.

[6]This result—presented in a somewhat different form—first appeared in Black et al. (2009).

prices. Without loss of generality, we can take $u^{\mathrm{HS}} = 1$, $u^{\mathrm{C}} = u$. Then $I_j = \left(e(p_j, u)\right) / \left(e(p_j, 1)\right) = g(u, 1)$, so we can write $e(p_j, u) = e(p_j, 1) \cdot g(u, 1)$. Setting $\psi(p) \equiv e(p, 1)$ and $f(u) \equiv g(u, 1)$, we find that the expenditure function has the form $e(p, u) \equiv \psi(p)\, f(u)$, which implies that the preferences are homothetic.

Homothetic preferences play a crucial role in many models in economics, and it is worth examining the intuition as it pertains here. Because agents in the two cities face different prices, optimal consumption bundles will typically differ between the two cities. Thus, for example, a high school–educated agent in Austin, where housing is relatively inexpensive, may purchase more housing than he would if he lived in New York, where housing is dear. The same is true for those who are college educated. Importantly, however, if preferences are homothetic, this guarantees that our high school graduate in Austin will have the same budget share for local and traded goods as our college graduate in Austin. And this same condition holds in New York. Thus, the only difference between low- and high-skill workers is the scale of consumption, which can then be summarized by a single number. When we are in a spatial equilibrium, this scale factor must be the same for Austin as for New York.

If preferences are not homothetic, the problem identifying an index of inequality becomes quite complicated. Indeed, many simple concepts in economics become problematic. For instance, when preferences are not homothetic, the cost of living varies by income level. Because the budget shares of high- and low-income households are different, unless prices change proportionately, there is no single cost of living. Similarly, as Black et al. (2009) note, if local goods, primarily housing, are income inelastic and local amenities are normal goods, then the measured returns to education will be negatively correlated with local price levels. To avoid these complexities, we maintain the assumption of homotheticity, although we briefly discuss some implications of its violation later in the article.

Given this setup, we now turn to our topic of interest—changes in local inequality over time. We start with two dates, 0 and 1, and look at the change in inequality between those dates—maintaining the assumption of equilibrium. Let wages of high school–educated workers in location j in periods 0 and 1 be $w_j^{\mathrm{HS},0}$ and $w_j^{\mathrm{HS},1}$, respectively, and corresponding wages of college graduates be $w_j^{\mathrm{C},0}$ and $w_j^{\mathrm{C},1}$. Define the inequality index in periods 0 and 1 as

$$I_j^0 = I_j^0\left(p_j^0, u^{\mathrm{HS},0}, u^{\mathrm{C},0}\right) = \frac{w_j^{\mathrm{C},0}}{w_j^{\mathrm{HS},0}} = \frac{e\left(p_j^0, u^{\mathrm{C},0}\right)}{e\left(p_j^0, u^{\mathrm{HS},0}\right)}, \tag{2}$$

and

$$I_j^1 = I_j^1\left(p_j^1, u^{\mathrm{HS},1}, u^{\mathrm{C},1}\right) = \frac{w_j^{\mathrm{C},1}}{w_j^{\mathrm{HS},1}} = \frac{e\left(p_j^1, u^{\mathrm{C},1}\right)}{e\left(p_j^1, u^{\mathrm{HS},1}\right)}, \tag{3}$$

respectively. Then, by Proposition 1, both I_j^0 and I_j^1 are independent of location if and only if preferences are homothetic. The following result immediately follows.

Proposition 2: The change in inequality $I_j^1 - I_j^0$ is independent of location if and only if preferences are homothetic.

It is important to point out that even though our model predicts that in both periods inequality is the same across locations, and therefore that the growth in inequality is the same

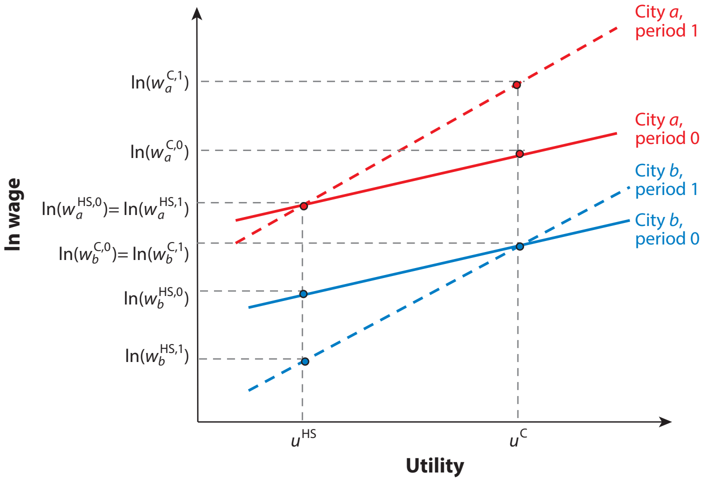

Figure 1

Relative earnings in two cities over time. Shown are the log(w)–education gradients in period 0 (*solid lines*) and in period 1 (*dashed lines*).

across locations, changes in inequality need not occur the same way in each location. **Figure 1** illustrates this point. Here we work with the log of the inequality measure, which is common in the literature; this is just the log(w)–education gradient. In this example, the solid lines shows such gradients in period 0. Under the assumption of homothetic preferences, these lines must be parallel. Now suppose that skill-biased technological change increases the relative demand of college-educated workers and decreases the relative demand of workers educated at the high school level. Suppose, moreover, that increased demand for high-skill workers happens primarily in city a, whereas the decreasing demand for low-skill workers is felt most prominently in city b.[7] Then the log(w)–education gradients in period 1 are given by the dashed lines. Again, these line are parallel. Equilibrium here requires of course that the price of housing drops in city b relative to city a. Importantly, real wage inequality changes by the same amount in each city.

As for the empirical assessment of inequality, **Figure 1** shows how misunderstandings might arise if analysts ignore variation across local labor markets. First, consider the issue of measuring inequality in the cross section, for example, in period 0. The object of interest is simply the return to a college education, relative to a high school education. Typically, estimates of the return to schooling are found using regression analysis that includes controls for experience or age via some function, g(age). Thus, a researcher who has a sample of workers who have a high school or college education typically employs a regression along the lines of

$$\log(w_i) = \beta_0 + \beta_1 \mathbb{I}_{C,i} + g(\text{age}_i) + \epsilon_i, \tag{4}$$

where $\mathbb{I}_{C,i}$ is an indicator variable equal to one if individual i has a college degree. As illustrated in **Figure 1**, however, to form an unbiased estimate, we typically also need to add location fixed effects; for individual i in city j,

[7]As noted above, and discussed in more detail below, New York is like city a, whereas Detroit is like city b.

$$\log(w_{ij}) = \beta_0 + \beta_1 \mathbb{I}_{C,ij} + g(\text{age}_{ij}) + \theta_j + \epsilon_{ij}. \tag{5}$$

If location fixed effects (θ_j) are omitted, estimates of β_1 are biased unless $\mathbb{I}_C$ is orthogonal to fixed effects (i.e., unless the education distribution is the same across cities).[8]

Second, when we evaluate changes in inequality, it is easy to see how incorrect inferences might arise if we fail to control for location. Typically, changes in the relative demand for workers across cities will also induce migration. Thus, for example, over time we might expect an outflow of workers (especially low-skill workers) from places like city b and an inflow of workers (especially high-skill workers) into places like city a. Failure to account for location could easily cause us to form biased estimates of changes in inequality over time.

We turn to empirical issues next. Our goal is to evaluate shifts in inequality from 1980 to 2010, using as our inequality measure the log wage gap between workers with college and high school education. Our theory tells us that the proper approach is to control for location when doing so. The clearest way of doing this is to first evaluate inequality changes at the local level and then aggregate to form estimates of national trends in inequality.

3. AN EMPIRICAL EVALUATION OF RISING INEQUALITY IN THE UNITED STATES, 1980–2010

Our empirical analysis is based on Public Use Microdata Samples (PUMS) of the US Census for 1980, 1990, and 2000 and the American Community Survey (ACS) for 2010. Empirical studies of inequality typically standardize on the race and/or gender distribution (see, e.g., Goldin & Katz 2008, Acemoglu & Autor 2011). We do this in a particularly simple way; we focus initially exclusively on white non-Hispanic men. We do so because we do not want our analysis of local labor markets to be affected by complications that might arise specifically for women or minority workers. (We return to these issues below.) As noted above, we further limit our sample to those who have either a high school education or a college-level education, by which we mean a four-year degree such as a bachelor of arts or science. This allows us to sidestep issues in the measurement of education among those with a low level of education and among those with graduate degrees.[9] We study wage and salary workers ages 25–64 who work 30 hours or more a week and 40 or more weeks per year.

3.1. Methods

Our estimation strategy is to use a two-stage procedure that first adjusts for nonresponse in the US Census data and then allows for a nonparametric regression. Because of item and unit nonresponse, the Census Bureau imputes missing responses. Although one might hope that the use of imputed data provides consistent estimates of the marginal distributions of relevant variables, in fact its use in regression analysis is generally problematic (for details, see Lillard et al. 1986, Bollinger & Hirsch 2006). However, dropping imputed data without further adjustment would potentially bias our estimates. Thus we employ inverse probability weighting to adjust the nonimputed data (see Wooldridge 2007). We first estimate the likelihood that an observation is imputed by the covariates we wish to use in our regression of interest. Letting Imp $= 1$ indicate that an observation has some (or all) imputed elements, we estimate

[8]Below we show that in fact the education distribution varies widely across cities.

[9]Black et al. (2003) discuss problems that arise in the measurement of higher education in the US Census.

$$\Pr(\text{Imp} = 1|\mathbf{X} = x) = P(x), \tag{6}$$

where $\mathbf{X}$ is the vector of covariates on which we wish to condition. We then drop the observations with imputed data and recalculate the sampling weights that the Census Bureau provides for observations without imputed data, using

$$\omega_{1,i} = \frac{\omega_{s,i}}{1 - P(x_i)}, \tag{7}$$

where $\omega_{s,i}$ are the Census Bureau weights, and $\omega_{1,i}$ are our imputation-adjusted weights. Although the use of $\omega_{1,i}$ will exactly reproduce the joint distribution of the vector $\mathbf{X}$ observed in the ACS or US Census data, this joint distribution uses the imputed data and so is only approximately correct.

When estimating the returns to schooling, we adopt the approach outlined in Black et al. (2006). Because our data have a limited number of distinct ages (25–64) with a college degree or high school diploma, we have discrete data, which allow us to proceed with a simple nonparametric approach. We use inverse probability weighting to implement the estimator proposed by DiNardo et al. (1996). The advantage of the nonparametric estimator is that it does not impose the restriction that all agents have the same return to a college education. Given that the return differs across agents, however, the measured return to a bachelor's degree will be a function of the selected distribution of covariates (ages). Our approach bases estimates on each location's age distribution of college degree holders, so our estimated parameter can be interpreted as the average return to a bachelor's degree in each city; in the nomenclature of the treatment effects literature, we estimate the average impact of "treatment" (a college education) on the "treated" (college graduates) in each city.[10]

The implementation of the estimator is quite simple. To begin, we estimate

$$\Pr(\mathbb{I}_{\mathrm{C}} = 1|\mathbf{X} = x) = q(x). \tag{8}$$

We then define the weights

$$\omega_{2,i} = \omega_{1,i} \frac{q(x_i)}{(1 - q(x_i))} \quad \text{if} \quad \mathbb{I}_{\mathrm{C},i} = 0, \tag{9}$$

$$\omega_{2,i} = \omega_{1,i} \quad \text{if} \quad \mathbb{I}_{\mathrm{C},i} = 1. \tag{10}$$

With these weights, we run weighted least squares on the equation

$$\log(w_i) = \beta_0 + \beta_1 \mathbb{I}_{\mathrm{C},i} + \varepsilon_i. \tag{11}$$

The estimated coefficient $\hat{\beta}_1$ will be the estimated average return to a college degree for a college-educated individual. This is equivalent to calculating the mean return of a bachelor's degree for each age (e.g., Δ_k) and then calculating $\beta_1 = \sum \omega_k \Delta_k$, where the weights are the relative frequencies of college degree holders at each age. In other words, we are implementing the regression in Equation 4 with the most flexible possible specification of g(age)—treating each age as an indicator variable.

[10]Although the terminology "treatment on the treated" helps focus attention on the intent of our empirical exercise, of course we must bear in mind that the "treatment" is not randomly assigned. College attendance is a choice, so our estimates incorporate both selection into college and the "treatment" for those who attend college.

Next, we repeat the exercise, but this time we add location fixed effects as Equation 5 suggests we should do to get an unbiased estimate.

Finally, in line with our theory, we undertake the analysis described above separately for large cities in the United States. Our goal of course is to examine how wage inequality—as measured by the log wage gap between college-educated and high school–educated men—evolves at the local level. In so doing, we are taking a streamlined approach that abstracts from any number of complicating factors. To mention just one, it is certainly possible that particular cities also have disproportionately more workers holding bachelor's degrees from elite institutions, and if the relative mix changes over time, this could account for some of the evolution in wage inequality.[11]

3.2. Findings

Table 1 presents estimated returns to a college degree in 1980, 1990, 2000, and 2010. In the first row of this table, we start with trends for the United States as a whole. Consistent with observations made by many analysts (e.g., Acemoglu & Autor 2011, and references therein), we estimate an upward trend in wage inequality between college-educated and high school–educated men—with the log wage gap growing from 0.248 in 1980 to 0.522 in 2010 (i.e., increasing by 0.274). Using our approach that controls for location, in the second row we find that the within-location wage inequality in the United States grew by somewhat less over this time period, by 0.253.[12] There are also interesting differences in the timing of inequality changes; overall increases in inequality appear to be smaller than the within-location increases in inequality from 1980 to 1990, but the opposite is true thereafter (see, in particular, the first two rows of the bottom half of **Table 1**).

Next, we turn our attention to 21 large metropolitan statistical areas (MSAs). We notice two interesting features. First, by comparing statistics in the third and fourth rows of **Table 1**, we see that in every year, within-location inequality in these urban areas is higher than urban inequality more generally.[13] Furthermore, within-location urban inequality is substantially higher than within-location inequality across the entire United States. For example, in 2010, inequality within our large 21 MSAs averaged 0.565 log points, compared to 0.504 log points in the United States generally—a difference of more than 10%. Second, when we focus attention on trends in inequality in urban areas, we see patterns that are similar to the entire United States. Specifically, the increase in wage inequality is moderately overstated when location controls are not used.

The evidence is at least moderately consistent with the notion that large cities have emerged as hot spots of wage inequality. Several observers have noted that inequality might be especially large in locations with high concentrations of high-skill workers. In an earlier paper, Black et al. (2009),

---

[11]Using detailed controls for academic ability, Monks (2000) and Black & Smith (2004, 2006) find that students attending more elite universities have higher returns than students attending institutions of lower quality. Hoekstra (2009) uses a regression discontinuity design to test whether attending a state's flagship school increases earnings relative to other options and finds substantial returns to the higher-quality institution. Similarly, using detailed controls, Andrews et al. (2012) find substantial returns to attending Texas flagship institutions relative to lower-quality institutions. [In contrast, Dale & Krueger's (2002) study of elite colleges finds no return to college quality.]

[12]In this analysis, we define location as a metropolitan statistical area (MSA) for those respondents who live in an MSA and as a state of residence for those who live outside any MSAs.

[13]Black & Sanders (2012) examine the growth in inequality from 1980 to 2000 for the 48 contiguous states. They find much faster growth in wage inequality in urban counties compared with rural counties. Indeed, they find that the value of the 90th percentile of wage actually fell in most rural counties; thus, the growth of wage inequality appears to be an urban phenomenon. Moretti (2011) finds that there are large wage differences across numerous MSAs in 1980 and that these differences persist between 1980 and 2000. Similarly, Lindley & Machin (2014) find evidence of substantial variation in the returns to a college degree and a high degree of spatial persistence in that variation using 1980–2000 US Census data and data from the ACS for 2010.

Table 1 Returns to a bachelor's degree, 1980–2010[a]

Weekly earnings premia for college graduates relative to high school graduates	1980	1990	2000	2010
United States, no location controls	0.248 (0.0038)	0.393 (0.0035)	0.463 (0.0033)	0.522 (0.0043)
United States, with location controls	0.251 (0.0032)	0.356 (0.0025)	0.428 (0.0026)	0.504 (0.0034)
21 MSAs, no location controls	0.245 (0.0079)	0.365 (0.0059)	0.451 (0.0060)	0.547 (0.0076)
21 MSAs, with location controls	0.269 (0.0068)	0.370 (0.0052)	0.462 (0.0048)	0.565 (0.0062)
Growth in returns to bachelor's degree		**1980–1990**	**1990–2000**	**2000–2010**
United States, no location controls		0.145	0.070	0.059
United States, with location controls		0.105	0.072	0.076
21 MSAs, no location controls		0.120	0.086	0.096
21 MSAs, with location controls		0.101	0.092	0.103

[a]Authors' calculations based on the 1980, 1990, and 2000 Census Public Use Microdata Samples and the 2010 American Community Survey. The sample consists of white, non-Hispanic men ages 25–64, who worked at least 40 weeks a year and at least 30 hours a week, have either a high school or bachelor's degree, and are wage and salary workers. The dependent variable is the natural logarithm of weekly earnings. Workers with any imputed data are dropped, but we use inverse probability weighting to ensure that the joint distribution of location, education, and age is maintained after we drop imputed data. We then reweight data to match the local distribution of workers with a bachelor's degree. Parameters are then estimated using weighted least squares and should be interpreted as the impact of treatment on the treated. Abbreviation: MSA, metropolitan statistical area.

who follow the work of Rauch (1993) and Berry & Glaeser (2005), suggest that high concentrations of human capital might be correlated with returns to education, specifically arguing that "the observed return to college would be higher in metropolitan areas with higher concentrations of college-educated individuals."[14]

With this in mind, consider **Table 2**, which provides the fraction of workers with a college degree or higher separately for our 21 large MSAs, all other MSAs, and the areas outside of MSAs. In 1980, the fraction of highly educated workers in 21 large MSAs varied between 20% and 57% and averaged 30%, compared to only 17% in rural areas and 25% in smaller MSAs. These rural-urban disparities in the concentration of skilled workers if anything increased over 30 years. By 2010, the proportion of those with at least a college degree varied from 30% to 62% in 21 large MSAs. At the same time, only 21% of workers in rural areas held a college degree or higher, and in smaller MSAs, that proportion was 33%. Notice that the proportion of workers with a bachelor's degree or higher grew on average by 45% between 1980 and 2010 in 21 large MSAs, but only by 24% in rural areas and by 32% in smaller MSAs. In any event, **Table 2** shows that the distribution of worker skills differs dramatically by location and confirms that the distribution of skills has been changing over time.

Tables 3 and **4** break down the cross-location evidence about inequality in more detail yet—at the city level. **Table 3** shows log wages. To facilitate comparisons across cities and over time, we normalize the 1980 average log wage of high school–educated men to be zero. In addition, for years other than 1980, we adjust wages using the CPI. Thus, all values in **Table 3** can be interpreted as log wages relative to wages of high school graduates in 1980 adjusting for national trends in inflation.

[14]Black et al. (2009) also provide a theoretical treatment when preferences are nonhomothetic, in which case we would expect differences across location in the nominal return to education.

Table 2 Fraction of workers with a bachelor's degree or higher, 1980–2010[a]

	1980	1990	2000	2010
Non-MSAs	0.166	0.186	0.192	0.209
21 MSAs	0.305	0.368	0.410	0.445
Highest	0.571	0.554	0.579	0.623
Lowest	0.204	0.246	0.280	0.295
All other MSAs	0.246	0.279	0.309	0.332

[a]Authors' calculations based on the 1980, 1990, and 2000 Census Public Use Microdata Samples and the 2010 American Community Survey. The sample consists of white, non-Hispanic men, ages 25–64. Abbreviation: MSA, metropolitan statistical area.

Inspection of **Table 3** reveals a number of interesting features of inequality in the United States. Consider, to begin, high school–educated men in the base year, 1980. Wages for these men differed substantially across cities. Detroit, with a large auto industry that still had strong employment in 1980, had by far the highest wage—0.23 log points higher than the national average wage (normalized to 0.00). By way of comparison, wages for this group were lower than the national average in several large cities (e.g., as low as −0.17 in Tampa). Continuing our focus on 1980, notice next that the log wage gap between college and high school graduates varied substantially across cities, ranging from 0.16 (Seattle) and 0.19 (Chicago) to 0.33 (New York) and 0.35 (San Diego).

Our observation about cross-location variation in the wage gap between college and high school graduates pertains to other years as well; the range was 0.27–0.46 in 1990, 0.36–0.55 in 2000, and 0.49–0.64 in 2010. As for trends over time, we have already seen that inequality increased broadly in the United States—the log wage gap between college and high school graduates increased from 0.25 in 1980 to 0.50 in 2010. **Table 3** shows that inequality increased in each of the 21 cities we study.

Interestingly, however, the 1980–2010 increases in inequality happened in quite different ways in different cities. Consider, for example, the two cities in Pennsylvania: In Philadelphia, real wages among high school graduates held roughly constant from 1980 to 2010, while wages among the college educated rose substantially. In contrast, in Pittsburgh, wages of high school graduates dropped precipitously, while the inflation-adjusted wages of college-educated men barely held steady. (Again, we remind readers that we are using the national CPI.)

Table 4 provides statistics that allow us to follow up on our last point above. In this table, we report relevant changes based on **Table 3**. The first set of columns shows how inequality increased in each city from 1980 to 1990; the second and third sets of columns do the same for 1990–2000 and 2000–2010, respectively; and the fourth set of columns gives changes for our entire period of study, 1980–2010. Inequality increases in every decade in every city. Over the entire time period, there is certainly some variation in the increase in inequality across cities, but if we set aside the three most extreme cities—San Francisco (a remarkable increase in inequality of 0.41), and San Diego and Baltimore (inequality increases of 0.19 and 0.20, respectively)—all other cities saw increases in inequality of between 0.24 and 0.35, a reasonably narrow range.

Figure 2 provides a graphical representation of these results. For each city j, the lower end of the bar corresponds to the 30-year change in real $\log(w_j^{\mathrm{HS}})$, and the upper end corresponds to the corresponding change in real $\log(w_j^{\mathrm{C}})$. The height of the bars thus gives the increase in the log wage

Table 3 Real wages relative to wages of high school graduates in 1980[a]

	1980			1990			2000			2010		
	BA wage	HS wage	Gap	BA wage	HS wage	Gap	BA wage	HS wage	Gap	BA wage	HS wage	Gap
United States	0.25	0.00	0.25	0.35	−0.01	0.36	0.42	−0.01	0.43	0.43	−0.08	0.50
Atlanta	0.32	0.02	0.31	0.46	0.05	0.42	0.56	0.04	0.52	0.54	−0.04	0.58
Baltimore	0.33	0.03	0.30	0.43	0.06	0.37	0.48	0.05	0.43	0.55	0.06	0.49
Boston	0.20	−0.07	0.27	0.43	0.08	0.35	0.55	0.05	0.50	0.57	−0.02	0.59
Chicago	0.38	0.18	0.19	0.45	0.11	0.34	0.53	0.15	0.38	0.52	0.00	0.52
Cleveland	0.34	0.13	0.21	0.39	0.02	0.37	0.42	0.00	0.43	0.37	−0.15	0.51
Dallas	0.33	0.02	0.31	0.41	−0.02	0.44	0.56	0.02	0.54	0.58	−0.02	0.60
Denver	0.27	0.00	0.26	0.29	−0.07	0.36	0.41	0.00	0.40	0.44	−0.08	0.52
Detroit	0.45	0.23	0.22	0.48	0.13	0.34	0.56	0.15	0.41	0.44	−0.08	0.52
Houston	0.46	0.16	0.29	0.50	0.04	0.46	0.64	0.09	0.55	0.72	0.08	0.64
Los Angeles	0.37	0.08	0.29	0.53	0.15	0.38	0.59	0.08	0.51	0.59	0.03	0.57
Minneapolis	0.31	0.09	0.22	0.33	0.03	0.31	0.45	0.03	0.41	0.45	−0.07	0.52
New York	0.37	0.04	0.33	0.55	0.18	0.37	0.62	0.16	0.46	0.68	0.08	0.60
Philadelphia	0.32	0.04	0.27	0.45	0.08	0.37	0.52	0.08	0.44	0.54	0.02	0.52
Phoenix	0.21	−0.02	0.22	0.34	−0.10	0.45	0.42	−0.04	0.46	0.42	−0.14	0.56
Pittsburgh	0.31	0.07	0.24	0.33	−0.14	0.47	0.34	−0.14	0.49	0.30	−0.22	0.52
San Diego	0.25	−0.10	0.35	0.35	−0.01	0.36	0.46	−0.04	0.49	0.51	−0.03	0.54
San Francisco	0.32	0.11	0.20	0.46	0.14	0.32	0.63	0.16	0.48	0.70	0.09	0.61
Seattle	0.30	0.14	0.16	0.34	0.08	0.27	0.47	0.11	0.36	0.58	0.07	0.51
St. Louis	0.30	0.06	0.24	0.36	0.01	0.34	0.40	−0.01	0.41	0.42	−0.11	0.52
Tampa	0.13	−0.17	0.30	0.28	−0.19	0.46	0.35	−0.16	0.50	0.37	−0.24	0.61
Washington	0.41	0.10	0.31	0.46	0.11	0.35	0.56	0.08	0.48	0.68	0.05	0.64

[a]Authors' calculations based on the 1980, 1990, and 2000 Census Public Use Microdata Sample and the 2010 American Community Survey. The sample consists of white, non-Hispanic men ages 25–64, who worked at least 40 weeks a year and at least 30 hours a week, have either a high school or bachelor's degree, and are wage and salary workers. Wages are adjusted for national trends in inflation using the Consumer Price Index. The 1980 average log wage of high school–educated men is normalized to zero. Abbreviations: BA, bachelor of arts; HS, high school.

Table 4 Growth of real wages and real earnings inequality[a]

	1980–1990			1990–2000			2000–2010			1980–2010		
	ΔBA	ΔHS	ΔGap	ΔBA	ΔHS	ΔGap	ΔBA	ΔHS	ΔGap	ΔBA	ΔHS	ΔGap
United States	0.10	−0.01	0.11	0.07	−0.01	0.07	0.01	−0.07	0.08	0.17	−0.08	0.25
Atlanta	0.14	0.03	0.11	0.10	−0.01	0.11	−0.02	−0.08	0.06	0.22	−0.06	0.28
Baltimore	0.11	0.03	0.08	0.05	−0.01	0.06	0.07	0.01	0.07	0.23	0.03	0.20
Boston	0.23	0.15	0.08	0.12	−0.03	0.15	0.02	−0.08	0.09	0.37	0.05	0.32
Chicago	0.07	−0.07	0.14	0.09	0.04	0.05	−0.02	−0.15	0.14	0.14	−0.19	0.33
Cleveland	0.05	−0.11	0.16	0.03	−0.02	0.06	−0.06	−0.14	0.09	0.03	−0.28	0.31
Dallas	0.09	−0.04	0.13	0.15	0.04	0.11	0.02	−0.04	0.06	0.25	−0.04	0.29
Denver	0.02	−0.08	0.10	0.12	0.07	0.04	0.03	−0.08	0.12	0.17	−0.09	0.26
Detroit	0.03	−0.09	0.12	0.09	0.02	0.07	−0.12	−0.24	0.11	−0.01	−0.31	0.30
Houston	0.04	−0.12	0.16	0.14	0.05	0.09	0.08	−0.02	0.10	0.26	−0.09	0.35
Los Angeles	0.16	0.06	0.10	0.06	−0.07	0.13	0.01	−0.05	0.06	0.22	−0.06	0.28
Minneapolis	0.02	−0.06	0.08	0.12	0.01	0.11	0.00	−0.11	0.11	0.14	−0.16	0.30
New York	0.19	0.14	0.04	0.07	−0.02	0.09	0.06	−0.08	0.14	0.31	0.04	0.27
Philadelphia	0.14	0.04	0.10	0.06	0.00	0.06	0.02	−0.06	0.08	0.22	−0.02	0.24
Phoenix	0.13	−0.09	0.22	0.08	0.07	0.01	0.00	−0.10	0.10	0.21	−0.12	0.33
Pittsburgh	0.02	−0.21	0.23	0.01	−0.01	0.02	−0.04	−0.07	0.03	−0.01	−0.29	0.28
San Diego	0.10	0.10	0.00	0.11	−0.03	0.13	0.05	0.00	0.05	0.26	0.07	0.19
San Francisco	0.15	0.03	0.12	0.17	0.01	0.16	0.07	−0.07	0.13	0.38	−0.02	0.41
Seattle	0.04	−0.06	0.10	0.13	0.04	0.09	0.11	−0.04	0.15	0.28	−0.06	0.34
St. Louis	0.06	−0.05	0.11	0.04	−0.03	0.07	0.02	−0.09	0.11	0.12	−0.16	0.29
Tampa	0.15	−0.01	0.16	0.07	0.03	0.04	0.03	−0.08	0.11	0.25	−0.07	0.32
Washington	0.05	0.01	0.04	0.10	−0.03	0.14	0.12	−0.03	0.15	0.28	−0.05	0.33

[a]Authors' calculations based on the 1980, 1990, and 2000 Census Public Use Microdata Samples and the 2010 American Community Survey. The sample consists of white, non-Hispanic men ages 25–64, who worked at least 40 weeks a year and at least 30 hours a week, have either a high school or bachelor's degree, and are wage and salary workers. Wages are adjusted for national trends in inflation using the Consumer Price Index. Abbreviations: BA, bachelor of arts; HS, high school.

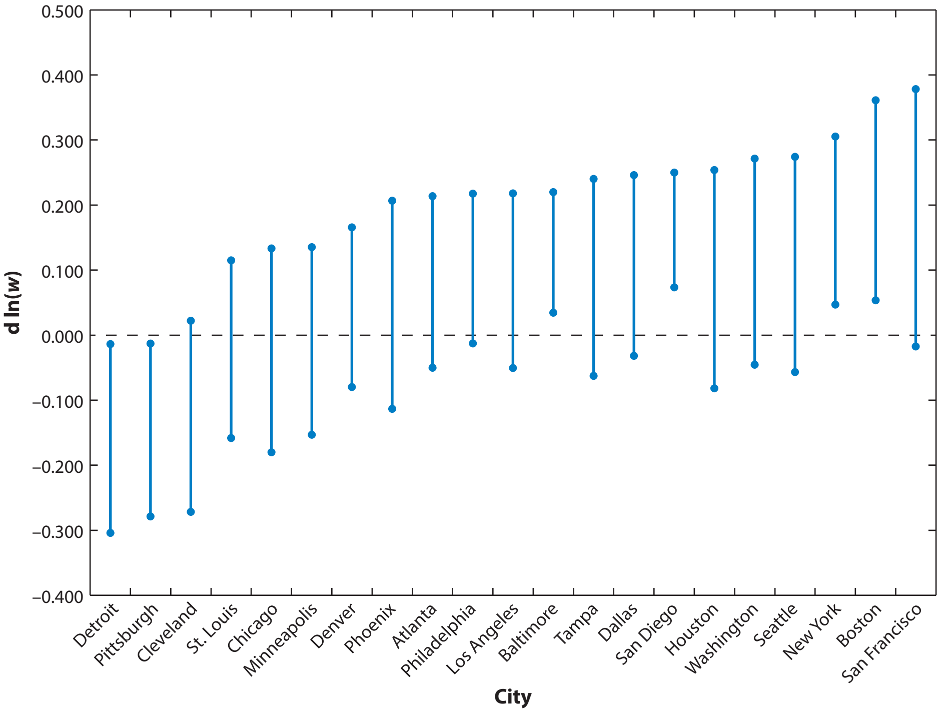

Figure 2

The growth in real earnings for high school– and college-educated men, 1980–2010. Calculations are the authors based on the 1980, 1990, and 2000 Census Public Use Microdata Samples and the 2009, 2010, and 2011 American Community Surveys. For each city, the point at the base of the bar corresponds to the change in the wage of high school graduates, and the point at the top of the bar is the change in the wage of college graduates.

gap in city *j*. As we see, inequality increases everywhere by a substantial amount. Aside from San Francisco, which had an unusually large increase in inequality, and San Diego and Baltimore, which had atypically small increases in inequality, the heights of the bars appear to be reasonably similar across cities. However, **Figure 2** underscores a key point made above: In some cities (e.g., Detroit, Pittsburgh, and Cleveland), increasing inequality is primarily a phenomenon of declining real wages among high school graduates, whereas in other cities (e.g., New York, Boston, and San Francisco), it is primarily a phenomenon of increasing real wages among the college educated.

These basic facts strike us as potentially useful for sorting out the causes of rising inequality. They are also surely important for understanding the economic impact of rising inequality. For the moment, we mention two points along these lines.

One implication has to do with disparities across location in the burden of federal taxation—an issue analyzed in the important work of Albouy (2009). Albouy provides careful documentation of the unsurprising fact that cities with high wages tend also to have high housing prices (doing so with 2000 data). Given the progressive nature of federal income taxation, and given that tax schedules are based on nominal income, this means that workers in expensive locations pay more in federal taxes than they would if they lived in inexpensive locations. Clearly this has been changing over time. For example, the tax bill for workers in San Francisco has been rising relative to the tax bill for workers in Pittsburgh.

A second observation follows from the fact that inequality changes in different ways in different cities, that is, that several cities (e.g., Detroit and Pittsburgh) have experienced increasing inequality via declining real wages for low-skill workers and declining real housing values. These forces in combination can have economic consequences for some individuals that extend well beyond simple wage stagnation. Consider, for instance, a middle-aged high school graduate working in the steel or auto industry in 1980. That individual may very well have invested heavily in industry-specific human capital that suddenly lost value, invested in a house in a market with declining real prices, and seen hoped-for pensions vanish.[15] Furthermore, moving to a new labor market might be difficult for any number of reasons (e.g., local ties to friends and family). In this sense, measured log wage gaps might substantially understate the impact of industrial shifts on real economic inequality.

Below we pursue a broader discussion of issues and puzzles raised by examining inequality in local labor markets. First, though, we have a few observations about our theoretical approach.

3.3. Discussion

Our theory sets out a simple and coherent way of thinking about the evolution of local inequality. As is standard in urban economics, we start with the premise that locations differ in terms of local amenities (e.g., weather, local beauty, cultural assets) and/or productivity (due to differences in resources or owing to production agglomerations), and in turn, prices and wages can be expected to vary across locations. Even so, if moving costs are negligible, people should be indifferent over where they live. Thus, real inequality—the utility differences between differentially skilled workers—should be approximately the same in all locations. Our theory says that if, in addition, preferences are homothetic, the standard metric for measured inequality, $\log(w^{C}) - \log(w^{HS})$, should be similar across locations. As for the evolution of inequality, we expect that over time our inequality measure should be rising (or declining) by roughly the same amount in all locations.

Before proceeding, it is worth asking if our theoretical reasoning about local inequality represents a step forward in conceptualizing the evolution of inequality. We can see why the answer to this query is not obvious.

A glass-half-empty view of our work would emphasize that the evidence lines up poorly with the theory. In particular, our model suggests that the return to college, $\log(w^{C}) - \log(w^{HS})$, should be roughly the same across locations—a prediction that is clearly counterfactual (**Table 3**). Furthermore, one might argue that controlling for location does not make too much of a difference in overall inferences about trends in inequality (**Table 4**). So why bother with all the complication?

We prefer a glass-half-full perspective to our approach. In defense of our work we note, first of all, that at a minimum, our theory clarifies the economic interpretation of empirical metrics such as $\log(w^{C}) - \log(w^{HS})$ for the use of tracking inequality (between locations or over time). As prices vary across locations and over time, such metrics can be interpreted as measuring real inequality only under the assumption that preferences are homothetic. Second, as long as homotheticity is being implicitly assumed, researchers may as well make the assumption explicit and then exploit the resulting implications in specifying estimation procedures. We can see no good reason for ignoring cross-location variation in inequality when measuring time trends in inequality. Having said all this, our third point is an acknowledgment that our theory is of course flawed. Economic

[15]To put this story into perspective, we note that in 1980 in Detroit, 45% of male employment generally, and 50% of black male employment specifically, was in manufacturing. Within two decades, the corresponding figures declined to 33% and 26%, respectively (see Black et al. 2010).

models are at best an imperfect approximation of a complicated reality. Pursuing this last point further, it is worth revisiting the two key assumptions in our theory—costless migration and the homotheticity of preferences.

In fact, although the costless migration assumption is widely adopted in the urban economics literature, it is surely wrong. Indeed, work by Kennan & Walker (2011) suggests that many people behave as if there is a very high latent cost of moving (the cost is latent because only some workers actually chose to move in equilibrium), perhaps in excess of $300,000 for the average worker.[16] This is equivalent to people having excessive attachment to the places where they currently live. With this in mind, it would be sensible to extend our approach to allow for heterogeneity in the extent to which people are attached to places where they live (e.g., as in Moretti 2011), although that would doubtlessly add considerable complication to the analysis of evolving inequality. We should expect high moving costs to lead to a great deal of stickiness in local labor markets so that our predictions hold only in the long run.[17] Interestingly, as noted above, when we look at changes in inequality over 30 years, we do see that our inequality measure increases by reasonably similar amounts in the majority of cities; perhaps our model does not do too badly in the very long run.

Similarly, the assumption that preferences are homothetic is very common in economic models (e.g., the assumption appears in classical trade models), but the assumption may be problematic for our purposes. Some headway can be made if we are willing to take a stand on the form of nonhomotheticity. For example, Black et al. (2009) show that if the income elasticity of demand for housing is less than one, $\log(w)$–education gradients will be flatter in expensive cities than in low-priced cities. Looking at the evidence presented in **Tables 3** and **4**, however, we see no clear evidence about how nonhomotheticity might play a role in driving the evolution of measured inequality; we leave it to future work (and other researchers) to pursue this possibility more carefully.

4. ISSUES AND UNANSWERED QUESTIONS

We hope that our theoretical approach provides insight into the evolution of local inequality, and we do think that it at least represents a helpful jumping-off point for improving usual practice. Here we discuss several paths of inquiry that might be clarified by our model, or by other approaches that treat local labor markets more carefully than is typical in the inequality literature.

4.1. Real Wage Inequality

The empirical work we conducted above—crude as it is—suggests that researchers might well be overestimating the increase in wage inequality in the United States over the past 30 years by ignoring that people work in local labor markets.

In an important paper that studies this issue with considerable care, Moretti (2013) argues that the differential sorting of high-skill workers into large expensive cities has indeed caused researchers to overestimate the growth of inequality over recent decades. Moretti makes valiant

[16]Kennan & Walker (2011) show, however, that costs are quite heterogeneous, and many observed moves actually have estimated "negative costs."

[17]Importantly, Cadena & Kovak (2013) show that during the Great Recession, low-skill Mexican-born immigrants demonstrated a high willingness to move in response to shifts in local labor demand, thereby contributing substantially to the equalization of spatial differences in labor market outcomes for other workers. This serves as a reminder that only marginal workers need to have low migration costs for the basic message of our model to pertain.

attempts to capture differences in the cost of living across the extremely heterogeneous cities in the United States, primarily using measures of the price of housing. He finds that at least 22% of the increase in inequality between 1980 and 2000 was the result of this sorting of high-skill workers into more expensive cities.[18]

In an earlier working paper version of the 2013 paper, moreover, Moretti (2008) offers an intriguing analysis that focuses on within-city variation in the prices faced by workers. In these specifications, Moretti finds that earnings inequality is overstated by 40–50% when accounting for local price differences paid by high-skill workers and low-skill workers. This is necessarily controversial; people in the same community do in general have the opportunity to consume the same housing, pay for the same services, and make purchases at the same shops. The law of one price would seem to imply that all consumers face the same prices. However, when preferences are not homothetic, consumers of differing income levels may purchase fundamentally different goods. For instance, suppose that houses come in two varieties—1,500-ft^2 homes and 3,000-ft^2 homes. If high-skill workers purchase only 3,000-ft^2 homes and low-skill workers purchase only 1,500-ft^2 homes, then factors affecting the relative demand for high-skill and low-skill workers in a community will also be reflected in the relative prices of these two types of housing. We think that an important goal for future research is to better understand the implications of differential consumption patterns of low-skill and high-skill workers for the measurement of inequality.

4.2. The Evolution of Black-White Disparity

Quite apart from the literature on inequality in labor markets, a significant body of research explores the evolution of disparity in wages between black and white workers. Examples include the seminal work of Smith & Welch (1989) and many subsequent analyses. Most studies rely on national-level data, with only limited attention to local variation (attention is typically only given to regional variation, at most). Black et al. (2013b) set up a theoretical structure similar to the one given above and argue that it is crucial to take location into account when assessing black-white economic disparity.

To see the issue at hand, suppose we begin our analysis in 1940. Restricting attention to black and white prime-age men (ages 25–55), we find from PUMS data that in 1940 southern residency was 2.87 times more likely for black men than for white men. In that same year, urban residency was 0.76 times more likely for black men (i.e., African Americans were more likely to live in rural areas). By 2000, the southern residency index had fallen to 1.59, and the urban residency index increased to 1.15. It seems clear that a comprehensive assessment of long-run changes in racial economic disparity should explore the impact of the disproportionate movement of the black population from low-wage, low-price rural locations in the south to higher-wage, expensive urban locations outside the south.

With this in mind, consider the first column of **Table 5**, which includes results from a non-parametric analysis similar to the one described above but comparing black and white men (rather than college graduates and high school graduates), matching exactly on age but using no other covariates. We estimate that for prime-age men, the black-white log wage gap declined from an

[18]Moretti (2013) uses all the cities identified in the 2000 US Census in his analysis, which requires him to approximate the cost of living in 315 different locations. To do this, he relies on some fairly strong assumptions about the nature of the differences in the cost of living across locations, but he shows that his results are robust to several different specifications of the local cost-of-living index.

Table 5 Black-white gaps in the log weekly wages[a]

Year	Controlling for age only	Controlling for age and location
1940	−0.741	−0.662
1950	−0.511	−0.485
1960	−0.510	−0.489
1970	−0.447	−0.448
1980	−0.308	−0.332
1990	−0.281	−0.323
2000	−0.259	−0.310

[a]Authors' calculations, 1940–2000 Census Public Use Microdata Samples. The dependent variable is the logarithm of weekly earnings. Readers are referred to Black et al. (2013b) for details.

astounding −0.741 in 1940 to −0.259 in 2000 (i.e., the gap dropped by approximately two-thirds). If we conduct this same analysis and match on age and location, we estimate that the gap dropped from −0.662 to −0.310 (i.e., by only a little more than one half) (see the second column of **Table 5**).[19] In the analysis that conditions on location, we see that the black-white wage gap is little changed since 1980. As many observers have noted, the increase in labor market inequality that began around 1980 has disproportionately hurt minority workers. It is true that from 1940 through 1980, there was a substantial drop in black-white wage disparity, likely because of improving educational opportunities for black students and reduced racial discrimination in labor markets. However, when we take account of the local labor markets, we see no evidence that this trend is continuing in an era of increasing labor market inequality more broadly (i.e., since 1980).[20]

4.3. Labor Market Outcomes for Women

Our analysis to this point has focused exclusively on men. However, women are an extremely important part of the labor market. Indeed, women currently comprise 47% of the US labor force. Because young women have substantially surpassed their male counterparts in many measures of educational attainment (e.g., at present 56% of college students are women), women at all skill levels are likely to be an increasing presence in labor markets.

Economists have often focused attention on men, rather than women, in the labor market because of complications associated with studying the labor supply of women specifically. A thorough analysis of the evolution in wage inequality among women over the past several decades, for instance, would require careful attention to the rapid increase in women's labor force participation and, moreover, the likely important differences between women who now are in the labor force and those women who do not work (and that selection into the labor force could be changing over time). This same observation holds for spatial variation. Black et al. (2014) show

[19]Black et al. (2013b) provide details. We conducted this same analysis matching also on schooling. The same basic observation pertains; measured black-white wage disparity declines by much less if we condition on location than if we do not.

[20]Moreover, racial gaps exist not only in wages but also in unemployment; black men are substantially more likely to be unemployed than comparable white men (see, e.g., Ritter & Taylor 2011). Black et al. (2010) show that from 1980 to 2000, black male participation in the labor force declined in every one of the 14 large cities in their analysis.

that there are dramatic differences in the labor supply of married women across US cities. In 2000, for example, 79% of married, non-Hispanic white, high school–educated women ages 25–55 were employed in Minneapolis, whereas the comparable percent was only 52 in New York. We argue that one factor driving these locational differences in the labor supply of married women (especially those with young children) is that married women are very sensitive to the commuting time in the city where they live. We show that the labor force participation rate of women is lower in cities with higher average commuting times. We also find that the growth of commuting time over recent decades is negatively associated with changes in the city's female labor force participation.

All of this suggests that a comprehensive analysis of inequality that includes women will require a great deal of care. In any event, that analysis surely must take account of the local labor markets in which women work. As a first step, consider **Figure 3**, which undertakes the same analysis as shown in **Figure 2** but for women. There are some interesting similarities between the two figures and also some important differences. As with men, there is a substantial wage gap as measured by the difference in real log wages between female college and high school graduates. **Figure 3** shows that this measure of inequality increased in every city from 1980 to 2010. That increase was typically on the order of 0.25–0.35, as it was for men, but there are several exceptions.

One important difference between men and women in the evolution of inequality is that real wages for women (as calculated by the national CPI), both high school and college graduates, were

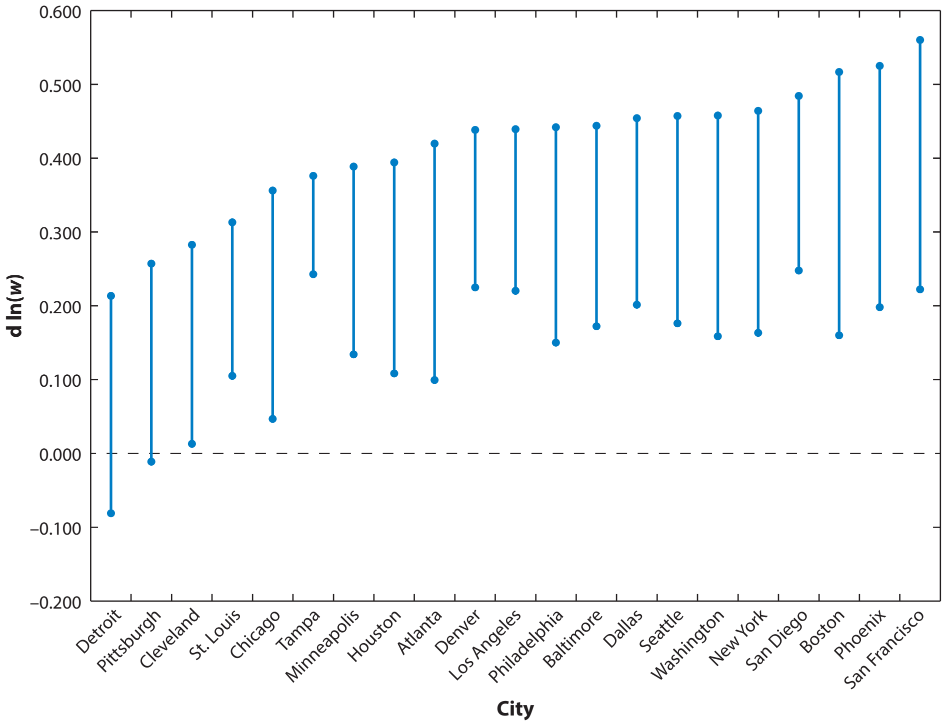

Figure 3

The growth in real earnings for high school– and college-educated women, 1980–2010. Calculations are the authors based on the 1980, 1990, and 2000 Census Public Use Microdata Samples and the 2009, 2010, and 2011 American Community Surveys. For each city, the point at the base of the bar corresponds to the change in the wage of high school graduates, and the point at the top of the bar is the change in the wage of college graduates.

increasing in almost all cities, the only exceptions being high school graduates in Detroit and Pittsburgh. In contrast, in most cities, the real wage of male high school graduates dropped during this period. Understanding the forces that shaped these differences is well beyond the scope of our work here. Our primary point for the moment, as mentioned above, is that for women as for men, there is an important spatial element to the evolution of inequality.

4.4. Families

Most people belong to families. As Blundell et al. (2012) argue, a comprehensive understanding of inequality requires attention to links between wage inequality and household consumption, including an understanding of family labor decisions. In this light, the findings we have just mentioned—that married women's labor supply decisions vary substantially across cities—are likely an important indicator of the necessity of studying the role of location for family decisions of all sorts. One particularly interesting feature of location, discussed in Black et al. (2013a), is that families at all income levels have fewer children when they live in high-price cities. The argument developed in that paper is that children are likely "normal" (the paper presents evidence that, all else equal, higher income leads people to have more children) and raising children is especially costly in high-price cities, as children require living space. This is a nice example of how nonhomotheticity in preferences might be important. Along these same lines, Black et al. (2002) suggest that nonhomothetic preferences might be key to understanding why gay men tend to disproportionately locate in high-amenity, high-price cities (e.g., San Francisco and Washington, DC).

Black et al. (2010) compare the family income of African American children ages 8–12 in 14 large cites to similarly aged white children from the same cities. In 1980, in each of the cities studied, the median black child lived in a family with an income that would have placed that family near the bottom of the corresponding white distribution (i.e., white families with similarly aged children). Specifically, in Los Angeles, the city with the least racial disparity, the median black child would have been in the 20th percentile of the white distribution. In Chicago, the city with the greatest disparity, he or she would have been in the 11th percentile of the white distribution. Between 1980 and 2000, this statistic improved for African Americans, albeit irregularly, in a few cities: Atlanta, Baltimore, Cleveland, Houston, New York, and New Orleans. However, it remained unchanged in Chicago, St. Louis, and Washington and actually declined in Detroit, Los Angeles, Memphis, Philadelphia, and San Francisco. By 2000, this statistic ranged from a high of the 23rd percentile of white families in New York to a low of 11% in both Chicago and San Francisco. Although this statistic nicely summarizes the immense differences in economic resources available to white and black children, it hides the complex process that contributed to these resource gaps: the improving wages of black women, the decline in employment rates of black men, and the drastic decline in the number of black children living with two parents.

Again, an evaluation of inequality that factors into family-level consumption lies well outside the scope of our article. We do want to point out, however, that such an analysis surely would focus on the local nature of labor and housing markets.

4.5. Life-Cycle Effects

Blundell et al. (2012) also emphasize that it is important for families to take a life-cycle perspective to understand wage inequality. Having more than one potential earner allows households to smooth consumption over time. A multiperson household will also have a more complex decision to make than a single individual when choosing where to live.

Costa & Kahn (2000) and Compton & Pollak (2007) discuss how assortative mating in the marriage market results in power couples (in which both have at least a four-year college degree). Costa & Kahn argue that such couples are more likely to migrate to large cities, whereas Compton & Pollak, using panel data, find that the concentration of power couples seems to be the result of higher rates of formation of power couples in large cities. At the other end of the life cycle, Chen & Rosenthal (2008) document that both high-skill and low-skill workers tend to migrate to high-amenity places with low housing prices. In the words of Chen & Rosenthal (2008, p. 530), "the composition of cities becomes increasingly tilted towards retirees in locations with improving consumer amenities and low cost of living."

Chen & Rosenthal's model, like ours, is a static model rather than a dynamic intertemporal optimization model. A fully dynamic model would allow workers to work in high-return cities during their work life and move after retirement (and might be modeled along the lines of Kennan & Walker 2011).

4.6. Can the Analysis of Local Labor Markets Shed Light on Causes of Increasing Inequality?

In an insightful paper, Card & DiNardo (2002) point out that although the broad evidence from the labor market is consistent with skill-biased technological change as a key driver of increasing wage inequality, much of the evidence on the matter is indirect and identification is difficult. As Card & DiNardo ask, should we think about the evidence as "technology or tautology?"

Many researchers have taken up the issues raised by Card & DiNardo, attempting to look more carefully at the relationship between technological change and labor market outcomes. The goal of this recent literature is to understand more precisely how technological changes map onto the demand for various skills and to address the concern that relationships can work in both directions—technological advancement affects the relative demand for labor of various types, but the development and adoption of technology also might be influenced by the availability of workers with various skills. Importantly, from our perspective, one of the most promising ways to make headway on these issues is the careful study of local labor markets.

One example of such research is important work by Beaudry et al. (2010), who build on the idea that the adoption of technology is an endogenous process [an idea examined in previous work by, e.g., Beaudry & Green (2005)]. They set up a model in which the adoption of new technology (the personal computer) is most advantageous in local labor markets in which the initial return to workers' skills is low. As an empirical matter, they then show that cities that had high concentrations of high-skill workers in 1980 tended also to be places that differentially adopted new computer-related information technologies thereafter. In this way, they exploit cross-city differences to shed light on how technological developments affected inequality (i.e., they argue that this is evidence of complementarities between high-skill workers and information technology).

A second example is recent work by Autor & Dorn (2013). That paper builds on the seminal work of Autor et al. (2003), who argue that technological developments in communication and information processing have substantially reduced the demand for workers who perform routine tasks that can be codified and performed by machines (e.g., clerical work and production that involves repetitive steps). Many of these tasks previously were performed by workers who might reasonably be called middle-skill workers. But these same technological developments caused a relative increase in complementary abstract tasks that involve intuition and creativity. At the same time, these innovations did little to displace demand for a class of relatively lower-skilled service workers whose tasks were not easily performed by machines (e.g., food service workers, security guards, janitors, and gardeners). The consequence is increased inequality that takes the form of a hollowing out of the wage distribution (or polarization)—a relative increase in the demand for skills of both the high and low ends of the skills distribution, accompanied by a decline in the middle. The distinctive contribution of Autor & Dorn (2013) is an empirical examination of this phenomenon based on an analysis of local labor markets over the 1980–2005 period. They find that "local labor markets that specialized in routine tasks differentially adopted information technology, reallocated low-skill labor into service occupations (employment polarization), experienced earnings growth at the tails of the distribution (wage polarization), and received inflows of skilled labor" [see also Goos et al. (2011), who study European labor markets].

We mention one final example of empirical work that sheds light on the recent evolution of wages and employment using a research design that focuses on local labor markets—analyses that seek to understand the role of increasing import competition. Autor et al. (2014) focus on the spectacular rise in imports from China to the United States, 1990–2007. They make the sensible argument that labor market effects from imports should be felt most heavily in local labor markets for which the industrial mix was initially concentrated in the production of goods for which China has the comparative advantage (e.g., particular labor-intensive industries). Empirical work based on this insight shows that Chinese import competition has recently had substantial labor market effects—lowering local wages and labor market participation and increasing transfer payments in the form of unemployment, disability, and retirement benefits.[21]

4.7. Final Thoughts

We see our article as having two central messages. First, we argue that the empirical analysis of long-run trends in wage inequality should take account of the wide variation in wages and prices that exists in local labor markets. We have set out one simple approach that has the advantage of corresponding closely with the canonical theory used in urban economics. This approach leads to tractable empirical implementation. No doubt there are serious limitations to our stripped-down theory; we nonetheless hope, as mentioned above, that it serves at least as a jumping-off point for more refined thinking about the evolution of inequality in local labor markets.

Second, we hope that our empirical analysis of wage inequality in US cities over the past three decades usefully serves to highlight the substantial and interesting variation that exists across local markets. We find it remarkable, for instance, that wage inequality measured in the usual way—as the difference between the log wages of college graduates and high school graduates—increased by a similar amount in most cities from 1980 to 2010, although in some cities primarily because of the declining wage of high school graduates (Detroit and Pittsburgh) and in other cities primarily because of the increasing wage of college graduates (San Francisco and Boston). We provide examples of recent studies that suggest that this city-specific variation in the evolution of inequality likely matters substantially for evaluating the welfare consequences of increasing labor market inequality and that it is also important for shaping our understanding of the forces that have caused inequality to increase.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

[21]Kovak (2013) builds the theory that underlies this line of work and provides an application to trade liberalization in Brazil.

LITERATURE CITED

Acemoglu D, Autor DH. 2011. Skills, tasks and technologies: implications for employment and earnings. In *Handbook of Labor Economics*, Vol. 4B, ed. O Ashenfelter, D Card, pp. 1043–171. Amsterdam: North-Holland

Albouy D. 2009. The unequal geographic burden of federal taxation. *J. Polit. Econ.* 117:635–67

Andrews RJ, Li J, Lovenheim M. 2012. *Quantile treatment effects of college quality on earnings: evidence from administrative data in Texas*. NBER Work. Pap. 18068

Autor DH, Dorn D. 2013. The growth of low-skill service jobs and the polarization of the US labor market. *Am. Econ. Rev.* 103:1553–97

Autor DH, Dorn D, Hanson GH. 2014. The China syndrome: local labor market effects of import competition in the United States. *Am. Econ. Rev.* 103:2121–68

Autor DH, Levy F, Murnane RJ. 2003. The skill content of recent technological change: an empirical exploration. *Q. J. Econ.* 118:1279–333

Beaudry P, Doms M, Lewis E. 2010. Should the personal computer be considered a technological revolution? Evidence from US metropolitan areas. *J. Polit. Econ.* 118:988–1036

Beaudry P, Green DA. 2005. Changes in U.S. wages, 1976–2000: ongoing skill bias or major technological change? *J. Labor Econ.* 23:609–48

Berry CR, Glaeser EL. 2005. The divergence of human capital levels across cities. *Pap. Reg. Sci.* 84:407–44

Black DA, Gates G, Sanders SG, Taylor LJ. 2002. Why do gay men live in San Francisco? *J. Urban Econ.* 51:54–76

Black DA, Haviland A, Sanders SG, Taylor LJ. 2006. Why do minority men earn less? A study of wage differentials among the highly educated. *Rev. Econ. Stat.* 88:300–13

Black DA, Kolesnikova N, Sanders S, Taylor LJ. 2013a. Are children “normal”? *Rev. Econ. Stat.* 95:21–33

Black DA, Kolesnikova N, Sanders S, Taylor LJ. 2013b. The role of location in evaluating racial wage disparity. *IZA J. Labor Econ.* 2:2

Black DA, Kolesnikova N, Taylor LJ. 2009. Earnings functions when wages and prices vary by location. *J. Labor Econ.* 27:21–47

Black DA, Kolesnikova N, Taylor LJ. 2010. The economic progress of African Americans in urban areas: a tale of 14 cities. *Fed. Reserve Bank St. Louis Rev.* 92(5):353–79

Black DA, Kolesnikova N, Taylor LJ. 2014. Why do so few women work in New York (and so many in Minneapolis)? Labor supply of married women across U.S. cities. *J. Urban Econ.*79:59–71

Black DA, Sanders SG. 2012. Inequality and human capital in Appalachia: 1960–2000. In *Appalachian Legacy: Economic Opportunity After the War on Poverty*, ed. JP Ziliak, pp. 45–80. Washington, DC: Brookings Inst.

Black DA, Sanders SG, Taylor LJ. 2003. Measurement of higher education in the Census and Current Population Survey. *J. Am. Stat. Soc.* 98:545–54

Black DA, Smith JA. 2004. How robust is the evidence on the effects of college quality? Evidence from matching. *J. Econom.* 121:99–124

Black DA, Smith JA. 2006. Estimating the returns to college quality with multiple proxies for quality. *J. Labor Econ.* 24:701–28

Blundell R, Pistaferri L, Saporta-Eksten I. 2012. *Consumption inequality and family labor supply*. NBER Work. Pap. 18445

Bollinger C, Hirsch B. 2006. Match bias from earnings imputations in the Current Population Survey: the case of imperfect matching. *J. Labor Econ.* 24:483–520

Cadena BC, Kovak BK. 2013. *Immigrants equilibrate local labor markets: evidence from the Great Recession*. NBER Work. Pap. 19272

Card D, DiNardo JE. 2002. Skill-biased technological change and rising wage inequality: some problems and puzzles. *J. Labor Econ.* 20:733–83

Card D, Heining J, Kline P. 2013. Workplace heterogeneity and the rise of West German wage inequality. *Q. J. Econ.* 128:967–1015

Chen Y, Rosenthal SS. 2008. Local amenities and life-cycle migration: Do people move for jobs or fun? *J. Urban Econ.* 64:519–37

Compton J, Pollak RA. 2007. Why are power couples increasingly concentrated in large metropolitan areas? *J. Labor Econ.* 25:475–512

Costa DL, Kahn ME. 2000. Power couples: changes in the locational choice of the college educated, 1940–1990. *Q. J. Econ.* 115:1286–314

Dale SB, Krueger AB. 2002. Estimating the payoff to attending a more selective college: an application of selection on observables and unobservables. *Q. J. Econ.* 117:1491–527

DiNardo J, Fortin N, Lemieux T. 1996. Labor market institutions and the distribution of wages, 1973–1992: a semiparametric approach. *Econometrica* 64:1001–44

Goldin C, Katz LF. 2008. *The Race Between Education and Technology*. Cambridge, MA: Harvard Univ. Press

Goos M, Manning A, Salomons A. 2011. *Explaining job polarization in Europe: the roles of technology, offshoring and institutions*. Discuss. Pap., Kathol. Univ., Leuven, Belg.

Haurin DR. 1980. The regional distribution of population, migration, and climate. *Q. J. Econ.* 95:293–308

Hoekstra M. 2009. The effect of attending the flagship state university on earnings: a discontinuity-based approach. *Rev. Econ. Stat.* 91:717–24

Kennan J, Walker JR. 2011. The effect of expected income on individual migration decisions. *Econometrica* 79:211–51

Kovak BK. 2013. Regional effects of trade reform: What is the correct measure of liberalization? *Am. Econ. Rev.* 103:1960–76

Lillard L, Smith JP, Welch F. 1986. What do we really know about wages? The importance of nonreporting and census imputation. *J. Polit. Econ.* 94:489–506

Lindley J, Machin S. 2014. Spatial changes in labour market inequality. *J. Urban Econ.* 79:121–38

Monks J. 2000. The returns to individual and college characteristics: evidence from the National Longitudinal Survey of Youth. *Econ. Educ. Rev.* 19:279–89

Moretti E. 2008. *Real wage inequality*. NBER Work. Pap. 14370

Moretti E. 2011. Local labor markets. In *Handbook of Labor Economics*, Vol. 4B, ed. O Ashenfelter, D Card, pp. 1237–313. Amsterdam: North-Holland

Moretti E. 2013. Real wage inequality. *Am. Econ. J.* 5:65–103

Rauch JE. 1993. Productivity gains from geographic concentration of human capital: evidence from the cities. *J. Urban Econ.* 34:380–400

Ritter JA, Taylor LJ. 2011. Racial disparity in unemployment. *Rev. Econ. Stat.* 93:30–42

Roback J. 1982. Wages, rents, and the quality of life. *J. Polit. Econ.* 90:1257–78

Smith JP, Welch F. 1989. Black economic progress after Myrdal. *J. Econ. Lit.* 27:519–64

Tinbergen J. 1974. Substitution of graduate by other labour. *Kyklos* 27:217–26

Wooldridge JM. 2007. Inverse probability weighted estimation for general missing data problems. *J. Econom.* 141:1281–301