

Real Wage Inequality

Enrico Moretti
UC Berkeley, NBER, CEPR and IZA

First Draft: May 2008
This Draft February 2011

Abstract. A large literature has documented a significant increase in the difference between the wage of college graduates and high school graduates over the past 30 years. I show that from 1980 to 2000, college graduates have experienced relatively larger increases in cost of living, because they have increasingly concentrated in metropolitan areas that are characterized by a high cost of housing. When I deflate nominal wages using a location-specific CPI, I find that the difference between the wage of college graduates and high school graduates is lower in real terms than in nominal terms and has grown less. At least 22% of the documented increase in college premium is accounted for by spatial differences in the cost of living. The implications of this finding for changes in well-being inequality depend on why college graduates sort into expensive cities. Using a simple general equilibrium model of the labor and housing markets, I consider two alternative explanations. First, it is possible that the relative supply of college graduates increases in expensive cities because college graduates are increasingly attracted by amenities located in those cities. In this case, the higher cost of housing reflects consumption of desirable local amenities, and there may still be a significant increase in well-being inequality even if the increase in real wage inequality is limited. Alternatively, it is possible that the relative demand for college graduates increases in expensive cities due to shifts in the relative productivity of skilled labor. In this case, the relative increase in skilled workers' standard of living is offset by the higher cost of living. The evidence indicates that changes in the geographical location of different skill groups are mostly driven by changes in their relative demand. I conclude that the increase in well-being disparities between 1980 and 2000 is smaller than previously thought.

I thank David Autor, Dan Black, David Card, Tom Davidoff, Ed Glaeser, Chang-Tai Hsieh, Matt Kahn, Pat Kline, Douglas Krupka, David Levine, Adam Looney and Krishna Pendakur for insightful conversations, and seminar participants at Banco de Portugal, Berkeley Economics, Berkeley Haas, Bocconi, Bologna, British Columbia, Chicago Harris, Collegio Carlo Alberto, Edinburgh, Federal Reserve Board of Governors, IZA, Milano, Missouri, NBER Summer Institute, Northwestern, Oxford, San Francisco Federal Reserve, Simon Fraser, Stanford, Stanford GSB, UCLA, UC Santa Cruz, Toronto, Tulane, UC Merced, Verona and Victoria for many useful comments. I thank Emek Basker for generously providing the Accra data on consumption prices. Issi Romen, Mariana Carrera, Justin Gallagher, Jonas Hjort, Max Kasy and Zach Liscow provided excellent research assistance.

1 Introduction

One of the most important developments in the US labor market over the past 30 years has been a significant increase in wage inequality. For example, the difference between the wage of skilled and unskilled workers has increased significantly since 1980. The existing literature has focused on three classes of explanations: an increase in the relative demand for skills caused, for example, by skill biased technical change; a slowdown in the growth of the relative supply of skilled workers; and the erosion of labor market institutions that protect low-wage workers.¹

In this paper, I re-examine how inequality is measured and how it is interpreted. I begin by noting that skilled and unskilled workers are not distributed uniformly across cities within the US, and I assess how existing estimates of inequality change when differences in the cost of living across locations are taken into account. I then discuss how to interpret these measures of real wage inequality when changes in amenities are different across cities.

I focus on changes between 1980 and 2000 in the difference in the average hourly wage for workers with a high school degree and workers with college or more. Using Census data, I show that from 1980 to 2000 college graduates have increasingly concentrated in metropolitan areas with a high cost of housing. This is due both to the fact that college graduates in 1980 are overrepresented in cities that experience large increases in housing costs and to the fact that much of the growth in the number of college graduates has occurred in cities with initial high housing costs. College graduates are therefore increasingly exposed to a high cost of living and the relative increase in their real wage may be smaller than the relative increase in their nominal wage.

To measure the wage difference between college graduates and high school graduates in *real* terms, I deflate nominal wages using a cost of living index that allows for price differences across metropolitan areas. I closely follow the methodology that the Bureau of Labor Statistics uses to build the official CPI, while allowing for changes in the cost of housing to vary across metropolitan areas. Since housing is by far the largest item in the CPI—accounting for more than a third of the index—geographical differences in housing costs have the potential to significantly affect the local CPI. In some specifications, I also allow for local variation in non-housing prices.

The results are striking. First, I find that between 1980 and 2000, the cost of housing for college graduates grows much faster than cost of housing for high school graduates. Specifically, in 1980 the difference in the average cost of housing between college and high school graduates is only 4%. This difference grows to 14% in 2000, or more than three times the 1980 difference. Second, consistent with what is documented by the previous literature, I find that the difference between the nominal wage of high school and college graduates has increased 20 percentage points between 1980 and 2000. However, the difference between the real wage of high school and college graduates has increased significantly less. Changes in the

¹A comprehensive survey is found in Katz and Autor (1999).

cost of living experienced by high school and college graduates account for about a quarter of the increase in the nominal college premium over the 1980-2000 period. This finding does not appear to be driven by different trends in relative worker ability or housing quality and is robust to a number of alternative specifications. Third, the difference between the wage of college graduates and high school graduates is smaller in real terms than in nominal terms for each year. For example, in 2000 the difference is 60% in nominal terms and 51% in real terms.

Overall, the difference in the real wage between skilled and unskilled workers is smaller than the nominal difference and has grown less.² Does this finding mean that the significant increases in wage disparities that have been documented by the previous literature over the last 30 years have failed to translate into significant increases in disparities in well-being? Not necessarily. Since local amenities differ significantly across cities, changes in real wages do not necessarily equal changes in well-being.

To understand the implications of my empirical findings for well-being inequality, I use a simple general equilibrium model of the housing and labor markets with two types of labor, skilled and unskilled.³ The model indicates that the implications of my empirical findings for well-being inequality crucially depend on why college graduates tend to sort into expensive metropolitan areas. I consider two possible explanations. First, it is possible that college graduates move to expensive cities because firms in those cities experience an increase in the relative demand for skilled workers. This increase can be due to localized skill-biased technical change or positive shocks to the product demand for skill intensive industries that are predominantly located in expensive cities (for example, high tech and finance are mostly located in expensive coastal cities). If college graduates increasingly concentrate in expensive cities such as San Francisco and New York because the jobs for college graduates are increasingly concentrated in those cities—and not because they particularly like living in San Francisco and New York—then the increase in their utility level is smaller than the increase in their nominal wage. In this scenario, the increase in well-being inequality is smaller than the increase in nominal wage inequality because of the higher costs of living faced by college graduates.

Alternatively, it is possible that college graduates move to expensive cities because the relative supply of skilled workers increases in those cities. This may be due, for example, to an increase in the local amenities that attract college graduates. In this scenario, increases in the cost of living in these cities reflect the increased attractiveness of the cities and represent the price to pay for the consumption of desirable amenities. This consumption arguably

²It is worth stressing that changes in cost of living, while clearly important, account only for a fraction of the overall increase in wage inequality in this period.

³The model clarifies what happens to employment, wages, costs of housing of skilled and unskilled workers and when a local economy experiences a shock to the productivity of skilled labor or a change in local amenities. Unlike Roback (1982), productivity and amenity shocks are not necessarily fully capitalized into land prices. This allows shocks to the relative demand and relative supply of skilled workers in a city to have different effects on the well-being of skilled and unskilled workers and landowners.

generates utility. If college graduates move to expensive cities like San Francisco and New York because they want to enjoy the local amenities—and not primarily because of labor demand—then there may still be a significant increase in utility inequality even if the increase in real wage inequality is limited.⁴ Of course, the two scenarios are not mutually exclusive, since in practice it is possible that both relative demand and supply shift at the same time.

To determine whether relative demand or relative supply shocks are more important in practice, I analyze the empirical relationship between changes in the college premium and changes in the share of college graduates across metropolitan areas. My model indicates that under the relative demand hypothesis, one should see a *positive* equilibrium relationship between changes in the college premium and changes in the college share. Intuitively, increases in the relative demand of college graduates in a city should result in increases in their relative wage there. Under the relative supply hypothesis, one should not see such a positive relationship. This test is related to the test proposed by Katz and Murphy (1992) to understand nationwide changes in inequality.

Consistent with relative demand shocks playing an important role, I find a strong *positive* association between changes in the college premium and changes in the college share. While this suggests that demand factors are important, it does not necessarily rule out supply factors. As a second piece of evidence, I present instrumental variable estimates of the relationship between changes in the college premium and changes in the college share based on a shift-share instrument.⁵ The IV estimate establishes what happens to the college premium in a city when the city experiences an increase in the number of college graduates that is driven purely by an increase in the relative demand for college graduates. By contrast, the OLS estimate establishes what happens to the college premium in a city when the city experiences an increase in the number of college graduates that may be driven by either demand or supply shocks. The comparison of the two estimates is therefore informative about the relative importance of demand and supply shocks.

Overall, the empirical evidence is more consistent with the notion that relative demand shocks are the main force driving changes in the number of skilled workers across metropolitan areas.⁶ If this is true, the increase in well-being disparities between 1980 and 2000 is significantly smaller than we previously thought based on the existing literature.⁷

⁴See also Kahn (1999).

⁵The instrument is a weighted average of nationwide relative skilled employment growth by industry, with weights reflecting the city-specific employment share in those industries in 1980.

⁶This finding does not imply that amenities do not affect worker location decisions in general, nor it implies that amenities do not affect location decisions of skilled and unskilled workers differently. Rather, it implies that the differential change in the geographical location of skilled and unskilled workers in this period is mostly driven by *changes* in the geographical *differences* in the availability of new skilled and unskilled jobs, not by *changes* in the geographical *differences* in local amenities that are relevant to skilled and unskilled workers.

⁷I note that the exact magnitude of the increase in well-being disparities remains unknown. While my estimates indicate that the increase in well-being disparities is smaller than suggested by existing estimates, a full account of changes in well-being disparities would require several additional pieces of information. For

My findings are consistent with previous studies that identify shifts in labor demand—whether due to skill-biased technical change or product demand shifts across industries with different skill intensities—as an important determinant of the increase in wage inequality (for example, Katz and Murphy, 1992). But unlike the previous literature, my findings point to an important role for the *local* component of these demand shifts. While in this paper I take these local demand shifts as exogenous, future research should investigate the economic forces that make skilled workers more productive in some parts of the country.⁸ The notion that demand shocks are important determinants of population shifts is consistent with the evidence in Blanchard and Katz (1992) and Bound and Holzer (2000).⁹ The specific finding that variation in the college share is mostly driven by demand factors is consistent with the argument made by Berry and Glaeser (2005) and Beaudry, Doms and Lewis (2008).

My results are also related to a series of papers by Pendakur (1998, 2002) and Lewbel and Pendakur (forthcoming) on the correct use of price indexes on the measurement of inequality. My approach is related to a paper by Black et al. (2010) which, along with earlier work by Dahl (2002), criticizes the standard practice of treating the returns to education as uniform across locations. They show that, in theory, the return to schooling is constant across locations only in the special case of homothetic preferences, and argue that the returns to education are empirically lower in high-amenity locations.¹⁰ My findings complement the literature on consumption inequality, which has documented that income inequality is higher and has grown faster than consumption inequality in many countries, including the US. See Krueger, Perri, Pistaferri, Violante (2010) for a recent review of the evidence. In principle, my estimates have the potential to provide an explanation for the slower increase in consumption inequality in this period.¹¹

From the methodological point of view, this paper illustrates the importance of accounting for general equilibrium effects when thinking about the effects of group specific labor market shocks. Labor economists often approach the analysis of labor market shocks using a partial equilibrium analysis. However, this paper shows that a partial equilibrium analysis can miss

example, it would require estimates of relative changes in features of jobs other than wages (job amenities, other forms of compensation, etc.) and estimates of the relative changes in housing wealth induced by changes in housing prices. A full empirical treatment of these issues is complicated and is beyond the scope of this paper.

⁸See for example Moretti (2004a and 2004b) and Greenstone, Hornbeck and Moretti (forthcoming).

⁹Chen and Rosenthal (forthcoming) document that jobs are the key determinant of mobility of young individuals. Mobility of older individuals seems more likely to be driven by amenities.

¹⁰In a related paper, Black et al. (2009) argue that estimates of the wage differences between blacks and whites need to account for differences in the geographical location of different racial groups and develop a theoretical model to understand when estimates of black-white earnings gap can be used to infer welfare differences.

¹¹See also Duranton (2002 and 2008) on spatial wage disparities; Broda and Romalis (2009) who document the distributional consequences of increased imports from China; Gordon (2009) and Gordon and Dew-Becker (2005, 2007 and 2008); and Aguiar and Hurst (2007a and 2007b) who focus on the role of differential changes in labor supply and leisure, by skill group.

important parts of the picture, since the endogenous reaction of factor prices and quantities can significantly alter the ultimate effects of a shock. Because aggregate shocks to the labor market are rarely geographically uniform, the geographic reallocation of factors and local price adjustments are empirically important. It is difficult to fully understand *aggregate* labor market changes—like changes in relative wages— if ignoring the spatial dimension of labor markets. This paper shows that labor flows across localities and changes in local prices have the potential to undo some of the direct effects of labor market shocks and this may alter the implications for policy.

The rest of the paper is organized as follows. In Section 2, I describe how the official CPI is calculated by the BLS and I propose two alternative CPI's that allow for geographical differences across skill groups. In Section 3, I present estimates of nominal and real college premia. In Section 4, I present a simple model that can help interpreting the empirical evidence. In Section 5, I discuss the different implications of the demand pull and supply push hypotheses and present empirical evidence to distinguish the two. Section 6 concludes.

2 Cost of Living Indexes and the Location of Skilled and Unskilled Workers

In this Section, I begin with some descriptive evidence on recent changes in the geographical location of skilled and unskilled workers and housing costs (subsection 2.1). I then describe how the Bureau of Labor Statistics computes the official Consumer Price Index and I propose two alternative measures of cost of living that account for geographical differences (subsection 2.2). Finally, I use my measures of cost of living to document the differential change in the cost of living experienced by high school and college graduates between 1980 and 2000 (subsection 2.3).

2.1 Changes in the Location of Skilled and Unskilled Workers

Throughout the paper, I use data from the 1980, 1990 and 2000 Censuses of Population.¹² The geographical unit of analysis is the metropolitan statistical area (MSA) of residence. Rural households in the Census are not assigned a MSA. In order to keep my wage regressions as representative and as consistent with the previous literature as possible, I group workers who live outside a MSA by state, and treat these groups as additional geographical units.

Table 1 documents differences in the fraction of college graduates across some US metropolitan areas. Specifically, the top (bottom) panel reports the 10 cities with the highest (lowest) fraction of workers with a college degree or more in 2000. Throughout the paper, college graduates also include individuals with a post-graduate education. The metropolitan area

¹²Because my data end in 2000, my empirical analysis is not affected by the run-up in home prices during the housing bubble years and the subsequent decline in home prices.

with the largest share of workers with a college degree among its residents is Stamford, CT, where 58% of workers has a college degree or more. The fraction of college graduates in Stamford is almost 5 times the fraction of college graduates in the city at the bottom on the distribution—Danville, VA—where only 12% of workers have a college degree. Other metropolitan areas in the top group include MSA’s with an industrial mix that is heavy in high tech and R&D—such as San Jose, San Francisco, Boston and Raleigh-Durham—and MSA’s with large universities— such as Ann Arbor, MI and Fort Collins, CO. Metropolitan areas in the top panel have a higher cost of housing—as measured by the average monthly rent for a 2 or 3 bedroom apartment—than metropolitan areas in the bottom panel. College share and the cost of housing vary substantially not only in their levels across locations but also in their changes over time. While cities like Stamford, Boston, San Jose and San Francisco experienced large increases in both the share of workers with a college degree and the monthly rent between 1980 and 2000, cities in the bottom panel experienced more limited increases.

The relation between changes in the number of college graduates and changes in housing costs is shown more systematically in Figure 1. The top panel shows how the 1980-2000 change in the share of college graduates relates to the 1980 share of college graduates. The size of the bubbles reflect population in 1980. The positive relationship indicates that college graduates are increasingly concentrated in metropolitan areas that have a large share of college graduates in 1980. This relationship has been documented by Berry and Glaeser (2005) and Moretti (2004), among others.¹³

The middle panel of Figure 1 shows how the 1980-2000 change in the share of college graduates relates to the average cost of housing in 1980. The positive relationship indicates that college graduates are increasingly concentrated in MSA’s where housing is initially expensive.¹⁴ The bottom panel plots the 1980-2000 change in college share as a function of the 1980-2000 change in the average monthly rental price. The positive relationship suggests that the share of college graduates has increased in MSA’s where housing has become more expensive.¹⁵

These relationships do not have a causal interpretation, but instead need to be interpreted as equilibrium relationships. Taken together, the panels in Figure 1 show that the metropolitan areas that have experienced the largest increases in the share of college graduates are the metropolitan areas where the average cost of housing in 1980 is highest and

¹³The regression of the 1980-2000 change in college share on the 1980 level in college share weighted by the 1980 MSA size yields a coefficient equal to .460 (.032), indicating that a 10 percentage point difference in the baseline college share in 1980 is associated with a 4.6 percentage point increase in college share between 1980 and 2000.

¹⁴The regression of the 1980-2000 change in college share on the 1980 cost of housing weighted by the 1980 MSA size yields a coefficient equal to .0011 (.00006), indicating that a 100 dollar difference in the baseline monthly rent in 1980 is associated with a 4.7 percentage point increase in college share between 1980 and 2000.

¹⁵The regression yields a coefficient equal to .0003 (.00001).

also the areas where the average cost of housing has increased the most.

2.2 Local Consumer Price Indexes

A cost of living index seeks to measure changes over time in the amount that consumers need to spend to reach a certain utility level or “standard of living.” Changes in the official Consumer Price Index between period t and $t + 1$ as measured by the Bureau of Labor Statistics are a weighted average of changes in the price of the goods in a representative consumption basket. The basket is the original consumption basket at time t , and the weights reflect the share of income that the average consumer spends on each good at time t .¹⁶

Table 2 shows the relative importance of the main aggregate components of the CPI for all urban consumers, CPI-U, in 2000. The largest component by far is housing. In 2000, housing accounts for more than 42% of the CPI-U. The largest sub-components of housing costs are “Shelter” and “Fuel and Utilities”. The second and third main components of the CPI-U are transportation and food. They only account for 17.2% and 14.9% of the CPI-U, respectively. The weights of all the other categories are 6% or smaller.

Although most households in the US are homeowners, changes in the price of housing are measured by the BLS using changes in the cost of renting an apartment (Poole, Ptacek and Verbugge, 2006; Bureau of Labor Statistics, 2007). The rationale for using rental costs instead of home prices is that rental costs are a better approximation of the user cost of housing. Since houses are an asset, their price reflects both the user cost as well as expectations of future appreciation.

Rental costs vary significantly across metropolitan areas. For example, in 2000, the average rental cost for a 2 or 3 bedroom apartment in San Diego, CA—the city at the 90th percentile of the distribution—is \$894. This rental cost is almost 3 times higher than the rental cost for an equally sized apartment in Decatur, AL, the city at the 10th percentile. Changes over time in rental costs also vary significantly across metropolitan areas. For example, between 1980 and 2000, the rental cost increased by \$165 in Johnstown, PA—one of the cities at the bottom of the distribution—and by \$892 in San Jose—one of the cities at the top of the distribution. The distribution of average rental costs and changes in average rental costs are shown in Figure 2.

Although the cost of living varies substantially across metropolitan areas, wage and income are typically deflated using a single, nation-wide deflator, such as the CPI-U calculated by the BLS. The use a nation-wide deflator is particularly striking in light of the fact that

¹⁶One well known problem with the CPI is the potential for substitution bias, which is the possibility that consumers respond to price changes by substituting relatively cheaper goods for goods that have become more expensive. While the actual consumption baskets may change, the CPI reports inflation for the original basket. Details of the BLS methodology are described in Chapter 17 of the Handbook of Methods (BLS, 2007), titled “The Consumer Price Index”.

more than 40% of the CPI-U is driven by housing costs (Table 2), and that housing costs vary so much across locations (Figure 2). To investigate the role of cost of living differences on wage differences between skill groups, I propose two alternative CPI indexes that vary across metropolitan areas. I closely follow the methodology that the Bureau of Labor Statistics uses to build the official Consumer Price Index, but I generalize two of its assumptions.

Local CPI 1. First, I compute a CPI that allows for the fact that the cost of housing varies across metropolitan areas. I call the resulting local price index “Local CPI 1”. Following the BLS methodology, I define Local CPI 1 as the properly weighted sum of local cost of housing—with the average across cities normalized to 1 in 1980—and non-housing consumption—normalized to 1 in 1980. I measure the cost of housing faced by an individual in metropolitan area c in two ways. In my preferred specification, I follow the BLS methodology and I use rental costs. I assign the cost of housing to residents in a metropolitan area based on the relevant average monthly rent. Specifically, I take the average of the monthly cost of renting a 2 or 3 bedroom apartment among all renters in area c . As an alternative way to measure cost of housing, in some models I use the price of owner occupied houses instead of rental costs. Specifically, I take the average reported value of all 2 or 3 bedroom owner occupied single family houses in area c . Both rental costs and housing prices are from the Census of Population. As I discuss later, empirical results are not sensitive to measuring housing costs using rental costs or housing prices. The price of non housing goods and services is assumed to be the same in a given year, irrespective of location. This assumption is relaxed in Local CPI 2.¹⁷

I describe the details of this approach in Appendix 1. It is important to note that this methodology ensures that the deflator that I use for a given worker does not reflect the increase in the cost of the apartment rented or the cost of the house owned by that specific worker. Instead, it reflects the increase in the cost of housing experienced by residents in the same city, irrespective of their own individual housing cost and irrespective of whether they rent or own.

Local CPI 2. In local CPI 1, changes in the cost of housing can vary across localities, but changes in the cost of non-housing goods and services are assumed to be the same everywhere. While the cost of housing is the most important component of the CPI, the price of other goods and services is likely to vary systematically with the cost of housing. In cities where land is more expensive, production and retail costs are higher and therefore the cost of many goods and services is higher. For example, a slice of pizza or a hair cut are

¹⁷The motivation for using 2 and 3 bedroom apartments is to keep the size of the apartment roughly similar. I have experimented with variants of this selection rule. Estimates based only on 2 bedroom apartments are similar to the ones presented below. Estimates based on data from the American Housing Survey that use information on square footage to hold constant the exact size of the apartment also yield similar results. See Section 3.3 below for details.

likely to be more expensive in New York city than in Indianapolis, since it is more expensive to operate a pizza restaurant or a barber shop in New York city than Indianapolis.

Local CPI 2 allows for both the cost of housing and the cost of non-housing consumption to vary across metropolitan areas. Systematic, high quality, city-level data on the price of non-housing good and services are not available for most cities over a long time period. To overcome this limitation, I use two alternative approaches. First, in my preferred specification, I use the fact that the BLS releases a local CPI for a limited number of metropolitan areas. This local CPI is not ideal because it is made available by the BLS only for 23 MSA’s in the period under consideration and there are 315 MSA’s in the 2000 Census. Additionally, it is normalized to 1 in a given year, thus precluding cross-sectional comparisons. However, it can still be used to impute the part of local non-housing prices that varies systematically with housing costs. The local CPI computed by the BLS for city c in year t is a weighted average of housing cost (HP_{ct}) and non-housing costs (NHP_{ct}): $BLS_{ct} = wHP_{ct} + (1 - w)NHP_{ct}$ where w is the CPI weight used by BLS for housing. Non-housing costs can be divided in two components:

$$NHP_{ct} = \pi HP_{ct} + v_{ct} \tag{1}$$

where πHP_{ct} is the component of non-housing costs that varies systematically with housing costs; and v_{ct} is the component that is orthogonal to housing costs. If $\pi > 0$ it means that cities with higher cost of housing also have higher costs of non-housing goods and services. I use the small sample of MSA’s for which a local BLS CPI is available to estimate π .¹⁸ I then impute the systematic component of non-housing costs to all MSA’s, based on their housing cost: $E(NHP_{ct}|HP_{ct}) = \hat{\pi}HP_{ct}$. Finally, I compute “Local CPI 2” as a properly weighted sum of the cost of housing, the component of non-housing costs that varies with housing ($\hat{\pi}HP_{ct}$), and the component of non-housing costs that does not vary with housing. See Appendix 1 for more details.

As an alternative strategy to measure local variation in non-housing prices, I use data on non-housing prices taken from the Accra dataset, which is collected by the Council for Community and Economic Research.¹⁹ The Accra data have both advantages and disadvantages. On one hand, the Accra data are available for most cities, and therefore do not require any imputation. Furthermore, the detail is such that price information is available at the level of specific consumption goods and the price is not normalized to a base year. On the other hand, the Accra data are available only for a very limited number of goods.²⁰ Importantly,

¹⁸To do so, I first regress changes in the BLS local index on changes in housing costs: $\Delta BLS_{ct} = \beta \Delta HP_{ct} + e_{ct}$. Estimating this regression in differences is necessary because BLS_{ct} is normalized to 1 in a given year. While cross-sectional comparisons based on BLS_{ct} are meaningless, BLS_{ct} does measure changes in prices within a city. Once I have an estimate of β , I can calculate $\hat{\pi} = \frac{\hat{\beta} - w}{1 - w}$. Empirically, $\hat{\beta}$ is equal to .588 (.001) and $\hat{\pi}$ is equal to .35 in 2000.

¹⁹The data were generously provided by Emek Basker. Basker (2005) and Basker and Noel (2007) describe the Accra dataset in detail.

²⁰Only 48 goods have prices that are consistently defined for the entire period under consideration. The BLS basket includes more than 1000 goods.

the sample size for each good and city is quite small, so that local price averages are noisy. Additionally, the set of cities covered changes over time. In practice, the empirical findings based on the version of local CPI 2 that uses the imputation and those based on the version of local CPI 2 that uses Accra data are similar.

In sum, local CPI 2 is more comprehensive than Local CPI 1 because it includes local variation in both housing and non-housing costs, but it has the limitation that non-housing costs are imputed or come from Accra data. For this reason, in the next Section I present separate estimates for Local CPI 1 and Local CPI 2.

2.3 Changes in the Cost of Living Experienced by Skilled and Unskilled Workers Between 1980 and 2000

I now quantify the changes in the cost of living experienced by high school and college graduates between 1980 and 2000. The top panel of Table 3 shows changes in the official CPI-U, as reported by the BLS, and normalized to 1 in 1980. This is the most widely used measure of inflation, and it is the measure that is almost universally used to deflate wages and incomes. According to this index, the price level doubled between 1980 and 2000. This increase is—by construction—the same for college graduates and high school graduates.

The next panel shows the increase in the cost of housing faced by college graduates and high school graduates. College graduates and high school graduates are exposed to very different increases in the cost of housing. In 1980 the cost of housing for the average college graduate is only 4% more than the cost of housing for the average high school graduate. This gap grows to 11% in 1990 and reaches 14% by 2000. Column 4 indicates that housing costs for high school and college graduates increased between 1980 and 2000 by 127% and 147%, respectively.

The third panel shows “Local CPI 1”, normalized to 1 in 1980 for the average household.²¹ The panel shows that in 1980 the overall cost of living experienced by college graduates is only 2% higher than the cost of living experienced by high school graduates. This difference increases to 6% by year 2000. The difference in Local CPI 1 between high school and college graduates is less pronounced than the difference in monthly rent because Local CPI 1 includes non-housing costs as well as housing costs.

The differential increase in cost of living faced by college graduates relative to high school graduates is more pronounced when the price of non-housing goods and services is allowed to vary across locations, as in the bottom panel. In the case of Local CPI 2, the cost of living is 3% higher for college graduates relative to high school graduates in 1980 and 9% in 2000. Column 4 indicates that the increase in the overall price level experienced by high school graduates between 1980 and 2000 is 108%. The increase in the overall price level experienced by college graduates between 1980 and 2000 is 119%.

²¹Here I use rental costs to measure housing costs. Using property values for owner occupied houses yields similar results.

The relative increase in the cost of housing experienced by college graduates between 1980 and 2000 can be decomposed into a part due to geographical mobility and a part due to the fact that already in 1980 college graduates are overrepresented in cities that experience large increases in costs. Specifically, the 1980-2000 nationwide change in the cost of housing experienced by skill group j (j =high school or college), can be written as

$$P_{j2000} - P_{j1980} = \sum_c \omega_{jc2000} P_{c2000} - \sum_c \omega_{jc1980} P_{c1980} \\ \sum_c (\omega_{jc2000} - \omega_{jc1980}) P_{c2000} + \sum_c \omega_{jc1980} (P_{c2000} - P_{c1980})$$

where ω_{jct} is the share of workers in skill group j who live in city c in year t and P_{ct} is the cost of housing in city c in year t . The equation illustrates that the total change in cost of housing is the sum of two components: a part due to the change in the share of workers in each city, given 2000 prices ($\sum_c (\omega_{jc2000} - \omega_{jc1980}) P_{c2000}$); and a part due to the differential change in the cost of housing across cities, given the 1980 geographical distribution ($\sum_c \omega_{jc1980} (P_{c2000} - P_{c1980})$). The change in cost of housing of college graduates *relative* to high school graduates is therefore the difference of these two components for college graduates and high school graduates.

Empirically, I find that both factors are important. About 43% of the total increase in cost of housing of college graduates relative to high school graduates is due to the first component (geographical mobility of college graduates toward expensive cities), and 57% is due to the second component (larger cost increase in cities that have many college graduates in 1980).

3 Nominal and Real Wage Differences

In this Section, I estimate how much of the increase in nominal wage differences between college graduates and high school graduates is accounted for by differences in the cost of living. In particular, in Section 3.1 I show estimates of the college premium in nominal and real terms. In Sections 3.2 and 3.3 I discuss whether my estimates are biased by the presence of unobserved worker characteristics or unobserved housing characteristics. In Section 3.4 I show estimates of the college premium in real terms based on an alternative local CPI that varies not just by metropolitan area, but also by skill level within metropolitan area.

3.1 Main Estimates

Model 1 in the top panel of Table 4 estimates the conditional nominal wage difference between workers with a high school degree and workers with college or more, by year. Estimates in columns 1 to 4 are from a regression of the log nominal hourly wage on an indicator for college interacted with an indicator for year 1980, an indicator for college interacted with an indicator for year 1990, an indicator for college interacted with an indicator for year 2000, years dummies, a cubic in potential experience, and dummies for gender and race.

Estimates in columns 5 to 8 are from models that also include MSA fixed effects. Entries are the coefficients on the interactions of college and year and represent the conditional wage difference for the relevant year. The sample includes all US born wage and salary workers aged 25-60 who have worked at least 48 weeks in the previous year.²²

My estimates in columns 1 to 4 indicate that the conditional nominal wage difference between workers with a high school degree and workers with college or more has increased significantly. The difference is 40% in 1980 and rises to 60% by 2000. Column 4 indicates that this increase amounts to 20 percentage points. This estimate is generally consistent with the previous literature (see, for example, Table 3 in Katz and Autor, 1999).

Models 2 and 3 in Table 4 show the conditional real wage differences between workers with a high school degree and workers with college or more. To quantify this difference, I estimate models that are similar to Model 1, where the dependent variable is the nominal wage divided by Local CPI 1 (in Model 2) or by Local CPI 2 (in Model 3). Two features are noteworthy. First, the level of the conditional college premium is lower in real terms than in nominal terms in each year. For example, in 2000 the conditional difference between the wage for college graduates and high school graduates is .60 in nominal terms and only .53 in real terms when Local CPI 1 is used as deflator. The difference is smaller—.51 percentage points—when Local CPI 2 is used as deflator. Second, the increase between 1980 and 2000 in college premium is significantly smaller in real terms than in nominal terms. For example, using Local CPI 1, the 1980-2000 increase in the conditional real wage difference between college graduates and high school graduates is 15 percentage points. In other words, cost of living differences as measured by Local CPI 1 account for 25% of the increase in conditional inequality between college and high school graduates between 1980 and 2000 (column 4).

The effect of cost of living differences is even more pronounced when the cost of living is measured by Local CPI 2. In this case, the increase in the conditional real wage difference between college graduates and high school graduates is 14 percentage points. This implies that cost of living differences as measured by Local CPI 2 account for 30% of the increase in conditional wage inequality between college and high school graduates between 1980 and 2000 (column 4).

When I control for fixed effects for metropolitan areas in columns 5-8, the nominal college premium is slightly smaller, but the real college premium is generally similar. The increase in the college premium is 18 percentage points when measured in nominal terms, and 14-15 percentage points when measured in real terms, depending on whether CPI 1 or CPI 2 is used as deflator. After conditioning on MSA fixed effects, cost of living differences account 22% of the increase in conditional inequality between college and high school graduates between 1980 and 2000 when CPI 2 is used as a deflator (column 8).

²²The sample includes both men and women. This may be a concern, since in a recent paper by Black et al. (2010) shows that female labor force participation is different in different cities. At the end of this subsection, I discuss a number of alternative specifications, including one when I estimate the college premium for men and women separately. Estimates by gender are similar to those obtained from the pooled sample.

In Tables 5 and 6 I present the results from several alternative specifications. I begin in the top panel of Table 5 by showing estimates where I deflate nominal wages based on local CPI's that measure housing costs using the average price of owner occupied houses instead of average rental costs. In particular, as discussed in Section 2.2, I measure local housing prices by taking the average reported property value of all 2 or 3 bedroom single family owner occupied houses in the relevant MSA. In the second panel, I compute Local CPI 2 using the Accra dataset described above to measure local variation in non-housing prices. (See Section 2.2 for details). In the third panel, I compute the Local CPI's allowing for the expenditure share of housing and non-housing goods to vary by metropolitan areas and skill level. (See Appendix 1 for more details). In the bottom panel, I consider the possibility that commuting distance may vary differentially for high school and college graduates. For example, it is possible that increases in the number of college graduates in some cities lead high school graduates to live farther away from job locations. To account for possible differential changes in commuting times, I re-estimate the baseline model where the dependent variable is wage per hour worked or spent commuting. In the baseline estimates, I calculate hourly wage by taking the ratio of weekly or monthly earnings over the sum of number of hours worked. By contrast, here I calculate hourly wage by taking the ratio of weekly or monthly earnings over the sum of number of hours worked *plus time spent commuting*.

In the top panel of Table 6, I show estimates based on a sample that includes all wage and salary workers 25-60, irrespective of the number of weeks worked in the previous year. In the middle panel, I show estimates that include workers born outside the US. In the bottom panel I drop rural workers (i.e. those who are not assigned an MSA).

In general, estimates in Tables 5 and 6 are not very different from the baseline estimates in Table 4. The inclusion of workers with less than 48 weeks of work results in a slightly larger percent of the nominal increase in inequality being accounted for by differences in cost of living. I have performed several additional robustness checks that are not reported in the Table due to space limitations and that are generally consistent with the estimates reported in the Table.²³

²³For example, when I allow for the effect of experience, race, and gender to vary over time by controlling for the interaction of year with gender, race and a cubic in experience, results are similar to Table 4. When I estimate separate models for male and females, results are generally similar. When I estimate separate models for workers with less than 20 years of experience and workers with more than 20 years of experience, I find that the college premium seems to be smaller, and to have grown less—both in nominal and real terms—for workers with higher levels of potential experience. Estimates where the dependent variable is the log of weekly or yearly earnings are also generally consistent with Table 4. Finally, my estimates are not very sensitive to the exclusion of outliers (defined as the top 1% and the bottom 1% of each year's wage distribution).

3.2 Worker Ability

One might be concerned about unobserved differences in worker ability. Models in Tables 4 and 5 control for standard demographics, but not for worker ability. Ability of college graduates and high school graduates is likely to vary across metropolitan areas (Combes, Duranton and Gobillon, 2008). Note that what may cause bias is not the mere presence of cross-sectional differences across cities in the relative average ability of college graduates and high school graduates. My estimates of the change in college premium in real terms are biased if the change over time in the average ability of college graduates relative to high school graduates in a given city is systematically related to changes over time in cost of living in that city. The direction of the bias is a priori not obvious. If the average unobserved ability of college graduates relative to high school graduates grows more (less) in expensive cities compared to less expensive cities, then the estimates of the real college premia in Table 4 are biased downward (upward).

While I can not completely rule out the possibility of unmeasured worker differences, in Figure 3 I provide some evidence on the relationship between one measure of worker ability and housing costs. Specifically, I use NLSY data to relate the difference in average AFQT scores between college graduates and high-school graduates across metropolitan areas to the cost of housing across metropolitan areas.²⁴

The top panel in the Figure shows average cost of renting a 2 or a 3 bedroom apartment in 1980 on the x-axis against the difference between college graduates and high school graduates in average AFQT score percentiles on the y-axis, across metropolitan areas. The level of observation is a metropolitan area. The size of the bubbles reflects the size of the metropolitan areas. Not surprisingly, the Figures shows that in most metropolitan areas college graduates have significantly higher average AFQT score than high school graduates. However, this difference does not appear to be systematically associated with housing costs. A weighted regression of the difference between college graduates and high school graduates in average AFQT scores on the average cost of renting a 2 or a 3 bedroom apartment yields a coefficient equal to .0203 (.0274).

The bottom panel of the Figure shows the same relationship in changes over time. Specifically, the graph shows the 1980-1990 change in average cost of renting a 2 or a 3 bedroom apartment and the 1980-1990 change in the difference between college graduates and high school graduates in average AFQT scores. A weighted regression yields a coefficient equal to .0010 (.0131).

In sum, the Figure indicates that both in a cross section of cities, as well as in changes over time for the same city, differences in ability between skill groups are generally orthogonal to housing costs. This finding is consistent with the evidence in Glaeser and Mare (2001).

²⁴My data contain AFQT score percentiles in 1980 and 1989. I merge these data with Census data on housing costs for 1980 and 1990. Like in Section 3.1, housing costs are measured using the average cost of renting a 2 or 3 bedroom apartment in the relevant MSA. I do not have AFQT scores in 2000.

3.3 Housing Quality

A second concern is the possibility that the changes in housing costs faced by skilled and unskilled workers reflect not just changes in cost of living, but also differential changes in the quality of housing. This could bias my estimates of the relative increase in the cost of living experienced by different skill groups, although the direction of the bias is not a priori obvious. On the one hand, the relative increase in the cost of housing experienced by college graduates may be overestimated if apartments in cities with many college graduates are subject to more quality improvements between 1980 and 2000 than apartments in cities with many high school graduates. In this case part of the additional increase in the rental cost in cities with many college graduates relative to cities with many high school graduates reflects differential quality improvements. Take, for example, features like the presence of a fireplace, or quality of the kitchen and bathrooms. If these features have improved more in cities with many college graduates, I may be overestimating the relative increase in cost of living experienced by college graduates.

On the other hand, the relative increase in the cost of housing faced by college graduates may be underestimated if apartments in cities with many high school graduates experience more quality or size improvements. Take, for example, features like the size of an apartment²⁵, or the availability of a garden, a garage, or a porch. The average apartment in New York or San Francisco is likely to be smaller than the average apartment in Houston or Indianapolis and it is also less likely to have a garden, a garage or a porch. Moreover, these features are less likely to have increased between 1980 and 2000 in New York or San Francisco than in Houston or Indianapolis. Since the share of college graduates has increased more in denser and more expensive cities, the true change in quality-adjusted per-square-foot price faced by college graduates can in principle be larger than the one that I measure.

While I can not completely rule out the possibility of unmeasured quality differences, here I present evidence based on a rich set of observable quality differences. I use data from the American Housing Survey, which includes richer information on housing quality than the Census of Population. Available quality variables include exact square footage, number of rooms, number of bathrooms, indicators for the presence of a garage, a usable fireplace, a porch, a washer, a dryer, a dishwasher, outside water leaks, inside water leaks, open cracks in walls, open cracks in ceilings, broken windows, presence of rodents, and a broken toilet in the last 3 months.²⁶

I begin by reproducing the baseline estimates that do not control for quality. Nominal estimates based on the American Housing Survey in the top panel of Table 7 are generally

²⁵Although my measure of housing cost is the average rent for apartments with a fixed number of bedrooms, exact square footage may vary.

²⁶Each year, the American Housing Survey has a sample size that is significantly smaller than the sample size in the Census. To increase precision, instead of taking only 1980, 1990 and 2000, I group years 1978-1984, 1988-1992 and 1998-2002 together.

similar to the corresponding baseline estimates based on the Census reported in Table 4.²⁷ These estimates indicate that the nominal college premium increases by 19 percentage points between 1980 and 2000. In the middle panel I estimate the real college premium, without controlling for housing quality. Finally, in the bottom panel I re-estimate the same model holding constant all available measures of housing quality. As before, I measure housing cost using the rental price for renters. But, unlike before, I first regress housing costs on the vector of observable housing characteristics. The residual from this regression represents the component of the cost of housing that is orthogonal to my measures of dwelling quality. The bottom panel of Table 7 shows how the baseline estimates change when I use the properly renormalized residual as a measure of housing cost in my local CPI 1 and CPI 2. The comparison of the middle and the bottom panels suggests that the 1980-2000 increase in real college premium estimated controlling for quality is *smaller* than the corresponding increase in the real college premium estimated without controlling for quality. Specifically, column 4 indicates that the increase in real college premium estimated controlling for quality is 15 percentage points. The corresponding estimate that does not control for quality is 16 percentage points.

In sum, though I can not completely rule out the possibility of unmeasured quality differences, Table 7 indicates that controlling for a rich vector of observable quality differences results in differences between nominal and real college premium that are slightly larger than the baseline differences. This result is consistent with estimates in Malpezzi, Chun and Green (1998).

3.4 An Alternative Measures of Local Cost of Living

My estimates in Section 3.1 are based on a definition of cost of living where the housing component of cost of living varies only by metropolitan area. In Appendix Table A1 I show how my estimates change when an alternative definition of cost of living is adopted. In particular, I allow for the cost of housing experienced by different individuals to vary depending not just on their city of residence, but also on their education level, family structure and race. The idea is that, within a city, not all households necessarily use the same type of housing. Allowing for the cost of housing faced by different demographic groups in a given city to be different may matter if tastes and budget constraints differ across groups, so that the type of housing that is used by some demographic groups in a city is not identical to the one that is used by other groups. In this case, the group-specific rental cost is measured as the predicted value from a regression of rental cost on identifiers for metropolitan area, education group, number of children, race and interactions, where the regression is estimated on the sample of renters of 2 or 3 bedroom apartments and the predicted values are calculated

²⁷Unlike Table 4, the dependent variable here is log of yearly earnings. In the American Housing Survey there is less information on number of hours worked than in the Census. Since college graduates work longer hours, the estimated nominal college premium is slightly smaller than in Table 4.

for all households. Local CPI 3 only uses local variation in cost of living that arises from variation in predicted cost of housing. Local CPI 4 uses local variation both in predicted cost of housing and cost of non housing good and services. Estimates in Appendix Table 1 indicate that, relative to Table 4, a larger share of the increase in nominal wage differences appears to be accounted for by cost of living differences.²⁸

4 A Simple Framework

In the previous Section, I have shown that over the 1980-2000 period, real wage inequality has grown less than nominal wage inequality. Does this mean that the large increases in nominal inequality have not translated into large increases in well being inequality? Not necessarily. If amenities differ across cities, changes in real wages do not necessarily equal changes in well-being. In this Section, I use a simple general equilibrium model to investigate the implications of my empirical findings for changes in well-being disparities. The implications are different depending on the reasons for the increase in the share of college graduates in expensive cities. I consider two alternative explanations for such an increase.

1. First, it is possible that skilled workers move to expensive cities because the relative demand of skilled labor increases in expensive cities, as firms located in these cities increasingly seek to hire skilled labor. This can be due to localized skill-biased technical change or positive shocks to the demand faced by industries that employ skilled workers and are located in expensive cities (for example, high tech, finance, etc.). In this case, the increase in utility disparity between skilled and unskilled workers is smaller than the increase in nominal wage disparity, because the higher nominal wage of skilled workers is in part off-set by higher cost of living in the cities where skilled jobs are located.
2. Alternatively, it is possible that skilled workers move to expensive cities because the relative supply of skilled labor increases in expensive cities, as skilled workers are increasingly attracted by amenities located in those cities. In this case, a higher cost of housing reflects consumption of desirable local amenities. Since this consumption arguably generates utility, it is possible to have large increases in utility disparities even when increases in real wage disparities are limited.

To formalize these two alternative hypotheses, and what they imply for inequality in utility and wages, I consider a simple general equilibrium model of the labor and housing market. The model is a generalization of the Roback (1982, 1988) model and has two types of workers, skilled workers (type H) and unskilled workers (type L). Like in Roback, workers and firms are mobile and choose the location that maximizes utility or profits. But unlike

²⁸An obvious concern is the possibility of differential changes in the unmeasured quality of housing for college graduates and high school graduates within a city. I have repeated the analysis of Table 7 and found results that are generally similar.

Roback, the elasticity of local labor supply is not infinite, so that productivity and amenity shocks are not always fully capitalized into land prices. This allows shocks to the relative demand and relative supply of skilled workers to have different effects on the utility of skilled and unskilled workers.

For simplicity of exposition, I model the two explanations as mutually exclusive. In the empirical tests that seek to distinguish between the two explanations (Section 5), I allow for the possibility that both demand and supply forces are at play at the same time.

4.1 Assumptions and Equilibrium

I assume that each city is a competitive economy that produces a single output good y which is traded on the international market, so that its price is the same everywhere and set equal to 1. Like in Roback, I abstract from labor supply decisions and I assume that each worker provides one unit of labor, so that local labor supply is only determined by workers' location decisions. The indirect utility of skilled workers in city c is assumed to be

$$U_{Hic} = w_{Hc} - r_c + A_{Hc} + e_{Hic} \quad (2)$$

where w_{Hc} is the nominal wage in the city; r_c is the cost of housing; A_{Hc} is a measure of local amenities. The random term e_{Hic} represents worker i idiosyncratic preferences for location c . A larger e_{Hic} means that worker i is particularly attached to city c , holding constant real wage and amenities. For example, being born in city c or having family in city c may make city c more attractive to a worker. Similarly, the indirect utility of unskilled workers is

$$U_{Lic} = w_{Lc} - r_c + A_{Lc} + e_{Lic} \quad (3)$$

In equations 2 and 3, skilled and unskilled workers in a city compete for housing in the same housing market and therefore face the same price of housing. This allows a shock to one group to be transmitted to the other group through its effect on housing prices.²⁹ While they have access to the same local amenities, different skill groups do not need to value these amenities equally: A_{Hc} and A_{Lc} represent the skill-specific value of local amenities.

Assume that there are two cities—Detroit (city a) and San Francisco (city b)—and a fixed number of workers is divided between the two cities. Tastes for location can vary by skill group. Specifically, skilled workers' and unskilled workers' relative preferences for city a over city b are, respectively

$$e_{Hia} - e_{Hib} \sim U[-s_H, s_H] \quad (4)$$

and

$$e_{Lia} - e_{Lib} \sim U[-s_L, s_L] \quad (5)$$

²⁹It is easy to relax this assumption by assuming some residential segregation by skill level within a city.

The parameters s_H and s_L characterize the importance of idiosyncratic preferences for location and therefore the degree of labor mobility. If s_H is large, for example, it means that preferences for location are important for skilled workers and therefore their willingness to move to arbitrage away real wage differences or amenity differences is limited. On the other hand, if s_H is small, preferences for location are not very important and therefore skilled workers are more willing to move in response to differences in real wages or amenities. In the extreme, if $s_H = 0$ skilled workers' mobility is perfect.

A worker chooses city a if and only if $e_{ia} - e_{ib} > (w_b - r_b) - (w_a - r_a) + (A_b - A_a)$. In equilibrium, the marginal worker needs to be indifferent between living in Detroit and San Francisco. This implies that skilled workers' labor supply is upward sloping, with the slope that depends on s . For example, the supply of skilled workers in San Francisco is:

$$w_{Hb} = w_{Ha} + (r_b - r_a) + (A_a - A_b) + s_H \left(\frac{N_{Hb} - N_{Ha}}{N} \right) \quad (6)$$

where N_{Hb} is the log of the number of skilled workers hired in San Francisco and $N = N_{Ha} + N_{Hb}$. If idiosyncratic preferences for location are not very important (s_H is small), then workers are very mobile and the supply curve is relatively flat. If idiosyncratic preferences for location are very important (s_H is large), then workers are rather immobile and the supply curve is relatively steep. Moreover, an increase in the real wage in Detroit, or an improvement in relative amenities shifts back the labor supply curve in San Francisco.³⁰

For simplicity, I focus on the case where skilled and unskilled workers in the same city work in different firms. This amounts to assuming away imperfect substitution between skilled and unskilled workers. This assumption simplifies the analysis, and it is not crucial (Moretti, 2010). The production function for firms in city c that use skilled labor is Cobb-Douglas with constant returns to scale: $\ln y_{Hc} = X_{Hc} + hN_{Hc} + (1 - h)K_{Hc}$, where K_{Hc} is the log of capital and X_{Hc} is a skill and city-specific productivity shifter. Firms are assumed to be perfectly mobile. If firms are price takers and labor is paid its marginal product, labor demand for skilled labor in city c is

$$w_{Hc} = X_{Hc} - (1 - h)N_{Hc} + (1 - h)K_{Hc} + \ln h \quad (7)$$

The labor market for unskilled workers is similar. I assume that there is an international capital market, and that capital is infinitely supplied at a given price i .³¹

Each worker consumes one unit of housing, so that demand for housing is determined by the number of skilled and unskilled workers in a city. Specifically, the the local demand for

³⁰An important difference between the Rosen-Roback setting and this setting is that in Rosen-Roback, all workers are identical, and always indifferent across locations. In this setting, workers differ in their preferences for location. While the marginal worker is indifferent between locations, here there are inframarginal workers who enjoy economic rents. These rents are larger the smaller the elasticity of local labor supply.

³¹In equilibrium demand for capital is equal to its supply and marginal product of capital is the same for firms that use skilled labor and those that use unskilled labor: $X_{Hc} - hK_{Hc} + hN_{Hc} + \ln(1 - h) = \ln i X_{Lc} - hK_{Lc} + hN_{Lc} + \ln(1 - h) = \ln i$.

housing is the sum the demand of skilled workers and the demand of unskilled workers. For example, in city b :

$$r_b = \frac{(2s_H s_L)}{(s_H + s_L)} - \frac{(2s_H s_L)(N_{Hb} + N_{Lb})}{N(s_H + s_L)} - \frac{s_L(w_{Ha} - w_{Hb} - r_a)}{(s_L + s_H)} - \frac{s_H(w_{La} - w_{Lb} - r_a)}{(s_L + s_H)} \quad (8)$$

To close the model, I assume that the supply of housing is

$$r_c = z + k_c N_c \quad (9)$$

where $N_c = N_{Hc} + N_{Lc}$ is the number of housing units in city c , which is the same as the number of workers. The parameter k_c characterizes the elasticity of the supply of housing. I assume that this parameter is exogenously determined by geography and local land regulations. In cities where geography and regulations make it easy to build new housing, k_c is small. In the extreme case where there are no constraints to building new houses, the supply curve is horizontal, and k_c is zero. In cities where geography and regulations make it difficult to build new housing, k_c is large. In the extreme case where it is impossible to build new houses, the supply curve is vertical, and k_c is infinite.³²

In period 1, the two cities are assumed to be identical. Equilibrium in the labor market is obtained by equating equations 6 and 7 for each city. Equilibrium in the housing market is obtained by equating equations 8 and 9. I consider two scenarios for period 2. In the first scenario, the relative demand of skilled workers increases in one of the two cities (Section 4.2). In the second scenario, the relative supply of skilled workers increases in one of the two cities (Section 4.3). The implications of the two scenarios for the empirical analysis are summarized in Section 4.4.

4.2 Increase in the Relative Demand of Skilled Labor

Here I consider the case where the productivity of skilled workers increases relative to the productivity of unskilled workers in San Francisco. Nothing happens to the productivity of unskilled workers in San Francisco and the productivity of skilled and unskilled workers in Detroit. In other words, the relative demand for skilled labor increases in San Francisco. The amenities in the two cities are identical and fixed. Formally, I assume that in period 2, the productivity shifter for skilled workers in San Francisco is higher than in period 1: $X_{Hb2} = X_{Hb1} + \Delta$, where $\Delta > 0$ represents a positive, localized, skill-biased productivity shock. I have added subscripts 1 and 2 to denote periods 1 and 2. The dot-com boom

³²A limitation of equation 9 is housing production does not involve the use of any local input. Roback (1982) and Glaeser (2008), among others, discuss spatial equilibrium in the case where housing production involves the use of local labor and other local inputs. Moreover, equation 9 ignores the durability of housing. Glaeser and Gyourko (2001) point out that once built, the housing stock does not depreciate quickly and this introduces an asymmetry between positive and negative demand shocks. In particular, when demand declines, the quantity of housing cannot decline, at least in the short run.

experienced by the San Francisco Bay Area is arguably an example of such a localized skill biased shock. Driven by the advent of the Internet and the agglomeration of high tech firms in the area, the demand for skilled workers increased significantly (relative to the demand for unskilled workers) in San Francisco in the second half of the 1990s.³³

Because skilled workers in San Francisco have become more productive, their nominal wage increases by an amount Δ/h , proportional to the productivity increase. Attracted by this higher productivity, some skilled workers leave Detroit and move to San Francisco. Following this inflow of skilled workers, the cost of housing in San Francisco increases by

$$r_{b2} - r_{b1} = \frac{s_L N k_b \Delta}{h(k_a N s_H + 2s_H s_L + k_a N s_L + k_b N s_H + k_b N s_L)} \geq 0 \quad (10)$$

In Detroit, the cost of housing declines by the same amount because of out-migration. In San Francisco, real wages of skilled workers increase by

$$(w_{Hb2} - r_{b2}) - (w_{Hb1} - r_{b1}) = \frac{k_a N s_H + k_b N s_H + k_a N s_L + 2s_H s_L}{h(k_a N s_H + 2s_H s_L + k_a N s_L + k_b N s_H + k_b N s_L)} \Delta \geq 0 \quad (11)$$

It is easy to see that, because of the increased cost of housing, the increase in real wages is smaller than the increase in nominal wages Δ/h . Moreover, this increase in the real wage of skilled workers is larger the more elastic is housing supply in San Francisco (small k_b). Intuitively, a more elastic housing supply implies a smaller increase in housing prices in San Francisco, and therefore a larger increase in real wage, for a given increase in nominal wage. The increase in the real wage of skilled workers is also larger the smaller the elasticity of local labor supply of skilled workers (large s_H). Intuitively, lower elasticity of labor supply implies less mobility. With less mobility, a larger fraction of the benefit of the productivity shocks is capitalized in real wages. In the extreme case of no mobility, ($s_H = \infty$), the entire productivity shock is capitalized in the real wage of skilled workers. The increase in the real wage of skilled workers is larger the larger the elasticity of local labor supply of skilled workers (small s_L). A higher elasticity of labor supply of unskilled workers implies that a larger number of unskilled workers move out in response to the inflow of skilled workers, so that the increase in housing costs is more limited.

In Detroit nominal wages don't change and housing costs decline, so that real wages for skilled workers increase by

$$(w_{Ha2} - r_{a2}) - (w_{Ha1} - r_{a1}) = \frac{s_L k_a N}{h(k_a N s_H + 2s_H s_L + k_a N s_L + k_b N s_H + k_b N s_L)} \Delta \geq 0 \quad (12)$$

Although the shock has increased productivity only in one city, the equilibrium real wages of skilled workers increase in both cities because of mobility. By comparing equation

³³Beaudry, Doms and Lewis (2008) argue that over the past 30 years, technological change resulted in increases in the productivity of skilled workers in cities that already had many skilled workers. These cities also happen to be cities with a higher than average initial share of college graduates and cost of housing. See also Berry and Glaeser (2005).

11 with 12, it is easy to see that the increase in real wages in the city directly affected by the productivity shock (San Francisco) is larger than the increase in real wages in the city not affected by the productivity shock (Detroit): $(w_{Hb2} - r_{b2}) - (w_{Hb1} - r_{b1}) \geq (w_{Ha2} - r_{a2}) - (w_{Ha1} - r_{a1})$. This is not surprising. While labor mobility causes real wages to increase in Detroit following a shock in San Francisco, real wages are not fully equalized because mobility is not perfect and only the marginal worker is indifferent between the two cities in equilibrium. With perfect mobility ($s_H = 0$), real wages are completely equalized.

What happens to the wage of unskilled workers? Because their productivity is fixed, their nominal wage does not change. However, housing costs increase in San Francisco and decline in Detroit. As a consequence, the real wage of unskilled workers in San Francisco decreases by

$$(w_{Lb2} - r_{b2}) - (w_{Lb1} - r_{b1}) = -\frac{s_L N k_b}{h(k_a N s_H + 2s_H s_L + k_a N s_L + k_b N s_H + k_b N s_L)} \Delta \leq 0 \quad (13)$$

Effectively, unskilled workers compete for scarce housing with skilled workers, and the inflow of new skilled workers in San Francisco hurts inframarginal unskilled workers through higher housing costs. Marginal unskilled workers leave San Francisco, since their real wage is higher in Detroit. Inframarginal unskilled workers (those who have a strong preference for San Francisco over Detroit) opt to stay in San Francisco, even if their real wage is lower. For the same reason, the real wage and utility of inframarginal unskilled workers in Detroit increases:

$$(w_{La2} - r_{a2}) - (w_{La1} - r_{a1}) = \frac{s_L N k_a}{h(k_a N s_H + 2s_H s_L + k_a N s_L + k_b N s_H + k_b N s_L)} \Delta \geq 0 \quad (14)$$

The equilibrium number of skilled workers increases in San Francisco, while the equilibrium number of unskilled workers decreases. Changes in employment in Detroit are exactly specular, by assumption. On net, the overall population of San Francisco increases because the number of skilled workers who move in is larger than the number of unskilled workers who leave.³⁴

³⁴In particular, the number of skilled workers in San Francisco increases by

$$N_{Hb2} - N_{Hb1} = \frac{\Delta N((k_a + k_b)N + 2s_L)}{2h(k_a N(s_H + s_L) + k_b N(s_H + s_L) + 2s_H s_L)} \geq 0 \quad (15)$$

The number of unskilled workers declines by

$$N_{Lb2} - N_{Lb1} = -\frac{N^2(k_a + k_b)}{2h(k_a N s_H + 2s_H s_L + k_a N s_L + k_b N s_H + k_b N s_L)} \Delta \leq 0 \quad (16)$$

San Francisco population increases by

$$(N_{Hb2} + N_{Lb2}) - (N_{Hb1} + N_{Lb1}) = \frac{\Delta N s_L}{h(k_a N(s_H + s_L) + k_b N(s_H + s_L) + 2s_H s_L)} \geq 0 \quad (17)$$

The productivity shock creates winners and losers. Skilled workers in both cities and landowners in San Francisco benefit from the productivity increase. Inframarginal unskilled workers in San Francisco are negatively affected, and inframarginal unskilled workers in Detroit are positively affected.³⁵ The exact magnitude of the changes in utility for skilled and unskilled workers and for landowners crucially depends on which of the three factors—skilled labor, unskilled labor or land—is supplied more elastically at the local level. Specifically, the incidence of the shock depends on the elasticities of labor supply of the two groups (which are governed by the preference parameters s_H and s_L) and the elasticities of housing supply in the two cities (which are governed by the parameters k_a and k_b). Moretti (forthcoming) provides detailed discussion of the incidence and welfare consequences of relative demand shocks.

The model also illustrates that a non-degenerate equilibrium is possible. After a shock that makes one group more productive, both groups are still represented in both cities. This conclusion hinges upon the assumption of a less than infinite elasticity of local labor supply.³⁶ Firms are indifferent between cities because they make the same profits in both cities. While labor is now more expensive in San Francisco, it is also more productive there. Because firms produce a good that is internationally traded, if skilled workers weren't more productive, employers would leave San Francisco and relocate to Detroit.³⁷

4.3 Increase in the Relative Supply of Skilled Labor

In the case of demand pull described above, the number of skilled workers in San Francisco increases because the relative demand of skilled workers increases. I now turn to the opposite case, where the number of skilled workers in San Francisco increases because the relative supply of skilled workers in San Francisco increases.

Specifically, I consider what happens when San Francisco becomes relatively more desirable for skilled workers compared to Detroit. I assume that in period 2, the amenity level

³⁵Although inframarginal unskilled workers in San Francisco are made worse off by the decline in their real wage, they are still better off in San Francisco than in Detroit because of their preference for San Francisco.

³⁶In the absence of individual preferences for location, no unskilled worker would remain in San Francisco and the equilibrium would be characterized by complete geographic segregation of workers by skill level. This is not realistic, since in reality we never observe cities that are populated by workers of only one type.

³⁷An assumption of this model is that skilled and unskilled workers are employed by different firms, so that the labor market is segregated by skill within a city. This assumption effectively rules out imperfect substitutability between skilled and unskilled labor. In a more general setting, skilled and unskilled workers work in the same firm. The qualitative results generalize, but the equilibrium depends on the degree of imperfect substitution between skilled and unskilled labor. Specifically, complementarity between skilled and unskilled workers implies that the marginal product of unskilled workers increases in the number of skilled workers in the same firm. Thus, the inflow of skilled workers in city b caused by the increase in their productivity endogenously raises the productivity of unskilled workers in city b . As a consequence, the real wage of unskilled workers declines less than in the case described above. This mitigates the negative effect on the welfare of unskilled workers in city b and it reduces the number of unskilled workers who leave the city.

increases for skilled workers in San Francisco: $A_{Hb2} = A_{Hb1} + \Delta'$, where $\Delta' > 0$ represents the improvement in the amenity. I assume that the productivity of both skilled and unskilled workers, as well as the amenity level in Detroit, do not change.³⁸

Unlike the case of demand, here the nominal wage of skilled workers in San Francisco and Detroit remains unchanged.³⁹ Attracted by the better amenity, some skilled workers move from Detroit to San Francisco and some unskilled workers leave San Francisco to Detroit.⁴⁰ On net, the population in San Francisco increases by

$$(N_{Hb2} + N_{Lb2}) - (N_{Hb1} + N_{Lb1}) = \frac{\Delta' N s_L}{h(k_a N (s_H + s_L) + k_b N (s_H + s_L) + 2s_H s_L)} \geq 0 \quad (18)$$

As a consequence, housing costs in San Francisco increase by

$$r_{b2} - r_{b1} = \frac{s_L N k_b \Delta'}{h(k_a N s_H + 2s_H s_L + k_a N s_L + k_b N s_H + k_b N s_L)} \geq 0 \quad (19)$$

and decline in Detroit by

$$r_{a2} - r_{a1} = -\frac{s_L N k_a \Delta'}{h(k_a N s_H + 2s_H s_L + k_a N s_L + k_b N s_H + k_b N s_L)} \leq 0 \quad (20)$$

Real wages of skilled workers in San Francisco decline by an amount equal to equation 19 (with a minus sign in front). This reflects the compensating differential for the better amenity in San Francisco. Real wages of skilled workers in Detroit increase by an amount equal to equation 20 (with a minus sign in front).

Similarly, the real wage for unskilled workers in San Francisco declines by

$$(w_{Lb2} - r_{b2}) - (w_{Lb1} - r_{b1}) = -\frac{s_L N k_b}{h(k_a N s_H + 2s_H s_L + k_a N s_L + k_b N s_H + k_b N s_L)} \Delta' \leq 0 \quad (21)$$

and it increases in Detroit.

³⁸For simplicity, I have assumed that supply shocks are driven by increases in amenities for given tastes. Glaeser and Tobio (2007) have a model that makes a similar assumption. Alternatively I could assume that (i) amenities are fixed, but the taste for those amenities increase; or (ii) both amenities and tastes are fixed, but amenities are a normal good so that college graduates consume more of them than high school graduates (Gyourko, Mayer, and Sinai, 2006).

³⁹This may be surprising at first. While one might expect wage *increases* in response to demand increases (indeed, this is what happens in subsection 4.2), one might expect wage *decreases* in response to supply increases. Why do nominal wages not decline in San Francisco? The reason is that in a model with capital, nominal wages do not move in San Francisco because capital flows to San Francisco and leaves Detroit, offsetting the changes in labor supply in the two cities. (In a model without capital nominal wages do decline.)

⁴⁰Specifically, the number of skilled workers who move to San Francisco is equal to $\frac{\Delta' N ((k_a + k_b) N + 2s_L)}{2h(k_a N (s_H + s_L) + k_b N (s_H + s_L) + 2s_H s_L)} \geq 0$. The number of unskilled workers who move to Detroit is equal to $\frac{\Delta' N^2 (k_a + k_b)}{2h(k_a N (s_H + s_L) + k_b N (s_H + s_L) + 2s_H s_L)} \geq 0$.

Like for the case of demand shocks, a supply shock generates winners and losers. Here inframarginal skilled workers benefit from the improvement in amenities. While the utility gain is larger for inframarginal skilled workers in San Francisco, inframarginal skilled workers in Detroit are also made better off, even if there is no change in amenity there. On the other hand, inframarginal unskilled workers in San Francisco are made worse off by the increase in housing prices. Similarly, inframarginal unskilled workers in Detroit are made better off by the decline in local housing prices.

4.4 Implications for Inequality in Wages and Utility

The model has three implications that are useful in guiding the interpretation of the empirical findings.

(A) First, the model clarifies the relationship between changes in relative wages and changes in relative utility in the two scenarios. The analysis in Sections 4.2 and 4.3 suggests that *for a given nation-wide increase in the nominal wage gap between skilled and unskilled workers*, the demand pull hypothesis implies a more limited increase in utility inequality, while the supply push hypothesis implies a larger increase in utility inequality.

More specifically, in the demand pull scenario the nominal wage difference between skilled and unskilled workers averaged across the two cities increases.⁴¹ The utility difference between skilled and unskilled workers averaged across the two cities also increases, but by an amount smaller than the increase in the nominal wage gap. It is possible to show that the larger is the increase in housing costs experienced by skilled workers relative to unskilled workers, the smaller is the increase in average utility experienced by skilled workers relative to unskilled workers.⁴²

The intuition is simple. The benefits of a higher nominal wage for skilled workers are in part eroded by the higher cost of housing in the cities where the new skilled jobs are created. Thus, the relative utility of skilled workers does not increase as much as their relative nominal wage. Put differently, if college graduates move to expensive cities like San Francisco and New York because of increases in the relative demand for college graduates in

⁴¹This average is a weighted average reflecting the size of the two cities.

⁴²To formally see this, consider the population-weighted average across the two cities of the change in the skilled-unskilled nominal wage difference and compare it with the population-weighted average across the two cities of the change in the skilled-unskilled utility difference. In the simple case where $k_a = k_b = k$, the difference between the two is

$$\frac{Nk\Delta^2 s_L (s_L + 2kN)}{2h^2 (kNs_H + s_H s_L + kNs_L)^2} \geq 0 \quad (22)$$

which is non-negative, indicating that the relative nominal wage of skilled workers grows more than their relative utility. In the more general case where $k_a \neq k_b$, the difference between the two remains positive as long as the elasticity of housing supply in the city affected by the demand shock is not too large compared with the elasticity of housing supply in the city not directly affected by the demand shock.

these cities—and not because they particularly like living in San Francisco and New York—then part of the benefit of higher nominal wage is offset by the higher cost of living. In this case, the increase in their real wage and utility level is smaller than the increase in their nominal wage.

By contrast, in the supply push scenario, the utility difference between skilled and unskilled workers averaged across the two cities increases more than the nominal and real wage difference between skilled and unskilled workers averaged across the two cities. Intuitively, if college graduates move to expensive cities like San Francisco and New York because improvements in amenities raise the relative supply of college graduates there—and not because of labor demand—then there may still be a significant increase in utility inequality even if the increase in real wage inequality is limited. In this case, increases in the cost of living in these cities simply reflect the increased attractiveness of these cities to skilled workers and represent the price to pay for the consumption of desirable amenities.⁴³

(B) Second, the equilibrium described in subsections 4.2 and 4.3 suggests a simple empirical test to distinguish between the two cases. If relative demand shifts are responsible for the geographical reallocation of labor, we should see that in equilibrium, cities that experience large increases in the relative number of skilled workers (in the model: San Francisco) also experience increases in the relative nominal wage of skilled workers, compared to cities that experience small increases (or declines) in the relative number of skilled workers (in the model: Detroit). By contrast, if relative supply shifts are responsible for the geographical reallocation of labor, we should see that in equilibrium, cities that experience an increase in the relative number of skilled workers experience no change in the relative nominal wage of skilled workers.

One might have expected that an increase in the relative supply of factor of production in a city should cause a *decline* in its equilibrium relative price. Why in the model the nominal wage of skilled workers in San Francisco remains constant following an increase in the relative supply of skilled workers? As discussed in Section 4.3, this is due to the endogenous reaction of capital. Because capital is supplied with infinite elasticity at a fixed interest rate, nominal wages do not move in San Francisco because capital flows to San

⁴³To formally see this, note that the simple case where $k_a = k_b = k$, the population-weighted average change in the skilled-unskilled nominal wage difference minus the population-weighted change in the skilled-unskilled utility difference is equal to

$$\frac{-\Delta'(-kNs_L\Delta' + 2kNs_Hs_L + 2kNs_Hs_L + s_Hs_L + s_L\Delta's_H + k^2N^2s_L^2 + k^2N^2s_H^2 + k^2N^2s_H\Delta' + 2kNs_H^2s_L + 2Nk\Delta's_Hs_L)}{(2(kNs_H + s_Hs_L + kNs_L)^2)} \quad (23)$$

which is non-positive unless the elasticity of local labor supply of skilled workers is too small compared with the elasticity of local labor supply of unskilled workers. In the more general case where $k_a \neq k_b$, the expression is considerably more complicated, but the difference remains non-positive unless the elasticity of local labor supply of skilled workers is too small.

Francisco and leaves Detroit, thus offsetting the effect of changes in labor supply in the two cities. In a model without capital, nominal wages of skilled workers decline in San Francisco following an increase in their supply.

(C) Finally, it is important to point out that, while the focus of the paper is on inequality related to labor market outcomes, the broader welfare consequences of the demand and supply shocks depend not just on changes in relative wages, but also on which of the two education groups originally owns the land in the cities that benefit from the demand and supply shocks. In the model, some landowners benefit from the demand and supply shocks (namely those in San Francisco), while other are hurt (namely those in Detroit). The relevant empirical question in this respect is which of the two skill groups owns more of the land in the neighborhoods whose land prices are raised by the inflow of new residents in cities that experience positive skill-biased shocks and the neighborhoods that are abandoned by the outflow of residents in cities that experience negative shocks. This is an important but complicated question. A full empirical treatment of this issue is beyond the scope of this paper and is left for future research.

5 Interpreting the Evidence: Demand Pull or Supply Push?

I now present empirical evidence that seeks to determine whether relative demand or relative supply shifts—or a combination of the two—drive changes in the geographical location of different skill groups. The analysis in subsections 4.2 and 4.3 suggests that the demand pull and the supply push hypotheses have similar predictions for equilibrium housing costs: under both hypotheses, cities that experience large increases in the share of college graduates should also experience large increases in housing costs.

But the demand pull and supply push hypotheses have different predictions for wage changes. Under the demand pull hypothesis, cities that experience large increases in the share of college graduates should experience large increases in the equilibrium relative wage of college graduates. By contrast, under the supply push hypothesis, there should be no positive relationship between increases in the share of college graduates and changes in the equilibrium relative nominal wages. (See Section 4.4, part B.) Intuitively, increases in the relative demand of a factor of production in a city should result in increases in its equilibrium relative price there. Increases in the relative supply of factor of production in a city can not cause an increase in its equilibrium relative price. A similar idea is used in Katz and Murphy (1992) to explain *nationwide* changes in relative wages.⁴⁴

⁴⁴In the literature, there are several existing measures of quality-of-life differences across metropolitan areas (see, for example, Chen and Rosenthal, 2008; and Albouy, 2009). One may be tempted to use these measures to provide an additional empirical test of my conclusions by estimating the relationship between

It is important to highlight that the two hypotheses are not mutually exclusive since it is possible that cities experience both demand and supply shocks. It is also possible that relative demand shifts endogenously generate relative supply shifts, and vice versa. For example, an increase in the relative demand for skilled labor in a city may result in an increase in the number of college educated residents in that city and this in turns may result in increases in the local amenities that are attractive to college graduates, such as good schools, good theaters, good restaurants, etc. Alternatively, an increase in the supply of skilled workers in a city may generate agglomeration spillovers that lead to increases in the productivity of firms and workers in that city (Moretti 2004a, 2004b). Similarly, Desmet and Rossi-Hansberg (2009) propose a model of the localization of service and manufacturing industries with agglomeration economies that can generate both supply shifts and demand shifts.

I present two pieces of empirical evidence. First, I look at the OLS relationship between changes in the college share and changes in the college premium across US metropolitan areas. The finding of a positive coefficient indicates that relative demand shifts are important, but does not rule out the existence of relative supply shifts. Second, to shed more light on whether relative supply shifts are important, I use an instrumental variable strategy.

(1) First, in Figure 4, I show the empirical relationship between the equilibrium college share and the equilibrium college premium across US metropolitan areas, both in the 2000 cross-section and in changes between 1980 and 2000. Demand pull would predict a positive slope, while supply push would predict zero slope. Note that that the relationship in the Figure is *not causal*. Rather, it is an *equilibrium* relationship between relative number of college graduates and their relative wage. This is in contrast with earlier work, including my own, that seeks to establish the *causal* effect of increases in college share on wages and therefore estimate different specifications.⁴⁵

The Figure shows a positive association between the college share and the college premium across US metropolitan areas, both in levels as well as in changes. Columns 1 and 2

 these existing measures of local amenities and the share of college graduates across cities. However, this type of test would be difficult to implement in practice. As explained above, what matters in my framework is the (change over time in) the difference in amenities that attract skilled and unskilled workers. By contrast, the existing measures of local amenities are typically city-wide, and do not differentiate between skilled and unskilled workers.

⁴⁵For example, in Moretti (2004), I try to establish the causal effect of increases in college share on wages. The econometric specification adopted here differs from the specification there, because in Moretti (2004) the econometric model seeks to *control for* shocks to the relative demand of skilled labor. To this end, I include in the regressions as controls several variables in order to absorb changes in the relative demand for college graduates. I also use instrumental variables to further control for relative demand shocks. By contrast, in this paper, I engage in a completely different exercise. I do not seek to hold constant demand shocks. Instead, I am interested in establishing the role played by demand shocks in affecting changes in college share across cities. What I am measuring in Figure 4 and Table 8 is the relationship between the wage gap and the college share, inclusive of any human capital spillover.

in Table 8 quantify the corresponding regression coefficients. The level of observation is the metropolitan area. The dependent variable is the city-specific college premium, defined as the city-specific difference in the log of hourly wage for college graduates and high school graduates conditional on all the controls used in the regressions (a cubic in potential experience, year effects, gender and race). Models are weighted by city size. The coefficient for the specification in column 2 is positive and statistically significant: .388 (.057).

This evidence is consistent with demand factors playing a significant role in driving variation in college share across cities. This conclusion is consistent with Berry and Glaeser (2005), who argue that demand factors play a more important role than supply factors in explaining the sorting of skilled workers across US metropolitan areas.

(2) The evidence in Figure 4 and Table 8 suggests that demand factors are important, but does not rule out that supply factors are also present. As a second piece of evidence that may shed more light on whether relative supply factors play any role in driving variation in college share across cities, I use observable shocks to the relative demand of skilled labor as an instrumental variable for college share.

This IV estimate isolates the effect on the college premium of changes in the college share that are driven exclusively by changes in relative demand. Put differently, the instrumental variable estimate establishes what happens to the college premium in a city when the city experiences an increase in the number of college graduates that is driven purely by an increase in the relative demand for college graduates. By contrast, the OLS estimate above establishes what happens to the college premium in a city when the city experiences an increase in the number of college graduates that may be driven by either demand or supply shocks. The comparison of the two estimates is therefore informative about the relative importance of demand and supply shocks.

To isolate relative demand shocks, I use as an instrument the weighted average of nationwide relative employment growth by industry, with weights reflecting the city-specific employment share in those industries:

$$\text{Change in Relative Demand in City } c = \sum_s \eta_{sc} (\Delta E_{Hs} - \Delta E_{Ls}) \quad (24)$$

where η_{sc} is the share of jobs in industry s in city c in 1980; ΔE_{Hs} is the nationwide change between 1980 and 2000 in the log of number of jobs for college graduates in industry s (excluding city c); ΔE_{Ls} is a similar change for high school graduates. If relative employment of skilled workers in a given industry increases (decreases) nationally, cities where that industry employs a significant share of the labor force will experience a positive (negative) relative shock to the labor demand of skilled workers (Katz and Murphy, 1992).

The first stage relationship between demand shocks and changes in college share is shown graphically in Figure 5. The figure shows that in cities that experience an increase in the relative demand of college graduates the share of college graduates increases and the relationship appears fairly tight. The regression coefficient is .42(.02), with R^2 of .44. This

is interesting because it means that this measure of demand shocks alone accounts for almost half of the variation in the geographical location of different skill groups. Since there are other demand shocks that are not captured by the instrument, this lends indirect support to the notion that demand shocks play an important role.

The instrumental variable estimate, in column 3 of Table 8, is .371 (.106) and is remarkably close to the OLS estimate. The similarity between the OLS and the IV estimates suggests that the increase in the college premium in a city caused by a demand shock (IV estimate in column 3) is not very different from the empirical correlation between the college share and the college premium that is observed in the data (OLS estimate in column 2). In other words, most of the empirical correlation between the college share and the college premium that is observed in the data seems to be driven by demand shocks.

(3) Discussion of the role of amenities. The key finding in this section is that most of the changes over time in the geographical location of skilled workers relative to unskilled workers are driven by changes in the relative demand for skilled labor, rather than changes in its relative supply.

It is important to clarify what this finding implies for the role of amenities in worker location decisions. My finding does **not** imply that amenities do not affect worker location decisions in general. Amenities are clearly an important determinant of where people decide to live. Furthermore, my finding does **not** imply that amenities do not affect location decisions of skilled and unskilled workers differently. It is possible that the relative importance of certain amenities (cultural amenities, school quality, crime, restaurants) is different for different skill groups. What my finding implies is that the change over time in the difference between skilled and unskilled workers in relevant local amenities did not play an important role in driving differential changes in the geographical location of skilled and unskilled workers in the period 1980-2000.⁴⁶ Instead, my finding suggests that differential changes in the geographical location of skilled and unskilled workers in this period were mostly driven by geographical differences in the availability of new skilled and unskilled jobs.

6 Conclusions

Because of their different geographical distribution, college graduates and high school graduates have experienced different increases in the cost of living over the past 30 years. One contribution of this paper is to document that, as a consequence, the conditional difference between the wage of college graduates and of high school graduates is significantly lower in real terms than in nominal terms and has grown less. In 2000, the level of the college

⁴⁶This would be true, for example, if the amenities that matter for skilled and unskilled workers have changed in a similar way within each city in this period. This would also be true if the amenities that matter for skilled workers have changed differently from the amenities that matter for unskilled workers, but this differential change is similar across metropolitan areas in the US.

premium is 60% in nominal terms and only 51% in real terms. More importantly, the *increase* in the college premium between 1980 and 2000 in real terms is significantly smaller than the increase in nominal terms. Specifically, at least 22% of the documented increase in the college premium between 1980 and 2000 is accounted for by differences in the cost of living.

The implications of this empirical finding for disparities in well-being depend on the reasons for the increase in the share of college graduates in expensive cities. Using a simple general equilibrium model of the labor and housing markets, I consider two broad classes of explanations. Under a demand pull hypothesis, the relative demand of college graduates increases in expensive cities because of localized skill-biased technical change or other demand shocks. In this case, college graduates move to expensive cities because the jobs for college graduates are increasingly located in those cities, and not because they particularly like living in those cities. The increase in their utility level is smaller than the increase in their nominal wage due to higher cost of living. Under a supply push hypothesis, the relative supply of college graduates increases in expensive cities because college graduates are increasingly attracted by amenities located in those cities. The increase in the cost of living in those cities reflects the attractiveness of the cities to skilled workers and is the price for the consumption of desirable amenities. In this case, there may still be a significant increase in utility inequality even if the increase in real wage inequality is limited. Of course, the two hypotheses are not mutually exclusive and it is possible that cities experience both demand and supply shocks.

To determine whether the variation in the relative number of college graduates across cities is driven by relative demand or relative supply shocks, I analyze the equilibrium relationship between changes in college premium and changes in the share of college graduates across metropolitan areas. Consistent with demand shocks playing an important role, I find a positive association between changes in college premium and changes in college share: cities that experience large increases in the fraction of college graduates also experience large increases in the relative wage of college graduates. I also present an instrumental variable estimate obtained by instrumenting changes in college share with a measure of arguably exogenous relative demand shocks.

The weight of the evidence seems consistent with the notion that changes in the geographical location of different skill groups are mostly driven by changes in their relative demand. I conclude that the increase in well-being disparities between 1980 and 2000 is significantly smaller than we previously thought based on the existing literature.⁴⁷

It is important to clarify that the exact magnitude of the increase in well-being disparities remains unknown. While my findings suggest that the increase in the difference in well-being between skilled and unskilled workers is smaller than we may have thought, a full account

⁴⁷My results have the potential to explain at least in part an outstanding puzzle in the inequality literature. Despite the increase in the return to education, the rate of growth in the number of college graduates is still low relative to earlier periods. The fact that their real wage has not increased as much as previously thought may explain why the number of college graduates has not increased as much as one would have expected.

of well-being disparities is difficult to accomplish here for at least two reasons. First, my analysis does not take into consideration features of jobs other than wages. Hamermesh (1999) shows that the amount of workplace disamenities (such as risk of death or workplace injury) born by low skill workers increased more than the amount of workplace disamenities born by high skill workers during the 1980's. This differential change implies a larger increase in well-being inequality than the one measured ignoring workplace disamenities, although this bias is likely to be limited for the typical worker.⁴⁸

Second, this paper leaves open the question of how changes in housing wealth affect the relative welfare of skilled and unskilled homeowners. Consistent with the previous literature on inequality, the main focus this paper is on differences between skilled and unskilled workers that are caused by labor market changes. However, the broader distributional consequences of the demand and supply shocks depend not just on changes in relative wages, but also on changes in wealth, as discussed above. Changes in the price of housing have the potential to affect the relative wealth of different skill groups depending on who originally owns the land in the cities that are affected by the demand and supply shocks. A full empirical treatment of this issue is complicated and is beyond the scope of this paper.

This paper leaves open the question of what ultimately causes the local relative demand shocks. In my theoretical setting, I take these shocks as exogenous. Future research should focus on exactly what generates the localized relative demand shifts that make college graduates more productive in some parts of the country. Localized skill-biased technical change is a potential explanation, as long as it is enriched by a theory of why demand shocks occur in some cities and not in others. Beaudry, Doms and Lewis (2008) and Berry and Glaeser (2005) propose realistic models and intriguing empirical evidence. Models with human capital spillovers or agglomeration spillovers also have the potential to explain localized demand shifts (Moretti, 2004a and 2004b; Greenstone, Hornbeck and Moretti, 2007). An alternative explanation centers on shifts in product demand across industries that have different skill intensities (Buera and Kaboski, 2009). For example, industries like finance and high tech that are skill intensive and are located in expensive coastal metropolitan areas, have been expanding during the 1980s and 1990s. Future research should determine the role of the local industrial mix in driving differential labor demand shifts for skilled and unskilled workers.

⁴⁸Similarly, Pierce (2001) finds that from 1981 to 1997 the increase in compensation inequality exceeded the increase in wage inequality by 15%.

7 Bibliography

Aguiar, Mark and Erik Hurst, "Measuring Trends in Leisure: The Allocation of Time Over Five Decades", *Quarterly Journal of Economics* (2007b).

Bayer, Patrick, Robert McMillan, and Kim Rueben "An Equilibrium Model of Sorting in an Urban Housing Market" NBER Working Paper 10865

Bayer, Patrick, Shakeeb Khan, Christopher Timmins "Nonparametric Identification and Estimation in a Generalized Roy Model", NBER Working Paper 13949

Bayer, Patrick, and Christopher Timmins "Estimating Equilibrium Models of Sorting Across Locations" *Economic Journal*, March 2007

Bayer, Patrick, and Robert McMillan "Racial Sorting and Neighborhood Quality," NBER Working Paper 11813

Basker, Emek "Selling a Cheaper Mousetrap: Wal-Mart's Effect on Retail Prices", *JUE* 2005, which uses the average prices of drugstore items in selected cities for 1982-2002 and

Baum-Snow Nathaniel and Ronni Pavan Inequality and City Size Brown University, mimeo, 2009.

Beaudry, Paul, Mark Doms and Ethan Lewis "Should the PC be Considered a Technological Revolution? City Level Evidence from 1980-2000", mimeo (2008)

Beeson, Patricia E., 1991. "Amenities and regional differences in returns to worker characteristics," *Journal of Urban Economics*, Elsevier, vol. 30(2), pages 224-241, September.

Black, Dan, Natalia Kolesnikova, and Lowell J. Taylor, "Earnings Functions When Wages and Prices Vary by Location", *Journal of Labor Economics*, 2009.

Black, Dan, Natalia Kolesnikova, and Lowell J. Taylor, Why do so Few Women Work in New York (and so many in Minneapolis)? Labor Supply of Married Women across U.S. Cities, mimeo, University of Chicago (2010)

Black, Dan, Natalia Kolesnikova, and Lowell J. Taylor, Local Price Variation and Labor Supply *Regional Economic Development* 4(1) October 2008 2-14.

Black, Dan, Natalia Kolesnikova, and Lowell J. Taylor, S. Sanders, and M. Wessel A Divergent View on Black-White Earnings Convergence, Mimeo, University of Chicago (2009).

Blanchard Olivier Jean and Lawrence F. Katz, 1992. "Regional Evolutions," *Brookings Papers on Economic Activity*, Economic Studies Program, The Brookings Institution, vol. 23(1992-1), pages 1-76.

Blundell, R, L. Pistaferri and I. Preston, "Consumption inequality and partial insurance" *American Economic Review* 2008

Bound, John and Holzer, Harry J, 2000. "Demand Shifts, Population Adjustments,

and Labor Market Outcomes during the 1980s,” *Journal of Labor Economics*, University of Chicago Press, vol. 18(1), pages 20-54, January.

Buera, Francisco J. and Joseph P. Kaboski ”The Rise of the Service Economy”, NBER Working Paper No. 14822 (2009)

Bureau of Labor Statistics, “Handbook of Methods”, June 2007.

Chen, Tong and Stuart S. Rosenthal “Local Amenities and Migration Over the Life Cycle: Do People Move for Jobs or Fun?” *Journal of Urban Economics* (2008).

Desmet, Klaus and Rossi-Hansberg, Esteban, 2009. ”Spatial growth and industry age,” *Journal of Economic Theory*, Elsevier, vol. 144(6), pages 2477-2502, November.

Dora L. Costa and Matthew E. Kahn, 2000. ”Power Couples: Changes In The Locational Choice Of The College Educated, 1940-1990,” *The Quarterly Journal of Economics*, MIT Press, vol. 115(4), pages 1287-1315, November.

Dahl, Gordon “Mobility and the Return to Education: Testing a Roy Model with Multiple Markets,” *Econometrica*, Vol. 70, No. 6, pp. 2367-2420, (2002).

Duranton, Gilles ”Spatial wage disparities: Sorting matters!” *Journal of Urban Economics*, 2008, 63(2), 723-742

Duranton, Gilles ” Mind the Gaps: The Evolution of Regional Inequalities in the U.K. 1982-1997” *Journal of Regional Science*, 2002, 42(2), 219-256

Duranton, Gilles ”Urban evolutions” *American Economic Review*, 2007, 97 (1), 197-221

Duranton, Gilles, Kristian Behrens and Frdric Robert-Nicoud ”Productive cities: Sorting, selection, and agglomeration”, mimeo, University of Toronto (2010).

Glaeser, E, Gyourko, J. and R. Saks (2005) Why Have Housing Prices Gone Up? *American Economic Review*, forthcoming

Glaeser, E. and C. Berry, “The Divergence of Human Capital Levels across Cities” *Harvard Institute of Economic Research Discussion Paper Number 2091* (2005)

Glaeser, Edward L and Mare, David C, 2001. ”Cities and Skills,” *Journal of Labor Economics*, University of Chicago Press, vol. 19(2), pages 316-42, April.

Goldin Claudia and Lawrence F. Katz, 2007. “Long-Run Changes in the U.S. Wage Structure: Narrowing, Widening, Polarizing,” NBER Working Papers 13568, National Bureau of Economic Research, Inc

Gordon, Robert and Ian Dew-Beker “Controversies About the Rise of American Inequality: A Survey”, NBER WP 13982, 2008.

Gordon, Robert and Ian Dew-Beker ”Selected Issues in American Inequality”. *Brookings Papers on Economic Activity*, 2007, no. 2, pp. 169-92.

Greenstone M., R. Hornbeck and E. Moretti "Identifying Agglomeration Spillovers: Evidence from Million Dollar Plants", *Journal of Political Economy* (forthcoming).

Hamermesh, Dan "Demand for Labor," in *International Encyclopedia of the Social and Behavioral Sciences*, Pergamon Press, 2001.

Hamermesh, Dan "Changing Inequality in Markets for Workplace Amenities," *Quarterly Journal of Economics*, November 1999.

Hamermesh, Dan "A General Model of Dynamic Labor Demand," *Review of Economics and Statistics*, November 1992.

Heckman James, Lance Lochner and Christopher Taber, 1998. "Explaining Rising Wage Inequality: Explanations With A Dynamic General Equilibrium Model of Labor Earnings With Heterogeneous Agents," *Review of Economic Dynamics*, vol. 1(1), pages 1-58, January.

Kahn Matthew E., 1995. "A Revealed Preference Approach to Ranking City Quality of Life," *Journal of Urban Economics*, Elsevier, vol. 38(2), pages 221-235, September.

Kahn, Matthew E. 2005. "Estimating Housing Demand With an Application to Explaining Racial Segregation in Cities," *Journal of Business and Economic Statistics*, American Statistical Association, vol. 23, pages 20-33, January.

Kahn, Matthew E., 1999. "Climate consumption and climate pricing from 1940 to 1990," *Regional Science and Urban Economics*, Elsevier, vol. 29(4), pages 519-539, July.

Katz, Lawrence F and Murphy, Kevin M, 1992. "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *The Quarterly Journal of Economics*, MIT Press, vol. 107(1), pages 35-78, February.

Katz L. and D. Autor "Changes in the Wage Structure and Earnings Inequality", *Handbook of Labor Economics*, Ashenfelter and Crad, Eds, Elsevier (1999).

Krupka Douglas J. and Kwame Donaldson, "Wages, Rents and Heterogeneous Moving Costs," *IZA Discussion Papers 3224*, Institute for the Study of Labor (2007)

Lewbel, Arthur and Krishna Pendakur, 2008, "Tricks with Hicks: The EASI Implicit Marshallian Demand System for Unobserved Heterogeneity and Flexible Engel Curves.", *American Economic Review*, forthcoming.

Lemieux, Thomas, "Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?," *American Economic Review* 96(3), June 2006, 461-98.

Lemieux, Thomas "Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis" *Quarterly Journal of Economics* 116, May 2001, pp. 705-46.

Krueger, Perri, Pistaferri, Violante "Cross Sectional Facts for Macroeconomists" *Review of Economic Dynamics* 2010.

Malpezzi, Stephen, Gregory Chun and Richard Green (1998) "New Place-to-Place Hous-

ing Price Indexes for US Metro Areas”, Real Estate Economics, 26.

Moretti, Enrico ”Workers’ Education, Spillovers and Productivity: Evidence from Plant-Level Production Functions”, American Economic Review 94(3), (2004a)

Moretti, Enrico “Workers’ Education, Spillovers and Productivity: Evidence from Plant-Level Production Functions” American Economic Review 94(3) (2004b)

Moretti, Enrico “Local Labor Markets” in Ashenfelter and Card Eds. ”Handbook of labor Economics (forthcoming)

Roback, Jennifer “Wages, rents and the quality of Life”, Journal of Political Economy, 90-6, pp. 1257-1278 (1982)

Roback, Jennifer ”Wages, rents, and amenities: differences among workers and regions” Economic Inquiry, (1988)

Pendakur, Krishna, 2002. ”Taking prices seriously in the measurement of inequality,” Journal of Public Economics, Elsevier, vol. 86(1), pages 47-69, October

Pendakur, Krishna, 1998. ”Changes in Canadian Family Income and Family Consumption Inequality between 1978 and 1992,” Review of Income and Wealth, Blackwell Publishing, vol. 44(2), pages 259-83, June.

Pendakur, Krishna, 2009. ”Testing and imposing Slutsky symmetry in nonparametric demand systems,” Journal of Econometrics, Elsevier, vol. 153(1), pages 33-50, November.

Pendakur, Krishna 2005. ”Semiparametric estimation of lifetime equivalence scales,” Journal of Applied Econometrics, John Wiley and Sons, Ltd., vol. 20(4), pages 487-507

Pierce, Brooks ”Compensation Inequality”, Quarterly Journal of Economics (2001)

Pistaferri, Luigi ”Consumption and Income Inequality” (2008)

Pistaferri, Luigi and Tulio Jappelli The consumption response to income changes, Annual Review of Economics (2010).

Poole R., F. Ptacek and R. Verbrugge “The treatment of Owner-Occupied Housing in the CPI” (2006), Office of Price and Living Conditions, Bureau of Labor Statistics.

Shapiro, Jesse, ”Smart cities: Quality of life, productivity, and the growth effects of human capital” Review of Economics and Statistics, May 2006

Appendix 1

Here I describe in more details on how I compute Local CPI 1 and Local CPI 2. As I mention in the main text, I follow closely the BLS methodology, and take the properly weighted sum of changes in the cost of housing and non-housing consumption. Cost of housing is measured either using rental costs or housing prices. In the first case, my measure of rent is the “gross monthly rental cost” of the housing unit. I limit the sample to 2 or 3 bedrooms rental units. This includes contract rent plus additional costs for utilities (water, electricity, gas) and fuels (oil, coal, kerosene, wood, etc.). This variable is considered by IPUMS as more comparable across households than “contract rent”, which may or may not include utilities and fuels. The Department of Housing and Urban Development (HUD) also uses the “gross monthly rental cost” measure of rent to calculate the federally mandated “Fair Market Rent”.⁴⁹

The housing costs relevant for a worker living in metropolitan area c —whether he rents or own—is the average of the monthly cost of renting a 2 or 3 bedroom apartment among all renters in area c . When cost of housing is measured using housing prices, I use the property value reported by homeowners of 2 or 3 bedroom single family houses. In this case, the housing costs relevant for a worker living in metropolitan area c is then the average of housing values reported by all homeowners of 2 or 3 bedroom homes in area c .

Note that measured changes in cost of housing do not reflect the change in rental cost or changes in property values at the individual level. Instead, measured changes in cost of housing reflect an average for the local housing market, irrespective of an individual own housing cost and irrespective of whether she rents or owns.

As weights, in my baseline specifications I use the expenditure shares that the BLS uses to compute the official CPI. Since the basket is updated periodically, the BLS weights vary by year. One concern is that housing expenditure shares may vary across metropolitan areas because differences in housing prices. Additionally, it is possible that housing expenditure shares vary across skill groups if preferences are non homotetic. I address these possibilities as follows.

First, I consider the possible differences in expenditure shares across metropolitan areas. Since housing costs vary across cities, it is in principle possible that the share of income spent on housing also vary, as consumers adjust their consumption bundles to local prices. Empirically, the demand for housing is not very price elastic and the share of income spent

⁴⁹Rents are imputed for top-coded observations by multiplying the value of the top code by 1.3. Results do not change significantly when no imputation is performed or when I multiply the value of the top code by 1.4. For Local CPI 1, the cost of non-housing consumption is obtained by subtracting changes in the cost of housing from the nationwide CPI-U computed by the BLS:

$$\text{CPI Non-Housing} = (\text{CPI-U}/(1 - w)) - (w/(1 - w))\text{Housing} \quad (25)$$

where “Housing” is the average nationwide increase in cost of housing (from Census data) and w is the BLS housing weight in the relevant year.

on housing appears to be higher in more expensive cities. In a recent AER paper, Lewbel and Pendakur (forthcoming) find that a housing price increase of 10 percent results in a 0.63 percentage points *higher* housing share, everything else constant. If this is true, it implies that the share of income spent on housing in expensive cities like New York is higher than the share of income spent on housing in less expensive cities like Indianapolis, everything else constant. Because college graduates are over-represented in expensive cities like New York and underrepresented in less expensive cities like Indianapolis, this should increase the housing share of college graduates relative to high-school graduates, everything else constant. (In this case, the use of constant housing shares across cities would lead me to *underestimate* the effect that cost of living adjustments have on wage inequality.)

Second, I consider the possibility that housing price elasticity vary by skill level (or income level). Lewbel and Pendakur find that high income individuals substitute less than low income individuals in the face of an increase in the price of housing. This should further increase the housing share of college graduates relative to high-school graduates, everything else constant.

Third, I consider the possibility of non homotetic preferences. Most empirical studies find that housing is a normal good, with an income elasticity just below 1 when income is measured as permanent income.⁵⁰ If this is true, the share of income spent on housing should be slightly lower for college graduates than high-school graduates.

To account for these possibilities, I have replicated my results using different expenditure shares for different cities and different skill groups in different years. In particular, I use available estimates in the literature of price elasticity and income elasticity to impute shares that vary as a function of local housing prices and individual income. For housing, I assume a permanent income elasticity equal to .85, which is the mid-point in the range of estimates provided by Polinsky and Ellwood (1979). I also assume that the percent difference in permanent income between skilled and unskilled workers is 40% in 1980, 53% in 1990 and 60% in 2000. (These figures reflect estimates of the the nominal college premium.) To allow for differences across cities as a function of local housing prices, I use estimates of demand elasticity from Lewbel and Pendakur (AER forthcoming).

As I discuss in the main text, estimates of the college premium based on expenditure shares that vary by MSA, skill group and year are similar to the ones obtained using BLS shares that vary only by year. Overall, using a common housing share for all individuals within a year appears not to be a bad approximation. This is consistent with what reported by Baum-Snow and Pavan (2009), who find that expenditures shares are generally similar across cities of different size (and therefore different price level).

⁵⁰For example, Polinsky and Ellwood (1979) uncover estimates of permanent income elasticity ranging from 0.80 to 0.87.

Table 1: Metropolitan Areas with the Largest and Smallest Share of College Graduates in the Workforce

	College Share in 2000	Change in College Share 1980-2000	Monthly Rent in 2000	Change in Monthly Rent 1980-2000
<u>Metropolitan Areas with the Largest College Share in 2000</u>				
Stamford, CT	.58	.26	1109	759
San Jose, CA	.48	.15	1231	892
Washington, DC/MD/VA	.48	.08	834	532
Boston, MA-NH	.45	.17	854	556
San Francisco-Oakland-Vallejo, CA	.44	.12	1045	724
Ann Arbor, MI	.43	.02	724	417
Columbia, MO	.43	.06	485	239
Raleigh-Durham, NC	.42	.12	669	427
Fort Collins-Loveland, CO	.42	.10	693	419
Trenton, NJ	.41	.14	776	494
<u>Metropolitan Areas with the Smallest College Share in 2000</u>				
Ocala, FL	.15	.02	514	285
Williamsport, PA	.15	.04	434	229
Lima, OH	.15	.05	444	226
Hickory-Morgantown, NC	.15	.02	486	286
Johnstown, PA	.14	.01	370	165
Flint, MI	.14	.01	481	217
Vineland-Milville-Bridgetown, NJ	.13	.01	617	368
Mansfield, OH	.13	.01	460	242
Visalia-Tulare-Porterville, CA	.13	.00	495	270
Danville, VA	.12	.02	401	231

Notes: Share of college graduates is the share of full-time workers between 25 and 60 years old with a college degree or more who live in the relevant city. Monthly rent refers to the average rent paid for a 2 or 3 bedroom apartment.

Table 2: Relative Importance of the Main Aggregate Components in the BLS Consumer Price Index (CPI-U)

Housing	42.7%
Shelter	32.8%
Fuels and Utilities	5.3%
Other Housing	4.6%
Transportation	17.2%
Food and Beverages	14.9%
Medical Care	6.2%
Education and Communication	6.0%
Recreation	5.5%
Apparel	3.7%
Other Goods and Services	3.5%

Notes: Entries are the share of the main aggregate components of the CPI-U. For more disaggregated categories see Appendix 4 in Chapter 17 of the Bureau of Labor Statistics’s “Handbook of Methods” (2007).

Table 3: Changes in the Cost of Living, by Education Group

	1980	1990	2000	Percent Increase 1980-2000
	(1)	(2)	(3)	(4)
<u>Official CPI</u>				
High-School	1	1.53	2.02	102%
College	1	1.53	2.02	102%
Percent Difference	0	0	0	
<u>Monthly Rent</u>				
High-School	247	432	563	127%
College	259	491	642	147%
Percent Difference	4%	11%	14%	
<u>Local CPI 1</u>				
High-School	0.99	1.49	1.95	96%
College	1.01	1.58	2.07	105%
Percent Difference	2%	4%	6%	
<u>Local CPI 2</u>				
High-School	0.98	1.57	2.04	108%
College	1.01	1.71	2.22	119%
Percent Difference	3%	7%	9%	

Notes: Monthly rent refers to the rent paid for a two or three bedroom apartment. Local CPI 1 allows for local variation only in the cost of housing. Local CPI 2 allows for local variation both in the cost of housing and the cost of non-housing goods and services.

Table 4: Nominal and Real Conditional Wage Difference Between Workers with a High School Degree and Workers With College or More, by Year - Baseline Estimates

	1980	1990	2000	1980-2000 Increase	1980	1990	2000	1980-2000 Increase
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<u>Model 1</u>								
Nominal Wage Difference	.40 (.011)	.53 (.012)	.60 (.013)	.20	.35 (.007)	.47 (.006)	.53 (.007)	.18
<u>Model 2</u>								
Real Wage Difference - Local CPI 1	.38 (.010)	.48 (.008)	.53 (.008)	.15	.37 (.008)	.46 (.006)	.52 (.007)	.15
Percent of Nominal Increase Accounted for by Cost of Living				25%				17%
<u>Model 3</u>								
Real Wage Difference - Local CPI 2	.37 (.009)	.45 (.008)	.51 (.008)	.14	.37 (.008)	.46 (.006)	.51 (.007)	.14
Percent of Nominal Increase Accounted for by Cost of Living				30%				22%
MSA Fixed Effects	No	No	No		Yes	Yes	Yes	

Notes: Standard errors clustered by metropolitan area in parentheses. The dependent variable in Model 1 is the log of nominal hourly wage. The dependent variable in Model 2 is the log of real hourly wage, where real hourly wage is the ratio of nominal wage and Local CPI 1. The dependent variable in Model 3 is the log of real hourly wage, where real hourly wage is the ratio of nominal wage and Local CPI 2. All models include dummies for gender and race, a cubic in potential experience, and year effects. Models in columns 5 to 8 also include MSA fixed effects. Sample size is 5,024,221.

Table 5: Additional Specifications: Part I

	1980	1990	2000	1980-2000 Increase	Percent of Nominal Increase Accounted for by Cost of Living
<u>Model 1: Housing Prices Instead of Rental Costs</u>					
Nominal Wage Difference	.40 (.015)	.53 (.009)	.60 (.010)	.20	
Real Wage - Local CPI 1	.36 (.009)	.43 (.010)	.50 (.009)	.14	30%
Real Wage - Local CPI 2	.35 (.010)	.42 (.011)	.50 (.009)	.15	25%
<u>Model 2: ACCRA Non-Housing Prices</u>					
Nominal Wage Difference	.40 (.015)	.53 (.009)	.60 (.010)	.20	
Real Wage - Local CPI 1	.38 (.008)	.48 (.006)	.53 (.007)	.15	25%
Real Wage - Local CPI 2 (ACCRA data)	.39 (.012)	.48 (.006)	.54 (.006)	.15	25%
<u>Model 3: Expenditure Shares Vary By MSA and Skill Group</u>					
Nominal Wage Difference	.40 (.015)	.53 (.009)	.60 (.010)	.20	
Real Wage - Local CPI 1	.38 (.010)	.49 (.008)	.53 (.009)	.15	25%
Real Wage - Local CPI 2	.37 (.010)	.46 (.008)	.51 (.008)	.14	30%
<u>Model 4: Account for Commuting Time</u>					
Nominal Wage Difference	.40 (.010)	.54 (.009)	.60 (.011)	.20	
Real Wage - Local CPI 1	.38 (.008)	.48 (.006)	.53 (.007)	.15	25%
Real Wage - Local CPI 2	.37 (.008)	.45 (.007)	.51 (.007)	.14	30%

Notes: Standard errors clustered by metropolitan area in parentheses. In the top panel I report estimates where I deflate nominal wages based on local CPI's that measure housing costs using the average price of owner occupied houses instead of average rental costs. In the second panel, I compute Local CPI 2 using the Accra dataset to measure local variation in non-housing prices. In the third panel, I compute the Local CPI's allowing for the expenditure share of housing and non-housing goods to vary by metropolitan areas and skill level. In the bottom panel, the dependent variable is hourly wage defined as the ratio of weekly or monthly earnings over the sum of number of hours worked *plus time spent commuting*. See text for details.

Table 6: Additional Specifications: Part II

	1980	1990	2000	1980-2000 Increase	Percent of Nominal Increase Accounted for by Cost of Living
	(1)	(2)	(3)	(4)	(5)
<u>Model 1: Include Workers with Less Than 48 Weeks</u>					
Nominal Wage Difference	.43 (.009)	.57 (.010)	.62 (.012)	.19	
Real Wage - Local CPI 1	.42 (.008)	.52 (.007)	.56 (.008)	.14	26%
Real Wage - Local CPI 2	.41 (.007)	.49 (.007)	.53 (.007)	.12	37%
<u>Model 2: Include Immigrants</u>					
Nominal Wage Difference	.40 (.011)	.54 (.012)	.61 (.013)	.21	
Real Wage - Local CPI 1	.39 (.010)	.49 (.009)	.55 (.010)	.16	24%
Real Wage - Local CPI 2	.38 (.010)	.46 (.010)	.52 (.010)	.14	33%
<u>Model 3: Only Urban Workers</u>					
Nominal Wage Difference	.40 (.011)	.52 (.008)	.60 (.010)	.20	
Real Wage - Local CPI 1	.39 (.010)	.49 (.007)	.55 (.007)	.16	20%
Real Wage - Local CPI 2	.38 (.010)	.47 (.007)	.53 (.007)	.15	25%

Notes: Standard errors clustered by metropolitan area in parentheses. In the top panel, I show estimates based on a sample that includes all wage and salary workers 25-60, irrespective of the number of weeks worked in the previous year. In the middle panel, I show estimates that include workers born outside the US. In the bottom panel I drop rural workers (i.e. those who are not assigned an MSA). See text for details.

Table 7: Nominal and Real Conditional Earnings Difference Controlling for Quality of Housing, by Year - American Housing Survey

	1980	1990	2000	1980-2000 Increase	Percent of Nominal Increase Accounted for by Cost of Living
	(1)	(2)	(3)	(4)	(5)
Nominal Earnings Difference	.37 (.019)	.47 (.008)	.56 (.010)	.19	
Real Earnings Difference - Not Controlling for Quality					
Real Earnings - Local CPI 1	.36 (.010)	.45 (.006)	.52 (.010)	.16	15%
Real Earnings - Local CPI 2	.35 (.013)	.44 (.006)	.51 (.010)	.16	15%
Real Earnings Difference - Controlling For Quality					
Real Earnings - Local CPI 1	.35 (.012)	.43 (.007)	.50 (.012)	.15	21%
Real Earnings - Local CPI 2	.34 (.014)	.42 (.009)	.49 (.014)	.15	21%

Notes: Standard errors clustered by metropolitan area in parentheses. Data are from the American Housing Survey. Available housing quality variables include square footage, number of rooms, number of bathrooms, indicators for the presence of a garage, a usable fireplace, a porch, a washer, a dryer, a dishwasher, outside water leaks, inside water leaks, open cracks in walls, open cracks in ceilings, broken windows, rodents, and a broken toilet in the last 3 months. The dependent variable is log of yearly earnings (top row) or log of yearly earnings divided by the relevant CPI (middle and bottom panel).

Table 8: The Relation between Share of College Graduates and College Premium

	2000	1980-2000 Change	
	Cross-section		
	OLS	OLS	IV
	(1)	(2)	(3)
College Share	.375	.388	.371
	(.031)	(.070)	(.106)
R^2	.30	.10	


Notes: Standard errors in parentheses. The dependent variable in column 1 is the city-specific college premium, defined as the city-specific difference in the log of hourly wage for college graduates and high school graduates conditional on gender, a cubic in potential experience, race and year. The dependent variable in columns 2 and 3 is the change in the city-specific college premium. Entries are the coefficient on college share in column 1 and change in college share in columns 2 and 3. All models are weighted by city size.

Appendix Table 1. Estimates Based on an Alternative Definition of Rental Cost

	1980	1990	2000	1980-2000 Increase	Percent of Nominal Increase Accounted for by Cost of Living
	(1)	(2)	(3)	(4)	(5)
Nominal Wage Difference	.39 (.008)	.53 (.012)	.59 (.013)	.20	
Real Wage - Local CPI 3	.32 (.006)	.41 (.005)	.44 (.004)	.12	40%
Real Wage - Local CPI 4	.28 (.006)	.34 (.006)	.38 (.005)	.10	50%

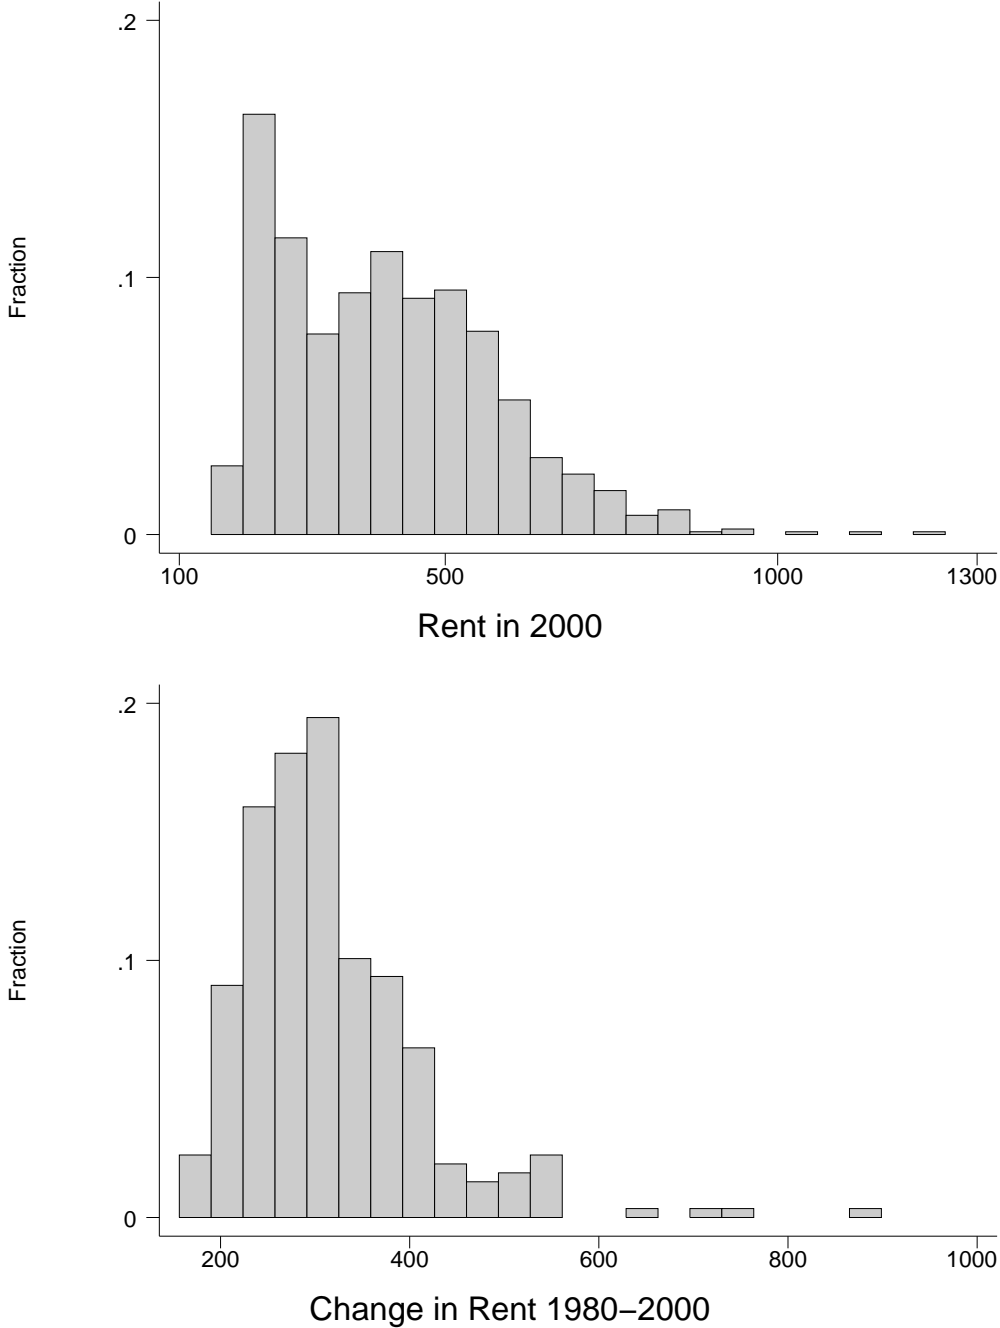
Notes: Standard errors clustered by metropolitan area in parentheses. The dependent variable in the first row is the log of nominal hourly wage. The dependent variable in the second and third row is the log of real hourly wage, where real hourly wage is the ratio of nominal wage and Local CPI 3 or Local CPI 4. In Local CPI 3 and 4, housing costs are allowed to vary by metropolitan area, skill group, race and number of children in the household. Local CPI 3 only uses local variation in cost of living that arises from variation in cost of housing. (The difference with Local CPI 1 is that in Local CPI 1 cost of housing varies only by MSA, while in Local CPI 3 cost of housing varies by MSA, education group, race and number of children.) Local CPI 4 uses local variation both in cost of housing and cost of non housing good and services. (The difference with Local CPI 2 is that in Local CPI 2 cost of housing varies only by MSA, while in Local CPI 4 cost of housing varies by MSA, education group race and number of children.) All models include dummies for gender and race, a cubic in potential experience, and year effects. Sample size is 4,920,703.

Figure 1: How Changes in the Share of College Graduates Relate to the Initial Share of College Graduates, the Initial Cost of Housing and Changes in Cost of Housing



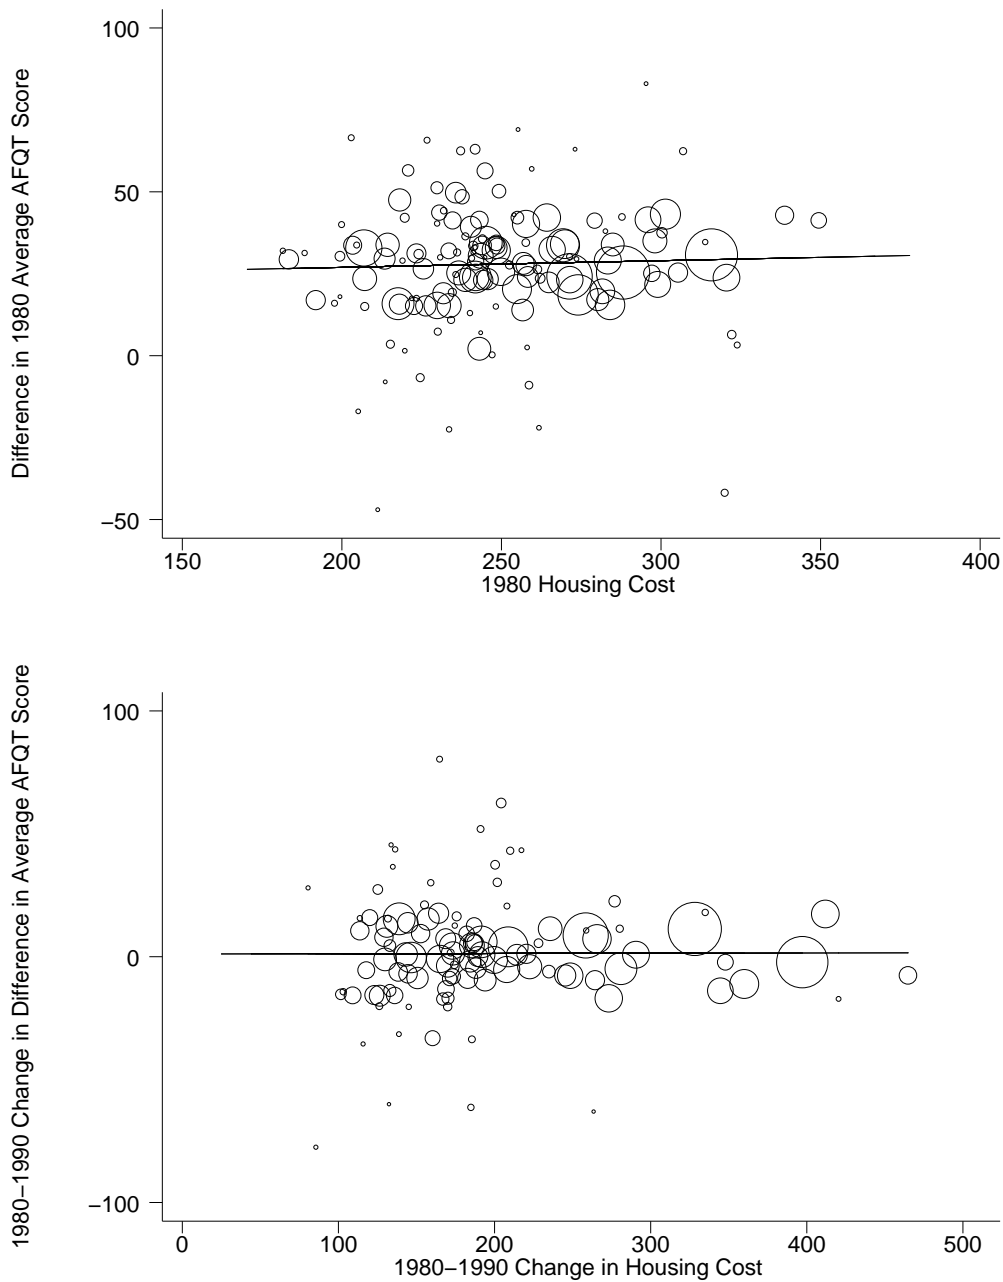
Notes: Each bubble is a metropolitan area. The size of the bubbles reflect population in 1980. Average rent is the average monthly rental price of a two or three bedroom apartment.

Figure 2: The Distribution of Average Rental Costs Across Metropolitan Areas: 2000 Cross-Section and 1980-2000 Change



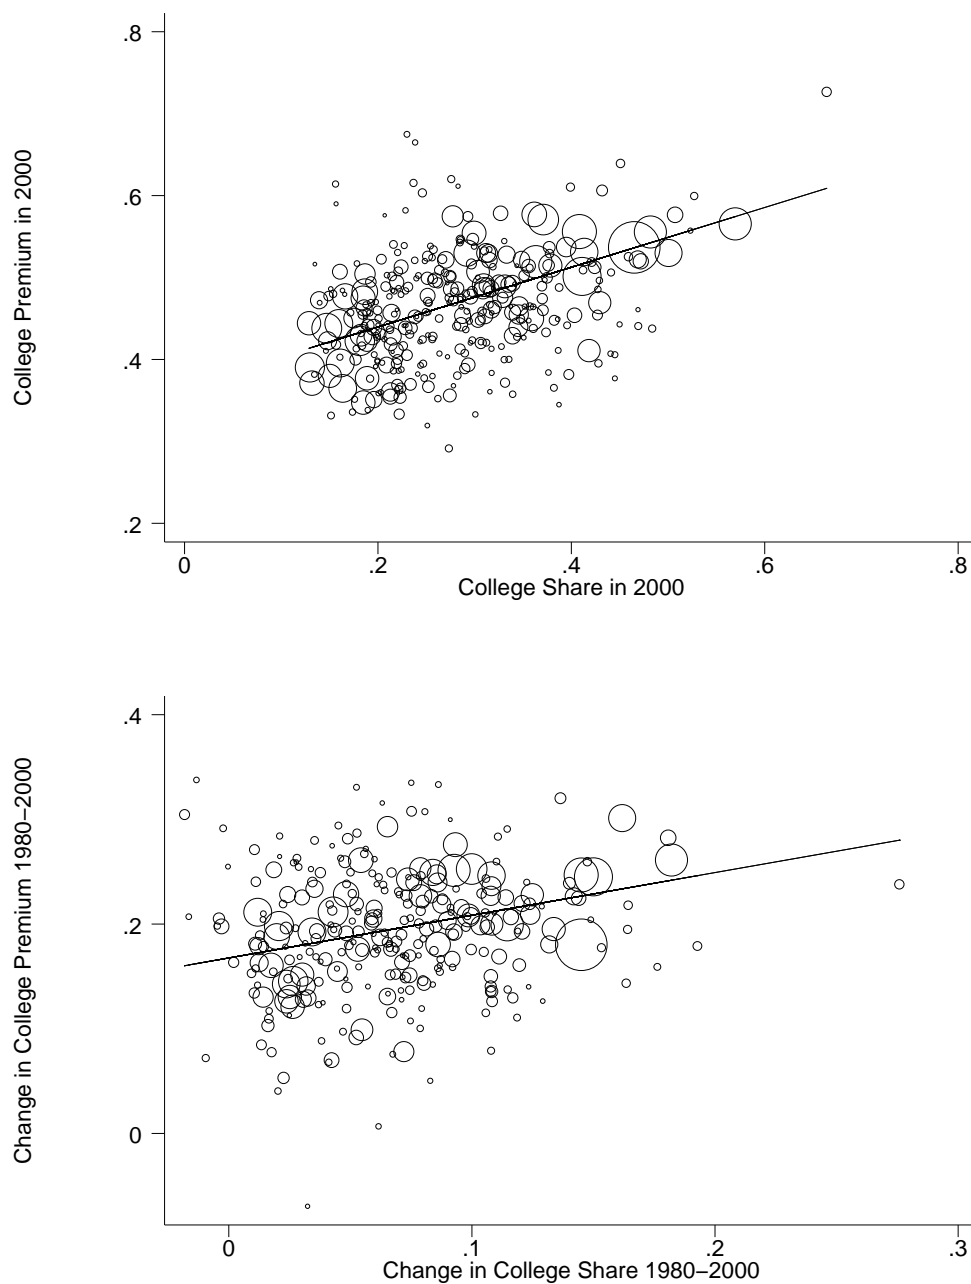
Notes: The top panel shows the distribution of the average cost of renting a 2 or a 3 bedroom apartment in year 2000. The bottom panel shows the distribution of the changes between 1980 and 2000 in the average cost of renting a 2 or a 3 bedroom apartment.

Figure 3: How the Difference Between College and High-School Graduates Average AFQT Scores Relates to Cost of Housing




Notes: The top panel shows average cost of renting a 2 or a 3 bedroom apartment in 1980 (x-axis) and the difference between college graduates and high school graduates in average AFQT scores (y-axis), across metropolitan areas. The size of the bubbles reflects the size of metropolitan areas. A weighted regression yields a coefficient equal to .0203 (.0274). The bottom panel shows the 1980-1990 change in average cost of renting a 2 or a 3 bedroom apartment (x-axis) and the 1980-1990 change in the difference between college graduates and high school graduates in average AFQT scores (y-axis). A weighted regression yields a coefficient equal to .0010 (.0131).

Figure 4: Share of College Graduates and College Premium, by City



Notes: The top panel plots estimates of the city-specific college premium in 2000 against the share of college graduates in 2000. The bottom panel plots the 1980-2000 change in college premium against the 1980-2000 change in the share of college graduates.

Figure 5: Share of College Graduates and Relative Demand Shocks, by City



Notes: The panel plots changes in the share of college graduates 1980-2000 on the y-axis against 1980-2000 shocks to the relative demand of college graduates due to 1980 differences in industry mix on the x-axis. Shocks to the relative demand are defined in equation 24.