Three Essays in Empirical Labour Economics

Wage Determination in Local Labour Markets

by

Benjamin MacLean Sand

B.Ec. (Hons), The University of Calgary, 2002
M.A., The University Calgary, 2003

A THESIS SUBMITTED IN PARTIAL FULFILMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

The Faculty of Graduate Studies

(Economics)

The University of British Columbia

(Vancouver)

July 2009

Abstract

This dissertation consists of three empirical essays that examine different aspects of wage determination in local labour markets. The first essay investigates whether or not there are human capital externalities or spill-overs from education. I find that the fraction of college graduates in U.S. cities is associated with higher wages in the 1980s but not in the 1990s. To rationalize this pattern, I empirically investigate a model of structural change by Acemoglu (1999) and find considerable support for it in a number of dimensions. Consistent with the notion that there has been a structural change in the labour market, increases in the supply of skilled labour in the 1990s induce a change in the composition of jobs, increase inequality, unemployment, the return to education, and the wages of high-skill workers and *harm* low-skill workers. The second essay, which is co-authored with Paul Beaudry and David Green, develops a multi-sector search and matching model of the labour market that illustrates a mechanism through which changes in local industrial composition can cause changes in wages in all sectors of the local economy. We empirically test this model using geographical variation in industrial composition across U.S. metropolitan areas from 1970 to 2000 and find that shifts in industrial composition that favor high-paying industries impact wages in other sectors in a manner that is consistent with the model. The third chapter, co-authored with Christopher Bidner, extends the model developed in chapter two to examine the impact of changes in industrial composition on the relative wages of men and women. We find that men lost representation in high-paying industries relative to women and that these losses can account for a substantial portion of the 'unexplained' gender pay gap. All three essays use data from the U.S. decennial Censuses and take U.S. metropolitan areas as local labour markets.

Table of Contents

Appendices

List of Tables

List of Figures

Acknowledgements

This thesis has benefited greatly from interaction and support from many individuals. I would like to thank my supervisor, David A. Green, for his support and friendship throughout my entire time at the University of British Columbia. His guidance and encouragement benefited me both academically and personally, and I am extremely grateful to have him as a friend and adviser. I would also like to thank Paul Beaudry and Thomas Lemieux for many helpful discussions and exceptional advice. I am indebted for their intellectual contributions and generosity in providing their time. Many friends have also helped along the way. I would like to thank Christopher Bidner, Kelly Foley, and Ken Jackson for many good times, support and encouragement, both personally and academically. Finally, I would like to thank my Mother and Father, and Allison Metzak for their unwavering support.

Chapter 1

Has There Been A Structural Change In the Labour Market: Evidence from U.S. Cities

1.1 Introduction

Are the benefits of acquiring additional education confined to those individuals that obtain it? Unfortunately, there is little empirical consensus regarding either the existence or the magnitude of human capital externalities. On one hand, studies that use variation in lower educational attainment to identify external returns to education find little evidence of education spill-overs (Acemoglu and Angrist 1999). On the other, studies using variation in the proportion of college graduates in localities find external returns (Moretti 2004). One emerging view is that there are education spill-overs to post-secondary education but not to secondary education (Iranzo and Peri 2006). However, I will show that, in spite of this building consensus, a close inspection of the data reveals that there are good reasons to re-examine the impact of changes in the supply of highly educated workers on local labor market outcomes.

As I demonstrate below, the external return to post-secondary education, estimated as the effect of the share of college graduates in a city on wages after controlling for the private effect of education, is remarkably unstable over time. In particular, using city level U.S. Census data, I find large, positive spill-overs from college education in the 1980s, which is consistent with previous literature. However, when I focus on the 1990s, I find that the relationship between city college share and wages changes dramatically. When looking at the impact of the share of college workers on the wages of all workers, it appears that there is no evidence of spill-overs. This result hides impacts that are heterogeneous by education level. In particular, I find that an increase in the number of high-skill workers in a locality negatively effects the wages of low-skill workers. One contribution of this paper is to bring attention to this pattern and document that it is robust to a variety of econometric concerns. Given that this result is robust and not in accordance with standard production functions or views regarding human capital spill-overs, it is natural to ask how best to explain it.

In order to understand the pattern of education spill-overs observed in the data, I examine Acemoglu's (1999) model of endogenous job composition in which firms design jobs based on the skill level of workers they expect to meet. This model is consistent with the observed over-time pattern of education spill-overs and also has a number of additional implications regarding labor market outcomes. Using my data on U.S. cities, I empirically test these predictions and find considerable support for them in a variety of dimensions. This analysis has implications for how we interpret human capital externalities and several recent labor market developments. The main finding of this paper is that it appears we have witnessed a structural change in the labor market between the 1980s and 1990s that has harmed low-skill workers. I argue that, while technical change has undoubtedly played a role, the key to understanding this process is the supply of skilled workers, which has induced a change in the composition of jobs, increasing inequality and unemployment.

In order to study the social returns to education, I use U.S. census data from 1980 to 2000 and exploit geographical variation in the fraction of college graduates across U.S. cities. I follow the literature in identifying education spill-overs by the impact of aggregate education on individual wages, while controlling for individual education. To do this, I examine decadal changes in city level wages using data from the 1980, 1990, and 2000 U.S. censuses for 286 consistently defined metropolitan areas. The key covariate in this analysis is the change in the fraction of a city's labor force with at least a college degree. After controlling for the individual effect of education, the coefficient on city college share will reflect the presence of human capital spill-overs in ideal conditions.

Working in decadal differences allows me to control for possible correlations between city time-invariant factors and changes in college share by using within-city variation. However, there are still potential endogeneity concerns arising from omitted time-varying variables that are correlated with changes in the fraction of college graduates in a city. I address these concerns by using an instrumental variables strategy to predict changes in the local proportion of college graduates. The first instrument that I use is based on Doms and Lewis (2006) and borrows from the immigration literature that uses enclaves, or previous immigration settlement, to predict future immigration flows. Specifically, I predict changes in city college share using the distribution of immigrants in 1970 as enclaves to allocate immigration flows from sending countries with different average levels of educational attainment. As a second source of exogenous identifying variation, I use a set of long-term city average climate variables as instruments. The intuition for this instrument set comes from the literature that shows that the migration of college graduates is at least partly influenced by local amenities. The validity of this instrument set hinges on the assumption productivity differences between cities arising from long-term climate patterns are addressed by the differencing procedure. I also address the potential non-random location of workers across cities using a method first developed by Dahl (2002) and extended by Beaudry, Green, and Sand (2007) to account for the self-selected location of foreign born workers.

Despite addressing several econometric concerns and focusing on variation in higher education, which has been argued to explain the mixed results found in the literature, I find that spill-overs from post-secondary education are remarkably unstable between the 1980s and the 1990s. In the 1980s, spill-overs to college education are large and benefit all workers, while in the 1990s, spill-overs vary by education level and are even negative for workers with the least education. This is a puzzling result when viewed from standard models of technological externalities, in which externalities are built directly into the aggregate production function. Externalities in these models are usually motivated by the sharing of knowledge and skills between workers, so it would be hard to understand why such worker interaction would become unimportant in the 1990s. In fact, it has been suggested that if increasing private returns to education reflects changes in the social value of human capital, we might even expect to see larger spill-overs (Moretti 2004). Indeed, both the returns to a college degree and the fraction of the workforce with a college degree grew over the 1990s (see, for example, Lemieux (2006)). In light of these observations, it appears that something has changed in the 1990s relative to the 1980s and that a new or expanded explanation is needed to reconcile these facts.

It is not surprising that models of technological externalities do not help us understand the change in estimated spill-overs between the 1980s and the 1990s because the equilibrium is unique making outcomes smooth on the parameter space. In a second class of theories of human capital externalities, spill-overs to education are pecuniary rather than being assumed to be part of the production function. In these models, externalities arise through market interaction and small changes in the underlying parameters can produce equilibria that display significant qualitative differences. This important difference has been overlooked by the empirical literature because the models produce otherwise similar empirical relationships. Given the apparent observed instability in the estimated social returns to education, this difference seems particularly important. This suggests that models of pecuniary externalities that incorporate structural breaks, such as Acemoglu (1999), might be a valuable way to interpret the patterns in the data.

Acemoglu's (1999) model is one of endogenous job composition in which firms create jobs and then search for workers. When the fraction of skilled workers or the productivity gap between skilled and unskilled workers is low, firms pool across workers and create only one type of 'middling' job. In this *pooling* equilibrium, the quality of the middling job depends on the fraction of skilled workers in the economy. In particular, in response to an increase in the probability of being matched with a skilled worker, firms invest more. Therefore, both skilled and unskilled workers end up working with more capital and there is a pecuniary externality as in Acemoglu (1996). However, if the fraction of skilled workers or the productivity gap between skilled and unskilled workers becomes large enough, firms may find it profitable to create jobs specifically designed for skilled and unskilled workers. In this *separating* equilibrium, the change in the composition of jobs increases the wages of skilled workers, lowers the wages of unskilled workers, and increases inequality and unemployment.

I examine the predictions of this model under the assumption that labor markets were characterized by a pooling equilibrium in the 1980s, but during the 1990s the supply of skilled workers and technical change had reached a point where cities moved toward a separating equilibrium. Under this assumption, finding positive education externalities only in the 1980s arises naturally in the context of this model. However, I also demonstrate that the additional predictions of this model are remarkably consistent with the data. In particular, I provide evidence that changes in the composition of jobs, inequality, wages and unemployment are non-monotonically related to the supply of skilled workers between the 1980s and the 1990s. I also present evidence that the most highly educated cities in the 1980s may have already switched to a separating equilibrium and estimate the threshold level of education in which this switch occurs.

The main message of this paper is that there appears to have been a structural change in the labor market between the 1980s and the 1990s that has harmed the least skilled workers. It should be emphasized that the driving forces behind this change are empirically plausible. For example, there has been a large increase in the fraction of skilled workers in the U.S., and there has been a great deal of technical change that is widely perceived to be skill-biased (Acemoglu 2001). In addition, a growing number of studies either view these two decades as distinct periods (Autor (2007), Autor, Katz, Kerney (2007), Lemieux (2007)) or as a period of transition (Beaudry and Green 2005). Taken as a whole, there is mounting support for the view that there has been a change in the organization or mode of production over this period, which is consistent with the findings presented here.

This paper is related to a recent and growing literature on changes in job composition or 'polarization' during the 1990s. Autor, Katz, and Kerney (2006) and Goos and Manning (2007) argue that increased employment in low skilled (non-routine manual) jobs and very high skilled (non-routine cognitive) jobs is caused by a form of skill biased technical change that 'polarizes' the demand for skills. In fact, Goos and Manning (2007) mention the model of Acemoglu (1999) as a potential alternative to their interpretation of the data. Autor and Dorn (2007), Manning (2004), and Mazzolari and Ragusa

(2007) suggest that growth in low-skilled service employment is related to the location of highly educated or high income workers. These works emphasize demand forces in explaining changes in employment composition in contrast to the non-competitive mechanism suggested here. In my empirical work, I take care to differentiate these competing explanations for the observed change in employment composition.

The remaining sections of this paper are organized as follows. In section 2, I give a brief review of the literature on the social returns to education and reevaluate the evidence on education spill-overs in U.S. cities from 1980-2000. In section 3, I present Acemoglu's (1999) model of endogenous job composition and in section 4 I empirically examine its predictions. Section 5 concludes.

1.2 A Reexamination of the Social Returns to Human Capital: 1980-2000

I begin this section by presenting the empirical observation that motivates the remainder of the paper. The data for this analysis comes from U.S. Censuses for the years 1970-2000 extracted from IPUMS.[1] All analysis is restricted to individuals between the ages of 18 and 64 with positive weeks worked in the year proceeding the Census. Real wages are obtained by dividing wage and salary income by weeks worked and deflating by the CPI. Appendix A contains additional details on the processing of the Census data, including the construction of consistent city definitions over time.

The observation that I want to emphasize is that the relationship between changes in average city wages and aggregate human capital has changed considerably between the 1980s and the 1990s. In particular, I find that in the 1980s there is a strong positive association between city average wages and college share, as previously documented by Moretti (2004). However, this relationship disappears in the 1990s. Below, I argue that this change is surprising in both magnitude and direction and is not consistent with technological externalities as they are normally interpreted.

Consider Figure (1.1) which plots city level changes in wages against changes in the fraction of workers with a college degree for each decade.[2] The first panel contains the results for the 1980s and shows the strong positive association between these two variables. The solid line is the fitted regression line and has a coefficient on college share of 2.01 with a standard error of 0.26, and, as can seen in the figure, this relationship is

[1]See Ruggles, Sobek, Alexander, Fitch, Goeken, Hall, King and Ronnander (2004) for a description of this data.

[2]Average wages in each year and city have been reweighted to hold the distribution of observable characteristics in each city at their 1990 levels, to avoid changes in average wages due to composition effects. See figure notes for additional details on its construction.

Figure 1.1: Change in Average City Wages and College Share

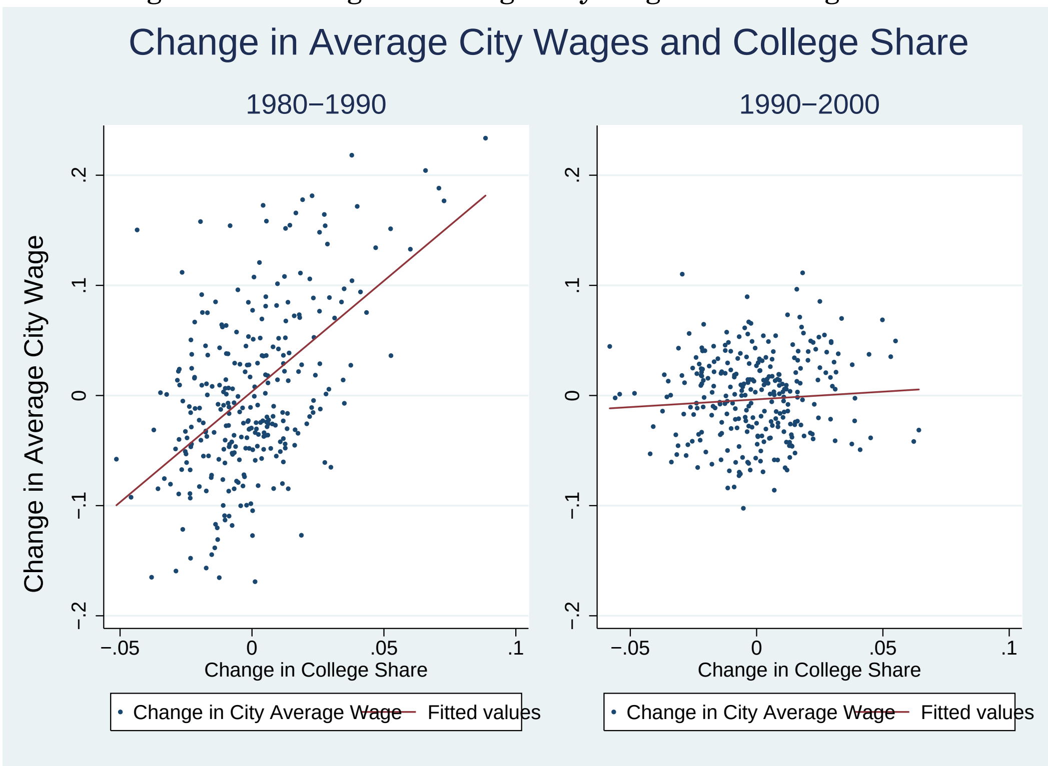

Notes: The change in average wage is calculated by holding the distribution of observable characteristics in each city constant at their 1990s levels via a propensity score reweighting method to avoid composition effects. All variables in the figure have been demeaned, and so represent deviations from the yearly average. The fitted OLS line uses the square root of city size as weights.

not driven by a subset of outlying cities. The next panel shows that in the 1990s the positive association between changes in city average wages and city college share is no longer present. The regression line in this panel has a coefficient on college share of 0.14 with a standard error of 0.13.

Given that this observation is rather striking, it seems necessary to assess its robustness. This is especially relevant since there are a number of well known difficulties in estimating relationships such as in Figure (1.1). Below I discuss these issues and my empirical strategies to deal with them. But first, in order to focus the discussion, I relate the above finding to what has been found in previous literature

1.2.1 Previous Literature

Much of the existing literature that estimates productivity spill-overs from human capital attempts to link a measure of aggregate education in a local labor market to wages. Most studies adopt a similar empirical specification that regresses wages on aggregate education while controlling for individual characteristics such as education. Consider the following representative estimating equation:

$$\ln w_{ict} = X'_{ict}\mu_t + \beta S_{ct} + Z'_{ct}\alpha_t + \delta_t + \delta_c + u_{ict}, \tag{1.1}$$

where the dependent variable is the log wage for individual i in city c in time t, X_{ict} is a vector of individual characteristics including individual education, Z_{ct} is a vector of city level time varying characteristics, δ_t and δ_c are year and city fixed effects, and u_{ict} is an error term. S_{ct} is a measure of city level aggregate education and β is the term reflecting human capital spill-overs under ideal conditions. This strand of literature has come to considerable disagreement regarding the magnitude of human capital spillovers.[3] For example, Lange and Topel (2004), Ciccone and Peri (2002), Rudd (2000) and Acemoglu and Angrist (1999) all find little evidence of external effects of education, while Moretti (2004) and Iranzo and Peri (2006) find positive externalities associated with the share of college graduates in a local labor market.

There are several econometric problems to estimating β, and the literature can be broadly grouped by how researchers deal with the issue of endogeneity of local human capital. Earlier works treat aggregate education as exogenous while later ones use instrumental variables to identify a causal effect of aggregate education on local wage outcomes. Two well known papers among this latter category are Acemoglu and Angrist (1999) and Moretti (2004).

Acemoglu and Angrist (1999) use U.S. states as their geographical unit of analysis and identify social returns to education by using variation from compulsory schooling laws. They find that more restrictive laws are associated with higher average educational attainment. When they instrument state average education using these laws their estimates of the external impact of an additional year of average state education are typically around 1-2% and not significantly different from zero. Acemoglu and Angrist interpret these findings as providing little support for sizable human capital externalities over the range of variation induced by compulsory schooling laws. In contrast to this, Moretti (2004) finds significant spill-overs to education using U.S. cities as his geographical unit of interest. In this study, the measure of local human capital is the proportion of workers in a city with a college degree. Moretti instruments this variable with lagged city age structure or the presence of a land grant college. He finds that a 1% increase in college share raises average city wages by 1.9% for high school gradutes and 0.4% for college graduates.

The difference between these findings might be attributable to several factors. First, Moretti (2004) finds that when states rather than cities are used as the level of analysis,

[3] See Moretti (2003) or Harmon and Oosterbeek (2000) for a summary of the literature.

the estimates are typically lower. Second, and the factor that has received the most attention in the literature, is that Acemoglu and Angrist (1999) identify spill-overs using variation from compulsory school laws that affect the lower part of the education distribution, while Moretti (2004) uses variation in the share of college graduates or the upper part of the education distribution. One emerging view is that externalities are associated with college graduates, while, for example, increasing the fraction of high school graduates has no discernible effect on wages.[4]

The last difference between these studies, and the one that I emphasize in this paper, is that they study different time periods. Acemoglu and Angrist (1999) mainly focus on the period from 1960-1980 due to data considerations[5], while Moretti (2004) studies only the 1980s. However, when Acemoglu and Angrist (1999) extend their data to include the 1980s, several of their estimates of education spill-overs are positive and statistically significant. They interpret this result as being caused by a change in the way that education is coded in the U.S. census, while Moretti (2004) argues that this had little impact (p. 37). This latter detail, the sensitivity to time period, has not been examined closely in the literature and I use this as a starting point in my empirical work below.

1.2.2 Setting up the Puzzle: A Reexamination of Social Returns to Higher Education

To fix ideas, consider a model suggested by Moretti (2003, 2004) to understand the effect of an increase in the fraction of college workers in a city. Each city is treated as a competitive economy that produces a single good that is traded on the national market. Output comes from a CES production function $F(\theta_U U, \theta_S S)$ where U and S refer to the quantities of unskilled and skilled labor, respectively, and the θs are labor augmenting factors. Spillovers are technological in nature and come from letting each groups' productivity shifter depend on the fraction of skilled workers: $\ln(\theta_j) = \psi_j + \beta(\frac{S}{U+S})$ for $j \in \{U, S\}$. If there are spill-overs to increasing the fraction of skilled workers, then $\beta > 0$. This model implies that an increase in the fraction of skilled workers will (1) increase average wages, and (2) increase the wages of unskilled workers more than for skilled workers, due to imperfect substitution.[6]

As mentioned previously, Moretti (2004) finds that $\beta > 0$ using the fraction of the workforce with a college degree as the measure for local human capital and focusing on

[4]For example, Iranzo and Peri (2006) estimate models in which they regress wages on both average years of education and the share of college graduates in U.S. states that supports this view. Wheeler and La Jeunesse (2006) take this as the point of departure in their study on income distribution in U.S. metropolitan areas.

[5]Acemolgu and Angrist (1999) also treat individual schooling as endogenous and use the quarter of birth as an instrument. This variable is not available in other census years.

[6]See Moretti (2004) for details.

the period from 1980 to 1990. At the time of writing that paper, the 2000 Census was not available. As a first step in assessing this view of human capital spill-overs, I estimate versions of (1.1) separately for the 1980s and the 1990s.

The actual estimation procedure that I use throughout is a common two-stage procedure. In the first stage, I regress log wages on a vector of individual characteristics separately for each Census year, forming city averages of the residuals from the regression.[7] I then use differences in these average regression adjusted wages in an estimating regression that takes the form

$$\Delta \log w_c = \Delta Z_c \cdot \alpha_t + \Delta BA_c \cdot \beta + \Delta \epsilon_c, \tag{1.2}$$

where Z_c is a set of time-varying city level controls and BA_c is the fraction of workers in a city with a college degree.[8]

Table (1.1) contains the estimates of equation (1.2) for various sample periods and education groups. The first two columns of Table (1.1) refer to the OLS estimates for each decade. Each row of the table corresponds to a different education subsample for the dependent variable: All workers, High School, or College Graduates. Each entry of the table displays the estimated β from separate regressions. For example, in row (1) column (1) the estimated β is 1.52 and is estimated using all workers for the period 1980-1990. This corresponds to the period studied by Moretti (2004), and corroborates his result: the fraction of college educated workers in a city is significantly correlated with average city wages in the 1980s.[9] According to this estimate, a one percentage point increase in the share of college graduates in a city is associated with an increase in wages of 1.52%.

The second column contains the estimate of β for the 1990s, which is small, negative, and not statistically significant. The implication is that if one only looked at the 1990s, they would conclude that there were no spill-overs to education - even when using college share as the human capital measure. If OLS provides consistent estimates of education spill-overs, these results imply that there were positive and significant spill-overs in the

[7] I take a very flexible approach in specifying the first stage equation. I include indicators for education (4 categories), black, female, and foreign born and a quartic in potential experience, as well as interactions between education and the black, female, and foreign born indicators, and potential experience. I have also experimented with controls for industry affiliation, but these have little impact on the estimated results.

[8] The time-varying city controls are the fraction of the workforce that are female, black, hispanic, and U.S. citizens and the unemployment rate. These are the controls used in Moretti (2004). All regressions use these controls unless otherwise noted. I have also experimented with industrial composition controls, including the Katz-Murphy index used in Moretti (2004), and they do not appreciably alter the patterns presented below.

[9] The estimate in Moretti (2004) that is the most comparable to the one reported here is 1.32 with a standard error of 0.11. Despite using the same dependent and control variables, my estimates are slightly higher. It is worth noting that when I restrict the age of my sample to include only those 25 or older, as in Moretti (2004), I obtain an estimate of 1.37 with a standard error of 0.16.

1980s but not in the 1990s.[10] This is a puzzling result if spill-overs are thought to arise from productivity linkages between individuals, such as the sharing of knowledge and skills. Since the private return to education continued to rise throughout the 1990s, it does not seem very compelling to think that the social value of knowledge and skills declined at the same time. In addition, as many economists believe, skill-biased technical change continued and may have even accelerated during the 1990s (Card and DiNardo 2002). If anything, *a priori* reasoning would suggest that these trends might go together to make education spill-overs even more important.

The second two rows of Table (1.1) display the results for high school and college graduate workers separately. According to the neoclassical model outlined above, the magnitude of the coefficient on college share should be larger for lower skilled (high school) than for highly skilled (college graduates) workers due to imperfect substitution. The second and third rows of Table (3.1) show that this is the case for the point estimates during the 1980s (although, they are not statistically different from each other). The OLS estimates indicate that a one percentage point increase in the fraction of college graduates is associated with a 1.6 percent increase in the wages of high school graduates, while the estimate is 1.5 for college graduates. In the 1990s, however, the estimated impact of college share on the wages of high school and college graduates is quite different. For high school workers, the results indicate that increases in college share were associated with *declines* in wages. That is, there appears to be a negative externality for this group which is a perverse result from the stand point of the model outlined above. For college graduates, the association between wages and college share is significant and positively signed. This pattern does not fit with the standard model above, at least for the 1990s.

1.2.3 Selection and Endogeneity

The OLS results above indicate that the nature of spill-overs from college education changed substantially between the 1980s and the 1990s, from benefiting all workers in the 1980s to potentially having negative impacts on low-skill workers in the 1990s. It could be the case, however, that any association between wages and college share reflects endogeneity or the effects non-random selection of workers across cities, rather than spill-over effects from highly educated workers. Thus, a key concern in being able to consistently identify β from equation (1.2) is omitted variables bias arising from correlation between changes in city college share and the error term, $\Delta\epsilon_c$. I consider two

[10]With the emphasis that the literature has placed on using variation in higher levels of education to identify spill-overs, one might wonder if this same pattern holds when focusing on a city's share of post-graduates as the relevant measurement of aggregate human capital. The same pattern does hold and, in fact, is even more extreme.

potential sources of this bias. The first is possible correlation stemming from city-year shocks to the labor market that could, for example, increase both wages and the demand for skilled workers. Also, to the extent that college educated workers are more mobile than less educated individuals, cities experiencing growth could attract relatively more highly educated workers seeking to take advantage of better economic opportunities.

Another prospective source of bias is self-selected migration of workers across cities. For example, if changes in unobserved (to the econometrician) ability across cities is correlated with city educational attainment as well as city average wages, then, in general, an OLS estimate of β will be inconsistent. Suppose, for instance, that cities differ in the demand for some unobserved ability and that workers locate based on comparative advantage such as in a Roy model. Then, as seems likely, if cities that have relatively higher demands for education also have higher demands for ability, correlation between college share and wages could reflect worker sorting instead of spill-overs (Moretti 2004).

Instrumental Variables: Methods and Results

I address the first issue by using an instrumental variable strategy. I have two arguably exogenous sources of variation in ΔBA_c. The first is an instrument based on Doms and Lewis (2006) that has its origins in the immigration literature that use "enclaves" to predict immigration flows. The idea is based on the observation that recent immigrants tend to locate in areas where earlier waves of immigrants from the same home country had settled. Doms and Lewis (2006) combine this with the fact that immigrants from different sending countries have different average education levels to create a variable that is the expected change in a city's college graduate share if immigrants continue to arrive in the same proportion as in the past, and there are no other changes to a city's population.

More specifically, as in Doms and Lewis (2006), let $\theta_{h,c,70}$ represent the fraction of immigrants from sending country h who have settled in city c as of 1970. Let $N_{I,h,t}$ be the number of immigrants from country h that arrived during time t (in practice, t corresponds to the last ten years) at the national level. Then the predicted inflow of immigrants to city c is $\hat{N_{I,c,t}} = \sum_h \theta_{h,c,70} N_{I,h,t}$.

Using the same idea, let $\hat{N}^{BA}_{I,c,t}$ be the expected number of immigrants into city c over the time period t with a bachelor's degree (BA) or greater that is predicted in the same way except using inflows of college graduates. Then the instrument can be written as

$$IV^{BA}_{c,t} = (\hat{S}^{BA}_{I,c,t} - S^{BA}_{c,t-1}) \cdot \frac{\hat{N}_{I,c,t}}{\hat{N}_{I,c,t} + N_{c,t-1}}, \tag{1.3}$$

where $\hat{S}^{BA}_{I,c,t} = \hat{N}^{BA}_{I,c,t} / \hat{N_{I,c,t}}$ is the expected share of immigrants with a BA, and $S^{BA}_{c,t-1}$

is the observed initial period BA share.[11] The first stage estimates perform well; the F-statistics, reported in the last row of table (3.1), are above 10 in both decades.

The use of immigration enclaves as instruments to predict local skill mix has a history in the literature studying the impact of immigration on local labor market outcomes (Card (2001) is a recent example) and skill mix on technology adoption (Doms and Lewis 2006, Lewis 2005). One argument for the validity of these instruments is that they represent a supply-push component of immigration flows; that is, enclaves predict the flows of immigration based on cultural or family reasons as opposed to local demand conditions (Lewis 2005). The validity of the enclave instrument in this paper rests on the assumption that the distribution of immigration across cities in 1970s is uncorrelated with time-varying demand shocks for skilled labor in later decades. This assumption would be violated if, for example, city-specific productivity shocks are highly persistent and the education composition of immigration flows from source countries is relatively constant, then past immigrant settlement may be correlated with a city's current skill specific demand shock (Card 2001). Also, recent evidence suggests that immigrant and native workers are imperfect substitutes within education groups. Thus, the effect on wages from changes in college share stemming from immigration may not have the same effects as changes in city college share generally. Although the differences are likely to be small given that estimates of the elasticity of substitution appear to be quite large (Peri and Ottaviano 2006, Ottaviano and Peri 2008, Card 2009, Borjas, Grogger, and Hanson 2008).

The second instrument set is based on city climate variables. The motivation for this instrument set is from Dahl (2002) who empirically tests a Roy (1951) model of self-selection of workers across U.S. states. He finds that while the migration patterns of college graduates are partially motivated by comparative advantage, amenity differences across states also play a large role in location decisions of college graduates relative to high school graduates. Building on this insight, I create an instrument set based on long-term climate patterns of U.S. cities. To do this, I collect data from a number of sources on average temperatures and precipitation for each city in my sample. Consistent with the idea that college graduates are more responsive to amenity factors, I find that indicators of mild climates are significant predictors of changes in city college graduate share. The specification that I use in practice uses average temperatures in July and January, and their squares. Additional details, including alternative specifications and robustness checks, are provided in the Appendix. The first stage estimates using the climate instruments indicate that the climate variables are good predictors of changes in

[11]See Appendix for details on the construction of this instrument. There are only 152 metropolitan areas that are identified in the 1970 Census. Since I use the 1970 distribution of immigrants as 'enclaves', analysis using this instrument is restricted to these 152 cities.

the city share of college graduates; the F-statistics are reported in the last row of Table (1.1).

The validity of the climate instruments depend on the assumption that the relationship between productivity and city specific climates is constant over time. In this case, these variables represent time-invariant city-specific factors that are differenced out of the estimating equation. This assumption may not be valid if technologies suited to city-specific climates are developed or adopted gradually over time.

The last four columns of Table (1.1) contain the results using instrumental variables. The middle column, which uses the enclave instrument, generates an estimate of β of 2.3 (s.e. = 0.52) for the 1980s and -0.65 (s.e. = 0.47) for the 1990s when focusing on all workers. This result is in agreement with the OLS result above and suggests that spill-overs from college education fell between the these two decades. The next two rows of the middle columns show the results for high school and college graduates separately. Again, the enclave IV results indicate that both groups benefited from college graduates in the 1980s but in the 1990s the effect was negative and insignificant. The last two columns display the results using the climate variables as instruments. The same pattern is apparent; there appears to be positive spill-overs from college education for all workers in the 1980s, but in the 1990s these results suggest the presence of college graduates negatively affects the wages of high school workers while there is weaker evidence that college graduates benefited. Overall, the instrumental variables results are very supportive of the OLS results above, and suggest that the relationship between wages and local college share has changed dramatically between the 1980s and the 1990s.

Selection: Methods and Results

Worker self-selection across cities also poses a potential problem for causal interpretation of the OLS estimates above. Changes in unobservable skills across cities that are correlated with changes in city college share could lead to biased OLS estimates of β. I address self-selected worker migration by using a semi-parametric correction procedure developed and implemented by Dahl (2002) in his paper on variation in college premia across U.S. states. The actual procedure that I use in this paper is an extension of Dahl's (2002) method that accounts for possible non-random location of foreign born workers (Beaudry, Green, and Sand 2007). The intuition behind this problem and its correction is as follows. Since we only observe the wage of a worker in the city that she chooses to live and work, city observations are a self-selected sample. Consider again the first stage equation for wages in year t:

$$\mathbb{E}\left[\ln w_{i,c,t}|x_{i,t,x}, M_{i,j,c}=1\right] = \delta_{c,t} + x'_{i,c,t}\mu_t + \mathbb{E}\left[u_{i,c,t}|x_{i,c,t}, M_{i,j,c}=1\right], \tag{1.4}$$

where $x_{i,c,t}$ is a vector of observable characteristics for individual i in city c and $\delta_{c,t}$ is a city-time intercept. $M_{i,j,c}$ is a dummy variable if individual i was born in state j and chose to live in city c. If the last expectation is non-zero, OLS potentially yields biased estimates of both $\delta_{c,t}$ and μ_t (Dahl 2002). Since I use the city-time intercepts as a dependent variable in my second stage regression, correlation between unobserved ability and city education can give rise to a systematic relationship between regression adjusted wages and city college share, even when there is no spill-over effect of the kind that I consider here.

Since the number of states of birth and cities is large, usual methods of correcting for selection bias that add a function of the probability of selecting each specific alternative into equation (1.4) are infeasible. Instead, Dahl (2002) suggests using an index sufficiency assumption that allows the error mean term to be written as unknown functions of a small number of selection probabilities, such as the probability of the first-best or observed choice. Dahl (2002) also suggests a method to estimate the relevant selection probabilities that groups observations into cells based on observable characteristics and uses cell means of different migration paths for the selection probabilities. The resulting selection correction procedure is carried out in two-steps, first estimating the selection probabilities and then adding a flexible function of them into equation (1.4). Identification of the selection correction function comes from variation in the probability of being observed in city c from similar individuals who originate from different states of birth. Intuitively, individuals with a lower probability of being observed in some city may have unobservable skills that are rewarded particularly well in that city, conditional on observable characteristics.

I implement the selection correction procedure in much the same way as Dahl (2002) except for modifications to account for the fact that I use cities instead of states and to incorporate foreign born workers. Additional details on the implementation and the estimation of selection probabilities are provided in the Appendix.[12] As a practical matter, I enter the estimated probabilities into equation (1.4) as quadratics. As before, I estimate equation (1.4) separately for each census year as a first stage and form city averages from the residuals. Differences in these regression adjusted average city wages are again used as a dependent variable in equation (1.2). The idea is that this procedure controls for differences in average city wages due to differences in unobserved ability as if it were another control variable and its effects are netted out in the first stage. It should be noted that in the first stage regressions, the selection correction functions are jointly highly significant, which is a necessary condition for the removal of selection bias (Dahl 2002).

[12]This procedure was also used in joint work with David Green and Paul Beaudry (Beaudry, Green, and Sand 2007). Additional details can be found there as well.

Table (1.2) replicates Table (1.1) using the selection correction procedure. The first two columns contain the OLS results. The selection correction has little impact on the OLS estimates; in the first two rows that use all workers or high school graduates the estimates are virtually unchanged from the uncorrected estimates, while the estimate for college graduates falls from 1.5 (s.e. = 0.15) to 1.34 (s.e. = 0.15) when using the selection corrected wages. This last result suggests that perhaps some of the estimated impact of city college share on the wages of college graduates might be attributable to self-selection. However, overall, accounting for possible non-random selection of workers across cities does not change the main conclusions from the results above. In the 1980s, the estimated β is positive for all workers, but in the 1990s it is negative for high school workers and positive for college graduates. It should be noted that these results are consistent with Moretti (2004) who also accounts for worker self-selection across cities by using longitudinal data from the National Longitudinal Survey of Youths (NLSY). Using panel data allows Moretti (2004) to control for both city and individual fixed effects; Moretti finds this has very little impact on the estimated spill-overs in comparison to models without individual fixed effects and concludes that omitted individual characteristics are not a significant source of bias. The last four columns of Table (1.2) display the results using both the selection correction and the instrumental variables. Again, the results are very similar to the estimates in Table (1.1) without the selection correction.

Summary

My overall assessment is that a simple model of technological externalities, as it is typically estimated, is of little use in understanding the patterns in the data that are observed above. To summarize, these patterns are: (1) a strong positive correlation of average wages and aggregate education in the 1980s but not in the 1990s, (2) a negative correlation between aggregate education and wages of unskilled workers in the 1990s, and (3) a positive correlation between the fraction of skilled workers and the wages of skilled workers in the 1990s. The first observation is surprising in both magnitude and direction. Dramatic parameter instability is not a standard feature in models of technological externalities and, as argued above, *a priori* reasoning would suggest that education spill-overs would become even more important in the 1990s, not less as observed in the data. The second two observations are not consistent with most competitive models. Then what are we left with? A second class of theories models human capital spill-overs as pecuniary externalities that can arise through frictions in the labor market. Given that several recent papers argue that labor market imperfections may not be stable over time (Lemieux 2007, Autor 2001, Levy and Temin 2007), this class of models may offer a more promising approach to understanding these otherwise puzzling findings.

1.3 An Alternative View - Acemoglu 1999

The goal of this section is to outline a model that is helpful in organizing or understanding the patterns in the data that are documented above. As argued previously, these patterns are not consistent with standard competitive models and so I focus on non-competitive models. A candidate model must also be able to explain the time pattern of spill-overs from education that are documented above. One particular model that includes equilibria with and without spill-overs to education is Acemoglu (1999). Acemoglu's (1999) model is one of endogenous job composition, in which firms first create jobs and then are matched with workers. In this model, an increase in the fraction of skilled workers and/or skill-biased technical change can generate the change in social returns that are observed in the data.

This model incorporates labor market frictions and *ex ante* investments that, while highly stylized, offers an interpretation of the above findings. Moreover, the model has a number of additional predictions on a range of labor market outcomes. In this section, I briefly describe a simple version of the model that is presented in Acemoglu (1999) and discuss its implications. Although this model is mainly for illustrative purposes, and should be interpreted as such, I nevertheless take several of its predictions quite seriously in the empirical section that follows. The main benefit n doing so is that it illustrates how the simple observation that motivated this paper could have broad implications for the labor market in general.

1.3.1 Acemoglu's Model of Endogenous Job Composition

Consider an economy with an equal mass of workers and profit-maximizing firms that lasts for one period.[13] A fraction θ of workers are high skilled and $1-\theta$ are low skilled. Workers can also be of low or high education, with high skilled workers more likely to be highly educated (although this correlation is not perfect). Note that in this model there is a distinction between education, which is observable, and skill, which is not observable.[14] Let the human capital or skill of each type of worker be h_j for $j \in \{Low, High\}$.

In this economy, each worker is matched with a firm and production takes place in one-worker–one-firm partnerships according to the production function:

$$y_{j,i}(h_j, k_i) = k_i^{1-\alpha} h_j^{\alpha}. \tag{1.5}$$

[13]This section borrows heavily from Acemoglu (1999). For derivations and proofs or for an extended version of this model, the interested reader is directed to that paper.

[14]The distinction between education and skill in the model is an artificial one that does not play an essential role in the model that follows. This distinction is essentially made in order to apply the model's insight and predictions to the observable component of skill we observe in the data (i.e. education).

As in Acemoglu (1996, 1999), firms chose k_i irreversibly before the skill of the worker is known. The investment decision is assumed costless. After this decision is made, firms and workers are randomly matched. A firm finds out the worker's skill and then decides whether or not to produce. If the firm does produce, it pays ck to install the capital and workers and firms bargain over wages, where the worker's bargaining power is assumed to be β.

The two main features of this model are that search is costly and the irreversible investment decision takes place before the firm finds out who they are matched with. When matching is random, firms with different capacity types are equally as likely to meet low or high skilled workers. Also, since the cost of the capital is sunk at the wage bargaining stage, workers and firms split output with workers getting $\rho \cdot y_{i,j}(k_i, h_j)$. These features imply that the type of job that a firm would wish to create depends on the amount of skilled workers in the economy, θ, and their relative productivities. To see this, let

$$V(k, x^H, x^L) = \theta x^H[(1-\rho)k^{1-\alpha}h_H^\alpha - ck] + (1-\theta)x^L[(1-\rho)k^{1-\theta}h_L^\alpha - ck] \quad (1.6)$$

be the expected value of a firm choosing k. Here, $x^j \in \{High, Low\}$ represents the probability in equilibrium that a firm matched with a worker of either high or low type decides to produce. The first part of this equation represents the value to the firm of producing with a high type worker ($(1-\rho)y(k, h_H) - ck$) and the second the value of producing with a low type. Since matching is random, these values are weighted by the proportion of each type of worker.

Acemoglu (1999) shows that in this economy, there are two possible equilibria - pooling and separating. Consider the case when firms pool across both types of workers, so that $x^H = x^L = 1$. In this case, firms maximize equation (1.6) by choosing

$$k^p = \left[[\theta h_H^\alpha + (1-\theta)h_L^\alpha] \left(\frac{(1-\rho)(1-\alpha)}{c} \right) \right]^{\frac{1}{\alpha}}. \quad (1.7)$$

Here, all firms choose the same amount of capital and hire both types of workers. The wages of each type of worker are

$$w^j = \rho\lambda \left[\theta h_H^\alpha + (1-\theta)h_L^\alpha\right]^{\frac{1-\alpha}{\alpha}} h_j^\alpha \quad \text{for} \quad j \in \{L, H\}, \quad (1.8)$$

where $\lambda = \left[\frac{(1-\rho)(1-\alpha)}{c}\right]^{\frac{1-\alpha}{\alpha}}$. Thus, in the pooling equilibrium, the wages of all workers are increasing in the fraction of skilled workers in the economy, θ. This happens because firms are willing to invest more when the probability of meeting a skilled worker increases (Acemoglu 1996, Acemoglu 1999). Also, in this equilibrium, the relative wages of high and low skilled workers are not related to θ. Thus, the fraction of skilled workers in an economy has no effect on inequality.

However, if either θ or $\frac{h_H}{h_L}$ are large enough, the economy will switch from a pooling to a separating equilibrium.[15] In this case $x^H = 1$ and $x^L = 0$, and firms maximize (1.6) by choosing $k^H = \left((1-\alpha)(1-\rho)c^{-1}\right)^{1/\alpha} h_H > k^p$. The wages of high skilled workers increase while the wages and employment of low skilled workers collapse. That is, in a separating equilibrium firms create high capital jobs and only hire skilled workers. This is in contrast to the pooling equilibrium where there were 'middling' jobs open to both skill groups (Acemoglu, 1999). It should be kept in mind that in Acemoglu's (1999) extended dynamic version of the model, the separating equilibrium has low capital jobs targeted to low skilled workers and positive employment and wages for this group as well.

To summarize the main predictions of this model with respect to the fraction of skilled workers in an economy, consider an increase in θ while holding skill-biased technical change (h_H/h_L) constant.[16] Starting from the pooling equilibrium, increasing the fraction of skilled workers will (1) increase the wages of all workers and (2) have no effect on inequality. If the fraction of skilled workers becomes large enough, the economy will switch from a pooling to a separating equilibrium and firms will create high capital jobs targeted at high skilled workers and low capital jobs targeted at low skilled workers. An increase in θ that triggers an equilibrium switch will (1) increase the wages of skilled workers, (2) decrease the wages of unskilled workers, (3) increase unemployment and (4) change the composition of employment from 'middling' jobs once opened to both high- and low-skill workers to jobs in extreme parts of the distribution. The switch to a separating equilibrium will also increase returns to education because highly skilled workers are also more likely to be highly educated. But since this correlation is not perfect, residual inequality will also increase.

To relate the above discussion to labor market outcomes in U.S. cities during the 1980s and 1990s, first assume that in the 1980s the fraction of skilled workers was below the threshold so that cities were in a pooling equilibrium. Additionally, assume that in the 1990s the fraction of skilled workers and/or skill-biased technical change were large enough to trigger a switch from a pooling to a separating equilibrium. To get an idea of the plausibility of these assumptions, Figure (1.2) a plots the distribution of city college share for each year. This figure illustrates that the distribution of skilled labour has shifted rightward over the 80s and by a slightly larger amount over the 90s. Since

[15] The condition is $\frac{h_H}{h_L} > \left(\frac{1-\theta}{\theta^\alpha - \theta}\right)^{1/\alpha}$, with the normalization $c = 1-\beta$. See Acemoglu (1999) for details.

[16] In the current model, increasing the relative productivity of high-skill workers is equivalent to skill-biased technical change. This is because firms only hire one worker and increases in $\frac{h_H}{h_L}$ increase the productivity of capital when combined with skilled workers relative to unskilled workers (Acemoglu 1999).

the cut-off between the pooling and separating equilibria in the model is determined by the fraction of skilled labour in the economy and the relative productivity of skilled-to-unskilled labour, this figure illustrates that one of the model's driving forces is present in the data. However, since there is a substantial overlap between the 2000 and 1990 distributions of city college share, we would require that this rightward shift in distribution of college share be accompanied by skill-biased technical change or a relative productivity increase of skilled labour. Many economists have interpreted the large increases in the relative wages of skilled to unskilled workers occurring simultaneously with growth in their relative supply during the 1980s as evidence in favor of skill-biased technical change (Acemoglu 2002). These co-movements can be represented as North-East movement in relative-productivity/fraction-of-skilled-worker space, where the cut-off can be represented as a downward sloping curve as in Acemoglu (1999). The assumption made in this paper is that their intersection occurred in the 1990s.

Under the above assumptions, the above model has the following implications. In the 1980s, the fraction of skilled workers in a city should be positively correlated with the wages of all workers and impact them in the same magnitude. Also, during this decade, the fraction of skilled workers should be uncorrelated to returns to education and residual inequality. However, in the 1990s, the fraction of skilled workers in a city should have different impacts on high- and low-skill workers. In particular, the model predicts that the fraction of skilled workers in a city will have positive impacts on high-skill workers, but negative impacts on low-skill workers. Thus, more skilled cities will have higher returns to education and greater residual inequality. The dynamic version of the model suggests that we should also expect unemployment to be positively correlated with city educational attainment in the 1990s in a separating equilibrium. Below, I take these predictions to the data and evaluate each one in turn.

1.4 Some Empirical Evidence

1.4.1 Empirical Method

The goal of this section is to examine empirically several predictions of the model presented above. I examine these implications in much the same way as I examined wages in section (2). In particular, I carry out the analysis at the city level in differenced equations such as

$$\Delta Y_c = \beta_0 + \Delta Z_c \cdot \alpha + \Delta BA_c \cdot \beta + \Delta\epsilon_c, \tag{1.9}$$

where Y_c is some outcome for city c and Z_c is a vector of city level controls. The main variable of interest is BA_c, which is the fraction of workers in a city with a college degree.[17]

1.4.2 Results

Before turning to additional results, it is useful to take another look at the patterns presented above. Table (1.1) shows that all education groups benefit from increases in city level college share during the 1980s. In contrast, during the 1990s, increases in city college share benefit higher educated workers and harm those with high school education. The model presented in the previous section assumes that high educated workers are more likely to be high-skill and low educated workers are more likely to be low-skill. Under this assumption, the patterns of the estimated spill-overs with respect to education level is consistent with the idea that there has been a change in the labor market that harmed low-skill workers, such as a move from a pooling to a separating equilibrium in the model described above.

Residual Inequality

The above results suggest that high- and low-skill workers are impacted differently by aggregate education in the 1990s. Another way to examine this is to look at different percentiles of the residual distribution, interpreting higher percentiles as wages of workers with higher (unobservable) skills. Consider again the first stage regression of the log of individual weekly wage on a set of individual characteristics, estimated separately for each year. I calculate the residuals of these regressions for each city at different deciles of the residual wage distribution. I then use differenced city residual wage deciles as dependent variables in various regressions such as equation (1.9).

Panel A of Table (1.3) reports the results of this exercise when estimated by OLS. Each row corresponds to a different decade and each column to a different decile of the residual wage distribution. Each entry in the table reports the coefficient on the fraction of college educated workers in each city (β) from a separate regression where the dependent variables are differenced wage deciles. All regressions control for the city level characteristics mentioned above. The first row shows the results for the 1980s, which indicate no strong pattern over the different deciles. For example, while there is a slight increase in the coefficient on college share from the 10th percentile to the 90th, there is a decrease from the 20th to the 80th percentiles. In fact, the magnitudes of the estimated effects are very similar over the entire residual wage distribution.

This is in contrast to the 1990s (the 2nd row), where the results are negative at the lower deciles and monotonically increase with each decile. For example, at the 10th

[17] All regressions are weighted by $1/\sqrt{1/n_t + 1/n_{t-1}}$.

percentile the estimated coefficient is -0.71 and statistically significant. At the 90th percentile, the result is positive, statistically significant, and estimated to be about 0.40.

Panel B of Table (1.3) shows the results for the restricted city sample using the 1970 enclave instrument. The patterns here are the same as the OLS results above. In the 1980s there does not seem to be any significant relationship between the share of college graduates in a city and the magnitude of the estimated impacts on different skill groups. During this decade increases in the fraction of skilled workers increase the wages of all workers, regardless of skill type. However, in the 1990s it is a different story - the fraction of skilled workers in a city negatively impacts low-skill workers and positively impacts workers with higher skills. In terms of the model presented above, these results are consistent with the 1980s being characterized by a pooling equilibrium and the 1990s by a separating equilibrium. However, these patterns are not easily reconciled in standard models of spillovers such as in Moretti (2003, 2004). Panel C of Table (1.3) contains the results using the city climate instruments. These results are very similar to the enclave instrument above.

Table (1.4) summarizes these results in terms of residual wage dispersion. The format of this table is the same as above, except that the dependent variables are various measures of residual wage inequality. For example, in column (1) the dependent variable is the differenced 90-10 residual log wage differential, while columns (2) and (3) show the 50-10 and 90-50 log wage differential, respectively. Again, each row corresponds to a different decade and each cell entry is the coefficient on the fraction of college graduates in a city. Looking across the first row, we can see that the fraction of college graduates in a city is not significantly associated with the measures of residual wage dispersion considered here. However, in the 1990s, the fraction of college graduates is positively and significantly associated with each measure. The last six columns repeat this exercise using instrumental variables. Again, the results suggest that residual inequality is not correlated with college share in the 1980s but is strongly, positively correlated to college share in the 1990s. This result is consistent with the idea that during the 1990s, cities switched from a pooling to a separating equilibrium.[18]

[18]Bernard and Jensen (1998) also examine patterns of residual wage inequality a the local level. In their paper, they run regressions similar to those shown above at the state level. They find, as do I, that residual inequality varies considerably between local labor markets at any given point in time. Among the most robust of their results is the strong negative impact of durable manufacturing employment on residual wage inequality. I find support for this result in specifications not reported here. However, the patterns reported above are not sensitive to controlling for the share of workers in durable manufacturing. Thus, the mechanism behind residual wage inequality emphasized here is robust to other hypothesized driving forces thought to explain inequality at the regional level.

Returns to Education

Since highly skilled workers are more likely to be highly educated, when an economy switches from a pooling equilibrium to a separating equilibrium, the return to education should also increase. However, in the pooling equilibrium, the fraction of skilled workers in a city should have no effect on college wage premia. The goal of this section is to examine to what extent college wage premia are correlated to city level educational attainment in each decade.

The first step in this exercise is to calculate local college premia. I calculate the high school graduate - exactly college log wage differential by running regressions separately by year and city of log wages on dummies for gender, race, and U.S. citizenship status, a quartic in potential experience and 4 indicators for completed education[19]. I take the high school - exactly college graduate wage premia to be the coefficient on the college graduate dummy from the above regressions. Consistent with previous literature, I find large differences in the return to college across local labor markets (Bernard and Jensen 1998, Dahl 2002). I also apply the selection correction procedure described in section (2) to the city level regressions when estimating the college premia to obtain estimates of the college premia that are not biased due to the self-selected migration of workers as in Dahl (2002).

Table (1.5) contains the results that use the change in college premia as the dependent variable in regressions of equation (1.9). Table (1.5) has the same layout as the previous tables: each row represents a different decade and each entry is the coefficient β on the share of college educated workers in a city from separate regressions that also control for the city level characteristics. In the first three columns the dependent variable refers to estimates of the college premia without using the selection correction procedure and the last three columns to the selection corrected estimates. The first column shows the OLS results which indicate that in the 1980s, an increase in the fraction of college graduates in a city is negatively (but not significantly) correlated with the returns to a college degree. This result is reversed in the 1990s. During this decade an increase in the fraction of college graduates is associated with an increase in college premia and this relationship is statistically significant at any conventional significance level. Using the enclave instrument preserves this pattern, but the coefficient on the change in college share becomes large, negative, and significant in the 1980s, while it is smaller and not well defined in the 1990s. However, in the 1990s, the estimated coefficient still has a positive point estimate. The third column displays the results using the climate variables

[19]These indicators are for less than high school, some post-secondary, exactly college graduate, and advanced degree. High school graduate is the omitted category. 'Exactly college' is the case where an individual has a college degree, but no post college education in the 1990-2000 Censuses and exactly 16 years of education in earlier years.

as instruments and indicates that increases in city college share lowered the returns to college education in cities during the 1980s but increased the returns to college in cities during the 1990s. The last three columns use the selection corrected returns to college as a dependent variable, which does not appreciably change the results. Overall, this section presents strong evidence that the relationship between the supply of college educated labor and the returns to college in local labor markets changed substantially between the 1980s and the 1990s.

The outcome of this section is not surprising given earlier results, and again is in opposition to predictions of standard production functions. Taken as a whole, these results are consistent with the theory presented above suggesting that there was a switch in regime from the 1980s to the 1990s.

Unemployment

Another implication of the extended dynamic model offered in Acemoglu (1999) is that when an economy switches from a pooling to a separating equilibrium the unemployment rates of all workers should increase. The intuition is that in a pooling equilibrium, firms design one type of job and hire all workers they meet. In contrast, in a separating equilibrium, firms that create high-skill jobs only hire high-skill workers and high-skill workers prefer to wait for high-skill jobs. The result is that unemployment for all workers increases in a separating equilibrium. In this section, I provide brief evidence that this occurred in the 1990s in my sample of U.S. cities. To carry out this analysis, I calculate the unemployment rate for different groups of workers in each city and year. I then use these as dependent variables in equation (1.9). Table (1.6) contains the results.

Panel A displays the OLS results while Panels B and C show the IV results. Each column shows the estimated coefficient on college share for a different group of workers. The first entry of Panel A in column (1) row (1), for example, shows that the fraction of college graduates in a city is uncorrelated with the unemployment rate in the 1980s. Row (2) shows the result for the 1990s, where city college share is highly and positively associated with unemployment. Columns (2)-(4) show that in the 1980s, unemployment rates for each group are not significantly correlated to college share. However, in the 1990s, and particularly for lower education groups, the unemployment rate and college share show positive association. The IV results in the second two panels indicate that in the 1980s, if anything, increases in BA share lower unemployment - particularly among lesser educated workers. The 1990s, however, indicate the opposite - increases in college share increase unemployment (particularly among the lower education groups). Again, this pattern is predicted by a switch from a pooling to a separating equilibrium.

Changes in Composition of Employment

The model presented above suggests that once an economy switches from a pooling equilibrium to a separating one, 'middling' jobs once open to all workers are replaced with jobs targeted at either high- or low-skill workers. Acemoglu (1999) presents evidence that employment has shifted from the middle to more extreme parts of the job quality distribution at the national level for the U.S. from the mid-1980s to the early 1990s. More recently, several authors argue that during the 1990s, the composition of employment has 'polarized'. That is, employment has increased in occupations at the very low and high ends of the occupation skill distribution (Autor and Dorn 2007, Autor, Katz, and Kearney 2006). However, these author's theoretical justification for these empirical trends is differnet from Acemoglu (1999). In this section, I examine whether there is a link between the share of college graduates in a city and the composition of jobs. In doing so, I take care to differentiate between the alternative theories attempting to explain the shifting employment structure.

In order to study changes in job composition across U.S. cities, it is first necessary to define a 'job' and rank them in terms of quality. As suggested by Acemoglu (1999), one way to rank jobs is according to average wages or residuals in some base year. Then, in terms of the model, we can think of higher ranked jobs as higher k jobs - jobs that pay more given a set of worker characteristics. In this framework, for example, jobs that pay close to the median are 'middling' jobs. After ranking jobs it is possible to examine whether employment has shifted away from jobs ranked near the middle toward jobs with more extreme rankings. Of particular interest in this section is whether or not these shifts are correlated to city characteristics - specifically, college share.

I choose 1980 as the base year to rank jobs and categorize them according to detailed occupation codes.[20] I follow Acemoglu (1999) by ranking jobs according to their residual wages. Specifically, I regress individual log wages on education (4 categories), a quartic in potential experience, and indicators for gender, race, and U.S. citizenship at the national level using 1980 census data. I then sort occupations into percentiles according to the occupation's average residual from the above regression. Jobs that are ranked below the 25th percentile are called 'low skill' jobs and jobs ranked above the 75th percentile are called 'high skill' jobs. I calculate employment shares as the share of annual hours supplied in each job for each city. Annual hours supplied are calculated by multiplying weeks worked in the previous year by the usual amount of hours worked.[21]

Table (1.8) displays the results of this exercise at the national level for different education groups. Each cell shows the percentage of employment for low, middle and high

[20]The occupation grouping I use is based on the IPUMS variable occ1990. This categorization is similar to Autor, Katz, Kerney (2006) and in the same spirit as Goos and Manning (2007)

[21]See data appendix for more details.

skilled jobs for each year. The last column, labeled Weight-at-Tails, shows the percentage of employment in low and high ranked jobs. Looking first at all workers, the weight-at-tails increases from 44.4% in 1980 to 47.9% in 2000 - most of this increase occuring in the 1990s. This result is consistent with Acemoglu (1999) who shows that from 1983 to 1993 the composition of employment increased in the top and bottom 25 percent job categories - half of this increase occuring in the 1990s. Looking separately at workers with high school or less education shows that, for this group, employment has moved away from jobs in the middle of the quality distrubtion (where quality is still defined over all workers) and mostly into low quality jobs. For the college degree or greater group, employment has also moved away from middling jobs and more toward jobs above the 75th percential.

As mentioned above, the goal of this section is to link a city's share of skilled workers to the composition of jobs in that city. One problem in examining this connection is that there is a possible mechanical relationship between the fraction of highly skilled workers in a city and low-skill employment as a share of total employment. To deal with this issue, I follow the literature by focusing on the employment of non-college workers (those with high school or less) where any mechanical relationship is much weaker (Autor and Dorn, 2007; Mazzolari and Ragusa 2007). Since most of the change in the composition of employment is due to increases at the lower tail for this group, I further restrict attention to changes in employment of high school or less workers in jobs ranked below the 25th job percentile. I keep the same job rankings as above and use residuals from a regression run for all workers at the national level in 1980. I then calculate the change in employment in the lower tail of the job distribution for non-college workers in each city. In addition, I also calculate the change in mean wages of the lower tail jobs for non-college workers (but I do not adjust them for other observable characteristics). I use these as dependent variables in equation (1.9).

Table (1.7) shows the number and type of occupations that are ranked below the 25th percentile using the above ranking procedure. The bulk of these low-skill jobs are service occupations, but several other occupations are also well represented. It should also be noted that in 1990, mid-way through my sample period, the average percentage of employment of non-college workers in low skilled jobs was 28.6 percent and ranged from 19% to below 50% among the cities in my sample. Thus, there is ample room for this fraction to rise or fall during the 1990s.

Table (1.9) contains results from the regressions that use the change in low-skill employment as a dependent variable. Each row shows the estimates with a different skill ranking procedure. Row (1) contains the baseline specification, described above, that ranks jobs according to the residuals from a regression at the national level using all workers in 1980. Each cell contains the coefficient on the change in city college share estimated from equation (1.9). The OLS results in the first two columns of row (1) indicate that in both periods, changes in college share are positively associated with changes in the composition of employment of non-college workers toward jobs at the lower end of the skill distribution. The coefficient in the 1990s is twice as large as in the 1980s, suggesting a stronger relationship. The next two columns show the results using the enclave instrument. For the 1980s, the point estimate is lower and no longer statistically significant. In the 1990s, the point estimate is larger and still significant. Using the city climate variables as instruments, as in the next two columns, reinforces this finding. These results indicate that increases in city college share caused a shift in employment among non-college workers toward lower quality jobs during the 1990s.

Table (1.10) contains the results using the change in the average wage of non-college workers in jobs ranked below the 25th percentile as the dependent variable and has the same layout as the previous table. The first row again shows the baseline specification. In the 1980s, increases in college share are associated with increases in wages of jobs in the lower part of the quality distribution - consistent with the results described elsewhere in this paper. However, in the 1990s, this relationship is reversed. Increases in college share are now negatively correlated to the wages of non-college workers in low-skill jobs. This striking pattern of the reversal of fortunes for low-skill workers in jobs with increases in employment would be hard to reconcile with standard competitive models. In Acemoglu (1999) this result comes about because firms create low quality jobs specifically designed for low-skill workers in a separating equilibrium; this is a prediction that appears to be supported by the data. The next four columns use the enclave and climate instruments and only make this pattern more stark. These data suggest that in the 1990s, the fraction of skilled workers in local economies changed the composition of employment in a way that harmed low-skill workers.

This analysis is related to a recent and growing literature on the changes in the composition employment. Two papers that examine this issue at the national level are Autor, Katz, and Kerney (2006) and Goos and Manning (2007). Autor, Katz, and Kerney (2006) show that in the 1990s employment was "polarizing into high-wage and low-wage work at the expense of middle-wage jobs." They argue that a more nuanced version of the skill-biased technical change hypothesis can rationalize these patterns as arising from changes in the demand for skills. Specifically, they argue for a version of the skill-biased technical change theory put forth by Autor, Levy, and Murnane (2003) which suggests that computers change the demand for skills non-monotonically. This theory predicts that computers replace 'routine' tasks while complementing 'non-routine' tasks. Since non-routine tasks are performed primarily by workers in either high- or low-paying occupations and routine tasks are intensive in middle-pay occupations, employment changes in these occupations can be thought of as being induced by technology that changes the demand for skills. The authors cite as evidence of their demand side interpretation the coinciding patterns of polarizing employment, falling inequality at the bottom and increasing inequality at the top of the wage distribution, in the U.S. over the 1990s.

Lemieux (2007) assesses the argument of Autor, Katz, and Kerney (2006) and suggests that there are two potential puzzles for this story. First, Lemieux (2007) shows that wages in occupations directly related to computers, such as computer scientists and other technology workers, faired relatively poorly compared to other high-skill occupations - even when several of these occupations had increases in their relative employment. Secondly, Lemieux (2007) argues that it is not clear why this form of skill-biased technical change took until the 1990s to have polarizing impacts on non-routine tasks. It should be emphasized that the driving forces for structural change in this paper are the state of technology and the supply of highly skilled workers, and so the timing of events are more easily explained - an issue that I explore in the following sub-section.

Goos and Manning (2007) also explore changes in employment composition using data that is mainly from the United Kingdom. They show that employment patterns in the U.K. exhibit job polarization, and argue that these trends are consistent with the 'routinization' hypothesis of Autor, Katz, and Murnane (2003). They find that employment has been increasingly concentrated in traditionally high- and low-paying occupations. Interestingly, though, they find that the wage patterns of low-skill workers do not seem consistent with the view that technology is changing the demand for skills. Specifically, they find that the rise in the number of 'lousy' or low-paying jobs has coincided with a decline in their pay, which is consistent with the evidence presented here. In fact, the authors also mention the model of Acemoglu (1999) as a possible explanation for this otherwise puzzling result. However, they dismiss that model and conclude that changes in demand induced by technology is a more "plausible" explanation for the employment trends they document. The results presented in this paper suggest that non-competitive models, such as Acemoglu (1999), might be more relevant than previously thought in explaining changes in employment composition.

Finally, there are several papers that examine changes in employment at the local level. Manning (2004) argues that the employment prospects of low-skill workers are increasingly dependent upon physical proximity to high-skill workers because they produce goods that high-skill/high-income workers consume. Using U.S. city level data, he shows that low-skill workers are increasingly employed in the non-tradables sector and that this relationship is strongly linked to the fraction of college educated workers in a city. Mazzolari and Ragusa (2007) also study changes in employment for low-skill workers; using U.S. census data from 1980-2005, they find that cities with a larger share of skilled workers have a higher share of employment of low-skill workers in outsourced home production services. They argue that this relationship is due to the demand of high-skill workers for these services. They support this by showing that city wage growth in the bottom of the distribution (with respect to the median) is correlated with the share of low-wage workers in the outsourced production sector. Autor and Dorn (2007) show that low-skill employment in service occupations grew most in U.S. commuting zones with larger increases in upper income inequality from 1980-2005. They argue that this is due to the collocation of high income earners who demand services. While the findings of these papers with respect to low-skill employment patterns are consistent with the findings of this paper, more research is needed to differentiate these demand side interpretations from non-competitive explanations such as Acemoglu (1999).

The next three rows of Table (1.9) and (1.10) are robustness checks and consider different methods of ranking and categorizing jobs. In the second row, jobs are ranked by residuals after netting out industry effects. In terms of the model, this can be thought of as firms choosing what type of job to open (laborer or manager) within an industry. For example, firms may choose their line of business, and then choose their production technique given the supply of skill and state of technology. The third row ranks jobs using the wages of high school workers only. This can be thought of as firms choosing to pool or separate within the market for workers with low education. The forth row considers a different categorization of jobs by using broad industry-occupation cells rather than the very detailed occupation codes used above. This is similar to the procedure used by Acemoglu (1999). The conclusions reached above are robust to these different specifications, and lend broad support to the idea that in the 1990s, cities moved to a separating equilibrium changing the composition of jobs away from the middle of the pay distribution.[22]

1.4.3 Education Threshold

Up to this point, the analysis has proceeded under the assumption that in the 1990s relative to the 1980s we have witnessed cities in a separating equilibrium rather than a pooling equilibrium, because education and skill-biased technology have reached some threshold. However, it is possible that a number of cities could have reached the education threshold in the 1980s. In this section, I ask whether a group of cities in the 1980s could have been passed this threshold, and, if so, behaved like cities in the 1990s.

To carry out this analysis, I assume that the coefficient β is a function of the initial fraction of skilled workers in the economy. In particular, I assume that $\beta(BA_{t-1})$ is a smooth function of initial BA share except at some discontinuity point where the economy switches from the pooling regime to the separating one. The idea is that if there

[22]In addition, in results not reported here, I performed the same analysis focusing on non-college employment in service occupations (instead of ranking jobs by residuals). The results using this procedure result in same patterns as those in Table (1.9) and (1.10) for both employment and wages.

were cities that began the decade with very high levels of skilled workers then increases in their share could push these cities past some critical threshold and into a separating equilibrium. The goal of this section is to estimate this cut-off or threshold value, if any.

We can think of the coefficient β as being indexed by the initial share of college workers. Then for each city c, there is a β_c that corresponds to that city's initial share of skilled workers. If we observed both average wages and β_c for each city, then we can apply methods used to estimate unknown discontinuity points such as those in Card, Rothstein and Mas (2007) or Loader (1996). However, in this case we do not observe the β_cs; they must be estimated from the data.

The method I use to estimate an education threshold is an approach based partly on Loader (1996). Consider the regression

$$\Delta w_c = \beta_0 + \Delta C_c \delta + \beta_c(BA_{t-1}) \cdot \Delta BA_c + \Delta \epsilon_c, \tag{1.10}$$

where the estimated impact of a change in college share depends on initial college share. Assume that $\beta(BA_{t-1})$ is left and right continuous except as some change point BA^*_{t-1}. To find this value, I estimate $\hat{\beta}_c$s with local one-sided weighting and choose as the threshold level the value that maximizes $\Delta(BA_{t-1}) = \text{abs}(\hat{\beta}^L_{BA_{t-1}} - \hat{\beta}^R_{BA_{t-1}})$.

Before proceeding with this method, I take the data on changes in average regression adjusted city wages and college graduate share and regress these variables on all of the city control variables and retain the residuals: $\Delta\tilde{w}_c$ and $\Delta B\tilde{A}_c$. This first step forces the coefficients on the city controls to be the same for all cities, regardless of initial skill level. I do this so that the local weighting procedure can be performed in the bi-variate case.

The next step is to estimate β_cs locally, by smoothing on initial BA share. I construct a weighting function $W\left(\frac{BA_{c,t-1} - BA_{t-1}}{h}\right)$ for each observed BA_{t-1}, where $W(\cdot)$ is a kernel weighting function and h is a bandwidth. Then for each city, I estimate $\Delta\tilde{w}_c = \gamma + \beta^L_c \Delta B\tilde{A}_c + \epsilon_c$ using these weights and data with $BA_{c,t-1} < BA_{t-1}$. β^R_c is defined similarly for data $BA_{c,t-1} > BA_{t-1}$. The result is a pair $\{\hat{\beta}^L_c, \hat{\beta}^R_c\}$ for each initial college share.

Figure (3.2) contains the results for this exercise using wages of all workers in a city. The horizontal axis represents the initial fraction of skilled workers, while the vertical axis represents the estimated β_cs. For each candidate threshold value, there is a green (circle) and orange (diamond) dot for the left sided estimate, β^L_c, and right sided estimate, β^R_c, respectively. The cut off value is estimated to be the value of BA_{t-1} that gives the largest difference between these estimates; in this case it is estimated to be 0.253.[23] At this value, there is a large discontinuity in the estimated impact of changes in the fraction of skilled workers. I interpret cities with initial shares of college

[23]See notes on the figure for additional details.

workers less than this cut-off to be in a pooling equilibrium during the 1980s. Cities with initial shares larger than this cut-off are predicted to have switched to a separating equilibrium. Below, I split the sample and test these predictions indirectly.

It should be noted that this estimated threshold value is not sensitive to alternative kernel functions, including rectangular, and is robust to reasonable alternative bandwidths. In addition, using alternative outcomes (instead of wages of all workers) produces similar results. When this procedure is applied to data in the 1990s, there is no obvious threshold value.[24] Change point methods, such as in Hansen (1998), produce similar results, but are more sensitive to specification. In contrast, the method outlined above is quite robust to alternative specifications to produce $\Delta\tilde{w}_c$ and $\Delta \tilde{BA_c}$.[25]

After estimating what appears to be a threshold value for the 1980s, I cut the sample of cities into two groups depending on whether they were above or below the estimated cut-off. 17 of the 286 cities are estimated to be above the threshold. Table (1.11) contains the results of several of the outcomes examined above regressed on changes in college share for each of the two groups. Panel A of this table refers to the 17 cities estimated to be past the threshold and Panel B to cities that were below the threshold. Each column is a different outcome measure. The hypothesis is that cities past the cutoff behave like cities in the 1990s. Out of the 7 different city level outcomes that I estimate, 6 have point estimates that are of correct sign, although few are significant. For example, changes in college share are positively associated with non-college employment in the bottom 25 percentile jobs but negatively associated with wages for this group. However, point estimates were sensitive to the inclusion of one or two cities so that caution is needed when interpreting these results. This is not surprising given the small number of cities estimated to be above the threshold. In addition, theoretically the threshold level depends both on the fraction of skilled workers and the level of technology - and there is increasing evidence that the adoption of technology can vary endogenously by city (Beaudry, Doms, and Lewis 2006). Nevertheless, the estimated discontinuity is striking and relatively robust, and does provide evidence that the impact of changes in college share on city level outcomes depends at least in on part the initial skill level. This

[24]The fact that a threshold value could not be identified in the 1990s is not surprising and is consistent with the maintained assumption of this paper that all cities had switched to a separating equilibrium in this decade. The common perception is that technical change continued and may have even accelerated during the 1990s and, thus, the education threshold would be lower making this a likely outcome. It should be noted, however, that technical change by itself is an unlikely driving force for the outcomes in the paper. An idea that is supported by the fact that a threshold is identified in the 1980s when technical advances were presumably lower.

[25]In addition, in results not reported in this paper, I have performed this exercise using the initial fraction of females, blacks, hispanics, and U.S. citizens in the workforce as well as the initial share of durable manufacturing employment as alternatives to the initial fraction of skilled workers. None of these alternative initial conditions generated discontinuities that displayed a pattern similar to the initial fraction of college graduates.

result is important because it suggests that even when holding technological possibilities constant, by looking within a decade, an increase in the proportion of skilled workers can have non-monotonic impacts on the city level outcomes examined in this paper.

1.5 Conclusion

The empirical literature on human capital externalities has often come to conflicting conclusions regarding the magnitude and existence of education spill-overs. One emerging view reconciles previous mixed results by suggesting that positive externalities only accrue to post-secondary education. However, in this paper, I show that this view of human capital externalities is far from complete. In particular, even when focusing on a city's share of college graduates as the measure of local human capital, spill-overs are remarkably unstable between the 1980s and the 1990s.

In the first part of this paper I document the over-time pattern of education spill-overs using U.S. Census data from 1980 to 2000, focusing on a set of 286 consistently defined metropolitan areas. Using an empirical specification that is commonly employed in the literature on social returns to education, I find that the external return to college education is large, positive, and benefits all workers in the 1980s. In the 1990s, however, this relationship changes dramatically; the overall impact of changes in college share on average city wages is essentially zero. This result hides heterogeneous impacts by education level. In particular, I find that more high-skill workers in a locality negatively effect the wages of low-skill workers in the 1990s. I address the potential endogeneity of city college share by using an instrumental variables approach. The first instrument I use is based on immigration enclaves and has a long history in the literature that studies the impacts of skill mix on local labor market outcomes. As a second instrument set, I use average city climate variables to proxy for local amenities and show that it successfully predicts changes in city's educational composition. I also address possible bias arising from self-selected migration of workers by using a selection correction procedure developed by Dahl (2002). Despite addressing these econometric concerns, the instability of the estimated spill-overs remains.

This analysis has important implications for the interpretation of human capital externalities. Documenting the instability of education spill-overs helps differentiate between two classes of models that give rise to human capital spill-overs. I argue that the observed patterns of education spill-overs are not consistent with models of technological human capital externalities.[26] In these models externalities are built directly into the production function and could reasonably be expected to be smooth functions

[26]This is not to say that technological externalities are unimportant, but that they may be more relevant at a more local or firm level (Bidner 2007).

of the supply of skilled labor. In addition, as mentioned above, I find that spill-overs not only fall for low-skilled workers but are actually negative. Since these models are usually motivated by the sharing of knowledge and skills between workers, it is hard to understand why such worker interaction would become unimportant in the 1990s when other dimensions of skills saw increased returns. Thus, these findings suggest that we should look toward alternative models of externalities that can potentially explain these patterns.

I examine and provide support for one potential explanation of the patterns of spill-overs observed in the data. In particular, I examine Acemoglu's (1999) model of endogenous job composition in which firms choose a production technique based on the skill level of the workforce. In this model, the observed pattern of education spill-overs can be rationalized by a structural change in the labor market driven either by an increase of in the supply of skilled workers and/or technical change. Despite being highly stylized, this model also has a rich set of testable implications which I take to the data in the second part of this paper and find a considerable match between theory and data.

The main finding of this paper is that there has been a structural change in the labor market between the 1980s and the 1990s that has harmed low-skill workers. While technical change has likely played a role in this process, *a priori* it is not clear how technology could induce the observed changes. This paper stresses the interaction between technology and the supply of skills as emphasized by the model. I find that the key to understanding city level outcomes is the supply of skilled workers. This alters our assessment both of the role of education in the economy and of the type of technical change we have been witnessing. Consistent with the notion that there has been a structural change in the labor market, increases in the supply of skilled labor in the 1990s induced a change in the composition of jobs, increasing inequality, unemployment, returns to education, and the wages of high-skill workers while *harming* low-skill workers.

Figures

Figure 1.2: Distribution of City College Share

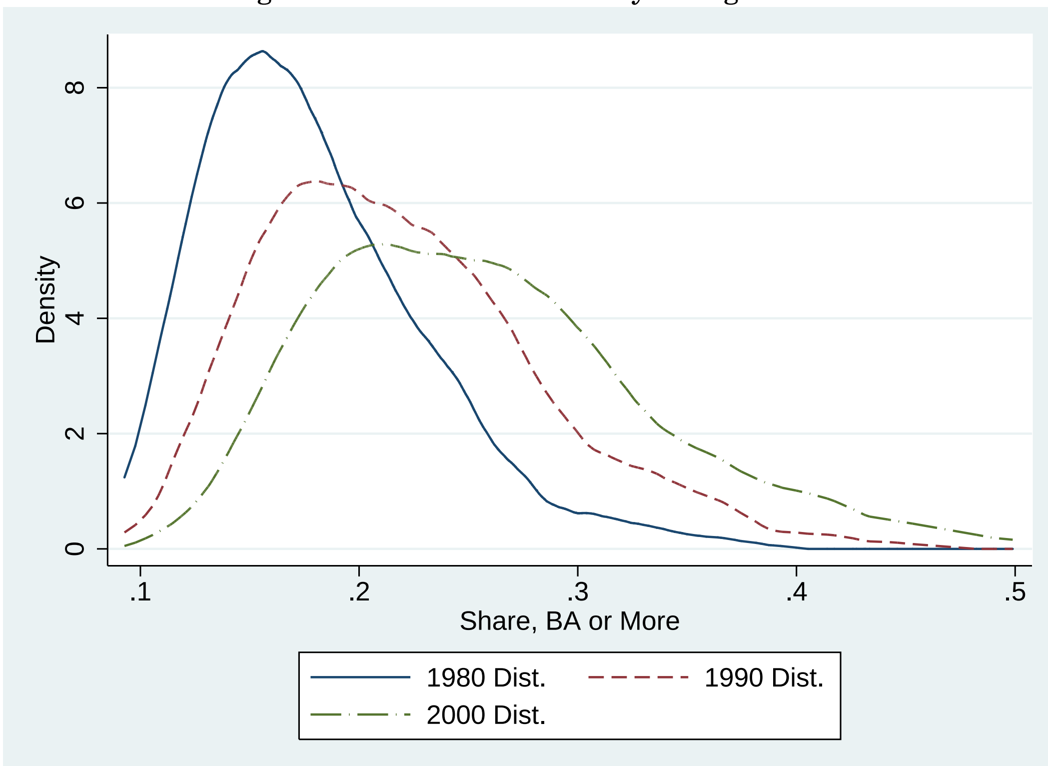

Figure 1.3: Smoothed Beta Estimates: High School Workers

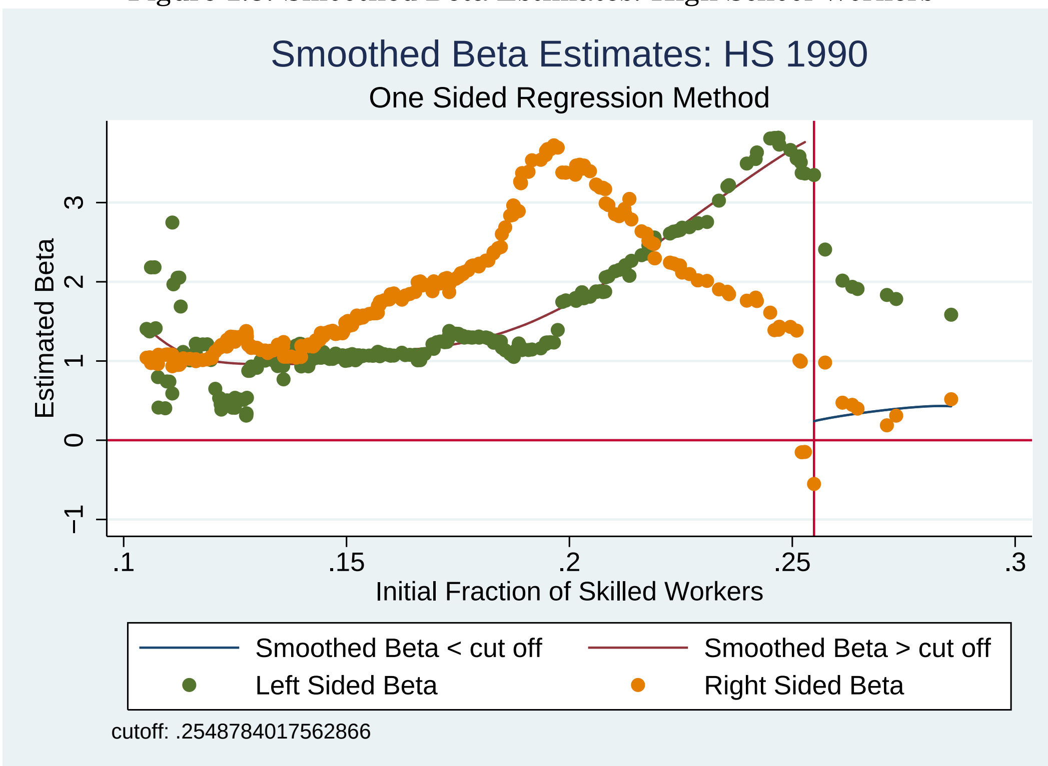

Notes: This figure is generated by the left- and right-sided smoothing procedure as described in the text. The orange dots (diamonds) represent the right-sided estimates and the green dots (circles) represent the left-sided estimates. The figure uses data on all cities with initial shares greater than 0.13, which captures around 85% of the variation. The cut-off estimate is sensitive to this restriction. The bandwith used is 0.065.

Tables

Table 1.1: The Effect of Changes in College Share on Wages of Education Groups

	OLS		Enclave IV		Climate IV	
	1980-1990	1990-2000	1980-1990	1990-2000	1980-1990	1990-2000
Subsample	(1)	(2)	(3)	(4)	(5)	(6)
All Workers	1.515	-.187	2.316	-.654	2.194	-.875
	(0.228)*	(0.143)	(0.515)*	(0.477)	(0.547)*	(0.299)*
High School	1.621	-.433	2.952	-.610	2.524	-1.306
	(0.277)*	(0.157)*	(0.604)*	(0.562)	(0.645)*	(0.396)*
BA or >	1.494	0.505	1.573	-.296	1.829	0.51
	(0.153)*	(0.197)*	(0.392)*	(0.485)	(0.405)*	(0.334)
Obs.	286	286	152	152	286	286
R^2	0.467	0.235				
F-Stat.			15.86	25.02	4.41	10.64
P-Val.			0.0	0.0	0.0	0.0

Notes: Each entry is a separate regression. Each cell contains the coefficient on the fraction of BA or > workers in a city. The dependent variables are regression adjusted wages of different education groups. All regressions control for city level characteristics. F-Stat refers to the F-statistic on the instruments in the first stage regression for the entire sample.

Table 1.2: The Effect of Changes in College Share on Wages of Education Groups: Selection Correction

	OLS		Enclave IV		Climate IV	
	1980-1990	1990-2000	1980-1990	1990-2000	1980-1990	1990-2000
Subsample	(1)	(2)	(3)	(4)	(5)	(6)
All Workers	1.522	-.127	2.366	-.673	2.257	-.733
	(0.224)*	(0.142)	(0.498)*	(0.463)	(0.551)*	(0.285)*
High School	1.663	-.404	3.074	-.628	2.677	-1.283
	(0.278)*	(0.158)*	(0.611)*	(0.558)	(0.654)*	(0.383)*
BA or >	1.343	0.519	1.463	-.432	1.884	0.665
	(0.154)*	(0.196)*	(0.369)*	(0.466)	(0.426)*	(0.331)*
Obs.	286	286	152	152	286	286
R^2	0.485	0.252				
F-Stat			15.93	25.22	4.37	10.62
P-Val.			0.0	0.0	0.0	0.0

Notes: Each entry is a separate regression. Each cell contains the coefficient on the fraction of BA or > workers in a city. The dependent variables are regression adjusted wages of different education groups that include the selection correction described in text. All regressions control for city level characteristics. F-Stat refers to the F-statistic on the instruments in the first stage regression for the entire sample.

Table 1.3: Deciles - OLS

		Residual Decile								
	Mean	10	20	30	40	50	60	70	80	90
Panel A:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
1980s	1.525	1.427	1.536	1.560	1.538	1.503	1.478	1.479	1.490	1.641
	(0.226)*	(0.281)*	(0.27)*	(0.248)*	(0.232)*	(0.222)*	(0.216)*	(0.21)*	(0.209)*	(0.218)*
1990s	-.190	-.714	-.492	-.360	-.253	-.179	-.081	0.004	0.125	0.399
	(0.149)	(0.212)*	(0.183)*	(0.16)*	(0.147)	(0.136)	(0.134)	(0.133)	(0.146)	(0.177)*
Obs. = 286										
R^2	0.915	0.879	0.895	0.905	0.913	0.918	0.921	0.924	0.925	0.921

Panel B:	Enclave IV									
1980s	2.310	2.492	2.424	2.353	2.371	2.288	2.195	2.263	2.287	2.442
	(0.507)*	(0.649)*	(0.599)*	(0.556)*	(0.533)*	(0.516)*	(0.494)*	(0.477)*	(0.472)*	(0.478)*
1990s	-.666	-1.887	-1.225	-1.007	-.764	-.578	-.349	-.305	-.095	0.395
	(0.488)	(0.828)*	(0.617)*	(0.548)	(0.487)	(0.441)	(0.412)	(0.376)	(0.383)	(0.37)
Obs. = 152										

Panel C:	Climate IV									
1980s	2.114	1.852	2.067	2.121	2.169	2.197	2.211	2.177	2.219	2.366
	(0.53)*	(0.645)*	(0.613)*	(0.575)*	(0.546)*	(0.529)*	(0.521)*	(0.51)*	(0.501)*	(0.517)*
1990s	-.884	-1.809	-1.528	-1.273	-1.075	-.871	-.663	-.479	-.276	0.351
	(0.302)*	(0.447)*	(0.387)*	(0.353)*	(0.322)*	(0.298)*	(0.281)*	(0.263)	(0.264)	(0.283)
Obs. = 286										

Notes: Each entry is a separate regression. Each cell contains the coefficient on the fraction of BA or > workers in a city. The dependent variables are different deciles of regression adjusted wages. All regressions control for city level characteristics.

Table 1.4: Residual Inequality and College Graduate Share by Decade

	OLS			Enclave IV			Climate IV		
	90-10	50-10	90-50	90-10	50-10	90-50	90-10	50-10	90-50
Years	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
1980s	0.215	0.076	0.139	-.050	-.204	0.154	0.514	0.345	0.168
	(0.149)	(0.096)	(0.082)	(0.381)	(0.255)	(0.18)	(0.365)	(0.277)	(0.169)
1990s	1.113	0.535	0.578	2.282	1.309	0.973	2.160	0.938	1.222
	(0.189)*	(0.125)*	(0.106)*	(0.646)*	(0.471)*	(0.264)*	(0.395)*	(0.259)*	(0.223)*
Obs.	286	286	286	152	152	152	286	286	286
R^2	0.218	0.269	0.172						

Notes: Each entry is a separate regression. Each cell contains the coefficient on the fraction of BA or > workers in a city. The dependent variables are different measures of residual wage inequality. The IV estimates utilize either immigrant enclave instrument or the climate instruments. All regressions control for city level characteristics.

Table 1.5: Returns to College and Franction of Skilled Worders by Decade

	Selection					
	Uncorrected			Corrected		
		IV			IV	
	OLS	Enclave	Climate	OLS	Enclave	Climate
	(1)	(2)	(3)	(4)	(5)	(6)
1980s	-.252	-1.474	-.882	-.607	-1.602	-1.505
	(0.197)	(0.646)*	(0.49)†	(0.221)*	(0.891)†	(0.628)*
1990s	0.88	0.383	1.805	0.608	0.517	1.268
	(0.199)*	(0.492)	(0.414)*	(0.254)*	(0.526)	(0.44)*
Obs.	286	152	286	286	152	286
R^2	0.464			0.346		

Notes: Each entry is a separate regression. Each cell contains the coefficient on the fraction of BA or > workers in a city. The dependent variable is the return to exactly a college degree, relative to a high school degree, estimated separately by year and city. The IV estimates utilize either the immigrant enclave instrument or the climate instruments. All regressions control for city level characteristics.

Table 1.6: Unemployment

Panel A:	All (1)	< High School (2)	High School (3)	Some Post. (4)	BA or > (5)
1980s	-.008	0.049	0.082	0.044	0.001
	(0.055)	(0.095)	(0.066)	(0.047)	(0.022)
1990s	0.069	0.192	0.097	0.09	-.006
	(0.039)†	(0.096)*	(0.058)*	(0.038)*	(0.018)
Obs.	572	572	572	572	572
R^2	0.235	0.301	0.288	0.145	0.113

Panel B:			Enclave Instrument		
1980s	-.424	-.748	-.324	-.144	-.060
	(0.201)*	(0.326)*	(0.217)	(0.099)	(0.057)
1990s	0.265	0.898	0.28	0.117	0.06
	(0.123)*	(0.318)*	(0.195)	(0.117)	(0.063)
Obs.	304	304	304	304	304

Panel C:			Climate Instrument		
1980s	-.780	-1.320	-.714	-.487	-.234
	(0.209)*	(0.379)*	(0.231)*	(0.162)*	(0.079)*
1990s	0.313	0.748	0.332	0.288	0.126
	(0.083)*	(0.217)*	(0.11)*	(0.08)*	(0.048)*
Obs.	572	572	572	572	572

Notes: Each entry is a separate regression. Each cell contains the coefficient on the fraction of BA or > workers in a city. The dependent variables are the city level unemployment rates of different education groups. The IV estimates utilize either the immigrant enclave instrument or the climate instruments. All regressions control for city level characteristics.

Table 1.7: Breakdown of Occupations in the Bottom 25 Skill Percentiles

Broad Occupation	**Percent in 25th or Less**	**N**
Managerial	2.11	2
Professional Specialty	10.5	10
Sales Occ.	4.21	4
Admin./Clerical	7.37	7
Service	29.5	28
Farm, Fish, Forestry	12.6	12
Precision Craft	12.6	12
Laborers and Operators	12.6	12
Transport/Material Handling	8.42	8

Source: U.S. Census

Table 1.8: Changes in the Composition of Employment: 1980-2000

Year	25th or <	25th to 75th	75th or >	Weight at Tail
All Workers				
1980	0.183	0.556	0.261	0.444
1990	0.191	0.555	0.254	0.445
2000	0.210	0.521	0.269	0.479
High School or Less				
1980	0.243	0.566	0.191	0.434
1990	0.286	0.556	0.159	0.444
2000	0.324	0.520	0.157	0.480
BA or More				
1980	0.074	0.518	0.408	0.482
1990	0.077	0.505	0.418	0.495
2000	0.094	0.483	0.423	0.517

Source: U.S. Census 1980, 1990, 2000

Table 1.9: Low-Skill Workers: Employment in Jobs Ranked Below the 25th Percentile

	OLS		Enclave IV		Climate IV	
	1980-1990	1990-2000	1980-1990	1990-2000	1980-1990	1990-2000
	(1)	(2)	(3)	(4)	(5)	(6)
Occ. Rank	1980 Residual					
Δ BA or >	0.089	0.34	0.03	0.628	0.15	0.707
	(0.053)	(0.086)*	(0.212)	(0.335)†	(0.149)	(0.18)*
Occ. Rank	1980 Residual including Industry Controls					
Δ BA or >	0.043	0.251	0.154	0.383	0.33	1.040
	(0.052)	(0.084)*	(0.225)	(0.31)	(0.153)*	(0.221)*
Occ. Rank	1980 High School or Less Log Wages					
Δ BA or >	0.041	0.286	0.049	0.648	0.07	0.753
	(0.055)	(0.078)*	(0.224)	(0.326)*	(0.145)	(0.166)*
Occ. Rank	1980 Grouped by Broad Industry and Occupation					
Δ BA or >	0.216	0.454	0.133	0.924	0.592	0.698
	(0.051)*	(0.077)*	(0.188)	(0.313)*	(0.18)*	(0.17)*
Obs.	286	286	152	152	286	286

Notes: Each entry is a separate regression. Each cell contains the coefficient on the fraction of BA or > workers in a city. The dependent variables are either the percent change in employment or change in average log wage of occupations in the bottom 25th percentile of the 1980 skill distribution. All regressions control for city level characteristics.

Table 1.10: Low-Skill Workers: Wages in Jobs Ranked Below the 25th Percentile

	OLS		Enclave IV		Climate IV	
	1980-1990	1990-2000	1980-1990	1990-2000	1980-1990	1990-2000
Variables	(1)	(2)	(3)	(4)	(5)	(6)
Occ. Rank	1980 Residual					
Δ BA or $>$	1.554	-.839	3.186	-1.848	3.032	-2.060
	(0.316)*	(0.228)*	(0.7)*	(0.825)*	(0.707)*	(0.487)*
Occ. Rank	1980 Residual including Industry Controls					
Δ BA or $>$	1.497	-.661	2.753	-1.080	2.961	-1.761
	(0.321)*	(0.235)*	(0.694)*	(0.832)	(0.736)*	(0.509)*
Occ. Rank	1980 High School or Less Log Wages					
Δ BA or $>$	1.608	-.834	3.222	-1.450	3.066	-2.361
	(0.3)*	(0.224)*	(0.704)*	(0.806)	(0.69)*	(0.529)*
Occ. Rank	1980 Residual Grouped by Broad Industry and Occupation					
Δ BA or $>$	1.727	-.953	3.426	-1.726	3.561	-2.232
	(0.334)*	(0.235)*	(0.789)*	(0.832)*	(0.777)*	(0.499)*
Obs.	286	286	152	152	286	286

Notes: Each entry is a separate regression. Each cell contains the coefficient on the fraction of BA or $>$ workers in a city. The dependent variables are either the percent change in employment or change in average log wage of occupations in the bottom 25th percentile of the 1980 skill distribution. All regressions control for city level characteristics.

Table 1.11: Cities Above Threshold: 1980s

	High School Wages	BA or > Wages	Unemployment Rate	Inequality 90-10	Returns to College	Employment in Jobs < 25th Pct. Ranked	Wages in Jobs < 25th Pct. Ranked
Panal A	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Δ BA or >	-.601	0.738	-.293	0.646	1.139	0.258	-1.342
	(0.747)	(0.491)	(0.091)***	(0.564)	(0.691)*	(0.25)	(0.726)*
Const.	-.002	-.009	0.005	-.007	-.020	0.004	0.005
	(0.018)	(0.012)	(0.002)**	(0.013)	(0.016)	(0.006)	(0.017)
Obs.	17	17	17	17	17	17	17
R^2	0.041	0.131	0.411	0.08	0.153	0.066	0.185
Panel B				Cities below Threshold: 1980s			
Δ BA or >	1.834	1.538	0.017	0.177	-.374	-.039	1.837
	(0.2)***	(0.152)***	(0.051)	(0.128)	(0.151)**	(0.057)	(0.208)***
Const.	0.0005	0.0007	-.0003	0.0004	0.001	-.0003	0.0001
	(0.004)	(0.003)	(0.001)	(0.003)	(0.003)	(0.001)	(0.004)
Obs.	269	269	269	269	269	269	269
R^2	0.24	0.278	0.0004	0.007	0.022	0.002	0.225

Notes: Each entry is a separate regression. Each cell contains the coefficient on the fraction of BA or > workers in a city. The dependent variables are either the percent change in employment or change in average log wage of occupations in the bottom 25th percentile of the 1980 skill distribution. All regressions control for city level characteristics.

Chapter 2

The Value of Good Jobs: A General Equilibrium Perspective on a Recurring Debate

Introduction

In popular discussion about labor market developments, whether it be at the local or national level, changes in the nature of jobs are often given a pre-eminent role. In particular, it is often claimed that labor market performance hinges on whether an economy is attracting or losing "good jobs": that is, jobs in industries which pay a premium relative to wages for similarly qualified workers in other industries. For example, in Bluestone and Harrison's highly cited 1982 book, *The Deindustrialization of America*, the authors argued that the switch away from highly paid manufacturing jobs was key to understanding the poor labor market performance of the U.S. economy during the 1970s and 1980s. However, based on simple accounting exercises aimed at assessing the potential importance of changes in industrial composition on average wages, most serious economic researchers have dismissed such views as being ill-informed. Interestingly, the populist view is now reemerging in economic discussions in many countries because of perceptions about the effects of globalization on industrial composition and wages. For this reason it appears to be an opportune time to re-visit the issue.

In this paper we present theory and evidence aimed at evaluating the effects of changes in industrial composition on wages. The distinguishing feature of our approach is an emphasis on non-Walrasian general equilibrium effects. In particular, building on a search and bargaining model of the labor market, we highlight why wages in different sectors may interact as strategic complements giving rise to rich general equilibrium interactions. Such interaction between sectors implies that wage determination takes the form of a classical social interaction problem –also known as a reflection problem (see Manski (1993), Moffitt (2001)). The object of our empirical work is to evaluate the relevance of such cross-sector links for the determination of wages.

It is helpful to begin by contrasting our approach with the common practice of using a simple accounting procedure for examining the effects of a shift in industrial composition on average wages. The accounting approach (also known as shift-share analysis) consists of multiplying the change in the proportion of workers in a given industry in a city by the average wage in that industry then summing across industries. The result from such an exercise almost always indicates that the average wage change directly accounted for by changes in sectoral composition of employment is small. The validity of the accounting approach hinges critically on the assumption that a change in employment composition between two sectors does not affect wages in other sectors if the change does not effect the employment level in these other sectors; i.e., there are no general equilibrium effects on wages of a shift in the composition of the industrial structure. Without this assumption, one would also have to account for changes in the within sector wages arising from the compositional shifts, destroying the clean break into "within" (premia change) and "between" (composition change) components that is a key feature of the accounting approach.[27]

Our goal is to examine whether changes in sectoral composition of employment, especially shifts in composition between high and low paying sectors, have important general equilibrium effects on the determination of within sector wages. To help clarify our question and to help direct the empirical strategy to answer it, we begin by showing how a standard search and matching model of the labor market,[28] extended to include many industries, implies general equilibrium effects of changes in industrial composition that would invalidate the use of the standard accounting approach. In particular, the search and matching framework implies that a change in industrial composition, through its effect on the bargaining position of workers, can affect wages in sectors not directly involved in the compositional change as long as workers are potentially mobile across sectors. In this case, an improved outside option for workers will place upward pressure on wages even if total employment is unaffected. This implication of search and matching is very basic and implies that wages in different sectors act as strategic complements, and that the general equilibrium determination of wages takes the form of a social interaction or reflection problem (Manski (1993), Moffitt (2001)). Although the insight is straightforward, to the best of our knowledge it is an empirical implication of this class of models which has not previously been extensively pursued.[29] One possible

[27] There are many ways to justify the no-GE effects assumption, which is part of the appeal of this approach. The easiest defense is to note that if wages are simply a function of productivity and returns to labor are close to constant, one just needs to assume that changes in industrial composition do not change productivity within sectors to arrive at the conclusion that there are no GE effects. The latter assumption might be viewed as innocuous by many economists, but it is the one we want to place into question.

[28] For an introduction to search and matching models of the labor market see Mortensen and Pissarides (1999)

[29] One side implication of our model is a relationship between wages and the employment rate in a local economy that echoes that explored in Blanchflower and Oswald (1995). In this sense, our results are

reason for this is that its empirical exploration involves many hurdles, including solving a reflection problem and controlling for the endogeneity of industrial composition. The theoretical framework we use is particularly helpful in providing guidance on how to discuss and address issues related to endogeneity, potential instrumental variables, and interpretation of estimates.

In order to examine the relevance of adopting a social interaction model for examining the general equilibrium effects associated with changes in the fraction of jobs that are "good jobs", we exploit geographical variation in industrial composition across U.S. cities over the period 1970-2000. Our approach is to look at whether wage changes within a given industry vary systematically across cities with changes in the predicted distribution of employment across other industries in each city. We are careful in our theoretical section to show that the implications we explore should be present even if worker mobility causes expected utility to be equalized across cities. While we do not impose the structure of our model on the data, our empirical specification is closely aligned with the theory we present. In particular, we examine whether wages in any given industry-city cell tend to increase in a city which has witnessed an (externally driven) change in industrial composition that is more concentrated toward high paying jobs. To do this, we examine 10 year changes in industry $\times$ city level wages using data from the 1970, 1980, 1990 and 2000 U.S. Censuses for 152 cities.[30] We devote considerable effort to addressing endogeneity issues associated with solving the reflection problem inherently related to strategic complementarities, and we also address possible non-random selection in unobservable worker characteristics across cities using the method proposed in Dahl (2002).

The main empirical result of the paper is that city level changes in industrial composition have effects on average wages that are 3 to 3.5 times that implied by a pure accounting approach, indicating the presence of substantial general equilibrium effects. It is important to keep in mind when considering this result that measured composition effects are often small, so that the effects we find are large but not extreme.[31] Impacts

in line with their argument that the wage curve they identify fits with a number of models, including a bargaining-based model. We differ from that work, and from other research on bargaining effects on wage setting, in that we emphasize the importance of industrial composition on setting wages within all sectors in an economy.

[30] There are both advantages and disadvantages of using city level observations to examine this issue. An attractive feature is there are over 150 metropolitan areas in the U.S. which gives us a sample with many different patterns of industrial composition. On the other hand, if labor markets across U.S. cities are all perfectly integrated, then we could find very small general equilibrium effects even if such effects are large at the national level. In this sense, this study may a priori be seen as biased toward finding small or no general equilibrium effects of industrial composition even if they are present.

[31] For example, even the seemingly large event of a city losing an industry that employed 10% of the workforce and paid a premium of 20% relative to other industries (roughly the situation facing Pittsburgh with the loss of the steel industry in the 1980s) implies only a 2% drop in the average wage using the pure accounting approach. Our result says that the total impact on city average wages would be a 6 to 7% decline.

of the size we identify give credence to explanations for changes in wage patterns that operate through changes in industrial structure; explanations which have largely been discounted because the pure accounting measures of their impacts are relatively small (e.g., Bound and Johnson (1992)). Furthermore, our results provide considerable support for search and matching models of the labor market as a framework for understanding wage determination in a decentralized environment. For example, we find that wages in different sectors act like strong strategic complements in precisely the manner predicted by decentralized bargaining. This result is important in its own right since it suggests that the general equilibrium determination of wages may best be viewed as a social interaction problem. While search and bargaining models are often cited as helping understand unemployment, we believe that our results provide very strong evidence in support of this type of model for understanding wage behavior in the U.S. and, potentially, other market economies.

It is important to emphasize that we are interested in longer term differences in wage levels associated with different industrial composition as opposed to short run adjustments to industrial change and the related change in the level of employment. In particular, we focus mainly on changes in industrial composition that are orthogonal to changes in city level employment rates. Thus, our focus is very different from, for example, that in Greenstone and Moretti (2003) or Blanchard and Katz (1992). While Greenstone and Moretti emphasize shorter run demand effects by focussing mainly on local real estate price and same-industry wage bill effects within three years of acquiring a large plant, we focus on changes arising over 10 year horizons and are investigating whether there are general equilibrium effects of changes in industrial composition holding direct demand effects constant. This focus also differentiates our work from studies of regional adjustment to demand changes such as in Blanchard and Katz (1992). In order to clarify this difference, we take care in our empirical work to control for the types of demand effects examined in these papers.

Our empirical results are also related to a set of papers which examine the causes of city level employment and wage growth. In that literature, strong city performance has been variously linked to city size, the diversity of employment in a city (Glaeser, Kallal, Scheinkman, and Shleifer 1992), and the concentration of educated workers in a city (Moretti 2004). In our empirical work, we introduce measures capturing each of these effects and show that these do not change our main findings. Thus, what we are identifying is over and above these other hypothesized driving forces. Our paper is also related to the voluminous literature aimed at understanding the effects of international trade on wages since much of that literature has debated the potential effects of trade-induced changes in industrial composition. The paper most closely related to ours is Borjas and Ramey (1995), which uses city level variation similar to ours to examine how trade induced changes in industrial composition may have affected returns to skill. Our focus is, nevertheless, very different since we focus on wage levels rather than on returns to skill, and we emphasize the relevance of the social interaction model for understanding the general equilibrium determination of wages.

The remaining sections of the paper are as follows. In Section 1, we present a model with search frictions and bargaining to illustrate how changes in industrial composition in an economy can affect wage setting within all sectors in the economy - including wages in sectors not directly involved in any composition shift. In Section 2, we use the model to derive a general empirical specification which embeds alternative views about the determination of wages. In particular, our empirical specification allows us to examine whether the data support the view that general equilibrium effects from industrial composition changes are small and that the accounting approach is a justifiable procedure. In Section 3 we discuss the data used in the study and report basic empirical results. In Section 4 and 5 we address issues related to endogeneity, selection, robustness and interpretation. Section 6 concludes.

2.1 The General Equilibrium Effects of Changes in Industrial Composition in a Search and Bargaining Model

In this section, we illustrate how a shift in the composition of jobs in a local economy can affect wages within all sectors in the economy, including sectors that are not part of the shift. We discuss this in the context of a general equilibrium search and bargaining model, showing that such a composition effect is separate from wage effects arising from shifts in the overall level of demand in the economy. For example, if a city experiences the loss of a high wage paying industry, one would expect an effect on local wages due to a reduction in overall demand. However, we show that, even holding the employment rate constant, a standard bargaining model also implies that the resulting shift in the composition of employment has an impact on wages in all sectors in the economy. Whether that composition effect is sizable is, of course, an empirical issue. The goal of this section is to derive an empirical strategy to explore whether changes in job composition have important general equilibrium effects on within sector wage determination.

In order to illustrate the mechanisms for why job composition may affect the determination of wages, we consider an economy with one final good, denoted Y, which is an aggregation of output from I industries as given by

$$Y = \left(\sum_{i=1}^{I} a_i Z_i^{\chi} \right)^{1/\chi}, \quad \text{where } \chi < 1. \tag{2.1}$$

The price of the final good is normalized to 1, while the price of the good produced by industry i is given by P_i. In this economy, we assume that there are C local markets called cities and that the industrial goods can be produced in any of these markets. The total quantity of the industrial good Z_i produced across the economy is equal to the sum across cities of X_{ic}, the output in industry i in city c.

To greatly simplify the exposition, we begin by examining the extreme case where workers are not mobile across cities. In the next subsection, we will show that the results derived in this section are robust to an extension of the model allowing for worker mobility and endogenous housing prices.

To create a job in industry i in city c, a firm must pay a cost of c_{ic}, the value of which will be endogenously determined in equilibrium. If a job is filled, it generates a flow of profits for the firm given by

$$p_i - w_{ic} + \epsilon_{ic},$$

where w_{ic} is the wage, ϵ_{ic} is a city-industry specific cost advantage and $\sum_c \epsilon_{ic} = 0$. If we denote by V^f the discounted value of profits from a filled position, and we denote by V^v the discounted value of a vacancy, then in steady state V^f and V^v must satisfy the standard Bellman relationship given by

$$\rho V_{ic}^f = (p_i - w_{ic} + \epsilon_{ic}) + \delta(V_{ic}^v - V_{ic}^f), \tag{2.2}$$

where ρ is the discount rate and δ is the exogenous death rate of matches. If a firm does not fill a job it faces a per-period cost of r_i to maintain the position. Thus, the discounted value of profits from a vacant position must satisfy

$$\rho V_{ic}^v = -r_i + \phi_c(V_{ic}^f - V_{ic}^v), \tag{2.3}$$

where ϕ_c is the probability a firm fills a posted vacancy. For simplicity, and without loss of generality, we set $r_i = 0$.

Workers in the economy can be either employed or unemployed in a given period. The discounted value of being employed in industry i in city c, denoted U_{ic}^e, must satisfy the Bellman equation:

$$\rho U_{ic}^e = w_{ic} + \delta(U_{ic}^u - U_{ic}^e), \tag{2.4}$$

where U^u_{ic} represents the value associated with being unemployed when the worker's previous job was in industry i. When an individual is unemployed, he gets flow utility from an unemployment benefit, b, plus a city specific amenity term, τ_c.[32] Letting ψ_c represent the probability that an unemployed individual finds a job, and $1-\mu$ represent the probability that an individual finding a job gets a random draw from jobs in all industries (including i) versus being assigned a match in his previous industry,[33] the value associated with being unemployed satisfies:

$$\rho U^u_{ic} = b + \tau_c + \psi_c(\mu U^e_{ic} + (1-\mu)\sum_j \eta_{jc} U^e_{jc} - U^u_{ic}). \tag{2.5}$$

The important aspect to note from equation (2.5) is that, as long as $\mu < 1$, the utility level associated with having lost one's job in industry i depends on the utility associated with jobs in other industries. The instantaneous probability of finding a job in industry j is given by $\psi_c(1-\mu)\eta_{jc}$, where η_{ic} represents the fraction of vacant jobs that are in industry j. This implies that a worker finding a job in another industry is assumed to find it in proportion to the relative size of that industry. In this formulation, we are assuming that workers can only search while being unemployed. While this is a strong assumption, it has the attractive feature of allowing the problem to be solved explicitly and thereby amenable to simple empirical exploration. For these reasons we maintain this assumption throughout and leave for future work the implications for industrial composition with on-the-job search.

Once a match is made, workers and firms bargain a wage,[34] which is set according to the bargaining rule:

$$(V^f_{ic} - V^v_{ic}) = (U^e_{ic} - U^u_{ic}) \times \kappa, \tag{2.6}$$

[32] We have not included an amenity term in the flow utility when employed to simplify notation. No results would change if we included such a term since what is important for wage determination is the difference in utility between being unemployed or employed. The amenity term τ should therefore be interpreted as the difference in flow utility associated with amenities when unemployed versus when someone is employed.

[33] We could easily add to an unemployed worker's trajectory the possibility of drawing an offer from another city. For example, if, with probability Γ, an unemployed worker from industry i in city c sampled jobs in industry i from other cities, then the value functions would need to satisfy:

$$\rho U^u_{ic} = b + \tau_c + (1-\Gamma)\psi_c(\mu U^e_{ic} + (1-\mu)\sum_j \eta_{jc} U^e_{jc} - U^u_{ic}) + \Gamma(\sum_c \psi_c \frac{N_{ic'}}{\sum_c N_{ic'}} U^e_{ic'} - U^u_{ic}),$$

where $N_{ic'}$ is the number of jobs in industry i in city c'. Since the additional term $\sum_{c'} \psi_c \frac{N_{ic'}}{\sum_c N_{ic'}} U^e_{ic'}$ would act as a common element for workers across all cities, it would not change the implications we are focusing on which relate to within city effects. For this reason, we can disregard this possibility without loss.

[34] We assume throughout that there are always gains from trade between workers and firms for all jobs created in equilibrium.

where κ is a parameter governing the relative bargaining power of workers and firms. The probability a match is made is determined by the matching function:

$$M\left((L_c - E_c), (N_c - E_c)\right),$$

where L_c is the total number of workers in city c, E_c is the number of employed workers (or matches) in city c, and $N_c = \sum_i N_{ic}$ is the number of jobs in city c, with N_{ic} being the number of jobs in industry i in city c. Then, given the exogenous death rate of matches, δ, and assuming that the match function is of Cobb-Douglas form, the steady state condition is given by

$$\delta ER_c = M\left((1 - ER_c), \left(\frac{N_c}{L_c} - ER_c\right)\right) = (1 - ER_c)^{\sigma}\left(\frac{N_c}{L_c} - ER_c\right)^{1-\sigma}, \tag{2.7}$$

where ER_c is the employment rate. It follows that the proportion of filled jobs and vacant jobs in industry i, can be expressed as $\eta_{ic} = \frac{N_{ic}}{\sum_i N_{ic}}$.

The number of jobs created in industry i in city c, N_{ic}, is determined by the free entry condition:

$$c_{ic} = V^{v}_{ic}. \tag{2.8}$$

The cost c_{ic} should be viewed as the cost of creating a marginal job. If this cost were fixed, then cities would generally specialize in only one industry. Since we want to generate cities with employment across a wide range of industries, we assume that c_{ic} is increasing in the number of new jobs being created locally in that industry, which, in equilibrium, is proportional to N_{ic}. We also want to allow cities to have a comparative advantage in creating certain types of jobs relative to others. We therefore assume that c_{ic} is a decreasing function of the industry-city specific measure of advantage denoted Ω_{ic}. For tractability, we assume that the relationship between c_{ic}, N_{ic} and Ω_{ic} is given by:

$$c_{ic} = \frac{N_{ic}}{\Upsilon_i + \Omega_{ic}}.$$

In the above equation we have assumed that c_{ic} is proportional to N_{ic}, but this assumption can easily be relaxed. The term Υ_i reflects any systematic differences in cost of entry across industries, which allows us to assume that $\sum_c \Omega_{ic} = 0$.

Finally, the probability an unemployed worker finds a match and the probability a firm fills a vacancy satisfy

$$\psi_c = \frac{\delta ER_c}{1 - ER_c} \quad \text{and} \quad \phi_c = \left(\frac{1 - ER_c}{\delta ER_c}\right)^{\frac{\sigma}{1-\sigma}}, \tag{2.9}$$

respectively.

At the city level, the price of industrial output is taken as given and an equilibrium consists of values of N_{ic}, w_{ic}, and ER_c that satisfy equations (2.6), (2.7) and (2.8). Note that these equilibrium values will depend upon (among other things) the city specific productivity parameters Ω_{ic} and ϵ_{ic}. An equilibrium for the entire economy has the additional requirement that the prices for industrial goods must ensure that markets for these goods clear.

At the city level, the equilibrium outcomes will reflect how a city adjusts its production mix to take advantage of the different prices of the industrial goods in relation to its comparative technological advantages in producing the goods. At the economy wide level, the change in prices will be caused by changes in demand for the industry level goods, captured in the a_is. As we will make clear, our focus will be on isolating the effect of a change in job composition on city level wages. To this end, we will take the above description of a steady state equilibrium as representing an equilibrium at a point in time and compare how this equilibrium changes in response to changes in the exogenous driving forces a_i, Ω_{ic}, and ϵ_{ic}.

2.1.1 The Interaction between Sectoral Wages and the Associated Reflection Problem

Our focus will be on the determination of wages, as implied by (2.6). Our goal is to highlight how sectoral level wages inherit a strategic complementarity property in the presence of wage bargaining that gives rise to potentially rich general equilibrium effects of changes in industrial composition on wages. For now we will treat η_{ic} and ER_c as given, leaving the discussion of their endogenous determination and its implications for later.

To understand the forces determining wages in an industry-city cell, we begin by using equations (2.4) and (2.5) to express the value of finding a job relative to being unemployed as

$$U_{ic}^e - U_{ic}^u = \frac{w_{ic} - b - \tau_c}{\rho + \delta + \psi_c \mu} - \frac{\psi_c(1-\mu)\sum_j \eta_{jc}(w_{jc} - b - \tau_c)}{(\rho + \delta + \psi_c \mu)(\rho + \delta + \psi_c)}. \tag{2.10}$$

A key feature of equation (2.10) is that, as long as $\mu < 1$, a worker's utility from being employed relative to being unemployed is a decreasing function of the wages paid in other sectors of the city's economy. This arises because higher wages elsewhere in the economy imply a greater value of staying unemployed and finding a job in another sector. Also note that a compositional change captured by a change in the ηs, holding ψ_c constant, will in general affect the utility of finding a job. For example, if the composition change implies more concentration of jobs in high wage sectors this decreases the value of a match in sector i.

To express the value of a match to a firm, we can use equations (2.2) and (2.3) in a similar fashion:

$$(V_{ic}^f - V_{ic}^v) = \frac{p_i - w_{ic} + \epsilon_{ic}}{\rho + \delta + \phi_c}. \tag{2.11}$$

From equations (2.10) and (2.11), we can use equation (2.6) to solve for w_{ic}. This is given by

$$w_{ic} = \gamma_{c0} + \gamma_{c1} p_i + \gamma_{c2} \sum_j \eta_{jc} w_{jc} + \gamma_{c1} \epsilon_{ic}, \tag{2.12}$$

where the coefficients in (2.12) are $\gamma_{c0} = \frac{(\rho+\delta+\psi_c\mu)(\rho+\delta+\phi_c)\kappa}{[(\rho+\delta+\psi_c\mu)+\kappa(\rho+\delta+\phi_c)](\rho+\delta+\psi_c)}(b+\tau_c)$, $\gamma_{c1} = \frac{\rho+\delta+\psi_c\mu}{(\rho+\delta+\phi_c)\kappa+(\rho+\delta+\psi_c\mu)}$ and $\gamma_{c2} = \frac{(\rho+\delta+\phi_c)\kappa}{[(\rho+\delta+\psi_c\mu)+\kappa(\rho+\delta+\phi_c)]}\frac{\psi_c(1-\mu)}{(\rho+\delta+\psi_c)}$. Note that these coefficients are implicitly functions of the employment rate and that $0 < \gamma_{c2} < 1$ as long as $0 < \mu < 1$. Equation (2.12) captures the notion that in a search and matching framework, sectoral wages act as strategic complements; that is, high wages in one sector are associated with high wages in other sectors. The strength of this strategic complementarity is captured by γ_{c2}. It is interesting to notice the effect of workers' bargaining power on the size of γ_{c2}. As can be verified, γ_{c2} is an increasing in κ, implying that sectoral wages are more strongly positively linked the weaker is workers' bargaining power. This suggests that even in environments where workers have minimal bargaining power, it is possible that γ_{c2} is high and that wages act as strong strategic complements. Notice also that equation (2.12) indicates why wages in one industry can vary across cities even when the productivity of the job is identical. In the case where $\mu = 1$ (i.e. workers are immobile across sectors), this effect disappears and wages are determined solely by value of marginal product. In such a case there would not be any effects of changes in industrial composition on within-sector wages. In this sense, the model nests more standard formulations where there are no general equilibrium effects. In contrast, with $\mu < 1$, a change in industrial composition can have an effect on within sector wages, even when the employment rate is unchanged.

If there is a pure industrial composition shift that causes a 1 unit increase in the average city wage, $\sum_j \eta_{jc} w_{jc}$, equation (2.12) indicates that within industry wages would then increase by γ_{c2}. But this is just a first round effect. Since the initial compositional change would affect all within industry wages, it would cause the average wage to increase by another γ_{c2} units, inducing a further round of adjustments. Hence, there are feedback dynamics which continue to multiply themselves out. The total effect of the change in industrial composition on the average wage would therefore be $\frac{1}{1-\gamma_{c2}}$.[35]

Equation (2.12) has the structure of the classic reflection or social interaction problem (Manski (1993), Moffitt (2001)) in that an outcome, here a sectoral wage, depends on the average of wages in a city. To make progress toward estimating such a relationship, it is necessary to overcome the simultaneity inherent to this type of interaction. To this end, we begin by rewriting equation (2.12) so that the sectoral wage is expressed only as a function of nationally determined prices and of the exogenous productivity terms ϵ:

$$w_{ic} = d_{ic} + \left(\frac{\gamma_{c2}}{1-\gamma_{c2}}\right)\gamma_{c1}\sum_j \eta_{jc}(P_j - P_1) + \gamma_{c1}\epsilon_{ic} + \gamma_{c1}\left(\frac{\gamma_{c2}}{1-\gamma_{c2}}\right)\sum_j \eta_{jc}\epsilon_{jc}, \qquad (2.13)$$

where $d_{ic} = \gamma_{c0}\left(1 + \frac{\gamma_{c2}}{1-\gamma_{c2}}\right) + \gamma_{c1}\left(\frac{\gamma_{c2}}{1-\gamma_{c2}}\right)P_1 + \gamma_{c1}P_i$. In equation (2.13), we have expressed prices in relation to the price of an arbitrarily chosen good denoted P_1 in order to help emphasize how a pure shift change in industrial composition affects wages (where by a pure shift we mean a change in the ηs that does not change the total number of jobs). From equation (2.13) we see that a city with an industrial composition biased toward goods with high value added will have higher wages within each sector. This relationship indicates that industry wage differentials ultimately arise from differences in value added (which themselves arise because of differences in the demand parameters, a_i, in equation (2.1)), interacting with increasing costs of creating jobs. If we had not set the costs of a vacancy r_i to zero, these would also be a determinant of wage premia.

We can now use equation (2.12) or (2.13) to relate the price differential $P_i - P_1$ to the average national level wage premium in industry i relative to the a baseline industry 1. This is achieved by noting that $w_{ic} - w_{1c}$ is equal to $\gamma_{c1}(P_i - P_1) + \gamma_{c1}(\epsilon_{ic} - \epsilon_{1c})$. If we now take the average of $w_{ic} - w_{1c}$ across cities, we obtain

$$w_i - w_1 = \gamma_1(P_i - P_1) + \hat{d}_i, \qquad (2.14)$$

where $w_i - w_1$ is the national level wage premium in industry i relative to to industry 1, γ_1

[35] In order to get an upper bound on this effect, one can consider the limit case where the bargaining power of workers goes to zero (κ goes to infinity) and ρ goes to zero. In this case, $\gamma_{c2} = \frac{\psi_c(1-\mu)}{(\delta+\psi_c)}$. Using the definition of ψ_c, this implies that $\frac{1}{1-\gamma_{c2}} = \frac{1}{1-ER_c(1-\mu)}$. If $1 - ER$ is assumed to be captured by the average unemployment rate, say 4%, and workers stay in the same industry 80% of the time, then $\frac{1}{1-\gamma_{c2}} = \frac{1}{1-.96*.80}$. This indicates that a change in industrial composition that would have a direct effect of 1 on the average wage due simply to the composition shift, would have a total effect of approximately 5 due to the strategic complementarity of wages. In this case, the general equilibrium effect would be four times the direct effect.

is the average of γ_{c1} across cities, and $\hat{d}_i$ is an industry specific constant.[36] Substituting (2.14) into (2.13), we get

$$w_{ic} = \tilde{d}_{ic} + \left(\frac{\gamma_{c2}}{1-\gamma_{c2}}\right)\left(\frac{\gamma_{c1}}{\gamma_1}\right)\sum_j \eta_{jc}(w_j - w_1) + \gamma_{c1}\epsilon_{ic} + \gamma_{c1}\left(\frac{\gamma_{c2}}{1-\gamma_{c2}}\right)\sum_j \eta_{jc}\epsilon_{jc}, \quad (2.15)$$

where $\tilde{d}_{ic} = d_{ic} + \left(\frac{\gamma_{c2}}{1-\gamma_{c2}}\right)\left(\frac{\gamma_{c1}}{\gamma_1}\right)\sum_j \eta_{jc}\hat{d}_{ic}$. Equation (2.15) provides a direct expression of how wages within an industry-city cell depend on the industrial composition of a city as captured by the index $\sum_j \eta_{jc}(w_i - w_1)$. We will denote this index by R_c and refer to it as average city rent. A high value of this index indicates that a city's employment is concentrated in high paying industries. Thus, the specific composition effect captured in (2.15) is one related to the proportion of good jobs in a city.

In interpreting the effect of R_c on wages from equation (2.15), we are performing a partial equilibrium exercise as we treat the employment rate in a city (as reflected in the γ parameters), and the sectoral composition, as given. In order to capture the dependence of wages on the city's employment rate more explicitly, it is useful to take a linear approximation of (2.15) around the point where cities have identical industrial composition ($\eta_{ic} = \eta_i$) and employment rates ($ER_c = ER$), which arises when $\epsilon_{ic} = 0$ and $\Omega_{ic} = 0$. For simplicity, we actually perform the expansion at the point, where the η_is are equal. Furthermore, to eliminate the city level fixed effects driven by the amenity, τ, we focus on the difference in wages within a city-industry cell across two steady state equilibria, denoted Δw_{ci}. This is given by equation (2.16):

$$\Delta w_{ic} = \Delta d_i + \left(\frac{\gamma_2}{1-\gamma_2}\right)\Delta\sum_j \eta_{jc}(w_i - w_1) + \gamma_{i5}\Delta ER_c + \Delta\xi_{ic}, \quad (2.16)$$

where Δd_i is an industry specific effect ($\Delta d_i = \gamma_1\frac{\gamma_2}{1-\gamma_2}\Delta P_1 + \gamma_1\Delta P_i$) that does not vary across cities, and hence can be captured in an empirical specification by including industry dummies, and $\Delta\xi_{ic} = \gamma_1\Delta\epsilon_{ic} + \gamma_1\frac{\gamma_2}{1-\gamma_2}\sum_j \frac{1}{I}\Delta\epsilon_{jc}$ is the error term, with I being the total number of industries. In (2.16), the γ coefficients are the same as presented after equation (2.12), except now they are evaluated at common match probabilities, ψ and ϕ. The added coefficient, γ_{i5}, reflects the effect of a change in the employment rate on wage determination; an effect which depends on all the parameters of the model. In particular, even in a first order approximation, this coefficient may vary across industries since the effects of a tighter labor market may affect the bargaining power of firms in an industry with a high value added product differently from the bargaining power of firms in an industry with a low value added product.

[36]The industry specific constant $\hat{d}_i$ is equal to $\sum_c(\gamma_{c1} - \gamma_1)(\epsilon_{ic} - \epsilon_{1c})$. When we later take a first order approximation around an equilibrium where the ϵs equal zero, this term will be equal to zero.

The coefficient that interests us most in equation (2.16) is the coefficient on $\Delta R_c = \Delta \sum_j \eta_{jc}(w_j - w_1)$. If we could estimate this coefficient consistently, we would obtain an estimate of the extent of (structural) city-level strategic complementarity between wages in different sectors. This would be done by inferring the value of γ_2 from the estimated effect of changes in R_c. However, the coefficient $\frac{\gamma_2}{1-\gamma_2}$ is of interest in its own right as it provides an estimate of the total - direct and feedback - effect of a change of industrial composition on within sector wages, as opposed to γ_2, which provides the partial uni-directional effect. Thus, if we examine wages in the same industry (say the textile industry) in different cities, a positive value for $\frac{\gamma_2}{1-\gamma_2}$ implies that those wages will be higher in cities where employment is more heavily weighted toward what we call high rent industries, where high rent industries are defined in terms of national level wage premia. This arises in the model because the workers in that industry have a better outside option to use when bargaining with firms in a higher average rent city. It is worth noting once again that this average rent effect is estimated conditioning on the employment rate. Thus, it is concerned with the composition of employment rather than the level of labor demand.

The coefficient $\frac{\gamma_2}{1-\gamma_2}$ relates directly to the question raised in the introduction concerning the extent to which the simple "composition adjustment" accounting procedure commonly used in the literature to calculate the effects of industrial composition on wages is appropriate. The coefficient $\frac{\gamma_2}{1-\gamma_2}$ provides an answer to this by quantifying the general equilibrium effect as multiples of the accounting procedure. To see this, note that the accounting procedure would calculate the effect on the average wage within a city of a change in industrial composition as being given by the change in R_c.[37] Given

[37] To see this more clearly, note that the standard accounting approach involves first estimating a wage equation that controls for individual characteristics such as education, experience and industry. Then one recovers the estimates on industry dummies, denoted ν_{it}. These estimated industry coefficients correspond to the inter-industry wage differentials. Then, the accounting approach consists of computing the term

$$A_{ct} = \sum_i \nu_{it}[\eta_{ict+1} - \eta_{ict}],$$

which shows how the average wage in a city changes with the change in industrial composition when using a given set of industry wage premia. If a change in the local industrial composition of employment does not alter within industry wages then A_{ct} is a reasonable way to measure the impact of industrial change on the average wage. However, if wages are determined instead according to Equation (2.16), then we must include the impact of changes in local composition on local wages and the total effect of a change in industrial composition on the average wage in a city becomes:

$$\sum_i \nu_{it}[\eta_{ict+1} - \eta_{ict}] + \frac{\gamma_2}{1-\gamma_2}(R_{ct+1} - R_{ct}) = A_{ct} + \frac{\gamma_2}{1-\gamma_2}(R_{ct+1} - R_{ct}).$$

Interestingly, the change in R_{ct} is very closely related to A_{ct} given that the change in R_{ct} can itself be written as

$$R_{ct+1} - R_{ct} = A_{ct} + \sum_i (\nu_{it+1} - \nu_{it})\eta_{ict+1}.$$

this, the total effect of a composition change is the accounting measures plus $\frac{\gamma_2}{1-\gamma_2}$ times that measure (or, $\frac{1}{1-\gamma_2}$ times that measure). If $\frac{\gamma_2}{1-\gamma_2}$ is estimated to be zero then the accounting procedure completely captures the effects of the composition shift. In terms of the model, this would fit with $\mu = 1$, a scenario in which there are no general equilibrium effects. If, instead, $\frac{\gamma_2}{1-\gamma_2}$ is not zero, it indicates by what percentage the accounting measure needs to be multiplied to capture the full effect of a change in industrial composition on the average wage.

Finally, as always, the potential to identify model parameters depends on the properties of the error term in the key estimating equation, which in our case is equation (2.16). Given that this error term is $\Delta\xi_{ic} = \gamma_1\Delta\epsilon_{ic} + \gamma_1\frac{\gamma_2}{1-\gamma_2}\sum_j \frac{1}{I}\Delta\epsilon_{jc}$, this effectively reduces to properties of ϵ_{ic}. To discuss identification issues, it is helpful to express ϵ_{ic} (and Ω_{ic}, as well) as the sum of a common component, which reflects absolute advantage, and a second component which captures relative advantage. For example, let $\hat{\epsilon}_{ct}$ represent the common component of the ϵs and let v^{ϵ}_{ict} represent the relative advantage component, with $\epsilon_{ict} \equiv \hat{\epsilon}_{ct} + v^{\epsilon}_{ict}$, where by definition the v^{ϵ}_{ict}s across industries within a city sum to zero. As we shall see, identification of the parameters of equation (2.16) will depend primarily on properties of the absolute advantage component $\hat{\epsilon}_{ct}$. A rather minimal assumption that can be placed on the common component $\hat{\epsilon}_{ct}$ is that it behaves as a random walk; that is, that the increments in $\hat{\epsilon}_{ct}$ are independent of the past. Since this appears to us as a reasonable assumption, we will examine whether it is possible to consistently estimate equation (2.16) under this assumption, recognizing that industrial structure and employment are endogenous. In particular, we will devise two instrumental variable strategies which will allow us to estimate (2.16) consistently under the sole assumption that the city-level common component of the ϵs behave like a random walk. Since this assumption is nevertheless questionable, we will also devise a test of whether this condition holds. A more stringent, but still possible condition on the ϵs is that the absolute advantage term be independent of all relative advantage terms (present and past). This implies that whatever drives general city performance is not related to a particular pattern of industrial structure. While this assumption is clearly more demanding, we will show that it is sufficient for the OLS estimation of equation (2.16) to give consistent estimates. By comparing OLS and IV estimates we will therefore be able to evaluate whether the data are supportive or not of this more demanding assumption.

Hence the average wage change in a city can be expressed as

$$\left(1 + \frac{\gamma_2}{1-\gamma_2}\right) A_{ct} + \frac{\gamma_2}{1-\gamma_2}\sum_i (\nu_{it+1} - \nu_{it})\eta_{ict+1}.$$

Thus, in the case where ν is not varying over time, the change in R_{ct} is exactly the amount the accounting approach attributes to changes in industrial composition and the complete effect of a composition shift is one plus the total estimated general equilibrium effect times the simple accounting measure.

2.1.2 Endogeneity of Industrial Composition

Although our use of economy wide wage premia in (2.15) and (2.16) provides a means of resolving the reflection problem inherent to local interactions, this does not resolve all endogeneity issues. Let us begin by focusing on the potential endogeneity of ΔR_c, acting for now as if there were no endogeneity issue related to ΔER_c. The issue that arises in this case is that changes in R_c may be correlated with the error term in (2.16) due to the fact that industrial composition, as captured by η_{ic}, is correlated with the ϵs. To see this potential endogeneity issue, we can use equation (2.7) to get the following expression for η_{ic}:

$$\eta_{i'c} = \frac{(\Upsilon_i + \Omega_{i'c})[p_i - d_{ci'} + (1-\gamma_{c1})\epsilon_{i'c} + (\frac{\gamma_{c2}}{1-\gamma_{c2}})\gamma_{c1}(\sum_j \eta_{jc}(Pj - P_1) - \sum_j \eta_{jc}\epsilon_{jc})]}{\sum_i (\Upsilon_i + \Omega_{ic})[p_i - d_{ci} + (1-\gamma_{c1})\epsilon_{ic} + (\frac{\gamma_{c2}}{1-\gamma_{c2}})\gamma_{c1}(\sum_j \eta_{jc}(Pj - P_1) - \sum_j \eta_{jc}\epsilon_{jc})]}. \tag{2.17}$$

We can express (2.17) more simply by taking a linear approximation:[38]

$$\eta_{ic} \approx \frac{1}{I} + \pi_1(\epsilon_{ic} - \frac{1}{I}\sum_j \epsilon_{jc}) + \pi_2(P_i\Omega_{ic} - \frac{1}{I}\sum_j P_i\Omega_{ic}), \tag{2.18}$$

where the πs represent positive coefficients obtained by the linear approximation.

The requirement for OLS to give consistent estimates of the coefficients in (2.16) can be expressed as follows:

$$\lim_{C,I\to\infty} \frac{1}{I}\frac{1}{C}\sum_{i=1}^{I}\sum_{c=1}^{C} \Delta R_c \Delta\xi_{ic} = \lim_{C,I\to\infty} \frac{1}{I}\frac{1}{C}\sum_{c=1}^{C} \Delta R_c \sum_{i=1}^{I} \Delta\xi_{ic} = 0, \tag{2.19}$$

where $\Delta R_c = \Delta \sum_j \eta_{jc}(w_j - w_1)$ and $\Delta\xi_{ic} = \gamma_1 \Delta\epsilon_{ic} + \gamma_1 \frac{\gamma_2}{1-\gamma_2}\sum_j \frac{1}{I}\Delta\epsilon_{jc}$, the error term in (2.16).[39] In Appendix A, we show that this condition is met under the assumption stated at the end of the previous section; that the absolute advantage components of the ϵs in all periods are independent of the relative advantage components of the ϵs and of the Ωs. This can be seen from Equation (2.19) since the city specific dimension of the η_{ic}s depends on relative advantage components, while $\sum_{i=1}^{I}\Delta\xi_{ic}$ depends only on the absolute advantage component of the ϵs. If this condition is met, one need not worry about any endogeneity issue that may involve the wage premia $(w_i - w_1)$ since the later do not vary with c. In intuitive terms, there must be no link between whether a city

[38] In deriving (2.17), we are again taking the linear approximation around the point where $\epsilon_{ic} = 0$ and $\Omega_{ic} = 0$, which implies that the linear approximation is taken around a point where the employment shares and the employment rate are identical across cities. For simplicity, we again assume the common shares equal $\frac{1}{I}$.

[39] Where appropriate, the symbol $\lim$ should be interpreted as a probability limit.

has experienced a shift in industrial composition toward high paying industries (a shift in ΔR) and other, general improvements affecting wages in all industries in the city. If this were not true, then the estimated effect of ΔR would partly reflect changes in the general conditions in a city rather than just capturing the impact of compositional shifts.

While we believe that the assumption ensuring consistency of the OLS estimates is possible, we are more comfortable with the much weaker assumption that the city-level common component of the ϵ has increments independent of the past. For this reason we will pursue two different instrumental variable strategies to address potential correlation between ΔR_c and the error term in (2.16) that are valid under this weaker assumption. The idea behind one of our instrumental variables for $\Delta \sum \eta_{ic}(w_i - w_1)$ is to use national level observations on changes in industrial employment shares to predict city level industrial shares. This approach can be seen as using elements contained in the second term in (2.18), that is, in $\Delta(P_i\Omega_{ic} - \frac{1}{I}\sum_j P_i\Omega_{ic})$, to predict changes in η_{ic}. In Appendix A, we show that this instrument and our alternate instrument provide consistent estimates under the condition that $\Delta\hat{\epsilon}_c$ is independent of v^{ϵ}_{ict} and Ω_{ict}; i.e., that changes in the absolute advantage of a city are independent of the past level of comparative advantage for the various industries in a city. This is a weaker assumption than that required for consistency under OLS because it allows for the possibility that a city that shifts in the direction of a higher wage industrial composition is also a city experiencing improvements in general. We just require that those general improvements must be independent of past comparative advantage. While this assumption can also be placed in question, our use of two instrumental variable strategies will allow to test the identification strategy. We leave the details of this discussion and the construction of our instrumental variables for ΔR_c until we get to the empirical section.

The last endogeneity issue relates to the potential correlation between the change in the employment rate, ΔER_c, and the error term in (2.16). To see this possibility, it is helpful to to use a linear approximation of (2.7) and (2.8) to obtain an expression for ER_c in terms of the ϵs and the Ωs. This is given in Equation (2.20), where the coefficients $\tilde{\pi}$ are again positive and obtained from the linear approximation:

$$ER_c \approx \tilde{\pi}_0 + \tilde{\pi}_1 \sum_j \epsilon_{jc} + \tilde{\pi}_2 \sum_j P_j\Omega_{ic}. \tag{2.20}$$

From (2.20) one can see that ΔER_c is likely correlated with the error term in (2.16) and, as a result, that OLS will generate inconsistent estimates. We will address this with an instrumental variable strategy in which we again use national level information on growth patterns to predict city level changes in employment rates. The approach is similar to that used by Blanchard and Katz (1992) in a closely related problem. Details of the construction of this instrument are again left to the empirical section.

2.1.3 Worker Mobility

In the model as presented so far, we have assumed that workers are not mobile across cities. At first blush, it may appear that the result that wages can vary systematically across cities due to the composition of employment will disappear once we allow for mobility of workers. However, this will not generally be the case even if we allow for directed search across cities.[40] To see this, consider allowing unemployed individuals to occasionally have the option of changing cities. In this case, an individual who lost his job in industry i would choose the city c' where U^u_{ic} is highest, that is, this individual would choose his location by solving $\max_{c'} U^u_{ic'}$. If we let μ_1 represent the probability that an unemployed individual has the option to change cities, then the value function associated with being unemployed from industry i in city c will need to be modified as to satisfy the following relationship:

$$\rho U^u_{ic} = b + \tau_c + \psi_c(1-\mu_1)(\mu U^e_{ic} + (1-\mu)\sum_j \eta_{jc} U^e_{jc} - U^u_{ic}) + \mu_1(\max_{c'} U^u_{ic'} - U^u_{ic}). \tag{2.21}$$

There are two reasons that this alternative Bellman equation does not greatly change our previous analysis. First, since $\max_{c'} U^u_{ic'}$ does not depend on initial location of the worker (that is, it does not depend on c), it can be treated as a common element across cities and therefore can be captured by industry specific dummies. Second, in equilibrium the $U^u_{ic'}$s are equalized across cities and therefore the term $\max_{c'}(U^u_{ic'} - U^u_{ic})$ in (2.21) will equal zero. Hence, in equilibrium, equations (2.5) and (2.21) will be identical, which implies that equations (2.12)-(2.16) remain valid with worker mobility.

When discussing mobility across cities, it would be appropriate to allow house costs (or land prices) into the equilibrating mechanism so as to have individuals choose a city that maximizes their expected utility taking into account housing prices and local amenities. To allow this extension, assume that workers care about wages, the price of housing in a city, p^h_{ct}, and about a local amenity, Ψ_c. In this case, a worker's (indirect) flow utility when employed in industry i in city c could be expressed as, $w_{ic} - \varsigma p^h_{ct} + \Psi_c$. Accordingly, his or her flow utility when unemployed will be given by $b + \tau_c - \varsigma p^h_{ct} + \Psi_c$. Note that in such a setup, wage negotiation will not be directly affected by housing prices since it is a cost that is incurred whether or not someone is employed. However, housing prices will need to adjust to equilibrate expected utility across cities. To capture this role, one can summarize the functioning of the housing market by assuming that housing prices can be expressed as a positive function of the population of a city and of

[40] As noted in footnote 7, it is easy to verify that allowing for random search across cities does not significantly change our previous analysis.

amenities, such that

$$p^h_{ct} = d_0 + d_1 L_{ct} + d_2 \Psi_c.$$

It is straightforward to show that in this world, a city with a higher employment rate and a better employment mix, as captured by a higher value of R_{ct}, will attract more workers. This immigration will drive up local housing prices, causing the migration to stop before wages are equalized across cities. Housing prices will adjust such that a city with a favorable composition of jobs (due to favorable ϵs and Ωs) has local benefits that are captured by local landowners. Thus, labor mobility implies that the determinants of wages emphasized in equation (2.16) should also be determinants of housing prices and inter-city migration flows. We will investigate, toward the end of the paper, whether changes in R_{ct} have these predicted impacts on inter-city migration flows and on housing prices. Nonetheless, the impact of a favorable job composition will still be seen in wages, as implied by equations (2.12)-(2.16) and, thus, worker mobility does not change the predictions of the model in terms of the complementarity of wages across sectors and the implied general equilibrium effects of industrial composition.

2.1.4 Worker Heterogeneity

Up to now, labor in our model has been treated as homogenous and not differentiated by skill. Since workers in reality do differ in skills (and are treated differently by employers as a result), we need to take this into account when exploring the relevance of the general equilibrium effects predicted by the model. In particular, our model should be seen as applying to labor markets defined by skills. Empirically, we can approach the issue of ensuring we are only estimating within skill-group implications of the model in two ways. The first is to work with sub-samples of the data intended to be homogeneous with respect to skill (e.g., using only young high school graduates). The other is to eliminate skill differentials by first controlling for a flexible set of "skill" indicators in the determination of wages (where a wage regression would now be seen as an individual level wage regression rather than one specified at the industry $\times$ city level as it has been to this point). This latter approach is justified if the setup costs and benefits associated with employing a more skilled worker are proportional to the increased productivity of the worker and if all the parameters of the model are identical across skill groups. To implement this approach, we would also need to control for the same set of skill indicators when calculating the national level wage premia used in constructing the R_{ct} measure. We pursue both approaches in the empirical section.

2.2 Empirical Implementation

Our baseline empirical specification is given by Equation (2.22), which is a simple rewrite of equation (2.16) where we have divided both sides by w_1 in order to focus on a log specification and where we have added a time subscript since we will pull data from different periods:

$$\Delta \log w_{ict} = \alpha_{1i} d_{it} + \alpha_2 \Delta R_{ct} + \alpha_{3i} \Delta ER_{ct} + \Delta \xi_{ict}. \tag{2.22}$$

In (2.22) the d_{it}s are time varying industry dummies, $\alpha_2 = \frac{\gamma_2}{1-\gamma_2}$ is our main coefficient of interest, $R_{ct} = \sum_j \eta_{jct}(\frac{w_{jt}}{w_1} - 1)$ is our index of industrial composition, α_{3i} are the coefficients capturing the effect of city level employment rates on wages (where these effects are allowed to vary across industries), and $\Delta\xi_{ict}$ is the error term defined by $\Delta\xi_{ict} = \frac{\gamma_1}{w_1}\Delta\epsilon_{ict} + \frac{\gamma_1}{w_1}\frac{\gamma_2}{1-\gamma_2}\sum_j \frac{1}{I}\Delta\epsilon_{jct}$.

Our goal is to investigate the null hypothesis that α_2 = 0 or, in other words, whether disregarding inter-sectoral wage interactions provides an appropriate description of wage determination in local economies. Support for this null hypothesis would indicate that the standard accounting procedure completely captures the effect of a local change in the composition of employment on the local average wage. Our alternative hypothesis is that $\alpha_2 > 0$. A finding of $\alpha_2 > 0$ would indicate both the presence of a significant general equilibrium effect of industrial composition on wages as predicted by the model and that the standard accounting approach is an inappropriate means of evaluating the effects of changes in industrial composition on average wages.

When estimating the effect of R_{ct} on wages, it is appropriate to worry about omitted variable bias, especially given existing alternative explanations for differences in wages across cities such as those related to city size, education levels (Moretti (2004), Acemoglu and Angrist (1999)), and diversity of employment in a city (Glaeser, Kallal, Scheinkman, and Shleifer 1992). To allow for such issues, we will add to equation (2.22) measures related to these explanations as additional covariates, Z_{ct}. As we suggested previously, we also address the potential issues of endogeneity of R_{ct} and ER_{ct}. In addition, we will address the potential for worker mobility to cause a selection bias.

2.3 Data and Basic Results

2.3.1 Data

The data we use in the following investigations come from the 1970, 1980, 1990 and 2000 U.S. Census Public Use Micro-Samples (PUMS). We focus on wage and salary earners, aged 20 to 65 with positive weekly wages who were living in a metropolitan area at the time of the Census. To form our dependent variable we use the log of weekly wages, calculated by dividing wage and salary income by annual weeks worked (we also verified the robustness of our results to using hourly wages). We create real wages (in 1990 dollars) using the national level CPI as the deflator. Given our use of multiple Censuses, an important part of our data construction is the creation of consistent definitions of cities, education groups and industries over time. We provide the details on how we address these issues in Appendix B.

As we described in the previous section, one approach to addressing worker heterogeneity is to control for observable skills in a regression context. Our actual approach is to use a common two-stage procedure. In the first stage, we run individual level regressions of log wages using all the individuals in our national sample on categorical education variables (4 categories), a quadratic in experience, interactions of the experience and education variables, a gender dummy, black, hispanic and immigrant dummy variables, and the complete set of interactions of the gender, race and immigrant dummies with all the education and experience variables. We run these regressions separately by Census year to allow for changes in returns to skills over time. The regressions also include a full set of industry-by-city cell dummies and it is the coefficients on those that are used to construct the dependent variable in the second stage regression (equation (2.22) above). We eliminate all industry-city cells with fewer than 20 included individuals in any of the years. We use the square root of the number of observations in each industry-city cell to form weights for the second stage estimation. For most of our estimates, we use decadal differences within industry-city cells for each pair of decades in our data (1980- 1970, 1990-1980, 2000-1990), pooling these together into one large dataset and including period specific industry dummies. In all the estimation results we calculate standard errors allowing for clustering by city and year.

The main covariate in our estimation is the ΔR_{ct} variable which is a function of the national industrial wage premia and the proportion of workers in each industry in a city. We estimate the wage premia in a regression at the national level in which we control for the same set of education, experience, gender, race and immigration variables described for our first stage wage regression and also include a full set of industry dummy variables. This regression is estimated separately for each Census year. The coefficients on the industry dummy variables are what we use as the wage premia in constructing our R measures.

2.3.2 OLS Results

We begin our presentation of results with the estimates of (2.22) without the inclusion of any additional control (Z_{ct}) variables. The first two columns of Table (3.1) contain the results from OLS estimation of the regression. These two regressions and all of those that follow include a full set of time varying industry dummy variables (3×144), thus allowing for changes in industry premia over time, but we do not present the long list of corresponding coefficients here. In column (1) we restrict the coefficient on the employment rate to be equal across industries. In column (2) we allow for the employment rate to vary across large industry groups.[41] The coefficient on the change in R_{ct} variable is 2.6 in both specifications and is statistically significantly different from zero at any conventional significance level. If OLS provides consistent estimates of this coefficient, the fact that this coefficient is both economically substantial and statistically significant implies a rejection of the null hypothesis that the impact of changes in the composition of employment in a city is completely captured in the standard accounting measure. Further, the coefficient fits with the alternative hypothesis that cities with employment structures that shift toward higher premia industries have better wage performance within industries. Recall from our discussion of the definition of the R_{ct} variable that the magnitude of the coefficient on this variable can be interpreted as a multiple of the standard accounting effect. Thus, the OLS estimate implies that the total effect on average wages of a shift in composition toward higher paying industries is approximately three and a half times what is reflected in a standard accounting measure. This total effect may initially sound overly large but it is worth recalling that the accounting measure effects tend to be quite small. For example, consider the average real weekly wage for men with a BA or higher education (examples with other education or gender groups give similar results). For this group, the average wage increased by 8% across the cities in our sample between 1980 and 1990. If we recalculate the 1990 average wage for this group holding the industrial composition constant, the increase becomes 7%, implying that the accounting measure of the impact of shifts in industrial composition is 1%. Our estimates suggest that the total impact of shifts in industrial composition would be 3.5% in this example. Such an increase is certainly larger than what is usually attributed to industrial shifts but is still only just over 40% of the overall increase. The fact that direct accounting measures of the impact of industrial shifts tend to be small has led to a discounting of explanations for changes in the U.S. wage structure that might show up through such shifts. Trade, for example, is usually relegated to a lower place in the list of potential explanations for this reason. An estimate of the size we report may imply that there is reason to re-examine those types of explanations.

One point of interest about this result is whether it is being driven by a subset of

[41] We are allowing the effect of the employment rate to vary across 16 industry groups. We also explored the effect of allowing employment to vary with all 144 industries, and found similar results. Since in most of our specifications, we could not reject that 16 interaction were sufficient, we adopted this specifications as our baseline case.

cities, such as those that faced particularly large re-adjustment after the difficulties in key manufacturing industries. To examine this, in Figure (3.1) we plot the change in city average wages (the average of our dependent variable across industries within a city) against ΔR_{ct}.[42] The key point from this figure is that there is a strong positive relationship between wage changes and changes in our R measure that is not driven by outliers.

We are also interested in whether the estimated effect stems from some particular set of industries for which wages are particularly sensitive to the presence of high premia industries. We investigate this by re-estimating our basic specification interacting the ΔR variable with a complete set of industry dummy variables. This is equivalent to re-writing the α_2 coefficient with an i subscript. In Figure (3.2), we present a histogram of the full set of these α_{2i} coefficients. What is noteworthy in this figure is the concentration of values around the mean. The implication is that workers in virtually all industries benefit from a shift in employment composition toward high paying sectors and benefit to much the same degree. It is worth recalling when considering this result that we estimate industry premia while controlling for observable skills. Thus, high premium industries are not necessarily high skill industries.

It is interesting to note that our OLS estimates of α_2 in Table 1 are very stable whether or not we allow the employment rate effects to vary across industries (a pattern that also holds true when we move to estimating the relationship by instrumental variables). In fact, the OLS estimate of α_2 is almost identical if we don't include the change in the employment rate in the regression. The reason for this can be easily understood from Figure 3, in which we provide a plot of the change in the employment rate against ΔR_{ct}. The flat relationship in this figure indicates that these two variables are almost orthogonal to each other. This implies that most of the shifts in industrial composition look like pure compositional shifts, with little co-movement in a city's employment rate.

2.4 Addressing Endogeneity and Selection Issues

2.4.1 Endogeneity: Methods and Results

As we discussed earlier, OLS estimation of (2.22) will provide consistent estimates if changes in a city's absolute advantage are independent of patterns (levels and changes) of relative advantage. While this may be the case, in this section we explore estimation of (2.22) using two instrumental variable strategies which rely on weaker identification

[42] We actually first regress $\Delta \log w_{ict}$ on industry-year dummies and plot the weighted average of the residuals from that regression in order to obtain a plot that replicates our actual regression.

assumptions.

To help motivate our instrumental variable strategies, it is useful to start by writing out a standard decomposition of ΔR_{ct}:

$$\Delta R_{ct} = \sum_i \Delta\eta_{ict}\nu_{it} + \sum_i \eta_{ict+1}\Delta\nu_{it}, \tag{2.23}$$

where $\nu_{it} = \frac{(w_{it}-w_{1t})}{w_{1t}}$ are national level industrial premia. Thus, movements in R_{ct} are decomposed into a component due to changes in the composition of the workforce (in the η_{ict}s) holding the wage premia constant and a component due to changes in the premia holding the local composition of the workforce constant.

Our instrumenting strategy for ΔR consists of building two instruments that emphasize variation in different components of the above decomposition, and which by construction are not functions of the $\Delta\epsilon$s. The first is constructed using the following procedure. We first predict a level of employment for industry i in city c in period $t+1$ using the formula:

$$\hat{l}_{ict+1} = l_{ict}\left(\frac{l_{it+1}}{l_{it}}\right).$$

That is, we predict future employment in industry i in city c using the employment in that industry in period t multiplied by the growth rate for the industry at the national level. Using these predicted values, we construct a set of predicted industry specific employment shares, $\hat{\eta}_{ict} = \frac{\hat{l}_{ict}}{\sum_i \hat{l}_{ict}}$, for the city in period $t+1$ and form a measure given by:

$$IV1_{ct} = \sum_i \nu_{it}(\hat{\eta}_{ict+1} - \eta_{ict}). \tag{2.24}$$

This variable is closely related to the first term in the decomposition of the R measure given in (2.23). Thus, this instrument isolates the variation in ΔR that stems from changes in the employment composition, but instead of using actual employment share changes we use predicted changes based on national level changes, breaking the direct link between city level employment and wage changes. Essentially, $IV1$ focuses attention on the question, "what is the impact on local wages of a national level demand shift (stemming from, for example, trade or preference shocks) if that shift is distributed across cities according to start-of-period employment shares?" Recall that use of this type of variation is implied by the model, where shifts in national level demand ($\{a_i\}$) are translated into local shifts in employment shares because of local differences in comparative advantage that will be reflected in initial period employment shares. In terms of the econometrics, since $IV1$ is not a function of $\Delta\epsilon_{ict}$, it is not correlated with the error term in (2.22), under the less demanding identification assumption that the change in the common component of the ϵ_{ict} (the absolute advantage component) is not predictable based on past information on comparative advantage. This contrasts with the condition for consistency of OLS which requires that changes in the absolute advantage component are also not correlated with changes in comparative advantage. A test of equivalence of the IV and OLS estimates is, therefore, a test of the stronger assumptions required for consistency of the OLS estimator.

Our second instrument is designed to isolate the variation inherent in the second term in the decomposition, (2.23): the variation stemming from changes in wage premia over time, weighted by the importance of the relevant industry in the local economy. Thus, our second instrument is given by:

$$IV2_{ct} = \sum_i \hat{\eta}_{ict+1}(\Delta\nu_{it}).$$

This instrument may initially seem less natural, since the discussion to this point has been almost entirely couched in terms of shifts in the concentration of employment. However, if our theoretical explanation, which emphasizes bargaining, is correct then it should not matter whether the average premium available in the city declines because a high paying industry shuts down or because the premium paid in that industry declines. In either case, workers in other industries end up with a less valuable outside option. This would imply that we should get similar results using $IV1$ and $IV2$. Hence, examining whether the results obtained using these two alternative instruments are the same can provide a means of evaluating whether the outside option effect outlined in Section 2 is the likely mechanism at play. As with $IV1$, $IV2$ is uncorrelated with the error term in (2.22) under our main identifying assumption that the common component in the $\Delta\epsilon$s is independent of the past comparative advantage.

Both instruments perform well in the first stage estimation. The F-statistic from the test of the significance of $IV1$ in the first stage regression of ΔR on the instrument equals 15.8 and has an associated p-value of 0.0. The same statistic for $IV2$ is 20.31 with a p-value of 0.

A non-zero covariance between ΔER_c and $\Delta\xi_{ic}$ is to be expected since a city with a set of large increases in its local, industry specific cost shocks (its ϵ_{ic}s) will have higher employment. We respond to that endogeneity problem using an instrument that is similar to IV1.[43] In particular, we use as an instrument $\sum_i \eta_{ic} g_i$, where g_i is the growth rate of employment in industry i at the national level. Thus, the instrument is the weighted average of national level industrial employment growth rates, where the weights are the start of period industrial employment shares in the local economy. A city that has a

[43] This instrument is similar to that used in Blanchard and Katz (1992).

strong weight on an industry that turns out to grow well at the national level will have a high value for this instrument. Because the ϵ_{ic}s are local demand shocks that sum to zero across cities, their movements are not correlated with the g_is by construction. Finally, under the assumption that $\Delta\epsilon_{ic}$ is independent of ϵ_{ic}, the changes in ϵ_{ic} that constitute $\Delta\xi_{ic}$ will be independent of the ϵ_{ic}s used as weights in the instrument, resulting in a zero correlation between the instrument and the error term.

We present results from instrumental variables estimation using $IV1$ and $IV2$ individually, in the third and fourth columns of Table (3.1). In all the instrumental variable estimates we present, the employment rate is treated as endogenous and instrumented using $\sum_i \eta_{ic} g_i$. The third column contains results from instrumental variable estimation in which we use $IV1$ to counter the potential endoegneity of ΔR_{ct}. The estimated coefficient of ΔR_{ct} is very similar to that obtained from OLS estimation and is again highly statistically significant. The fourth column contains results when we use $IV2$, the instrument that uses changes in industry premia over time, and in column 5 we use both $IV1$ and $IV2$ simultaneously as instruments. All of these results, again, imply large and statistically significant effects of ΔR on local average wages.

The similarity in the estimated coefficients obtained using $IV1$ and $IV2$ is striking and points to two interesting implications. First, as we discussed earlier, this outcome is supportive of theories of the impact of R that are based on changes in the value of outside options in bargaining. The results in columns 3 and 4 of Table (3.1) indicate that both moves away from high paying jobs and reductions in the premia associated with high paying jobs have approximately the same impact on within industry wages; a result implied by our bargaining story. We can check this implication in a somewhat more transparent way by estimating our main specification, replacing ΔR with its two components entered separately. We present results from an OLS estimation of this adjusted specification in column 6 of the table. In column 7, we present results associated with estimating the effects of each component of ΔR_{ct} by instrumental variables, where we use $IV1$ and $IV2$ as instruments. In both the OLS and IV specifications, the two components of ΔR have coefficients that are close in size and not statistically significantly different from one another. In the case of OLS, the estimated effect of $\sum_i \Delta\eta_{ict}\nu_{it}$ is slightly lower than the estimated effect of $\sum_i \eta_{ict+1}\Delta\nu_{it}$. However, when estimating by IV the two coefficients as extremely close. Thus, as the theory suggests, variation in R stemming either from changes in industrial composition or from changes in industrial premia appear to have the same impact on local wages.

The second conclusion stemming from the similarity of the $IV1$ and $IV2$ estimates is related to its implications for the statistical assumptions underlying our consistency proofs. The consistency of estimates based on both $IV1$ and $IV2$ rely on the same assumption since the instruments are essentially different weighted averages of a city's initial employment shares and our key identifying assumption relates to those shares. Given that the weighting schemes are different, one would expect that if the identification assumption does not hold, these two IV strategies would produce different estimates since they would weight any departures from the assumption differently. This intuition can be formalized by considering a test of whether the estimates of α_2 obtained using $IV1$ and $IV2$ are similar. The idea of the test is similar to a Hausman test: under the null both are consistent and should provide similar estimates. The only issue is that the variance-covariance matrix for the differences in estimates based on the two instruments needs to be calculated differently from a standard Hausman test since neither IV estimator is efficient. We perform this test,[44] and not surprisingly given the similarity in the estimates in Column 3 and 4 in Table (3.1) or (3.2), it passes easily. We believe that the fact that the two IV approaches, which focus on very different data variation, give very similar results provides considerable support for the assumption that the common city level component in the ϵ acts like a random walk. Moreover, the fact that OLS also gives similar results suggests, further, that even the strong assumption needed for the consistency of OLS likely holds in these data.[45]

2.4.2 Selection: Methods and Results

A second key concern in estimating (2.22) is with selection of workers across cities. The R variable varies at the city level over time. Thus, changes in unobserved skills in a city that are correlated with movements of R will imply a non-zero coefficient on R that does not reflect general equilibrium effects of the type we are considering. For example, suppose that there are unobserved skills (which we will call ability) and that high premia industries can choose higher ability workers from lines of applicants. Suppose, further that the most able workers move out of a city if it loses a high paying industry, regardless of the industry in which they themselves are employed, because they want to live in a place where they have a chance of getting into a higher paying job. In that case, shifts in R may actually pick up the effects of shifts in the unobserved ability distribution.

We address selection concerns in a number of ways. First, we control for observable skill variables (education and experience) both when estimating the wage premia in the national level wage regression and when obtaining the industry-city average wages that form our key dependent variable. Our second approach is to implement the selection correction estimator that Dahl (2002) proposes and implements in his examination of

[44] We also performed a more standard over-identification test for column 5 and have not found any evidence against our maintained assumption.

[45] The conditions that allow OLS to provide consistent estimates of (2.22) can also be shown to imply that estimating the relationship in levels instead of in differences should give similar results. Interestingly, when we estimate (2.22) in level we do get similar results.

regional variation in returns to education.

To understand the nature of Dahl's approach, consider a model in which each worker has a (latent) wage value that he would earn if he lived in each possible city and chooses to live in the city in which his wage net of moving costs is highest. If we explicitly introduce individual heterogneity, this implies that we should write the regression corresponding to observed wages as

$$E(\log w_{kict}|d_{kct} = 1) = \alpha_{0t} + \beta_{1t}x_{kct} + \alpha_1 ER_{ct} + \alpha_2 R_{ct} + \nu_i + \nu_c + E(e_{kct}|d_{kct} = 1), \quad \textbf{(2.25)}$$

where k indexes individuals and d_{kct} is a dummy variable equaling one if worker k is observed in city c at time t. The last, error mean, term is non-zero if worker city selection is not independent of the unobserved component of wages. If one were to estimate equation (2.22) not taking account of this error mean term then the estimated regression coefficients will suffer from well-known consistency problems.

In situations such as in the union wage premium literature where there are only two options facing a worker, it is well known that the error mean term can be expressed as a function of the probability of selecting the given option (Heckman 1979, Lee 1983). In our case, with multiple possible destinations to choose from, the error mean term will potentially be a function of characteristics of all of them, making estimation complicated. Dahl (2002) argues that under specific sufficiency conditions, the error mean term is only a function of the probability that a person born in the same state as k would make the choice that k actually made, greatly simplifying the problem. In his examination of the impact of selection of location across states on returns to education, however, he argues that this sufficiency assumption is overly restrictive and that one can effectively account for selection using functions of the probability k did not move from his state of birth and the probability he moved to the state in which he is observed at the time of the Census. Following work such as Ahn and Powell (1993) and Heckman and Robb (1985) for the binary choice case, he also proposes a non- parametric estimator for the relevant probabilities and the function of them that enters the regression of interest. We follow his approach with a few adjustments to account for the facts that we include immigrants in our analysis and that we are dealing with cities. Details on our selection estimation are provided in Appendix C. In essence, this estimator identifies the error mean (selection) effect using differences in the probabilities of being observed in a given city between two people who are identical in education, experience, race and gender but are born in different states. The idea is that, for example, people born in Oregon are more likely to be observed in Seattle than people born in Pennsylvania because Oregon is so much closer. If both are in fact observed living in Seattle then we are assuming that the person from Pennsylvania must have a larger Seattle specific "ability" (a stronger earnings related reason for being there) and this is what is being captured when we include functions of the relevant probabilities of being observed in Seattle for each of them. Identification in this approach is based on the exclusion of state of birth by current city of residence interactions from the wage regression. That is, we assume that being born in a state close to your city of residence (or, more generally, a state with a high associated probability of moving to that city) does not directly determine the wage a worker receives.[46]

In practical terms, this approach to the potential selection problem again involves two estimation steps. In the first, as before, we estimate individual level regressions of log wages on the same complete set of education and experience variables, indicators for race, immigrant status, and gender, as well as a full set of city-by-industry dummies but now also add our proxies for the error mean term. We again retain the coefficients on the city-industry dummy variables and then proceed with the second stage regressions as before. The coefficients on the error mean proxy variables are jointly highly significant in the first stage regressions, implying that there are significant sample selection issues being addressed with this estimator.

In Table (3.2), we recreate the results from Table (3.1) while implementing Dahl (2002)'s selection correction. The resulting estimates for α_2 are very similar in magnitude to those obtained when we did not correct for sample selection. Thus, we do not believe that movements in unobserved ability across cities is strongly contaminating our estimates. The implication of the selection analysis is that while workers do select themselves across cities in a manner that is non-random with respect to earnings outcomes, changes in their selection pattern are not correlated with changes in the average industrial premium paid in a city.

2.4.3 Observing Strategic Complementarity

The results from the previous section suggest that changes in industrial composition toward higher paying jobs have the effect of increasing wages across all sectors. The key mechanism that leads to the amplification of an initial industrial composition change on wages in our model is that wages across sectors play the role of strategic complements. This effect was expressed most clearly in Equation (2.12). As previously emphasized, Equation (2.12) implies that a change in industrial composition that initially increases average wages in a city by 1%, would lead to a cumulative increase in average wages by a factor of $(\frac{1}{1-\gamma_2})\%$ due to the strategic complementarity of wages. Given that α_2 in Equation (2.22) relates to γ_2 according to $\alpha_2 = \frac{\gamma_2}{1-\gamma_2}$, our estimates of α_2 that are around

[46]Note that this is different from assuming that state of birth does not affect current wages since, even if we include a set of state of birth dummy variables in our first stage estimation, our approach remains identified off interactions between city-of-employment and state-of-birth.

2.5 in Table (3.1) suggest that γ_2 should be in the range of $0.71 = (\frac{2.5}{3.5})$. This implication of the model can also be examined directly by estimating Equation (2.12) using an IV estimator. Since the coefficients in Equation (2.12) depend on the employment rate, we proceed as with Equation (2.15) to take a linear approximation of (2.12) around the point where the ϵs and Ωs are zero and then taking first differences, to get a linear equation of the form

$$\Delta \ln w_{ict} = \psi_1 d_{it} + \psi_2 \Delta \sum_j \eta_{jct} \ln w_{jct} + \psi_{3i} \Delta ER_{ct} + \tilde{U}_{ict}, \tag{2.26}$$

where ψ_2 corresponds to γ_2 in the model, and the error term corresponds to $\gamma_1 \Delta\epsilon_{ict}$. In the case of Equation (2.26), estimation by OLS would definitely be expected to give upward biased estimates of $\psi_2 = \gamma_2$ since the relationship suffers from the reflection problem. However, it can be verified that instrumental variable estimation of Equation (2.26) using our previous set of instruments should give consistent estimates under the same assumption as before, that is, that the common component of the ϵs (a city's absolute advantage) is independent of the past. It is worth emphasizing that the difference between Equation (2.22) and (2.26) pertains only to the main variable of interest. In (2.26) this variable is the average city wage,[47] while in (2.22) it is a city level average of national wage premia.

We present estimates of Equation (2.26) that parallel Tables (3.1) and (3.2) in Tables (3.3) and (3.4). In all cases, we control for industry specific effects of the employment rate and include a full set of time varying industry dummies as before; the difference between Tables (3.3) and (3.4) being that in Table (3.4) we implement the selection correction. The first thing to note from these two tables is that, as one should expect, there is now a large and significant difference between estimates of ψ_2 (denoted Δ Av. City Wage in the tables) obtained by OLS versus IV. The OLS estimate is .853, which if translated to compare with α_2 would imply an $\alpha_2 = 5.80 = \frac{.853}{1-.853}$. However, in this case there are no conditions under which we should expect OLS to give consistent estimates. In contrast, when we estimate by IV, we get an estimate of ψ_2 equal approximately to .69, which implies a value for $\alpha_2 = 2.22 = \frac{.69}{1-.69}$, which is very close to that obtained Tables (3.1) and (3.2) using different approaches. In particular, the estimation of (2.22) by OLS provides one means of overcoming the reflection problem by focusing on national level wage premia, while the IV estimation of (2.26) provides a conceptually quite different approach. In the remaining sections of the paper, we focus on verifying the robustness of the results presented in Table (3.1) and (3.2). In order to save space, we do not provide further results on the estimation of (2.26), but we have verified that they are as robust

[47] Calculated as the city average of the regression adjusted city-industry wages that form our dependent variable.

as those for Equation (2.22). Also, in all subsequent sections of the paper, we present results incorporating Dahl's sample selection correction.

2.5 Further Explorations of the Wage Premia Effects

2.5.1 Other Driving Forces for City Level Wage Changes

Ours is certainly not the first attempt to examine the determinants of city level wage changes and/or city-level growth. The literature on what makes for a high performing city has produced a number of hypotheses. In this section, we introduce measures corresponding to some of the more prominent hypotheses to see whether our R_{ct} measure may be spuriously capturing one of these alternative driving forces.

One prominent explanation for city level growth is provided in Glaeser *et al.* (1992). They examine city level growth over time in the U.S., comparing the impact of measures of city size, which would be important determinants of growth if agglomeration type models were driving growth patterns, and measures of the industrial diversity of the economy. They argue that the importance of diversity is implied by, for example, Jane Jacob's theorizing. They find that industrial diversity is a stronger determinant of city specific growth than city size. In response to this view, in the first three columns of Table (3.5), we introduce a measure of the "fractionalization" of employment in a city at the start of each decade. The measure of fractionalization we use is one minus the Herfindahl index, or one minus the sum of squared industry shares. This measure itself tends not to be very significant in our estimates and, more importantly, does not change our estimates of the α_2 coefficient.

Another possibility relates to the recent literature on education externalities which examines the claim that having a larger proportion of workers in a city being highly educated benefits all workers in the city. Moretti (2004), for example, in an examination of wages in U.S. cities in the 1980s finds that cities with a greater increase in the proportion of workers with a BA or higher education have higher wage gains. Acemoglu and Angrist (1999) find weaker results for the impact of education using average years of education in a state. Again, we are interested in whether our R_{ct} measure is actually picking up this alternative effect. It is worth re-emphasizing, though, that we control for education in the regressions from which we estimate our national level wage premia and, thus, the R_{ct} measure does not reflect cities that have high wages because they have high levels of education. In the middle set of three columns in Table (3.5), we introduce the change in the proportion of workers with a BA or higher education (the College Share) as an additional regressor. The college share variable itself enters significantly, supporting Moretti (2004)'s findings, but introducing this variable has very little impact on our estimates of the effect of changes in R_{ct}.[48] Although not reported, we also examined the the effect of using average years of education as an alternative measure of the education level of a city. This latter variable does not enter significantly, supporting results in Acemoglu and Angrist (1999) and fitting with the often contradictory results in this literature. Moreover, including average years of education does not affect the estimates of our R_{ct} effects.

Finally, in columns (7), (8) and (9), we introduce the change in the log of the size of the city's labor force. This city size variable is intended to capture the type of agglomeration effects tested in, for example, Glaeser *et al.* (1992). It can also be viewed as a direct control for overall demand effects, allowing us to check whether the effects of shifts in industrial composition we are measuring are truly pure compositional effects or whether they are partly capturing shifts in overall demand that may accompany the loss or gain of major industries. In the IV columns, we instrument for the labor force variable using the same type of instrument we have used in various ways to this point; that is, one based on predicting labor force growth in a city from national level growth for each industry combined with the initial industrial composition in the city. Whether instrumented or not, this variable has small and statistically insignificant effects and, more importantly, its inclusion does not alter the estimates of the ΔR effect. Again, this fits with our interpretation of α_2 as capturing a pure compositional effect. Moreover, our finding that predicted changes in labor force across cities has little effect on wages is consistent with Blanchard and Katz's [1992] finding that local wages react little to change in demand or migration. [49]

Overall, our conclusion is that while some of the other hypothesized factors we have considered may affect city level wage growth, we are not inadvertently picking any of them up with our R_{ct} measure. Moreover, the impact of the shift in industrial composition toward higher paying industries is much stronger than any of the effects from these

[48]It is worth noting, though, that Sand (2006) finds that this positive and significant impact is observed in the 1980s but not in the 1970s or 1990s when estimation is carried out separately by decade.

[49]In order to further explore the potential for our estimates to be capturing aggregate demand effects instead of industrial composition effects, we divided our cities into four groups of equal size based on how national level movements in employment predict city outcomes in terms of increase in labor demand and changes in R_{ct}. The first group is the set of cities with predicted increases in R_c and predicted increases in employment (where increases and decreases represent differences with mean behavior by year). The second group represents the set of cities with predicted increases in employment, but decreases in R_c, which arises if the city is predicted to increase by expanding the low wage sector. The third group are cities with predicted decreases in both R_c and in employment. Finally, group 4 are cities with predicted decreases in employment and increases in R_c. Not surprisingly, we observe that cities in group 1 experience higher growth in wages that the overall mean growth, while cities in group 3 experience below average growth. The interesting cases are group 2 and 4. Here we observed that within-industry wages performed below the mean for cities in group 2 even though demand is expanding. In contrast, we found that within-industry wages increase above the mean in group 4 even if demand was declining. Interestingly, this is precisely the pattern implied by our model whereby it is the change in R_c which is the dominant determinant of local wages, not local demand.

competing explanations.

2.5.2 Robustness Checks

One possible explanation for the patterns we observe is that while standard trade forces are affecting wages in tradeable goods sectors, wages associated with skills that are used in the non-tradeable sector are moving for standard demand-induced reasons. Thus, a shift in employment in a city toward having more workers with high levels of unobserved skills (perhaps because of their pursuit of local amenities) could lead to an increase in R_{ct} that would not affect wages in the tradeable sector for standard factor price equalization reasons but, because the higher skilled workers have more income to spend on locally produced non- traded goods, it could affect wages in the non-traded sector. Under this explanation, we should see smaller impacts of changes in R_{ct} on wages in tradeable sector industries than on wages in non-tradeable sector industries.

To define tradeable and non-tradeable sectors, we rely on an approach suggested in Jensen and Kletzer (2005). They argue that the share of output or employment in tradeable goods should vary widely across regional entities (cities in our case) since different cities will be more heavily concentrated in producing different goods which they can then trade. For non-tradeable goods, on the other hand, assuming that preferences are the same across cities, one should observe similar proportions of workers in their production across cities. We rank industries by the variance of their employment shares across cities and call the industries in the top third, high trade industries, those in the middle third, medium trade industries, and those in the bottom third, low trade industries.[50] In Table (3.6), we present estimates of our basic model carried out separately for the low, medium and high trade industries. While the estimated effect of changes in R_{ct} do tend to be slightly higher for the low trade industries, the effects for the medium and high trade industries continue to be strongly significant and of the same order of magnitude as the estimated effects we obtained from the overall sample. Thus, our results do not appear to be arising simply because of spill-overs into the non-traded goods sector labor market. In Table (3.7) we also present results of estimating α_2 separately for twelve industry groupings. As can be seen, the results are very similar across these industries, with the effect estimated for manufacturing industries corresponding closely to that observed for the overall sample. This is further indication that the effects do not seem concentrated in non-trade good sectors since most manufactured goods are tradeable across cities.[51]

[50]The actual observations in the low trade industries is much lower than those in the medium and high trade industries because the low trade industries tend to be small and so tend to be disproportionately dropped when we impose our restriction that a given industry-city cell must contain at least 20 observations.

[51]Interestingly, the sector where wages exhibit the least response to change in R_{ct} is public administration. This may be due to the fact that wage for federal employees are generally set at the national level.

2.5.3 Education Breakdowns

To this point, we have established that: shifts in industrial composition toward higher paying sectors is associated with growth in wages in almost all industries; this result holds up to corrections for endogeneity and sample selection; and it is not proxying for explanations for city growth based on overall demand, diversity of the industrial structure, or changes in the education level of the workforce. We are interested, now, in examining whether the effects are ubiquitous across different education groupings. As previously indicated, the model presented in section 2 conceptually applies to workers of a given skill group. Accordingly, in this section we report results associated with estimating equation (2.22) for different skill groups defined by education and experience. In particular, we consider three education groups: workers with at most a high school education, workers with some post secondary education but without a BA, and finally a group with at least a BA. We further divide each of these groups into young workers (those with less than 10 years of experience[52]) and older workers (those with more than ten years of experience). For each of these groups we calculate an R_{ct} variable that is specific to the group.[53] The estimates of the effects of changes in R_{ct} on each of these groups is presented in Table (3.8). All the results in this table also control for changes in the employment rate interacted with 16 industry grouping as before.[54]

The results in Table (3.8) indicate substantial, though not identical, effects of shifts in industrial composition on within-industry wages in all the skill groups. The estimates reveal a slight education gradient, with ΔR effects being smaller for those with a BA or more in both experience groups. However, these differences are not very significant. The main dimension on which there are significant differences is with respect to experience. As shown in the table, the effects of a change in industrial composition on wages is much greater for younger workers than for older workers.[55] While this observation is not consistent with a strict interpretation of the model, it is potentially consistent with an internal labor market view in which more experienced workers tend to be partially shel-

[52]The measure of experience we use is potential experience, defined by Age - Years of School - 6. Appendix B contains details.

[53] We create the national level wage premia, and thus the R_{ct} measures, separately for each skill group. We, again, estimate in two steps, with the first step individual wage regression (including Dahl's selection correction) as well as the second step run separately for the 6 skill groups.

[54] For the results of this table, the effects of employment rates on different industries are constrained to be identical across the 6 skill groups. We also estimated the model without imposing such constraints. This gave similar, but less precise estimates.

[55] Since we use different R_{ct} measures for each skill group, the group-specific results in Table (3.8) do not bear an easy mechanical relationship to the overall results in Table (3.2). This explains why the estimates for some groups in Table (3.2) that are substantially below the Table (3.2) estimates are not counterbalanced by substantially higher estimates for other groups. In contrast, if we use the single, overall R_{ct} measure for all groups, we find the same education and experience patterns as in Table (3.8) but with coefficients that are all larger. For example, the estimates for young workers with high school education is above 4, while the result for older post-secondary workers is above 2.

tered from direct market comparisons. For example, if workers in ongoing relationship only renegotiate their wages infrequently, this is the pattern we would expect. Hence, we believe that the results based on focusing on different skill groups paint a similar picture to that documented previously, that is, changes in industrial composition appear to have substantial general equilibrium effects on within industry wages.

2.5.4 Additional Effects Associated with Changes in Industrial Composition

Up to this point, our empirical investigation has focused on evaluating how shifts in industrial composition toward high paying jobs affects wages. In this subsection, we briefly explore the effects of such a change on other city level outcomes. Given our interpretation of the wage effects, it would be natural to expect - as documented in section 1.3 - that a change in industrial composition that favors high paying jobs should be associated with in-migration and, potentially, an increase in the price of housing. In Table (2.9) we investigate this possibility. In the first three columns of the table, we examine whether changes in R_{ct} are associated with increases in the price of housing, as measured by the rent for one bedroom apartments. As can be seen from the table, we again observe a positive association. It is worth noting that while the estimated coefficient on the change in R_{ct} varies substantially across our different estimation strategies, in all three cases we find that housing prices capitalize a large faction of the changes in wages.[56]

It is interesting, in addition, to consider the effect of changes in R_{ct} on labor force growth, as we do in the specifications in columns 4, 5 and 6 in Table (2.9). The results in these columns show that a change in industrial composition in favor of high paying jobs has a robust positive association with labor force growth. Together, the observations on the effects of shifts in industrial composition on labor force growth and housing costs suggest that a city that experiences a positive increase in R_{ct} becomes a more attractive city, as we would expect given the impact in terms of higher average wages that we demonstrated in the rest of the paper. [57]

[56] Since rent makes up only a fraction of total consumption, under perfect mobility of workers across cities we could expect that the effect of a change in R_{ct} on housing prices to be greater than the effect on wages.

[57] Our analysis suggest that changes in industrial composition affects wages, which affects mobility and finally affects housing prices. Alternatively, one may think that the main causality runs in the opposite direction, with housing prices (for some unspecified reason) being the initial driving force. Although we have trouble formalizing this line of reasoning, we nevertheless explored its empirical relevance by including changes in housing costs as an additional regressor in our baseline wages regressions. When we do not instrument housing costs, we do find that they enter significantly in the determination of wages, but we do not find that they substantially change our estimated effects of industrial composition. Moreover, when we use weather patterns to instrument changes in housing prices, as would be appropriate if housing prices reflect a change in the desire to locate in easier climates, then we find no significant effect of housing prices on wages. Together these results support the direction of causality suggested by our model, and provide little support for the alternative direction of causality.

2.6 Conclusion

Policy forums often involve discussions of the effects and desirability of attracting or retaining jobs in high paying industries. In the popular press, it is common to hear statements claiming that economic success is closely tied to favoring employment growth in sectors that pay high wages for comparable individuals. In contrast, the most prevalent view among economic researchers is that changes in industrial composition generally contribute very little to labor market performance and therefore a focus on the effects of different policies with respect to the creation or destruction of better paying jobs is likely misplaced. This consensus position is based primarily on evaluating the economic impact of changes in industrial composition using a simple accounting approach which assumes away general equilibrium effects from the loss of jobs in one sector on wages in other sectors. Although traditional economic theory may provide good reason to believe that such general equilibrium effects should be absent or small, in this paper we build on a bargaining model to highlight an empirical strategy for evaluating the general equilibrium effects of changes in job composition on industry level wage payments. We implement that strategy using U.S. census data from 1970 to 2000. Our main finding is that general equilibrium effects appear pervasive, persistent and large. In particular, at the city level we find that having jobs more concentrated in high paying industries has an effect on the average wage within the city that is 2.5 to 4 times larger than that implied by the common composition adjustment accounting approach. We show that these results are robust to using different instrumental variable strategies, controlling for worker selection and focusing on sectors producing highly tradeable goods.

Our results suggest that policies or events which affect industrial composition should not be evaluated simply using the standard accounting approach but instead should explicitly take account of substantial general equilibrium effects implied by a social interaction model of wage determination. For example, it is common for the opening up of trade relationships to involve a reallocation of high and low paying jobs across trading partners. Our results suggest that a proper evaluation of the effects of increased trade needs to incorporate the potential general equilibrium effects on wages in other sectors. In general, recognizing and quantifying these feedback effects may lead to much more variable assessments of the gains from trade since markets that attract high paying industries may benefit more than traditionally thought, while markets that lose such jobs should benefit less.[58]

[58] Beaudry, Collard and Green (2005) and Beaudry and Collard (2006) find that increased openness to international trade over the period 1978-98 had very uneven effects across countries. In particular, countries that attracted high-capital-high-wage industries gained dis-proportionally relative to countries that increased employment in low-capital intensive industries. The general equilibrium effects found in this paper offer a potential explanation for the size of the effects found in these two papers.

Table 2.1: Basic Results

	OLS	OLS	IV1	IV2	IV1&IV2	OLS	IV1&IV2
Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ΔR_{ct}	2.580	2.574	2.940	2.895	2.916		
	(0.193)*	(0.193)*	(0.424)*	(0.41)*	(0.4)*		
$\sum_i \nu_{it-1}(\eta_{ict} - \eta_{ict-1})$						2.352	2.952
						(0.215)*	(0.443)*
$\sum_i \eta_{ict}(\nu_{it} - \nu_{it-1})$						2.911	2.814
						(0.468)*	(0.634)*
ΔER_{ct}	0.444						
	(0.089)*						
$\Delta ER_{ct} \times$ Ind.	No	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	28425	28425	28425	28425	28425	28425	28425
R^2	0.55	0.552				0.553	

Notes: Each column is an estimate of equation (2.22). Stars (*) denote significance at the 5% level. Standard errors, in parentheses, are clustered at the year-city level. All regressions contain a full set of industry-by-year dummies.

Table 2.2: Basic Results with Selection Correction

	OLS	OLS	IV1	IV2	IV1&IV2	OLS	IV1&IV2
Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ΔR_{ct}	2.626	2.619	2.996	2.999	2.997		
	(0.201)*	(0.201)*	(0.451)*	(0.449)*	(0.435)*		
$\sum_i \nu_{it-1}(\eta_{ict} - \eta_{ict-1})$						2.338	2.995
						(0.218)*	(0.472)*
$\sum_i \eta_{ict}(\nu_{it} - \nu_{it-1})$						3.046	3.005
						(0.491)*	(0.687)*
ΔER_{ct}	0.433						
	(0.097)*						
$\Delta ER_{ct} \times$ Ind.	No	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	27945	27945	27945	27945	27945	27945	27945
R^2	0.595	0.597				0.598	

Notes: Each column is an estimate of equation (2.22) using the selection correction procedure described in section (4.2) of the text. Stars (*) denote significance at the 5% level. Standard errors, in parentheses, are clustered at the year-city level. All regressions contain a full set of industry-by-year dummies.

Table 2.3: Reflection Specification

	OLS	IV1	IV2	IV1&IV2
Variables	(1)	(2)	(3)	(4)
Δ Av. City Wage	0.853	0.689	0.698	0.693
	(0.013)*	(0.038)*	(0.038)*	(0.036)*
$\Delta ER_{ct} \times$ Ind.	Yes	Yes	Yes	Yes
Obs.	28425	28425	28425	28425
R^2	0.626			

Notes: Each column is an estimate of equation (2.26). Stars (*) denote significance at the 5% level. Standard errors, in parentheses, are clustered at the year-city level. All regressions contain a full set of industry-by-year dummies. Average city wage is equal to $\sum_i \eta_{jct} \ln w_{ict}$, as described in equation (2.26).

Table 2.4: Reflection Specification with Selection Correction

	OLS	IV1	IV2	IV1&IV2
Variables	(1)	(2)	(3)	(4)
Δ Av. City Wage	0.857	0.689	0.701	0.695
	(0.012)*	(0.039)*	(0.039)*	(0.038)*
$\Delta ER_{ct} \times$ Ind.	Yes	Yes	Yes	Yes
Obs.	27945	27945	27945	27945
R^2	0.667			

Notes: Each column is an estimate of equation (2.26) using the selection corrected wages. Stars (*) denote significance at the 5% level. Standard errors, in parentheses, are clustered at the year-city level. All regressions contain a full set of industry-by-year dummies. Average city wage is equal to $\sum_i \eta_{jct} \ln w_{ict}$, as described in equation (2.26).

Table 2.5: Alternative Explanations

	OLS	IV1	IV2	OLS	IV1	IV2	OLS	IV1	IV2
Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
ΔR_{ct}	2.624	2.928	2.900	2.579	2.701	2.664	2.552	2.947	2.948
	(0.201)*	(0.437)*	(0.447)*	(0.192)*	(0.415)*	(0.402)*	(0.218)*	(0.65)*	(0.65)*
1 - Herfindahl	0.283	0.291	0.29						
	(0.205)	(0.217)	(0.215)						
Δ BA or $>$				0.527	0.539	0.538			
				(0.16)*	(0.163)*	(0.163)*			
Δ log labour force							0.017	0.007	0.007
							(0.016)	(0.033)	(0.033)
$\Delta ER_{ct} \times$ Ind.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	27945	27945	27945	27945	27945	27945	27945	27945	27945
R^2	0.598			0.601			0.597		

Notes: Each column is an estimate of equation (2.22) with an added regressor which reflects an alternative hypothesis regarding wage outcomes as discussed in section (5.1). Stars (*) denote significance at the 5% level. Standard errors, in parentheses, are clustered at the year-city level. All regressions contain a full set of industry-by-year dummies.

Table 2.6: By Trade and Non-Trade Industries

	Low Trade			Medium Trade			High Trade		
	OLS	IV1	IV2	OLS	IV1	IV2	OLS	IV1	IV2
Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Δ Total Rent	3.165	4.081	4.496	2.777	3.167	2.972	2.469	2.791	2.872
	(0.422)*	(1.162)*	(0.962)*	(0.232)*	(0.539)*	(0.545)*	(0.18)*	(0.392)*	(0.386)*
$\Delta ER_{ct} \times$ Ind.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	4518	4518	4518	11527	11527	11527	11900	11900	11900
R^2	0.555			0.57			0.645		

Notes: Each column is an estimate of equation (2.22) estimated separately for low, medium, and high trade industries. Definitions are provided in section (5.2) of the text. Stars (*) denote significance at the 5% level. Standard errors, in parentheses, are clustered at the year-city level. All regressions contain a full set of industry-by-year dummies.

Table 2.7: Breakdown by Industry Group

	OLS	IV1	IV2	IV1&IV2
Industry Group	(1)	(2)	(3)	(4)
Durable Manufacturing	3.322	2.585	3.471	3.145
	(0.225)*	(0.448)*	(0.393)*	(0.389)*
Non-Durable Manufacturing	3.062	2.348	2.856	2.577
	(0.293)*	(0.665)*	(0.668)*	(0.672)*
Construction	3.363	4.044	3.120	3.561
	(0.262)*	(0.612)*	(0.534)*	(0.489)*
Transport and Utilities	2.383	2.549	2.720	2.681
	(0.256)*	(0.499)*	(0.459)*	(0.453)*
Wholesale and Retail Trade	3.152	3.502	3.439	3.472
	(0.271)*	(0.547)*	(0.562)*	(0.532)*
Professional	2.268	2.907	2.749	2.822
	(0.2)*	(0.432)*	(0.41)*	(0.397)*
Agriculture and Mining	2.856	2.567	2.426	2.561
	(0.442)*	(0.898)*	(0.927)*	(0.932)*
Finance, Insurance, Real Estate	2.399	4.039	4.281	4.198
	(0.314)*	(0.755)*	(0.742)*	(0.716)*
Public Administration	1.279	1.773	1.974	1.883
	(0.23)*	(0.545)*	(0.52)*	(0.504)*
Others (Private, Business, Entertainment)	3.256	3.954	4.055	3.990
	(0.338)*	(0.761)*	(0.774)*	(0.744)*
Obs.	27946	27946	27946	27946
R^2	0.601			

Notes: Each column is an estimate of equation (2.22) estimated separately by industry group. Stars (*) denote significance at the 5% level. Standard errors, in parentheses, are clustered at the year-city level. All regressions contain a full set of industry-by-year dummies.

Table 2.8: Breakdown By Education and Potential Experience

		OLS	IV1	IV2	IV1&IV2
Experience	**Ed. Group**	(1)	(2)	(3)	(4)
	HS	2.944	2.508	3.214	2.894
		(0.197)*	(0.294)*	(0.299)*	(0.254)*
< 10	SP	2.791	2.861	3.540	3.303
		(0.235)*	(0.495)*	(0.492)*	(0.441)*
	BA	1.951	2.033	2.504	2.431
		(0.192)*	(0.639)*	(0.382)*	(0.377)*
	HS	1.412	3.681	1.358	1.848
		(0.172)*	(1.037)*	(0.348)*	(0.349)*
> 10	SP	1.708	1.943	1.959	2.042
		(0.377)*	(3.422)	(0.732)*	(0.684)*
	BA	1.171	1.756	1.572	1.596
		(0.246)*	(1.660)	(0.403)*	(0.395)*
$\Delta ER_{ct} \times$ Ind.		Yes	Yes	Yes	Yes
	Obs.	50306	50306	50306	50306
	R^2	0.323			

Notes: Each column is an estimate of equation (2.22) stacked by education and potential experience group, where we have restricted common employment $\times$ industry effects but is otherwise fully interacted. Stars (*) denote significance at the 5% level. Standard errors, in parentheses, are clustered at the year-city level. All regressions contain a full set of industry-by-year dummies.

Table 2.9: Housing Price and Labor Force Growth

	Log 1BR Rent			Log Labor Force		
	OLS	IV1	IV2	OLS	IV1	IV2
Variables	(1)	(2)	(3)	(4)	(5)	(6)
ΔR_{ct}	3.500	4.439	2.772	2.400	6.193	5.990
	(0.639)*	(1.023)*	(0.992)*	(0.772)*	(1.748)*	(1.722)*
ΔER_{ct}	1.633	-1.004	1.699	1.836	12.904	13.232
	(0.327)*	(1.877)	(1.170)	(0.476)*	(3.582)*	(2.860)*
Obs.	456	456	456	456	456	456
R^2	0.172			0.205		

Notes: Stars (*) denote significance at the 5% level. Standard errors, in parentheses, are clustered at the year-city level. All regressions contain a full set of year dummies.

Figure 2.1: Average City Wages and Total Rent

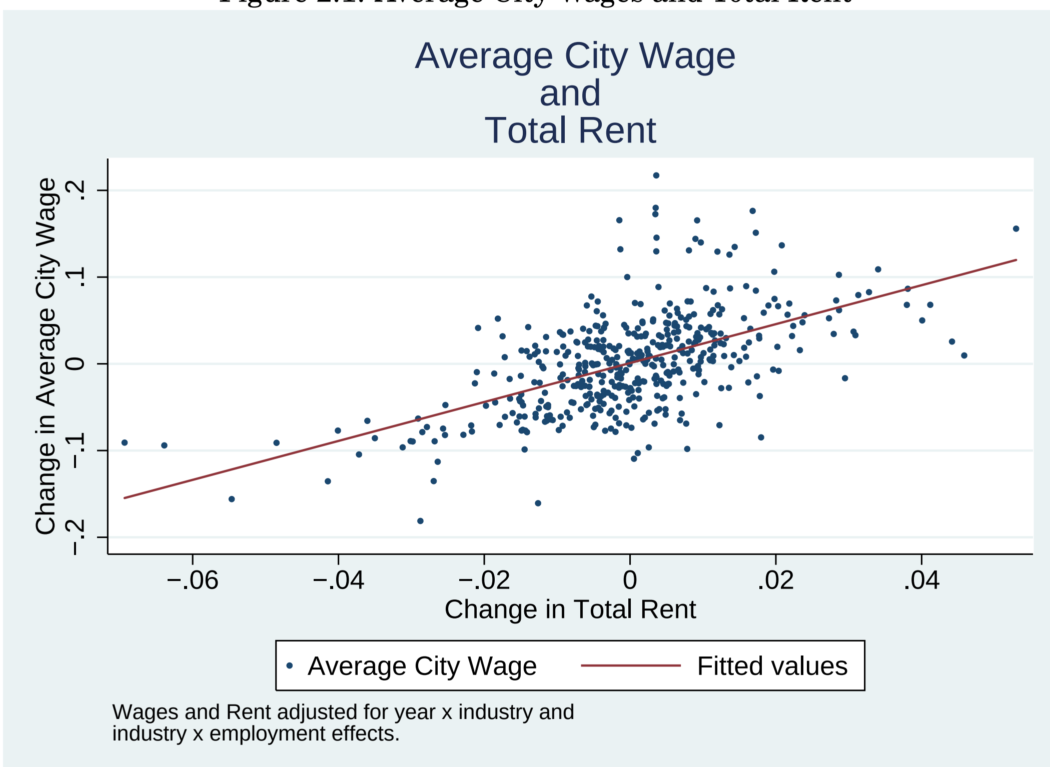

Figure 2.2: Estimates By Industry

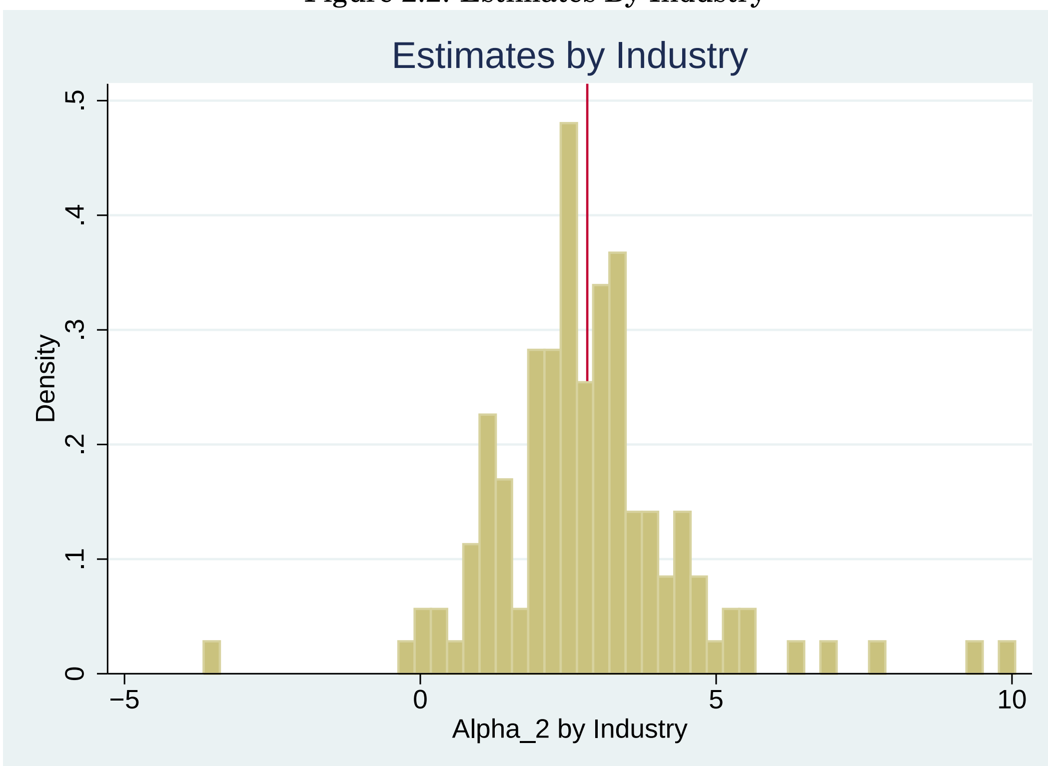

Figure 2.3: Change in Rent vs. Change in Employment Rates

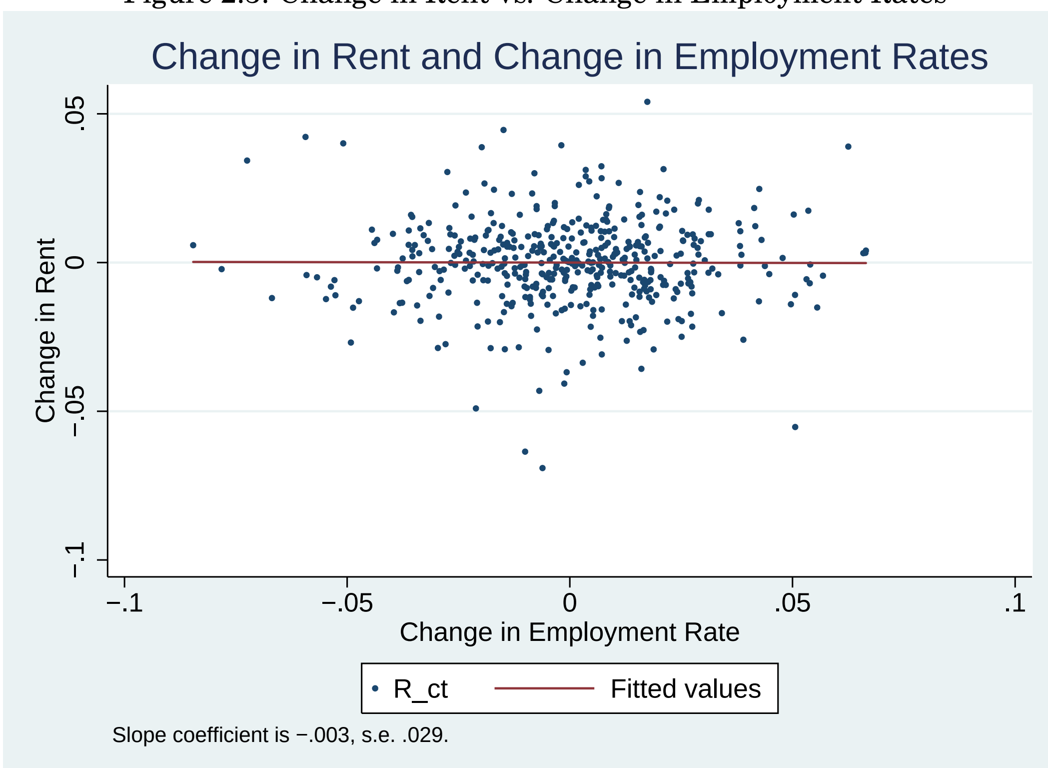

Chapter 3

Industrial Composition and the Gender Wage Gap: Evidence from U.S. Cities

3.1 Introduction

A great deal of research in labor economics attempts to explain both the levels and trends of male-female log wage differentials. The basic facts, documented in Blau and Kahn (1997, 2006), Card and DiNardo (2002) among others, are reproduced in Figure 3.1 using CPS data from 1970 - 2007.[59] As Card and DiNardo (2002) point out, trends in the raw male-female pay gap, shown as the solid line, seem to fall into three distinct periods: the gap was relatively stable in the early 1970s before declining sharply through the 1980s and early 1990s before the gap returns to being stable in the mid-1990s. For example, the gender pay gap in log hourly wages was 0.457 in 1970 and fell by 0.027 to 0.43 in 1980. Over the 1980s, the gap fell by 0.193 to 0.237 in 1993, and then only declined by an additional 0.047 by 2007.

In order to understand the gender pay gap, researchers often decompose mean log wage differentials into 'explained' and 'unexplained' components (Altonji and Blank 1999). The explained portion of the gap is due to differences in (productivity-related) characteristics while the unexplained portion may be due to a number of factors, including discrimination and differences in unobserved characteristics or omitted variables. Most studies find that a substantial pay gap remains even after controlling for observable characteristics, such as education, experience, occupation and industry. To analyze changes in the gender pay gap, the explained and unexplained gap components can be further decomposed over time. For example, Blau and Kahn (2006) use PSID data to study the slowdown of convergence of female and male wages over the 1990s compared to the 1980s. They find that women's relative wage gains in the 1980s were due to

[59] This figure uses March Current Population Survey (CPS) data on full-time full-year wage and salary workers aged 18-65. Full-time full-year workers are defined to be those who work at least 40 weeks in the previous year and report working at least 35 per week. Additional details on the construction of this series are contained in the Data Appendix.

relative human capital improvements, occupational upgrading, deunionization - which negatively impacted men relative to women - and a substantial improvement in the 'unexplained' gap (Blau and Kahn 1997, Blau and Kahn 2006). In comparison, only a small part of the slowdown in gender wage convergence in the 1990s can be explained by human capital characteristics or occupational upgrading; the largest explanation for the slowdown in wage convergence is slower narrowing of the 'unexplained' gap.[60]

In an alternative strand of literature, several observers have noted that the trend in the male-female wage gap appears to be closely linked to the male wage structure. For example, Fortin and Lemieux (2000) suggest that there is a systematic link between male wage inequality and gender pay ratios. This helps to explain why the gender gap fell more in the 1980s than the 1990s, since male wage inequality followed a similar pattern. Welch (2000) notes a similar pattern and offers a model in which changes in the relative prices of "brains" versus "brawn" explain the co-movement in male inequality and male-female pay ratios.

In this paper, we suggest an alternative explanation for the trends in male-female wage differences that is closely related to the male wage structure and present a search and matching model to rationalize this link. In particular, we argue that changes in the industrial structure during the 1980s adversely affected men relative to women and impacted the gender gap beyond pure compositional effects. For motivation, consider figure (3.2), which plots the gender gap after controlling for standard observable characteristics, including industry, (the dashed line) for each year from 1970 to 2005.[61] It should be noted that even after controlling for observable characteristics, the gap is still substantial and follows the same pattern as the raw wage gap. The solid line in the figure is male-female differences in "rent", calculated as the difference in the employment weighted sum of inter-industry wage differentials in year-by-year regressions.[62]

[60] Blau and Kahn (2006) consider four potential explanations for the slowdown in the unexplained gender gap: changing labor market selectivity, changes in supply and demand shifts, and differences in unmeasured characteristics and discriminations. They present evidence that each of these factors may have played a role.

[61] The adjusted male-female differential is calculated from the 1970-2007 March Current Population Surveys. In each of these cross-section surveys, we restrict the calculation to the sample of working men and women who worked full-time full-year in the civilian sector in the year prior to the survey, and who were not self-employed or working without pay. A worker is considered to be full-time full-year if they work at least 40 weeks per year and at least 35 hours per week. We calculate the adjusted wage differential in hourly earnings by estimating the following regression model in each of the cross sections:

$$\log w_i = X_i'\beta + \delta \cdot \text{female}_i + \epsilon_i$$

where X_i includes years of school, potential experience and its square, a black indicator, controls for occupation (12 categories) and industry (144 categories). The adjusted wage differential is given by the negative of the coefficient δ on the female dummy variable.

[62] Inter-industry wage differentials are the coefficients on the industry dummy variables in the adjusted wage gap regression described above. Gender specific employment shares are calculated from the CPS data. The difference in rent measure is given by $\sum_j (\pi_{Mj} - \pi_{Fj}) \cdot \xi_j$ where π_{gj} is the employment share for gender

This measure can be thought of as an index of relative male-female value of average employment; higher values of this index indicate that male employment is relatively more heavily concentrated in higher paying industries. These data indicate that the male-female wage differential and the rent series track each other closely, including the slowdown of gender wage convergence in the 1990s. The coefficient on the rent variable is positive and significant (t-statistic of around 3.5) in a regression of the adjusted gender gap on the rent variable and a time trend.[63] It should be emphasized that the gender gap reported in the figure controls for industrial composition and so any relationship is not mechanical. The implication is that male losses in relative employment in high paying industries or changes in industrial premia that favor women compared to men are associated with lower male-female wage differentials, even after accounting for differences in male-female composition of employment across industries.

To interpret this time-series evidence, we develop a search and matching model of the labor market and test it using a panel of data of relative wages across U.S. metropolitan areas for the years 1970 to 2000. Our model extends the basic search and matching model to include many sectors and two different types of workers: male or female. We take as a primitive that the distribution of male and female workers differs across industries and emphasize how these differences combined with search frictions and bargaining can generate differences in equilibrium wages between men and women in the absence of productivity disparities or discrimination.[64] In particular, access to high-paying typically-male industries can create upward pressure on male-female wage differentials by increasing the relative value of mens' outside options in a bargaining framework. We argue that these forces should operate in local economies and use geographical variation in cities' industrial structure to assess their relevance.

We show that there is a tight link between our model and the well-known Blinder-Oaxaca decomposition. In fact, our model nests this common decomposition and one of our objectives is to assess its usefulness in presence of bargaining. We find that the standard approach of omitting general equilibrium effects due to bargaining overstates the component of the gender gap that is attributed to 'unexplained' or discriminatory factors. For example, we estimate that the total effect - composition plus general equilibrium effects - of a relative change in employment composition between men and women among industries is approximately twice that given by the standard accounting procedure used in Blinder-Oaxaca decompositions. To quantify this, we find that the gender gap among full-time full-year workers in our sample of 152 U.S. cities fell from 0.42 in

$g \in \{M, F\}$ in industry j and the ξ_js are the industrial premia.

[63]This relationship is robust to including higher orders of time trends and the log ratio of men to women in the economy.

[64]It should be noted that differences in employment across industries maybe partly due to discrimination. We will return to this below.

1979 to 0.29 in 1989. About half, or 0.067 log points, of this change remains unexplained after accounting for observable factors such as education, potential experience, industry and occupation. 0.024 log points of this decline can be explained by industrial composition alone, and we estimate that bargaining accounts for just over 1/3 (.024/.067) of the 'unexplained' change. Of the total decline in the gender gap over this period of .127 log points, then, the total effect of changes in the industrial composition account for just under 40% ($\frac{2\times 0.024}{.127}$).

Our paper is related to Borjas and Ramey (1995) who investigate the link between trends in the college premium and employment in trade-impacted concentrated industries. They argue that the market structure of trade impacted industries is important for understanding wage differences between college educated workers and those with high school or less education. In particular, highly concentrated industries typically earn rents that are shared with workers, who are disproportionally lower-skilled, in the form of higher-than-average wages. Employment losses and trade pressures during the 1980s lowered rents in these industries and lead to a relative decline in the wages of less educated workers. They show that employment in a set of highly concentrated industries can explain a significant portion of the trends in the college premium over time and across metropolitan areas. While our approach and source of identification is very similar to that used in Borjas and Ramey (1995), our model and empirical approach emphasize a quite different mechanism for the effects of industrial composition on relative wages. In particular, we are interested in the effects of industrial composition on relative wages *after* controlling for composition effects in a bargaining framework. Our paper is also closely related to Beaudry, Green, and Sand (2008) who use the same data set and show that industrial composition can have important general equilibrium effects on wage levels, whereas we highlight the impact of the same mechanism on relative wages.

The remainder of the paper is organized as follows. In section 3.2 we present our search model and illustrate how search frictions can lead to wage differentials in the absence of productivity differences or discrimination, and in section 3 we discuss our empirical methodology. Section 4 describes our data and section 5 discusses our results and assess robustness. Section 6 concludes.

3.2 Model

The model presented here is general in the sense that it applies to an arbitrary number of worker types. Although we are only interested in two worker types, male and female, the model is presented in the more general form for two reasons. First, it provides a foundation for an empirical exercise performed below, where the number of types increases by dividing each gender into education sub-groups. Second, it suggests that the empirical exercise developed here is applicable to the analysis of a wide variety of observed wage differentials (e.g. based on race, age, immigrant status, etc.).

3.2.1 Fundamentals

The economy consists of firms and workers, and unfolds in continuous time. All agents are infinitely lived and discount the future at a rate, r. Matches separate for exogenous reasons at a rate, s. We focus on a steady state of an economy that exhibits search frictions (it takes time for workers and firms to meet one another).

Workers

There are G types of workers, $g = 1, ..., G$, each existing with associated population size of P_g. While unemployed, workers of type g encounter vacancies in industry n at the Poisson arrival rate q_{gn} (which, from the workers' perspective, is taken as given). If we let w_{gn} be the wage paid to a worker of type g in industry n, then their value function associated with being employed in industry n, E_{gn}, satisfies the Bellman equation:

$$rE_{gn} = w_{gn} + s \cdot (U_g - E_{gn}), \tag{3.1}$$

where U_g is the value function associated with being an unemployed worker of type g, and satisfies:

$$rU_g = b + \sum_n q_{gn} \cdot \max\left\{E_{gn} - U_g, 0\right\}, \tag{3.2}$$

where b is flow utility associated with being unemployed.

Firms

Each firm belongs to one of N industries, $n = 1, ..., N$. The value of the output produced by a worker when matched with a firm depends on the industry, but not on their type. In particular, each worker in industry n produces a unit of output valued at p_n. Thus, the value function associated with filling a vacancy in industry n with a worker of type g, J_{gn}, satisfies:

$$rJ_{gn} = p_n - w_{gn} + s \cdot \left[V_n - J_{gn}\right], \tag{3.3}$$

where V_n is the value function associated with holding a vacancy in industry n. Firms in industry n encounter workers at an arrival rate of $\hat{q}_{gn}$, implying that V_n satisfies:

$$rV_n = \sum_{g=1}^{G} \hat{q}_{gn} \cdot \max\left\{J_{gn} - V_n, 0\right\}. \tag{3.4}$$

Free Entry

Opening a vacancy in industry n requires k_n units of capital (the numeraire good). In the presence of free entry, the net value associated with opening a new vacancy, $V_n - k_n$, is zero. Thus, free entry implies:

$$V_n = k_n \text{ for all } n = 1, ..., N. \tag{3.5}$$

Bargaining

Once a worker and firm meet, they determine whether they want to form an employment relationship. No relationship forms if $E_{gn}+J_{gn} < U_g+V_n$, since there is no way of dividing the surplus from the relationship in such a way that both parties are better off than they would be if they instead continued to search. On the other hand, if $E_{gn} + J_{gn} \geq U_g + V_n$, then wages are determined by generalized Nash bargaining. By letting $\phi \in [0, 1]$ be the workers' relative bargaining power, the wage, w_{gn}, is such that the value functions satisfy:

$$E_{gn} - U_g = \frac{\phi}{1-\phi} \cdot \left[J_{gn} - V_n\right]. \tag{3.6}$$

Arrival Rates

We conceive of each industry as having their own associated job market. If the unemployment rate among workers of type g is u_g, then the measure of workers that type that are unemployed is $u_g P_g$. For some some $\lambda_{gn} \geq 0$, we say that $\tilde{u}_{gn} = \lambda_{gn} \cdot u_g \cdot P_g$ is the *effective* measure of unemployed workers of type g searching in the job market associated with industry n. The total effective measure of unemployed workers searching in this job market is therefore $\tilde{u}_n = \sum_g \tilde{u}_{gn}$. The total measure of matches formed per unit of time between workers and vacancies in industry n is assumed to depend on $\tilde{u}_n$ and v_n (the measure of vacancies in industry n), and is given by $M(\tilde{u}_n, v_n)$, where M is constant returns to scale. This implies that workers of type g encounter vacancies in industry n

at the rate of

$$q_{gn} = M(\tilde{u}_n, v_n) \cdot \frac{\tilde{u}_{gn}}{\tilde{u}_n} \cdot \frac{1}{u_g P_g} = \lambda_{gn} \cdot \frac{M(\tilde{u}_n, v_n)}{\tilde{u}_n}. \tag{3.7}$$

Similarly, vacancies in industry n encounter workers of type n at the rate of

$$\hat{q}_{gn} = \frac{M(\tilde{u}_n, v_n)}{v_n} \cdot \frac{\tilde{u}_{gn}}{\tilde{u}_n}. \tag{3.8}$$

Notice that random matching within an industry's search market corresponds to the case in which $\lambda_{gn} = \lambda_n$ for all g.

Steady State

In steady state, the measure of workers leaving the unemployed pool equals the measure entering the unemployed pool. If we define:

$$\tilde{q}_{gn} = \begin{cases} q_{gn} & \text{if } E_{gn} + J_{gn} \geq U_g + V_n \\ 0 & \text{otherwise,} \end{cases}$$

then the rate at which workers of type g leave employment is $Q_g \equiv \sum_j \tilde{q}_{gj}$. This then implies that the aggregate steady state condition is:

$$Q_g \cdot u_g = s \cdot (1 - u_g). \tag{3.9}$$

Furthermore, the steady-state proportion of workers of type g in industry n, π_{gn}, satisfies

$$s \cdot (1 - u_g) \cdot \pi_{gn} = \tilde{q}_{gn} \cdot u_g.$$

Using (3.9), this is equivalent to:

$$\pi_{gn} = \frac{\tilde{q_{gn}}}{\sum_j \tilde{q}_{gj}}. \tag{3.10}$$

3.2.2 Equilibrium

The wage profile, $w = \{w_{gn}\}$, and employment distribution, $\{\pi_{gn}\}$, constitutes an equilibrium if we can find value functions, arrival rates, unemployment rates, and vacancy levels that simultaneously satisfy equations (3.1)-(3.9).

One set of equilibria that we may be particularly interested in are what we call *full support equilibria*. These are equilibria that support the employment of a positive measure of workers in each industry. Focusing on this class is not particularly restrictive in the case of a single type, and greatly aids the calculation of equilibria in the case of multiple worker types. Given that we are interested in studying equilibrium wages, and not, say, employment distributions, the benefit associated with this reduced focus seems to outweigh the loss of generality.

To be more precise, a full support equilibrium is an equilibrium in which $E_{gn} + J_{gn} \geq U_g + V_n$ for all n. This immediately implies that $\tilde{q}_{gn} = q_{gn}$, and via bargaining, ensures that $\max\{E_{gn} - U, 0\} = E_{gn} - U$ and $\max\{J_{gn} - V_n, 0\} = J_{gn}$. These facts allow us to solve for equilibrium wages in a relatively standard way, as follows.

Solving

Solving (3.1) and (3.2) for E_{gn} and U_g, and using the full support conditions, allows us to express the workers' side of the bargaining condition (3.6) as:

$$E_n - U = \frac{1}{r+s}\left[w_{gn} - (1-\delta)\cdot b - \delta_g \cdot \sum_n \pi_{gn} w_{gn}\right], \tag{3.11}$$

where $\delta_g \in (0,1)$ is a measure of search frictions facing workers of type g (high values imply lower frictions), and is given by:

$$\delta_g \equiv \frac{Q_g}{r+s+Q_g}. \tag{3.12}$$

Using (3.3) and (3.5), along with the full support conditions, allows us to express the firms's side of (3.6) as:

$$\frac{\phi}{1-\phi}\cdot[J_{gn} - V_n] = \frac{\phi}{1-\phi}\cdot\frac{1}{r+s}\left[\psi_n - w_{gn}\right], \tag{3.13}$$

where $\psi_n \equiv p_n - rk_n$ is the surplus derived from a match in industry n. Using (3.11) and (3.13) in (3.6) gives us an expression for equilibrium wages:

$$w_{gn} = \phi\cdot\psi_n + (1-\phi)\cdot\left[(1-\delta)\cdot b + \delta_g \cdot \overline{w}_g\right], \tag{3.14}$$

where $\overline{w}_g \equiv \sum_j \pi_{gj} w_{gj}$ is the average wage of among workers of type g.

Equation (3.14) indicates a reflection problem since the wage of a worker in industry n depends on the average wage of workers across industries. To solve for the average wage, we can take the expecation of both sides of (3.14) and re-arrange to get:

$$\overline{w}_g = \left[\frac{(1-\phi)(1-\delta_g)}{1-(1-\phi)\delta_g}\right]\cdot b + \left[\frac{\phi}{1-(1-\phi)\delta_g}\right]\cdot\sum_j \pi_{gj}\cdot\psi_j. \tag{3.15}$$

Using this in (3.14) gives wages in terms of exogenous variables:

$$w_{gn} = b_{0g} + (\phi \cdot \psi_n) + \gamma_g \cdot \sum_{j=1}^{N} \pi_{gj} \cdot (\phi \cdot \psi_j), \tag{3.16}$$

where the constant is:

$$b_{0g} \equiv \frac{(1-\phi)(1-\delta_g)(1+\phi\delta_g)}{1-(1-\phi)\delta_g} \cdot b, \tag{3.17}$$

and the parameter γ_g is given by:

$$\gamma_g \equiv \frac{(1-\phi) \cdot \delta_g}{1-(1-\phi) \cdot \delta_g}. \tag{3.18}$$

This parameter plays a large role in the analysis since it quantifies the relevance of general equilibrium effects, as described shortly.

Full Support Conditions

Since we are interested in full support equilibria, we need to determine the conditions under which this type of equilibrium exists. Fixing some (g, n), the Nash bargaining condition ensures that $E_{gn} - U_g$ is non-negative if and only if $J_{gn} - V_n$ is non-negative. Thus, the full surplus condition, $E_{gn} + J_{gn} \geq U_g + V_n$, will be satisfied if and only if $J_{gn} - V_n \geq 0$. Since $J_{gn} - V_n$ is a positive multiple of $\psi_n - w_{gn}$, it follows that we can simply verify that:

$$\psi_n \geq w_{gn} \text{ for all } (g, n). \tag{3.19}$$

Using (3.16) to obtain an expression for w_{gn}, this can be expressed in the following way that takes unemployment and vacancy levels as given:

$$\psi_n \geq \frac{b_{0g}}{1-\phi} + \frac{1}{1-\phi} \cdot \gamma_g \cdot \sum_{j=1}^{N} \pi_{gj} \cdot (\phi \cdot \psi_j).$$

Since the right side of this is independent of industry, it is sufficient to check that this holds for the lowest-surplus industry. Since the final term reflects a multiple of the average surplus encountered by some worker type, this can be thought of as a condition that industries are not too different in terms of their associated surplus (for given unemployment and vacancy levels).

3.2.3 Interpreting γ_g

Consider one type of worker for which we observe a distribution of employment $\pi = \{\pi_1, ..., \pi_N\}$ and wages $w = \{w_1, ..., w_N\}$ (the type subscript dropped for convenience). Naturally, the average wage is given by:

$$\overline{w} \equiv \sum_n \pi_n w_n = b_0 + \left[\sum_j \pi_j \cdot \phi \cdot \psi_j\right] + \gamma \cdot \left[\sum_j \pi_j \cdot \phi \cdot \psi_j\right] \tag{3.20}$$

$$= b_0 + (1+\gamma) \cdot \phi \cdot \left[\sum_j \pi_j \cdot \psi_j\right]. \tag{3.21}$$

Suppose that we want to predict the effect a change in the average industry premium, $\sum_j \pi_j \cdot \psi_j$, on average wages for this particular type. If general equilibrium effects were ignored, then average wages would be predicted to increase by ϕ times the change in the average industry premium. This exercise understates the true impact on average wages, however, since all wages are affected by the average in the presence of bargaining. The extent of this is captured by $\gamma \cdot \phi$ times the change in the average industry premium. Thus, as pointed out in Beaudry, Green, and Sand (2008), the presence of bargaining magnifies any mechanical change in average wages induced by a change in the average industry premium. As such, γ becomes our focus in the empirical section.

3.2.4 Relative Wages

In order to sharpen the focus on differences in *relative* search patterns, suppose that aggregate arrival rates are equal across types: $Q_g = Q$ for all g. This implies that all types face the same level of search frictions: $\delta_g = \delta$ for all g. This in turn implies that $\gamma_g = \gamma$ and $b_{0g} = b_0$ for all g. In this case, the difference in wages between two types, g and g', in industry n is:

$$w_{gn} - w_{g'n} = \gamma \cdot \sum_{j=1}^{N} (\pi_{gj} - \pi_{g'j}) \cdot (\phi \cdot \psi_j). \tag{3.22}$$

This expression indicates that different types will, in general, be paid different amounts even within the same job. To be sure, this is not because of differences in productivity, nor because of differences in search frictions. Rather, it is driven by differences in relative search patterns: facing any potential employer, those types that are more likely to encounter lower-surplus jobs have a worse bargaining position compared with types that are more likely to encounter high-surplus jobs.

3.2.5 Empirical Connections

The estimating equations are based on the wage equation (3.16). In order to estimate γ_g, we conceptualize the above economy as operating at the city level. In particular, we see the relevant variation as occurring in the different search patterns across cities. That is, the π_{gn} terms now become city specific: π_{gnc}. That is, wages for workers of type g, that work in industry n in city c, are given by:

$$w_{gnc} = b_{g0} + (\phi \cdot \psi_n) + \gamma_g \cdot \sum_{j=1}^{N} \pi_{gjc} \cdot (\phi \cdot \psi_j). \tag{3.23}$$

Note that differences in employment patterns across cities will induce different entry patterns across industries in general equilibrium, and therefore will also influence unemployment rates across cities.

Note that we can write the wage as an industry-specific term plus a type×city-specific term:

$$w_{gnc} = \xi_n + \xi_{gc}, \tag{3.24}$$

where $\xi_n \equiv b'_{0g} + \phi \cdot \psi_n$, and $\xi_{gc} \equiv b''_{0g} + \gamma_g \cdot \sum_{j=1}^{N} \pi_{gjc} \cdot (\phi \cdot \psi_j)$, where $b'_{0g} + b''_{0g} = b_{0g}$. These two components can be estimated as the coefficients on industry and type×city dummies in an individual level wage regression. In practice, this individual-level regression will control for other individual characteristics, such as education, experience, race, and occupation.

Notice how ξ_{gc} can be expressed in terms of ξ_n in the following way:

$$\xi_{gc} = b_{g0} + \gamma_g \cdot \sum_{j=1}^{N} \pi_{gjc} \cdot \xi_j. \tag{3.25}$$

We can estimate the employment distribution (the π_{gjc} terms) from the observed employment distribution within cities. These estimates, along with the estimates of ξ_n allow us to generate the following 'rent' variable:

$$R_{gc} \equiv \sum_{j=1}^{N} \hat{\pi}_{gjc} \cdot \hat{\xi}_j. \tag{3.26}$$

Then, an estimate of γ_g is obtained by estimating:

$$\hat{\xi}_{gc} = \gamma_{g0} + \gamma_g \cdot R_{gc} + \nu_{gc}, \tag{3.27}$$

where γ_{0g} is a type-specific constant and ν_{gc} is a mean-zero error term. For OLS to be appropriate, we require ν_{gc} to be uncorrelated with R_{gc}. There are at least two reasons why this may not hold. First, there may be city-specific characteristics that are correlated with average industry premia within cities: e.g. workers in a city near a port may have a productivity advantage over workers in other cities, and this city may tend to have employment concentrated in high-surplus industries. This type of problem can be dealt with in two ways. First, we can simply include city dummies. Second, if the aggregate arrival rate (Q_g) is constant across types, then the coefficients γ_{0g} and γ_g are independent of type. This allows us to estimate γ by taking differences across types (with cities). In particular, for some pair of types (g, g'), let $\Delta X_z \equiv X_{gc} - X_{g'z}$, so that the estimating equation of interest is:

$$\Delta\hat{\xi}_c = \Delta\gamma_0 + \gamma \cdot B_c + \Delta\nu_c, \tag{3.28}$$

where $B_c \equiv \Delta R_c$ represents a 'bargaining' variable. Estimating this city-level specification will also deal with the omitted city effect, since this effect is canceled once the difference is taken.

Second, and more importantly, there may be city-gender specific factors that are correlated with average industry surplus. For instance, if some type (e.g. females) tend to be discriminated against, and the extent of discrimination i) varies across cities, and ii) is correlated with that types' average industry surplus within the city: e.g. if the cities that tend to discriminate against females also tend to be the ones in which females find it difficult to obtain employment in high-surplus industries. This type of problem may be overcome by exploiting the time dimension of the data. By pooling observations, it becomes possible to include city$\times$type dummies (or, city dummies if working in type differences). Below, we discuss in detail our empirical implementation.

3.3 Empirical Implementation

The main empirical implication to take way from the above model is that differences between genders in employment across industries can generate differences in equilibrium wages between men and women in the absence of productivity disparities or discrimination. In particular, the model predicts that in cities where men's employment is relatively more heavily weighted toward higher paying industries compared to women, gender wage gaps will also be higher after accounting for pure compositional effects of differential employment patterns. We view our model as applying to men and women with homogeneous skills and differing only with respect to employment across industries. In reality, however, workers of both genders have different observable characteristics that affect their productivities. Our approach to dealing with this heterogeneity is to control for observable characteristics in a regression context. In addition, while the model emphasizes a cross-sectional relationship, one might worry about city-specific factors that influence both the gender gap and B_c. Thus, we pull data from several cross-sections and pool them together to control for city fixed-effects. Specifically, consider the following wage equation for the pooled sample of men and women at the individual level:

$$w_{i,g,c,t} = X'_{i,g,c,t} \cdot \mathrm{B}_t + Z'_{c,t}\beta_g + \gamma \cdot R_{g,c,t} + d_c + d_g + d_t + d_{c,g} + d_{g,t} + d_{c,t} + \epsilon_{i,g,c,t}, \quad (3.29)$$

where the subscripts i, g, c and t index individuals, gender, city and year, respectively. $X'_{i,g,c,t}$ is a vector of individual characteristics whose effects are allowed to vary by year, B_t. $Z_{g,t}$ is a vector of city level characteristics whose effects may vary by gender, β_g. $R_{g,c,t}$ is our measure of city-gender rent, and γ is the main coefficient of interest. If $\gamma = 0$, we do not expect differences in the distributions of genders across industries to have effects on the gender gap beyond that captured by pure compositional effects. d_c, d_g and d_t are unrestricted city, gender, and year fixed-effects, respectively. We also allow unrestricted effects for city-gender, $d_{c,g}$; gender-year, $d_{g,t}$; and city-year, $d_{c,t}$. It is well known that γ can be estimated by a two-step approach by first fitting:

$$w_{i,g,c,t} = X'_{i,g,c,t} \cdot \mathrm{B}_t + d_{c,g,t}, \quad (3.30)$$

where $d_{c,g,t}$ are a complete set of unrestricted city-gender-year dummies. The coefficients on the city-gender-year dummies can be used as a dependent variable in the following model:

$$d_{c,g,t} = Z'_{c,t}\beta_g + \gamma \cdot R_{g,c,t} + d_c + d_g + d_t + d_{c,g} + d_{g,t} + d_{c,t} + \bar{\epsilon}_{g,c,t}. \quad (3.31)$$

We further transform the equation by taking gender differences within city-year cells, so that we are working with city-year level gender gaps:

$$\Delta w_{c,t} = Z'_{c,t}\Delta\beta + \gamma \cdot \Delta R_{c,t} + \Delta d_g + \Delta d_c + \Delta d_t + \Delta\bar{\epsilon}_{c,t}, \quad (3.32)$$

where the symbol Δ represents gender differences of the variable (or coefficient) that directly follows it. For example, $\Delta w_{c,t}$ represents the gender gap in wages for city c and year t. Δd_g is the transformed constant, and Δd_c and Δd_t can be estimated using city and year dummies. Identification of γ comes from within-city time variation in $\Delta R_{c,t}$. In other words, the model is identified by the differential evolution of rent across groups defined by city-gender (i.e., Gender/City/Time variation). γ identifies how wages respond to changes in rents by comparing the responses between genders across cities that had different evolutions of rents. This estimation strategy is effectively a triple difference-in-differences estimator, and assumes that the three way interaction of gender-city-year effects, $d_{c,g,t}$, are validly excluded (uncorrelated to other regressors) from equation (3.29). For (3.32) to yield consistent estimates of γ, we require that city-year shocks that impact males and females differently are uncorrelated with $\Delta R_{c,t}$ profiles.

Equation (3.32) is our empirical specification of the relationship given by equation (3.28) where $\Delta R_{c,t}$ corresponds to the value of B_c in each period. As noted earlier, the values of γ in general depends on the arrival rate of jobs, Q, which will also vary by city. To capture this dependence, our final empirical specification also conditions on city employment rates, $E_{c,t}$:

$$\Delta w_{c,t} = Z'_{c,t}\Delta\beta + \gamma \cdot B_{c,t} + \varphi E_{c,t} + \Delta d_g + \Delta d_c + \Delta d_t + \Delta\bar{\epsilon}_{c,t}, \tag{3.33}$$

where $B_{c,t}$ is the 'bargaining' variable (gender difference in city rents).

To implement this procedure, we estimate (3.30) separately by year and include a full set of city dummies as well as female-city interactions. The coefficient on the female-city interactions are used as a dependent variable in estimation of equation (3.33). In all estimates below, we weight observations by the inverse of the estimated standard error on the coefficients on city-female dummies from the first stage regressions. The vector $X'_{i,g,c,t}$ includes controls for education, potential experience and its square, black and hispanic indicators, 12 occupation indicators, and 144 industry indicators. Controlling for industrial composition is crucial for our interpretation of γ, as we wish to examine the impact of $B_{c,t}$ after accounting for composition effects. In controlling for education and occupation, we hope to eliminate any relationship that might arise through the possibility that men and women perform different tasks within industries.

The main covariate of interest is $B_{ct} \equiv \Delta R_{c,t} = \sum_j(\pi^M_{j,c,t} - \pi^F_{j,c,t}) \cdot \xi_{j,t}$, where $\xi_{j,t}$ are national level industrial premia and $\pi^g_{j,c,t}$ are employment shares for industry j in city c in year t for $g \in \{M, F\}$. We estimate the $\xi_{j,t}$s as the coefficients on industry dummy variables in regressions run separately by year at the individual level for the entire country, using the same set of covariates $X'_{i,g,c,t}$ that are described above and controlling for industry-female, city, and city-female interactions.

3.4 Data and Descriptive Statistics

The data for our main empirical analysis comes from the U.S. Census Public Use Micro-Samples (PUMS) for the years 1970, 1980, 1990, and 2000. These data provide labor market and earnings information for the previous calendar year. Our analysis focuses on civilian wage and salary workers residing in a metropolitan area and between the ages of 18 and 65 at the time of the Census, with positive annual hours of work and positive wage and salary earnings in the previous year, and with at least one year of potential experience. We construct an hourly wage measure by dividing annual wage and salary earnings by annual hours worked and use only observations with hourly wages greater then 1 and less than 250 1980 dollars. An important part of our analysis is having consistent coding of metropolitan areas, years of education and completed schooling, and industrial classification across the four Census years. Details on how we construct these variables are left to Appendix A.

Due to extremely large sample sizes, much of our analysis is carried out at the city level. To do this, we construct a panel of cities across the four census years. Between 1970 and 2000, we are able to consistently identify 152 cities.

We supplement our empirical work with data from the March Current Population Surveys (CPSs) for the years 1970 - 2007. We use CPS data to analyse movements in relative wages over time, and employ the same selection criteria as with the U.S. Census data described above.

3.5 Results

3.5.1 Motivation

Before turning to our main empirical exercise, we first illustrate how the distribution of male employment across industries can affect the gender gap over and above pure compositional effects. Consider the regression equation:

$$\Delta w_{c,t} = I'_{c,t}\beta_1 + H'_{c,t}\beta_2 + \mu_c + \tau_t + e_{c,t}, \tag{3.34}$$

where $\Delta w_{c,t}$ is the city-level gender gap after controlling for human capital and industrial composition effects as described above, $I'_{c,t}$ is a vector of variables indicating the industrial composition of cities, $H'_{c,t}$ is a vector of city-level variables that potentially affect the gender gap, μ_c and τ_t are city and period fixed-effects, and $e_{c,t}$ is a disturbance term assumed to be uncorrelated with the other variables in the model.

The vector of industrial composition variables, $I'_{c,t}$, describes the structure of employment between high- and low-paying jobs for both men and women. To identify which industries are good- or high-paying industries, we run a regression at the individual level using observations from the entire nation in 1980 and control for education, potential experience and its square as well as indicators for race and gender. We then take the mean residual by industry from this regression and sort industries into quintiles based on an industry's mean residual. High paying industries are industries that were in the top quintile of the residual wage distribution in 1980. We then calculate the employment share of men and women in each quintile for each year and city, which form elements of $I'_{c,t}$.[65]

Table (3.1) contains the results from this exercise when $I'_{c,t}$ includes only men's share of employment in the industries identified to be in the top quintile in 1980. Row (1) columns (1-4) show the estimated coefficient on male employment in the top industries under various identification assumptions. In column (1), both period and city fixed effects are excluded, in column (2) and (3) we added either period or city fixed effects and in column (4) we add both. In all specifications, there is a strong positive relationship between male employment in high-paying industries and the city gender gap, although the magnitudes differ. For instance, a ten-point increase in the fraction of men employed in the top industries is associated with a relative increase in men's wages of 1.3%, according to column (4). It should be emphasized that the gender gap is calculated after netting out industrial composition effects so that this relationship is not being driven by mechanical forces.

In order to illustrate this point more clearly, consider columns (5-8) of Table (3.1). In these columns, the dependent variable is gender gap calculated the exact same way as described above, except we only include observations on workers employed *outside* the top quintile of jobs. In this case, there will be no composition effects of employment in the top ranked industries and the gender gap in the rest of the industries. However, as indicated in row (1), the relationship between male employment in high-paying industries is still significant and positively correlated to the city-level gender gap in the rest of the industries, as would be predicted by our search and matching model.[66]

While this analysis captures the idea that the concentration of male employment in high paying industries can have impacts on the gender gap that are above and beyond pure composition effects, the model we presented in section (2) emphasized that the what should matter for the gender gap is the entire distribution of male and female employment. Before turning to analysis more closely aligned to the theory, it is useful to consider again the estimation of equation (3.34) where $I'_{c,t}$ contains other aspects of the industrial composition aside from employment in the top ranked industries.

Table (3.2) contains the results from this estimation when $I'_{c,t}$ includes male and female employment shares in the top four quintiles of jobs. Since the share of employment

[65]Allowing the high-paying jobs do differ across years, instead of using the 1980 classification as above, does not significantly affect the results.

[66]The relationship between the proportion of males in top jobs and the gender gap in the rest of the industries is virtually unchanged when we adjust for the industrial composition in these industries or not, as would be expected since these estimates should be free from pure composition effects. It is also worth noting that the estimates in column (8) and column (4) are very similar, which gives some confidence in the effectiveness the adjustment for industrial composition effects by including a large number of industry dummies in the first stage.

across all quintiles sums to one, the bottom quintile of jobs should be thought of as the omitted category and the estimated parameters on the other shares interpreted in reference to this omitted group. All specifications include period and city fixed-effects. In addition, we also consider the effect of other city-level variables on the gender gap. Column (1) includes male employment shares in different categories of jobs. The top three categories are individually significant as well all four categories being highly jointly significant. Note the pattern of the estimated coefficients on these shares; employment share in higher ranked industries are associated with larger gender gaps. Columns (2) - (4) consider including various city level characteristics, such as male and female employment rates, the log ratio of men to women in the local economy, and various demographic characteristics.[67] Most of these additional regressors do not come in significantly and, even more importantly, they have little impact on the magnitude or relative magnitudes of the coefficients on male employment shares; that is, they do not alter the conclusion that male employment in higher paying industries is associated with larger gender gaps. The theoretical model outlined above suggests that what is important is the differences in male and female employment across industries. To shed some light on this importance, columns (5-8) include female employment shares in addition to male employment shares. These columns indicate the same relationship for male employment shares and a symmetric relationship between female employment shares and city gender pay gaps. That is, female employment more heavily concentrated in higher paying industries is associated with smaller wage gaps.

While the analysis in this section is certainly consistent with the theoretical model presented in section (2), there are several limitations. First, the grouping of industries into quintiles based on 1980 wage residuals is admittedly arbitrary. In addition, the magnitudes of the coefficients are not easily interpretable in terms of general compositional shifts because they are in reference to an omitted category, holding the rest of the distribution constant. Lastly, the OLS estimate of relationship between the employment shares and the gender gap is not necessarily causal and any attempt to overcome the potential endogeneity of industrial composition is complicated by the fact that it would require eight valid exclusion restrictions. In light of these problems, below we focus on the bargaining variable $B_{c,t}$ defined above.

3.5.2 Empirical Evidence

Returning again to our bargaining model which emphasizes that what is important is the relative outside options of males and females as captured by the bargaining variable B_c.

[67] In columns (4) and (8), in addition to the coefficients listed in the table, we control for the percent of immigrants, black, hispanic origin, college graduates, some post-secondary, and high school graduates in a city. Most of these coefficients are not significant and are not reported to save space.

In this section we wish to more formally test our theory by estimating equation (3.33). Recall that our bargaining model predicts that $\gamma > 0$, which would represent general equilibrium effects from the relative distribution of male and female employment across industries. OLS estimates of equation (3.33) are contained in table (3.3). All estimates reported in Table (3.3) contain a full set of unrestricted city and year dummy variables, although we do not list the corresponding coefficients to save space.

The first column of Table (3.3) displays the OLS estimate without additional city level regressors, Z_{ct}. The coefficient on the bargaining variable, B_c, is 0.43 and is statistically significant from zero at the one percent level. If OLS provides consistent estimates of γ, this results is supportive or our search and matching model which predicts that when men's outside options are high relative to women's, relative wages will also be higher. This estimate implies that the standard accounting or Blinder-Oaxaca approach used to decompose male-female wage differentials into 'explained' and 'unexplained' components will underestimate the contribution of different industrial employment patterns across genders to the explained portion of the gender gap.

When estimating and interpreting γ we want to make sure that what we are capturing is the effect of a pure compositional shift in relative employment patterns between males and females. Thus, we want to make sure that we are accounting for different employment rates across local labor markets, as suggested in section 3.2. Column (2) of Table 3.3 adds to the regression the city employment rate, E_{ct}. Including this regressor increases the estimate of γ to 0.54, implying that gender differences in rent and employment rates are negatively correlated.

Upon further inspection of the data, it was evident that this correlation is particularly strong among a number of "Rust Belt" cities.[68] To make sure our results are not sensitive to the inclusion of these cities, in columns (3) and (4) we re-estimate (3.33) including a dummy variable for 1980 rust belt cities, column (3), and excluding these cities altogether, column (4). In both cases, the estimate of γ increases further implying that our results are not being driven by this subset of cities. We also want to make sure that different regional trends, such as the decline in manufacturing in the North West and the increase in manufacturing in the South or differential female employment rates by region, are not biasing our results. In column (5) we include interactions between Census division dummies and year dummies, which again increases the estimate of γ. As a final specification check, we allow male and female rents to have separate effects on relative wages, which relaxes the assumption that γ is the same for both sexes. Column (6) contains the results and indicates that the assumption of a common γ is not overly restrictive, as the estimated coefficient on each variable is roughly the same magnitude

[68] We identify 'rust belt' cities as those in table I of Feyrer, Sacerdote, and Stern (2007) in their study of outcomes of cities who excessively lost steel and auto jobs in the early 1980s.

with opposite signs.

3.5.3 Addressing Potential Endogeneity

Since adding additional regressors in the previous section tended to increase the estimate of γ, this might be viewed as evidence of a potential correlation of $B_{c,t}$ with the error term. We address the potential endogeneity problem by an using instrumental variables approach. The inclusion of city fixed-effects removes any bias arising from time-invariant city factors that differentially affect males and females. However, OLS estimates will still be biased if the B_c variable is contemporaneously correlated to the error term arising from time-varying omitted variables or simultaneity. For example, city-year shocks that differentially affect males and females that also influence average city rents.

To overcome this problem we require an instrument that predicts $B_{c,t}$ that is not correlated with current city-year shocks to a particular gender. Several observers have noted that trade pressures in the 1980s may have lead to substantial loss of rents in trade impacted industries (Card and DiNardo 2002, Borjas and Ramey 1995). In this case, we would expect that cities with greater employment in tradable goods sectors lost rents relative to cities with lower employment in the tradable goods sectors. This is essentially the variation used in Borjas and Ramey (1995) who analyze the impact of employment in highly concentrated industries on the college premium. In this paper, we use variation in national trade patterns weighted by the local importance of tradable sectors to a particular gender in the local economy to predict average rents for each gender. More accurately, since we include city fixed effects, we expect that in years where males in a given city are more impacted by trade relative to the city's average to predict negative deviations in male rents from a city's average.

To create this instrument, we use national level trade data obtained from NBER's data library.[69] We download data on imports and exports by SIC industrial codes for each of our Census years and for each industry and year, we calculate net imports.[70] We assign these net imports to each city based on a city's *total* employment in that industry in 1970 and multiply this number by the fraction of a particular gender employed in that industry in 1970 *nationally*. We then sum over industries for each city to create an index

[69] We download international trade data on exports from http://cid.econ.ucdavis.edu/data/sasstata/usxss.html and imports from http://cid.econ.ucdavis.edu/data/sasstata/usiss.html. These data contain information on the dollar amount of imports or exports for detailed SIC codings from 1972 to 2001.

[70] We use trade data for 1972 for the 1970 Census, since the trade data is unavailable in earlier years. For the other Census years, we use trade data for the year prior to the Census to correspond to the wage and employment data from the Census. We convert all dollar amounts of trade into 1982 prices based on a CPI deflater. We code the SIC codes to correspond to the `ind1950` variable available from IPUMs. The trade data and code to produce it is available from the authors upon request.

of the impact of trade for each city and gender. Thus, for each gender, this instrument is:

$$IV_{g,c,t}^{NI} = \sum_j \pi_{c,j,1970}^T \cdot \theta_{j,1970}^g \cdot NI_{j,t}, \tag{3.35}$$

where $\pi_{c,j,1970}^T$ denotes the share of total employment in city c in industry j in 1970, $\theta_{j,1970}^g$ is the fraction of gender g employed in industry j nationally in 1970, and $NI_{j,t}$ is the national level net imports in industry j. In effect, for each city and gender, this instrument is an index of 'sensitivity' to trade. We expect high values of this instrument to be associated with low values of average city rent for each gender. Each city-year observation has two instruments: one for each gender. We let these two variables have an independent effect on $B_{c,t}$ in the first-stage which has two main benefits: (1) it allows for greater precision, and (2) it leaves an over-identifying exclusion restriction. This latter benefit will enable us to test our identifying assumption that the distribution of total employment across cities in 1970 is uncorrelated to current city-gender shocks to the local economy.

Table (3.4) contains the results from this exercise. As argued earlier, the employment rate should directly affect the gender wage gap. We can either control for this measure or assume that the instrumented $B_{c,t}$ variable is uncorrelated with employment rates. We pursue both approaches below. Starting with column (1), we exclude the employment rate. The first-stage estimates show that our instruments are good predictors of differences in rent between genders and both instruments have the predicted sign. The F-statistic on the test that our instruments are jointly equal to zero is 15.67 with a p-value of 0.0, easily passing conventional 'rule of thumb' levels of significance. Figure 3.3 shows the resulting first-stage for the male version of the instrument, $IV_{M,c,t}^{NI}$. As predicted, high values of this instrument are associated with lower values of gender difference in rents, $B_{c,t}$. Figure 3.4 shows the first stage for the female version of the instrument. Again, this instrument has the predicted sign and is strongly, positively correlated to difference in rents.

Column (1) of Table 3.4 contains the second-stage estimates using the net import instruments and shows an estimated γ of 1.26 and is statistically significant from zero at conventional significance levels. Column (2) adds to the regression the employment rate. While this regressor is statistically significant, its inclusion has little effect on the estimated γ. It should be noted that the over-identification test, reported in the last row of table 3.4, indicates that we cannot reject the over-identifying restrictions at any reasonable level, which adds confidence to the validity of our instruments.

The instrumental variables estimates using the net import instruments are considerably larger than the OLS estimates presented in the previous section, although the standard error is also considerably larger and so we are unable to statistically distinguish the IV and OLS point estimates. Nevertheless, there are reasons to believe that the IV estimates would be higher in the case. One potential reason could be that $B_{c,t}$ suffers from classical measurement error and the fixed-effect estimates suffer from attenuation bias. A second potential reason could be that the IV estimates estimate a local average treatment effect (LATE) (Imbens and Angrist 1994). For example, the net import instruments use mainly variation from manufacturing industries (tradabe goods) which employ disproportionately low-skilled workers. To the extent that bargaining is more important for low-skilled workers, variation stemming from manufacturing may have higher impacts than shifts in the composition of employment generally. In a later section, we provide evidence that suggests that bargaining is indeed more important for low-skilled workers. Yet another explanation may be that simultaneity or omitted time-varying variables correlated with gender difference in rent biases the OLS coefficient downward. The fact that adding additional regressors to the OLS estimates tends to increase the magnitude of γ, makes this explanation seem reasonable.

The remaining two columns of Table (3.4) provide two additional robustness checks. The net import instruments use city-level employment shares in 1970 as weights in the construction of the instruments. Despite the fact that the over-identification tests provide support for our assumption that the distribution of employment across industries in 1970 is not correlated with current city-year shocks, one may wonder if this is true for 1970. Thus, in column (3) we use only data from 1980 to 2000 in the estimation. The point estimate is larger in this case but it is also less precisely estimated. Nevertheless, our results are not sensitive to the inclusion of 1970. Thus far, we have treated the employment rate as though it were exogenous. In column (4) we address this by instrumenting the employment rate and B_c using our net import instruments. While our net import instruments are not strong predictors of employment rates, the fact that the estimates of our coefficient of interest, γ, are not sensitive to how we handle the employment rate variable is reassuring.

3.5.4 Robustness

In this section we assess the robustness of our results by pursuing several alternative identification strategies and investigating heterogeneous responses by education groups.

Alternative Instruments

As a first robustness check, we construct instruments used in Beaudry, Green, and Sand (2008) who investigate the same type of general equilibrium mechanism we study here in a different context. They note that the change in rent for a gender can be decomposed over time:

$$\sum_j \pi^G_{c,j,t}\xi_{j,t} - \sum_j \pi^G_{c,j,t-1}\xi_{j,t-1} = \sum_j (\pi^G_{j,c,t} - \pi^G_{j,c,t-1}) \cdot \xi_{j,t} + \sum_j \pi^G_{j,c,t-1} \cdot (\xi_{j,t} - \xi_{j,t-1}). \tag{3.36}$$

Using variation from each of these components, Beaudry, Green, and Sand (2008) create instruments:

$$IV1^G = \sum_j \xi_{j,t-1} \cdot (\hat{\pi}^G_{c,j,t} - \pi^G_{c,j,t}) \quad \text{and} \quad IV2^G = \sum_j \pi^G_{c,j,t-1} \cdot (\xi_{j,t} - \xi_{j,t-1}), \tag{3.37}$$

where $\hat{\pi}^G_{c,j,t} = \frac{\hat{L}^G_{c,j,t}}{\sum_j \hat{L}^G_{g,c,j,t}}$ and $\hat{L}^G_{c,j,t} = L^G_{j,c,t-1}\left(\frac{L^G_{j,t}}{L^G_{j,t-1}}\right)$ and $L^G_{c,j,t}$ is the number of workers of gender G in city c in year t working in industry j and $L^G_{j,t}$ is the national level counter part. Thus, we use national growth in industry j to predict city level growth and construct predicted industry shares in city c to form our instruments. The result is that our instruments are only a function of initial industry shares and national level trends that are presumably orthogonal to city-year shocks. To use these instruments, we estimate equation (3.33) in first-differences rather than fixed-effects to eliminate the city-specific component of the error term. Working in first-differences also allows a construction of an additional instrument for the change employment rate that has been used, for example, by Blanchard and Katz (1992), that also works off of national level industry growth. This instrument is weighted average of national growth in each industry where the weights are the start of the period employment shares for each industry:

$$IV^E_{c,t} = \sum_j \pi^T_{c,j,t-1} \cdot g_{j,t}, \tag{3.38}$$

where $\pi^T_{c,j,t-1}$ is the share of total employment in industry j in city c at the start of the period and $g_{j,t}$ is the national growth in industry j between years t and $t-1$.

We present the first-difference results in Table 3.5. The first four columns of the table use data from all three decades in our sample, while the last four use only data from the 1980s. The first column shows the first-differenced OLS estimate of γ of about 0.41 and is similar to the OLS results obtained by using city-fixed effects. Columns (2-4) show the instrumental variables results using our alternative set of instruments described above. Column (2) shows the estimate from using $IV1^M_{c,t}$ and $IV1^F_{c,t}$ as will as the employment instrument. While the instruments perform well in the first-stage, they are unable to identify the coefficients on either $B_{c,t}$ or $\Delta E_{c,t}$. Column (2) uses $IV2^M_{c,t}$ and $IV2^F_{c,t}$ as well as the employment rate instrument, and, again, the first-stage performs well. The estimated γ is 1.46 and significant at the ten percent level and is now comparable in magnitude to earlier IV estimates using the net import instrument. Column (4) combines all of the instruments used in column (2) and (3) and the estimated coefficient on $\Delta B_{c,t}$ is 0.76 and significant at the five percent level.

Since most of the identification is coming from the 1980s, the decade with the largest decline in the gender gap and gender differences in rent, in the next four columns of Table (3.5) we focus solely on this decade. Column (5) shows that the OLS estimate is slightly larger in magnitude than its counter-part in column (1) that uses all years. The next three columns show the IV results in the same format as columns (2-4). The IV estimates are better defined when focusing on this decade and estimates of γ rage from 1.2 - 1.6, which are again comparable to earlier IV estimates. We conclude this section by noting that our results are robust to alternative instrument strategies.

Heterogeneous Responses

Our model conceptually applies to male and female workers of homogeneous skills. In the previous sections, we hold skill constant by controlling for education and potential experience in a regression framework. In this section, we take the alternative strategy of stratifying workers by education group. In particular, we divide the sample of workers into low education (high school or less) and high education (college degree or more) groups. We precede in the same way as before by calculating city-level adjusted wage gaps controlling for age, race, occupation and industry for each education group. We also calculate the industrial premia separately by education group in the same way as describe above so that the $B_{c,t}$ is specific to education group.

Table 3.6 contains the results broken down by education group. The first two columns display the results for the high school or less education group, while the las two contain the results for the college degree or more group. In all regressions, we control for unrestricted year and city fixed-effects and weight each observation by the inverse of the sampling distribution of the dependent variable (the adjusted city-level gender gap). In column (1) the OLS estimate of γ for the low education group is about 0.7, which is considerably higher than the corresponding estimate that pools all workers. Column (2) shows the IV estimate of γ using the net import instruments is about 2.0, which is also higher than the pooled sample.

The next two columns display the results for the college degree or more group. The OLS estimate of γ, column (3), is wrong signed and insignificant. The next column shows the IV estimate using the net import instruments. Unfortunately, the first-stage performs poorly and the second-stage estimate is wrong signed and insignificant. Overall, the results indicate that bargaining mechanism emphasized in the model is likely more appropriate for lower skilled groups.

Alternative Identification Strategy

Up to now, we have been using within-city over-time variation to identify the coefficient on $B_{c,t}$, instead of between-city variation, out of concern for the potential that there are city-specific factors that influence both the gender gap and the bargaining variable. For example, if a city is particularly discriminating against women they may pay them less in a given job and also prevent them from having access to certain industries, confounding estimates based on cross-city variation. However, not using between-city variation also throws out variation in the data that ideally would be helpful in identifying γ. Dividing the sample into two education groups, as in the above section, may provide a way of addressing both concerns.

Consider the regression equation for a given Census cross-section:

$$\Delta w_{s,c} = \gamma \cdot B_{s,c} + \beta \Delta E_{s,c} + \xi_c + \xi_s + \epsilon_{s,c}, \tag{3.39}$$

where $\Delta w_{s,c}$ is the adjusted gender gap for skill group $s \in \{HS, BA\}$ in city c, ξ_c are city-specific fixed-effects, ξ_s is a skill-group fixed-effect, and $\epsilon_{s,c}$ is an idiosyncratic error term. Including the city-wide fixed-effect will absorb any city-specific factors correlated with $B_{s,c}$ and the gender wage gap that are common between skill groups. For instance, suppose, as in the example above, a city is particularly discriminating against women in pay and access to industries. As long as this discrimination against women is the same for high- and low-educated women, then the inclusion of city dummy variables will account for this. With city fixed-effects included, γ is identified from covariation between relative gender wage gaps in different still groups and relative differences in the outside option of men and women as captured by $B_{s,c}$.

Table (3.7) contains the results from estimates of equation (3.39) for each of the four cross-sections of Census data. In column (1), the estimate from 1970 is shown to be slightly negative and not significantly different from zero. In columns (2-4), estimates of γ rage from 0.5 - 1.4, which is very close to the range of estimates using instrumental variables earlier, and slightly higher than using the within-city variation.

One reason for the odd estimate in 1970 is that this specification assumes that the γ coefficient is common between skill groups; an assumption that is unlikely to be valid given the heterogeneous responses estimated in the previous section. In Table (3.8), we include an interaction between the high-educated skill group and the bargaining variable $B_{s,c}$. Estimates using the interaction generally indicate that net effect of $B_{s,c}$ for highly educated workers is zero, while it is between 0.8 and 1.7 for low-skill workers. An interesting observation from this analysis is that while combining all years produces an estimate of γ that is about the same as instrumental variable estimates, the trend over time seems to indicate that bargaining is becoming more important. That is, larger estimates of γ are obtained in latter years. In the next sub-section, we briefly comment on a potential explanation for this observation.

3.5.5 Potential Selectivity Bias

Over the time period of interest, the fraction of women in the labor force grew substantially, from about 50 percent in 1970 to nearly 70 percent in 2000 - almost all of the increase coming before 1990. One point of interest is to gauge whether the changing participation of women in paid employment over the time period we study could have confounding effects on the analysis above. Instead of attempting to estimate or control for selection bias in wage equations, we present a simple model of selection which draws heavily upon Card (2001) that suggests that the effects of selection bias, if anything, are working against finding evidence of bargaining. Thus, we view the estimates presented elsewhere in this paper as conservative in the presence of significant selection bias.

To see this, following Card (2001), we write log wages for individual i of gender g in city c as:

$$w_{i,g,c} = w_{g,c} + \zeta_{i,g,c}, \tag{3.40}$$

where $w_{g,c}$ is the average log wage for gender g in city c, and $\zeta_{i,g,c}$ is an individual's deviation from this average that is assumed to be normally distributed with a standard deviation of $\sigma(\zeta)$. Assume further that an individual's employment status is determined by a latent index model:

$$H^*_{i,g,c} = d_g + \chi \cdot \zeta_{i,g,c} + \nu_{i,g,c} = d_g + \tilde{\nu}_{i,g,c}, \tag{3.41}$$

where d_g is a gender specific factor and $\nu_{i,g,c}$ is another normally distributed disturbance term. A non-negative value of this index means that i is observed working. Under these assumptions, we can write the unconditional mean log wage for gender g in city c as:

$$\mathbb{E}\left[w_{i,g,c} | H^*_{i,g,c} \geq 0\right] = w_{g,c} + \rho \cdot \sigma(\zeta) \cdot \lambda(p_{g,c}), \tag{3.42}$$

where ρ is the correlation between the normally distributed error terms in equations (3.40) and (3.41), $\lambda(\cdot)$ is the inverse Mill's ratio, and $p_{g,c}$ is the employment rate of gender g in city c. The last term of the above expression, $\rho \cdot \sigma(\zeta) \cdot \lambda(p_{g,c})$, is positive and decreasing in $p_{g,c}$ as long as workers are positively selected from the population - that is, we assume

$\rho > 0$.[71]

Under the above assumptions, and noting that $\lambda(\cdot)$ is approximately linear over most of its range, the selectivity bias is an approximately linear function of a groups employment rate (Card 2001). For $0.4 \leq p_{g,c} \leq 0.99$, $\lambda(p_{g,c}) \approx 1.5 + 1.5 \cdot p_{g,c}$. Assuming that $\sigma(\zeta)$ and ρ are approximately equal between genders, we can write the gender difference in selection bias for a city as a function of the gender difference in employment rates.:

$$\Delta\text{Bias}_c = G(p_{M,c} - p_{F,c}) < 0 \quad \text{if} \quad p_{M,c} - p_{F,c} > 0, \tag{3.43}$$

where $G(\cdot)$ is negative and a decreasing function of its argument. That is, under these assumptions, the greater the difference in participation rates between men and women, the greater the amount of downward bias in observed gender wage gaps. For a given observed gender wage gap in a city, then, we under estimate the unconditional (on participation) gender wage gap. Any correlation between a gender's rent and employment rate would be expected to be positive, implying a downward bias on γ when estimating a relationship such as equation (3.33) using observed wage data. Given the observed pattern of employment and $B_{c,t}$, we find it likely that this bias is more important in earlier Census years.

3.6 Conclusion

This paper presents and tests the hypothesis that differences in employment of men and women across industries has impacts on the gender gap beyond pure mechanical composition effects. We present time-series evidence using CPS data from 1970 to 2007 that show that the largest declines in the gender gap coincided with declines in men's relative representation in high-paying industries. We then present a search and matching model that rationalizes this link, and test it using a panel of U.S. cities from 1970 to 2000 constructed from the U.S. decennial Censuses. Our theoretical framework suggest that cities where men's employment is relatively more heavily weighted in high-paying industries will have a greater gender wage gap, even after controlling for composition effects. The mechanism underlying these general equilibrium effects in our model is bargaining: the greater men's relative chance of employment in a high paying industry the greater their outside option, and, thus, this places upward pressure on men's wages relative to women's in a bargaining framework. We show that our approach is closely related to the common Blinder-Oaxaca type decompositions and that if bargaining is important that these decompositions over estimate the 'unexplained' portion of the gender gap.

[71] Blau and Kahn (2007) suggest that this is the case for women over the time period we study.

We find that there are significant general equilibrium effects or spill-overs from relative male-female employment across industries. Using our data on U.S. cities, we show that the total effect of relative representation across industries is about twice what the standard accounting procedures attribute to industrial composition. We estimate that over the 1980s, changing relative employment of genders across industries can account for almost 40% of the change in the gender pay gap for full-time full-year workers in our sample of U.S. cities.

Figures

Figure 3.1: Trends in Male-Female Wage Gap: CPS 1970-2007

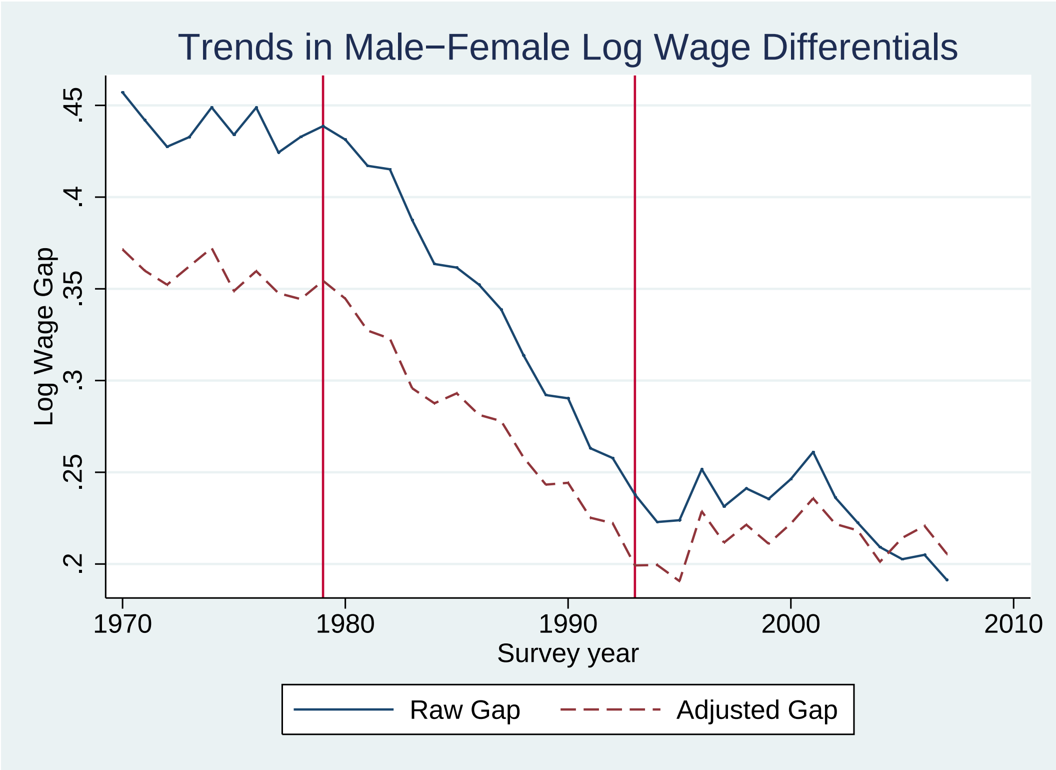

Figure 3.2: Male-Female Wage Gap and Industrial Composition

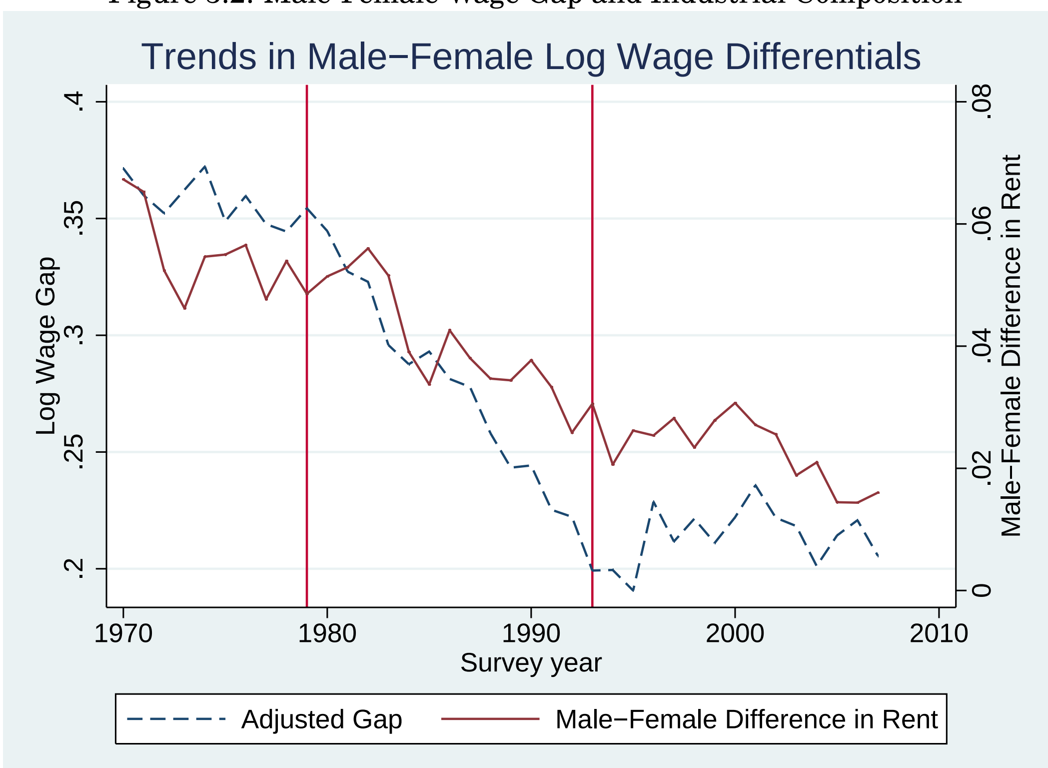

Figure 3.3: First-Stage: Male Net Imports IV

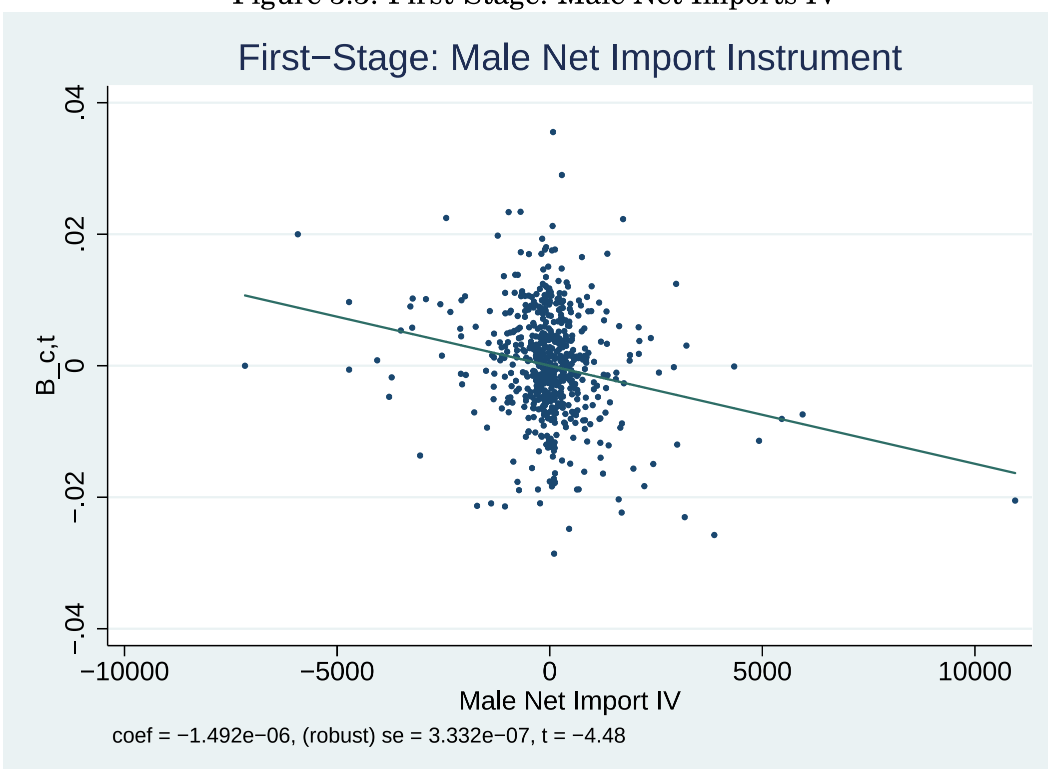

Figure 3.4: First-Stage: Female Net Imports IV

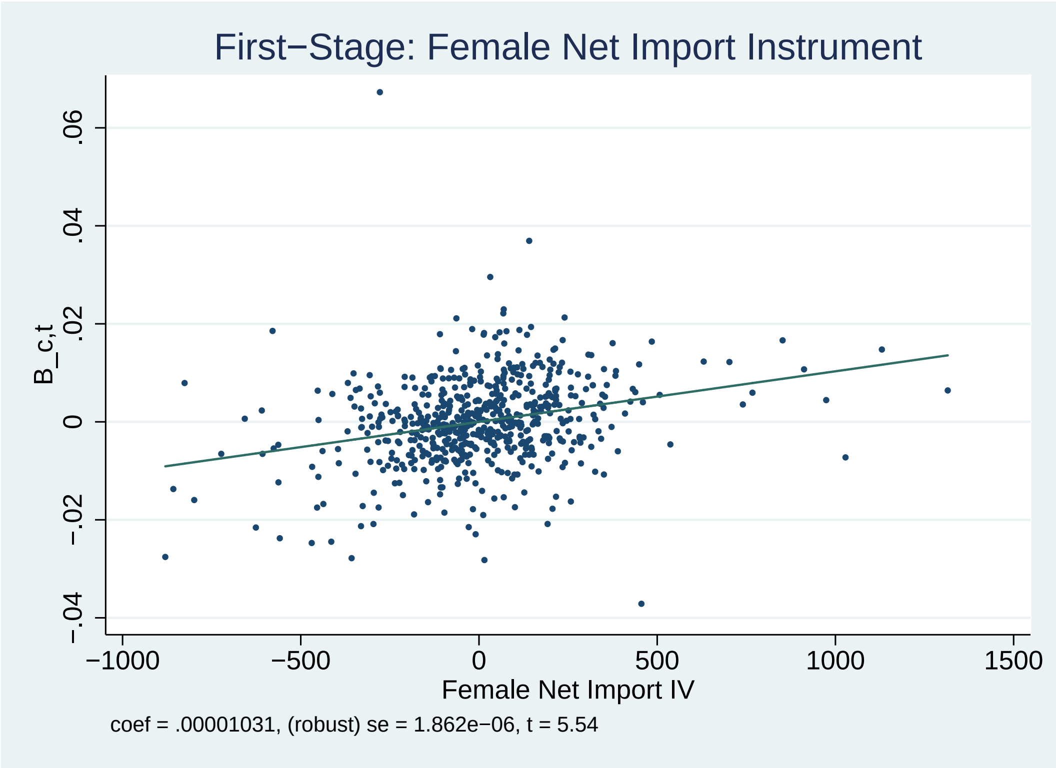

Tables

Table 3.1: The Effect of Male Employment in High Paying Industries on the Controlled Gender Gap

	All Workers				Not Top Workers			
Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
% Male Emp. in Top Quintile	0.564	0.315	1.292	0.149	0.557	0.293	1.346	0.18
	(0.038)*	(0.03)*	(0.063)*	(0.039)*	(0.039)*	(0.03)*	(0.073)*	(0.049)*
Unempr. Rate	-.748	-.203	-.902	-.302	-.861	-.248	-1.019	-.302
	(0.159)*	(0.127)	(0.148)*	(0.088)*	(0.166)*	(0.129)	(0.164)*	(0.11)*
Year	No	Yes	No	Yes	No	Yes	No	Yes
City	No	No	Yes	Yes	No	No	Yes	Yes
Obs.	608	608	608	608	608	608	608	608
R^2	0.305	0.674	0.692	0.925	0.283	0.662	0.68	0.908

Notes: Standard errors are in parentheses. Stars (*) denote significance at the 5% level. The observation is at the city level and weighted by the inverse of the sampling variance of the dependent variable.

Table 3.2: The Effect of Composition of Employment on the Controlled Gender Gap

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
% Male Emp. in Quintile 2	0.195	0.212	0.23	0.049	0.279	0.289	0.301	0.13
	(0.124)	(0.127)	(0.128)	(0.121)	(0.128)*	(0.13)*	(0.133)*	(0.131)
% Male Emp. in Quintile 3	0.391	0.399	0.405	0.267	0.302	0.314	0.323	0.259
	(0.13)*	(0.131)*	(0.13)*	(0.132)*	(0.139)*	(0.139)*	(0.141)*	(0.137)
% Male Emp. in Quintile 4	0.486	0.492	0.478	0.324	0.528	0.534	0.523	0.409
	(0.099)*	(0.103)*	(0.101)*	(0.101)*	(0.107)*	(0.109)*	(0.108)*	(0.109)*
% Male Emp. in Top Quintile	0.476	0.479	0.448	0.327	0.56	0.562	0.542	0.463
	(0.103)*	(0.106)*	(0.104)*	(0.106)*	(0.11)*	(0.112)*	(0.11)*	(0.111)*
% Female Emp. in Quintile 2					-.130	-.130	-.130	-.157
					(0.083)	(0.083)	(0.083)	(0.086)
% Female Emp. in Quintile 3					0.108	0.101	0.093	-.006
					(0.089)	(0.087)	(0.088)	(0.094)
% Female Emp. in Quintile 4					-.081	-.079	-.079	-.134
					(0.081)	(0.081)	(0.082)	(0.084)
% Female Emp. in Top Quintile					-.246	-.229	-.236	-.319
					(0.106)*	(0.105)*	(0.105)*	(0.106)*
Female Emp. Rate		-.084				-.035		
		(0.058)				(0.057)		
Male Emp. Rate		-.041				-.038		
		(0.099)				(0.096)		
$\log \frac{\text{Male}}{\text{Female}}$			0.04				0.025	
			(0.02)*				(0.02)	
% of Female Workforce				-.140				-.118
				(0.07)*				(0.072)
$\frac{\text{Year}}{\text{City}}$								
Obs.	608	608	608	608	608	608	608	608
R^2	0.932	0.933	0.933	0.937	0.935	0.935	0.935	0.939

Notes: Standard errors are in parentheses. Stars (*) denote significance at the 5% level. The observation is at the city level and weighted by the inverse of the sampling variance of the dependent variable.

Table 3.3: The Effect of Bargaining on Relative Wages of Males and Females

Variables	(1)	(2)	(3)	(4)	(5)	(6)
$B_{c,t}$	0.436	0.541	0.600	0.651	0.745	
	(0.122)*	(0.125)*	(0.13)*	(0.153)*	(0.126)*	
Employment Rate		0.346	0.300	0.261	0.561	0.355
		(0.104)*	(0.102)*	(0.107)*	(0.118)*	(0.102)*
Rust Belt dummy			-.015			
			(0.007)*			
$\sum_i \pi^M_{i,c,t} \cdot \xi_{i,t}$						0.594
						(0.147)*
$\sum_i \pi^F_{i,c,t} \cdot \xi_{i,t}$						-.482
						(0.159)*
Obs.	608	608	608	552	608	608
R^2	0.916	0.918	0.918	0.914	0.932	0.918

Notes: Standard errors are in parentheses. Stars (*) denote significance at the 5% level. The observation is at the city level and weighted by the inverse of the sampling variance of the dependent variable.

Table 3.4: Instrumental Variables Results

	IV	IV	IV80-00	IV
Variables	(1)	(2)	(3)	(4)
$B_{c,t}$	1.263	1.236	1.936	1.260
	(0.426)*	(0.397)*	(0.768)*	(0.425)*
Employment Rate		0.492	0.643	-.016
		(0.127)*	(0.188)*	(0.488)
Obs.	608	608	456	608
F-stat $B_{c,t}$	15.67	17.897	6.81	15.67
F-stat $E_{c,t}$				4.46
Over ID P-val	0.974	0.162	0.999	

Notes: Standard errors are in parentheses. Stars (*) denote significance at the 5% level. The observation is at the city level and weighted by the inverse of the sampling variance of the dependent variable. IV results use the Net Imports instrument described in the text.

Table 3.5: Alternative IV Results

	All Years				1980s			
	OLS	IV1	IV2	IV1&IV2	OLS	IV1	IV2	IV1&IV2
Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta B_{c,t}$	0.412	-.979	1.462	0.758	0.917	1.205	1.658	1.353
	(0.123)*	(1.748)	(0.828)	(0.371)*	(0.237)*	(0.526)*	(0.523)*	(0.414)*
$\Delta E_{c,t}$	0.472	-2.947	1.954	1.191	0.533	2.441	2.373	2.335
	(0.1)*	(4.448)	(0.966)*	(0.504)*	(0.157)*	(0.621)*	(0.672)*	(0.6)*
Obs.	456	456	456	456	152	152	152	152
R^2	0.228				0.138			
F-Stat $B_{c,t}$		26.946	58.346	81.456		23.1	19.797	22.208
Over ID P-val		0.868	0.349	0.153		0.136	0.208	0.304

Notes: Standard errors are in parentheses. Stars (*) denote significance at the 5% level. The observation is at the city level and weighted by the inverse of the sampling variance of the dependent variable. IV results use the Net Imports instrument described in the text.

Table 3.6: Breakdown By Education Groups

	HS or $<$		BA or $>$	
	OLS	IV	OLS	IV
Variables	(1)	(2)	(3)	(4)
$B_{c,t}$	0.831	2.074	-.23	-2.170
	(0.128)*	(0.591)*	(0.142)	(1.259)
$E_{c,t}$	0.312	0.373	0.714	0.186
	(0.109)*	(0.13)*	(0.514)	(0.717)
Obs.	608	608	608	608
R^2	0.912		0.622	
F-Stat $B_{c,t}$		12.271		2.021
Over ID P-val		0.026		0.211

Notes: Standard errors are in parentheses. Stars (*) denote significance at the 5% level. The observation is at the city level and weighted by the inverse of the sampling variance of the dependent variable. IV results use the Net Imports instrument described in the text.

Table 3.7: Basic Results by Census Year with City Fixed effects

	1970	1980	1990	2000	70-00
Variables	(1)	(2)	(3)	(4)	(5)
$B_{c,t}$	-.461	0.557	1.071	1.401	.417
	(0.288)	(0.195)*	(0.231)*	(0.212)*	(0.105)*
$E_{c,t}$	-1.812	-1.334	-.794	-.682	-.421
	(0.701)*	(0.232)*	(0.227)*	(0.184)*	(0.106)*
Obs.	304	304	304	304	1216
R^2	0.744	0.815	0.791	0.72	0.715

Notes: Standard errors are in parentheses. Stars (*) denote significance at the 5% level. The observation is at the city level and weighted by the inverse of the sampling variance of the dependent variable. IV results use the Net Imports instrument described in the text.

Table 3.8: Basic Results by Census Year with City Fixed effects

	1970	1980	1990	2000	70-00
Variables	(1)	(2)	(3)	(4)	(5)
$B_{c,t}$	0.156	0.826	1.227	1.710	1.002
	(0.364)	(0.189)*	(0.209)*	(0.185)*	(0.105)*
$BA \times B_{c,t}$	-.928	-.948	-1.162	-1.436	-1.289
	(0.378)*	(0.233)*	(0.227)*	(0.311)*	(0.142)*
$E_{c,t}$	-1.730	-.941	-.576	-.643	-.258
	(0.685)*	(0.232)*	(0.217)*	(0.171)*	(0.101)*
Obs.	304	304	304	304	1216
R^2	0.751	0.835	0.819	0.753	0.739

Notes: Standard errors are in parentheses. Stars (*) denote significance at the 5% level. The observation is at the city level and weighted by the inverse of the sampling variance of the dependent variable. IV results use the Net Imports instrument described in the text.

Bibliography

ACEMOGLU, D. (1996): "A Microfoundation for Social Increasing Returns in Human Capital Accumulation," *The Quarterly Journal of Economics*, 111(3), 779–804.

——— (1999): "Changes in Unemployment and Wage Inequality: An Alternative Theory and Some Evidence," *American Economic Review*, 89(5), 1259–1278.

——— (2002): "Technical Change, Inequality, and the Labor Market," *Journal of Economic Literature*, 40(1), 7–72.

ACEMOGLU, D., AND J. ANGRIST (1999): "How Large are the Social Returns to Education? Evidence from Compulsory Schooling Laws," NBER Working Papers 7444, National Bureau of Economic Research, Inc.

AHN, H., AND J. L. POWELL (1993): "Semiparametric estimation of censored selection models with a nonparametric selection mechanism," *Journal of Econometrics*, 58(1-2), 3–29.

ALTONJI, J. G., AND R. M. BLANK (1999): "Race and gender in the labor market," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 3 of *Handbook of Labor Economics*, chap. 48, pp. 3143–3259. Elsevier.

AUTOR, D. (2007): "Structural Demand Shifts and Potential Labor Supply Responses in the New Century," Discussion paper, Prepared for the Federal Reserve Bank of Boston.

AUTOR, D., AND D. DORN (2007): "Inequality and Specialization: The Growth of Low-Skill Service Jobs in the United States," Working papers, MIT.

AUTOR, D., F. LEVY, AND R. MURNANE (2003): "The Skill Content of Recent Technological Change: An Empirical Exploration," *The Quarterly Journal of Economics*, 118(4), 1279–1333.

AUTOR, D. H. (2001): "Wiring the Labor Market," *Journal of Economic Perspectives*, 15(1), 25–40.

AUTOR, D. H., L. F. KATZ, AND M. S. KEARNEY (2006): "The Polarization of the U.S. Labor Market," NBER Working Papers 11986, National Bureau of Economic Research, Inc.

BEAUDRY, P., AND F. COLLARD (2006): “Globalization, Returns to Accumulation, and the World Distribution of Income,” *Journal of Monetary Economics*.

BEAUDRY, P., F. COLLARD, AND D. GREEN (2005): “Demographics and Recent Productivity Performance: Insights from Cross-Country Comparisons,” *Canadian Journal of Economics*, (2), 309–344.

BEAUDRY, P., M. DOMS, AND E. LEWIS (2006): “Endogenous Skill Bias in Technology Adoption: City-Level Evidence from the IT Revolution,” NBER Working Papers 12521, National Bureau of Economic Research, Inc.

BEAUDRY, P., AND D. A. GREEN (2005): “Changes in U.S. Wages, 1976-2000: Ongoing Skill Bias or Major Technological Change?,” *Journal of Labor Economics*, 23(3), 491–526.

BEAUDRY, P., D. A. GREEN, AND B. SAND (2007): “Spill-Overs from Good Jobs,” NBER Working Papers 13006, National Bureau of Economic Research, Inc.

——— (2008): “The Value of Good Jobs: A General Equilibrium Perspective on a Recurring Debate,” Working papers, National Bureau of Economic Research, Inc.

BERNARD, A. B., AND J. B. JENSEN (1998): “Understanding Increasing and Decreasing Wage Inequality,” NBER Working Papers 6571, National Bureau of Economic Research, Inc.

BIDNER, C. (2007): “A Spillover-Based Theory of Credentialism,” Working paper, University of British Columbia.

BLANCHARD, O. J., AND L. F. KATZ (1992): “Regional Evolutions,” *Brookings Papers and Economic Activity*, 23(1992-1), 1–76.

BLAU, F., AND L. KAHN (2007): “The Gender Pay Gap,” *The Economists’ Voice*, 4(4), 5.

BLAU, F. D., AND L. M. KAHN (1997): “Swimming Upstream: Trends in the Gender Wage Differential in 1980s,” *Journal of Labor Economics*, 15(1), 1–42.

BLAU, F. D., AND L. M. KAHN (2006): “The U.S. gender pay gap in the 1990s: slowing convergence,” *Industrial and Labor Relations Review*, 60(1), 45–66.

BORJAS, G. J., J. GROGGER, AND G. H. HANSON (2008): “Imperfect Substitution between Immigrants and Natives: A Reappraisal,” Working Paper 13887, National Bureau of Economic Research.

Borjas, G. J., and V. A. Ramey (1995): “Foreign Competition, Market Power and Wage Inequality: Theory and Evidence,” *The Quarterly Journal of Economics*, 110(4), 1075–1110.

Bound, J., and G. Johnson (1992): “hanges in the Structure of Wages in the 1980’s: An Evaluation of Alternative Explanations,” *American Economic Review*, pp. 371–392.

Card, D. (2001): “Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration,” *Journal of Labor Economics*, 19(1), 22–64.

——— (2009): “Immigration and Inequality,” Working Paper 14683, National Bureau of Economic Research.

Card, D., and J. E. DiNardo (2002): “Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles,” *Journal of Labor Economics*, 20(4), 733–783.

Card, D., A. Mas, and J. Rothstein (2007): “Tipping and the Dynamics of Segregation,” NBER Working Papers 13052, National Bureau of Economic Research, Inc.

Ciccone, A., and G. Peri (2002): “Identifying Human Capital Externalities: Theory with an Application to US Cities,” CEPR Discussion Papers 3350, C.E.P.R. Discussion Papers.

Dahl, G. B. (2002): “Mobility and the Return to Education: Testing a Roy Model with Multiple Markets,” *Econometrica*, (6), 2367–2420.

Deaton, A., and D. Lubotsky (2001): “Mortality, Inequality and Race in American Cities and States,” NBER Working Papers 8370, National Bureau of Economic Research, Inc.

Doms, M., and E. Lewis (2006): “Labor supply and personal computer adoption,” Discussion paper.

Feyrer, J. D., B. Sacerdote, and A. D. Stern (2007): “Did the Rust Belt Become Shiny? A Study of Cities and Counties That Lost Steel and Auto Jobs in the 1980s,” .

Fortin, N. M., and T. Lemieux (2000): “Are Women’s Wage Gains Men’s Losses? A Distributional Test,” *American Economic Review*, 90(2), 456–460.

Glaeser, E., H. Kallal, J. Scheinkman, and A. Shleifer (1992): “Growth in Cities,” *Journal of Political Economy*, 100(6), 1126–1152.

GOOS, M., AND A. MANNING (2007): "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain," *The Review of Economics and Statistics*, 89(1), 11–133.

GREENSTONE, M., AND E. MORETTI (2003): "Bidding for Industrial Plants: Does Winning a 'Million Dollar Plant' Increase Welfare?," NBER Working Papers 9844, National Bureau of Economic Research, Inc.

HARMON, C., AND H. OOSTERBEEK (2000): "The Returns to Education: A Review of Evidence, Issues and Deficiencies in the Literature," CEE Discussion Papers 0005, Centre for the Economics of Education, LSE.

HECKMAN, J. J. (1979): "Sample Selection Bias as a Specification Error," *Econometrica*, 47(1), 153–162.

HECKMAN, J. J., AND R. J. ROBB (1985): "Alternative methods for evaluating the impact of interventions," in *Longitudinal Analysis of Labor Market Data*, ed. by J. Heckman, and B. Singer. Cambridge University Press.

IMBENS, G. W., AND J. D. ANGRIST (1994): "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 62(2), 467–75.

IRANZO, S., AND G. PERI (2006): "Schooling Externalities, Technology and Productivity: Theory and Evidence from U.S. States," NBER Discussion Papers 12400, NBER.

JENSEN, B., AND L. KLETZER (2005): "Tradable Services: Understanding the Scope and Impact of Services Outsourcing," Working Paper 05-9, Institute for International Economics.

LANGE, F., AND R. TOPEL (2004): "The Social Value of Education and Human Capital," Mimo, University of Chicago.

LEE, L. (1983): "Generalized Econometric Models with Selectivity," *Econometrica*, 51(2), 507–512.

LEMIEUX, T. (2006): "Postsecondary Education and Increasing Wage Inequality," *American Economic Review*, 96(2), 195–199.

——— (2007): "The Changing Nature of Wage Inequality," NBER Working Papers 13523, National Bureau of Economic Research, Inc.

LEVY, F., AND P. TEMIN (2007): "Inequality and Institutions in 20th Century America," NBER Working Papers 13106, National Bureau of Economic Research, Inc.

LEWIS, E. (2005): "Immigration, skill mix, and the choice of technique," Discussion paper.

LOADER, C. (1996): "Change point estimation using nonparametric regression," *Annals of Statistics*, (4), 1667–1678.

MANNING, A. (2004): "We can work it out: the impact of technological change on the demand for low skill workers," CEP Discussion Papers 640, Center for Economic Performance.

MANSKI, C. F. (1993): "Identification of Endogenous Social Effects: The Reflection Problem," *Review of Economic Studies*, 60(3), 531–42.

MAZZOLARI, F., AND G. RAGUSA (2007): "Spillovers from High-Skill Consumption to Low-Skill Labor Markets," IZA Discussion Papers 3048, Institute for the Study of Labor (IZA).

MOFFITT, R. A. (2001): "Policy Interventions, Low Level Equilibria, and Social Interactions," in *Social Dynamics*, ed. by S. N. Durlauf, and H. P. Young. MIT Press.

MORETTI, E. (2003): "Human Capital Externalities in Cities," NBER Discussion Papers 9641, NBER.

——— (2004): "Estimating the social return to higher education: evidence from longitudinal and repeated cross-sectional data," *Journal of Econometrics*, 121(1-2), 175–212.

MORTENSEN, D. T., AND C. A. PISSARIDES (1999): "New developments in models of search in the labor market," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 3 of *Handbook of Labor Economics*, chap. 39, pp. 2567–2627. Elsevier.

OSWALD, A. J., AND D. G. BLANCHFLOWER (1995): *The Wage Curve*, vol. 1 of *MIT Press Books*. The MIT Press.

OTTAVIANO, G. I., AND G. PERI (2008): "Immigration and National Wages: Clarifying the Theory and the Empirics," Working Paper 14188, National Bureau of Economic Research.

PARK, J. H. (1994): "Estimation of Sheepskin Effects and Returns to Schooling Using he Old and the New CPS Measures of Educational Attainment," Papers 338, Princeton, Department of Economics - Industrial Relations Sections.

PERI, G., AND G. L. P. OTTAVIANO (2006): "Rethinking the Effects of Immigration on Wages," Working Papers 06-34, University of California at Davis, Department of Economics.

ROY, A. D. (1951): "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers*, pp. 136–146.

RUDD, J. (2000): "Empirical evidence on human capital spillovers," Discussion paper.

RUGGLES, S., M. SOBEK, T. ALEXANDER, C. A. FITCH, R. GOEKEN, P. K. HALL, M. KING, AND C. RONNANDER (2004): *Integrated Public Use Microdata Series: Version 3.0 [Machine-readable database]*. Minneapolis, MN: Minnesota Population Center [producer and distributor].

SAND, B. (2006): "Estimating Labour Supply Responses Using Provincial Tax Reforms," University of British Columbia Working Paper.

WELCH, F. (2000): "Growth in Women's Relative Wages and in Inequality among Men: One Phenomenon or Two?," *American Economic Review*, 90(2), 444–449.

WHEELER, C. H., AND E. A. L. JEUNESSE (2006): "Trends in the distributions of income and human capital within metropolitan areas: 1980-2000," Discussion paper.

Appendix A

Appendix to Chapter 1

A.1 Data

The paper uses data from the 1970, 1980, 1990, and 2000 Census Public Use Mirco-Samples (PUMS) extracted from the IPUMS system.[72]

The initial extract included all individuals aged 18-64 not living in group quarters who had positive annual hours worked in previous year (calculated as weeks worked in previous year times the usual hours worked. The main analysis in the paper is limited to wage and salary workers with positive weekly wages who were living in a metropolitan area at the time of the Census. All calculations were made using the sample weights provided.[73]

The wage variable is the log of weekly wages, calculated by dividing wage and salary income by annual weeks worked. Top coded wage and salary incomes were imputed by multiplying the top code value by 1.5 in each year. Since top codes vary by state in 1990, 2000 and 2005, I use the cut-offs of \$140,000, \$175,000 and \$200000, respectively.[74] All dollar figures were converted to real terms (year 1999 dollars) using the national Consumer Price Index.

A consistent measure of education is not available for these Census years. I use indicators based on the IPUMS recoded variable `EDUCREC` that computes comparable categories from the 1980 Census data on years of school completed and later Census years that report categorical schooling only. To calculate potential experience (age minus years of education minus six), I assigned group means to the categorical education variable reported in 1990 and 2000 from Table (5) of Park (1994).

Census definitions of metropolitan areas are not comparable over time, since, in general, the geographic areas of MSAs increase over time and their definitions are updated to reflect this expansion. The definition of cities I use attempts to maximise geographical comparability over time and roughly correspond to the 1990 definition of MSAs provided

[72]See Ruggles et al. (2004). The files used were the 1980 5% State (A Sample), 1990 5% State, and the 2000 5% Census PUMS. For 1970, Forms 1 and 2 were used for the Metro sample.

[73]For the 1970 data, we adjust the weights for the fact that we combine two samples.

[74]All wages greater than these values were coded as State medians (1990) or means (2000 and 2005) in the raw data.

by the U.S. Office of Management and Budget (OMB).[75] To create geographically consistent MSAs, I follow a procedure based largely on Deaton and Lubotsky (2001) that uses the geographical equivalency files for each year to assign individuals to MSAs or PMSAs based on FIPs state and PUMA codes (in the case of 1990 and 2000) and county group codes (for 1980 and 1970). Each MSA label I use is essentially defined by what PUMAs they span in 1990. Once we have this information, the equivalency files tell us what counties to include in each city for the other years. Since the 1970 country group definitions are much more coarse than later years, we have less cities in the 1970-80-90-2000 sample (152 MSAs) compared to the 1980-90-2000 sample (286 MSAs).[76]

A.2 Instrumental Variables

A.2.1 Enclave Instrument

The construction of the enclave instrument follows Doms and Lewis (2006). The country of origin groups are (1) Mexico, (2) Central America, (3) South America, (4) Central Europe and Russia, (4) Caribbean, (5) China, (6) South East Asia, (7) India, (8) Canada, U.K., and Australia, (10) Africa, (11) Korea and Japan, (12) Pacific Islands, (13) Israel and NW Europe, (14) Middle East, (15) Central Asia, (16) Cuba, and (17) Souther Europe and can be identified from the IPUMS variable `bpl "Birthplace [general version]"` . To identify the inflows of immigrants, I use the IPUMS variable `yrimmig "Year of immigration"`.

A.2.2 Climate Instrument

The city level climate variables were extracted from ”Sperling's Best Places to Live” (http://www.bestplaces.net/docs/DataSource.aspx). Their data is compiled from the National Oceanic and Atmospheric Administration. The variables I use in this paper are the average daily high temperatures for July and January in degrees Fahrenheit. Alternative variables available from the same source are annual rainfall in inches and a comfort index. The comfort index is a variable created by ”Sperling's Best Places to Live” that uses afternoon temperature in the summer and local humidity to create an index in which higher values reflect greater ”comfort”. I have also compiled climate data from an alternative source to use as a robustness check. These data come from ”CityRating.com's” historical weather data, and include variables on average annual temperature, number of extreme temperature days per year, humidity, and annual precipitation.

[75] See http://www.census.gov/population/www/estimates/pastmetro.html for details.

[76] Code for this was generously provided by Ethan G. Lewis. My definitions differ from Deaton and Lubotsky only slightly to improve the 1970-80-90-2000 match.

Data from this source could only be collected for 106 cities.

Table (A.1) contains regression estimates of the change in city college share on several climate variables. The purpose of this exercise is to establish suggestive evidence that college graduates value mild climates. The first two columns of this table show the estimates of the change in college share on the comfort index for each decade. Consistent the idea that college graduates locate at least partially based on the amenity value of city climate, the comfort index is positively and significantly associated with changes in city college share. Columns (3) and (4) repeat this exercise using the number of extreme temperature days. These estimates indicate that the number of extreme temperature days, measured either by the average number of days below 32° or above 90° Fahrenheit, are negatively associated with the change in college share and are jointly significant. The last two columns of Table (A.1) show estimates using the average high temperatures in July and January and their squares. The quadratic specification indicates that mild climates are a good predictor of changes in city college share. It should also be noted that the coefficients from these regressions are quite similar between the two decades in each specification and suggest that the preference for climate amenities are relatively constant over the time period studied. In other specifications that can be obtained from my website, I show that the inclusion of the college premia does not greatly affect the coefficients reported in Table (A.1). This is consistent with Dahl (2002) who shows that while college graduates locate at least partly on comparative advantage, local amenities also play a large role in migration decisions of college graduates.

Tables (A.2) and (A.3) show several first and second stage estimates where the dependent variable is the change in average city wage (of all workers). Table (A.2) uses the Sperling's climate data used elsewhere in this paper and Table (A.3) uses the alternative climate data source to assess robustness. Columns (1) and (2) of Table (A.2) replicate the specification used in Table (1) in the main text. Columns (3)-(6) show that alternative specifications yield roughly similar results. Table (A.3) shows that alternative indicators of mild climates produce very similar results to those using the Sperling's data, even on the restricted city sample.

Table A.1: Changes in College Share and Climate Variables

	1980-1990	1990-2000	1980-1990	1990-2000	1980-1990	1990-2000
Variables	(1)	(2)	(3)	(4)	(5)	(6)
Comfort Index	0.028	0.021				
	(0.01)***	(0.011)*				
# days $< 32°$			-.00009	-.00005		
			(0.00005)*	(0.00006)		
# days $> 90°$			-.0002	-.0002		
			(0.00007)***	(0.00008)***		
Ave. High: July					0.004	0.006
					(0.005)	(0.006)
Ave. High: July^2 / 1000					-.028	-.040
					(0.031)	(0.037)
Ave. High: January					0.001	0.0002
					(0.0004)**	(0.0003)
Ave. High: January^2 / 1000					-.015	-.004
					(0.006)**	(0.004)
Obs.	286	286	106	106	286	286
F-Stat.	7.13	3.53	6.95	9.34	5.48	5.85
P-val.	0.01	0.06	0.00	0.00	0.00	0.00

Notes: The dependent variable is the change in the fraction of college graduates in a city's workforce. Regressions are weighted by $1/\sqrt{1/n_t + 1/n_{t-1}}$ where n is the size of a city's workforce. Robust standard errors are in parenthesis. Stars (*) denote significance at the 99% (***), 95% (**), and 90% (*) levels.

Table A.2: 2SLS using Climate Variable Instruments

	First Stage					
	1980-1990	1990-2000	1980-1990	1990-2000	1980-1990	1990-2000
Variables	(1)	(2)	(3)	(4)	(5)	(6)
Ave. High: July	0.004	0.004			-.001	-.003
	(0.005)	(0.005)			(0.005)	(0.005)
Ave. High: July^2 / 1000	-.032	-.030			0.007	0.008
	(0.031)	(0.03)			(0.026)	(0.027)
Ave. High: January	0.0007	0.00009			0.0008	0.00009
	(0.0004)	(0.0003)			(0.0004)**	(0.0003)
Ave. High: January^2 / 1000	-.012	-.0009			-.012	-.0009
	(0.007)*	(0.003)			(0.005)**	(0.003)
Comfort Index			0.236	0.174	0.231	0.142
			(0.069)***	(0.064)***	(0.064)***	(0.057)**
Comfort Index^2			-.022	-.018	-.020	-.016
			(0.007)***	(0.008)**	(0.006)***	(0.006)***
Precipitation (in.)					0.002	0.0001
					(0.0005)***	(0.0004)
Precipitation (in.)/100					-.002	0.00005
					(0.0006)***	(0.0005)
F-Statistic	4.41	10.64	6.65	7.42	4.6	9.17
P-Value	0.0	0.0	0.0	0.0	0.0	0.0
	Second Stage					
Δ BA or $>$	2.194	-.875	2.347	-2.363	2.899	-1.481
	(0.547)***	(0.299)***	(0.834)***	(0.948)**	(0.643)***	(0.293)***
Obs.	286	286	286	286	286	286

Notes: The dependent variable is the change in average regression adjusted city wage, as described in the main text. Regressions are weighted by inverse of the standard error of the regression adjusted average city wages. All regressions use the city level controls described in the main text. Robust standard errors are in parenthesis. Stars (*) denote significance at the 99% (***), 95% (**), and 90% (*) levels.

Table A.3: 2SLS using Climate Variable Instruments

	First Stage					
	1980-1990	1990-2000	1980-1990	1990-2000	1980-1990	1990-2000
Variables	(1)	(2)	(3)	(4)	(5)	(6)
# days $< 32°$	-.00004	-.00008	-.00005	-.00009	-.0002	-.0002
	(0.00004)	(0.00003)**	(0.00004)	(0.00003)***	(0.00006)***	(0.00005)***
# days $> 90°$	-.0002	-.0003	-.0001	-.0002	-.00009	-.0002
	(0.00008)**	(0.00004)***	(0.00008)	(0.00004)***	(0.00008)	(0.00004)***
Precipitation (in.)			0.001	0.0006	0.002	0.001
			(0.0005)***	(0.0004)	(0.0005)***	(0.0004)***
Precipitation (in.)/100			-.002	-.0008	-.003	-.002
			(0.0007)***	(0.0006)	(0.0007)***	(0.0006)***
Temp. $\times$ humidity (/100)					-1.104	-.879
					(0.505)**	(0.281)***
F-Statistic	2.06	17.00	3.8	8.32	4.59	9.65
P-Value	0.12	0.0	0.0	0.0	0.0	0.0
	Second Stage					
Δ BA or $>$	1.777	-1.130	2.142	-1.069	2.696	-1.235
	(1.554)	(0.474)**	(0.921)**	(0.429)**	(1.074)**	(0.439)***
Obs.	106	106	106	106	106	106

Notes: The dependent variable is the change in average regression adjusted city wage, as described in the main text. Regressions are weighted by inverse of the standard error of the regression adjusted average city wages. All regressions use the city level controls described in the main text. Robust standard errors are in parenthesis. Stars (*) denote significance at the 99% (***), 95% (**), and 90% (*) levels.

A.3 Implementing the Selection Estimator

The approach I use to address the issue of selection on unobservables of workers across cities follows Dahl (2002). Dahl argues that, under a sufficiency assumption, the error mean term in equation (1.4) for individual i can be expressed as a flexible function of the probability that a person born in i's state of birth actually chooses to live in city c in each Census year.[77] Dahl's approach is a two-step procedure that first requires estimates of the probability that i made the observed choice and then adds functions of these estimates into equation (1.4) to proxy for the error mean term. Dahl also presents a flexible method of estimating the migration probabilities that groups individuals based on observable characteristics and uses mean migration flows as the probability estimates. I closely follow Dahl's procedure aside from several small changes to account for the fact that I use cities rather than states and to account for the location of foreign born workers. The procedure used in this paper was first employed in joint work with David Green and Paul Beaudry (Beaudry, Green, and Sand 2007).

Dahl's approach first groups observations based on whether they are "stayers" or "movers". Dahl defines stayers as individuals that reside in their state of birth in the Census year. Since I use cities instead of states, I define stayers as those individuals that reside in a city that is at least partially located in individual's state of birth in a given Census year. Movers are defined as individuals that reside in a city that is not located in that individual's state of birth in a given Census year. I also retain foreign born workers, whereas Dahl drops them. For these workers, I essentially treat them as "movers" and use their country of origin as their "state of birth".[78] Within the groups defined as stayers, movers, and immigrants, I additionally divide observations based on gender, education (4 groups), age (5 groups), black, and hispanic indicators. Movers are further divided by state of birth. For stayers, I further divide the cells based on family characteristics.[79] Immigrants are further divided into cells based on country of origin as described above.

As in Dahl (2002), I estimate the relevant migration probabilities using the proportion of people within cells, defined above, who made the same move or stayed in their birth state. For each group, I calculate the probability that an individual made the observed choice and for movers, I follow Dahl in also calculating the retention probability (ie. the probability that individual i was born in a given state, and remained in a city

[77] This sufficiency assumption essentially says that knowing the probability of an individual's observed or "first-best" choice is all that is relevant for determining the selection effect, and that the probabilities of choices that were not made do not matter in the determination of ones wage in the city where they actually locate.

[78] I use the same country of origin groups as for the enclave instrument.

[79] Specifically, I use single, married without children, and married with at least one child under the age of 5.

situated at least partly in that state in general). For movers, the estimated probabilities that individuals are observed in city c in year t differ based on individuals' state of birth (and other observable characteristics). Thus, identification of the error mean term comes from the assumption that the state of birth does not affect the determination of individual wages, apart from through the selection term. For stayers, identification comes from differences in the probability of remaining in a city in ones birth state for individuals with different family circumstances. For immigrants, I assign the probability that an individual was observed in city c in a given Census year using the probabilities from immigrants with the same observable characteristics in the preceding Census year.[80] This follows the type of ethnic enclave assumption used in several recent papers on immigration, essentially using variation based on the observation that immigrants from a particular region tend to migrate to cities where there are already communities of people with their background.

Having estimated the observed choice or "first-best" choice of stayers, movers, and immigrants and the retention probability for movers, I can then proceed to the second step in adjusting for selection bias. To do this, I add functions of these estimated probabilities into the first stage individual-level regressions used to calculate city-level regression adjusted average city wages. For movers, I add a quadratic of the probability that an observationally similar individual was born in a given state and was observed in a given city and a quadratic of the probability that an observationally similar individual stayed in their birth state. For stayers, I add a quadratic of the probability that an individual remained in their state of birth. For immigrants, I add a quadratic of the probability that an similar individual was observed in a given city in the preceding Census year. Dahl allows the coefficients on these functions to differ by state, whereas I assume that they are the same across all cities.

[80]For cities in the 1980 Census not observed in the 1970 Census, I use the 1980 probabilities.

Appendix B

Appendix to Chapter 2

B.1 Examining Consistency

As described in the text, we are interested in the condition:[81]

$$\lim_{C,I\to\infty} \frac{1}{I}\frac{1}{C}\sum_{i=1}^{I}\sum_{c=1}^{C}\Delta R_c \Delta\xi_{ic}, \tag{B.1}$$

which, using $R = \sum_j \eta_{jc}(w_j - w1)$, can be written as

$$\lim_{C,I\to\infty} \frac{1}{I}\frac{1}{C}\sum_{i=1}^{I}\sum_{c=1}^{C}\left[\sum_{j}^{I}\Delta\eta_{jc}(w_j - w_1) + \sum_{j}^{I}\eta_{jc}\Delta(w_j - w_1)\right]\Delta\xi_{ic}$$

or

$$\lim_{C,I\to\infty} \frac{1}{I}\frac{1}{C}\left[\sum_{j}(w_j - w_1)\sum_{c}\Delta\eta_{jc}\sum_{i}\Delta\xi_{ic} + \sum_{j}\Delta(w_j - w_1)\sum_{c}\eta_{jc}\sum_{i}\Delta\xi_{ic}\right]. \tag{B.2}$$

We will handle the limiting arguments sequentially, allowing $C \to \infty$ first. Then, we are concerned with two components in (A2), which we will handle in turn. The first is

$$\lim_{C\to\infty} \frac{1}{C}\sum_{c}\Delta\eta_{jc}\sum_{i}\Delta\xi_{ic}. \tag{B.3}$$

Given the decomposition $\epsilon_{ic} = \hat{\epsilon}_c + v_{ic}^{\epsilon}$, where $\sum_i v_{ic}^{\epsilon} = 0$, we get

$$\Delta\eta_{jc} = \pi_1(\Delta v_{jc}^{\epsilon}) + \pi_2(\Delta P_j \Omega_{jc} - \Delta P\bar{\Omega}_c), \tag{B.4}$$

where $\bar{x_c}$ equals the simple average of x_{ic} across i within a city.

Also,

$$\sum_{i}\Delta\xi_{ic} = \left(\gamma_1 + \frac{\gamma_1\gamma_2}{1-\gamma_2}\right) I\Delta\hat{\epsilon}_c. \tag{B.5}$$

[81] Throughout this appendix we omit the t subscript for simplicity.

Then, given that $E(\Delta\hat{\epsilon}_c) = 0$ (again, recalling that we have removed economy-wide trends) and if $\Delta\hat{\epsilon}_c$ is independent of Δv^{ϵ}_{ic} and $(\Delta P_j \Omega_{jc} - \Delta P\bar{\Omega}_c)$, it is straightforward to show that (A3) equals zero.

The second component is

$$\lim_{C\to\infty} \frac{1}{C} \sum_c \eta_{jc} \sum_i \Delta\xi_{ic}, \tag{B.6}$$

where $\sum_i \Delta\xi_{ic}$ is again given by (A5), while η_{jc} is given by (20). For (A6) to be zero we require in addition that $\Delta\hat{\epsilon}_c$ be independent of past values of v^{ϵ}_{ic} and of $(P_j\Omega_{jc} - P\bar{\Omega}_c)$. Thus, if $\Delta\hat{\epsilon}_c$ is independent of the past and is independent of Δv^{ϵ}_{ic} and $(\Delta P_j\Omega_{jc} - \Delta P\bar{\Omega}_c)$, then (A1) equals zero and OLS is consistent.

We are also interested in the conditions under which our instruments can provide consistent estimates. Apart from the instruments being correlated with ΔR_{ic}, the condition we require for a given instrument, Z_c is,

$$\lim_{C,I\to\infty} \frac{1}{I}\frac{1}{C} \sum_{i=1}^{I}\sum_{c=1}^{C} Z_c \Delta\xi_{ic}. \tag{B.7}$$

For what we call IV1,

$$Z_c = \sum_j \eta_{jc}(g*_j - 1)(w_j - w_1),$$

where $g*_j = \frac{1+g_j}{\sum_k \eta_{kc}(1+g_k)}$ and g_j is the growth rate in employment in industry j at the national level. Given this, (A7) becomes:

$$\lim_{C\to\infty} \frac{1}{C} \sum_j (w_j - w_1) \sum_c \eta_{jc}(g*_j - 1) \sum_i \Delta\xi_{ic}. \tag{B.8}$$

Thus, (A8) equals zero under the same conditions under which (A6) equaled zero, i.e., that $E(\Delta\hat{\epsilon}_c) = 0$ and $\Delta\hat{\epsilon}_c$ to be independent of past values of v^{ϵ}_{ic} and $(P_j\Omega_{jc} - P\bar{\Omega}_c)$. Obviously, this condition will be satisfied is $\hat{\epsilon}_c$ beaves as a random walk with increments independent of the past.

Similarly, the relevant condition when using IV2 is given by

$$\lim_{C\to\infty} \frac{1}{C} \sum_j \Delta(w_j - w_1) \sum_c \eta_{jc} \sum_i \Delta\xi_{ic}, \tag{B.9}$$

and the same conditions ($\Delta\hat{\epsilon}_c$ to be independent of past values of v^{ϵ}_{ic} and of $(P_j\Omega_{jc} - P\bar{\Omega}_c)$) ensure that this condition equals zero.

Several points follow from this discussion. First, OLS can provide consistent estimates and for it to do so requires the assumptions needed for the IV's to provide consistent estimates (that changes in the absolute advantage for a city are independent of the initial set of comparative advantage factors for that city) plus the stronger assumption that changes in absolute advantage and changes in comparative advantage are independent. Thus, if OLS and the two IV estimates are equal then this is a test of the stronger assumption about independence in changes. Second, if the key identifying assumption underlying the IVs is not true (i.e., changes in absolute advantage are not independent of past comparative advantage) then the two IV's weight the problematic correlation (between $\Delta\hat{\epsilon}_c$ and v_{ic}^{ϵ}) differently (in particular, IV1 weights using the weights $(w_j - w_1)$, while IV2 uses the weights $\Delta(w_j - w_1)$) and estimates based on the different IVs should differ.

B.2 Data Construction

The Census data was obtained with extractions done using the IPUMS system (see Ruggles *et al.* (2004). The files were the 1980 5% State (A Sample), 1990 State, and the 2000 5% Census PUMS. For 1970, Forms 1 and 2 were used for the Metro sample. The initial extraction includes all individuals aged 20 - 65 not living in group quarters. All calculations are made using the sample weights provided. For the 1970 data, we adjust the weights for the fact that we combine two samples. We focus on the log of weekly wages, calculated by dividing wage and salary income by annual weeks worked. We impute incomes for top coded values by multiplying the top code value in each year by 1.5. Since top codes vary by State in 1990 and 2000, we impose common top-code values of $140,000$ in 1990 and $175,000$ in 2000.

A consistent measure of education is not available for these Census years. We use indicators based on the IPUMS recoded variable `EDUCREC` that computes comparable categories from the 1980 Census data on years of school completed and later Census years that report categorical schooling only. To calculate potential experience (age minus years of education minus six), we assign group mean years of education from Table 5 in Park (1994) to the categorical education values reported in the 1990 and 2000 Censuses.

Census definitions of metropolitan areas are not comparable over time since, in general, the geographic areas covered by them increase over time and their definitions are updated to reflect this expansion. The definition of cities we use attempts to maximize geographic comparability over time and roughly correspond to 1990 definitions of MSAs provided by the U.S. Office of Management and Budget.[82] To create geographically con-

[82] See http://www.census.gov/population/estimates/pastmetro.html for details.

sistent MSAs, we follow a procedure based largely on Deaton and Lubotsky (2001) which uses the geographical equivalency files for each year to assign individuals to MSAs or PMSAs based on FIPs state and PUMA codes (in the case of 1990 and 2000) and county group codes (for 1970 and 1980). Each MSA label we use is essentially defined by the PUMAs it spans in 1990. Once we have this information, the equivalency files dictate what counties to include in each city for the other years. Since the 1970 county group definitions are much courser than those in later years, the number of consistent cities we can create is dictated by the 1970 data. This process results in our having 152 MSAs that are consistent across all our sample years. Code for this exercise was generously provided by Ethan G. Lewis. Our definitions differ slightly from those in Deaton and Lubotsky (2001) in order to improve the 1970-1980-1990-2000 match.

We use an industry coding that is consistent across Censuses and is based on the IPUMS recoded variable `IND1950`, which recodes census industry codes to the 1950 definitions. This generates 144 consistent industries.[83] We have also replicated our results using data only for the period 1980 to 2000, where we can use 1980 industry definitions to generate a larger number of consistent industry categories.[84] We are also able to define more (231) consistent cities for that period.

B.3 Implementing the Selection Estimator

As described in the paper, our main approach to addressing the issue of selection on unobservables of workers across cities follows Dahl (2002). Dahl argues that the error mean term in equation (2.25) for person j can be expressed as a function of the full set of probabilities that a person born in j's state of birth would choose to live in each possible city in the Census year. Further, he presents a sufficiency assumption under which the error mean term is a function only of the probability of the choice actually made by j. That sufficiency condition essentially says that two people with the same probability of choosing to live in a given city have the same error mean term in their regression: knowing the differences in their probabilities of choosing other options is not relevant for the size of the selection effect in the process determining the wage where they actually live. Dahl, in fact, presents evidence that this assumption is overly restrictive and settles on a specification in which the error mean term is written as a function of the probability

[83] See http://usa.ipums.org/usa-action/variableDescription.do?mnemonic=IND1950 for details.

[84] The program used to convert 1990 codes to 1980 comparable codes is available at http://www.trinity.edu/bhirsch/unionstats . That site is maintained by Barry Hirsch, Trinity University and David Macpherson, Florida State University. Code to convert 2000 industry codes into 1990 codes was provided by Chris Wheeler and can be found at http://research.stlouisfed.org/publications/review/past/2006. See also a complete table of 2000-1990 industry crosswalks at http://www.census.gov/hhes/www/ioindex/indcswk2k.pdf

of making the migration choice actually observed and the probability that the person stayed in their birth state.

Implementing Dahl's selection correction approach requires two further decisions: how to estimate the relevant migration probabilities and what function of those probabilities to use as the error mean term. For the first, Dahl proposes a non-parametric estimator in which he divides individuals up into cells defined by discrete categories for education, age, gender, race and family status. He then uses the proportion of people within the cell that is relevant for person j who actually made the move from j's birth state to his destination and the proportion who stayed in his birth state as the estimates of the two relevant probabilities. This is a flexible estimator which does not impose any assumptions about the distribution of the errors in the processes determining the migration choice. For the second decision, Dahl uses a series estimator to provide a non-parametric estimate of the error mean term as a function of these probabilities.

We essentially implement Dahl's approach in the same manner apart from several small changes. First, we are examining the set of people who live in cities in the various Census years but we only know the state, not the city of birth. We form probabilities of choosing each city for people from each state of birth. People who live in a city in their state of birth are classified as "stayers" and those observed in a city not in their state of birth are classified as "movers".[85] We estimate the error mean term as a function of the probability that a person born in j's state of birth moved to j's city of residence and the probability that a person born in j's state of birth still resided in that same state. Stayers have an error mean term which is a function only of the probability that the person stayed in their state of birth (since the probability of their actual choice and the probability of staying are one and the same).

As in Dahl (2002), we estimate the relevant probabilities using the proportion of people within cells defined by observable characteristics who made the same move or who stayed in their birth state. Similar to Dahl (2002), we define the cells using 4 education categories, 8 age categories, gender and a black race dummy. For stayers, we also use extra dimensions based on family status.[86] This is possible because of the larger number of stayers than movers. The full interaction of these various characteristics defines 80 possible person types for the movers and 240 for stayers. For the movers in a particular city (i.e., for the set of people born outside the city in which that city is situated), the probabilities will also differ based on where the person was born. Thus, identification of the error mean term comes from the assumption that where a person was born does not affect the determination of their wage, apart from through the error mean term. Intu-

[85] For cities that span more than one state, we call a person who is observed in a city that is at least partly in their birth state a stayer.

[86] Specifically, we use single, married without children, and married with at least one child under age 5.

itively, a person born in Pennsylvania has a lower probability of being observed in Seattle than a person born in Oregon. If both are in fact observed living in Seattle then we are assuming that the person from Pennsylvania must have a larger Seattle specific ”ability” (a stronger earnings related reason for being there) and this is what is being captured when we include functions of the relevant functions of being observed in Seattle for each of them. For stayers, we do not have this form of variation and, hence, identification arises from the restriction that family status affects the decision to stay in one's state of birth but not (directly) the wage.

Our main difference relative to Dahl (2002) is that while he drops immigrants, we keep them in our sample. We essentially treat them as if they are born in a different state from the city of residence except that we do not include a probability of their remaining in their place of birth. We divide the rest of the world into 11 regions (or ”states” of birth). As with other movers, we divide them into cells based on the same education, age, gender and race variables and assign them a probability of choosing their city of residence. Contrary to other movers, however, we do not assign them the probability that immigrants from their region of birth are observed in their own city in the current Census year. Instead, we assign them the probability that a person with their same education was observed in their city in the previous Census. This follows the type of ethnic enclave assumption used in several recent papers on immigration, essentially using variation based on the observation that immigrants from a particular region tend to migrate to cities where there are already communities of people with their background.

Having obtained the estimated probabilities of following observed migration paths and of staying in state of birth, we need to introduce flexible functions of them into our regressions. In practice, we introduce these functions in our first estimation stage. The specific functions we use are quadratics in the estimated probabilities. For movers born in the U.S., we introduce a quadratic in the probability of moving to the actual city from the state of birth and a quadratic in the probability of remaining in the state of birth. For stayers, we introduce a quadratic in the probability of remaining in the state in general. For immigrants, we introduce a quadratic in the probability that people from the same region and with the same education chose the observed city. This represents a restriction on Dahl (2002), who allowed for separate functions for each destination state. We, instead, assume the parameters in the functions representing the error mean term are the same across all cities.

Appendix C

Appendix to Chapter 3

C.1 Data

C.1.1 Census Data

The Census data was obtained with extractions done using the IPUMS system (see Ruggles *et al.* (2004). The files were the 1980 5% State (A Sample), 1990 State, and the 2000 5% Census PUMS. For 1970, Forms 1 and 2 were used for the Metro sample. The initial extraction includes all individuals aged 18 - 65 not living in group quarters. All calculations are made using the sample weights provided. For the 1970 data, we adjust the weights for the fact that we combine two samples. We focus on the log of hourly wages, calculated by dividing wage and salary income by annual hours worked (usual hours worked $\times$ annual weeks worked). We impute incomes for top coded values by multiplying the top code value in each year by 1.5. Since top codes vary by State in 1990 and 2000, we impose common top-code values of $140,000$ in 1990 and $175,000$ in 2000.

A consistent measure of education is not available for these Census years. We use indicators based on the IPUMS recoded variable `EDUCREC` that computes comparable categories from the 1980 Census data on years of school completed and later Census years that report categorical schooling only. To calculate potential experience (age minus years of education minus six), we assign group mean years of education from Table 5 in Park (1994) to the categorical education values reported in the 1990 and 2000 Censuses.

Census definitions of metropolitan areas are not comparable over time since, in general, the geographic areas covered by them increase over time and their definitions are updated to reflect this expansion. The definition of cities we use attempts to maximize geographic comparability over time and roughly correspond to 1990 definitions of MSAs provided by the U.S. Office of Management and Budget.[87] To create geographically consistent MSAs, we follow a procedure based largely on Deaton and Lubotsky (2001) which uses the geographical equivalency files for each year to assign individuals to MSAs or PMSAs based on FIPs state and PUMA codes (in the case of 1990 and 2000) and county group codes (for 1970 and 1980). Each MSA label we use is essentially defined by the PUMAs it spans in 1990. Once we have this information, the equivalency files dictate

[87] See http://www.census.gov/population/estimates/pastmetro.html for details.

what counties to include in each city for the other years. Since the 1970 county group definitions are much courser than those in later years, the number of consistent cities we can create is dictated by the 1970 data. This process results in our having 152 MSAs that are consistent across all our sample years. Code for this exercise was generously provided by Ethan G. Lewis. Our definitions differ slightly from those in Deaton and Lubotsky (2001) in order to improve the 1970-1980-1990-2000 match.

We use an industry coding that is consistent across Censuses and is based on the IPUMS recoded variable `IND1950`, which recodes census industry codes to the 1950 definitions. This generates 144 consistent industries.[88]

C.1.2 Current Population Survey

We supplement our empirical work with data from the March Current Population Surveys (CPSs) for the years 1970 - 2007. We use CPS data to analysis movements in relative wages over time, and employ the same selection criteria as with the U.S. Census data described above.

[88]See http://usa.ipums.org/usa-action/variableDescription.do?mnemonic=IND1950 for details.

Additional Information

This thesis was typeset using the `ubcthesis` LaTeX class developed by Michael McNeil Forbes.