

NBER WORKING PAPER SERIES

REAL WAGE INEQUALITY

Enrico Moretti

Working Paper 14370

<http://www.nber.org/papers/w14370>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

September 2008

I thank David Card, Tom Davidoff, Ed Glaeser, Chang-Tai Hsieh, Pat Kline, Douglas Krupka and David Levine for insightful conversations, and seminar participants at Berkeley Economics, Berkeley Haas, Collegio Carlo Alberto in Torino, IZA, San Francisco Federal Reserve and UC Merced 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. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2008 by Enrico Moretti. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Real Wage Inequality
Enrico Moretti
NBER Working Paper No. 14370
September 2008
JEL No. J01,J2,J31,R00

ABSTRACT

A large literature has documented a significant increase in the return to college over the past 30 years. This increase is typically measured using nominal wages. I show that from 1980 to 2000, college graduates have increasingly concentrated in metropolitan areas that are characterized by a high cost of housing. This implies that college graduates are increasingly exposed to a high cost of living and that the relative increase in their real wage may be smaller than the relative increase in their nominal wage. To measure the college premium in real terms, I deflate nominal wages using a new CPI that allows for changes in the cost of housing to vary across metropolitan areas and education groups. I find that half of the documented increase in the return to college between 1980 and 2000 disappears when I use real wages. This finding does not appear to be driven by differences in housing quality and is robust to a number of alternative specifications.

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, 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, 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 of 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 higher cost of living. The empirical evidence indicates that relative demand shifts are more important than relative supply shifts, suggesting that the increase in well-being inequality between 1980 and 2000 is smaller than the increase in nominal wage inequality.

Enrico Moretti
University of California, Berkeley
Department of Economics
549 Evans Hall
Berkeley, CA 94720-3880
and NBER
moretti@econ.berkeley.edu

1 Introduction

A large literature in labor economics has documented a significant increase in wage inequality over the past 30 years.¹ Wage inequality is often measured as the difference between the wage of skilled and unskilled workers, or between the wage of workers at the top and the bottom of the wage distribution.² The existing literature has proposed three broad classes of explanations for the increase in inequality: an increase in the relative demand for skills caused by skill biased technical change and product demand shifts across sectors with different skill intensities; a slowdown in the growth of the relative supply of skilled workers; and the erosion of labor market institutions, such as unions and the minimum wage, that protect low-wage workers.

In this paper, I re-examine how inequality is measured and how it is interpreted. The increase in inequality is typically measured using nominal wages. However, skilled and unskilled workers are not distributed uniformly across cities within the US, and changes in housing costs vary significantly across cities. I assess how existing estimates of inequality change when differences in the cost of living across locations are taken into account. 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 data from the Census of Population, I show that from 1980 to 2000 college graduates have increasingly concentrated in metropolitan areas that are characterized by a high cost of housing. Indeed, much of the growth in the number of college graduates has occurred in metropolitan areas that both have a high initial cost of housing and also experience large increases in the cost of housing. This implies that college graduates are increasingly exposed to a high cost of living and that the relative increase in their real wage may be smaller than the relative increase in their nominal wage.

Although the cost of housing varies substantially across metropolitan areas, changes in the cost of living are almost universally measured using the single *nation-wide* Consumer Price Index (CPI) computed by the Bureau of Labor Statistics (BLS). Changes in this official CPI are a weighted average of changes in the price of the goods in a representative consumption basket. The weights reflect the share of income that the average consumer spends on each good. Housing is by far the largest single item in the CPI, accounting for more than a third of the index.

To measure the wage difference between college graduates and high school graduates in *real* terms, I deflate nominal wages using a new CPI that allows for geographical differences. I closely follow the methodology that the BLS uses to build the official CPI, while allowing for increases in the cost of housing to vary across metropolitan areas and skill groups. In some specifications, I also allow for the price of non-housing goods and services to vary across

¹Comprehensive surveys are found in Katz and Autor (1999) and Goldin and Katz (2007).

²Wage dispersion within skill groups has also increased. In this paper I focus on inequality between groups.

metropolitan areas.

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 19%. This difference grows to 44% in 2000, or more than double 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 only 8 to 10 percentage points. This implies that changes in the cost of living experienced by high school and college graduates account for at least half of the increase in the nominal college premium over the 1980-2000 period. Third, the college premium is significantly smaller in real terms than in nominal terms for each year. For example, in 2000 the average difference between the wage of college graduates and high school graduates is 60% in nominal terms and only 37%-43% in real terms. These findings do not appear to be driven by different trends in relative housing quality and are robust to a number of alternative specifications.

Overall, the difference in the real wage between skilled and unskilled workers is smaller than the nominal difference and has grown significantly less. The implications of this finding for well-being inequality are not straightforward and crucially depend on why college graduates tend to sort into expensive metropolitan areas. I consider two possible explanations and use a simple general equilibrium model to illustrate their different implications.

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 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. The two scenarios are not mutually exclusive, since in practice it is possible that both relative demand and supply shifts take place.

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. 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. Increases in the relative supply of college graduates in a city should cause their relative wage to decline, or at least not to increase. Consistent with demand shocks playing an important role, I find a strong positive association between changes in the college premium and changes in the college share. I also present instrumental variable estimates obtained by instrumenting changes in the college share with a measure of arguably exogenous demand shocks. The instrument—a weighted average of nationwide relative employment growth by industry, with weights reflecting the city-specific employment share in those industries—isolates the effect of changes in the college share that are driven exclusively by changes in relative demand.

Although I can not completely rule out the existence of supply shocks, the empirical evidence is more consistent with the notion that demand shocks are the main force driving changes in the number of skilled workers across metropolitan areas. If this is true, it implies that the increase in well-being inequality between 1980 and 2000 is smaller than the increase in nominal wage inequality.

This result has the potential to explain 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. More generally, the evidence in this paper indicates that general equilibrium effects can undo some of the effects of relative demand shifts.³

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; Krueger, 1993; Autor, Katz and Krueger, 1998). 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.⁴ The notion that demand shocks are important determinants of population shifts is consistent with the evidence in Blanchard and Katz

³See also Heckman et al (1998).

⁴See for example Moretti (2004a and 2004b) and Greenstone, Hornbeck and Moretti (2007).

(1992) and Bound and Holzer (2000).⁵ 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 on differential housing costs complement the findings on non-housing consumption in a contemporaneous paper by Broda and Romalis (2008) which documents the distributional consequences of increased imports from China. They find that, because of international trade, poor households are exposed to lower inflation for non-housing goods than rich households. Taken together, my findings and their findings suggest that the overall difference in real income between the rich and the poor is smaller than previously thought. My results are also related to a paper by Black, Kolesnikova, and Taylor (2007) 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.

The rest of the paper is organized as follows. In Section 2, I describe some recent changes in the geographical distribution of skilled and unskilled workers. In Section 3, 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 4, I present the main evidence on nominal and real inequality. In Section 5, I present a simple model that can help interpreting the empirical evidence. In Section 6, I discuss the different implications of the demand pull and supply push hypotheses and present empirical evidence to distinguish the two. Section 7 concludes.

2 The Cost of Living and Location of Skilled and Unskilled Workers

I begin with some descriptive evidence on recent changes in the geographical location of skilled and unskilled workers and housing costs. Throughout the paper, I use data from the 1980, 1990 and 2000 Censuses of Population. 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

⁵Chen 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.

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 top panel of Figure 1 shows how the 1980-2000 change in the share of college graduates relates to the 1980 share of college graduates. 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 Moretti (2004) and Berry and Glaeser (2005).⁶ 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.⁷ 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.⁸ These relationships do not have a causal interpretation, but instead need to be interpreted as equilibrium relationships.

Probably a better measure of the cost of housing experienced by college graduates in a given city is the average rent *paid by college graduates* in that city. Allowing for the cost of housing faced by different skill groups in a given city to be different is potentially important, since tastes and budget constraints might differ across skill groups. This implies that the type of housing that is used by college graduates is not necessarily identical to the one that is used by high school graduates. The top left panel in Figure 2 shows the 1980-2000 increase

⁶The 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.

⁷The 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.

⁸The regression yields a coefficient equal to .0003 (.00001).

in the share of college graduates, by quintile of 1980 cost of housing for college graduates, as measured by the average monthly rent paid by college graduates for a 2 or 3 bedroom apartment in the relevant metropolitan area. Metropolitan areas where college graduates pay high rent in 1980 experience a larger increase in college share between 1980 and 2000, and this increase is monotonic. The top right panel shows the 1980-2000 increase in the share of college graduates, by quintile of the 1980-2000 increase in cost of housing. Metropolitan areas where the rent paid by college graduates increases the most between 1980 and 2000 experience a larger increase in college share between 1980 and 2000.

Even within a skill group, not all households may use the same type of housing. For example, the housing needs of a family with many children might not be identical to those of a couple with the same education in the same metropolitan area but no children. I consider an alternative measure of the cost of housing that allows the cost faced by households in a given education group and city to vary depending on family size and race. To do so, I take the fitted value from a regression of rental cost on identifiers for metropolitan area, education group, number of children, race and interactions. This regression is estimated on the sample of renters of 2 or 3 bedroom apartments, and the predicted values are calculated for all households. The bottom panels of Figure 2 show the 1980-2000 increase in the share of college graduates by quintile of predicted 1980 cost of housing (left panel) and by quintile of predicted 1980-2000 increase in cost of housing (right panel). The pattern is similar to the one in the top panel.

Taken together, the four panels in Figure 2 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 for college graduates in 1980 is highest and also the areas where the average cost of housing for college graduates has increased the most.

3 Cost of Living Indexes

In this Section, I briefly describe how the Bureau of Labor Statistics computes the official Consumer Price Index (subsection 3.1), and I propose a new measure of cost of living that accounts for geographical differences (subsection 3.2). In the next Section, I use this new measure to estimate how much of the difference in wages between skilled and unskilled workers can be attributed to geographical differences in the cost of living.

3.1 The Official Cost of Living Index

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 .⁹

Table 2 shows the relative importance of the main aggregate components of the 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.

3.2 Local Consumer Price Indexes

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. To investigate the role of cost of living differences on wage differences between skill groups, I propose two CPI indexes that are skill-group specific. 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 metropolitan area of residence of skilled and unskilled workers and the type of housing used by skilled and unskilled workers to differ. Specifically, to measure the cost of housing faced by an individual in metropolitan area c and skill group j , I take the average of the monthly cost of renting a 2 or 3 bedroom apartment among all individuals in group j in area c . Consistent with BLS methodology, I assign the cost of housing to homeowners based on changes in the relevant average monthly rent. It is important to note that this methodology ensures that the deflator that I use for

⁹One 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”.

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 workers in the same city and education group, irrespective of their own individual housing cost and irrespective of whether they rent or own.

Following the BLS methodology, I then take the properly weighted sum of the cost of housing—with the average across cities and skill groups normalized to 1 in 1980—and non-housing consumption—normalized to 1 in 1980. The weights are the weights used by BLS in the relevant year. The cost of non-housing consumption is assumed to be the same for all individuals in a given year and is obtained by subtracting changes in the cost of housing from the nationwide CPI-U computed by the BLS:

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

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.¹⁰

I call the resulting local price index “Local CPI 1”. Note that Local CPI 1 varies by MSA as well as by education group. This is because the type of housing that is relevant for college graduates in a given MSA does not have to be identical to the type of housing that is relevant for high school graduates in the same MSA. For a given education group, the correct measure of housing cost inflation for the period between t and $t + 1$ should reflect how much the cost of purchasing the housing bundle consumed *by that education group* at time t increases between t and $t + 1$. Differences in the housing bundle (type of dwellings, type of neighborhoods, etc.) consumed by different education groups are potentially important if tastes are different, or if tastes are the same but preferences are non-homothetic. Whatever the housing consumption bundle chosen by college graduates and high school graduates in 1980, my CPI seeks to measure how the price of that bundle changes between 1980 and 2000. While college graduates may in principle be able to afford buying the housing bundle consumed by high school graduates, what matters to college graduates given their own tastes and their own budget constraint is how the price of *their own* housing consumption bundle changes over time. How the price of the housing consumption bundle of high school graduates changes over time is not relevant to college graduates.¹¹ Broda and Romalis (2008) make a

¹⁰In practice, my measure of rent is the “gross monthly rental cost” of the housing unit. 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”. The weights are the BLS weights for the relevant year for “Shelter” and “Fuel and Utilities”. Since the basket is updated periodically the weights vary over time. The weight for year 2000 is 0.381 (see Table 2). The weight in 1980 is .355 and in 1990 is .356. Rents 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.

¹¹Consider, for example, the difference between Palo Alto and the neighboring East Palo Alto. They both

similar point and argue that the basket of non-durable goods consumed by the poor differs from the one consumed by the rich, so that the inflation rate faced by the poor is not necessarily the same as the inflation rate faced by the rich. Of course, an important concern is the possibility of differential changes in the unmeasured quality of housing for college graduates and high school graduates. To address this concern, in Section 4.2 I measure changes in a rich set of observable housing quality variables, and show that the relative changes are not substantial.¹²

It is possible that the relevant housing market for individuals in a given city may depend not just on education, but also on other household characteristics, like number of children or race. For example, the housing needs of a family with many children might not be identical to those of a couple with no children with the same education in the same metropolitan area. The former might be more interested, for example, in larger houses and neighborhoods with better schools. To account for this possibility, in some specifications I allow for the cost of housing faced by different individuals to vary depending not only on metropolitan area and skill level, but also on number of children and race. In this case, the relevant cost of housing is obtained as the predicted value from a regression of rental cost on identifiers for metropolitan area, education group, number of children, race and interactions (see the bottom panels in Figure 2). This regression is estimated on the sample of renters of 2 or 3 bedroom apartments and the predicted values are calculated for all households.

Local CPI 2. In 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 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 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.¹³

belong to the same MSA, but the former is one of the richest and better educated cities in Silicon Valley, while the latter is much poorer, considerably less educated, and is characterized by very high crime rates and a high fraction of minority residents. The housing market that is relevant for college graduates in that area is more likely to be Palo Alto than East Palo Alto. As a consequence, the housing price dynamics in Palo Alto are more relevant for college graduates than the ones in East Palo Alto. Indeed, during the dot com boom of the late 1990s that attracted an enormous influx of skilled workers to Silicon Valley, housing prices in Palo Alto skyrocketed, while housing prices in East Palo Alto were generally more stable.

¹²A limitation is that I use the same consumption shares for all individuals in a given year, irrespective of where they live. While it is possible that the consumption shares differ across metropolitan areas, the BLS only publishes nation-wide shares. Because college graduates are over-represented in expensive cities, if the share of income spent on housing is higher in more expensive cities then my estimates of the fraction of the college premium explained by the cost of living may be conservative and should be considered a lower bound.

¹³The cost of leasing a store is certainly higher in New York; labor costs are also likely to be higher in

In a second departure from the standard BLS CPI, I propose an index that 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. The BLS releases a local CPI for some metropolitan areas. This local CPI is far from ideal. First, it is available only for a limited number of MSA's. Of the 315 MSA's in the 2000 Census, the metropolitan-level CPI is made available by the BLS only for 23 MSA's in the period under consideration. Second, 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}):

$$BLS_{ct} = wHP_{ct} + (1 - w)NHP_{ct} \quad (2)$$

where w is the CPI weight used by BLS for housing. Non-housing costs can be divided in two components:

$$NHP_{ct} = \pi HP_{ct} + v_{ct} \quad (3)$$

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 .¹⁴ 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. I use as weights the weights used by BLS in the relevant year.¹⁵

While Local CPI 2 is more comprehensive than Local CPI 1 because it includes local variation in both housing and non-housing costs, it is also arguably less reliable because non-housing costs are imputed. For this reason, in the next Section I present separate estimates

New York, since workers need higher wages to live in New York.

¹⁴To 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.

¹⁵A limitation of this methodology is that while I allow for within-metropolitan area differences in the type of housing consumed by workers belonging to different skill groups, I need to assign the same price of non-housing consumption to all groups. In reality, it is possible that different skill groups consume different types of non-housing goods. Indeed, Broda and Romalis (2008) document that the inflation rate for non-housing goods faced by the rich is on average higher than the inflation rate for non-housing goods faced by the poor. If the price of non-housing goods consumed by college graduates grows faster than the price of non-housing goods consumed by high school graduates, my measure of the difference in the cost of living increases between college and high school graduates is conservative.

for Local CPI 1 and Local CPI 2. Moreover, in Section 4, I show how my estimates change when I compute Local CPI 2 using data on non-housing prices taken from the Accra dataset collected by the Council for Community and Economic Research.

4 Empirical Evidence

In this section I estimate how much of the difference in wages between skilled and unskilled workers can be attributed to geographical differences in the cost of living. I begin in sub-section 4.1 by documenting empirical differences in the cost of living. The main empirical results of the paper are in sub-section 4.2, where I show estimates of the college premium in nominal and real terms, by year. I also discuss to what extent differences in the cost of housing may reflect differences in the quality of housing, and I investigate the robustness of my findings.

4.1 Differences in the Cost of Living

Table 3 compares changes in the official, nation-wide CPI from the BLS to changes in my two local CPI's. Specifically, the top of the table reports 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 19% more than the cost of housing for the average high school graduate. This gap grows to 36% in 1990 and reaches 44% by 2000. Column 4 indicates that housing costs for high school and college graduates increased between 1980 and 2000 by 113% and 156%, respectively.

The next panel shows “Local CPI 1”, normalized to 1 in 1980 for the average household. The panel shows that in 1980 the overall cost of living experienced by college graduates is 7% higher than the cost of living experienced by high school graduates. This difference increases to 18% 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 11% higher for college graduates relative to high school graduates in 1980 and 27% in 2000. Column 4 indicates that the increase in the overall price level experienced by high school

graduates between 1980 and 2000 is 98%. The increase in the overall price level experienced by college graduates between 1980 and 2000 is 127%.

4.2 Nominal and Real Wage Differences

In this subsection I measure 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.

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.

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 significantly 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 .43 in real terms when Local CPI 1 is used as deflator. The difference is even smaller—.37 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 10 percentage points, or half of the increase in the nominal wage difference. In other words, cost of living differences as measured by Local CPI 1 account for 50% 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 8 percentage points. This implies that cost of living differences as measured by Local CPI 2 account for 60% of the increase in conditional wage inequality between college and high school graduates between 1980 and 2000 (column 4). The effect of using real wages instead of nominal wages as a dependent variable in the regression is shown graphically in Figure 3.¹⁶

When I control for fixed effects for metropolitan areas in columns 5-8, the nominal college premium is smaller in all years, but the effect of using real wages is similar. In particular, the increase in the college premium is 18 percentage points when measured in nominal terms, and only 8-10 percentage points when measured in real terms, depending on whether CPI 1 or CPI 2 is used as deflator. Put differently, after conditioning on MSA fixed effects, cost of living differences account for 44% to 55% of the increase in conditional inequality between college and high school graduates between 1980 and 2000 (column 8).

Housing Quality. An important concern is the possibility that the differential changes in housing costs faced by skilled and unskilled workers reflect not just changes in cost of living, but also changes in the quality of housing. It is in principle possible that part of the relative increase in the measured cost of housing for college graduates reflects better unobserved quality. This could occur if housing quality is a normal good and features of the apartments inhabited by college graduates—such as the number of bathrooms, the quality of the kitchen, the availability of a fireplace or a garage, etc.—improve more than the features of the apartments inhabited by high school graduates. In this case, the relative increase in the cost of housing of college graduates documented above may be overestimated. Additionally, the size of the apartment may have changed differently for high school and college graduates. Although my measure of housing cost is the average rent for apartments with a fixed number of bedrooms, exact square footage may vary. For example, a 2 bedroom apartment in New York or San Francisco is likely to be smaller than a 2 bedroom apartment in Houston or Indianapolis. To the extent that the share of college graduates has increased more in more expensive cities, the true per-square-foot price faced by college graduates may be higher than the one that I measure. In this case, the relative increase in the cost of housing of college graduates documented above may be underestimated.

While I can not completely rule out the possibility of unmeasured quality differences,

¹⁶One might be concerned about worker selection. Models in Table 4 control for standard demographics, but not for worker ability. Ability of college graduates and high school graduates may vary across metropolitan areas. For example, a corporate lawyer in New York may have more unobserved ability than a lawyer in Indianapolis. Similarly, a software engineer in San Jose may be of better quality than one in Grand Rapids. Without knowing the exact type of selection, one can only speculate on the type of bias that may be caused by the failure to account for unobserved heterogeneity. If the average unobserved ability of college graduates relative to high school graduates is higher in expensive cities, then the estimates of the real college premia in Table 4 are biased upward. The quality-adjusted college premia would be even smaller than the one reported in the Table.

here I present evidence on some quality variables that I can measure. I use data from the American Housing Survey, which includes much richer information on housing quality than the Census of Population. Available quality variables include square footage, number of rooms, number of bathrooms, indicators for the presence of a garage, a usable fireplace, 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.¹⁷

I begin by reproducing the baseline estimates that do not control for quality. Estimates based on the American Housing Survey in the top panel of Table 5 are qualitatively similar to the corresponding baseline estimates based on the Census reported in Table 4, although the effect of controlling for cost of living is smaller than in Table 4. I then re-estimate my models 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 5 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. Entries suggest 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 only 9 or 6 percentage points, depending on whether Local CPI 1 or CPI 2 is used. The corresponding estimates that do not control for quality are 14 and 13.

In sum, though I can not completely rule out the possibility of unmeasured quality differences, Table 5 indicates that controlling for a rich vector of observable quality differences result in differences between nominal and real college premium that are even larger than the baseline differences.

Robustness. I now investigate the robustness of my findings. First, I relax the assumption made in Table 4 that all individuals within a given skill group and city experience similar changes in the cost of living. In Model 1 of Table 6 I allow for the cost of housing faced by different individuals to vary depending not just on skill level and metropolitan area of residence, but also on some household characteristics, including race and number of children. It is plausible that the type of housing and the type of neighborhoods that are relevant for, say, a college graduate in a given city may differ depending on her family structure and race. For the sample of renters in 2 or 3 bedrooms apartments, I run a regression where the dependent variable is the monthly rent and the independent variables include dummies

¹⁷Each 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 1978-1982, 1988-1992 and 1998-2002 together.

for MSA, skill group, year, race, and number of children and interactions. I then assign the predicted values to all individuals, irrespective of their own individual housing cost and irrespective of whether they rent or own.

In Model 2 of Table 6, I compute Local CPI 2 using data on non-housing prices from the Accra dataset collected by the Council for Community and Economic Research.¹⁸ The Accra data have both advantages and disadvantages when compared to the BLS local price index. On one hand, the Accra data are available for a much larger set of cities. 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.¹⁹ Moreover, the sample size for each good and city is quite small and the set of cities covered changes over time. With these limitations in mind, I follow the same methodology used to compute Local CPI 2, but use Accra data instead of the local BLS for non-housing goods.

In Model 3, 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. (I calculate hourly wage by summing number of hours worked and time spent commuting.) In Model 4, I show estimates that include workers born outside the US. In Model 5 I drop rural workers (i.e. those who are not assigned an MSA).

In general, estimates in Table 6 are not very different from the baseline estimates in Table 4. Indeed, estimates in Model 2 actually indicate an increase in real wage inequality of only 5 percentage points, which is smaller than the corresponding increase in Table 4.²⁰

¹⁸The data were generously provided by Emek Basker. Basker (2005) and Basker and Noel (2007) describe the Accra dataset in detail.

¹⁹Only 48 goods have prices that are consistently defined for the entire period under consideration. The BLS basket includes more than 1000 goods.

²⁰I have estimated several additional robustness checks that are not reported in the Table due to space limitations. When I include all US born workers—irrespective of the number of weeks worked in the previous year—I find that the nominal wage difference grows 19 percentage points, from 43% in 1980 to 62% in 2000. The real wage difference grows only 9 or 6 percentage points, depending on whether CPI 1 or CPI 2 is used. 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. For example, the 1980-2000 increase in wage gap (column 4) for males is .22 percentage points in nominal terms, .12 in real terms using local CPI 1 and .09 using local CPI 2. The corresponding figures for females are .20, .11, .09. 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. For example, the 1980-2000 increase in wage gap (column 4) for workers with less than 20 years of experience are .20 in nominal terms, .14 using local CPI 1 and .11 using local CPI 2. The corresponding figures for workers with more than 20 years of experience are .15, .06 and .05. Estimates where the dependent variable is the log of weekly or yearly earnings are also

5 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 finding mean that the significant increases in wage disparities that have been documented over the last 25 years have failed to translate into significant increases in disparities in well-being?

Not necessarily. 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 broad class of explanations for such an increase. 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. In this case, a higher cost of housing reflects consumption of local amenities. Since this consumption arguably generates utility, the increase in utility disparity is larger than the increase in real wage disparity. Under a demand pull hypothesis, the relative demand of college graduates increases in expensive cities because their relative productivity increases there so that 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 college graduates and are located in expensive cities (for example, high tech, finance, etc.). In this case, a higher cost of housing does not reflect better amenities, and the increase in utility inequality is smaller than the increase in nominal wage inequality.

To formalize these two alternative hypotheses, I consider the simplest possible general equilibrium model of the labor and housing market. The model is a generalization of the Roback (1982) model and has two types of workers, skilled and unskilled. Like in Roback, in equilibrium workers and firms are indifferent between cities. But unlike Roback, housing supply is not necessarily fixed, so that 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 to have different effects on the utility of skilled and unskilled workers.

5.1 Assumptions and Equilibrium

Assume that there are two cities: Detroit (city a) and San Francisco (city b). 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. There are two type of workers: skilled workers (type H) and unskilled workers (type L).

generally consistent with Table 4. For example, the 1980-2000 increase in wage gap for weekly earnings is .23 in nominal terms, .13 in real terms using local CPI 1 and .09 using local CPI 2. 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).

I assume that workers and firms are perfectly mobile. This implies that in equilibrium workers need to be indifferent between living in Detroit and San Francisco. Similarly, firm profits need to be equalized across locations. Since in my empirical analysis I focus on long run changes (over a 20 year period), this assumption does not appear to be unrealistic. For simplicity, I also assume no human capital externalities and that the owners of land and capital live abroad.

I first focus on the case where skilled and unskilled workers in the same city work in different firms and live in different neighborhoods. This amounts to assuming away imperfect substitution between skilled and unskilled workers and the effect that shocks to skilled workers have on unskilled workers. This assumption greatly simplifies the analysis. Later, I show that results generalize when I relax this assumption. The production function for firms in city c that hire skilled workers is Cobb-Douglas with constant returns to scale:

$$\ln y_{Hc} = X_{Hc} + hN_{Hc} + (1 - h)K_{Hc} \quad (4)$$

where N_{Hc} is the log of the number of skilled workers hired in city c ; $c = a, b$; K_{Hc} is the log of capital and X_{Hc} is a productivity shifter. If firms are price takers and labor is paid its marginal product, the log of the wage of skilled workers, w_{Hc} , is

$$w_{Hc} = X_{Hc} - (1 - h)N_{Hc} + (1 - h)K_{Hc} + \ln h \quad (5)$$

Equation 5 represents the labor demand for skilled labor in city c . I assume that there is an international capital market, so that capital is infinitely supplied at price i . In equilibrium demand for capital is equal to its supply:

$$X_{Hc} - hK_{Hc} + hN_{Hc} + \ln(1 - h) = \ln i \quad (6)$$

To keep things simple, I do not consider worker labor supply decisions and I assume that each worker provides one unit of labor. Similarly, I assume that each worker consumes one unit of housing. The indirect utility of skilled workers in city c is

$$U_{Hc} = w_{Hc} - r_{Hc} + A_{Hc} \quad (7)$$

where r_{Hc} is the cost of housing in city c in the neighborhoods where skilled workers live, and A_{Hc} is a local amenity.

In equilibrium it has to be the case that workers have the same level of utility in San Francisco and Detroit. This implies that skilled labor in San Francisco is supplied with infinite elasticity at the wage level

$$w_{Hb} = w_{Ha} + (r_{Hb} - r_{Ha}) - (A_{Hb} - A_{Ha}) \quad (8)$$

and that the (inverse of) the demand curve for housing in San Francisco is

$$r_{Hb} = (w_{Hb} - w_{Ha}) + r_{Ha} + (A_{Hb} - A_{Ha}) \quad (9)$$

An increase in the cost of housing in San Francisco or in the wage in Detroit lowers the supply of skilled workers in San Francisco. An increase in the cost of housing in Detroit or an increase in the amenity in San Francisco increases the supply of workers there.

I assume that the supply of housing is

$$r_{Hc} = z + k_c N_{Hc} \quad (10)$$

The slope parameter, k_c , represents how elastic the supply of housing is in city c . I assume that this parameter is exogenously determined by geography and local land regulations. In cities where geography and regulations make it is 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.²¹ Finally, I assume that the number of workers in the economy is fixed.

In period 1, the two cities are identical. Equilibrium in the labor market for skilled workers is obtained by equating equation 5 and 8. Equilibrium in the housing market for skilled workers is obtained by equating equation 9 and 10. Because of the assumptions on the technology, profits are always zero, so that firms are indifferent between cities. The labor and housing markets for unskilled workers are similar. For example, the city-level production for firms that hire unskilled workers is $y_{Lc} = X_{Lc} N_{Lc}^h K_{Lc}^{1-h}$.

5.2 Demand Pull

I consider two scenarios. In the first scenario, 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 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 and San Jose in the second half of the 1990s.

How the Equilibrium Changes. Attracted by higher labor demand, skilled workers

²¹Equation 10 ignores the durability of housing—i.e. the fact that once built, housing does not depreciate quickly (Glaeser and Gyourko, 2001).

move to San Francisco. The number of skilled workers in San Francisco increases by

$$N_{Hb2} - N_{Hb1} = \frac{\Delta}{k_a + k_b} \quad (11)$$

The number of skilled workers in Detroit declines by the same amount. What happens to wages and rents? In San Francisco, the nominal wage of skilled workers increases by an amount Δ equal to the productivity increase

$$w_{Hb2} - w_{Hb1} = \Delta \quad (12)$$

while rents increase by a fraction of Δ :

$$r_{Hb2} - r_{Hb1} = \frac{k_b}{k_a + k_b} \Delta \quad (13)$$

In Detroit, nominal wages for skilled workers do not change.²² Because of the decline in the number of workers, the cost of housing in Detroit declines:

$$r_{Ha2} - r_{Ha1} = -\frac{k_a}{k_a + k_b} \Delta \quad (14)$$

In equilibrium workers are indifferent between cities. Real wages and utility of skilled workers increase by the same amount in San Francisco and Detroit:

$$(w_{Hb2} - r_{Hb2}) - (w_{Hb1} - r_{Hb1}) = (w_{Ha2} - r_{Ha2}) - (w_{Ha1} - r_{Ha1}) = \frac{k_a}{k_a + k_b} \Delta \quad (15)$$

Firms are also indifferent between cities. Because of the assumptions on technology, firms have zero profits in both cities. While skilled 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.

Who Benefits? In this model, the benefit of the increase in workers' productivity, Δ , is split between workers and landowners. The fraction of Δ that goes to workers depends on the relative elasticity of housing supply in the two cities. To see this, note that the change in real wages in equation 15 depends on k_a and k_b , which are the elasticities of supply of housing in Detroit and San Francisco. Two special cases are of interest.

1. In the extreme case where the supply of housing in San Francisco is perfectly inelastic ($k_b = \infty$), all the benefit of the productivity increase goes to landowners in San Francisco. Workers' utility does not change. This case is the one described in Roback (1982).

²²This may look surprising at first. Given that the supply of skilled workers has declined, and that the demand curve is downward sloping, one might expect an increase in wages of those workers who stay in Detroit. Indeed, this would be true in a model without capital. But in a model that includes capital, the amount of capital increases in San Francisco ($K_{Hb2} - K_{Hb1} = \frac{\Delta}{k_a + k_b}$) and decreases in Detroit ($K_{Ha2} - K_{Ha1} = -\frac{\Delta}{k_a + k_b}$). This capital flow off-sets the decline in labor supply in Detroit, so that there is no change in the wage in Detroit.

Housing costs and nominal wages in San Francisco increase by the same amount Δ . Because the supply of housing is fixed, the number of skilled workers in San Francisco can not change. Since there is no migration from Detroit to San Francisco, nothing happens to labor or housing prices in Detroit. (Another special case where all the rent resulting from the productivity shock accrues to landowners in San Francisco is when the elasticity of supply of housing in Detroit is infinite so that $k_a = 0$.)

2. At the other extreme is the case where the elasticity of supply of housing in San Francisco is infinite ($k_b = 0$). In this case, all the benefit of the productivity increase goes to workers. Real wages in San Francisco and Detroit grow by Δ . Landowners in San Francisco are indifferent, while landowners in Detroit experience a loss equal to Δ due to the decline in demand for housing. (Another special case where all the rent resulting from the productivity shock accrues to workers is when the housing supply in Detroit is perfectly inelastic so that $k_a = \infty$.)

Distribution of the Shocks. Consistent with the empirical evidence in Section 4, in the demand pull scenario the nominal wage averaged across the two cities increases for skilled workers.²³ Furthermore, this increase is larger than the increase in the real wage averaged across the two cities for skilled workers, unless in period 1 skilled workers in San Francisco have much lower productivity than skilled workers Detroit.²⁴ In other words, for localized skill biased demand shocks to be consistent with my empirical evidence in Section 4, these shocks can not be concentrated in cities with a small initial share of college graduates. Consistent with this notion, Figure 1 has shown that increases in the number of college graduates are more concentrated among cities that have a large initial share of college graduates. Also consistent with this notion, Beaudry, 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. Similarly, Berry and Glaeser (2005) show evidence consistent with a model of urban agglomeration where the number of entrepreneurs is a function of the number of skilled people working in an area. If skilled people are more likely to innovate in ways that employ other skilled people, this creates an agglomeration economy where skilled people want to be around each other.

Before proceeding, it is important to highlight that the model in this Section focuses on the case where the housing market for skilled workers is separated from the housing market for unskilled workers in the same city. This assumption has the advantage of making the model very simple and transparent. It has the disadvantage that skilled and unskilled workers in a city do not compete for the same set of houses, and therefore shocks to the

²³This increase is equal to $\frac{\Delta(2X_{Hb} + hNk_a - 2X_{Ha} + h\Delta)}{h(k_b + k_a)N}$.

²⁴Formally, $X_{Ha} < X_{Hb} + (h\Delta)/2$.

relative demand or supply of skilled workers have no effect on unskilled workers in the city. I also assume that the labor supply in each city is infinitely elastic. In Appendix 1, I show that the qualitative results of the model hold when the housing market is integrated—i.e. skilled and unskilled workers in a city compete for the same set of houses—and the labor supply in a city is not infinitely elastic but is upward sloping.²⁵

5.3 Supply Push

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 more desirable for skilled workers relative to Detroit. I assume that in period 2, the amenity level increases in San Francisco: $A_{Hb2} = A_{Hb1} + \Delta'$, where $\Delta' > 0$ now 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.²⁶

How the Equilibrium Changes Like for the case of demand pull above, $\frac{\Delta'}{k_a + k_b}$ skilled workers move from Detroit to San Francisco. As before, the cost of housing increases in San Francisco (by the amount in equation 13) and declines in Detroit (by the amount in equation 14). Also, similar to before, the nominal wage in Detroit does not change.

A key difference with the case of demand pull is that the nominal wage of skilled workers in San Francisco remains unchanged. This may be surprising at first. While one expects wage *increases* in response to demand increases (this is exactly what happens in subsection 5.2), one expects wage *decreases* in response to supply increases. Why nominal wages do not decline in San Francisco after it has become more attractive? After all, skilled workers should be willing to pay a compensating differential in the form of lower nominal wages to live in the more desirable city. Indeed, this is what the Roback (1982) model would predict. But the Roback model ignores the endogenous reaction of capital. 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.²⁷ Workers in both

²⁵Another assumption of the model is that skilled and unskilled workers are employed by different firms. The model generalizes to the case where skilled and unskilled worker can work in the same firm. In this case, the productivity of unskilled workers may increase when the number of skilled workers in the same firm increases because of complementarities between skilled and unskilled labor.

²⁶For 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).

²⁷The amount of capital increases in San Francisco by $K_{Hb2} - K_{Hb1} = \frac{\Delta'}{k_a + k_b}$ and decreases in Detroit by

cities experience an increase in utility equal to $\frac{k_a}{k_a+k_b}\Delta'$. The two special cases described in subsection 5.2 apply to this scenario as well. As mentioned above, in Appendix 1, I discuss a more general model where the housing market is integrated, and labor supply in a city is not infinitely elastic. The results are similar.

Distribution of the Shocks. In this scenario, the mean nominal wage across the two cities increases for skilled workers only if improvements in amenities are concentrated in cities with high initial wages for skilled workers. The intuition is simple. Under the supply push hypothesis, the nominal wage of skilled workers in San Francisco and Detroit does not change. However, if San Francisco has higher nominal wages to begin with, the shift of skilled workers from Detroit to San Francisco results in higher mean nominal wage across the two cities for skilled workers. In other words, for the supply push scenario to be consistent with the nationwide increase in nominal wage inequality documented in Section 4, cities that experience increases in the relative supply of college graduates in the 1980-2000 period need to be cities with high initial nominal wages in 1980.

6 Interpreting the Evidence: Demand Pull or Supply Push?

6.1 Different Implications for Inequality

The analysis in subsections 5.2 and 5.3 suggests that *for a given nation-wide increase in the nominal wage gap*, the demand pull hypothesis implies a more limited increase in utility inequality, while the supply push hypothesis implies a larger increase in utility inequality.²⁸ The intuition is simple. If college graduates move to expensive cities like San Francisco and New York because of increases in the relative demand for college graduates in these cities—and not because they particularly like living in San Francisco and New York—then part of the increase in nominal wage is offset by the higher cost of living. In this case, the increase in their utility level is smaller than the increase in their nominal wage.

On the other hand, 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.

$$K_{Ha2} - K_{Ha1} = -\frac{\Delta'}{k_a+k_b}.$$

²⁸The ratio of the increase in worker's utility over the increase in mean nominal wage across the two cities is $\frac{Nk_a h}{2(X_{Hb}-X_{Ha})+h(Nk_a+\Delta')}$ for the demand pull case and $\frac{Nk_a h}{X_{Hb}-X_{Ha}}$ for the supply push case. It is clear that the latter is larger than the former.

The two hypotheses are not mutually exclusive since it is possible that cities experience both demand and supply shocks. Moreover, it is even possible that relative demand shifts endogenously generate relative supply shifts, and viceversa. For example, it is possible that an increase in the relative demand for skilled labor in a city results in an increase in the number of college educated residents in that city and this in turns results in increases in the local amenities that are attractive to college graduates, such as good schools, good theaters, good restaurants, etc. Alternatively, it is possible that an increase in the supply of skilled workers in a city generates agglomeration spillovers that lead to increases in the productivity of firms and workers in that city (Moretti 2004a, 2004b). Ultimately, what matters for the interpretation of my findings is whether the combined relative demand shocks are empirically more or less important in driving changes in college share across metropolitan areas than the combined relative supply shocks.

It is important to point out that, while the focus of the paper is on wage inequality, 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 the land in the marginal neighborhood that is gentrified by the inflow of college graduates in cities that experience positive shocks and the marginal neighborhood that is abandoned by the outflow of college graduates from cities that experience negative shocks. A full empirical treatment of this issue is beyond the scope of this paper and is left for future research.

6.2 Demand or Supply?

I now present empirical evidence that seeks to determine which of the two forces—demand or supply—dominates in practice. The analysis in subsections 5.2 and 5.3 suggests that the demand pull and the supply push hypotheses have similar predictions for 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 relative nominal wage of college graduates. Under the supply push hypothesis, there should be no relationship between increases in the share of college graduates and changes in the relative nominal wages. Intuitively, increases in the relative demand of a factor of production in a city should result in increases in its relative price there. Increases in the relative supply of factor of production in a city should cause its relative price to decline, or at least not to increase. (See subsections 5.2 and 5.3 for details.)

I present two pieces of empirical evidence.

(1) First, in Figure 4, I show the empirical relationship between the college share and the 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. 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 in Table 7 quantify the corresponding regression coefficients. 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. While Figure 4 and Table 7 indicate that demand factors are important, they can not rule out that supply shocks are also present.

It is important to point out that the relationship between college premium and college share is *not causal*. Rather, it is an *equilibrium* relationship. 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.²⁹ What I am measuring in Figure 4 and Table 7 is the relationship between the wage gap and the college share, inclusive of any human capital spillover.

(2) As a second piece of evidence that may shed more light on the importance of demand factors 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

²⁹For 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.

the number of college graduates that may be driven by either demand or supply shocks. The comparison of the two estimates is therefore informative on 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:

$$\text{Change in Relative Demand in City } c = \sum_s \eta_{sc}(\Delta E_{Hs} - \Delta E_{Ls}) \quad (16)$$

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 means that at least 44% of the variation in changes in college share can be attributed to *observable* demand shocks. (Of course, there are other demand shocks that are not captured by the instrument.)

The instrumental variable estimate, in column 3 of Table 7, 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.

7 Conclusions

Over the past 25 years, college graduates and high school graduates have experienced different increases in the cost of housing and thus different increases in the overall cost of living. The main contribution of this paper is to document that, as a consequence, the wage difference between college and high school graduates measured in real terms is significantly smaller than the wage difference measured in nominal terms.

Specifically, I show that much of the growth between 1980 and 2000 in the number of college graduates has occurred in metropolitan areas that have a high initial cost of housing and experience large increases in the cost of housing. I propose two skill-group specific cost

of living indexes as an alternative to the nation-wide CPI that is typically used to deflate wages and incomes. My first local cost of living index allows for differential variation in the cost of housing across metropolitan areas and skill groups. My second local cost of living index also allows for variation in the cost of non-housing goods and services.

Using these two local CPI's, I find that the level of the college premium is significantly lower in real terms than in nominal terms. For example, in 2000 the conditional difference between the wage of college graduates and of high school graduates is 60% in nominal terms and only 37%-43% in real terms. Second, and most importantly, the increase in the college premium between 1980 and 2000 in real terms is significantly smaller than the increase in nominal terms. Specifically, the increase in nominal terms is 20 percentage points. The increase in real terms is between 8 and 10 percentage points.

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. 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 relative demand or relative supply shocks are more important empirically, 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.

While I can not completely rule out that supply shocks may play a role, the weight of the evidence seems consistent with the notion that a significant part of the variation in the relative fraction of college graduates across cities is driven by demand factors. If this is true, the increase in well-being disparities between 1980 and 2000 is smaller than the increase in nominal wage disparities and the problem of inequality is less severe than previously thought.

This paper leaves open the question of what ultimately causes the local 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). Another potential explanation centers on shifts in product demand across industries that have different skill intensities. 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 generality of these industries' experience.

In interpreting the findings of this paper, three points are worth highlighting. First, consistent with the previous literature on inequality, the main focus this paper is on *wage* differences. However, the broader distributional consequences of the demand and supply shocks depend not just on changes in relative wages, but also on who originally owns the land in the cities that benefit from the demand and supply shocks. A full empirical treatment of this issue is beyond the scope of this paper.

Second, the return to college is but one measure of wage inequality. Although it has received much attention in the literature on inequality, future research should determine whether the results in this paper extend to other measures of inequality. Wage differences between blacks and whites should also be assessed using local cost of living deflators, since it is possible that differences in geographical distribution expose blacks and whites to different inflation rates.

Finally, while my estimates are based on pre-tax earnings, accounting for federal taxation is likely to make my results stronger. Since the tax system is progressive and is based on nominal income, workers in expensive cities pay more in federal taxes than otherwise identical workers in less expensive cities (Albouy, 2008). Furthermore, state taxes may accentuate this effect, as expensive coastal states in New England and California tend to have more progressive tax systems.

8 Bibliography

Autor, David H., Frank Levy and Richard J. Murnane, 2003. "The Skill Content Of Recent Technological Change: An Empirical Exploration," The Quarterly Journal of Economics, MIT Press, vol. 118(4), pages 1279-1333, November.

Autor, David, Lawrence F. Katz and Alan B. Krueger, 1998. "Computing Inequality: Have Computers Changed The Labor Market?," The Quarterly Journal of Economics, MIT Press, vol. 113(4), pages 1169-1213, November

Albouy, David "The Unequal Geographic Burden of Federal Taxation", NBER WP 13995 (2008).

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

Basker, Emek and Mike Noel, "The Evolving Food Chain: Competitive Effects of Wal-Mart's Entry into the Supermarket Industry" (2007)

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

Black, Dan, Natalia Kolesnikova, and Lowell J. Taylor, "Earnings Functions When Wages and Prices Vary by Location", Federal Reserve Bank of St Louis Working Paper 031A (2007)

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.

John Bound and George Johnson, Changes in the structure of wages in the 1980s: an evaluation of alternative explanations, American Economic Review 82 (1992).

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.

Broda, Christian and John Romalis "Inequality and Prices: Does China Benefit the Poor

in America? mimeo, University of Chicago GSB, (2008).

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*, forthcoming.

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

Glaeser, Edward L. and Gyourko, Joseph E., “Urban Decline and Durable Housing” NBER Working Paper No. W8598 (2001)

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 and Tobio “The Rise of the Sunbelt”, *Harvard University* (2007)

Goldin Claudia and Lawrence F. Katz, “Decreasing (and then Increasing) Inequality in America: A tale of Two Half-Centuries” in “Increasing Income Inequality in America”, 1999.

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

Goldin, Claudia and Lawrence F. Katz, 1999. “The Returns to Skill in the United States across the Twentieth Century,” NBER Working Papers 7126, 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 Where Did the Productivity Growth Go? Inflation Dynamics and the Distribution of Income, *Brookings Papers on Economic Activity*, vol. 36 (2005, no. 2), pp. 67-150 . NBER Working Paper 11842, December 2005.

Greenstone M., R. Hornbeck and E. Moretti “Identifying Agglomeration Spillovers: Evidence from Million Dollar Plants”, *University of California, Berkeley* (2007).

Grogger, Jeff “Changes in College Skills and the Rise in the College Wage Premium.” *Journal of Human Resources* 30, no. 2, Spring 1995, 280-310.

Gyourko, Joseph, Christopher Mayer, and Todd Sinai. *Superstar Cities*, Mimeo, The Wharton School, Revised, June 2006

Gyourko, Joseph “Housing Supply,” *Wharton Real Estate Center Working Paper*, August 2008

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

Hamermesh, Dan "Spatial and Temporal Aggregation in the Dynamics of Labor Demand," in J. van Ours, G. Pfann and G. Ridder, eds. Labor Demand and Equilibrium Wage Formation, North-Holland, 1993.

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.

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)

Krupka, Douglas J. "Location-Specific Human Capital, Location Choice and Amenity Demand," IZA Discussion Papers 2987, Institute for the Study of Labor (IZA) (2007).

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, "Postsecondary Education and Increasing Wage Inequality,. American Economic Review (Papers and Proceedings) 96(2), May 2006, 195-199.

Lemieux, Thomas "The Changing Nature of Wage Inequality" forthcoming in the Journal of Population Economics, 2007.

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.

Lockner, Lance and Christopher Taber, 1999. "Human capital formation and general equilibrium effects" Fiscal Studies, Institute for Fiscal Studies, vol. 20(1), pages 25-40.

Moretti, Enrico "Estimating the Social Return to Higher Education: Evidence From Longitudinal and Repeated Cross-Sectional Data" Journal of Econometrics 121(1-2) (2004a)

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

Neal, Derek The Measured Black-White Wage Gap Among Women Is Too Small Journal of Political Economy, February, 2004, vol. 112, pp. S1-S28.

Neal, Derek and Sherwin Rosen Theories of Labor Earnings, in the Handbook of Income Distribution, edited by A. B. Atkinson and F. Bourguignon, 2000, North Holland.

Neal, Derek, William Johnson and Yuichi Kitamura Evaluating a Simple Method for Estimating Black-White Gaps in Median Wages, American Economic Review, May, 2000, pp. 339-43.

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, 26-1, pp. 23-41 (1988).

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.

Topel, Robert H. "Factor proportions and relative wages: the supply-side determinants of wage inequality", Journal of Economic Perspectives 11 (1997)

Appendix 1

In this appendix I change two assumptions of the model in Section 5. In Section 5, I assume that the housing market for skilled workers is separated from the housing market for unskilled workers in the same city. I also assume that labor supply in a city is infinitely elastic. Here I consider the case where the housing market is completely integrated—i.e. skilled and unskilled workers in a city compete for the same set of houses—and labor supply in a city is not infinitely elastic but rather upward sloping.

Assume that the indirect utilities of skilled and unskilled workers in city c are, respectively: $U_{Hic} = w_{Hc} - r_c + A_{Hc} + \epsilon_{ic}$ and $U_{Lic} = w_{Lc} - r_c + A_{Lc} + \epsilon_{ic}$, where ϵ_{ic} is i.i.d and $\epsilon_{ic} \sim U[-s, s]$. This utility differs from Section 5 in two respects. First, the cost of housing is now the same for skilled and unskilled workers living in city c : r_c . This implies that skilled and unskilled workers face the same housing market within a city, and that shocks to the relative demand or supply of skilled workers will affect unskilled workers in the city through housing costs. Second, the term ϵ_{ic} allows for idiosyncratic preferences for a city. All the other assumptions remain unchanged.

A skilled worker chooses to live in city b if $w_{Hb} - r_b + A_{Hb} + \epsilon_{ib} > w_{Ha} - r_a + A_{Ha} + \epsilon_{ia}$. A similar expression holds for unskilled workers. If there are N unskilled workers and N skilled workers in the economy, the inverse labor supplies of the skilled and unskilled workers in city b are, respectively:

$$w_{Hb} = \frac{N_{Hb}}{q} - \frac{N}{2q} + w_{Ha} - r_a + r_b \quad (17)$$

and

$$w_{Lb} = \frac{N_{Lb}}{q} - \frac{N}{2q} + w_{La} - r_a + r_b \quad (18)$$

where $q = N/2s$. Unlike Section 5, labor supply is upward sloping.

Demand Pull. Consider what happens when the productivity of skilled workers increases relative to the productivity of unskilled workers in San Francisco, and nothing happens to the productivity of skilled and unskilled workers in Detroit. Some skilled workers move from Detroit to San Francisco:

$$N_{Hb2} - N_{Hb1} = \frac{q(q(k_a + k_b) + 1)}{h(2q(k_a + k_b) + 1)} \Delta \quad (19)$$

where I have added subscripts 1 and 2 to denote periods 1 and 2. Some unskilled workers leave San Francisco for Detroit:

$$N_{Lb2} - N_{Lb1} = -\frac{q^2(k_a + k_b)}{h(2q(k_a + k_b) + 1)} \Delta \quad (20)$$

As in Section 5, both total population and the fraction of skilled workers in San Francisco (Detroit) increase (decrease). As before, the nominal wage of skilled workers in San Francisco increases

$$w_{Hb2} - w_{Hb1} = \frac{\Delta}{h} \quad (21)$$

The wage of unskilled workers in San Francisco and the wage of skilled and unskilled workers in Detroit does not change. Rents in San Francisco increase by

$$r_{b2} - r_{b1} = \frac{k_b q}{h(1 + 2q(k_a + k_b))} \Delta \quad (22)$$

while rents in Detroit decline by

$$r_{a2} - r_{a1} = -\frac{k_a q}{h(1 + 2q(k_a + k_b))} \Delta \quad (23)$$

As in Section 5, it is possible to show that the difference between the nominal wage of skilled workers averaged across the two cities and the nominal wage of unskilled workers averaged across the two cities increases. Furthermore, such an increase is larger than the increase in the difference between the real wage of skilled workers averaged across the two cities and the real wage of unskilled workers averaged across the two cities.³⁰

Supply Push. Consider now what happens when San Francisco becomes more desirable for skilled workers relative to Detroit because the amenity that attracts skilled workers increases there. The productivity of both skilled and unskilled workers, as well as the amenity level in Detroit, do not change. The number of skilled workers who move from Detroit to San Francisco is

$$N_{Hb2} - N_{Hb1} = \frac{q(q(k_a + k_b) + 2)}{2q(k_a + k_b) + 1} \Delta' \quad (24)$$

The number of unskilled workers who leave San Francisco to Detroit is

$$N_{Lb2} - N_{Lb1} = -\frac{q(q(k_a + k_b) - 1)}{2q(k_a + k_b) + 1} \Delta' \quad (25)$$

The nominal wages of skilled and unskilled workers do not change. Rents in San Francisco increase by

$$r_{b2} - r_{b1} = \frac{(q(4k_b - 2k_a) - 1)}{2(1 + 2q(k_a + k_b))} \Delta' \quad (26)$$

while rents in Detroit decline by

$$r_{a2} - r_{a1} = -\frac{(q(4k_a - 2k_b) - 1)}{2(1 + 2q(k_a + k_b))} \Delta' \quad (27)$$

As before, the difference between the nominal wage of skilled and unskilled workers averaged across the two cities increases more than the difference between the real wage of skilled and unskilled workers averaged across the two cities, if the elasticity of housing supply in San Francisco is not much larger than the elasticity of housing supply in Detroit.

³⁰As before, this is true if the elasticity of housing supply in San Francisco is not much larger than the elasticity of housing supply in Detroit.

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	249	418	532	113%
College	298	570	761	156%
Percent Difference	19%	36%	44%	
<u>Local CPI 1</u>				
High-School	0.97	1.45	1.86	91%
College	1.04	1.65	2.19	111%
Percent Difference	7%	14%	18%	
<u>Local CPI 2</u>				
High-School	0.96	1.49	1.91	98%
College	1.07	1.83	2.42	127%
Percent Difference	11%	23%	27%	

Notes: Monthly rent refers to the rent paid for a two or three bedroom apartment. CPI 1 allows for local variation only in the cost of housing. 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

	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 (.001)	.47 (.001)	.53 (.001)	.18
<u>Model 2</u>								
Real Wage Difference - Local CPI 1	.33 (.006)	.41 (.005)	.43 (.004)	.10	.32 (.001)	.39 (.001)	.42 (.001)	.10
Percent of Nominal Accounted for by Cost of Living	15%	23%	27%	50%	9%	17%	17%	44%
<u>Model 3</u>								
Real Wage Difference - Local CPI 2	.29 (.005)	.34 (.005)	.37 (.005)	.08	.29 (.001)	.34 (.002)	.37 (.003)	.08
Percent of Nominal Accounted for by Cost of Living	27%	36%	39%	60%	11%	27%	30%	55%
MSA Fixed Effects	No	No	No	No	Yes	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: Nominal and Real Conditional Wage Difference Controlling for Quality of Housing, by Year - American Housing Survey

	1980	1990	2000	1980-2000 Increase
	(1)	(2)	(3)	(4)
Nominal Wage Difference	.36 (.024)	.48 (.008)	.55 (.011)	.19
Real Wage Difference - Not Controlling for Quality				
Real Wage - Local CPI 1	.29 (.010)	.39 (.008)	.43 (.013)	.14
Real Wage - Local CPI 2	.25 (.010)	.33 (.008)	.38 (.012)	.13
Real Wage Difference - Controlling For Quality				
Real Wage - Local CPI 1	.30 (.016)	.37 (.011)	.39 (.014)	.09
Real Wage - Local CPI 2	.27 (.03)	.32 (.015)	.33 (.019)	.06

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 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. Sample size is 237,210.

Table 6: Robustness

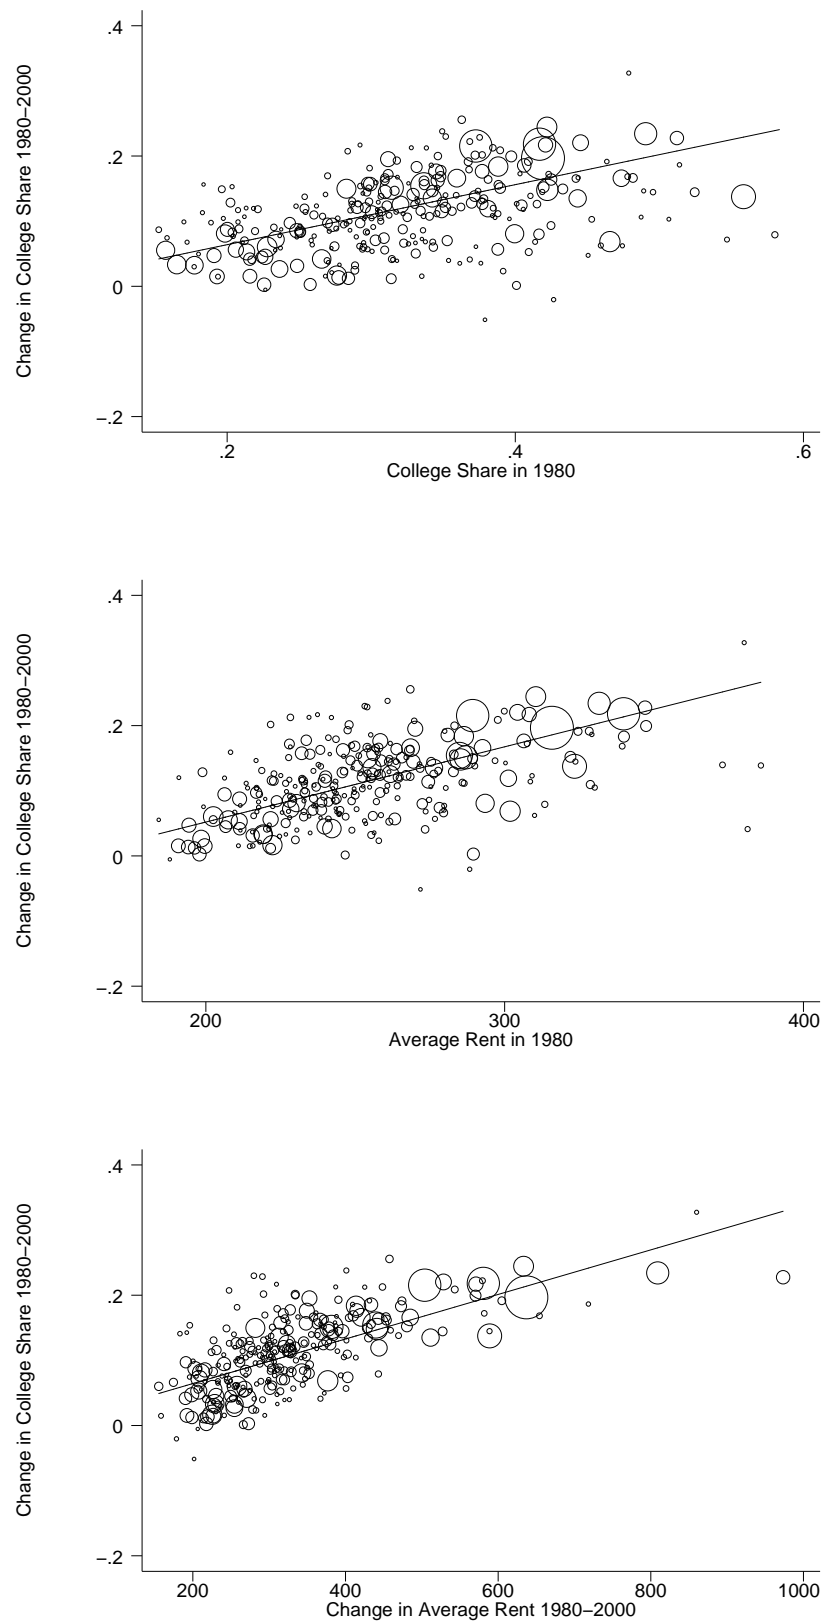
	1980	1990	2000	1980-2000 Increase
	(1)	(2)	(3)	(4)
<u>Model 1: Housing Cost by MSA, Skill, Race, N. of Children</u>				
Nominal Wage Difference	.39 (.008)	.53 (.012)	.59 (.013)	.20
Real Wage - Local CPI 1	.33 (.006)	.41 (.005)	.44 (.004)	.11
Real Wage - Local CPI 2	.29 (.006)	.34 (.006)	.38 (.005)	.09
<u>Model 2: ACCRA Non-Housing Prices</u>				
Nominal Wage Difference	.40 (.015)	.53 (.009)	.60 (.010)	.20
Real Wage - Local CPI 1	.33 (.006)	.41 (.005)	.43 (.004)	.10
Real Wage - Local CPI 2	.24 (.005)	.27 (.008)	.29 (.009)	.05
<u>Model 3: Commuting Time</u>				
Nominal Wage Difference	.40 (.011)	.54 (.011)	.60 (.013)	.20
Real Wage - Local CPI 1	.33 (.006)	.41 (.004)	.44 (.004)	.11
Real Wage - Local CPI 2	.31 (.006)	.37 (.005)	.40 (.004)	.09
<u>Model 4: Include Immigrants</u>				
Nominal Wage Difference	.40 (.011)	.54 (.012)	.61 (.013)	.21
Real Wage - Local CPI 1	.34 (.007)	.41 (.006)	.45 (.005)	.11
Real Wage - Local CPI 2	.30 (.006)	.34 (.006)	.38 (.005)	.08
<u>Model 5: Only Urban Workers</u>				
Nominal Wage Difference	.40 (.011)	.52 (.008)	.59 (.010)	.19
Real Wage - Local CPI 1	.33 (.007)	.41 (.005)	.44 (.004)	.11
Real Wage - Local CPI 2	.29 (.006)	.35 (.006)	.38 (.006)	.09

Table 7: 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	

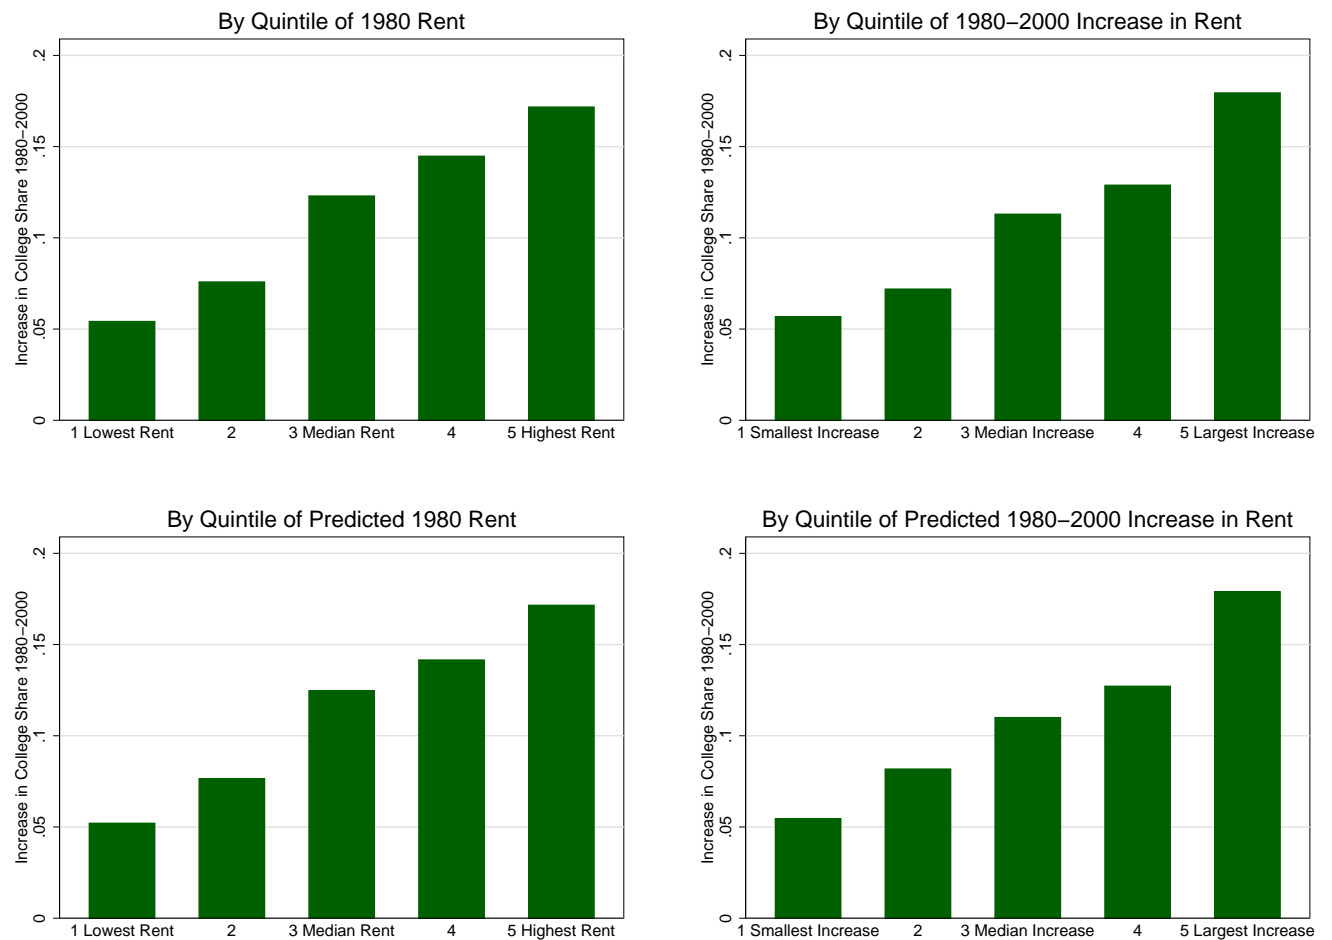
Notes: Standard errors in parentheses. The dependent variable in column 1 is the city-specific college premium, defined as the city-specific difference in the log of hourly wage for college graduates and high school graduates conditional on gender, a cubic in potential experience, race and year. The 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.

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



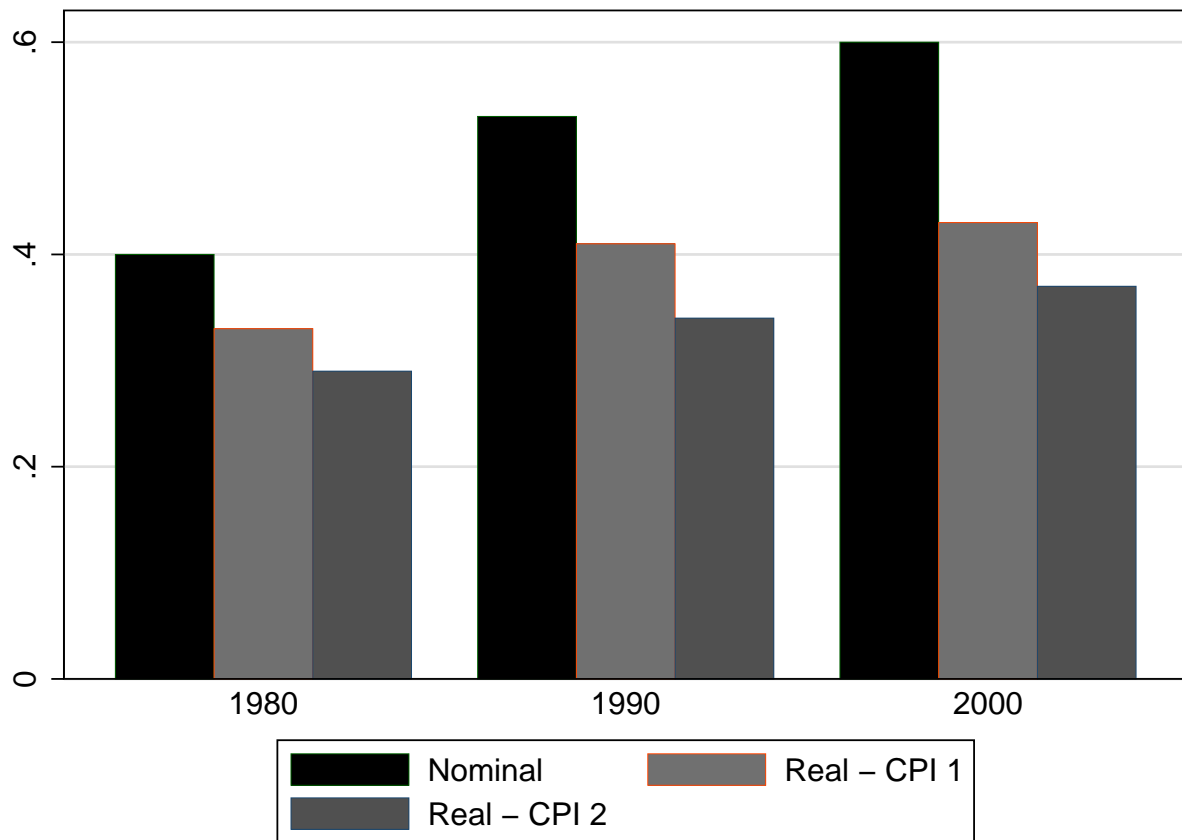
Notes: Average rent is the average monthly rental price of a two or three bedroom apartment.

Figure 2: How Changes in the Share of College Graduates Relates to the Cost of Housing for College Graduates and Changes in the Cost of Housing for College Graduates



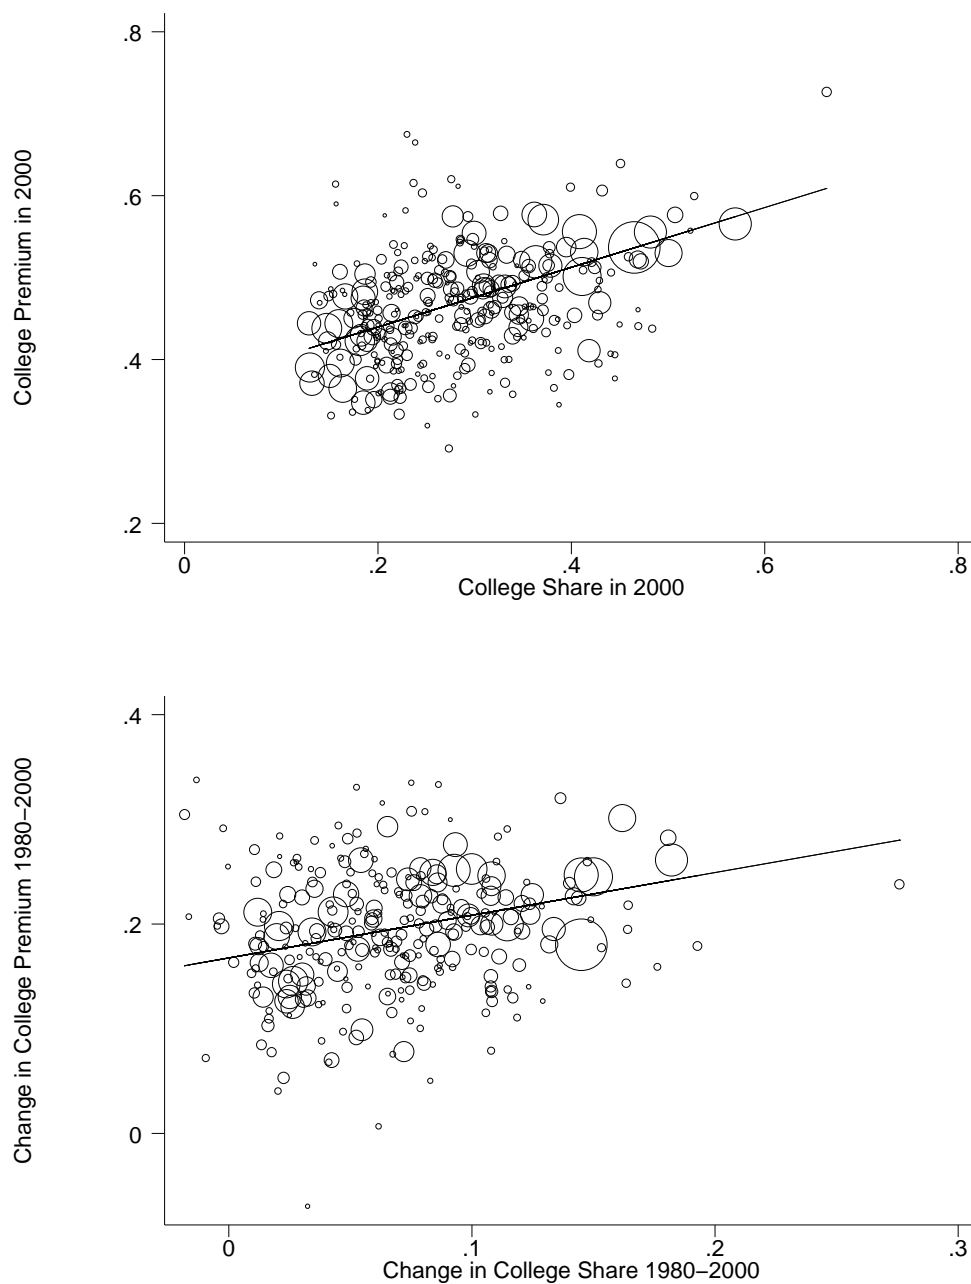
Notes: The top panels show the 1980-2000 increase in the share of college graduates, by quintile of the 1980 cost of housing (left panel) and by quintile of the 1980-2000 increase in the cost of housing (right panel). Cost of housing is measured as the average monthly rent paid by college graduates for a 2 or 3 bedroom apartment in the relevant city. The bottom panels show the 1980-2000 increase in the share of college graduates, by quintile of the predicted 1980 cost of housing (left panel) and by quintile of predicted 1980-2000 increase in cost of housing (right panel). The predicted cost of housing is the fitted value from a regression of rental cost on identifiers for metropolitan area, education group, number of children, race and interactions. This regression is estimated on the sample of renters of 2 or 3 bedroom apartments, and the predicted values are calculated for all households.

Figure 3: Nominal and Real Returns to College



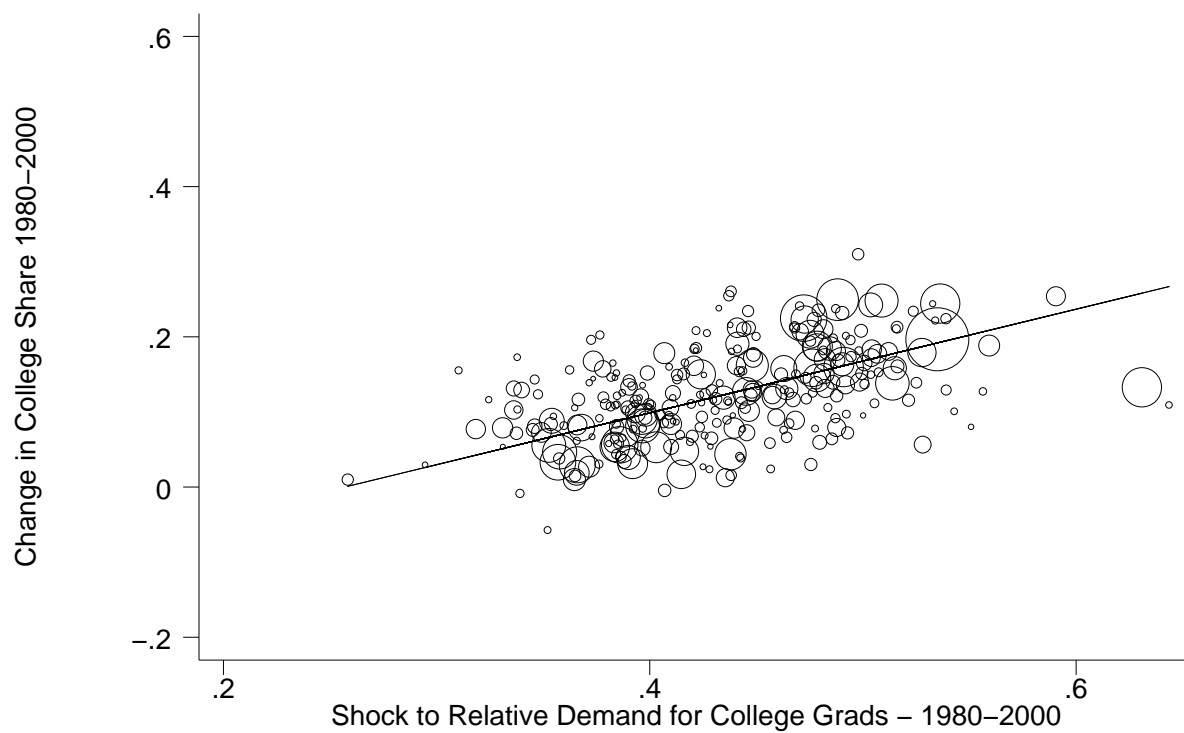
Notes: The figure plots the conditional difference between the average earnings of college graduates and the average earnings of high school graduates measured in nominal and real terms. The height of the bars is equal to the regression coefficients from models in Table 4 (columns 1 to 3). CPI 1 allows for local variation only in the cost of housing. CPI 2 allows for local variation both in the cost of housing and the cost of non-housing goods and services.

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 16.