Inequality and City Size*

Nathaniel Baum-Snow, Brown University & NBER
Ronni Pavan, University of Rochester

July, 2009

Abstract

Between 1969 and 2007 a strong monotonic relationship between wage inequality and city size has developed. In this paper, we investigate the causes of the city size inequality premium and its relationship with the growth in overall wage inequality. We find that about one-third of the overall increase in wage inequality in the United States from 1979 to 2007 is explained by city size independent of observable demographic characteristics correlated with city size. While this influence has occurred throughout the wage distribution, city size's effect on inequality in the lower half of the wage distribution is about twice as large as that in the upper half of the wage distribution. Growth in within group inequality has been the most important force driving these city size specific patterns in the data but greater increases in observed skill premia in larger cities have also played an important role. However, we find almost no role for changes in the sorting of population sub-groups with higher wage inequality towards larger cities for generating higher inequality in larger cities. Our evidence on the evolution of wage inequality in cross sections of cities improves our understanding of the mechanisms behind the expansion of inequality in marginal products of labor over time.

*We gratefully acknowledge financial support for this research from National Science Foundation Award SES - 0720763. Cemal Arbatli provided excellent research assistance. The paper has greatly benefited from discussions at the University of Chicago Booth School micro lunch and the North American Regional Science Council meetings. In addition, we thank Satyajit Chatterjee for his helpful comments.

1 Introduction

Juhn, Murphy & Pierce (1993), Card & DiNardo (2002), Lemieux (2006), and Autor, Katz & Kearney (2008) among others have documented a sharp rise in U.S. wage inequality since 1980, especially at the top end of the wage distribution. Among full time full year prime-age male workers, the variance of log hourly wages increased from 0.21 to 0.39 between 1979 and 2007. It is less widely recognized that over this same time period, a strong positive relationship between wage inequality and city size has developed. In the 2004 to 2007 period, the variance of log hourly wages was 0.28 in rural areas and roughly monotonically increasing to 0.53 in the largest three metropolitan areas. This pattern is seen in most measures of inequality and has become steeper over time. In 1979, the variance of log hourly wages for rural areas and the largest metropolitan areas were 0.19 and 0.24 respectively.

Figure 1 presents basic facts about the evolution of hourly wage inequality for our sample of white men in the United States from 1979 to 2007. As has been documented elsewhere, the 1980s saw a dramatic increase in wage dispersion at all points in the distribution and a decline in real wages for most workers. During the 1990s, those in the top quarter of the wage distribution saw wage increases and increasing inequality, with relatively stability throughout most of the rest of the distribution. Between 1999 and 2007, the wage distribution evolved similarly to during the 1980s with increased inequality emerging at all points in the distribution. The only difference is that wages in this later period kept up slightly better with inflation.

In this paper, we investigate the causes of the emergence of the city size inequality premium from 1979 to 2007 and its relationship with the growth in overall wage inequality. We find that about one-third of the overall increase in inequality in the United States from 1979 to 2007 is explained by city size independent of observable demographic characteristics correlated with city size. While this influence has occurred throughout the wage distribution, the portion of the increase in inequality in the lower half of the wage distribution that can be attributed to city size-specific factors is about twice as large as that in the upper half of the wage distribution. Commensurate with Autor, Katz & Kearney's (2008) evidence using national data, we demonstrate that growth in within group inequality has been the most important force driving these city size specific patterns in the data. That is, most of the impact city size has had on the increase in inequality nationwide has come because wages have become more unequal within demographic groups in larger cities than smaller cities. While this could reflect increased ability dispersion within observable groups in larger cities, we think it is more likely to reflect more rapid increases in the return to unobserved ability in larger cities. In addition, consistent with evidence in Glaeser et al. (2008), we demonstrate that greater increases in observed skill premia in larger cities also played an important role in generating the strengthening relationship between inequality and city size. However, we find almost no role for changes in the sorting of population sub-groups with higher wage inequality towards larger cities for generating higher inequality in larger cities. It is for this reason that we suspect increased sorting on unobserved skill is also not a serious issue.

Understanding the emerging relationship between city size and inequality is potentially invaluable for improving our understanding of the contemporaneous rise in inequality nationwide. Indeed, data on factor prices and quantities in cross-sections of cities at recent points in time exhibit many of the same features as factor price and quantity data in the time-series since 1979. In particular, even though larger cities have greater log wage gaps between skill groups, they also have larger quantities of skilled workers relative to unskilled workers. Katz & Murphy (1992) and Bound & Johnson (1992) interpret an analogous pattern in the time series as being driven by skill-biased technical change. Their empirical observations have precipitated an extensive theoretical literature including Acemoglu (1998) and Galor & Moav (2000) that attempts to better rationalize the sources of skill-biased labor demand shifts.

We argue that while cross-sectional patterns in the data are consistent with a consumer equilibrium in a Roback (1982) type spatial equilibrium model, typically assumed production technologies do not generate equilibrium patterns of labor demand across locations that are matched by the data. As such, cross-sectional patterns in inequality across metropolitan areas potentially provide important clues about the form of the production technology needed to justify the temporal changes in inequality widely documented in other research.[1] We argue that in order to capture these patterns in the data with a single production technology, this technology must incorporate an agglomeration force and some source of complementarity with skilled workers.

This paper proceeds as follows. Section 2 discusses in more detail the evolution of the city size inequality premium since 1979. Section 3 explains how we construct the

[1] Card (2009) makes a similar observation in the context of analyzing the extent to which immigration has led to increased inequality. Card also provides valuable empirical evidence that the data do not justify models with more than two labor factors of production.

data. Section 4 reports results of log wage variance decompositions. Section 5 investigates the independent role of city size in generating growth in several additional measures of wage inequality. Section 6 sketches a general equilibrium model of residential location and discusses its implications for mechanisms behind the facts documented in previous sections. Finally, Section 7 concludes.

2 Patterns of Wage Dispersion and City Size

Figures 2, 3 and 4 demonstrate the emergence of the positive relationship between wage inequality and city size over the full distributions of wages and city size since 1979. For the purpose of this paper, we index metropolitan area size to be 0 in rural areas and 1 to 10 to represent deciles of the urban population distribution in year 2000. That is, in 2000 approximately 10 percent of the metropolitan area population nationwide resides in each of our city size categories. For other years, we maintain the same assignment of metropolitan areas to categories based on populations in 2000. We experimented with other similar indexes of metropolitan area size, including using contemporaneous deciles and/or fixed cutoff populations over time, and they all generate very similar results. We prefer our measure because it eliminates the possibility that changes in the relationship between city size and inequality could have been generated by a few metropolitan areas that changed locations in the city size distribution. In addition, our measure provides a clear way to assign metropolitan areas to city size categories in 2005-2007 for which there is no reliable MSA level population data. Incidentally, our measure also generates the steepest relationship between city size and inequality in 1979 of the four measures that we examined.[2]

Figure 2 Panel A shows that while the variance of weekly wages was essentially flat as a function of city size in 1969, differing by 0.02 between rural areas and the largest metropolitan areas, its slope increased in each subsequent decade. In the 2004-2007 period, while the variance of weekly wages was 0.32 in rural areas, it was 0.62 in the largest three metropolitan areas. Figure 2 Panel B shows that a similar pattern emerged for hourly

[2]Others including Ciccone & Hall (1996) use density rather than metropolitan area population as a way of capturing the extent of agglomeration forces. We take no firm stand on which one of these measures is better. Depending on the importance of local transportation and communication costs, each measure can be justified by standard urban theory. We find population to be a more natural empirical measure as it does not require measuring developed area. The correlation between the two measures in our data set is 0.44.

wages. In 1979, the gap in the variance of hourly wages between the largest cities and rural areas was 0.05 whereas by 2004-2007 it had increased to 0.25. Figure 2 Panel C shows that the 90-10 percentile gap evolved similarly between 1979 and 2004-2007.

Figure 3 presents graphs of the 10th, 50th and 90th percentiles of the hourly wage distribution by year and city size. Panel A shows that the 10th percentile of the wage distribution declined during the 1980s and after 1999 throughout the city size distribution. In addition, its profile with respect to city size became less steep over time. This is evidence that over time agglomeration forces have become a less important component of the marginal product of very low skilled workers' labor. Panel B shows that from 1989 to 2007, the level and profile of median wages with respect to city size changed very little. Between 1979 and 1989, though median wages dropped everywhere, they dropped more in smaller cities. Panel C shows that while the 90th percentile of the wage distribution changed little in rural areas and small cities between 1979 and 2007, it grew in large cities, generating a much steeper relationship between the 90th percentile of wages and city size over time.

Figure 4 shows the evolution of hourly wage distribution percentile gaps by city size over time. Its examination allows us to draw some parallels between cross-sectional patterns across cities of different sizes and evolution of the top and bottom portions of the wage distribution nationwide. Figure 4 Panel A shows that the 50-10 percentile gap increased in all types of locations and changed very little in slope with respect to city size during the 1980s. It saw its greatest increase in slope with respect to city size during the 1990s, even though the average level actually changed very little. Both slope and level increased after 1999.

Comparison of Figure 4 Panel A and Figure 1 shows a striking evolution of the relationship between overall inequality at the bottom end of the wage distribution and that as a function of city size. Over the course of the 1980s, the increase in inequality in the bottom part of the wage distribution occurred within cities of all sizes simultaneously. During the 1990s, while inequality changed little in the bottom part of the wage distribution nationwide, the 50th-10th percentile gap became strongly increasing in city size. A decline in wage inequality in rural areas and small cities and an increase in larger cities generated this pattern. Finally, after 1999 both the level and slope of the 50-10 gap with respect to city size increased.

Comparison of Figure 4 Panel B and Figure 1 reveals a different story about the upper part of the wage distribution. It shows that the city size inequality premium in the top part of the wage distribution increased in every period studied after 1979. This is consistent with the evolution of the overall 90-50 gaps seen in Figure 1. In sum, evidence in Figures 1-4 shows that the city size inequality premium is evident in the data and emerged throughout the wage distribution, particularly after 1989.

3 Data

3.1 Demographic and Wage Data

Our primary data source for demographic information and wages is the census public use microdata 5 percent samples from 1980, 1990 and 2000 plus the 2005-2007 American Community Surveys (ACS). We choose these data sets so as to achieve large enough samples within metropolitan areas in order to precisely estimate and decompose wage distributions by metropolitan area size categories. We limit our analysis to white men ages 25-54 who report working at least 40 weeks, 35 usual hours per week and who earn at least 75 percent of the federal minimum wage in each year. The full time full year limitation allows us to measure marginal products of labor for individuals who are less likely to be constrained in their residential locations for family reasons. We use white men only to limit the possibility that changes in discrimination and patterns of labor market attachment for women and non-whites influence our estimates. Our earnings measure is the hourly wage calculated by dividing annual income by weeks times usual hours worked.[3] Annual income from the census is for the previous calendar year while that from the ACS is for the year ending in the (unobserved) survey month. Therefore, we sometimes report ACS wages as being for the period 2004-7. Many other studies that examine trends in inequality in the United States use the Current Population Surveys instead. We found that the CPS does not provide sufficient sample sizes and geographic detail to be the optimal data set for our purposes.

We consistently use year 1999 definition county based metropolitan area (MSA) geography throughout the analysis. Unfortunately, the most disaggregated census micro data geography of County Groups in 1980 and Public Use Microdata Areas in 1990 and 2000 in many cases does not match up to MSA geography. As such, our spatial allocation of individuals reported as living in regions that straddle MSA boundaries is imperfect. We al-

[3] Measurement error is an additional justification for using full time full year workers. Baum-Snow and Neal (2009) demonstrate that there exists significant measurement error in hourly wages for part time and part year workers in the census.

locate those living in straddling county groups or PUMAs to the subregion with the largest population. We assign each metropolitan area to one of ten population size categories and all other areas to the remaining non-metropolitan size category.

An examination of the population distribution across city size categories reveals some shifts over time. The distribution of total U.S. population across size categories in 1980, 1990 and 2000 built using aggregates of county populations reported in the census shows that while the population of rural areas declined from 22 percent in 1980 to 20 percent in 2000, the fraction of the population living in one of the top three metropolitan area size categories increased from 24 percent to 26 percent during this period. While this shift of population toward larger metropolitan areas is interesting, we find no evidence that it can explain any part of the increase in the city size inequality premium.

3.2 Spatial Price Deflator

One of the difficulties in analyzing residential location choice between metropolitan areas is that cost of living differences are an important input into household location choice. Moretti (2009) also argues that differences in cost of living may be an important element for understanding inequality in well-being and how it has changed over time. To handle this issue, we have constructed a deflator that takes into account differences in prices over time across space. This price deflator accounts for the possibility that consumers' consumption bundles may differ across metropolitan areas in response to variation in relative prices.

We denote the exogenously given price of good i in time/location j as p_i^j. We assume that consumers have Cobb-Douglas utility over I goods, meaning expenditure shares for each good are the same in each location. Our price index measures the relative expenditure required across time and locations to hold utility constant given observed price differences. The expenditure function in every location j must thus achieve the same level of utility U_0 as that achieved in an arbitrarily chosen base time period and location 0.

$$\sum_i \alpha_i \ln\left(\frac{\alpha_i E^j}{p_i^t}\right) = ln(U_0) \tag{1}$$

Equating utility in time/locations j and 0, we obtain the ideal index relating prices in time/location j to those in 0, capturing the percent increase in expenditure required to keep an individual at the same utility.

$$INDEX_j = \frac{E^j}{E^0} = \exp \sum_i \alpha_i \ln(p_i^j / p_i^0) = \prod_i (p_i^j / p_i^0)^{\alpha_i} \tag{2}$$

This is the index we use to deflate wages across locations and over time.[4]

Building this index requires price data by time and location for different goods and information on expenditure shares. We obtain prices by location from the American Chamber of Commerce Research Association (ACCRA) data sets from 2000 to 2002. These data report prices in six broad expenditure categories for most metropolitan areas and some rural counties nationwide. When possible, we take data from 2001. For the few regions not sampled in 2001 we take data from either 2000 or 2002. ACCRA reports provide us with price data for 244 metropolitan areas and 179 rural counties.[5] We impute price data for remaining areas as follows. Metropolitan counties are assigned the average prices from other MSAs in the same size category and state when possible. If there are no others of the same size in their state, we impute using data from MSAs of the same size by census division. Price data for rural counties are imputed analogously.

For time series variation in prices, we use regional and metropolitan price index data from the BLS disaggregated into the same six categories used for the ACCRA data. We assign each county to be represented by the most geographically specific index possible in each year. Together, the ACCRA and regional CPI data allow us to calculate the relative price in each expenditure category for location/time period j relative to the base location/time period. The base time/location we define as the average ACCRA location from 2001 but deflated to be index value 100 in 1999.[6]

Rather than take expenditure shares α_i directly from the CPI-U, we build expenditure shares for households including white men working full time using data from the biannual Consumer Expenditure Surveys (CEX) starting in 1982. We build shares directly from the

[4] Albouy (2009) uses a more general methodology to account for cost of living differences across locations. Absent federal income taxes, however, our two methods generate the same expression for adjusted wages. Albouy differentiates a generalized spatial indifference condition similar to (1) generating a deflated log wage in location j absent taxes and nonlabor income of $\ln(w_j) - s_p \ln(p_j)$ where s_p is the share of income spent on local goods and p_j is the price of nontraded goods in location j. This is very similar to the expression we use.

[5] ACCRA reports prices separately for different counties within some large metropolitan areas. In these cases, we allow our price index to differ accordingly within MSA. Otherwise, we assign the ACCRA reported prices to all counties in a given MSA.

[6] We choose 1999 as the base year so that 2000 census data does not have to be temporally deflated.

CEX in order to best capture preferences of those in our sample and because the weights used for the CPI-U sometimes fluctuate significantly from year to year. We found that expenditure shares implied by the CEX are very similar for different education groups and in different city sizes. As such, we prefer to use the sample from the CEX that best matches our full census sample to calculate one set of expenditure weights that we apply to all individuals in our sample.

We recognize that there are several ways to build a price index. Our key goal is to allow for substitutability across expenditure categories for individuals facing different price ratios because of their residential locations. In doing so, we are hampered by the limitations of the ACCRA price index in that it only includes prices for a few items in each expenditure category. This lack of data makes it nearly impossible to account for substitution between narrowly defined varieties. Indeed Broda & Weinstein (2007) argue that the CPI overstates inflation by about 0.8 percent per year because of quality upgrading, a bias that would be reduced if price data from a more complete set of products were used to calculate the index. Nevertheless, we feel that for our purposes the ACCRA data set is a superior source to the BLS's underlying price data used to build the CPI because its wide geographic coverage greatly outmatches the 38 metropolitan areas in which the BLS currently collects price data.[7]

As is evident in Figure A1, there is a distinct price gradient with respect to city size in each year. In 1979 rural areas had prices within 5 percent of the 50 percent of the population living in the smallest metropolitan areas. The next 40 percent of metropolitan area residents faced prices about 10 percent higher than rural areas while the largest 3 metropolitan areas faced prices 27 percent higher. This grouping remained stable over time with prices in each metropolitan area size category growing most between 1979 and 1989. After 1989, residents in the group of smaller metropolitan areas experienced price premia of about 5 percent, residents in the medium to large group experienced an 11 percent price premium and those in the largest group experienced a 32 percent price premium.[8]

[7] Broda & Weinstein (2007) build a temporal price index using a nested CES utility function. Unfortunately, we do not have the requisite data to estimate elasticities of substitution between items within broad commodity groups. Moretti (2009) employs a quantity weighted index to measure trends in wage inequality deflated for cost of living differences across space and time. His index potentially suffers from an assumption that college and high school educated individuals operate in separate housing markets and that these markets can be captured by exclusively examining 2 and 3 bedroom apartments.

[8] The price premium in the largest metropolitan area size category is not entirely driven by New York City. Indeed, in 2000 price premia go in order of MSA size with Chicago at 23 percent, Los Angeles at 32 percent and New York at 42 percent.

3.3 Wage Inequality Accounting for Cost of Living

If firms are mobile and produce tradeable goods, wages not adjusted for cost of living represent marginal products of labor. However, wages adjusted for cost of living measure something about utility, though they do an imperfect job of accounting for potential amenity differences across locations. In Section 6 we sketch a model that clarifies these points.

Table 1 reports distribution measures of wages unadjusted (Panel A) and adjusted (Panel B) for cost of living. The left three columns show residual wage inequality while the right three columns show total wage inequality. Table 1 shows that in 1979, adjusting for cost of living differences across locations leads to almost no change in any listed measure of wage inequality. This result makes sense given that there was only a very weak relationship between inequality and city size, which is strongly correlated with local prices, at this time. By 2007 however, each reported measure of wage inequality once adjusted for cost of living is somewhat smaller than unadjusted wage inequality. This pattern is driven almost exclusively by the increase in dispersion between skill groups since adjusted and unadjusted residual wage inequality are very similar in all years.

Adjusting for cost of living reduces the growth in wage dispersion over time. This adjustment reduces the growth in the variance of wages by 8 percent, the 90-50 percentile gap by 14 percent and the 50-10 percentile gap by 6 percent between 1979 and 2007. However, because the spatial price index is well predicted by observable skill and city size, residual inequality actually increases slightly more quickly once accounting for cost of living. Therefore, even if one is interested in understanding inequality in well-being rather than inequality in marginal products of labor, a more careful investigation of potential mechanisms behind the increase in residual inequality over time is in order. One goal in the remainder of this paper is to present the results of such an investigation.

4 Variance Decompositions

In this section, we propose a wage process to use as a basis for the remainder of our analysis. We use this wage process to examine the extent to which cities have influenced between or within skill-group increases in wage dispersion and to evaluate whether shifts in skill prices by location and skill level or shifts in the distribution of the population across locations could have generated the observed increase in inequality. We recognize that while the classification of "prices" versus "quantities" facilitates a convenient statistical decomposition, it imperfectly describes the underlying economic reasons for increases in inequality.[9] We find this statistical decomposition useful, however, because it reveals that quantities play almost no role in generating the increase of the city size inequality premium, or any increase in inequality at all, since 1979.

As a starting point we adopt the statistical model of the wage process employed in Chay & Lee (2000) and Lemieux (2006).

$$w_{ijt} = \sum_{j} (\mu_{jt} + \varepsilon_{ijt}) \tag{3}$$

In this equation, the log wage of individual i in demographic group j at time t has mean and error components. The μ_{jt} parameters are mean log wages and ε_{ijt} are residuals which for now we assume to be distributed with mean 0 and variance σ_{jt}^2. These two elements imply a decomposition of the variance of log wages into within and between components. We assign individuals to one of 825 demographic cells j using 5 education categories, 15 age categories and 11 city size categories.

Figure 5 examines the evolution of the between cell component of the variance of wages calculated using cell means μ_{jt}. Panel A shows that this component exhibited a positive relationship between variance and city size even in 1979. While its profile exhibited a parallel shift up during the 1980s, it became markedly steeper in each of the two subsequent study periods. In 1979, between variance in the largest MSAs was 0.06 relative to 0.02 in rural areas. By 2007 these numbers had grown to 0.17 and 0.05 respectively.

Using results from regressions of log wages on demographics and city size dummies, Panel B examines the extent to which this emerging slope is generated by city size-specific effects that are independent of age and education. The only decade in which the within-cell gradient of coefficients with respect to city size increased was the 1980s while the city size wage premium remained fairly constant after 1989. These results indicate that the increase in the slope of variance with respect to city size in the 1990s and after 1999 has primarily been due to systematic shifts in the prices or quantities of labor that happen to be more concentrated in certain city size categories. That is, Figure 5 provides evidence that most

[9] Indeed, a simple supply-demand framework indicates that quantities and prices must be determined together in equilibrium. The typical interpretation that changes in prices come from labor demand shifts and changes in quantities come from labor supply shifts is clearly not correct, as is emphasized by Katz and Murphy (1992).

of the increase in the city size wage premium driven by observables can be explained with supply and demand factors that are independent of city size.

Figure 6 depicts the evolution of the within component of the variance over time across the city size distribution. Note that in Panel A the left scale is almost twice as long as in Figure 5, consistent with evidence in Table 1 that the residual component of the variance is almost twice as large as the between component. Table 1 shows that the greater importance of residual variance has persisted over time. While the between variance of wages grew by 0.07 between 1979 and 2007, within cell variance grew by 0.11. As with the between component of the variance, the positive relationship between residual variance and city size increased the most during the 1990s with much smaller changes in the other two study periods.

Results in Figure 6 Panel B indicate that an important portion of the increase in the slope of the residual variance with respect to city size can be explained with city size effects independent of demographic factors. Regressions of residual variance on demographic cells and city size generate lines reported in Panel B that diverge dramatically between 1989 and 1999, though they change little in other decades. This is the same pattern of divergence as seen in overall residual variance by city size reported in Panel A. Therefore, the pattern in Panel B provides evidence that it is something about city size rather than demographic groups disproportionately located in larger cities that drove the increase in the city size residual wage inequality premium. This is in marked contrast to the results in Figure 5 indicating little change over the course of our study period in the portion of the between component of wage variance explained by city size.

To determine whether quantities or prices generate the increases in between and within components of the variance, we employ a standard variance decomposition. We specify Equation (3) as

$$w_{ijt} = \sum_j X_{ijt}\beta_{jt} + \varepsilon_{ijt} \tag{4}$$

where X_{ijt} are dummy variables that equal 1 if individual i at time t is in demographic group j. This leads to the following expression for the variance of the wage

$$V(w_{ijt}) = \sum_j \theta_{jt}\beta_{jt}^2 - \sum_j \sum_i \theta_{jt}\beta_{jt}\theta_{it}\beta_{it} + \sum_j \theta_{jt}\sigma_{jt}^2 \tag{5}$$

where θ_{jt} is the share of individuals in cell j at time t and β_{jt} is the mean log wage in cell j at time t. This expression facilitates determination of whether changes in prices or quantities influence changes in components of the variance of wages. We decompose the growth in variance into the four components given by (6).

$$
\begin{aligned}
V(w_{ijt}) - V(w_{ijs}) \quad = \quad & \sum_j \beta_{jt}^2[\theta_{jt} - \theta_{js}] + \sum_j \sum_i \beta_{it}\beta_{jt}(\theta_{is}\theta_{js} - \theta_{it}\theta_{jt}) \\
& + \sum_j \theta_{js}(\beta_{jt}^2 - \beta_{js}^2) + \sum_j \sum_i \theta_{js}\theta_{is}(\beta_{is}\beta_{js} - \beta_{it}\beta_{jt}) \\
& + \sum_j \sigma_{jt}^2(\theta_{jt} - \theta_{js}) \\
& + \sum_j \theta_{js}(\sigma_{jt}^2 - \sigma_{js}^2) \qquad (6)
\end{aligned}
$$

These components are between quantities, between prices, within quantities, and within prices respectively. We name the final component within prices even though it may also reflect changes in unobserved skill or ability distributions within cells. Other researchers have primarily interpreted this term as the change in the price of unobserved skill which we find reasonable given the lack of evidence that observed skill quantities have driven an appreciable part of the growth in inequality.

Table 2 presents the results of this decomposition. It shows that the 1990s was the only decade in which changes in quantities across demographic cells accounted for more than 10 percent of the increase in variance. Consistent with evidence by Lemieux (2006) and Autor, Katz & Kearney (2008) using different data sets, we find that shifts in quantities across cells accounted for 30 percent of the increase in residual variance and 26 percent of the increase total variance during the 1990s. On the whole, however, changes in prices rather than the composition of the population across locations and demographics was a much more important force generating the increase in wage inequality.

5 Measuring the Role of City Size

In this section, we evaluate the effects of city size independent of observed skill on changes in various measures of wage inequality. To do so, we begin with the nonparametric statistical decomposition of quantity and price components of changes in the wage distribution proposed by DiNardo et al. (1996) and adopted by Autor, Katz & Kearney (2008) for analysis of U.S. data and Dustmann et al. (2009) for analysis of German data. We then generate counterfactual nonparametric wage distributions holding the elements of prices and quantities influenced by city size constant at their 1979 levels. We use these counterfactual distributions to calculate counterfactual inequality measures absent city size effects. Our key result is that while city size has independently generated 24 to 33 percent of the growth in the variance of wages, this influence is more than twice as large for the bottom half of the wage distribution than the top half and is primarily driven by growth in the city size gradient of residual inequality.

In the construction of counterfactual wage distributions absent city size effects, we impose that skill price distributions in rural areas are unaffected by city size effects in each year. As such, we construct these counterfactual price distributions to maintain their 1979 shapes relative to rural distributions within demographic cells but impose changes in spread over time that are unconditional on location. This benchmarking to rural locations is a natural choice. As seen in Figure 2, rural wage inequality consistently increased the least of all location types since 1979. Additionally, the industrial structure of the economy has changed the least in rural areas. Therefore, our counterfactual distributions capture how inequality would have evolved had larger cities' technological gaps with rural areas not expanded beyond their 1980 levels.[10]

Calculation of the following modified Herfindahl index by location reveals that rural areas have experienced the most stable industrial composition since 1979. S_{kt} denotes the share of employment in industry k at time t.

$$H_{st} = \frac{1}{2}\sum_{k=1}^{K}(S_{kt} - S_{ks})^2 \tag{7}$$

This index can be thought of as a monotonic transformation of the distance in K dimensional space between industry compositions at times t and s. It is straightforward to show that the index is bounded between 0 and 1. We calculate the index separately by education group using data in each study year.

Table 3 presents our modified Herfindahl index calculated using 1 digit industries for each of 5 education groups within city size categories. Each panel shows results for a separate education group. Table 3 shows that for every education group, rural areas experienced the smallest change in their industry composition between 1979 and 2007 of

[10] One can imagine a model in which new technologies are available for adoption in all city sizes but certain new technologies are increasingly productive in larger agglomerations, and thus are only adopted there in equilibrium. In this world, our counterfactual distributions capture the impacts on inequality of agglomeration biased technical change.

any size category. This pattern holds in all intervening periods as well except for dropouts in the 1990s and more than college in the 1980s. The pattern looks very similar when the index is instead calculated using 3 digit industries. Given that the industrial composition of the workforce changed the least in rural areas, we argue that it is reasonable to use rural areas as an anchor location in calculation of counterfactual wage distributions.

5.1 Construction of Counterfactual Distributions

This subsection demonstrates how we build counterfactual residual and wage distributions absent city size effects. In each time period t, the distribution of residuals $f_t(\varepsilon)$ is the set of unobserved skill price distributions conditional on demographic group x and city size category s weighted by population shares integrated over x and s. In Equation (8), $g_t(\varepsilon|x,s)$ is the distribution of residuals while $h_t(x,s)$ is the joint distribution of demographics and locations at time t.

$$f_t(\varepsilon) = \int g_t(\varepsilon|x,s)h_t(x,s)dsdx \tag{8}$$

This decomposition allows us to calculate counterfactual wage and residual distributions given any combination of prices and quantities from any pair of years. We use this framework to examine counterfactual wage distributions absent price and quantity changes associated with city sizes. Calculating counterfactual distributions amounts to replacing the components of $f_t(\varepsilon)$ with different values in each year.

Table 4 presents inequality measures calculated from the actual and two counterfactual residual distributions. Columns headed "1980 Quantities" show measures of counterfactual inequality holding demographics by location category at its 1980 distribution. This amounts to reweighting the residuals in other years to 1980 shares. Therefore, the underlying distribution used to calculate constant quantity inequality measures in year t is given by the following expression.

$$f_t^q(\varepsilon) = \int g_t(\varepsilon|x,s)h_{1980}(x,s)dsdx \tag{9}$$

While not the focus of our study, we present these results to serve as an additional benchmark for our city size effect adjusted estimates and to connect our analysis to the existing literature on this topic which largely uses CPS data and does not account for sorting across locations.

Columns marked "1980 City Sizes" remove changes from the price and quantity components of the wage that occurred independently because of city size after 1980. These plots maintain the 1979/1980 distributions of prices and quantities across city sizes within demographic cells but allow the distribution of prices and quantities to change between demographic groups. To calculate the quantity components of these counterfactual distributions, we use the decomposition $h_t(x,s) = h_{at}(s|x)h_{bt}(x)$.

Our goal is to generate constant city size effect counterfactual residual distributions that maintain the 1979 profiles of residuals as functions of city size within each demographic group in other years. To achieve this goal, we hold the relative price distributions fixed across city size categories within demographic cells over time while allowing price distributions by demographic group unconditional on location to change. We center these counterfactual distributions around the percentiles at which residuals equal 0 in the actual $g_t(\varepsilon|x,s)$ distributions. As discussed above, we treat rural areas as the reference location.

The mathematical construction of our counterfactual price distributions is as follows. Because the key adjustment is done by percentile, it is convenient to initially express the counterfactual distribution in terms of its inverse cumulative distribution function $G_t^{c^{-1}}(\phi|x,s)$. We nonparametrically calculate

$$\begin{aligned}\widetilde{G}_t^{c^{-1}}(\phi|x,s) &= G_t^{-1}(\phi|x,0) + [G_{1980}^{-1}(\phi|x,s) - G_{1980}^{-1}(\phi|x,0)] \\ &\quad -[G_{1980}^{-1}(\overline{\phi}_{xst}|x,s) - G_{1980}^{-1}(\overline{\phi}_{xst}|x,0)] \qquad (10)\end{aligned}$$

for each percentile ϕ. We denote the percentile at the mean of the residual distribution for demographic group x in location s at time t as $\overline{\phi}_{xst}$, such that $\overline{\phi}_{xst}$ solves $G_t^{-1}(\overline{\phi}_{xst}|x,s) = 0$.

Equation (10) shows how we spread out the $G_t^{-1}(\phi|x,0)$ distributions for locations other than category 0 in each year. We expand (or contract) these distributions by adding on the growth in the difference between residuals at the given percentile ϕ and the reference percentile for the location s and location 0 distributions in 1980. The second bracketed term accounts for the possibility that the percentile of the mean of each residual distribution is not the same in 1980 as it is in subsequent study years. In order to maintain mean 0 residual distributions within demographic group and city size category, we demean the resulting probability distribution functions $\widetilde{g}_t^c(\varepsilon|x,s) = \frac{d}{d\varepsilon}\widetilde{G}_t^c(\varepsilon|x,s)$ to generate counterfactual residual price distributions $g_t^c(\varepsilon|x,s)$.[11]

[11] Maintaining counterfactual mean 0 residual price distributions is important for our calculation of counterfactual wages described below. The choice of the reference percentile $\overline{\phi}_{xst}$ makes little difference as long

Putting these price and quantity elements together, we have an expression for the counterfactual distribution of the residuals.

$$f_t^c(\varepsilon) = \int g_t^c(\varepsilon|x,s)h_{a1980}(s|x)h_{bt}(x)dsdx \tag{11}$$

This is the formula used to generate the columns marked "1980 City Sizes" in Table 4.

Table 5 presents analogous measures based on wages instead of residuals. To calculate counterfactual wage distributions, we use the fact that the unconditional distribution of log wages at time t is given as follows.

$$a_t(w) = \int [m_t(w|x,s) + g_t(\varepsilon|x,s)]h_t(x,s)dsdx \tag{12}$$

Note that this expression is the same as $f_t(\varepsilon)$ above with the additional intercept term or group mean $m_t(w|x,s)$. The calculation of the quantity constant wage distribution follows analogously to above by replacing $h_t(x,s)$ with $h_{1980}(x,s)$.

The calculation of the city size effect constant distribution additionally requires calculation of a counterfactual group mean wage absent city size effects. To calculate this, we use the same logic as for calculation of the counterfactual distribution of people across cells above absent city size effects. In particular, we calculate $m_t^c(w|x,s)$ to have the same gradient with respect to city size as existed in 1980 within city size category but a new mean for each demographic group x. This leads us to the following expression for counterfactual wages.

$$a_t^c(w) = \int [m_t^c(w|x,s) + g_t^c(\varepsilon|x,s)]h_{a1980}(s|x)h_{bt}(x)dsdx \tag{13}$$

The mean component $m_t^c(w|x,s)$ can be calculated several different ways, depending on whether the distribution of individuals within group x is allowed to change across s or not. In columns entitled "Composite Base" we report counterfactual inequality measures based on distributions constructed assuming that the composition of the population within demographic group evolves as it did in equilibrium across city size categories. For columns entitled "1980 Base" we constrain the distribution of the population across demographic groups to remain as it was in 1980. In all cases, the composite base generates group mean wages that imply smaller effects of city size on inequality than does the 1980 base.

as it is near the center of the relevant price distribution.

5.2 Results

Table 4 presents the evolution of actual residual inequality and two counterfactual inequality measures. Consistent with Figure 1, it shows that the variance, 50th-10th percentile gap and 90th-50th percentile gap all increased over every time interval studied except the 50-10 percentile gap during the 1990s.

The second set of columns, "1980 Quantities", shows the evolution of residual wage inequality holding the distribution of the population across demographic groups and city sizes at its 1980 level and is calculated using Equation (9). Overall, shifts in quantities across demographic and city size cells have hardly accounted for any of the increase in any of the three measures of residual inequality during the 1980s. However, as reported by Lemieux (2006), during the 1990s shifts in quantities appear important in generating about half of the increase in the variance of residual wages.[12] We again see no evidence of quantities' importance after 1999.

Consistent with evidence reported by Autor, Katz & Kearney (2008), we find that this pattern masks countervailing forces at the top and bottom ends of the residual wage distribution. At the bottom end, 1980 quantity adjusted residual inequality actually exceeds actual residual inequality in 1990 and then falls precipitously during the 1990s. However, behavior at the top end of the wage distribution is more consistent with Lemieux's evidence. Taken together, the "1980 Quantity" results in Table 4 show that between 1979 and 2007, demographic and location shifts account for no more than 20 percent of the increase in residual inequality. This result indicates that changes in the sorting equilibrium played a small role in generating increases in the city size inequality premium. Indeed, when city size is not considered in the construction of the constant quantity numbers, the implied distributions are so similar that all reported inequality measures change by less than 0.01.

The third set of columns in Table 4 is constructed using counterfactual residual distributions that take out the effects of city size. These results indicate that an effect of city size on residual inequality independent of observed skill exists in all three time periods examined and throughout the residual wage distribution. However, the role of city size is about twice as large in the bottom half of the distribution as in the top half of the distribution.

[12] This result differs from that for the variance decomposition in the previous section due to differences in base years used. Using data from the CPS, Lemieux (2006) actually claims that quantities account for the full increase in residual wage inequality during the 1990s. Our use of a different data source means that we should not expect to get exactly the same answer.

City size effects generated 36 percent of the variance growth, 29 percent of the growth in the 90-50 percentile gap and 60 percent of the growth in the 50-10 percentile gap between 1979 and 2007. For each inequality measure/year combination, counterfactual measures calculated using only the price component of city size effects are almost identical to those using both price and quantity components. Replacing $h_{a1980}(s|x)$ with $h_{at}(s|x)$ changes the variance, 50-10 and 90-50 percentile gaps in all years by less than 0.01.[13]

Table 5 presents results analogous to those in Table 4 but for wages rather than residuals. The numbers in Table 5 are calculated using wages constructed with the residuals used for Table 4 plus intercepts. Patterns in Table 5 are very similar to those in Table 4. In addition to actual inequality, Table 5 lists two measures of counterfactual inequality. The first, "Composite Base", uses calculations of the intercept portion of wage distributions from contemporaneous spatial distributions of the population. That is, it does not account for changes in the sorting equilibrium after 1980 within demographic groups. The columns entitled "1980 base" assume that the population remains at its 1980 distribution across location types in each skill group.

Results in Table 5 indicate that as with residual inequality, up to one-third of the increase in wage inequality from 1979 to 2007 can be attributed to city size-specific factors. It also shows that the change in the sorting of individuals across locations generated one-third to one-half of this effect. That is, variance of wages would have grown one-third slower absent all city-size specific effects but 9 percentage points of that is because we restrict individuals to remain in their 1980 locations. This restriction only matters for the overall variance between 1999 and 2007. As with residual inequality, our estimates indicate that overall wage inequality at the lower end of the distribution is much more greatly influenced by city size specific factors than that at the upper end of the distribution, with impact on the lower part of the distribution about twice as large as that on the upper part of the distribution.

[13] Figure A2 shows that in each year the fraction college educated has a monotonicly increasing profile that shifts up over time but does not change in slope. This pattern is consistent with our claim that changes in the sorting equilibrium had nothing to do with the emerging relationship between inequality and city size. (The slight level shift down between 1999 and 2007 comes because college graduates had large declines in participation during this period.)

6 Prices and Wages by Location and Skill

We have presented evidence indicating an important role for city size independent of observable demographic factors in generating increased wage inequality since 1979. In this section, we lay out a theoretical framework that is a first step toward understanding potential mechanisms behind this correlation in the data. We first outline the consumer side of a standard spatial equilibrium model in the spirit of Roback (1982) with the addition of worker skill heterogeneity. The central insight of the consumer theory is that an increase in the market price of skill causes a greater increase in log wage gaps between skill groups in more expensive cities than cheaper cities provided that the income elasticity of demand for local goods is greater than one. That is, the conditions for a consumer equilibrium require that the increase in the city size inequality premium is a direct by-product of the increase in the price of skill, provided that local goods are a luxury. We then show that cross-sectional data on wages and local prices by skill and location size fit this predicted pattern. Finally, we discuss potential mechanisms on the production side that can generate observed patterns in the data.

6.1 Consumers

Each worker has utility function U over composite traded consumption goods x of price 1, composite local goods R of price r and a local amenity level A. Each worker of skill level s receives a wage w_s to spend on consumption and local nontradeable goods. In equilibrium, all workers of skill level s must achieve the same utility level V^s no matter where they live. As in the Roback (1982) framework, this indifference condition generates a possible set of equilibrium wages, local prices and amenities across locations for each skill group, with the key constraint that all skill groups in each location experience the same prices and amenities. We express the indifference condition using the indirect utility function as follows.

$$V^s = \max_{x,R} \left[U(x, R, A) + \lambda(w_s - rR - x)\right] \tag{14}$$

Using this expression, it is straightforward to derive the shape of equilibrium relationships between observed wages w_s, local prices r and unobserved amenities A.

Total differentiation of (14) derives the standard equilibrium relationship between these three objects for small differences between locations.[14]

$$dV^s = 0 = \frac{U_A}{\lambda} dA + d(w_s) - R dr \tag{15}$$

Rearranging generates the following expression.

$$d\ln(w_s) - \frac{rR}{w_s} d\ln r = -\frac{U_A}{\lambda w_s} dA \tag{16}$$

This model thus predicts a constant elasticity between wages and rents across locations with the same amenity value for those in a given skill group. Equation (16) also shows that comparison of log wages net of local prices times expenditure share across locations for individuals with the same skill level allows for recovery of locations' relative amenity values. Albouy (2009) uses a similar framework to back out quality of life estimates across metropolitan areas.

We consider the comparative static in which log wages of high skilled workers increase relative to those of low skilled workers in all locations.[15] We focus on analyzing how the equilibrium elasticity between local prices and wages across locations given skill group and amenity level changes with an increase in overall log wage gaps between skill groups. Our goal is not to understand why log wage gaps between skill groups have increased over time, a topic of considerable existing research, but instead to derive conditions under which these widening gaps are accompanied by a stronger relationship between inequality and local prices. Analysis of this elasticity is informative about the relationship between log wage inequality and city size because local prices faced by residents of a given location are the same regardless of skill and local prices are monotonically increasing in city size. Therefore, examination of the variation in this elasticity across skill groups allows us to describe log wage gaps between low and high skilled workers as a function of city size and how they change with the skill price. This analysis foreshadows empirical evidence depicted graphically in log wage-log rent space in the following subsection.

The model predicts that log wage gaps are increasing in local prices provided that the income elasticity of demand for nontraded goods is greater than 1. For similar reasons described below, an increase in log wage gaps in every location, which could come about

[14] This condition ignores higher order effects and thus only holds in general for small differences. Nevertheless, it has been used extensively.

[15] In the data presented in the following subsection, we see that while college educated workers' wages have increased, those of high school educated workers have actually decreased in real terms over time.

because of nationwide shifts in labor demand for example, requires that these wage gaps grow more in higher price locations than in lower price locations.

These two implications follow directly from the sign of the partial derivative of the equilibrium elasticity between local prices and wages with respect to the wage, where amenities and local prices are held fixed. Denoting the Marshallian demand for nontradeables as $\widehat{R}$, we have

$$\frac{\partial}{\partial w_s}\left(\frac{d\ln r}{d\ln(w_s)}|_{A=\overline{A}}\right)=\frac{1}{Rr}[1-\frac{d\ln\widehat{R}(\cdot)}{d\ln(w_s)}]. \tag{17}$$

If the income elasticity of demand for nontradeables is greater than 1, the elasticity of rents with respect to wages across locations for a given skill group is predicted to decrease over time for each skill group. That is, at each location the slope of the indifference relationship between log rents and log wages is smaller for those with higher wages. This directly indicates that wage gaps must be greater in higher price locations in the cross-section. Similarly, since an increase in the wage for the high skill group at all locations implies a reduction in this elasticity of interest, a stronger relationship between increasing overall wage dispersion and local prices results.

The key to understanding (17) is to consider how much more nontradeables high skilled people facing a given price consume relative to low-skilled individuals. If both high and low skilled individuals consume the same amount of nontradeables (or the income elasticity of demand is 0), both groups require the same dollar increase in wages to maintain their utility level when faced with a price increase. This amounts to a smaller percentage increase for the high skilled group. If the income elasticity is greater than 1, however, the high skilled group requires a greater percentage increase in the wage to compensate them for the price increase than does the low skilled group.[16]

This brief treatment shows that standard consumer theory generates four lessons for empirical work using spatially disaggregtaed data on prices and wages by skill level. First, we can determine locations' relative amenity values by examining profiles of wages adjusted for cost of living differences. Second, we should see a constant elasticity relationship between local prices and wages for each skill group. Third, as a skill price increases, the

[16]In all cases the level wage gap is weakly increasing in rents conditional on amenities. The equilibrium slope between levels of rents and wages is $\frac{1}{R}$ for a given skill group. If nontradeables R is a normal good, the demand for R is increasing in skill. Therefore, this equilibrium slope is less steep for more highly skilled individuals. The logic is that since higher skilled individuals consume more R, they require a larger wage increase to compensate them for a rent increase than do lower skilled individuals. Therefore, in the context of the standard Roback (1982) setup presented in levels, as reviewed in Moretti (2004), indifference curves for different skill groups are convex and diverging upwards conditional on the amenity level.

elasticity of local prices with respect to wages for that skill group declines if the income elasticity of demand for nontradeables is greater than unity. Finally, the relationship between log local prices and log wages diverges as a function of skill if nontradeables are a luxury, conditional on amenity value. This divergence is increasing in the relative skill price.

6.2 Empirical Observations

In this subsection, we first use lessons from the previous subsection to determine the relative amenity values of city size groups. We then examine patterns in log local prices and log wages for high school and college graduates over time by city size for locations with the same amenity value.

To calculate relative amenity values, we regress log wages adjusted for cost of living on city size dummy variables and interacted age/education dummy variables. If these controls successfully capture skill heterogeneity across locations, (16) indicates that the coefficients on city size reflect location types' average relative amenity values. We graph these coefficients in Figure 7 Panel A. They show a distinct inverse U shape in each year indicating that rural areas and very large metropolitan areas have the most valuable amenities. Of particular interest is that this profile of real wages with respect to city size has hardly changed over time. We can think of these results as representing something about the relative labor supply conditions to cities of different sizes.

Consistent with the literature on local labor markets (see Card (2009) for a recent review), we break out the sample for high school and college educated individuals separately in Figure 7 Panels B and C. These two graphs show that while the profile of wages adjusted for cost of living is similar for both groups, it is shifted down for the high school graduates. That is, if our controls capture differences across individuals in potential wages, high school graduates' amenity value is greatest for the largest cities while college graduates derive the greatest amenity value from small and large locations. These figures also show that location categories 1 through 9 all have similar amenity values in each year for both groups.

Now that we have a sense of the relative amenity values of locations of different sizes, we can examine log wage gaps in cities of different sizes holding the amenity value constant. Figure 8 depicts log local prices and log wages for high school and college graduates plotted with numbers indicating city size category. Each panel presents these relationships for a different point in time. Underlined numbers indicate observations for college graduates.

For a given amenity level, each of these plotted points should fall on the same indifference curve given by (14). We index average prices across all skill groups in rural locations in each year to 0. Predicted indifference curves for location types 1-9 generated by simple linear regressions are plotted as dashed lines for high school graduates and solid lines for high school graduates. While not perfect, straight lines, or constant elasticities, fit the data quite well.[17]

In each of the decades between 1980 and 2000, the slope of the college indifference curve became less steep while it changed little for high school graduates. This matches the theoretical case in which the income elasticity of demand for nontradeables is greater than 1 and the college skill price was rising. In the most recent study period, the elasticity for high school graduates became markedly steeper while that for college graduates changed little. This is consistent with a declining high school skill price. Also consistent with a high income elasticity is that log wage gaps between high school and college graduates are increasing in local prices and city sizes. Furthermore, these wage gaps increased more over time in large (high price) locations than in small (low price) locations.

In summary, Figure 8 shows patterns that can be understood in the context of standard consumer theory provided that the income elasticity of demand for nontraded goods is greater than 1. One may wonder whether such a high value for this parameter is reasonable. One explanation is that the composition of nontradeable goods may vary by location size. Larger cities have a greater variety of nontraded goods, (restaurants, sports teams, etc.) and the demand for variety is partly what drives the high implied income elasticity of demand for these goods. An alternative explanation for the patterns in Figure 8 is that the distribution of unobserved skill within observed skill groups differs by location, a standard ability bias story. However, in Baum-Snow & Pavan (2009) we find scant evidence of sorting across location size on unobserved skill.

One additional consideration merits attention. We have treated as exogenous the strong positive empirical relationship between city size and local prices. This relationship has been justified theoretically in models of the internal structure of cities going back to Alonso (1964), Mills (1969) and Muth (1969). As cities grow in population, the increased competition for central locations generates an increase in the equilibrium price of space. In addition, lengthier travel times and greater congestion increases shipping costs and prices for many consumer goods. Our price index is designed to capture these cost of living

[17] Because location types 0 and 10 have greater amenity values from the rest, they fall above the plotted lines.

differences associated with larger cities. The exact mechanisms linking higher prices and populations are not important for our purposes as long as we can take as given that there is a structural relationship between these two objects.

6.3 Producers

The producer equilibrium is more difficult to rationalize than the consumer equilibrium as doing so requires a model generating both higher factor and wage ratios between high and low skilled workers in larger cities relative to smaller cities. Further, there is a greater number of workers of both types in larger cities indicating a role for agglomeration economies. This pattern is similar to that documented extensively in national time series data after 1980.

Given the considerable evidence in the time series of similar patterns, it is useful to consider whether mechanisms successfully proposed to rationalize time series data can also be used to understand the cross-sectional data. The over-arching observation, widely attributed to Katz & Murphy (1992), is that because the increase in wage inequality across skill groups coincided with an increase in relative quantities of college workers, at least part of this inequality increase must be attributable to relative shifts in labor demand toward the more highly skilled. However, in the cross-sectional context one would have to interpret this observation as heterogeneity in production technologies across cities of different sizes. This interpretation is not entirely satisfying as it assumes away the answer to the interesting follow-up question: Why are technologies different in cities of different sizes and how does this difference manifest itself in greater inequality in larger cities? To answer this question, we need a unified production technology that can generate the patterns observed in the data.

While there is no study to date that proposes such a technology, the study that comes closest is Krusell et al. (2000). These authors argue trends in aggregate factor ratios and input costs can be rationalized with a 3-factor model that incorporates capital-skill complementarity. The implication of their model is that as the price of capital fell with the introduction of computers, the demand for skill increased because of this complementarity in production. Given extensive empirical evidence by Ciccone & Peri (2005) and Card (2009) that low and high skilled labor are substitutes, this explanation for relative factor and input price ratios is more satisfying than a skill complementarity explanation. Therefore, we argue that a model with both capital-skill complementarity and agglomeration economies could capture both cross-sectional and time-series patterns in the data. We leave it to future research to more carefully develop and estimate such a model.[18]

7 Conclusions

In this paper, we demonstrate that cities have played an important role in the rise of inequality over time. In 1979, there was only a weak positive relationship between inequality and city size while by 2007 a much stronger relationship between these two variables had developed. We demonstrate that city size specific factors can explain about one-third of the overall increase of the variance in wages between 1979 to 2007 independent of other observable factors. City size is about twice as important for generating increased inequality in the bottom half of the wage distribution as in the top half.

The most important factor generating the city size specific component of inequality growth is that demographic groups disproportionately located in larger cities experienced larger increases in their wage dispersion in larger cities than smaller cities. That is, city size has become more complementary with within observed skill group wage dispersion. It is also true that the skill premium has grown more in larger cities than smaller cities and rural areas. However, the emergence of the city size inequality premium is not due to systematic migration of demographic groups with greater wage dispersion toward larger cities. We demonstrate that these results are what standard consumer theory predicts in response to a divergence in the prices of high and low skilled labor provided that the income elasticity of demand for nontraded goods is greater than 1.

We hope that our analysis sparks further research examining reasons for changes in the structure of labor demand using metropolitan area level data. While the empirical patterns that we document are consistent with equilibria generated by standard consumer theory, it is less clear how and why firm behavior generates these patterns. The considerable cross-sectional variation available in metropolitan area data facilitates investigation of determinants of changes in labor demand that have contributed to growing wage and income inequality. We believe that patterns in the cross-section in addition to the time-series provide useful clues that can help to further unlock the mysteries of the determinants of changes in the structure of labor demand.

[18] A production technology with only labor factors but incorporating a skill-biased agglomeration force as in Dalmazzo & De Blasio (2007) would also match the cross-sectional patterns in the data. However, it is less clear how one would use such a technology to understand trends over time.

References

Acemoglu, Daron. 1998. "Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality." Quarterly Journal of Economics, 108, 1055-1089

Albouy, David. 2009. "Are Big Cities Really Bad Places to Live? Improving Quality of Life Estimates Across Cities." manuscript

Alonso, William. 1964. Location and Land Use. Harvard University Press, Cambridge.

Autor, David H., Lawrence F. Katz & Melissa S. Kearney. 2008. "Trends in U.S. Wage Inequality: Revising the Revisionists." Review of Economics and Statistics, 30:2, 300-323

Autor, David, Katz, Lawrence and Alan B. Kreuger. 1998. "Have Computers Changed the Labor Market?" Quarterly Journal of Economics, 113:4, 1169-1213

Baum-Snow, Nathaniel and Derek Neal. 2009. "Mismeasurement of Usual Hours Worked in the Census and ACS" Economics Letters, 102:1, 39-41

Baum-Snow, Nathaniel and Ronni Pavan. 2009. "Understanding the City Size Wage Gap" manuscript

Bound, John and George Johnson. 1992. "Changes in the Structure of Wages in the 1980s: An Evaluation of Alternative Explanations" American Economic Review, 82:3, 371-392

Broda, Christian and David Weinstein. 2007. "Product Creation and Destruction: Evidence and Price Implications" NBER Working Paper #13041

Card, David. 2009. "Immigration and Inequality" NBER Working Paper #14863

Card, David and John E. DiNardo. 2002. "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles" Journal of Labor Economics, 20:4 733-783

Kenneth Chay and David S. Lee. 2000. "Changes in Relative Wages in the 1980s: Returns to Observed and Unobserved Skills and Black-White Wage Differentials" Journal of Econometrics, 99:1, 1-38

Ciccone, Antonio & Robert Hall. 1996. "Productivity and the Density of Economic Activity" American Economic Review, 86:1, 54-70

Ciccone, Antonio & Giovanni Peri. 2005. "Long-Run Substitutability Between More and Less Educated Workers: Evidence from U.S. States 1950-1990" The Review of Economics and Statistics, 87:4, 652-663

Costa, Dora & Matthew Kahn. 2000. "Power Couples: Changes in the Locational Choice of the College Educated: 1940-1990" Quarterly Journal of Economics

Dalmazzo, Alberto and Guido de Blasio. 2007. "Skill-Biased Agglomeration Effects and Amenities: Theory With an Application to Italian Cities" manuscript

DiNardo, John, Nicole Fortin and Thomas Lemieux. 1996. "Labor Market Institutions, and the Distribution of Wages 1973-1992: A Semiparametric Approach," Econometrica, 64, 1001-1044

Duranton, Gilles and Diego Puga. 2004 "Micro-Foundations of Urban Agglomeration Economies" chapter in Handbook of Regional and Urban Economics, Vol. 4, Henderson & Thisse eds, Elsevier, North Holland

Dustmann, Christian, Johannes Ludsteck & Uta Schoenberg. 2009. "Revisiting the German Wage Structure," Quarterly Journal of Economics, 124:2, 843-881

Galor, Oded and Omer Moav. 2000. "Ability-Biased Technological Transition, Ability Bias and Economic Growth," Quarterly Journal of Economics, 115:2, 469-497

Glaeser, Edward L. and David C. Mare. 2001. "Cities and Skills" Journal of Labor Economics, 19:2 316-342

Glaeser, Edward L., Matt Resseger and Kristina Tobio. 2008. "Urban Inequality", NBER Working Paper #14419

Gyourko, Joseph & Joseph Tracy. 1989. "The Importance of Local Fiscal Conditions in Analyzing Local Labor Markets." Journal of Political Economy, 97:5, 1208-1231

Juhn, Chinhui, Kevin Murphy & Brooks Pierce. 1993. "Wage Inequality and the Rise in Returns to Skill" Journal of Political Economy, 101:3, 410-442

Katz, Lawrence and Kevin Murphy. 1992. "Changes in Relative Wages 1963-1987: Supply and Demand Factors" Quarterly Journal of Economics, 107:1 35-78

Krusell, Per, Ohanian, Lee E., Rios-Rull, Jose-Victor and Giovanni L. Violante. 2000. "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis." Econometrica, 68:5 1029-1053

Lemieux, Thomas. 2006. "Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?" American Economic Review, 96:3 461-498

Machado, Jose and Jose Mata. 2005. "Counterfactual Decomposition of Changes in Wage Distributions Using Quantile Regression," Journal of Applied Econometrics, 20:4, 445-465

Manning, Alan. 2008. "The Plant Size-Place Effect: Agglomeration and Monopsony in Labour Markets" manuscript

Mills, Edwin. 1967. "An Aggregative Model of Resource Allocation in a Metropolitan Area," American Economic Review, 57:2, 197-210

Moretti, Enrico. 2004. "Human Capital Externalities in Cities" chapter in Handbook of Regional and Urban Economics, Vol. 4, Henderson & Thisse eds., Elsevier, North Holland

Moretti, Enrico. 2009. "Real Wage Inequality." manuscript

Muth, Richard. 1969. "Cities and Housing," University of Chicago Press, Chicago

Roback, Jennifer. 1982. "Wages, Rents and Quality of Life" Journal of Political Economy, 90:6 1257-1278

Figure 1: Log Hourly Wage Growth by Percentile 1979-2007

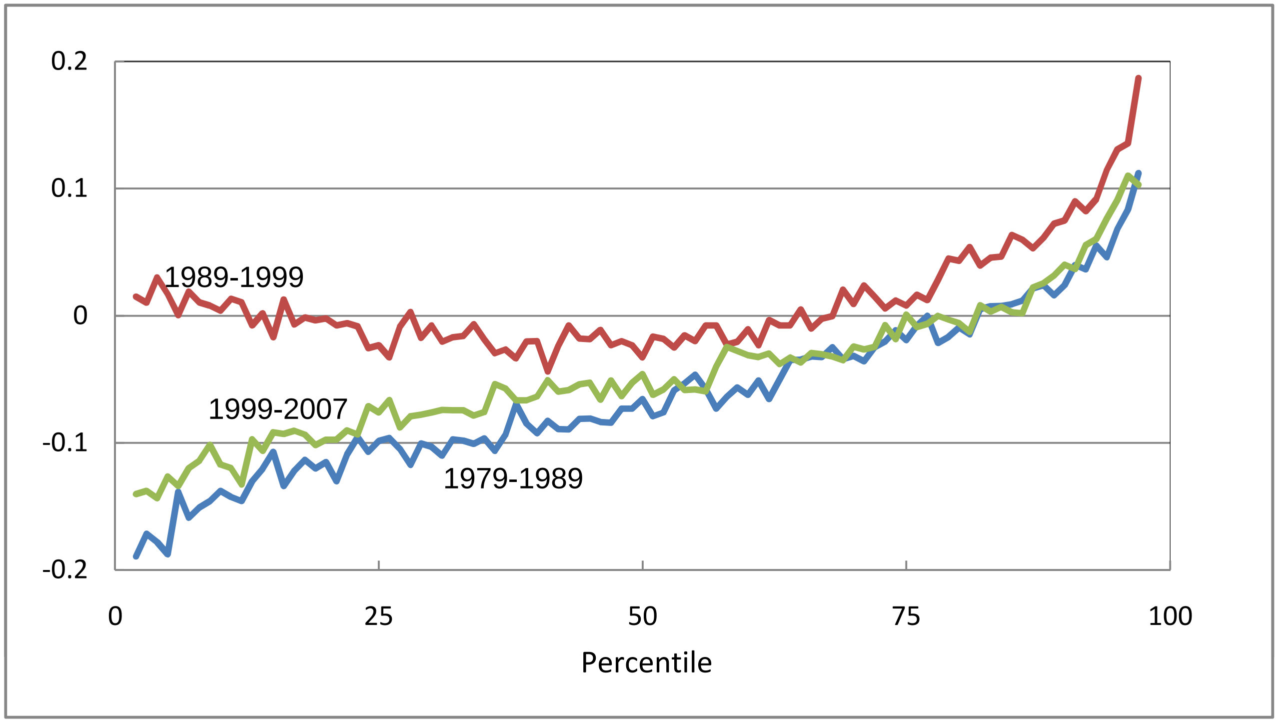

Notes: The sample includes all full-time white male workers ages 25-54 working at least 40 weeks in the listed years. Data is from the census 5% PUMS in 1980, 1990 and 2000 and 1% American Community Surveys (ACS) in 2005, 2006 and 2007. Hourly wages are deflated by the CPI-U and calculated as the logarithm of wage & salary income divided by the product of weeks worked and usual hours worked per week. Observations with imputed demographics, labor supply or wages, the self-employed and those who earned less than 75% of the federal minimum wage in the earnings year are excluded from the sample. Calculations are weighted by sampling weights except for those using the 1980 census which is an unweighted sample. Data listed as being for 2007 actually represents average wages from full years ending in 2005, 2006 or 2007.

Figure 2: Wage Inequality by City Size

Panel A: Variance of Weekly Wages

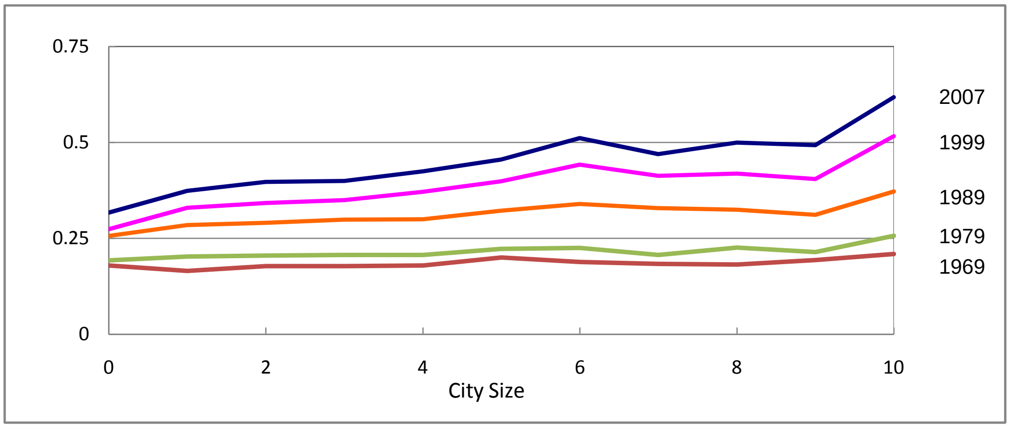

Panel B: Variance of Hourly Wages

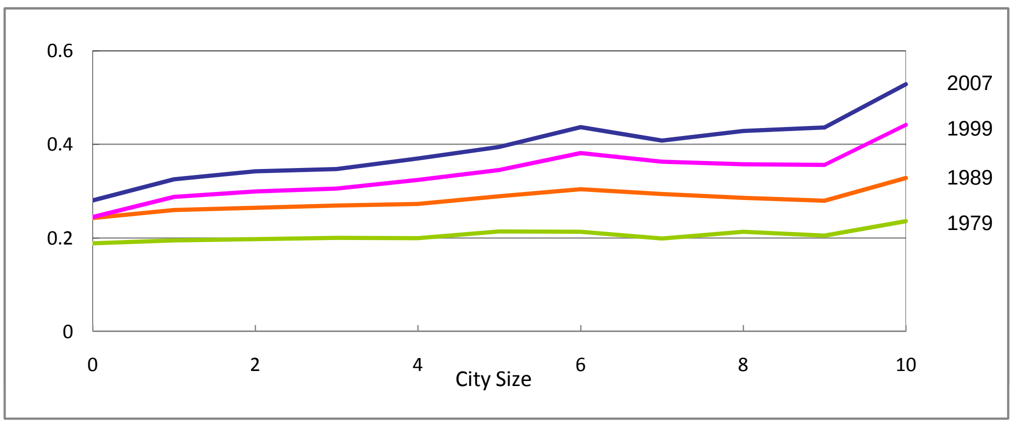

Panel C: 90th-10th Percentile Gap in Hourly Wages

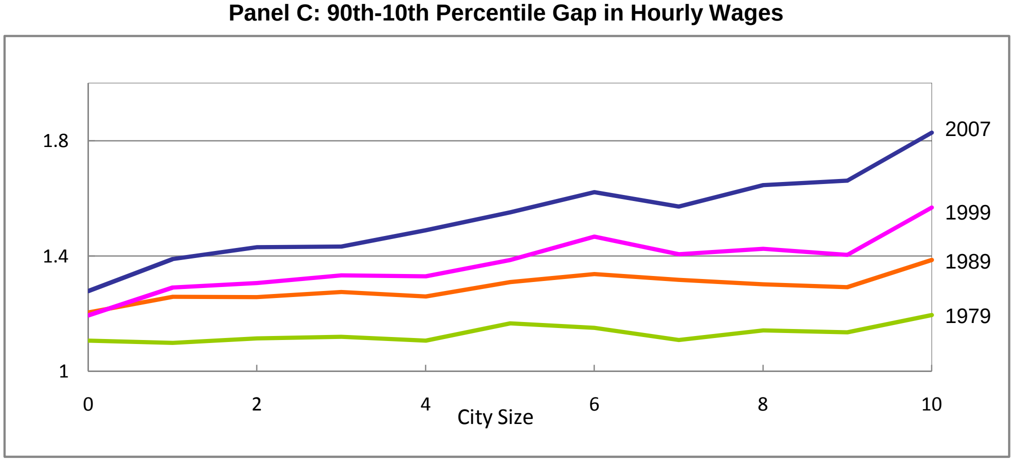

Notes: See the notes to Figure 1 for a description of the sample used for Panels B and C. Because usual hours worked is not observed in the 1970 census, the sample for Panel A additionally includes part time workers. City size categories are based on 2000 metro area populations. Size 0 corresponds to non-MSA locations. Sizes 1-10 correspond to 10-percentile bins from the year 2000 MSA population size distribution.

Figure 3: Percentiles of the Hourly Wage Distribution by City Size and Year

Panel A: 10th Percentile

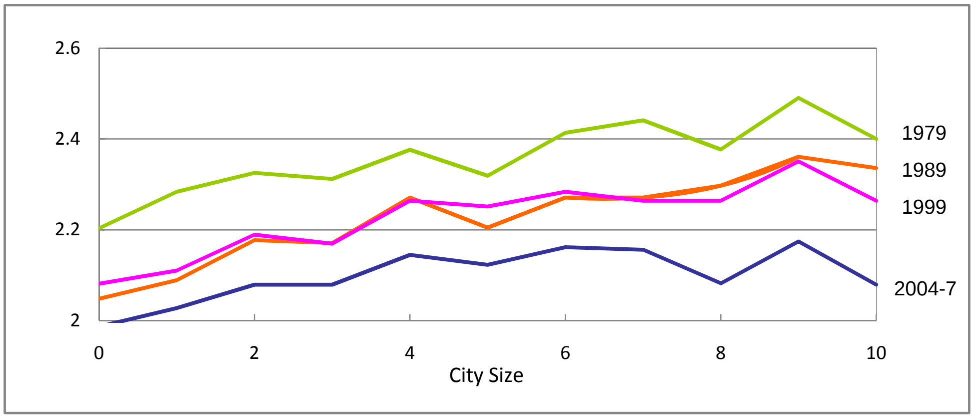

Panel B: 50th Percentile

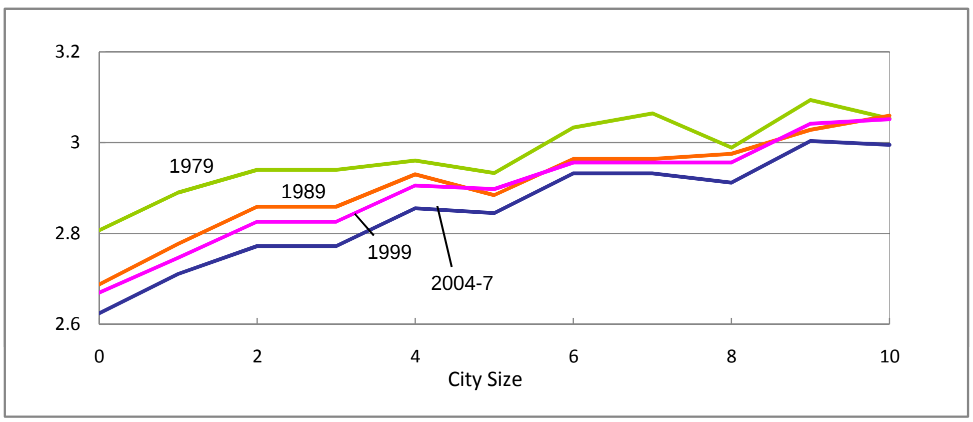

Panel C: 90th Percentile

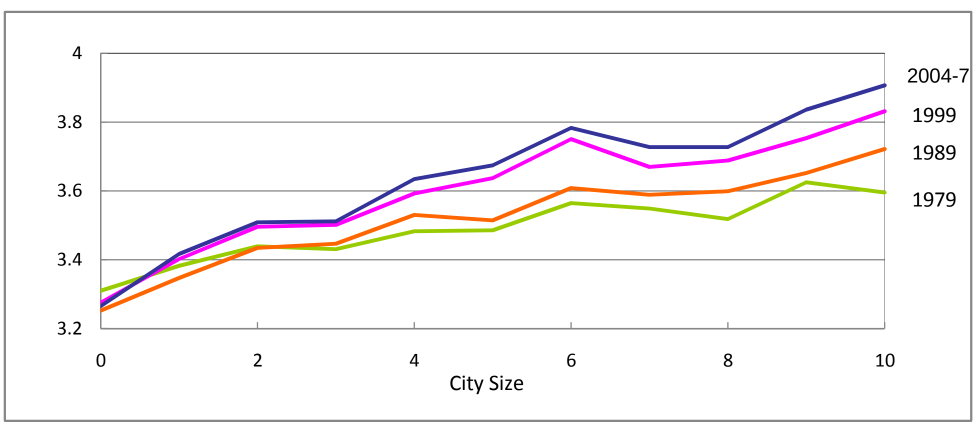

Notes: See the notes to Figures 1 and 2 for explanations of the sample and construction of city size categories.

Figure 4: Hourly Wage Percentile Gaps by City Size and Year

Panel A: 50th-10th Percentile Gap

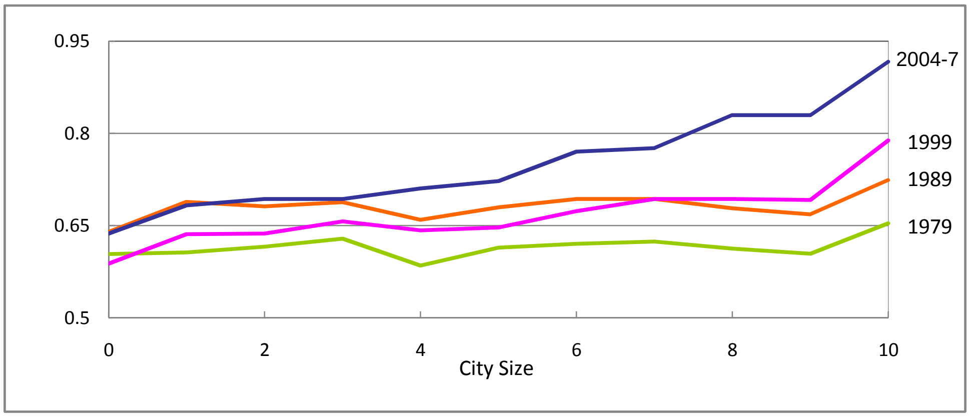

Panel B: 90th-50th Percentile Gap

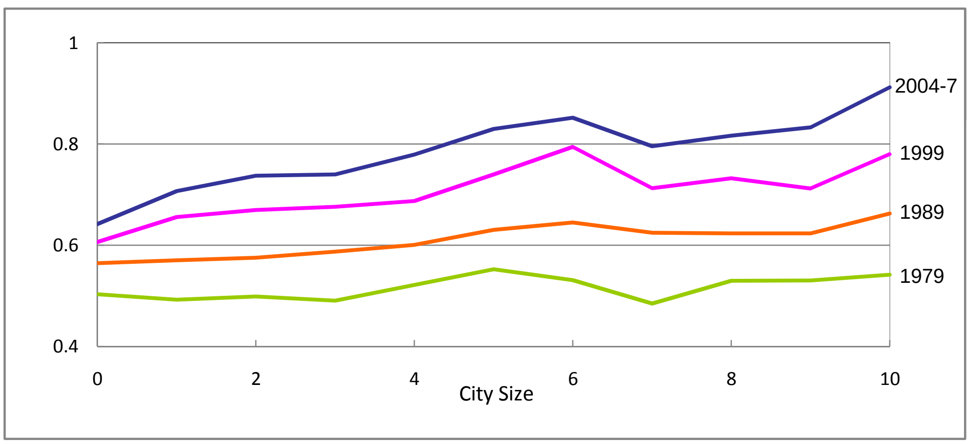

Notes: See the notes to Figures 1 and 2 for explanations of the sample and construction of city size categories.

Figure 5: Between Cell Variance of Hourly Wages by City Size and Year

Panel A: Between Variance of Hourly Wages

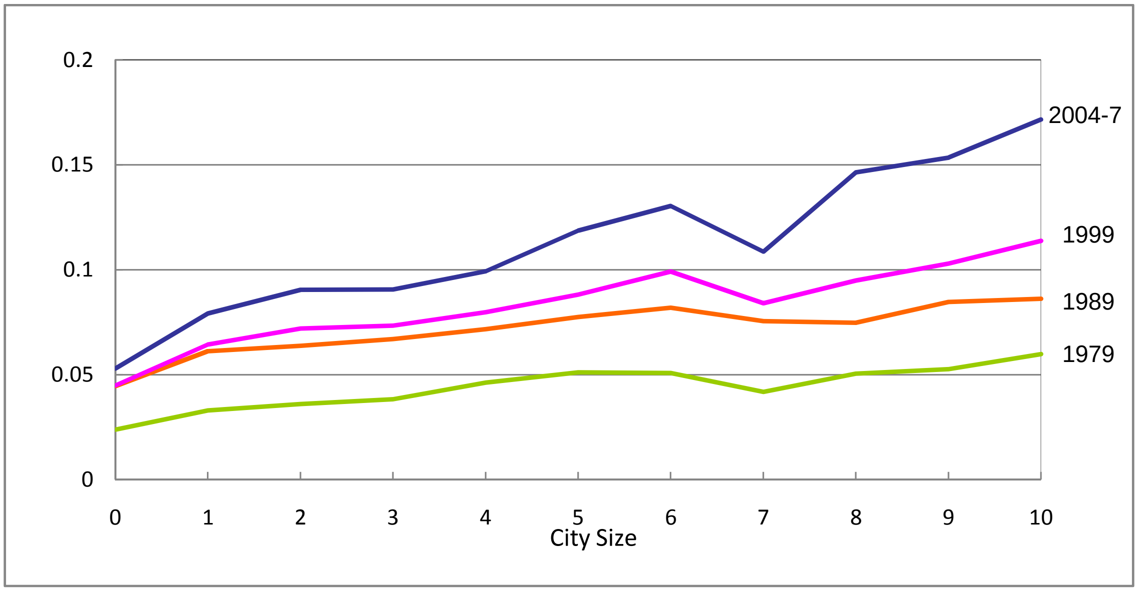

Panel B: Independent Effect of City Size on Between Variance

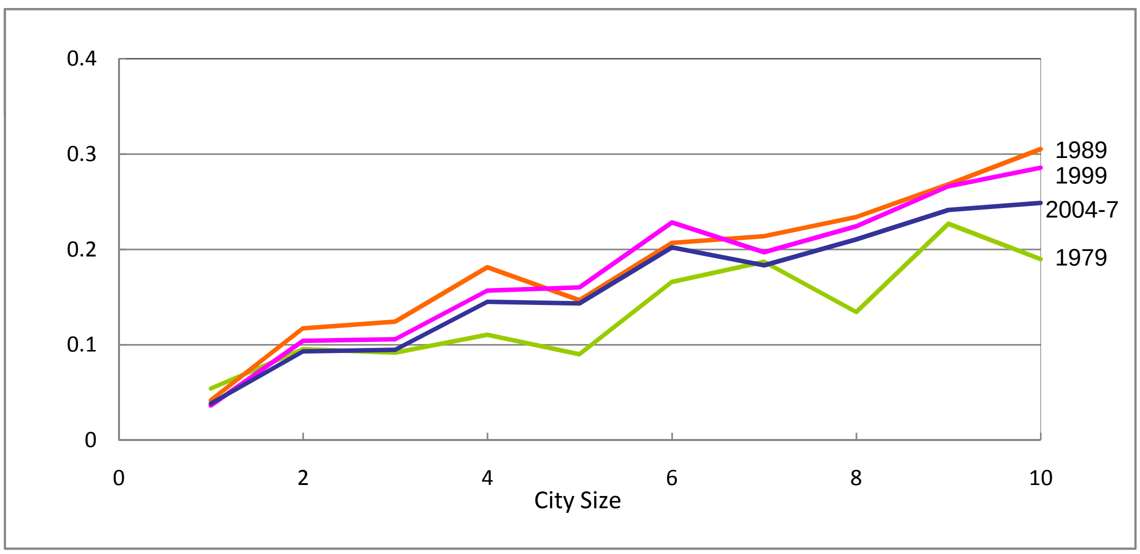

Notes: Panel A shows variances of age/education cell mean log-wages within city size categories. Panel B plots the coefficients on city size dummies from regressions of cell means on age/education cells and city size category indicator variables weighted by cell size. That is, Panel B shows city size premia for the between component of the variance of log wages in each year.

Figure 6: Within Cell Variance of Hourly Wages by City Size and Year

Panel A : Within Variance in Hourly Wages

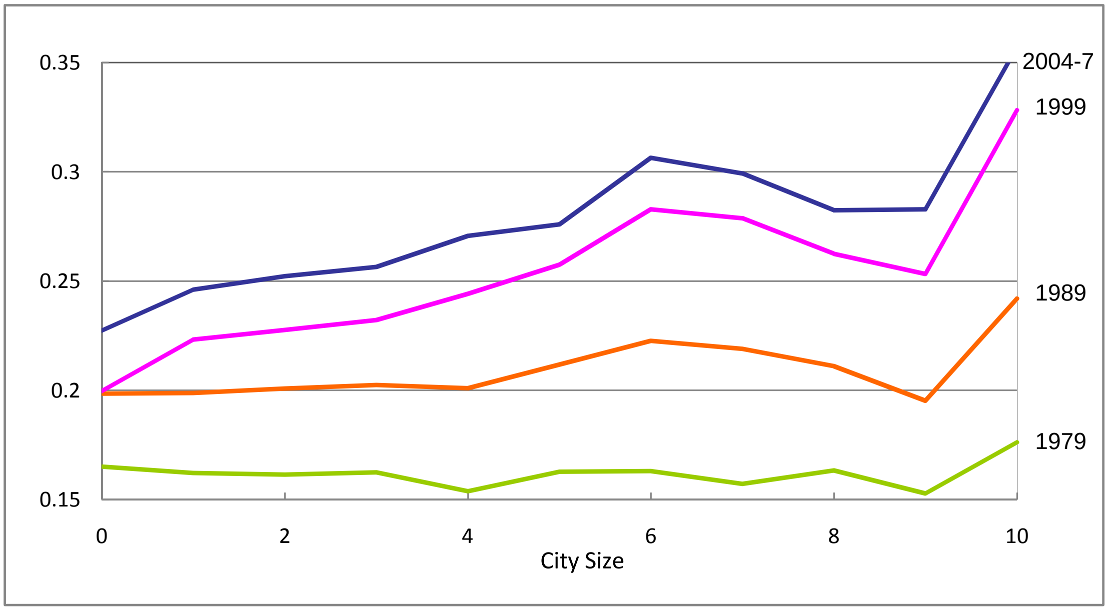

Panel B: Independent Effect of City Size on Within Cell Variances

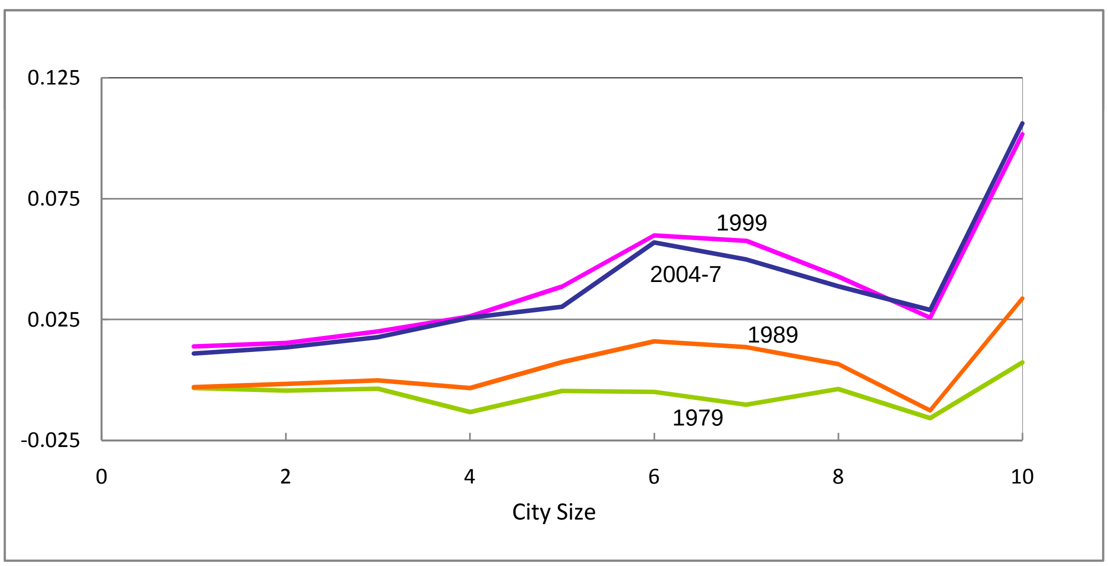

Notes: Panel A shows the variance of residual log-wages in age and education cells by city size category. Panel B plots the coefficients on city size dummies from regressions of these variances on age/education cells and city size category indicator variables weighted by cell size. That is, Panel B shows the average relationship between city size and within group variance.

Figure 7: Cost of Living Adjusted Wages by City Size

Panel A: Full Sample

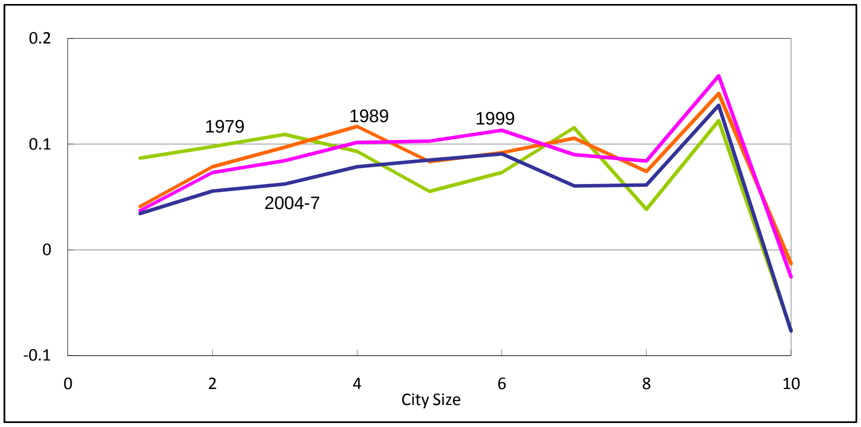

Panel B: High School Graduates Only

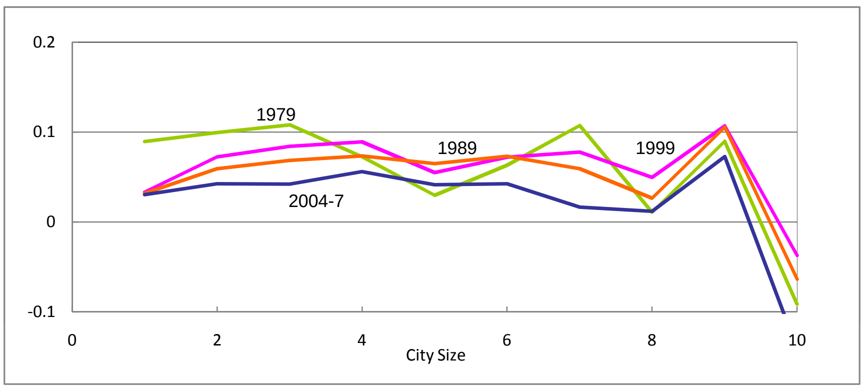

Panel C: College Graduates Only

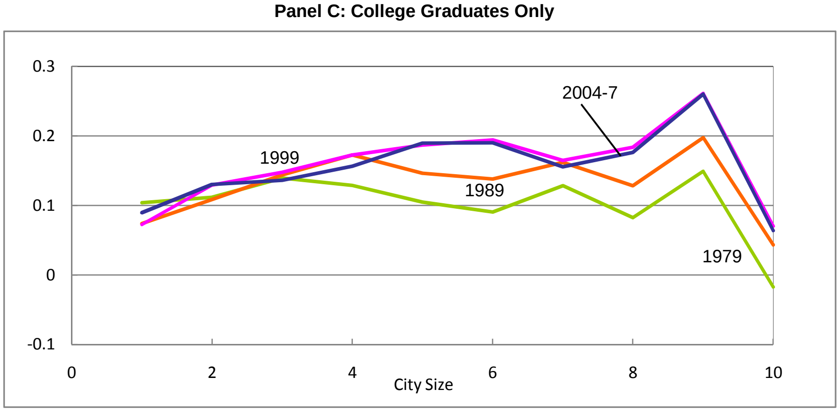

Notes: Each graph shows estimated coefficients on city sizes from regressions of wages adjusted for cost of living on interacted age/education dummy variables and city size dummy variables. Panel B restricts the sample to high school graduates only and Panel C restricts the sample to college graduates only. Locations with higher adjusted log wages have lower amenity values.

Figure 8: Log Prices and Log Hourly Wages by City Size and Education, 1979-2007

Panel A: 1979

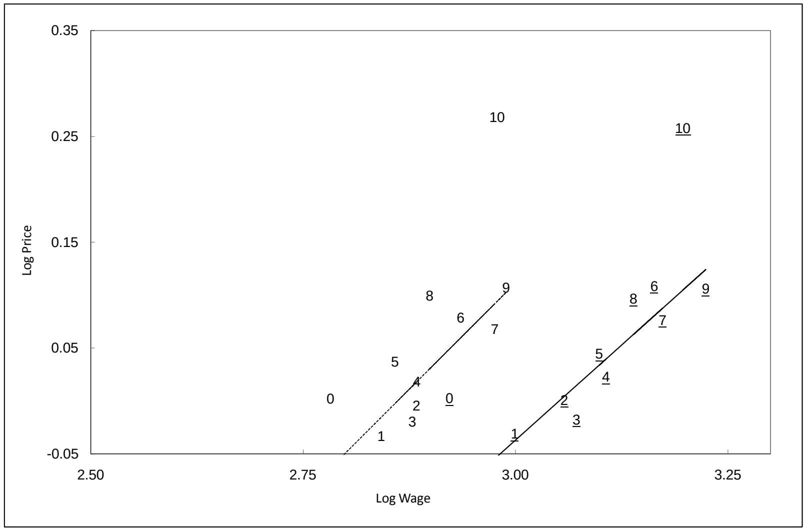

Panel B: 1989

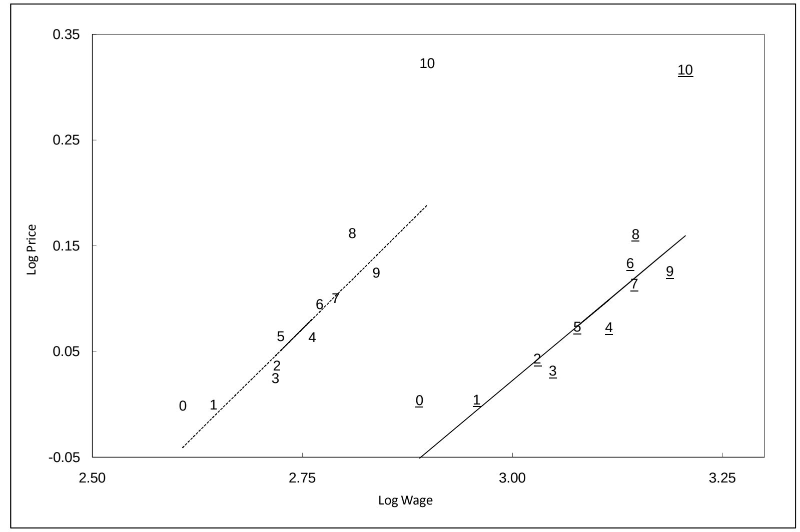

Panel C: 1999

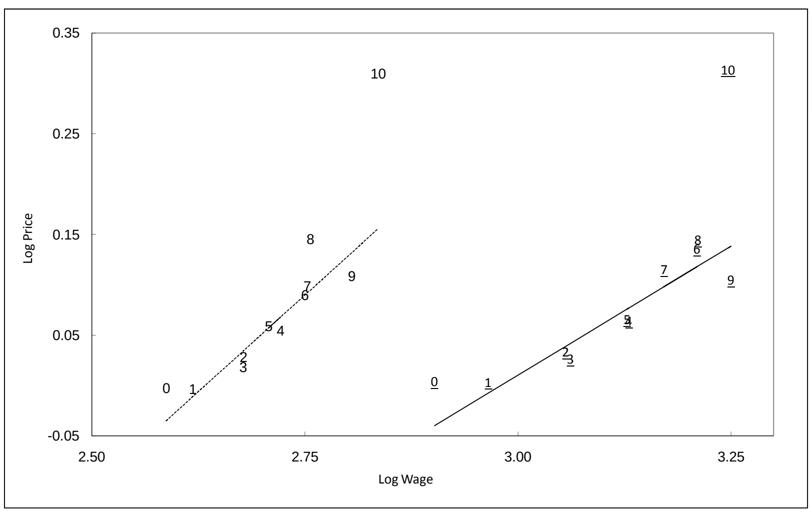

Panel D: 2004-7

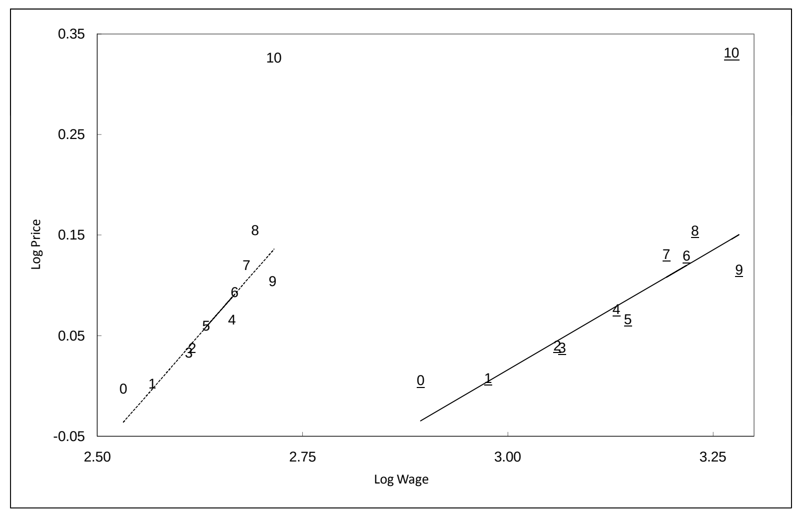

Notes: Each plotted point represents average log prices and log wages for the listed city size category. Underlined numbers are for high school graduates while other numbers are for college graduates. Log prices are the mean for the education group in the listed size category minus the overall mean for rural areas. High school and college log prices differ slightly because of different distributions of individuals across locations by education within size categories. Solid lines are indifference maps for college graduates in locations 1-9 while dashed lines are indifference maps for high school graduates in locations 1-9. See the notes to Table 1 for a description of the sample.

Figure A1: Prices by City Size by Year

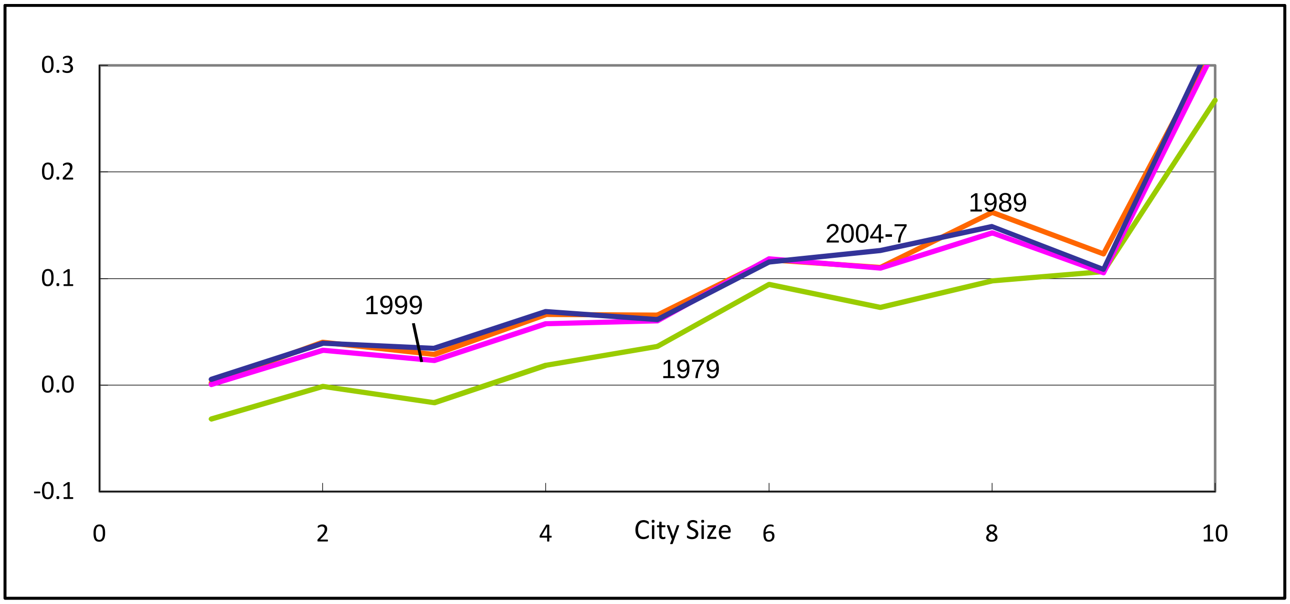

Notes: Each line shows the average prices across individuals relative to rural areas in each year.

Figure A2: Fraction College or More by City Size and Year

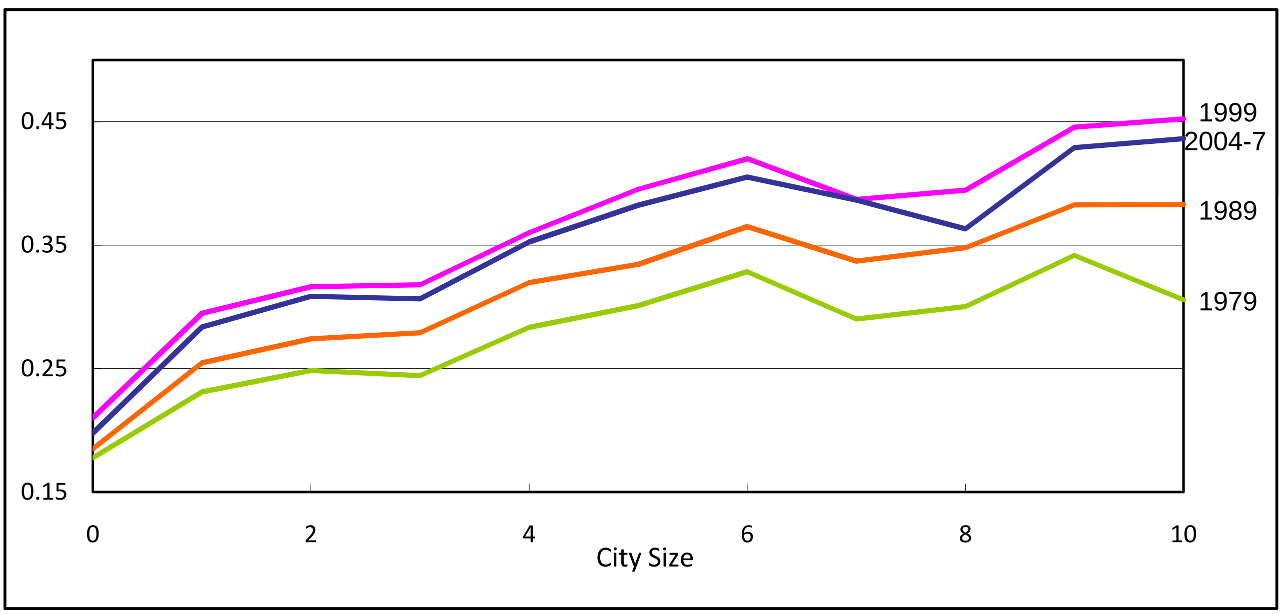

Notes: Each line shows the fraction of the sample with a college education or more by city size category in each year.

Table 1: Trends in Hourly Wage Inequality

	Residual Variance	Residual 90-50 Gap	Residual 50-10 Gap	Total Variance	Total 90-50 Gap	Total 50-10 Gap
Panel A: Nominal Wages						
1979	0.16	0.45	0.55	0.21	0.53	0.63
1989	0.21	0.50	0.60	0.29	0.62	0.70
1999	0.25	0.56	0.60	0.34	0.72	0.67
2004-7	0.27	0.61	0.64	0.39	0.81	0.74
Panel B: Wages Adjusted for Cost of Living						
1979	0.17	0.46	0.56	0.21	0.52	0.63
1989	0.21	0.50	0.60	0.28	0.61	0.68
1999	0.25	0.56	0.60	0.33	0.70	0.66
2004-7	0.28	0.62	0.64	0.38	0.77	0.73
Panel C: Fraction Reduction in Growth Due to Cost of Living Adjustment						
1979 to 1989	0.02	-0.02	0.08	0.17	0.08	0.34
1979 to 1999	-0.02	-0.01	-0.05	0.10	0.13	0.16
1979 to 2004-7	-0.04	-0.03	-0.04	0.08	0.14	0.06

Notes: See the notes to Figure 1 for a description of the sample. Residuals are calculated using age, education and city size cell means of hourly wages.

Table 2: Components of the Variance of Hourly Wages

	Between Cells			Within Cells			Total		
	Quantities	Prices	Sum	Quantities	Prices	Sum	Quantities	Prices	Sum
1979 to 1989	0.000	0.034	0.034	0.000	0.046	0.046	0.000	0.080	0.080
1989 to 1999	0.001	0.011	0.012	0.012	0.027	0.039	0.013	0.038	0.051
1999 to 2004-7	0.004	0.022	0.026	0.000	0.025	0.025	0.004	0.047	0.051
1979 to 2004-7	0.008	0.065	0.073	0.014	0.096	0.110	0.022	0.160	0.183

Notes: See the notes to Figure 1 for a description of the sample. The text explains our procedure for separating the price and quantity components of the growth of each portion of the variance.

Table 3: Shifts in Industry Composition by City Size

City Size Index	Sum of Squared Industry Share Changes 1980-1990	1990-2000	2000-2005-7	1980-2005-7
Panel A: High School Dropouts				
Rural	0.001	0.002	0.003	0.012
1	0.003	0.003	0.004	0.021
2	0.004	0.005	0.006	0.038
3	0.005	0.004	0.004	0.037
4	0.006	0.004	0.004	0.039
5	0.004	0.009	0.004	0.047
6	0.007	0.007	0.004	0.045
7	0.007	0.007	0.004	0.047
8	0.005	0.003	0.009	0.041
9	0.007	0.008	0.004	0.050
10	0.008	0.001	0.004	0.031
Panel B: High School Only				
Rural	0.000	0.001	0.002	0.007
1	0.002	0.001	0.002	0.014
2	0.003	0.002	0.002	0.019
3	0.004	0.001	0.002	0.019
4	0.004	0.002	0.002	0.021
5	0.002	0.003	0.002	0.020
6	0.004	0.003	0.002	0.022
7	0.005	0.002	0.002	0.023
8	0.003	0.002	0.003	0.021
9	0.004	0.002	0.002	0.023
10	0.004	0.002	0.002	0.022
Panel C: Some College				
Rural	0.000	0.001	0.001	0.005
1	0.002	0.001	0.001	0.010
2	0.001	0.002	0.001	0.011
3	0.002	0.001	0.001	0.011
4	0.002	0.002	0.002	0.014
5	0.001	0.002	0.001	0.011
6	0.002	0.003	0.001	0.014
7	0.002	0.002	0.001	0.015
8	0.002	0.003	0.002	0.016
9	0.003	0.002	0.001	0.012
10	0.002	0.003	0.001	0.016

Panel D: College Only

Rural	0.001	0.001	0.000	0.006
1	0.004	0.002	0.000	0.012
2	0.003	0.002	0.000	0.015
3	0.002	0.003	0.000	0.014
4	0.003	0.003	0.001	0.017
5	0.003	0.004	0.001	0.017
6	0.003	0.005	0.000	0.020
7	0.003	0.005	0.001	0.019
8	0.002	0.005	0.001	0.018
9	0.004	0.005	0.000	0.017
10	0.003	0.006	0.001	0.021

Panel E: More Than College

Rural	0.004	0.000	0.000	0.004
1	0.006	0.000	0.000	0.009
2	0.002	0.000	0.000	0.005
3	0.002	0.001	0.000	0.007
4	0.003	0.000	0.000	0.006
5	0.004	0.001	0.000	0.006
6	0.003	0.002	0.001	0.012
7	0.003	0.002	0.000	0.008
8	0.004	0.002	0.001	0.012
9	0.004	0.001	0.000	0.009
10	0.004	0.002	0.001	0.011

Notes: Each entry is the sum of the squared difference in 1-digit industry shares between the years in the column headers for the city size category indicated in the row and the education group listed in each panel header.

Table 4: Impacts of City Size on Residual Inequality

Year	Variance			90 - 50 Percentile Gap			50 - 10 Percentile Gap		
		Counterfactual			Counterfactual			Counterfactual	
	Actual	1980 Quantities	1980 City Sizes	Actual	1980 Quantities	1980 City Sizes	Actual	1980 Quantities	1980 City Sizes
Panel A: Measures of Residual Distributions									
1979	0.16	0.16	0.16	0.45	0.45	0.45	0.55	0.55	0.55
1989	0.21	0.21	0.20	0.50	0.50	0.49	0.60	0.61	0.60
1999	0.25	0.23	0.20	0.56	0.55	0.52	0.60	0.59	0.55
2004-7	0.27	0.26	0.23	0.61	0.62	0.57	0.64	0.64	0.59
Panel B: Fraction Reduction in Growth Due to City Size Effects									
1979 to 1989	0.00	0.00	0.24	0.00	-0.03	0.23	0.00	-0.14	0.05
1979 to 1999	0.00	0.19	0.51	0.00	0.11	0.38	0.00	0.19	1.06
1979 to 2004-7	0.00	0.13	0.36	0.00	-0.01	0.29	0.00	-0.03	0.60

Notes: The two columns in Panel A under the heading of Counterfactual are measures of inequality formed using counterfactual residual distributions. "1980 Quantities" holds the quantities of individuals by demographic group and location at their 1980 distribution but allows prices to change over time as they did in equilibrium. "1980 City Sizes" uses the 1980 profile of distribution spread as a function of city size. See the text for a more complete explanation of the construction of these counterfactual measures.

Table 5: Impacts of City Size on Measures of Wage Inequality

Year	Variance				90 - 50 Percentile Gap				50 - 10 Percentile Gap			
		Counterfactual				Counterfactual				Counterfactual		
	Actual	1980 Quantities	Composite Base	1980 Base	Actual	1980 Quantities	Composite Base	1980 Base	Actual	1980 Quantities	Composite Base	1980 Base
Panel A: Measures of Wage Inequality												
1979	0.21	0.21	0.21	0.21	0.53	0.53	0.53	0.53	0.63	0.63	0.63	0.63
1989	0.29	0.29	0.27	0.27	0.62	0.63	0.62	0.62	0.70	0.71	0.69	0.69
1999	0.34	0.32	0.29	0.29	0.72	0.71	0.69	0.67	0.67	0.66	0.63	0.63
2004-7	0.39	0.37	0.35	0.33	0.81	0.78	0.76	0.73	0.74	0.72	0.70	0.68
Panel B: Fraction Reduction in Growth Due to City Size Effects												
1979 to 1989	0.00	0.00	0.21	0.20	0.00	-0.21	-0.04	-0.03	0.00	-0.14	0.18	0.16
1979 to 1999	0.00	0.14	0.37	0.42	0.00	0.07	0.18	0.24	0.00	0.25	0.84	1.03
1979 to 2004-7	0.00	0.12	0.24	0.33	0.00	0.12	0.18	0.27	0.00	0.19	0.35	0.52

Notes: Measures of distribution inequality listed in the "Composite Base" and "1980 Base" columns are built using counterfactual hourly wage distributions absent changes in city size effects after 1980. The residuals for these distributions are the same as those used in Table 4. The intercepts for the Composite Base columns are built holding the slope with respect to city size the same as its 1980 level in each demographic group and weighting by contemporaneous population shares. The intercepts for the "1980 Base" columns instead weight by 1980 population shares. See the text for a more complete explanation.