

SPILL-OVERS FROM GOOD JOBS

Paul Beaudry
David A. Green
Benjamin Sand

Working Paper 13006
<http://www.nber.org/papers/w13006>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2007

The authors thank M. Bombardini, G. Dalh, M. Doms, J. Fernald, N. Fortin, G. Galipolli, R. Gordon, M. Greenstone, S. Kortum, T. Lemieux, K. Milligan and B. Meyer for helpful discussions. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2007 by Paul Beaudry, David A. Green, and Benjamin Sand. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Spill-Overs from Good Jobs
Paul Beaudry, David A. Green, and Benjamin Sand
NBER Working Paper No. 13006
April 2007
JEL No. J31

ABSTRACT

Does attracting or losing jobs in high paying sectors have important spill-over effects on wages in other sectors? The answer to this question is central to a proper assessment of many trade and industrial policies. In this paper, we explore this question by examining how predictable changes in industrial composition in favor of high paying sectors affect wage determination at the industry-city level. In particular, we use US Census data over the years 1970 to 2000 to quantify the relationship between changes in industry-specific city-level wages and changes in industrial composition. Our finding is that the spill-over (i.e., general equilibrium) effects associated with changes in the fraction of jobs in high paying sectors are very substantial and persistent. Our point estimates indicate that the total effect on average wages of a change in industrial composition that favors high paying sectors is about 3.5 times greater than that obtained from a commonly used composition-adjustment approach which neglects general equilibrium effects. We interpret our results as being most likely driven by a variant of the mechanism recently emphasized in the heterogenous firm literature whereby changes in competitive pressure cause a reallocation of employment toward the most efficient firms.

Paul Beaudry
Department of Economics
University of British Columbia
Room 997-1873 East Mall
Vancouver, B.C. Canada V6T 1Z1
and NBER
paulbe@interchange.ubc.ca

Benjamin Sand
Department of Economics
University of British Columbia
997-1873 East Mall
Vancouver, B.C.
Canada, V6T 1Z1
bmsand@interchange.ubc.ca

David A. Green
Department of Economics
University of British Columbia
997-1873 East Mall
Vancouver, B.C.
Canada, V6T 1Z1
green@econ.ubc.ca

Introduction

In popular discussion about labor market developments, whether it be at the national or international 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”: 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 loss of highly paid manufacturing jobs was key to understanding the poor labor market performance of the US 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 now dismiss such views as being ill-informed.

The accounting approach used to dispel the Bluestone-Harrison [1982] view consists of mechanically computing the fraction of a wage change that can be directly attributed to the loss or gain of employment in high versus low wage paying industries. The result from such exercises almost always indicates that the wage change directly accounted for by changes in sectoral composition of employment is small. This accounting approach has seen widespread use in the profession since it is easy to implement and it is theoretically defensible. Notwithstanding these attractive features, the starting point of this paper is a questioning of the validity of this approach.

The accounting approach hinges critically on the assumption that a change in employment opportunities in one sector does not affect the wages paid in other sectors, i.e., that there are no general equilibrium effects from shifts in industrial composition. Because of this assumption, the impact of a shift in industrial composition can be computed mechanically by multiplying industry- specific wage premia by changes in the share of employment in corresponding industries. Without the assumption, one would also have to account for changes in the wage premia 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. There are many ways to justify the no-GE effects assumption, which is part of the appeal of this approach. The easiest defence is to note that if wages are simply a function of productivity and returns to labour 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. In particular, there is now a substantive and growing literature – mainly developed in the context of international trade – which suggests that average productivity in a sector is potentially endogenous to the competitive environment and, accordingly, responds to outside forces. For example, Melitz [2003] and Bernard, Eaton, Kortum and Jensen [2003] show how the opening up of trade can cause sectoral productivity to increase by forcing inefficient firms to exit, resulting in a reallocation of production towards more productive firms. Although these ideas have not yet played a large role in the labor literature, they appear very relevant. Specifically, if international trade can have important general equilibrium effects on productivity then it follows that changes in industrial composition may have important effects on sectoral level wages (to the extent that wages are generally related to productivity).

The object of this paper is to examine whether changes in sectoral composition of employment, especially shifts in composition between high paying sectors and low paying sectors, have important general equilibrium effects on the determination of within sector wages. To address this question, we first present a simple theory that demonstrates how a change in industrial composition, through its effect on the bargaining environment, can affect wages in sectors not directly involved in the compositional change. In such sectors, we argue that an improved outside option for workers will place upward pressure on wages, forcing inefficient firms to exit the market and thereby favoring a reallocation of employment toward more productive firms. Since workers manage to appropriate part of the gain in average productivity through the bargaining process, their wages increase. This mechanism echoes closely that discussed in the recent heterogenous firm trade literature. However, in much of the related trade literature, wages are treated as exogenous, while our focus is on implications for wages.

The model we present is one with substantial frictions. In particular, we adopt a random matching set up in the labor market to allow firms with different productivities to co-exist in a market; an occurrence for which there is substantial corroborating evidence. Although the model is very stylized, and must be interpreted as such, it provides a concise illustration of how changes in industrial composition could have important spill-over effects that would be missed by adopting an accounting approach. Moreover, the model clearly illustrates how a change in industrial composition in a small sub-set of sectors has the potential to affect wages across the entire range of industries.

In order to examine the relevance of spill-over effects associated with changes in the fraction of jobs that are “good jobs”, we exploit geographical variation in industrial composition across

US cities over the period 1970-2000. More specifically, we exploit the fact that aggregate changes in US industrial composition have not been evenly distributed across US cities. Our approach is to look at whether wages paid in a given industry systematically differ across cities depending on the distribution of employment across the other industries in the city and, in particular, whether wages paid in any given industry tend to be higher in a city which has an industrial composition that is more concentrated toward high paying jobs. To do this, we examine 10 and 20 year changes in industry \times city level wages using data from the 1970, 1980, 1990 and 2000 US Censuses for 152 cities.¹ Our key covariate is the change in the weighted average industrial wage premium in the given city, where the weights are local employment shares for each industry and the industrial wage premia are national level premia estimated holding education, experience, race and gender constant. This variable will increase in value if the industrial composition of employment in a city shifts toward higher premia industries. Working in differences in this way allows us to control for general industry \times city fixed effects, and we also include a full set of industry dummy variables, effectively allowing national industrial premia to differ over time. Given this estimation framework, we are working with within-industry variation in wages. If the null hypothesis that there are no GE effects from composition shifts (and, thus, that a mechanical decomposition provides a theoretically appropriate measure) is true then our key covariate, which reflects changes in industrial composition, should have a coefficient of zero in our regressions: wages within industries should not change in response to local changes in composition. Under the alternative we consider in our theoretical model, this covariate should have a positive sign.

We address potential endogeneity concerns with the average industrial premium variable by using an instrumenting strategy in which we predict changes in local industrial employment shares using a combination of initial local shares and changes in employment shares at the national level. Thus, we effectively examine the impact of national level industrial shifts (arising, perhaps, due to preference or trade shocks) on local wages, where we apportion the impact of the national shifts to localities based on which cities were most heavily concentrated in the large shift industries before the shifts occurred. We also devote considerable effort at addressing the possible non-random selection in unobservable worker characteristics across cities using the method proposed in Dahl [2002].

¹ There are both advantages and disadvantages of using city level observations to examine this issue. On the one hand, an attractive feature is there are over 150 metropolitan areas in the US which gives us a sample with many different patterns of industrial composition. On the other, if labor markets across US cities are all perfectly integrated, then we could find no spill-over even if such effects exist 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.

The main empirical result of the paper is that city level changes in industrial composition induced by national level changes in demand patterns have effects on wages that are 3 to 3.5 times greater than what would be predicted by a pure accounting approach. It is important to keep in mind when considering this result that measured composition effects are often small. Thus, 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: a large, though not extreme, effect. Impacts of this size have the potential to place back on the table explanations for changes in the wage structure that may operate through changes in industrial structure and which have largely been discounted because the pure accounting measures of their impacts are relatively small (e.g., Bound and Johnson [1992]).

Having identified a substantial spillover effect of changes in industrial composition on the wage structure, we attempt to narrow down the range of potential explanations for what is driving it. We find that the effect of composition is present over long (20 year) horizons and is present in wages in both the tradeable and non-tradeable goods sectors. We argue that these results do not fit with models in which the composition effect reflects spill-overs into the non-tradeable sector or ones in which workers are simply taking rents from quasi-fixed factors. What remains is induced productivity changes in all sectors of the type illustrated in our theoretical model. Thus, we interpret these results as supporting the idea that changes in workers' outside options may have important effects on rationalization of production within an industry.²

It is important to emphasize that we are interested in longer term differences in wage structure associated with different industrial composition as opposed to short run adjustments to industrial change. Phrased in terms of a concrete example, Pittsburgh suffered from the loss of jobs in the steel sector with high unemployment in the immediate aftermath of those losses. In the longer run, though, Pittsburgh's unemployment rate has declined to levels below the national average and it has emerged with a new industrial composition [Briem 2002]. We are interested in the impact on Pittsburgh's wage structure of the long term shift in its industrial structure rather than the short run wage effects occurring as the direct aftermath

² While it may be possible to examine the implications of our model using productivity data, as opposed to wage data, there are tradeoffs. In particular, it is very difficult to create measures of sectoral productivity that properly control for human capital differences. For this reason, and others, we believe that focusing on wages implications is more desirable.

of the closing of the steel mills. Thus, our focus is different from, for example, that in Greenstone and Moretti [2003] or Blanchard and Katz [1992]. While Greenstone and Moretti focus mainly on local real estate price and same-industry wage bill effects within three years of acquiring a large plant, emphasizing shorter run demand effects, we focus on changes arising over 10 to 20 year horizons and are investigating whether there are wage effects of changes in industrial composition holding direct demand effects, or in other words the size of the city, 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 their paper.

Our paper most closely resembles 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 without substantially changing the estimated effect from our concentration of good jobs measure. Thus, whatever we are identifying, it is over and above these other hypothesized driving forces. In our estimates, the impact of good job concentration is much larger than the estimated impacts of any of these other forces. Our paper is also related to the voluminous literature aimed at understanding the effects of international trade on wages since much of this 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. As we shall show, this perspective appears important since we find little effect of changes in industrial composition on the returns to skill but very important effects on wage levels.

The remaining sections of the paper are as follows. In Section 1, we present a heterogeneous firms setup to illustrate how changes in industrial composition in a sub-set of industries can have substantial spill-over effects of wage determination in other sectors. 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 supports the wide spread 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, 5 and 6 we address issues related to endogeneity, selection, robustness and interpretation. Section 7 concludes.

1 A Simple Model of the Spill-over Effects of Industrial Composition.

The object of this section is to illustrate how changes in the composition of jobs between high and low paying sectors in a sub-set of the economy can lead to wide-spread changes in wages in the remaining sectors. The model combines elements from the holdup literature with the literature on heterogenous firms. To make the analysis tractable, the model economy is comprised of two distinct sub-sets of industries. In the first subset (which we will call the MC sector), firms are monopolistically competitive and face a traditional holdup problem. The second set of industries (which we will call the FE, or free entry, sector) has free entry of firms and, thus, a more competitive structure. However, due to matching frictions in the labor market, firms with different levels of productivity will co-exist in equilibrium in the FE sector. Our goal is to show how changes in demand patterns in the first subset of industries can affect productivity and wages in the second set. In particular, the model will illustrate how a shift in demand toward industries facing more extensive holdup problems will cause a rationalization of production and higher wages in other sectors.

Although the model is highly stylized, it will be sufficiently explicit to offer a clear interpretation for our empirical strategy. The main dividend of the model is its ability to illustrate how the determinants of industrial composition can be summarized into one index, and how variations in this index can affect industry level wages. It is this feature of the model – the potential link between the industrial composition index and industry level wages – which will form the basis of our empirical exploration of the spill-over effects of industrial composition on wages.

1.1 Model

Consider an environment with a set, C , of local labor markets, called cities. In each city there is a mass, L , of potential workers and there is a set, $M = M_1 + M_2$, of industries. For now, we will assume that workers are not mobile across cities. We will also assume

that workers are equally productive and that all goods are tradeable. However, to preview an extension we will make when we derive our empirical specification, it is useful to note that units of labor in the model can be interpreted as effective units. Accordingly, one can interpret all wage implications of the model as implications for wages controlling for human capital differences.

The output from the different industries, denoted X_i , are combined to form a final good y as follows:

$$y = \left[\sum_{i=1}^{M_1+M_2} (a_i X_i)^\sigma \right]^{\frac{1}{\sigma}}, \quad 0 < \sigma < 1$$

where differences in a_i capture differences in demand across industries.

The set of industries is divided into two groups differentiated by their technology and competitive structure. In each of the first group of industries, $i = 1$ to M_1 , the good, $Z_{i,c}$, is produced by local monopolists. Before hiring workers, the employers in these locally monopolized industries must incur a capital cost of k_i per worker. The local monopolists may have different levels of productivity, with $\Omega_{i,c}$ denoting the amount of $Z_{i,c}$ produced by one worker hired by industry i in city c . As will become clear, the main characteristic which will drive local outcomes relates to whether a city's comparative advantage is in industries with high or low capital costs, k_i .

The remaining industries, $i = M_1 + 1$ to M_2 , are assumed to be competitive in the sense that there is free entry into these sectors in each city. However, as in Melitz [2003], in order to produce output in these industries, a firm must first pay a fixed setup cost F to uncover its productivity level θ , where θ is draw from a distribution $G(\theta)$ with support $[\theta_1, \theta_2]$. For a firm with productivity level θ , one unit of effective labor produces θ units of output. Upon learning its productivity, a firm decides whether to stay in the market or exit; with the exit decision based on whether the firm will be able to attract and retain workers.

The total output $Z_{i,c}$ produced in city c for industry i is combined with outputs from other cities to form the industry level output according to,

$$X_i = \left[\sum_c (Z_{i,c})^{\tilde{\sigma}} \right]^{\frac{1}{\tilde{\sigma}}}, \quad 0 < \tilde{\sigma} < 1$$

The timing of actions for the determination of wages and employment is as follows. Given

the relative demand conditions captured by $\{a_i\}$, the capital costs $\{k_i\}$, and the pattern of city level productivity $\{\Omega_{i,c}\}$, the monopolists in the MC sector decide how many workers they would potentially like to hire by choosing to invest in multiples of k_i . At the same time, masses $N_{i,c}$ of firms decide to enter into the different competitive industries, and every entering firm receives a productivity draw. After observing θ , firms decide whether to stay or exit the market. Firms will stay in the market only if they expect to attract and retain workers. Once entry decisions are made, workers can apply to one of the firms in the FE sector (among those that have not exited) or they can wait to see whether they will be hired in the MC sector. Since it is costless to apply to a FE sector firm and the decision is reversible, all workers apply. When a worker applies for a job, he is matched with an employer with probability $1 - \mu$, ($0 < \mu < 1$). When deciding where to apply, workers know the identity of the industry but do not know each firm's realization of θ and therefore apply randomly to one of the active firms in the industry of their choice. If a worker is matched with a firm, they bargain a wage. If a worker does not obtain a match with a firm in the FE sector, or if the bargained wage is not sufficiently attractive to the worker, the worker can try his luck in the monopolistically competitive industries. Bargaining between a worker and a firm in all sectors is governed by the rule of equal division of the surplus.

Since wages and employment possibilities in the MC sector will play the role of an outside option for workers in the competitive sectors, it is best to first solve for wage and employment determination in these sectors. If a worker is hired by a firm in a MC sector, the wage is set such that the surplus to the firm, given by $P_{i,c}\Omega_{i,c} - w$, is equal to the surplus of the worker, given by $w - b$; where $P_{i,c}$ is the price of good $Z_{i,c}$ and b is the outside option of workers who are waiting for employment in the monopolistically competitive industries.³ This bargaining process implies that the wage paid in industry i , $i = 1, \dots, M_1$, in city c is equal to $\frac{P_{i,c}\Omega_{i,c} + b}{2}$.⁴ Foreseeing the outcome of wage bargaining, the monopolistically competitive firm will choose his level of employment, $l_{i,c}$, and price, $P_{i,c}$, as to solve the following maximization problem.

$$\max_{l_{i,c}, P_{i,c}} \left(\frac{P_{i,c}\Omega_{i,c} - b}{2} - k_i \right) l_{i,c}$$

subject to the perceived demand for good $Z_{i,c}$, given by $P_{i,c} = Z_{i,c}^{\tilde{\sigma}-1} X_i^{1-\tilde{\sigma}}$, and the production

³ Here we are assuming that if a worker is not employed in the MC sector then he cannot go back to find employment in the competitive sector. This assumption is purely for simplicity. If we allowed wage determination in the competitive sector to affect the outside option of workers in the MC sector, this would simply amplify the results emphasized here.

⁴ For the bargaining outcome to take this precise form under an equal division rule, we are implicitly assuming that the price $P_{i,c}$ is determined at the same time as capacity.

function $Z_{i,c} = \Omega_{i,c} l_{i,c}$. This maximization problem implies that wages and employment in the MC sectors are given by

$$w_{i,c} = \left[\frac{b + 2k_i}{\tilde{\sigma}} \right] + b$$

$$l_{i,c} = \left[\frac{\tilde{\sigma} \Omega_{i,c}^{\tilde{\sigma}} (X_i)^{1-\tilde{\sigma}}}{b + 2k_i} \right]^{\frac{1}{1-\tilde{\sigma}}} \quad (1)$$

From the above, we can see that wages will be highest in industries where k_i is high. This is a direct reflection of the standard holdup problem. Moreover, these wages do not vary by city. Hence, there is an unambiguous sense in which an industry with a high setup cost, k_i , is a high wage industry. In contrast to wages, we see that city level employment within an industry will vary depending on the city's productivity. In order to focus on comparative advantage instead of absolute advantage, we will assume that differences in $\Omega_{i,c}$ across cities is such that $\sum_{i=1}^{M_1} l_{i,c}$ can be expressed as ϕL for all cities, $0 < \phi < 1$. Thus, an increase in employment in one of the MC industries shifts the composition of employment but not the overall level of employment. In this sense, a city with higher employment in one of the MC industries is a city with a comparative advantage in this industry not one with an absolute advantage in terms of a higher employment rate.

The employment equation in (1) also says that the preferred level of employment of a firm in the MC sector can be determined independently of supply. Because of this, jobs in these industries will generally be rationed. Accordingly, let ρ_c represent the probability of employment of a worker who queues for a job in the MC industries. Assuming that $\rho_c < 1$, and assuming that workers that queue up for jobs in the MC sector are matched randomly to industries, the expected payoff associated with queuing for a job in the MC sector, denoted w_c^A , can be expressed as

$$w_c^A = \rho_c \sum_{i=1}^{M_1} w_i \frac{l_{i,c}}{\sum_{i=1}^{M_1} l_{i,c}} + (1 - \rho_c)b. \quad (2)$$

The payoff w_c^A plays the role of outside option for workers applying to jobs in the FE sectors. The bargained wage in industry $i \in [M_1 + 1, M_1 + M_2]$, for a job with productivity θ , denoted $w_{i,c}(\theta)$, is again set to create an equal division of the surplus between the firm and the worker. The equal division condition can be written as

$$P_{i,c}\theta - w_{i,c}(\theta) = w_{i,c}(\theta) - w_c^A$$

which implies that

$$w_{i,c}(\theta) = \frac{P_{i,c}\theta}{2} + \frac{\rho_c}{2} \sum_{i=1}^{M_1} w_i \left(\frac{l_{i,c}}{\sum_{i=1}^{M_1} l_{i,c}} \right) + \frac{(1-\rho_c)b}{2} \quad (3)$$

Knowing the outcome of bargaining, a firm will decide to stay in the market only if $P_{i,c}\theta > w_c^A$ or, alternatively, if $\theta > \frac{w_c^A}{P_{i,c}}$. Using this observation, and noting that workers will allocate themselves randomly across firms within a sector, we can express the average wage in competitive industry i for city c as follows ⁵

$$w_{i,c} = \frac{P_{i,c}}{2} \int_{\frac{w_c^A}{P_{i,c}}}^{\theta_2} \frac{\theta dG(\theta)}{1 - G\left(\frac{w_c^A}{P_{i,c}}\right)} + \frac{w_c^A}{2}, \quad i = M_1 + 1, \dots, M_1 + M_2 \quad (4)$$

Since workers can choose the industry to which they apply among the M_2 competitive industries, their application decisions will cause the expected wage to be equalized across those industries, which in turn implies that they apply to industry i in proportion to a_i . This equilibrium mechanism will lead the prices, $P_{i,c}$, in the competitive industries to equal a price, P , which is independent of c and i . Therefore in Equation (4), the only city level variable affecting wages is w_c^A . Furthermore, since all workers apply for jobs in the competitive sector and remain there if they are matched, the only workers that queue for jobs in the monopolistically competitive sectors are the μL workers that are not matched to a competitive firm. This allows us to write the probability of obtaining a job in the MC industries as a function of exogenous parameters, $\rho_c = \frac{\phi}{\mu}$, where we assume that $\mu \geq \phi$.⁶ Equations (4) and (2) can

⁵ We can also state how firms' entry decisions is determined. Because of free entry, it must be the case that the mass of firms entering each industry, denoted $N_{i,c}$, is such that there are zero expected profits. The zero expected profit condition can be expressed as follows:

$$\frac{(1-\mu)L_c}{2M_2N_{i,c}(1-G(\frac{w_c^A}{P_{i,c}}))} [P_{i,c}\tilde{\theta}(\frac{w_c^A}{P_{i,c}}) - w_c^A] = F$$

where $\tilde{\theta}(\frac{w_c^A}{P_{i,c}}) = \int_{\frac{w_c^A}{P_{i,c}}}^{\theta_1} \theta dG(\theta)$, and where L_i represents the number of workers who apply to jobs in industry i . In the above zero profit condition, the term $\frac{(1-\mu)L}{M_2N_{i,c}(1-G(\frac{w_c^A}{P_{i,c}}))}$ represents the number of employees a firm will have since $\frac{(1-\mu)L_c}{M_2}$ is the number of employees that will be matched to sector i and $N_{i,c}(1-G(\frac{w_c^A}{P_{i,c}}))$ represents the number of active firms in a sector.

⁶ In the current setup, the fraction of worker queuing for jobs in the MC sector is determined mechanically through the friction in the matching process given by μ . A less mechanical mechanism can be added to the

therefore be written as

$$w_{i,c} = \frac{P}{2} \int_{\frac{w_c^A}{P}}^{\theta_2} \frac{\theta dG(\theta)}{1 - G(\frac{w_c^A}{P})} + \frac{w_c^A}{2}, \quad i = M_1 + 1, \dots, M_1 + M_2 \quad (5)$$

$$w_c^A = \frac{1}{\mu} \sum_{i=1}^{M_1} w_i \frac{l_{i,c}}{L_c} + (1 - \frac{\phi}{\mu})b. \quad (6)$$

The main element to notice from Equations (5) and (6) is how the distribution of employment in the first M_1 industries, $\{l_{i,c}\}_{i=1}^{M_1}$, (which itself is driven by differences in $\Omega_{i,c}$) spills over and affects wages in the remaining industries through its effect on w_c^A . This arises through two channels. The first channel simply relates to the increased bargaining position of workers induced by an improved outside option. In particular, when a city has a distribution of employment in the first M_1 industries which is heavily weighted towards higher paying jobs then workers in the remaining M_2 industries can extract more from any given employer by threatening to leave and opt for the outside option. The second channel captures the effects of improved productivity of matches, echoing effects emphasized in recent trade models such as Melitz [2003] and Bernard, Eaton, Kortum and Jensen [2003]. Since relatively more high paying jobs in the first M_1 industries causes an increase in wages in the other industries, firms react to this change by staying in the market only if they are sufficiently productive. This effect reflects a within-industry reallocation of employment toward more productive firms, and since workers capture a fraction of the increased productivity, they obtain even higher wages. This mechanism suggests that small changes in outside options could potentially create substantial changes in wages due to a type of multiplier effect whereby the first increase in wages causes increased productivity which generates further increases in wages.⁷

In order to render this model empirically relevant, it is useful to recognize how differences in w_c^A across cities can be captured by a simple industrial composition index which we will denote by R_c . To this end, let us define the employment weighted average of national level inter-industry wage premia as follows:

$$R_c = \sum_{i=1}^{M_1+M_2} \left(\frac{\bar{w}_i}{\bar{w}_1} - 1 \right) \left(\frac{l_{i,c}}{\sum_{i=1}^{M_1+M_2} l_{i,c}} \right)$$

model by adding a cost to workers of applying to FE sector jobs. This will cause the fraction of workers that queue for MC sector jobs to endogenously respond to w_c^A , and thereby ensure that ρ is less than 1 even if μ is smaller than ϕ .

⁷ It should be emphasized that frictions in the labor market are central to obtaining these results.

where, $\bar{w}_i = \sum_C \frac{w_{i,c}}{c}$ is the national level average wage in industry i ⁸, and we normalize the wage premia to some industry, arbitrarily labeled as industry 1. Note that in the construction of R_c , wage realizations in city c play a negligible role since the wages used in the index are national averages. Since R_c is the average wage premium (relative to industry 1) across all industries and w_c^A is the weighted expected wage across a subset of industries, it is possible to write one as a positive linear function of the other as follows.

$$w_c^A = d_0 + d_1 R_c \quad (7)$$

where $d_0 = \bar{w}_1 \left[\frac{1-\mu+\phi}{\mu} - \frac{(1-\mu)}{\mu\bar{w}_1} \sum_{i=M_1+1}^{M_1+M_2} \frac{w_i}{M_2} \right] + (1 - \frac{\phi}{\mu}) \frac{b}{\bar{w}_1}$ and $d_1 = \frac{\bar{w}_1}{\mu} (1 - \mu + \phi)$.

We are now in a position to state the feature of this model which most interests us, that is the relationship between industry level wages and a city's industrial composition, as captured by the index R_c . In particular, based on Equations (5), (6) and (7), we see that a city with a higher value of R_c is predicted to exhibit higher wages across a whole set of industries since in such a city workers have a strong bargaining position and this strong bargaining position forces a rationalization of production in heterogenous firm industries. Such a property implies that changes in industry mix has spill-over effects on wages that go beyond the sectors where the composition of employment has changed.⁹ This relationship can be expressed in reduced form as

$$w_{ic} = Q(R_c), \quad Q'(R_c) > 0 \quad , i = 1, \dots, M_1 \quad (8)$$

Up to now, we have been implicitly focusing on how changes in a city's pattern of comparative advantage, working through R_c (or equivalently through w_c^A), affects industry level wages. However, in the model there are actually two sources of variation – which work through R_c – that can cause changes in a city's industry level wage patterns. The first regards changes in across industry demand patterns induced by changes in $\{a_i\}$. Such changes will have differential impacts on the composition of employment across cities since cities have different concentrations in the various industries based on their comparative advantages. Thus, changes in the national demand pattern will change R_c (and w_c^A), and hence local wages, by changing the local composition of employment. A second channel relates to changes

⁸ For simplicity, we are assuming that cities have the same size of population.

⁹ Formally, the model offers a mapping between comparative advantage patterns as driven by $\{\Omega_{i,c}\}_{i=1}^{M_1}$ and industry level wages with a city. The key feature of this mapping is that the effect of $\{\Omega_{i,c}\}_{i=1}^{M_1}$ on wages can be conceptualized as transiting through a simple index of industrial composition.

in holdup problems as captured by changes in $\{k_i\}$. These changes will lead to changes in the industrial wage premia in the MC sector as well as changes in the local composition of employment, implying, again, changes in R_c (and w_c^A) that affect local wages. In the empirical section, we will exploit these two potential channels of aggregate level changes in either $\{a_i\}$ or $\{k_i\}$ to develop an instrumental variable strategy for examining the presence of spill-over effects of changes in R_c on industry level wages.

In the model as set out, we assume that workers cannot move between cities to take advantage of differences in wages. It is important to emphasize that this assumption can easily be relaxed without losing the main implications of the model. In particular, consider the case where workers are mobile across cities but cities with higher population have higher land prices. In this case, if land enters a worker's utility function, worker mobility will not entirely overturn the model's implication that cities with more jobs in the high paying sectors will pay higher wages in many sectors. This is true because workers will not continue to move until wage parity is obtained. Instead, they will move only to the point where the marginal worker has the same expected utility across cities. Hence, with labor mobility, much of the increased productivity induced by better outside options for workers may largely (or entirely) be captured by land prices, but it should nevertheless first be reflected in wages.

2 Deriving an Empirical Specification

The main empirical implication we take away from the model is that the presence of a set of high paying jobs in a local labor market can have spill-over effects on wages in other sectors, even when the price of the goods produced in these other sectors are identical across markets and firms make zero expected profits. As we will make clear, the prediction of a spill-over effect is in direct conflict with assumptions needed to justify an accounting approach when evaluating the impact of goods jobs on average wages.

The first issue that arises in deriving a usable empirical specification is to allow for individual worker heterogeneity. If we assume that workers differ in terms of human capital and that this human capital augments a worker's productivity in all sectors then worker heterogeneity can be introduced easily in the above model. In particular, consider a worker, indexed by k , who has a vector of amounts of observable characteristics, x_{kct} , and an amount of unobserved (to the researcher) human capital given by ϵ_{kct} . Then, given returns on the observed characteristics, β_{1t} , the productivity of this worker in a firm with a productivity

draw of θ is given by $\theta \exp(\beta_{1t}X_{kct} + \epsilon_{kct})$. In this case, the log wage paid to a worker can be expressed as,

$$\ln W_{kict} = \beta_{1t}x_{kct} + \ln(w_{ict}) + \epsilon_{k,c,t}, \quad (9)$$

where w_{ict} is the wage per efficiency unit of labor as given in (8) except that we have now introduced a time subscript. Note that the lower case w_{ict} 's implied by the model should now be interpreted as industry specific wages after controlling for human capital differences.

As noted earlier, the main implication we take away from the model is the positive effect of industrial composition, R_{ct} , on industry specific wages, w_{ict} . Using a linear approximation to this relationship, we can re-write (9) as a function of R_{ct} , yielding our main estimating equation:

$$\ln W_{kict} \approx \beta_{0t} + \beta_{1t}x_{kct} + \beta_2 R_{ct} + \nu_{i,t} + \nu_c + \epsilon_{kct} \quad (10)$$

In this specification, we are allowing for the possibilities that cities face different transportation costs for the aggregate good, y_c , and that workers may have systematic preferences over jobs in different industries by including city fixed effects, ν_c and industry effects, $\nu_{i,t}$. Moreover, we allow for general time effects and for changes in returns to human capital over time.

Our main interest in (10) is in the sign and size of β_2 . The model in the previous section points to a positive effect of $R_{c,t}$ on individual wages (i.e., $\beta_2 > 0$) both because it provides workers with a higher outside option in wage bargaining and because of efficiency enhancing effects arising from low productivity firms dropping out of production. That is, cities with an industrial composition that is more heavily weighted toward high wage premia industries will have higher wages across other industries. It is worth emphasizing that the industrial premia used in constructing $R_{c,t}$ are estimated using national, not city, level data. Hence, even if a city had higher wages in all industries, this would not translate into a higher $R_{c,t}$.¹⁰ Note, moreover, that the composition effect could be sizeable even if total employment in the MC sector is quite small since what is important is not the size of these industries but the fact that they are present in the market and thereby affect the bargaining position of workers in other industries. Even though an outside option would not remain very attractive

¹⁰One implication of this is that high $R_{c,t}$ cities are not necessarily high average skill cities.

if all workers were to simultaneously exploit it; in the bargaining process, all workers can use it as a default option without it losing value.

It is helpful to contrast Equation (10) with the analogous equation that would be derived from a more standard model. To this end, consider a situation where wages are equal to the value of marginal product in each industry and where there are no differences in technology within an industry across firms or cities. If these industrial technologies require labor and capital then, assuming firms have access to a common capital market, this implies that there is effectively constant returns to scale for labor at the city level. To see this, consider the case where the production function for good i is of the form $y_i = K_{ict}^\gamma (\phi_{it} \hat{l}_{ict})^{1-\gamma}$, where K_{ict} is the capital stock used to produce good i , \hat{l}_{ict} is the human capital weighted sum of labor used in sector i ($\hat{l}_{ict} = \sum \exp(\beta_{1t} x_{kict} + \epsilon_{kict})$), ϕ_{it} represents labour augmenting technology and $0 \leq \gamma < 1$. In this case, if firms can rent capital on a common market at price r_{it} then the wage paid to individual k in industry i in city c is of the form

$$\ln W_{kict} = \nu_{it} + \beta_{1t} x_{kict} + \epsilon_{kct} \quad (11)$$

where the time and industry varying intercept ν_{it} corresponds to a combination of rental cost, technology and price effects defined by $\nu_{it} = \ln(1 - \gamma)\gamma^{\frac{\gamma}{(1-\gamma)}} - \frac{\gamma}{(1-\gamma)} \ln r_{it} + \ln \theta_{it}$. The main difference between Equation (11) and Equation (10) is that the industrial composition of employment does not determine industry specific wages in Equation (11), while it does in Equation (10). In fact, in this more standard setup, there is no systematic source for across city variation in industry specific wages. The systematic sources of variation in wages are concentrated at the industry-cross-time level. It is worth pointing out that this same conclusion would be reached using a standard Heckscher-Ohlin model with trade and assuming the local economies are within the cone of diversification: with the implied factor price equalization, industry specific wages should not vary with local employment distributions [Bernard, Redding, and Schott 2005].

This example is of special interest since it delivers exactly the type of economic structure which validates the simple accounting approach to the evaluation of the effects of change in industrial policy on average wages. To see this, note that the standard accounting approach involves first estimating an equation similar to Equation (11) and recovering the industry specific intercepts ν_{it} . These estimated industry coefficients correspond to the inter-industry wage differentials. Then, the accounting approach consists of computing the term,

$$A_{ct} = \sum_i \nu_{it} \left[\left(\frac{l_{ict+1}}{\sum_i l_{ict+1}} \right) - \left(\frac{l_{ict}}{\sum_i l_{i,c,t}} \right) \right]$$

which shows how the average (log) wage in a city changes with the change in industrial composition when using a given set of industry wage premia. The structure underlying (11) implies that changes in the local industrial composition of employment do not alter the industry premia and thus, 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 (10) 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 (log) wage in a city becomes,

$$\sum_i \nu_{it} \left[\left(\frac{l_{ict+1}}{\sum_i l_{ict+1}} \right) - \left(\frac{l_{ict}}{\sum_i l_{ict}} \right) \right] + \beta_2 (R_{ct+1} - R_{ct}) = A_{ct} + \beta_2 (R_{ct+1} - R_{ct}).$$

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

$$R_{ct+1} - R_{ct} = A_{c,t} + \sum_i (\nu_{it+1} - \nu_{it}) \left(\frac{l_{ict+1}}{\sum_i l_{ict+1}} \right) \quad (12)$$

where we have exploited the fact that $\frac{\bar{w}_{it}}{\bar{w}_{1t}} - 1 = \nu_{it}$. Hence the average (log) wage change in a city can be expressed as

$$(1 + \beta_2) A_{c,t} + \beta_2 \sum_i (\nu_{it+1} - \nu_{it}) \left(\frac{l_{ict+1}}{\sum_i l_{ict+1}} \right)$$

It is worth emphasizing that in the case where ν is not varying over time, then the change in $R_{c,t}$ is exactly the amount the accounting approach attributes to changes in industrial composition. This observation provides insights on how to interpret the β_2 coefficient one would retrieve from estimating Equation (10) across the whole set of industries. This relationship highlights that the estimated β_2 represents the spill-over effects of industrial composition on average wages – above that accounted for by the accounting approach– calculated in multiples of the accounting effect. For example, if β_2 were estimated to be equal to .5 then the total effect of changes in industrial composition on average wages should be taken to be 1.5 times the effect implied by a simple accounting approach that neglects any possible spill-over effects. This observation will be useful for evaluating the economic importance of any estimates of spill-over effects that we may obtain.

2.1 Empirical Implementation

Our baseline empirical specification is given by the first difference of Equation (10). Our goal is to investigate the null hypothesis that $\beta_2 = 0$ or, in other words, whether a specification such as Equation (11) (and the factor price equalization theory that underlies it) provides an appropriate description of wage determination in local economies. Support for this null hypothesis would indicate that the standard accounting procedure can be used to properly evaluate the effects of a local change in the composition of employment on the local average wage. Our alternative hypothesis is that $\beta_2 > 0$. A finding of $\beta_2 > 0$ would indicate the presence of a spill-over from industrial composition to wages as predicted by the model and would indicate that the standard accounting approach is an inappropriate means of evaluating the effects of changes in industrial composition on average wages. However, we recognize that such spill-over effects may also arise for reasons other than those emphasized in the model, and we will explore such possibilities.

When estimating the effect of R_{ct} on wages, we need 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 control for such issues, we introduce measures related to these explanations as additional covariates, $z_{c,t}$, in most of our estimations. As we discuss in detail below, we also address the important issues of endogeneity of $R_{c,t}$ and the potential for worker mobility to cause a sample selection bias.

The actual estimation procedure we use throughout is a common two stage procedure. In the first stage, we regress log wages on a vector of individual characteristics separately for each Census year, forming industry \times city group averages of the residuals from the regression. We then use differences in those averages as the dependent variable in our main estimating regression, which takes the following form

$$\Delta \log W_{ict} = \Delta \nu_{it} + \beta \Delta R_{ct} + \Delta z_{ct} \delta + e_{ict} \quad (13)$$

where Z_{ct} is a set of additional city level control variables.

3 Data and Basic Results

3.1 Data

The data we use in the following investigations come from the 1970, 1980, 1990 and 2000 US 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 report results using hourly wages). We deal with 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 A.

As we described in the previous section, we carry out our estimation in two stages. 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 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. We then calculate averages of the residuals for each industry/city combination in each year and use those as the dependent variable in the second stage regression (equation (13) above). We eliminate all industry-city cells with fewer than 20 included individuals in any of the years. We use standard deviations for the constructed industry-city means to form weights for the second stage estimation. For most of our estimates, we used 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. In all the estimation results we calculate standard errors allowing for clustering by city and year.

The main covariate in our estimation is the R_{ct} variable which is a function of the 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.

3.2 OLS Results

We begin our presentation of results with the estimates from specification (13) without the inclusion of any additional control (Z_{ct}) variables. The first column of Table (1) contains the results from OLS estimation of the regression. This regression and all of those that follow include a full set of industry dummy variables (144), thus allowing for changes in industry premia over time, but we do not present the long list of corresponding coefficients here. The coefficient on the change in R variable is 2.62 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, A_{ct} . 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 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 measured from 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, let us 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 US 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 cities, such as those that faced particularly large re-adjustment after the difficulties in the domestic automobile industry. To examine this, in Figure (1) we plot the change in city average wages (the average of our dependent variable across industries within a city) against ΔR_{ct} .¹¹ 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 with which wages are particularly sensitive to the presence of high premia industries. For example, service sector workers might be particular beneficiaries of having more high paid workers in a city through a simple demand route. 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 β_2 coefficient with an i subscript. In Figure (2), we present a histogram of the full set of these β_{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. Thus, any explanation for the impact of shifting the industrial composition toward higher premium industries must have the character that it predicts wide ranging impacts, not just spillovers to industries that are in some sense close to the high premium industries. 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.

4 Addressing Endogeneity and Selection Issues

4.1 Endogeneity: Methods and Results

The simple OLS regression in the previous section points to an estimate of β_2 that is substantial and positive. However, that coefficient may not equal zero even when there are no

¹¹We actually first regress ΔW_{ict} on industry dummies and plot the weighted average of the residuals from that regression in order to obtain a plot that replicates our actual regression.

general equilibrium effects of the kind we are considering because of some combination of worker selection across locations and/or endogeneity issues related to the R_{ct} measure.

To understand the potential endogeneity issues, return to the decomposition of the over-time movements in R_{ct} given in (12). The R_{ct} measure moves with shifts in the set of local employment shares, l_{ict} . One might expect local demand composition to affect the set of local industrial wages since the wage in a given industry would likely rise with increased demand in other industries that employ a closely related set of skills. This would imply a non-zero coefficient on ΔR even in the absence of broader general equilibrium effects of the kind described in the model in the first section.

We respond to this potential problem using three sets of instrumental variables for ΔR . 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, $\hat{s}_{ict} = \frac{\hat{l}_{ict}}{\sum_i \hat{l}_{ict}}$, for the city in period $t + 1$ and form a measure given by:

$$IV1_{ct} = \sum_i \nu_{it} (\hat{s}_{ict+1} - s_{ict}) \quad (14)$$

where, s_{ict} is the share of employment in city c in time t that is in industry i . This variable is closely related to the first term in the decomposition of the R measure given in (12). 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 shift 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.

Our second instrument is designed to isolate the variation inherent in the second term in the decomposition, (12): 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:

$$IV2_{ct} = \sum_i s_{ict}(\Delta\nu_{it}).$$

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 power, is correct then it should not matter whether the average premia available in the city declines because a high paying industry shuts down or because the premium paid in that industry declines.¹² 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 provides a means of evaluating whether the outside option effect outlined in Section 2 is the likely mechanism at play.

We also implement a third instrumenting strategy based on the fact that both *IV1* and *IV2* are functions of initial period employment shares, s_{ict} . Rather than restrict ourselves to specific functions of those shares, as we do in the first two instruments, our third approach is to use the set of shares in 21 aggregate industries in the initial period as our instruments.¹³

All three 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 takes a value of 81.9 and has an associated p-value of 0.0. The same statistic for *IV2* is 140 with a p-value of 0 and for *IV3* is 5.5 also with an associated p-value of 0.0.

We present results from instrumental variables estimation using each of our instruments individually, in the last three columns of Table (1). Thus, the second column contains results from instrumental variable estimation in which we use *IV1*, our instrument which uses national level variation in employment shifts. The estimated coefficient is very similar

¹²Indeed, in the model in the first section, changes in the holdup problem stemming from $\{k_i\}$ affect the alternative wage for workers in the FE sector both through changes in the national level wage premia and in the local composition of employment.

¹³As we described earlier, our estimation is based on 144 industries and 152 cities. Aggregating to 21 industries for this instrument allows for a reasonable number of degrees of freedom in the first stage regression.

to that obtained from OLS estimation and is again highly statistically significant. The third column contains results when we use $IV2$, the instrument that uses changes in industry premia over time. The similarity in the estimated coefficients obtained using $IV1$ and $IV2$ is striking; an outcome which we have argued fits well with theories of the impact of R that are based on changes in bargaining power. Thus, for example, the decline in the steel industry in what came to be called the Rust Belt resulted both in lower employment in steel jobs and lower wages within the steel industry over time [Beeson, Shaw, and Shore-Sheppard 2001]. The results in columns 2 and 3 of Table (1) indicate that both types of change had the same impact on wages within other industries, as implied by our bargaining story.

The final column of Table (1) provides estimates when we use our least restrictive instrument, $IV3$. The estimated β_2 coefficient is somewhat smaller than from the other specifications but still implies that the total effect of a change in composition is 2.8 times that obtained from a simple accounting approach and is still highly statistically significant.

4.2 Selection: Methods and Results

Our second key concern 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.¹⁴

¹⁴It is interesting to consider this type of selection issue in the context of a model with no GE effects such as a simple Hecksher-Ohlin model. In such a model, shifts in the distribution of an unobserved (to the researcher) skill across cities (perhaps because of changing valuations of amenities) does not imply that the R measure should enter our regressions significantly since factor price equalization should still hold and, thus, local shares of any skills will not change local wages. Suppose, however, that one of these skills, which we will call the comparative advantage ability, is unobserved (so that our estimated industry premia partly reflect returns to that skill) and, in addition, there is an absolute advantage skill which allows a worker to be better in all jobs. Now suppose that there is an amenity in a given city that attracts both high comparative advantage and high absolute advantage workers since they can afford to pay for the amenity. In that situation, the industrial mix in the high amenity city will shift toward industries that use the comparative advantage ability intensively (implying an increase in R) and, at the same time, the number of high absolute advantage workers will increase. This is a case where the estimated β_2 coefficient will be non-zero even though the

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 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. This implies that we should write the regression corresponding to observed wages as,

$$E(\ln W_{kict} | d_{kct} = 1) = \beta_{0t} + \beta_{1t}x_{kct} + \beta_2R_{ct} + \nu_i + \nu_c + E(\epsilon_{kct} | d_{kct} = 1) \quad (15)$$

where 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 (13) not taking account of this error mean term then the estimated regression coefficients will suffer from well-known consistency problems.

In situations such as 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 the 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

basic trade model (with the addition of the absolute advantage ability) holds, simply because the R measure is correlated with unobserved ability.

that we include immigrants in our analysis and that we are dealing with cities. Details on our selection estimation are provided in Appendix B. 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.¹⁵

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, experience, race, immigrant status and gender variables as before but now also add our proxies for the error mean term. We then form averages by city and industry from the residuals from that regression 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 (2), we recreate the results from Table (1) while implementing Dahl (2002)’s selection correction. The resulting estimates for β_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. Nonetheless, in all subsequent sections of the paper, we present results incorporating Dahl’s sample selection correction. 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.

¹⁵Note 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.

5 Further Explorations of the Wage Premia Effects

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 capturing one of these alternative driving forces.

Possibly the most intuitive explanation of differential wage growth across cities is differences in aggregate demand. Our estimated β_2 coefficient, for example, may just be picking up spill-overs in demand for products of related industries when a particular high wage industry re-locates to a town. Or, more generally, such a re-location will increase the general demand for workers in an area, driving up wages across the local economy. Blanchard and Katz [1992] examine the implications of demand shifts for local wages and employment using US data and argue that negative demand shocks in a locality lead to a permanent shift down in employment and a negative adjustment in wages that eventually dissipates. Greenstone and Moretti [2003] examine the impact on wages and housing prices of attracting a large manufacturing plant to a town by comparing outcomes in towns that won a competition for such a plant to towns that made the short list in the same competition but ultimately lost. Their results indicate that both the wage bill in the industry in which the plant is situated and local housing prices rise in the three year period following the arrival of the plant. While it is not the focus of the paper, Greenstone and Moretti also show that the arrival of the plant has a positive impact on the wage bill in other industries in the receiving town. It is difficult to compare our results to those in Greenstone and Moretti [2003] since the wage bill effects they estimate are a combination of employment effects and the wage effects we are measuring. Nonetheless, it is interesting that they also find evidence of spill-overs.

Our approach differs from both Blanchard and Katz [1992] and Greenstone and Moretti [2003] in that we focus on longer term impacts of differences in the composition of employment while both of those papers focus on adjustments to demand shocks. One way in which we try to insure we are capturing composition effects is through the construction of our R_{ct} measure. In particular, our R_{ct} measure is based on employment shares and so does not directly increase with increases in employment. Nevertheless, changes in our R_{ct} measure may

be correlated with local demand shifts, implying that this is really what we are capturing. This turns out not to be the case: as we show in Section 6, in a simple regression of changes in employment rates on the change in R_{ct} , the latter has a coefficient that is essentially zero and is statistically insignificant at any conventional significance level. Nonetheless, since this is such a prominent alternative explanation, we adopt several strategies to examine it more closely.¹⁶

Our first strategy is to construct a measure of changes in local demand which equals the sum of industry specific growth rates in each city, with each industry's growth rate weighted by its share in city employment at the start of the decade. This essentially states that total growth in employment reflects general increases in demand regardless of the sector in which the increase occurs. In contrast, the R_{ct} measure emphasizes movements that favour higher premium industries. In the specifications where we instrument for R_{ct} , we also instrument for this change in local demand measure in a manner similar to *IV1*. In particular, we use growth rates for the industry at the national level weighted by initial period employment shares in the particular city as an instrument.

We also include the change in own-industry employment share as a separate regressor in the specification in order to control for potential problems that might arise if movements in R_{ct} are dominated by a particular industry in a city, with our estimated coefficient then picking up the relationship between that industry's shifts and its own wage. We again instrument for it using national level changes in the industry share in the specifications where we instrument for the other two variables.

The results from this exercise are given in the first three columns of Table (3). The first column is estimated with OLS. In the second column we use an instrumental variable estimator in which the instruments used are *IV1* and the instruments just described. Columns 2 and 3 contain results from specifications in which we use *IV2* or *IV3* and the other instruments. The results indicate that including the overall demand measure does affect the coefficient on ΔR . This fits with the finding that changes in R_{ct} and changes in employment are almost orthogonal.

In the second set of columns in Table (3) we use a simpler measure of demand, replacing the overall demand measure with the local employment rate. This differs from using the

¹⁶An adjustment to the model from section 1 in which we allow shifts in employment in the MC industries to affect the employment rate in a city as well as the industrial composition of employment implies further, positive impacts on wages from an expansion of high paying industries. Thus, including the employment rate as an overall demand measure can be seen as capturing this added dimension.

overall demand measure in that the latter reflects changes in city size that are not captured in the employment rate. Again, controlling for the employment rate helps emphasize the point that the estimated coefficient on R_{ct} corresponds to the impact of shifts in employment composition not the level of employment. The employment rate itself enters positively and statistically significantly, implying that wages increase in cities with increasing employment. Such an effect can be easily shown to be consistent with our bargaining framework. Moreover, the β_2 coefficient remains the same size and significance when this measure is introduced. The coefficient on own employment share is again small and typically insignificant.

An alternative explanation for city level growth is provided in Glaeser *et al.* [1992]. They examine city level growth over time in the US, 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 the first three columns of Table (4), 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 significant in our estimates and, more importantly, does not change our estimates of the β_2 coefficient.

Finally, a recent literature on education externalities 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 US 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 city. 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 (4), 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} .¹⁷ In the next set of columns in Table (4), we

¹⁷It is worth noting, though, that Sand [2006] finds that this positive and significant impact is observed

use 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.

In the last set of columns of Table (4) we introduce the overall demand measure, the college share measure and the fractionalization measure at the same time. Again, there is little impact of introducing these variables on the estimated β_2 coefficient. Our conclusion is that while some of the other hypothesized factors 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 larger than any of the effects from these competing explanations.

5.2 Robustness Checks

In this sub-section, we examine the robustness of our estimates to a series of variations in our estimation approach. The first of these relates to our definitions of cities and industries. The use of data from 1970 through 2000 is a strength in that it allows us to see whether the effects we are measuring show up over a long period of time. However, the downside of using such a long period is that we are forced to be restrictive in our definitions of cities and industries in order to have consistency over the whole period. In particular, we are forced to drop some cities because we cannot create a consistent definition for them over the whole time period¹⁸ and we are forced to base our industry categories on a 1950 classification scheme that might not be well suited to capturing shifts across industries in a more recent era. To check on the sensitivity of our estimates to these issues, we re-estimated our main specifications using data only from the period 1980 to 2000. Dropping 1970 allows us to increase the number of consistently defined cities from 152 to 231 and allows to increase the number of industry categories from 144 to 221. We also shift to a 1980 Census definition of industries. Using the 1980 - 2000 data, we again estimate in two stages, with the initial stage being individual wage regressions run separately by Census. We then form 10 year differences in the regression- adjusted city-industry average log wages and use those in the second stage regression, pooling the 1980-90 and 1990-2000 differences. The results are found in Table (5), which is a recreation of Table (1) based on this alternate sample. The estimated

in the 1980s but not in the 1970s or 1990s when estimation is carried out separately by decade.

¹⁸This restriction is not related to rapidity of growth of cities but rather to inconsistencies in the definitions of our building block geographical units (counties and PUMA's) across Censuses.

effects are very similar in magnitude to those in Table (1), with the 1970-2000 sample based estimates being larger in some instances and the 1980-2000 being larger in others. The key conclusion, though, is that our results are not sensitive to the industry and city restrictions imposed due to using our longer time period.

In Table (6), we provide estimates from our basic specification run separately for each decade. To save on space, we provide estimates using only OLS, IV1 and IV3. The estimated effects do differ somewhat by decade. In particular, the β_2 estimates are larger for the 1980-90 decade than the 1970-80 decade, but both are still substantial. The OLS and IV3 estimates for the 1990-2000 are similar in size to those from the 1970-80 decade and, thus, indicate that there continues to be substantial impacts from shifts in the composition toward higher paying industries across time. However, the IV1 estimate for the most recent decade is very poorly defined and actually takes a negative sign. Underlying this is smaller variation in ΔR_{ct} across cities in the 1990s, and a less good fit is the first stage regression.

We also consider two types of non-linearities in our estimation. First, our R_{ct} measure is a linear function of the national level wage premia. It seems possible, however, that an increase in R_{ct} stemming from gaining a very high wage industry might have a differential impact relative to the same size increase generated from adding a larger industry which pays a wage only slightly above the existing average wage in the city. In particular, adding a very high wage industry may have greater salience in the thinking and bargaining of workers. We address this possibility in Figure (3). To construct this figure, we divide industries into thirds based on their associated, national level wage premia in 1970.¹⁹ In the top left panel of the figure, we plot changes in average wages in the top third industries in each city against the change in the share of employment in the top third industries. We also plot the simple regression line through the cloud of observations. The regression line is strongly positively sloped, with a slope coefficient of 0.69 and an associated standard error of 0.08. This strong impact of an increase in the share of employment in top industries is also seen in the bottom left panel, which shows the relationship between the bottom third industry wage changes and top third employment share changes. Thus, as emphasized in Figure (2), increasing the share of employment in top paying industries has similar impacts across industries, but Figure (3) makes it clearer that changing share in one part of the wage structure can affect wages in another part. The right two panels show the associations of wage changes with changes in the share in the lowest third industries. The key point here is that changes in

¹⁹The results are not sensitive to using alternative base years to divide industries into thirds.

this part of the structure produce strong negative impacts on wages. Thus, it is not the case that our main results are being driven just by gains or losses within one subset of industries: as the slope coefficients shown in the figure demonstrate, changing average wages either by gaining top paying industries or losing low paying industries has similar effects.

We also investigated possible non-linearities in the impact of R_{ct} on wages. In particular, one might imagine that the marginal impact of adding one more high premium industry in an already high-wage city is less than the impact of adding the same industry to a low-wage city. In Table (7), we present estimates from our basic specification but adding a squared change in R_{ct} term. The squared term enters statistically significantly and with a positive sign in all cases. In Figure (4), we plot the predicted profiles based on the OLS, IV1, IV2, and IV3 estimates and actual changes in R_{ct} from our sample.²⁰ The plot indicates that the estimated relationship is actually very close to linear in the relevant range in spite of the significance of the quadratic terms. This is particularly true for the IV estimates.²¹

6 Narrowing Down the Set of Explanations for the Effects of Changing Industrial Composition

To this point, we have established that shifts in industrial composition toward higher paying sectors has an impact on wages in almost all industries, it 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 investigating what might underlie the estimated effect.

As a first step, it is interesting to investigate the educational dimension of any response. In particular, it would be useful to know whether “good job” impacts are located entirely in the low education labor market. To examine this, we re-estimate our main specification separately for workers with a high school or less education and for workers with BA or higher education.²² We create the national level wage premia, and thus the R_{ct} measures,

²⁰More specifically, to derive this plot, we estimate the regressions in Table 7 then set all coefficients to zero except for the R_{ct} variables and the constant. We then predict values for the whole sample based only on those coefficients and plot them.

²¹ As an additional check, in Table (12) we report results obtained when using an hourly measures of wages as our dependent variable, as opposed to an weekly measure. As can be seen, results are robust to using either measure.

²²The R_{ct} measure is the same as that in earlier tables, i.e., it is constructed using industrial premia

separately for each education 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 two education groups. The results from this exercise are presented in Table (8). While the estimated β_2 coefficient varies slightly between those with a high school education and those with a college education, the differences across education groups are not large. Whatever we are measuring, it is not restricted to a subset of labour markets defined by skill.

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 (under a standard Heckscher-Ohlin model) 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.²³ In Table (9), 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.

estimated for all workers in a regression including education, experience, gender, race and immigrant status controls.

²³The 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.

Thus, our results do not appear as arising simply because of spill-overs into the non-traded goods sector labour market. In Table (10) we also present results of estimating β for twelve industry grouping. As can be seen, results are very similar across these industries, with wages in manufacturing industries responding close 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.²⁴

Another potential explanation for what we are measuring is that workers' bargaining power increases when the average wage premia in other industries goes up and that, as a result, they are able to take rents away from quasi-fixed factors. This would imply, for example, that returns to capital should decline locally when better firms come to town. This differs from the explanation in our initial model in that we would not expect such effects to persist: essentially it should correspond to inducing a redistribution in the short run. Unfortunately, we do not have data to examine this directly but we can get at it indirectly by estimating using longer differences. The idea is that the quasi-fixed factors should be able to readjust in response to the wage increases over a 20 year horizon even if they cannot respond effectively within 10 years. In Table (11), we present results from estimating our base specification using 1970-1990 differences and 1980-2000 differences. The estimated composition effects are somewhat smaller in the 1970-90 case than in our 10 year difference results but they still imply large and significant effects of nearly the same size as those seen in the 10 year difference estimates. Thus, a simple redistribution story does not seem to be driving what we find.²⁵

²⁴Interestingly, the sector where wage 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.

²⁵The case of Pittsburgh over the period from 1980 to 2000 is a good example of the patterns we are finding over both a ten year period and a 20 year period. During the period from 1980 to 1990, the average wage in Pittsburgh fell by over 11% relative to the average wage across US cities. When we control for changes in the educational attainment of the population, we obtain a similar fall in wages, indicating that educational attainment in the Pittsburgh labour market resembled that of the nation. As is well known, Pittsburgh has traditionally been highly concentrated in the steel industry, and during the 1980-90 period this industry did very poorly. In large part due to these changes in the steel industry, our measure of labour market rents for Pittsburgh, $R_{c,t}$, fell by 3.3% over the eighties relative to that experience by the average US city. Hence, from a purely compositional view, changes in the fraction of high paying jobs in Pittsburgh appear to account for a small fraction of the change in wages, leaving the majority of the fall (7.7%) unexplained. Our results suggest that the pure compositional approach is invalid since it omits the spill-over effects of good jobs on wages in other industries. In contrast to the compositional view, our estimates of the effect of industrial composition indicate that virtually all the fall in wages in Pittsburgh over the eighties can be attributable to the change in industry structure ($3.5 \times 3.3 = 11.5$). Interestingly, we can then look beyond the 1980-90 period to see whether wages in Pittsburgh reversed themselves. During the nineties, wages and industrial composition changes fared much better for Pittsburgh. In terms of changes in the rent measure, $R_{c,t}$, Pittsburgh experienced a change very close to the national average over the nineties. Accordingly, the average wage growth was also much better, but it still did not surpass the national average (it was

What, then, are we left with? The effect we measure is not restricted to one educational group, is present for tradeable goods and does not appear to be just a short term redistribution toward workers. The only way for these all to be true - for the effect to be long lasting even in tradeable goods markets - is if it is accompanied by an increase in productivity for workers. The message of the model we presented at the outset is that models of the type found in the recent heterogenous firm literature, with labour market frictions and firm heterogeneity, provide a potential explanation for how these productivity effects can arise as a market equilibrium. Further, as we argued earlier, the fact that we find that our measured effects are the same whether we use variation from changes in the share of employment in high wage industries or changes in the premia paid in high wage industries fits with the type of bargaining-based mechanism emphasized in the model.

6.1 Effects of changes in industrial composition on the distribution of wages

In addition to the implications of industrial composition on average wages, the theory presented in section 2 also has implications for the distribution of wages. In particular, the theory suggests that in a city that experiences a change in industrial composition towards higher paying jobs, the improved bargaining position of workers should cause the least productive firms to leave the market, thereby eliminating the lowest paying jobs. Workers are then reallocated to the more productive jobs, leading to a decrease in wage inequality. We explore this implication by examining the relationship between changes in industrial composition and changes in the different deciles of the city level residual wage distribution. The results are reported in Table (13), where we again pool data for the three decade differences, 70-80, 80-90 and 90-00. The dependent variable is the change in the the residual wage decile over a decade, where the wage residuals are obtained after regressing individual level wages on our set of individual level variables discussed previously (age, schooling, gender,..) and a full set of industry dummies. In contrast to our previous results, our dependent variable is not observed at the city-industry level since many city-industry cells are too small to use in an analysis of deciles. Instead, in this section, we rely on city level variation, controlling for industry effects by adding industry dummies in the first stage.

approximately 2% below the national average). Hence, there is no indication that in the nineties Pittsburgh experienced a reversal of the wage effects experienced in the eighties; Pittsburgh simply experience a period of growth similar to the rest of the country. This anecdote is consistent with the fact that we estimate similar effects of changes in industrial composition when we look over a 10 year period or a 20 year period.

As can be seen in Table (13), the estimated impact of a change in R_{ct} on the median wage is very similar to the results in previous tables where we used the mean wage. The estimates are close to 2.5 and highly significant. The more interesting aspect to note is the pattern of the coefficients across deciles. For all four series of β s, the coefficients are substantially larger for the lower deciles than for the higher deciles, indicating that changes in R_{ct} lead to a wage distribution with lower inequality. For example, when we use IV1 as an instrument, the estimated impact of a change in R_{ct} on the 10th decile is more than twice the impact on the 90th decile. Results based on IV3 shown even more compression. Although not reported, we have verified the robustness of this pattern with respect to adding additional regressors and dividing the sample into education sub-groups. In all these cases, we find that cities which experience a change in industrial composition in favor of better paying jobs also experience a reduction in residual wage inequality, which is consistent with the theoretical mechanism presented earlier. It is worth noting that most of the decrease in city level inequality we observe is concentrated in the bottom half of the distribution, with the 50-10 decile difference declining by much more than the 90-50 decile difference in response to a change in R_{ct} . When we divide the sample between education groups, we find the compression effects to be slightly greater for high school educated workers than for college educated workers. Although such reduced inequality could be driven by several different mechanisms,²⁶ the reallocation of workers across jobs with different levels of productivity offers a simple interpretation.

6.2 Additional effects associated with changes in industrial composition

Up to this point, our empirical investigation has focussed 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 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 (14) 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. We again observe a positive

²⁶ For example, in addition to the reallocation effects emphasized by the model, improvements in worker bargaining power could lead to within firm increases in productivity if firms are not always using the cost minimizing technology.

association, although it is less robust than that for weekly wages. 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 fraction of the changes in wages.²⁷

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 10, 11 and 12 in Table (14). 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 $R_{c,t}$ 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.

The model we presented in Section 2 focussed exclusively on the potential wage effects of changes in industrial composition. However, in models without firm heterogeneity, the common adjustment mechanism associated with a change in industrial composition would be through changes in the unemployment rate (or the employment rate). For example, a more conventional adjustment mechanism associated with an increase in high paying jobs is an increase in the fraction of individuals who choose to queue for the better jobs. If such an adjustment mechanism is present, we would expect an increase in $R_{c,t}$ to be associated with increases in a city's unemployment rate, or a decrease in its employment rate. In columns 4-9 of Table (14), we explore the relevance of this alternative adjustment mechanism. We find very little evidence in support of the view that adjustments in queue lengths for goods jobs – as measured by either changes in a city's rate of unemployment or the rate of employment – is an important equilibrium mechanism associated with a change in industrial composition. Such a finding is quite surprising and is in sharp contrast with the wage effect we have documented. Together our findings, therefore, suggest that it is wages that adjust following a change in outside options induced by a change in industrial composition, while almost none of the adjustment takes the form of higher unemployment.

²⁷ Since rents makes up only a fraction of total consumption, under perfect mobility of workers across cities we would expect that the effect of a change in R_{ct} on housing prices would be greater than the effect on wages.

7 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 focussing 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 spill-over effects from the loss of jobs in one sector to wages in other sectors. Although traditional economic theory provides good reason to believe that such spill-over effects should be absent or small, recent empirical and theoretical developments emphasizing firm heterogeneity and labour market frictions suggests a need to reexamine the issue.²⁸ In this paper, we build on recent theory to highlight an empirical strategy for evaluating the spill-over effects of changes in job composition on industry level wage payments. We implement that strategy using US census data from 1970 to 2000. Our main finding is that spill-overs 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 affects industrial composition should not be evaluated simply using the standard accounting approach but instead should explicitly take account of substantial spill-over effects. For example, it is common for 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 spill-over effects on wages in other sectors. In general, recognizing and quantifying these feedback effects will lead to much more variable assessments of the gains from trade since markets that attract high paying industries will benefit more than traditionally thought, while markets that lose such jobs should benefit less.²⁹

²⁸The mechanism presented in Acemoglu [2001] provides an alternative rationale to explore such a issue.

²⁹ Beaudry, Green and Collard [2005] and Beaudry and Collard [2006] find that increased openness to international trade over the period 1978-98 had vary 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 spill-over effects found in this paper offer a potential explanation to the size of the effects found in these two papers.

References

- ACEMOGLU, D. (2001): “Good Jobs versus Bad Jobs,” *Journal of Labor Economics*, pp. 1–22.
- 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.
- 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.
- BEESON, P., K. SHAW, AND L. SHORE-SHEPPARD (2001): “Industrial Change and Wage Inequality: Evidence from the Steel Industry,” *Industrial and Labor Relations Review*, pp. 466–483.
- BERNARD, A., S. J. REDDING, AND P. SCHOTT (2005): “Factor Price Equality and the Economies of the United States,” CEPR Discussion Papers 5111, C.E.P.R. Discussion Papers.
- BERNARD, A. B., J. EATON, J. B. JENSEN, AND S. KORTUM (2003): “Plants and Productivity in International Trade,” *American Economic Review*, 93(4), 1268–1290.
- BLANCHARD, O. J., AND L. F. KATZ (1992): “Regional Evolutions,” *Brookings Papers and Economic Activity*, (1), 1–75.
- BLUESTONE, B., AND B. HARRISON (1982): *The Deindustrialization of America*. New York: Basic Books.
- BORJAS, G. J., AND V. A. RAMEY (1995): “Foreign Competition, Market Power, and Wage Inequality: Theory and Practice,” *Quarterly Journal of Economics*, pp. 1075–110.

- BOUND, J., AND G. JOHNSON (1992): “Changes in the Structure of Wages in the 1980’s: An Evaluation of Alternative Explanations,” *American Economic Review*, pp. 371–392.
- BRIEM (2002): “It Was 20 Years Ago Today,” *Pittsburgh Post Gazette*.
- 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.
- GLAESER, E., H. KALLAL, J. SCHEINKMAN, AND A. SHLEIFER (1992): “Growth in Cities,” *Journal of Political Economy*, (6), 1126–1152.
- 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.
- 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.
- JENSEN, B., AND L. KLETZER (2005): “Tradable Services: Understanding the Scope and Impact of Services Outsourcing,” Working Paper 05-9, Institute for International Economics.
- LEE, L. (1983): “Generalized Econometric Models with Selectivity,” *Econometrica*, 51(2), 507–512.
- MELITZ, M. J. (2003): “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, (6), 1695–1725.
- MORETTI, E. (2004): “Estimating the social return to higher education: evidence from longitudinal and repeated cross-sectional data,” *Journal of Econometrics*, 121(1-2), 175–212.

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

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. M. (2006): “The Social Returns to Education,” University of British Columbia Working Paper.

A 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 US Office of Management and Budget.³⁰ 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 what counties to include in each city for the other years. Since the 1970 county group definitions are much coarser 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

³⁰See <http://www.census.gov/population/estimates/pastmetro.html> for details.

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.³¹ 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.³² We are also able to define more (231) consistent cities for that period.

B 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 (15) 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 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

³¹See <http://usa.ipums.org/usa-action/variableDescription.do?mnemonic=IND1950> for details.

³² 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>

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".³³ 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.³⁴ 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. Intuitively, 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)

³³For 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.

³⁴Specifically, we use single, married without children, and married with at least one child under age 5.

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 US, 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.

Table 1: 1970-2000

	OLS	IV1	IV2	IV3
Variables	(1)	(2)	(3)	(4)
$\Delta R_{c,t}$	2.622 (0.185)*	2.690 (0.343)*	2.604 (0.356)*	1.843 (0.33)*
Const.	0.287 (0.028)*	0.297 (0.049)*	0.285 (0.052)*	0.175 (0.048)*
Obs.	28142	28142	28142	28142
R^2	0.537	0.537	0.537	0.534

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

Table 2: 1970-2000

	OLS	IV1	IV2	IV3
Variables	(1)	(2)	(3)	(4)
$\Delta R_{c,t}$	2.365 (0.178)*	2.457 (0.32)*	2.565 (0.346)*	1.554 (0.38)*
Const.	0.192 (0.028)*	0.205 (0.047)*	0.221 (0.052)*	0.076 (0.056)
Obs.	28083	28083	28083	28083
R^2	0.468	0.468	0.467	0.464

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

Table 3: 1970-2000

Variables	OLS	IV1	IV2	IV3	OLS	IV1	IV2	IV3	OLS	IV1	IV2	IV3
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$\Delta R_{c,t}$	2.352 (0.198)*	2.453 (0.337)*	2.564 (0.361)*	1.554 (0.329)*	2.406 (0.214)*	2.439 (0.39)*	2.558 (0.331)*	3.196 (0.448)*	2.336 (0.192)*	2.469 (0.328)*	2.744 (0.343)*	1.626 (0.33)*
Δl_{ic}	0.702 (0.1)*	0.086 (0.196)	0.084 (0.195)	0.04 (0.21)	0.693 (0.098)*	0.084 (0.201)	0.087 (0.199)	0.009 (0.196)	0.704 (0.1)*	0.061 (0.201)	0.058 (0.199)	0.034 (0.216)
Δ Labour Demand					-0.11 (0.01)	0.002 (0.042)	-0.004 (0.039)	-0.088 (0.027)*				
Δ Emp. Rate									0.276 (0.085)*	0.273 (0.084)*	0.269 (0.085)*	0.288 (0.088)*
Const.	0.191 (0.029)*	0.205 (0.047)*	0.221 (0.052)*	0.076 (0.048)	0.203 (0.033)*	0.202 (0.067)*	0.222 (0.057)*	0.348 (0.07)*	0.18 (0.028)*	0.199 (0.046)*	0.238 (0.05)*	0.077 (0.048)
Obs.	28083	28083	28083	28083	28083	28083	28083	28083	28083	28083	28083	28083
R^2	0.469	0.468	0.468	0.464	0.469	0.468	0.468	0.455	0.47	0.469	0.468	0.466

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

Table 4: 1970-2000

Variables	OLS (1)	IV1 (2)	IV3 (3)	OLS (4)	IV1 (5)	IV3 (6)	OLS (7)	IV1 (8)	IV3 (9)	OLS (10)	IV1 (11)	IV3 (12)
$\Delta R_{c,t}$	2.374 (0.198)*	2.367 (0.307)*	1.552 (0.32)*	2.344 (0.195)*	2.254 (0.326)*	1.600 (0.32)*	2.365 (0.197)*	2.460 (0.339)*	1.553 (0.328)*	2.364 (0.213)*	2.251 (0.358)*	3.001 (0.432)*
Δ BA or >		0.459 (0.152)*	0.473 (0.159)*							0.428 (0.154)*	0.436 (0.149)*	0.308 (0.133)*
Δ Ave. Years School				-0.003 (0.012)	-0.003 (0.012)				-0.003 (0.011)			
1 - Herfindahl	0.278 (0.149)	0.277 (0.151)	0.253 (0.14)							0.13 (0.148)	0.166 (0.149)	0.178 (0.175)
Δ Labour Demand										-0.006 (0.011)	-0.005 (0.037)	-0.076 (0.027)*
Δl_{ic}										0.649 (0.099)*	0.015 (0.211)	-0.099 (0.204)
Const.	-0.074 (0.141)	-0.074 (0.136)	-0.168 (0.135)	0.161 (0.031)*	0.148 (0.049)*	0.054 (0.05)	0.195 (0.03)*	0.209 (0.051)*	0.079 (0.051)	0.044 (0.144)	-0.008 (0.149)	0.125 (0.176)
Obs.	28083	28083	28083	28083	28083	28083	28083	28083	28083	28083	28083	28083
R^2	0.468	0.468	0.464	0.471	0.471	0.468	0.468	0.468	0.464	0.472	0.471	0.461

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

Table 5: 1980-2000: Alternative Sample

Variables	OLS	IV1	IV2	IV3
	(1)	(2)	(3)	(4)
$\Delta R_{c,t}$	2.325 (0.276)*	2.403 (0.277)*	2.941 (0.441)*	1.977 (0.217)*
Const.	-.135 (0.012)*	-.133 (0.012)*	-.124 (0.014)*	-.141 (0.01)*
Obs.	31268	31268	31268	31268
R^2	0.596	0.596	0.594	0.595

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

Table 6: Results by Decade

Variables	1970-1980			1980-1990			1990-2000		
	OLS (1)	IV1 (2)	IV3 (3)	OLS (4)	IV1 (5)	IV3 (6)	OLS (7)	IV1 (8)	IV3 (9)
$\Delta R_{c,t}$	1.639 (0.258)*	3.290 (1.667)*	2.308 (0.307)*	3.050 (0.308)*	3.057 (0.352)*	3.200 (0.353)*	2.499 (0.429)*	-4.167 (2.785)	2.423 (0.635)*
Const.	0.216 (0.04)*	0.45 (0.234)	0.311 (0.045)*	-.122 (0.013)*	-.122 (0.013)*	-.120 (0.013)*	0.134 (0.012)*	-.008 (0.062)	0.132 (0.016)*
Obs.	6391	6391	6391	10810	10810	10810	10941	10941	10941
R^2	0.34	0.305	0.334	0.275	0.275	0.275	0.157	.	0.157

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

Table 7: 1970-2000: Non-linear Specification

Variables	OLS	IV1	IV2	IV3
	(1)	(2)	(3)	(4)
$\Delta R_{c,t}$	3.152 (0.308)*	2.806 (0.479)*	3.626 (0.488)*	2.402 (0.385)*
$(\Delta R_{c,t})^2$	5.471 (1.439)*	3.687 (5.098)	6.602 (2.719)*	7.452 (1.856)*
Const.	0.19 (0.029)*	0.178 (0.068)*	0.235 (0.048)*	0.041 (0.05)
Obs.	28083	28083	28083	28083
R^2	0.47	0.469	0.469	0.464

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

Table 8: 1970-2000: By Education Group

Variables	HS or <			BA or >		
	OLS	IV1	IV3	OLS	IV1	IV3
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta R_{c,t}$	2.591 (0.193)*	2.822 (0.345)*	1.522 (0.41)*	2.193 (0.258)*	2.299 (0.435)*	2.564 (0.528)*
Const.	0.246 (0.03)*	0.279 (0.05)*	0.093 (0.06)	0.116 (0.044)*	0.131 (0.067)	0.17 (0.081)*
Obs.	21564	21564	21564	9233	9233	9233
R^2	0.488	0.488	0.483	0.39	0.39	0.39

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

Table 9: 1970-2000: By Trade and Non-Trade Industries

Variables	Low Trade			Medium Trade			High Trade		
	OLS	IV1	IV3	OLS	IV1	IV3	OLS	IV1	IV3
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$\Delta R_{c,t}$	2.614 (0.405)*	2.970 (0.893)*	1.973 (0.857)*	2.171 (0.228)*	2.487 (0.38)*	1.964 (0.538)*	2.376 (0.16)*	2.352 (0.275)*	1.296 (0.367)*
Const.	0.299 (0.076)*	0.11 (0.166)	-.039 (0.17)	0.153 (0.041)*	0.199 (0.06)*	0.101 (0.098)	0.197 (0.025)*	0.193 (0.041)*	0.042 (0.054)
Obs.	4607	4607	4607	11680	11680	11680	11796	11796	11796
R^2	0.444	0.444	0.443	0.415	0.415	0.415	0.53	0.53	0.523

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

Table 10: 1970-2000: Breakdown by Industry Group

Industry	OLS	IV1	IV2	IV3
	(1)	(2)	(3)	(4)
Manufacturing (Durables)	3.120 (0.196)*	2.548 (0.368)*	3.968 (0.374)*	1.462 (0.482)*
Manufacturing (Non-Durables)	2.471 (0.242)*	2.240 (0.3349)*	2.620 (0.423)*	1.809 (0.415)*
Public Administration	1.292 (0.210)*	1.617 (0.398)*	0.747 (0.339)*	1.741 (0.464)*
Wholesale and Retail Trade	2.969 (0.267)*	3.149 (0.475)*	2.411 (0.505)*	2.350 (0.514)*
Professional	2.128 (0.194)*	2.627 (0.365)*	1.665 (0.358)*	2.128 (0.408)*
Financial, Real Estate, Insurance	2.208 (0.278)*	3.007 (0.506)*	2.057 (0.508)*	1.211 (0.564)*
Agriculture and Mining	2.850 (0.392)*	2.557 (0.565)*	3.258 (0.742)*	2.167 (0.895)*
Construction	3.398 (0.285)*	3.717 (0.502)*	2.542 (0.484)*	2.558 (0.583)*
Transportation and Utilities	2.171 (0.231)*	2.246 (0.397)*	2.179 (0.437)*	1.963 (0.501)*
Others (Private, Business, Entertainment)	3.050 (0.348)*	3.558 (0.665)*	2.180 (0.725)*	3.186 (0.779)*

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

Table 11: 1970-2000: Long Difference

Variables	1970-1990				1980-2000			
	OLS	IV1	IV2	IV3	OLS	IV1	IV2	IV3
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta R_{c,t}$	1.222 (0.302)*	2.007 (0.735)*	7.392 (2.295)*	1.987 (0.541)*	1.944 (0.215)*	2.373 (0.299)*	2.681 (0.396)*	2.251 (0.361)*
Const.	-.087 (0.054)	0.036 (0.114)	0.88 (0.357)*	0.033 (0.086)	-.087 (0.014)*	-.071 (0.016)*	-.059 (0.019)*	-.075 (0.018)*
Obs.	6268	6268	6268	6268	10302	10302	10302	10302
R^2	0.202	0.197	.	0.198	0.272	0.269	0.263	0.27

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

Table 12: 1970-2000: Hourly Wages

Variables	OLS	IV1	IV2	IV3
	(1)	(2)	(3)	(4)
$\Delta R_{c,t}$	2.351 (0.174)*	2.433 (0.328)*	2.332 (0.348)*	1.634 (0.326)*
Const.	0.48 (0.026)*	0.492 (0.047)*	0.478 (0.051)*	0.377 (0.047)*
Obs.	28143	28143	28143	28143
R^2	0.641	0.641	0.641	0.639

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

Table 13: 1970-2000: City Wage Deciles

Estimate	Decile								
	(10)	(20)	(30)	(40)	(50)	(60)	(70)	(80)	(90)
OLS	3.812 (0.347)*	3.395 (2.94)*	2.976 (0.279)*	2.703 (0.267)*	2.487 (0.265)*	2.364 (0.262)*	2.293 (0.267)*	2.197 (0.273)*	2.136 (0.312)*
IV1	5.059 (0.728)*	4.283 (0.480)*	3.631 (0.408)*	3.245 (0.401)*	2.966 (0.408)*	2.806 (0.425)*	2.743 (0.452)*	2.649 (0.479)*	2.656 (0.562)*
IV2	2.294 (0.577)*	2.475 (0.491)*	2.365 (0.467)*	2.268 (0.445)*	2.156 (0.442)*	2.126 (0.435)*	2.136 (0.436)*	2.090 (0.429)*	2.100 (0.497)*
IV3	5.060 (0.488)*	3.944 (0.341)*	3.095 (0.313)*	2.564 (0.293)*	2.168 (0.296)*	1.859 (0.297)*	1.660 (0.300)*	1.404 (0.335)*	1.205 (0.416)*

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

Table 14: 1970-2000: Alternate Dependent Variables

Variables	Log 1BR Rent		Employment Rate		Unemployment Rate		Log Labour Force	
	OLS (1)	IV1 (2)	OLS (4)	IV1 (5)	OLS (7)	IV1 (8)	OLS (10)	IV1 (11)
$\Delta R_{c,t}$	3.736 (0.722)*	6.337 (0.887)*	2.703 (1.399)	2.703 (1.399)	2.703 (1.399)	2.703 (1.399)	2.703 (1.399)	2.703 (1.399)
Const.	0.411 (0.104)*	0.793 (0.140)*	0.260 (0.196)	0.260 (0.196)	0.014 (0.021)	0.029 (0.021)	0.014 (0.021)	0.014 (0.021)
Obs.	456	456	456	456	456	456	456	456
R^2	0.536	0.535	0.53	0.453	0.452	0.451	0.221	0.188

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

Figure 1:




Figure 2:




Figure 3:




Figure 4:

