

ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

*An International Society for the Advancement of Economic
Theory in its Relation to Statistics and Mathematics*

<http://www.econometricsociety.org/>

Econometrica, Vol. 80, No. 3 (May, 2012), 1063–1104

DOES INDUSTRIAL COMPOSITION MATTER FOR WAGES? A TEST OF SEARCH AND BARGAINING THEORY

PAUL BEAUDRY

University of British Columbia, Vancouver, B.C., Canada, V6T 1Z1 and NBER

DAVID A. GREEN

*University of British Columbia, Vancouver, B.C., Canada, V6T 1Z1 and
IFS, London*

BENJAMIN SAND

York University, Toronto, ON, Canada, M3J 1P3

The copyright to this Article is held by the Econometric Society. It may be downloaded, printed and reproduced only for educational or research purposes, including use in course packs. No downloading or copying may be done for any commercial purpose without the explicit permission of the Econometric Society. For such commercial purposes contact the Office of the Econometric Society (contact information may be found at the website <http://www.econometricsociety.org> or in the back cover of *Econometrica*). This statement must be included on all copies of this Article that are made available electronically or in any other format.

DOES INDUSTRIAL COMPOSITION MATTER FOR WAGES? A TEST OF SEARCH AND BARGAINING THEORY

BY PAUL BEAUDRY, DAVID A. GREEN, AND BENJAMIN SAND¹

Does switching the composition of jobs between low-paying and high-paying industries have important effects on wages in other sectors? In this paper, we build on search and bargaining theory to clarify a key general equilibrium channel through which changes in industrial composition could have substantial effects on wages in all sectors. In this class of models, wage determination takes the form of a social interaction problem and we illustrate how the implied sectoral linkages can be empirically explored using U.S. Census data. We find that sector-level wages interact as implied by the model and that the predicted general equilibrium effects are present and substantial. We interpret our results as highlighting the relevance of search and bargaining theory for understanding the determination of wages, and we argue that the results provide support for the view that industrial composition is important for understanding wage outcomes.

KEYWORDS: Wages, industrial composition, social interaction, search, bargaining.

1. INTRODUCTION

CHANGES IN THE NATURE OF JOBS are often given a preeminent role in popular discussions about labor market developments. In particular, it is often claimed that labor market performance hinges on whether an economy is attracting or losing “good jobs”; that is, jobs in industries that pay a premium relative to wages for similarly qualified workers in other industries. These claims are often translated into substantial efforts by governments to attract good-jobs industries to specific localities. Moretti (2010a), quoting Bartik (2002), stated that U.S. state and local governments spend \$30–40 billion dollars per year on policies aimed at attracting certain businesses.² While many serious economic researchers dismissed such views as ill-informed, the populist view is now reemerging in economic discussions in many countries because of perceptions about the effects of globalization and technological change on industrial composition and wages. For this reason, it appears to be an opportune time to revisit the issue of the impact of changes in industrial composition on the wage structure.

Whether altering the industrial composition of an economy has substantial effects on wages depends heavily on the nature of general equilibrium effects,

¹The authors thank A. Bowlus, M. Bombardini, D. Card, G. Dahl, M. Doms, J. Fernald, N. Fortin, G. Galipoli, J. Gelbach, R. Gordon, M. Greenstone, S. Kortum, I. King, T. Lemieux, K. Milligan, B. Meyer, F. Pelgrin, J. Pencavel, J.-M. Robin, and F. Wolak for helpful discussions.

²One example is the case of the state of Georgia offering Kia over \$400 million to build a new plant there (OECD (2010)). At a more general level, in Bluestone and Harrison’s (1982) highly cited book, *The Deindustrialization of America*, the authors argued that the switch away from highly paid manufacturing jobs was key to understanding the poor labor market performance of the U.S. economy during the 1970s and 1980s.

or “spillovers,” from directly affected industries to other sectors. To see this, consider an example in which a city replaces a high-wage industry that pays a 20% premium relative to the average wage in the city and that employs 10% of the workforce with industries that employ the same proportion of workers but do not pay any premium. If there are no spillovers, then the only affected workers are those who have to switch from the departing, high-wage industry to other industries. One can then calculate the impact of this composition shift on the average wage by multiplying the proportion of workers affected (0.1) times their proportionate wage decline (0.2), implying a 2% decline in the average city wage. These types of calculations are the same as calculating the “between” component in standard decomposition exercises (sometimes called shift-share analysis)³ and have been performed many times by researchers trying to understand the wage impact of forces such as trade and technology that could be argued to have their impact through changing the industrial mix. By the very nature of these exercises (multiplying together relatively small proportions), one is bound to conclude that composition changes have relatively small impacts even though the event (shifting out of a sector employing 10% of the workforce at a 20% premium) seems large.

Recent research in the literature on local labor markets suggests, however, that general equilibrium effects are likely to be important for understanding wage and employment outcomes. Reviews by Glaeser and Gottlieb (2009) and Moretti (2010a) argue that changes in the size of a sector induced by a productivity shock can have substantive effects on workers outside that sector through agglomeration effects, shifts in skills, and migration-induced house price changes.⁴ Our goal in this paper is to bring together the insights from the local labor market literature with those from search and bargaining models to investigate an additional source of potential spillovers; one that is distinct from demand and agglomeration effects. In particular, we begin the paper by showing how a change in industrial composition in a standard search and bargaining model of the labor market⁵—extended to include many industries—affects the bargaining position of workers by changing their outside options. In this setup, an improved outside option for workers places upward pressure on wages, even if employment in the sector is unaffected. This implication of a search and bargaining model is very basic and implies that wages in different sectors act as strategic complements. Although the insight is straightforward and has potentially important ramifications for our understanding of the functioning of the labor market, as far as we know it is an empirical implication of this class of models which has not previously been extensively pursued.

³See, for example, Bound and Johnson (1992). See Lemieux (2002) for a thorough discussion of decomposition methods.

⁴Also see Moretti (2010b) for a discussion of how changes in the demand for tradeable goods affects the production of nontradeable goods at the city level.

⁵For an introduction to search and bargaining models of the labor market, see Mortensen and Pissarides (1999).

The central problem in analyzing the effects of industrial composition changes is that they are not exogenous with respect to wage formation, implying that identification of the causal impact of changes in industrial composition on wages is inherently difficult. There are many potential approaches to such an issue, including the exploitation of natural experiments and structural modeling. Our approach is structural in nature, but does not involve imposing theoretical restrictions on the data during estimation. Since different theoretical models often imply different routes to identification, different interpretations of what is being identified, and different evaluations of policy, we view it as important first to establish the validity of our model through testing. We show that our extended search and bargaining model implies tight and intuitive testable implications in the form of overidentifying restrictions. We then show that these restrictions are not rejected in our data, which we view as evidence in favor of using search and bargaining models to understand aggregate labor market phenomena. We use the model to derive our empirical specification and to justify our identification strategy. This approach fits with a tradition (e.g., [Blundell, Duncan, and Meghir \(1998\)](#)) that brings together structural modeling and instrumental variables estimation, with the objective to avoid a common criticism of structural estimation that relates to the difficulty of establishing the identifying variation that is being exploited.

In our empirical work, we exploit geographical variation in wages and industrial composition across U.S. cities over the period 1970–2007. Our key identification problem arises from the fact that while wages in sector j act as an outside option for workers in sector i , the opposite is also true. Thus, the theory implies that wage determination takes the form of what [Manski \(1993\)](#) named a reflection problem (this property is also sometimes referred to as a social interaction problem ([Moffitt \(2001\)](#))). We devote considerable effort to addressing the endogeneity issues associated with solving the reflection problem. We also pay close attention to lessons from the local labor market literature on the importance of accounting for migration and housing price changes, and explore the robustness of our results to these considerations both theoretically and empirically.

We find two main empirical results. First, co-movements between city-level outcomes in terms of wages, employment, and industrial composition conform closely to the patterns implied by search and bargaining theory. Second, we find that our parameter estimates imply that city-level changes in industrial composition have effects on average wages that are 3–4 times that implied by the pure accounting (or shift-share) approach described earlier. Since measured composition effects are often small, the effects we find are large but not unreasonable.⁶ Together these results add to the local labor market liter-

⁶Thus, while the pure accounting approach indicates that a city losing an industry that employed 10% of the workforce and paid a wage premium of 20% (roughly the situation facing Pittsburgh with the loss of the steel industry in the 1980s) would face only a 2% decline in its average wage, our result suggests that the total impact would be a 6–7% decline.

ature by pointing to an important extra general equilibrium effect of changes in industrial composition that has been missed in more standard neoclassical models. Our results also indicate the extent to which standard decomposition techniques can be misleading in the presence of general equilibrium effects. In this particular case, these results may indicate a need to revisit conclusions from the literature that suggest that factors related to changes in industrial composition played a relatively minor role in the observed wage changes in the United States over the last few decades.

Finally, it is important to emphasize that the question we aim to answer differs from the focus in papers on regional demand shocks such as [Feyrer, Sacerdote, and Stern \(2007\)](#), [Greenstone and Moretti \(2003\)](#), or [Blanchard and Katz \(1992\)](#). While Greenstone and Moretti examine shorter run effects of acquiring a large plant in a region, such a change is composed of both an aggregate demand component and a composition change. In contrast, we focus on changes that arise over 10 year horizons and aim to isolate effects of changes in industrial composition, holding direct demand effects constant. This focus also differentiates our work from studies of regional adjustment to aggregate labor demand changes such as in [Blanchard and Katz \(1992\)](#) and [Feyrer, Sacerdote, and Stern \(2007\)](#). To clarify this difference, we take care to control for the types of demand effects examined in their papers.

The remaining sections of the paper are as follows. In Section 2, we present a model with search frictions and bargaining to illustrate a particular channel through which changes in industrial composition in an economy can affect wage setting within all sectors in the economy, including wages in sectors not directly involved in any composition shift. In Section 3, we use the model to derive a general empirical specification which embeds alternative views about the determination of wages. In Sections 4.1 and 4.2, we discuss the data used in the study and report basic empirical results. In Section 5, we address issues related to endogeneity, and in Section 6, we explore the robustness of our results as well as additional implications of our model. In Section 7, we use the model and the estimated parameters to discuss the magnitude of the effects of industrial composition on wages. Section 8 concludes.

2. GENERAL EQUILIBRIUM EFFECTS OF INDUSTRIAL COMPOSITION: A SEARCH AND BARGAINING PERSPECTIVE

In this section, we present a multisector extension of a standard search and bargaining model. Our aim is to clarify how a shift in the composition of labor demand can affect wages in sectors that are not directly part of the shift. One key objective of this section is to derive an empirical strategy for evaluating whether search and bargaining theory offers a pertinent framework for understanding the potential links between industrial composition and wages.

The model economy we consider has one final good, denoted Y , which is an aggregation of output from I industries and is given by

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

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

To greatly simplify the exposition, we begin by examining the extreme case where workers are not mobile across cities. In Section 2.3, we show that the results derived in this section are robust to an extension of the model that allows for worker mobility and endogenous housing prices.

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

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

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

$$(1) \quad \rho V_{ic}^f = (p_i - w_{ic} + \varepsilon_{ic}) + \delta(V_{ic}^v - V_{ic}^f),$$

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

$$(2) \quad \rho V_{ic}^v = -r_i + \phi_c(V_{ic}^f - V_{ic}^v),$$

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

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

$$(3) \quad \rho U_{ic}^e = w_{ic} + \tau_c + \delta(U_{ic}^u - U_{ic}^e),$$

where U_{ic}^u represents the value associated with being unemployed when the worker's previous job was in industry i and τ_c is a city-specific amenity.

When an individual is unemployed, he gets flow utility from an unemployment benefit, b , plus the city-specific amenity τ_c . Unemployed individuals find jobs at rate ψ_c . When an unemployed worker from industry i finds a job, with probability μ the job draw is from industry i , and with probability $(1 - \mu)$ it is a random draw from jobs in all industries (including i). The value associated with being unemployed satisfies

$$(4) \quad \rho U_{ic}^u = b + \tau_c + \psi_c \cdot \left(\mu U_{ic}^e + (1 - \mu) \sum_j \eta_{jc} U_{jc}^e - U_{ic}^u \right).$$

The important aspect of equation (4) is that as long as $\mu < 1$, the utility level associated with losing a job in industry i depends on the utility associated with jobs in other industries. The instantaneous probability of finding a job in industry j is given by $\psi_c \cdot (1 - \mu) \cdot \eta_{jc}$, where η_{jc} represents the fraction of vacant jobs that are in industry j . This formulation assumes that workers can only search while unemployed. While this is a strong assumption, it allows the problem to be solved explicitly and thereby to be amenable to clear empirical exploration. For these reasons, we maintain this assumption throughout. However, it must be noted that omitting on-the-job search is an important limitation of our analysis and, therefore, we hope to conduct future research aimed at seeing how the relaxation of this assumption affects our results.⁷ Similarly, throughout our analysis we maintain the assumption that job search is random as opposed to being directed. Exploring the implications of our model under directed search is another direction that warrants future research.

Once a match is made, workers and firms bargain a wage,⁸ which is set according to the bargaining rule

$$(5) \quad (V_{ic}^f - V_{ic}^v) = (U_{ic}^e - U_{ic}^u) \times \kappa,$$

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

$$M((L_c - E_c), (N_c - E_c)),$$

where L_c is the total number of workers in city c , E_c is the number of employed workers (or matches), and $N_c = \sum_i N_{ic}$ is the number of jobs, with N_{ic} being the number of jobs in industry i in city c . Then, given the exogenous death rate

⁷We make a first step in Appendix S.4 of the Supplemental Material (Beaudry, Green, and Sand (2012)) where we explore some empirical implications of allowing on-the-job search.

⁸We assume throughout that there are always gains from trade between workers and firms for all jobs created in equilibrium.

of matches, δ , and assuming that the matching function is of Cobb–Douglas form, the steady state condition is given by

$$(6) \quad \delta ER_c = (1 - ER_c)^\sigma \left(\frac{N_c}{L_c} - ER_c \right)^{1-\sigma},$$

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

The number of jobs created, N_{ic} , is determined by the free-entry condition

$$(7) \quad c_{ic} = V_{ic}^v,$$

where c_{ic} is the cost of creating a marginal job. If this cost were a constant (independent of c and i), this would imply that an increase in flow profits in an industry would be met by the creation of an infinite number of jobs. Moreover, such a specification would rule out the possibility of across-city comparative advantage in job creation. Although such a situation is a possibility, we want to consider a slightly more general setting. Hence, we assume that c_{ic} is potentially increasing in the number of jobs created in a city—perhaps due to limited local entrepreneurial talent. We also want to allow for the possibility that cities differ in their cost of creating certain types of jobs. Therefore, we let c_{ic} be a decreasing function of the industry–city-specific measure of advantage denoted Ω_{ic} . For tractability, we assume that the relationship is given by

$$c_{ic} = \frac{(N_{ic})^q}{Y_i + \Omega_{ic}},$$

where q controls the extent to which there are decreasing returns to job creation at the industry–city level. If q is close to zero, we are back to the case where any change in flow profits leads to massive change in the supply of jobs.⁹ The term Y_i reflects any systematic differences in cost of entry across industries, allowing us to assume that $\sum_c \Omega_{ic} = 0$.

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

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

At the city level, the price of industrial output is taken as given and an equilibrium consists of values of N_{ic} , w_{ic} , and ER_c that satisfy equations (5), (6), and (7). Note that these equilibrium values depend on (among other things)

⁹In the case where q is exactly equal to zero, cities should specialize entirely in one sector. Since we want to consider an environment without complete specialization, we assume $q > 0$.

the city-specific productivity parameters Ω_{ic} and ε_{ic} . An equilibrium for the entire economy has the additional requirement that the prices for industrial goods must ensure that markets for these goods clear.

At the city level, equilibrium outcomes reflect how cities adjust their production mix to take advantage of the price of industrial goods in relation to their comparative technological advantages. At the economywide level, changes in prices are caused by changes in demand for the industry-level goods, captured in the a_i s. As we make clear, our focus is on isolating the effect of a change in job composition on city-level wages. To this end, we take the above description of a steady state equilibrium as representing an equilibrium at a point in time and we examine how this equilibrium changes in response to changes in the exogenous driving forces a_i , Ω_{ic} , and ε_{ic} .

2.1. Derivation of the Wage Equation

Our initial focus is on the determination of wages, as implied by (5). For now we treat η_{ic} and ER_c as given.

To understand the forces that determine wages in an industry–city cell, we begin by using equations (3) and (4) to express the value of finding a job relative to being unemployed as

$$(8) \quad U_{ic}^e - U_{ic}^u = \frac{w_{ic} - b}{\rho + \delta + \psi_c \mu} - \frac{\psi_c(1 - \mu) \sum_j \eta_{jc}(w_{jc} - b)}{(\rho + \delta + \psi_c \mu)(\rho + \delta + \psi_c)}$$

A key feature of equation (8) is that as long as $\mu < 1$, a worker’s utility from being employed relative to being unemployed is a decreasing function of the wages paid in other sectors of the city’s economy. This arises because unemployed workers get a draw from the distribution of all wages in the city when an offer arrives. Also note that a compositional change captured by a change in the η s, holding ψ_c constant, in general affects the utility of finding a job. For example, if the composition change implies a greater concentration of jobs in high-wage sectors, this decreases the value of a match in any sector i .

To express the value of a match to a firm, we can use equations (1) and (2) in a similar fashion:

$$(9) \quad (V_{ic}^f - V_{ic}^v) = \frac{p_i - w_{ic} + \varepsilon_{ic}}{\rho + \delta + \phi_c}$$

From equations (8) and (9), we can use equation (5) to solve for w_{ic} . This is given by

$$(10) \quad w_{ic} = \gamma_{c0} + \gamma_{c1} p_i + \gamma_{c2} \sum_j \eta_{jc} w_{jc} + \gamma_{c1} \varepsilon_{ic}$$

where the coefficients in (10) are functions of the model parameters and, implicitly, of the employment rate.¹⁰ It is easy to verify that $0 < \gamma_{c2} < 1$ as long as $0 < \mu < 1$. Equation (10) captures the notion that in a search and matching framework, sectoral wages act as strategic complements; that is, high wages in one sector are associated with high wages in other sectors. The strength of this strategic complementarity is captured by γ_{c2} . In the case where $\mu = 1$ (i.e., workers are immobile across sectors), this effect disappears and wages are determined solely by the value of marginal product (i.e., by p_i and ε_{ic}). In such a case, there are not any effects of changes in industrial composition on within-sector wages. In this sense, the model nests more standard formulations where there are no general equilibrium effects. In contrast, with $\mu < 1$, a change in industrial composition can have an effect on within-sector wages, even when the employment rate is unchanged.

If there is an industrial composition shift that causes a 1 unit increase in the average city wage, $\sum_j \eta_{jc} w_{jc}$, equation (10) indicates that within-industry wages then increase by γ_{c2} . But this is just a first-round effect. Since the initial compositional change affects all within-industry wages, it causes the average wage to increase by another γ_{c2} units, inducing a further round of adjustments. Multiplying these feedback effects out, the total effect of the change in industrial composition on the average wage is therefore $\frac{1}{1-\gamma_{c2}}$.

Equation (10) has the structure of the classic reflection or social interaction problem (Manski (1993), Moffitt (2001)) in that the sectoral wage depends on the average of such wages in a city. To estimate such a relationship, it is necessary to overcome the simultaneity inherent to this type of interaction. In addition, we need to take explicit account of the fact that the γ_c coefficients, being functions of ψ_c and ϕ_c , are functions of the city's employment rate, ER_c . To address these issues, we take a weighted sum of (10) across industries, so as to replace $\sum_j \eta_{jc} w_{jc}$ with $\sum_j \eta_{jc} \nu_j$, where ν_j is the national-level wage premium in industry j relative to an arbitrarily chosen industry. These steps are laid out in the Appendix. We then take a linear approximation and eliminate any city-level fixed effects by focusing on the difference in wages within a city-industry cell across two steady state equilibria, denoted Δw_{ci} . The result is our key estimating equation

$$(11) \quad \Delta w_{ic} = \Delta d_i + \left(\frac{\gamma_2}{1 - \gamma_2} \right) \Delta \sum_j \eta_{jc} \nu_j + \gamma_3 \Delta ER_c + \Delta \xi_{ic},$$

where the Δ operator corresponds to $\Delta x_i = (x_{it+1} - x_{it})$, $\Delta d_i = \gamma_1 \frac{\gamma_2}{1-\gamma_2} \Delta p_1 + \gamma_1 \Delta p_i$ is an industry-specific effect that does not vary across cities and, hence, can be captured in an empirical specification by including industry dummies,

¹⁰Specifically, the coefficients are given by $\gamma_{c0} = \frac{(\rho+\delta+\psi_c\mu)(\rho+\delta+\phi_c)\kappa}{[(\rho+\delta+\psi_c\mu)+\kappa(\rho+\delta+\phi_c)](\rho+\delta+\psi_c)} b$, $\gamma_{c1} = \frac{\rho+\delta+\psi_c\mu}{(\rho+\delta+\phi_c)\kappa+(\rho+\delta+\psi_c\mu)}$, and $\gamma_{c2} = \frac{(\rho+\delta+\phi_c)\kappa}{[(\rho+\delta+\psi_c\mu)+\kappa(\rho+\delta+\phi_c)]} \frac{\psi_c(1-\mu)}{(\rho+\delta+\psi_c)}$.

and $\Delta\xi_{ic} = \gamma_1\Delta\varepsilon_{ic} + \gamma_1\frac{\gamma_2}{1-\gamma_2}\sum_j\frac{1}{I}\Delta\varepsilon_{jc}$ is the error term, with I being the total number of industries.

In (11), the γ coefficients are the same as discussed after equation (10), except now they are evaluated at common (across-cities) match probabilities, ψ and ϕ . The added coefficient, γ_3 , reflects the effect of a change in the employment rate on wage determination, an effect which depends on many parameters of the model. In Section 2.2, we discuss how we address the potential endogeneity concerns with estimating (11).

The variable $R_c \equiv \sum_j \eta_{jc} \nu_j$ is an index of the industrial composition of a city, where the weights on the industrial proportions are the national-level wage premia. We refer to R_c as average city rent. Note that the coefficient on R_c in (11) is $\frac{\gamma_2}{1-\gamma_2}$, as the steps between equations (10) and (11) resolve the reflection problem. A high value of this index indicates that a city's employment is concentrated in high-paying industries. A positive coefficient on ΔR_c in (11) implies that if we examine wages in the same industry in different cities, those wages will be higher in cities where employment is more heavily weighted toward what we call high-rent industries. This arises in the model because the workers in that industry have a better outside option to use when bargaining with firms in a higher average rent city.¹¹ At first glance, it might appear that almost any model predicts a higher wage in a city with a high-rent industrial composition, but it is important to emphasize that we are considering wage movements within a given industry, not the overall average wage in a city. We are comparing, for example, wages in construction in different cities to see if they are higher in cities with high-paying manufacturing jobs. It is worth noting once again that this average rent effect is estimated conditioning on the employment rate. Thus, it is concerned with the composition of employment rather than the level of labor demand.

In the most basic economic models, the wage in an industry-city cell should depend only on the productivity of workers in that cell and that productivity should not depend on the industrial composition of the rest of the workforce. In that case, the coefficient on ΔR_c should equal zero. One can think of extensions of that model where composition can affect productivity (even holding ER_c and thus the overall level of demand constant), with the most prominent possible reasons being due to demand for output from parts and service suppliers when particular downstream industries arrive or leave. What separates our model from those other potential spillover mechanisms is that it implies a specific composition index with wage premia as weights. This, in turn, implies a test that separates our model from models that imply generic spillover effects.

¹¹Note that in the model, these industrial premia are rents based on differential productivity in different industries. If they were, for example, compensating differentials of some kind, then they would not be relevant as outside options.

To understand the test, consider using a simple decomposition of ΔR_{ct} to rewrite (11) as

$$(12) \quad \Delta w_{ict} = \Delta d_{it} + \left(\frac{\gamma_2}{1 - \gamma_2} \right) \sum_j \Delta \eta_{jct} \nu_{jt} \\ + \left(\frac{\gamma_2}{1 - \gamma_2} \right) \sum_j \Delta \nu_{jt} \eta_{jct+1} + \gamma_3 \Delta ER_{ct} + \Delta \xi_{ict},$$

where we reintroduce t subscripts to clarify the decomposition. The second term on the right side of (12) corresponds to changes in the average rent in a city due to shifts in industrial composition toward or away from high-wage industries. We denote this term by ΔR_{ct}^B , since it is the “between” industry component of the change in city rent, ΔR_{ct} . The third right hand side term corresponds to changes in the average rent because of changes in the national-level wage premia, weighted by the concentration of a city in particular industries, which we call ΔR_{ct}^W . These two components reflect very different forms of data variation: shifts in the local composition of employment versus changes in national-level wage premia. Crucially, though, the model implies that both types of changes in average city rent should have the same effect on within-industry wage movements, since it does not matter in bargaining whether a worker’s outside option is weakened because a high-wage industry, j , shuts down or because industry j ceases to pay high relative wages while maintaining its employment share. Thus, estimation of (12) implies a test of the search and bargaining model: under this model, the coefficients on the second and third terms in (12) should be equal. Given the very different types of variation underlying each term, we view this as a stringent test. Moreover, this test stems ultimately from the fact that the R_c index implied by the model weights industrial proportions by wage premia. Under other models of spillovers (with other implied composition indexes), there is no reason for the “between” and wage premia components to have the same coefficients.

2.2. Endogeneity of Industrial Composition

Although use of equilibrium conditions to transform equation (10) into equation (11) provides a means to resolve the reflection problem inherent to the local interactions induced by bargaining, this does not resolve all endogeneity issues. Let us begin by focusing on the potential endogeneity of ΔR_c , acting for now as if there were no endogeneity issue related to ΔER_c . The issue that arises in this case is that changes in R_c may be correlated with the error term in (11) because R_c is a function of the proportions of workers in different industries (the η_{ics}) and one would expect those to be correlated with the industry–city-specific productivity shocks (the ε_s) that appear in the error.

To see this more concretely, note that using equation (6) and taking a linear approximation, we can get an expression for η_{ic} as¹²

$$(13) \quad \eta_{ic} \approx \frac{1}{I} + \pi_1 \left(\varepsilon_{ic} - \frac{1}{I} \sum_j \varepsilon_{jc} \right) + \pi_2 \left(p_i \Omega_{ic} - \frac{1}{I} \sum_j p_j \Omega_{jc} \right),$$

where the π s are positive coefficients obtained from the linear approximation.

The requirement for ordinary least squares (OLS) to give consistent estimates of the coefficients in (11) can be expressed as

$$(14) \quad \text{plim}_{C,I \rightarrow \infty} \frac{1}{I} \frac{1}{C} \sum_{i=1}^I \sum_{c=1}^C \Delta R_c \Delta \xi_{ic} = \text{plim}_{C,I \rightarrow \infty} \frac{1}{I} \frac{1}{C} \sum_{c=1}^C \Delta R_c \sum_{i=1}^I \Delta \xi_{ic} = 0,$$

where $\Delta R_c = \Delta \sum_j \eta_{jc}(w_j - w_1)$ and $\Delta \xi_{ic} = \gamma_1 \Delta \varepsilon_{ic} + \gamma_1 \frac{\gamma_2}{1-\gamma_2} \sum_j \frac{1}{I} \Delta \varepsilon_{jc}$, the error term in (11). At first glance, it might appear impossible that the change in a city's average rent could be uncorrelated with a function of changes in industry times city-level productivities, but inspection of the second equality in (14) indicates this is not the case. To discuss this, it is helpful to express ε_{ic} (and Ω_{ic} , as well) as the sum of a component that is common across all industries in a city (which we call the absolute advantage component), and a second component that captures the relative advantage of the specific industry times city cell. In particular, let $\bar{\varepsilon}_{ct}$ represent the common component of the ε s and let v_{ict}^ε represent the relative advantage component, with $\varepsilon_{ict} \equiv \bar{\varepsilon}_{ct} + v_{ict}^\varepsilon$, where, by definition, the v_{ict}^ε s across industries within a city sum to zero.

We can prove consistency just by considering variation across cities within an industry (i.e., taking C to infinity in a mathematical derivation). In practical terms, the inclusion of industry dummies in specifications (11) and (12) means we get identification using within-industry, cross-city variation. With $\Delta R_c = \Delta \sum_j \eta_{jc} \nu_j$, its variation across cities comes only from variation in the η_{jc} s, since the ν_j s are defined at the national level and so do not vary across cities. On the other hand, $\sum_{i=1}^I \Delta \xi_{ic}$ varies across cities only because of differences in the $\bar{\varepsilon}_{ct}$ s (since the relative advantage components of the cost advantage shocks sum to zero within a city). Thus, equation (14) is satisfied if changes in the relative industrial composition of city employment are uncorrelated with changes in the absolute advantage level of the city.¹³ In intuitive terms, for OLS to be consistent, there must be no link between whether a city experiences a shift in industrial composition toward high-paying industries (a shift in ΔR_c) and other general improvements that affect wages in all industries in the city.

¹²See the Supplemental Material for the expression before the approximation.

¹³More technically, in the Appendix, we show that consistency condition (14) is met under the assumption that the absolute advantage component, $\bar{\varepsilon}_{ct}$, follows a random walk and is independent of the relative advantage components of the ε s and of the Ω s (past and present).

While we believe that the assumption ensuring consistency of the OLS estimates is possible, we are more comfortable with the much weaker assumption that the common component of the city–industry cost advantage, $\bar{\varepsilon}_{ct}$, has increments that are independent of past values of the relative advantage components v_{ict}^e and v_{ict}^Ω (details in Section A.2). Based on this assumption, we can construct two instruments, each targeting one of the components of $\Delta R_c = \Delta R_{ct}^B + \Delta R_{ct}^W$, as represented by the second and third right hand side terms in (12). To construct the first instrument, we first predict a level of employment for industry i in city c in period $t + 1$ as

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

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

$$IV1_{ct} \equiv \sum_i v_{it} (\hat{\eta}_{ict+1} - \eta_{ict}).$$

This instrument isolates the variation in ΔR_{ct} that stems from changes in the employment composition (ΔR_{ct}^B), 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 (arising from, for example, trade or preference shocks) if that shift is distributed across cities according to start-of-period employment shares?” Note that use of this type of variation is implied by the model, where shifts in national-level demand ($\{a_i\}$) result in shifts in local employment shares because of differences in comparative advantage that are reflected in initial period employment shares.

Our second instrument isolates the variation in ΔR_{ct} 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

$$IV2_{ct} \equiv \sum_i \hat{\eta}_{ict+1} \Delta v_{it}.$$

These instruments are functions of v_{ict}^e , but not Δv_{ict}^e . In Section A.2, we show that they provide consistent estimates under the condition that $\Delta \bar{\varepsilon}_{ct}$ is independent of v_{ict}^e and Ω_{ict} . In other words, changes in the absolute advantage of a city are independent of the past level of relative advantage for the various

industries in a city. This is a weaker assumption than that required for consistency under OLS because it allows for the possibility that a city that shifts in the direction of a higher-wage industrial composition is also a city experiencing improvements in general. For our instruments to be valid, we require that those general improvements be independent of past comparative advantage. Note that we do not require $\Delta \bar{\varepsilon}_{ct}$ to be independent of $\bar{\varepsilon}_{ct-1}$ and, thus, rich types of persistence in the common component are allowable. Nonetheless, this type of orthogonality assumption is certainly questionable, and for this reason we provide a means to test its implication in Section 5.

The last endogeneity issue relates to the potential correlation between the change in the employment rate, ΔER_c , and the error term in (11). To see this possibility, it is helpful to use a linear approximation of (6) and (7) to obtain an expression for ER_c in terms of the ε s and the Ω s. This is given by

$$(15) \quad ER_c \approx \tilde{\pi}_0 + \tilde{\pi}_1 \sum_j \varepsilon_{jc} + \tilde{\pi}_2 \sum_j p_j \Omega_{ic},$$

where the coefficients $\tilde{\pi}$ are again positive and obtained from the linear approximation.

From (15) one can see that ΔER_c is likely correlated with the error term in (11) and, as a result, that OLS generates inconsistent estimates. We address this with an instrumental variable strategy in which we again use national-level information on growth patterns to predict city-level changes in employment rates. The approach is similar to that used by [Blanchard and Katz \(1992\)](#) in a closely related problem. Details of the construction of this instrument are left to the empirical section.

2.3. Worker Mobility

In the model as presented so far, we have assumed that workers are not mobile across cities. At first blush, it may appear that the result that wages can vary systematically across cities due to the composition of employment will disappear once we allow for mobility of workers. However, this will not generally be the case, even if we allow for mobility across cities.¹⁴ To see this, consider allowing unemployed individuals to occasionally have the option of changing cities. In this case, an individual who lost his job in industry i would choose the city c' that maximizes his expected utility (that is, this individual would choose his location by solving $\max_{c'} \sum_i \eta_{ic'} U_{ic'}^u = U_{i,\max}^u$). If we let μ_1 represent the probability that an unemployed individual has the option to change cities, then the value function associated with being unemployed from industry i in

¹⁴It is easy to verify that allowing for random search across cities does not significantly change our previous analysis.

city c needs to be modified to satisfy the relationship

$$\begin{aligned} \rho U_{ic}^u &= b + \tau_c + \psi_c \cdot \left(\mu U_{ic}^e + (1 - \mu) \sum_j \eta_{jc} U_{jc}^e - U_{ic}^u \right) \\ &+ \mu_1 \cdot (U_{i,\max}^u - U_{ic}^u). \end{aligned}$$

This alternative Bellman equation does not greatly change our previous analysis because $U_{i,\max}^u$ does not depend on the initial location of the worker and so can be treated as a common element across cities and captured by a time dummy. More to the point, if mobility is sufficiently high (i.e., μ_1 is sufficiently great) such that utility is equalized across cities, then the term $U_{i,\max}^u - U_{ic}^u$ is equal to zero and we return to our previous equation without mobility.

When the model is extended to allow for mobility across cities, it is appropriate to introduce housing costs (or land prices) into the equilibrating mechanism. In such a setup, when mobility leads to equal utility across cities, wage negotiation is not directly affected by housing prices, since housing costs are incurred whether or not someone is employed. Nonetheless, housing prices do adjust to changes in industrial structure. A city with a higher employment rate or a better employment mix, as captured by a higher value of R_{ct} , attracts more workers. This immigration drives up local housing prices, delivering benefits to local landowners and causing the migration to stop before wages are equalized across cities. Thus, the impact of a favorable change in job composition is still seen in wages, as implied by equations (10) and (11), and so, when mobility is sufficient to equate expected utility across cities, the presence of mobility does not change the predictions of the model in terms of the implied general equilibrium effects of industrial composition. We discuss the implications of mobility in more depth in Section 6.2, examining the case where mobility is not sufficient to assure the equalization of expected utility across cities.

2.4. Worker Heterogeneity

Up to now, labor in our model has been treated as homogeneous and not differentiated by skill. Since workers do differ in skills (and are treated differently by employers as a result), we need to take this into account. Empirically, we can approach the issue of insuring we are only estimating within skill-group implications of the model in two ways. The first is to work with subsamples of the data intended to be homogeneous with respect to skill (e.g., using only young high school graduates). The other is to eliminate skill differentials by first controlling for a flexible set of “skill” indicators in the determination of wages (where a wage regression is now seen as an individual-level wage regression rather than one specified at the industry times city level as it has been to this point). This latter approach is justified if the setup costs and benefits associated with employing a more skilled worker are proportional to the increased

productivity of the worker and if all the parameters of the model are identical across skill groups. To implement this approach, we also need to control for the same set of skill indicators when calculating the national-level wage premia used in constructing the R_{ct} measure. We pursue both approaches in the empirical section.

3. EMPIRICAL IMPLEMENTATION

Our baseline empirical specification is given by equation (16).¹⁵ This equation is a simple rewrite of equation (11), where we have divided both sides by w_1 so as to focus on a log specification and have added a time subscript since we pull data from different periods:

$$(16) \quad \Delta \log w_{ict} = d_{it} + \alpha_2 \Delta R_{ct} + \alpha_3 \Delta ER_{ct} + \Delta \xi_{ict}.$$

In (16) the d_{it} s are time-varying industry dummies, $\alpha_2 = \frac{\gamma_2}{1-\gamma_2}$ is our main coefficient of interest, $R_{ct} = \sum_j \eta_{jc} (\frac{w_j}{w_1} - 1)$ is our index of industrial composition, α_3 captures the effect of city-level employment rates on wages, and $\Delta \xi_{ict}$ is the error term defined by $\Delta \xi_{ict} = \frac{\gamma_1}{w_1} \Delta \varepsilon_{ict} + \frac{\gamma_1 - \gamma_2}{w_1 (1-\gamma_2)} \sum_j \frac{1}{l} \Delta \varepsilon_{jct}$. We also implement a specification in which we allow the two components of ΔR_{ct} to enter separately, as in (12).

Our goal is to investigate the null hypothesis that $\alpha_2 = 0$ or, in other words, whether disregarding intersectoral wage interactions provides an appropriate description of wage determination in local economies. Our alternative hypothesis is that $\alpha_2 > 0$. A finding of $\alpha_2 > 0$ supports the notion of general equilibrium effects of industrial composition on wages as predicted by the model. The existence of such general equilibrium effects calls into question commonly used accounting approaches for evaluating the effects of changes in industrial composition on average wages.

When estimating the effect of R_{ct} on wages, it is appropriate to worry about omitted variable bias, especially given alternative explanations for differences in wages across cities such as those related to city size, education levels (Moretti (2004a), Acemoglu and Angrist (1999)), and diversity of employment in a city (Glaeser, Kallal, Scheinkman, and Shleifer (1992)). To allow for such issues, we add to equation (16) measures related to these explanations as additional covariates, Z_{ct} .¹⁶

¹⁵In Appendix S.10 of the Supplemental Material, we also provide estimates that relate to equation (10). When estimating equation (10), as the reflection problem is not solved, OLS is never consistent. As can be seen in the Supplemental Material, the IV estimates of γ_2 found when estimating equation (10) are very close to those implied from our estimates of $\frac{\gamma_2}{1-\gamma_2}$ when the estimation is based on equation (11).

¹⁶Another potential point of concern is the linear approximation used in deriving (16). Problems could arise if higher order terms assigned to the error are correlated with our instruments.

4. DATA AND BASIC RESULTS

4.1. *Data*

The data we use in the following investigations come from the 1970, 1980, 1990, and 2000 U.S. Census Public Use Micro-Samples (PUMS) and the 2007 American Community Survey (ACS). We focus on wage and salary earners, aged 20–65 with positive weekly wages who were living in a metropolitan area at the time of the Census. To form our dependent variable, we use the log of weekly wages, calculated by dividing wage and salary income by annual weeks worked (we also verified the robustness of our results to using hourly wages). We create real wages (in 1990 dollars) using the national-level consumer price index (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 S.1 in the Supplemental Material (Beaudry, Green, and Sand (2012)).

As we described in the previous section, one approach to addressing worker heterogeneity is to control for observable skills in a regression context. Our actual approach is to use a common two-stage procedure. In the first stage, we run individual-level regressions of log wages using all the individuals in our national sample on categorical education variables (four categories), a quadratic in experience, interactions of the experience, and education variables, a gender dummy, black, hispanic, and immigrant dummy variables, and the complete set of interactions of the gender, race, and immigrant dummies with all the education and experience variables. We run these regressions separately by Census year to allow for changes in returns to skills over time. The regressions also include a full set of industry-by-city cell dummies and it is the coefficients on those that are used to construct the dependent variable in the second-stage regression (equation (16) above). We eliminate all industry–city cells with fewer than 20 included individuals in any of the years. We use the square root of the number of observations in each industry–city cell to form weights for the second-stage estimation. For most of our estimates, we use decadal differences within industry–city cells for each pair of decades in our data (1980 – 1970, 1990 – 1980, 2000 – 1990) plus the 2007 – 2000 difference, pooling these together into one large data set and including period-specific industry dummies. In all the estimation results we calculate standard errors, allowing for clustering by city and year. Note that our procedure of pooling data both across industries and across time relies on assumptions of homogeneity of parameters across industries and stability of these parameters over time. In the Supplemental Material, we explore empirically how our results change when we relax

We explored this possibility through a Monte Carlo simulation in which we constructed wage data using our (nonlinear) theoretical model and data on ER_c , η_{ic} , and v_i . We then estimated (16) on the simulated data. The results indicated very little bias from our approximations. Details are given in Appendix S.5 in the Supplemental Material.

these assumptions. We find that the pooling procedure used here provides a good summary description of the data.

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

4.2. OLS Results

We begin our presentation of results with the estimates of our base specifications without the inclusion of any additional control (Z_{ct}) variables. One key concern we choose to address immediately is the issue of selection. In particular, changes in unobserved skills in a city may be correlated with movements of R_{ct} , which imply a nonzero coefficient on R_{ct} that does not reflect general equilibrium effects of the type we are considering. For example, suppose 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 where they have a chance to get a high-paying job. In that case, shifts in R_{ct} may pick up the effects of shifts in the unobserved ability distribution. To address this potential problem, we implement a Heckman two-step type procedure as proposed in Dahl (2002). A detailed discussion of our selection correction procedure is presented in Appendix S.2. We find that the selection correction terms from this procedure enter our estimating equations significantly (indicating that some selection of this type exists), but we also find that their inclusion does not have a substantive effect on our parameters of interest. Nonetheless, since selection seems to be present, all of the results we present include the selection correction unless stated otherwise.¹⁷

The first columns of Table I contain the results from OLS estimation of (16). These regressions and all of those that follow include a full set of time-varying industry dummy variables (4×144), thus allowing for changes in industry effects over time, but we do not present the long list of corresponding coefficients here. In column 1, the coefficient on ΔR_{ct} is 2.47 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 a change in the composition of employment in a city

¹⁷A sample of the results without the selection correction is available in Table S.VIII.

TABLE I
BASIC RESULTS^a

	OLS		IV				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ΔR_{ct}	2.47*			2.82*	2.95*	2.91*	2.86*
	(0.18)			(0.35)	(0.33)	(0.31)	(0.29)
ΔR_{ct}^W		2.07*	2.78*				
		(0.20)	(0.41)				
ΔR_{ct}^B		3.04*	3.05*				
		(0.42)	(0.45)				
ΔER_{ct}	0.42*	0.48*	0.74	0.70	0.58	0.63	0.63
	(0.078)	(0.076)	(0.47)	(0.44)	(0.48)	(0.44)	(0.43)
Year \times ind.	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	33,984	33,984	33,984	33,984	33,984	33,984	33,984
R^2	0.51	0.51					
Instrument set			IV1-IV2-IV3	IV1-IV3	IV2-IV3	IV1-IV2-IV3	IV1-IV2-IV3
<i>F</i> -statistics							
ΔR_{ct}^W			81.84				
ΔR_{ct}^B			583.37				
ΔR_{ct}				70.14	159.26	222.81	222.81
ΔER_{ct}			10.99	10.41	15.78	10.99	10.99
AP <i>p</i> -values							
ΔR_{ct}^W			0.00				
ΔR_{ct}^B			0.00				
ΔR_{ct}				0.00	0.00	0.00	0.00
ΔER_{ct}			0.00	0.00	0.00	0.00	0.00
Over-id. <i>p</i> -value			.	.	.	0.65	0.81

^aStandard errors, in parentheses, are clustered at the city-year level. The asterisk (*) denotes significance at the 5% level. All models are estimated on a sample of 152 U.S. cities using Census and ACS data for 1970–2007. The dependent variable is the decadal change in regression-adjusted city–industry wages.

is completely captured in the standard accounting measure. Further, the coefficient fits with the alternative hypothesis that cities with employment structures that shift toward high-premia industries have better wage performance within industries. To see that this result is not driven by outliers, in Figure 1 we plot the relationship between the average within-industry change in wages of a city against ΔR_{ct} (this nets out the time-varying industry effects). As can be seen, the positive relationship is very salient and is not driven by outliers.

In the third column of Table I, we present estimates of the specification in which we allow the component parts of ΔR_{ct} to have separate effects. The estimated coefficients are quite similar in size to one another, with the hypothesis that they are the same not being rejected at the 5% level of significance. Recall that the equality of these effects is implied by the search and bargaining framework, as both components of R_c should affect bargaining power in a symmetric fashion. Moreover, while one might view the $\sum_j \Delta \eta_{jc} \nu_j$ as spuriously picking

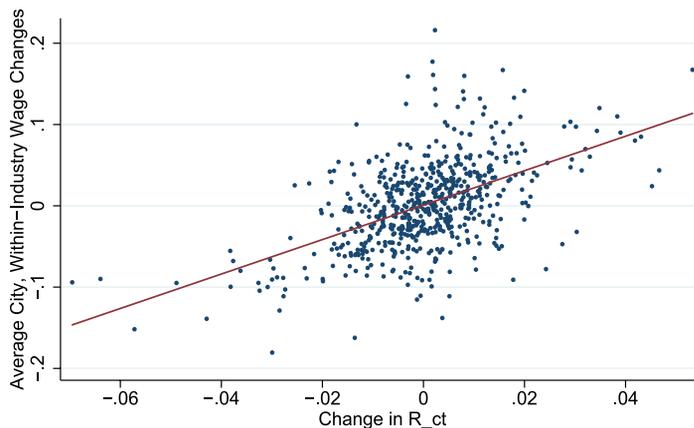


FIGURE 1.—The average, within-industry wage change plotted against changes in average city rent, once time-varying industry effects have been removed. The data contain information on 152 U.S. cities and 144 industries, and were obtained from the U.S. Census 1970–2000 and the 2007 ACS. Appendix S.1 contains details.

up some other effect of industrial composition change, it is hard to see how one could so argue for $\sum_j \eta_{jc} \Delta \nu_j$, since its variation comes from changes in national-level wage premia not changes in local composition. It is the fact that this latter component both enters and enters symmetrically with the first that, in our view, provides strong support for a bargaining interpretation.

5. ADDRESSING ENDOGENEITY ISSUES

As we discussed earlier, OLS estimation of (16) provides consistent estimates if changes in a city's absolute advantage are independent of levels and changes of relative advantage. While this may be the case, in this section we explore estimation of (16) using an instrumental variable strategy which relies on the weaker identification assumption that changes in the absolute advantage of a city are independent only of the past level of relative advantage for that city.

Recall from Section 2.2 that our instruments, IV1 and IV2, allow us to identify effects from each of the components of ΔR_{ct} in the specification given in (12). Alternatively, we can use either one as an instrument for ΔR_{ct} in (16). Using IV1 would then provide an estimate based on variation arising from changes in industry shares, while using IV2 provides an estimate based on changes in national-level industrial premia. As discussed earlier, in our bargaining setup, both types of variation should have the same effect on within-

industry wage growth. Thus, an equivalent means of testing this implication is to compare the coefficients obtained using IV1 versus IV2 in (16).

Both instruments perform well in the first stage. The t -statistic from the test of the significance of IV1 in the first-stage regression of ΔR_{ct} equals 11.6 and has an associated p -value of 0.0. The same statistic for IV2 is 16.24 with a p -value of 0.0.¹⁸

As we stated earlier, ΔER_{ct} is also likely to be endogenous and we respond to this problem using an instrument that is similar to IV1. In particular, we use $\sum_i \eta_{ict} g_{it}$, where g_{it} is the growth rate of employment in industry i at the national level in year t . Thus, the instrument is the weighted average of national-level industrial employment growth rates, where the weights are the start-of-period industrial employment shares in the local economy. We call this instrument IV3. A city that has a strong weight on an industry that turns out to grow well at the national level will have a high value for this instrument. Because the ε_{ict} s that drive the error term are local cost shocks that sum to zero across cities, their movements are not correlated with the g_{it} s by construction. Finally, under the assumption that changes in the common component of the city–industry cost shocks ($\Delta \bar{\varepsilon}_{ct}$) are independent of past relative advantage components (the v_{ict}^e s), the η_{ict} s used as weights in IV3 will be uncorrelated with error term in equation (16). To see this, recall from Section 2.2 that the η_{ict} s depend only on the relative city–industry advantages, while consistency requires that these be unrelated to $\sum_i \Delta \xi_{ict}$, which depends only on changes in the common component of ε_{ict} , $\Delta \bar{\varepsilon}_{ct}$. Therefore, this instrument is valid under the same conditions as IV1 and IV2.

We present results from the specification where we enter the components of ΔR_{ct} separately and instrument for both in the third column of Table I. Columns 4 and 5 contain results instrumenting for ΔR_{ct} using IV1 and IV2 individually. In all the instrumental variable estimates we present, the employment rate is instrumented for using IV3. All of these results, again, imply large and statistically significant effects of ΔR_{ct} or its components on local average wages.

The similarity in the estimated coefficients on the two subcomponents of ΔR_{ct} in column 3 is striking. Not surprisingly, a test of the null hypothesis that these two coefficients are equal cannot be rejected at any conventional significance level. The results in column 3 indicate that moves away from high-paying jobs and reductions in the premia associated with high-paying jobs have approximately the same impact on within-industry wages—a result implied by our bargaining story.

¹⁸At the bottom of the table, we also report results from standard F -statistics of the excluded instruments as well as Angrist–Pischke tests for weak instruments. The Angrist–Pischke (AP) tests are designed to take into account situations such as ours, with multiple endogenous variables. Again, the associated p -values are essentially zero, indicating a lack of a weak IV problem.

The same result is portrayed in an equivalent but less transparent way in the nearly identical results obtained when we instrument for ΔR_{ct} separately using IV1 and IV2 in columns 4 and 5. Examining the results in this way is useful because it allows for a discussion of their implications for the statistical assumptions underlying our consistency proofs. The consistency of estimates based on IV1 and IV2 each rely on the same assumption, since the instruments are essentially different weighted averages of a city's initial employment shares and our key identifying assumption relates to those shares. Given that the weighting schemes are different, one would expect that if the identification assumption does not hold, these two IV strategies would produce different estimates since they would weight any departures from the assumption differently. This intuition can be formalized by performing a standard Hansen- J overidentification test. We present the results from this test at the bottom of column 6. Not surprisingly given the similarity in the estimates in columns 4 and 5 in Table I, it passes easily. We believe that the fact that the two IV approaches, which focus on very different data variation,¹⁹ give very similar results provides considerable support for the search and bargaining model, and for the assumption that the common city-level component in the ε_{ct} acts like a random walk. Moreover, the fact that OLS also gives similar results suggests, further, that even the strong assumption needed for the consistency of OLS may hold in these data. In the last column of Table I, we reestimate the specification in column 6, but as a dependent variable use city-industry wages that are not corrected for selection across cities using the Dahl (2002) method described above. As can be seen from the table, we get very similar results with and without the selection correction.

6. ALTERNATIVE INTERPRETATIONS AND ROBUSTNESS

6.1. *Competing 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 that correspond to some of the more prominent hypotheses to see whether our R_{ct} measure may be spuriously capturing one of these alternative driving forces.²⁰

¹⁹The fact that IV1 and IV2 only have a correlation of 0.18 in our data (after removing aggregate time affects as we do in all our estimations) emphasizes that the two instruments highlight different variation.

²⁰In the Supplemental Material, we also examine alternative specifications where we allow for heterogeneity in various dimensions. In particular, we estimate our main regression separately by industry, resulting in 143 estimates of α_{2i} . The distribution of those estimates has a median of 2.50, with 25th and 75th percentiles of 1.92 and 2.97, respectively. Thus, our results are not driven by a subset of industries. We also allowed the job destruction rate and the returns to education

One prominent explanation for city-level growth is provided in Glaeser et al. (1992). Based on Jane Jacob's theorizing, they argued for the importance of industrial diversity and found that diversity is a stronger determinant of city-specific growth than city size. In response to this view, in the first three columns of Table II, 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 1 minus the Herfindahl index, or 1 minus the sum of squared industry shares. This measure itself tends not to be very significant in our estimates and, more importantly, does not change our estimates of the α_2 coefficient.

Another possibility relates to the recent literature on education externalities which examines the claim that having a larger proportion of workers in a city being highly educated benefits all workers in the city (e.g., Moretti (2004a), Acemoglu and Angrist (1999)). In considering this, it is worth reemphasizing 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 columns 4, 5, and 6 in Table II, 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's (2004a) findings, but introducing this variable has very little impact on our estimates of the effect of the components of the change in R_{ct} .²¹

In columns 7, 8, and 9, we introduce the change in the log of the size of the city's labor force. This city size variable is intended to capture the type of agglomeration effects tested in, for example, Glaeser et al. (1992). In the columns using instrumental variables (8 and 9), we instrument for the labor force variable using the same instrument that we have used in various ways to this point; that is, one based on predicting labor force growth in a city from national-level growth for each industry combined with the initial industrial composition in the city. Whether instrumented or not, this variable has small and statistically insignificant effects, and, more importantly, its inclusion does not alter the estimates of the ΔR_{ct} effect. The result that predicted that changes in labor force across cities have little effect on wages is consistent with Blanchard and Katz's (1992) finding that local wages react little to changes in demand or migration.

to vary by industry (though with a smaller number of industry groups) and found that our main results persist.

²¹It is worth noting, though, that Sand (2006) found that this positive and significant impact is observed in the 1980s but not in the 1970s or 1990s. In results not reported here, we used 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, its inclusion does not affect the estimates of our R_{ct} effects.

TABLE II
ALTERNATIVE EXPLANATIONS^a

	OLS		IV		OLS		IV		OLS		IV	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
ΔR_{ct}	2.48*	2.77*	2.97*	2.44*	2.65*	2.89*	2.45*	2.37*	2.70*	2.45*	2.65*	2.78*
	(0.18)	(0.33)	(0.33)	(0.18)	(0.33)	(0.31)	(0.20)	(0.43)	(0.34)	(0.19)	(0.64)	(0.36)
1 – Herfindahl	0.17	0.17	0.19									
	(0.18)	(0.19)	(0.20)									
ΔBA or >				0.39*	0.39*	0.38*						
				(0.14)	(0.14)	(0.14)						
$\Delta \log$ labor force							0.024	0.090	0.065			
							(0.015)	(0.061)	(0.057)			
$\Delta \log \bar{E}_{ict}$										0.031*	0.078	0.071
										(0.0045)	(0.13)	(0.095)
ΔER_{ct}	0.42*	0.74	0.55	0.43*	0.67	0.45						
	(0.077)	(0.43)	(0.46)	(0.078)	(0.41)	(0.46)						
Year \times ind.	Yes	Yes	Yes									
Observations	33,984	33,984	33,984	33,984	33,984	33,984	33,984	33,984	33,984	33,984	33,984	33,984
R^2	0.513			0.515			0.510			0.513		
Instrument set		IV1–IV3	IV2–IV3		IV1–IV3	IV2–IV3		IV1–IV3	IV2–IV3		IV1– $\Delta \hat{n}_{ict}$	IV2– $\Delta \hat{n}_{ict}$

^aStandard errors, in parentheses, are clustered at the city–year level. The asterisk (*) denotes significance at the 5% level. All models are estimated on a sample of 152 U.S. cities using Census and ACS data for 1970–2007. The dependent variable is the decadal change in regression-adjusted city–industry wages.

An alternative route for city-level changes in composition to affect within-industry wages is through shifts in labor supply. The departure of a large sector from a city implies that the former workers from that sector move to other sectors, shifting out their labor supply curves. This, in turn, implies moves down the labor demand curves and consequent wage declines in those sectors. To examine this possibility, in columns 10–12, we control for the change in employment in the industry–city cell. In this case, with employment held constant, any ΔR_{ct} effect should not be reflecting movements along a marginal product curve induced by supply shifts. As before, our instrument for the employment change uses predicted changes based on national-industry growth rates. The effect of the employment change is small and, again, does not alter our estimates of α_2 .

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

6.2. Local Demand Effects

In our model and interpretation of the data, we assumed that housing is the only nontraded good across cities. This assumption allows us to treat the price of goods as being common across cities and, therefore, to fully capture their effects through time-varying industry effects. When some produced goods are nontradeable across cities, this induces a city-specific component in prices to appear in the error term of our wage regressions. If changes in industrial composition cause changes in the prices of these nontradeables, perhaps through affecting the wealth of the city, this would lead our estimates of the effects of ΔR_{ct} on wages to capture a combination of the mechanisms emphasized in the model and more standard demand effects. As can be verified, in the presence of a nontradeable good, we can derive the analogue to equation (25), but now this equation is only applicable for the goods that are easily tradeable across cities. Accordingly, we can look for biases induced by local price changes by examining whether our results are very sensitive to being estimated on more versus less tradeable goods.

We define tradeable and nontradeable sectors using an approach from [Jensen and Kletzer \(2005\)](#). They argued that the share of employment in tradeable goods should vary widely across regional entities (cities in our case) since different cities concentrate in producing different goods which they can then trade. For nontradeable 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 therefore rank industries by the variance of their employment shares across cities and label the industries in the top, mid-

TABLE III
LOCAL DEMAND EFFECTS: TRADE AND NONTRADE INDUSTRIES^a

	Low Trade			Medium Trade			High Trade		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
ΔR_{ct}	3.13*	3.44*	4.10*	2.61*	2.93*	2.91*	2.30*	2.67*	2.83*
	(0.36)	(0.75)	(0.64)	(0.21)	(0.40)	(0.37)	(0.16)	(0.31)	(0.29)
ΔER_{ct}	0.49*	1.04	0.13	0.52*	0.54	0.56	0.33*	0.75*	0.62
	(0.13)	(1.01)	(1.31)	(0.093)	(0.53)	(0.55)	(0.071)	(0.38)	(0.41)
Year \times ind.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5230	5230	5230	14,078	14,078	14,078	14,676	14,676	14,676
R^2	0.45			0.48			0.56		
Instrument set		IV1–IV3	IV2–IV3		IV1–IV3	IV2–IV3		IV1–IV3	IV2–IV3

^aStandard errors, in parentheses, are clustered at the city–year level. The asterisk (*) denotes significance at the 5% level. All models are estimated on a sample of 152 U.S. cities using Census and ACS data for 1970–2007. The dependent variable is the decadal change in regression-adjusted city–industry wages.

dle, and bottom third as high-, medium-, and low-trade industries. In Table III, 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} is 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. This suggests that local demand effects on prices are not important enough to explain the relationship we observe between industrial composition and within-industry wages.

Another way demand effects may be confounding our inferences is through input–output linkages. For example, when one industry expands in a city, this likely creates demand for goods produced by the upstream industries present within the same city. To explore whether our ΔR_c measure reflects such linkages, we used input–output tables to create a measure of “input distance” between two industries i and j , denoted d_{ij} , as the fraction of inputs used by industry j which come from industry i .²² Note that d_{ij} is a measure calculated using national-level data and that $d_{ij} \neq d_{ji}$. We then create the index $M_{ict} \equiv \sum_j \eta_{jct} d_{ij}$. A high value for this index reflects an industrial structure that incorporates high demand for goods produced by industry i . We then explore the relevance of input–output demand linkages for our results by includ-

²²To construct d_{ij} , we use national-level input–output tables provided by the Bureau for Economic Analysis (<http://www.bea.gov/>). Following Moretti (2004b), we use the “Use” tables and construct our distance measure based on the value of inputs used in production in industry j that come from industry i . The data appendix in the Supplemental Material contains additional details.

TABLE IV
LOCAL DEMAND EFFECTS: INPUT-OUTPUT LINKAGES^a

	OLS		IV				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ΔR_{ct}	2.46*			2.80*	2.95*	2.84*	2.95*
	(0.18)			(0.35)	(0.32)	(0.35)	(0.33)
ΔR_{ct}^W		2.04*	2.75*				
		(0.20)	(0.41)				
ΔR_{ct}^B		3.06*	3.07*				
		(0.42)	(0.45)				
ΔER_{ct}	0.41*	0.47*	0.74	0.69	0.55	0.71	0.60
	(0.078)	(0.076)	(0.47)	(0.44)	(0.48)	(0.45)	(0.48)
ΔM_{ict}	1.08*	1.18*	0.90*	0.89*	0.94*	-0.86	-0.75
	(0.35)	(0.35)	(0.42)	(0.41)	(0.42)	(0.53)	(0.45)
Year \times ind.	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	33,984	33,984	33,984	33,984	33,984	33,984	33,984
R^2	0.51	0.51					
Instrument set			IV1-IV2-IV3	IV1-IV3	IV2-IV3	IV1-IV3-IV4	IV2-IV3-IV4
<i>F</i> -statistics							
ΔR_{ct}^W			81.36				
ΔR_{ct}^B			583.66				
ΔR_{ct}				70.49	159.60	61.41	106.13
ΔER_{ct}			10.93	10.45	15.55	20.83	20.02
ΔM_{ict}						6.95	6.77
AP <i>p</i> -values							
ΔR_{ct}^W			0.00				
ΔR_{ct}^B			0.00				
ΔR_{ct}				0.00	0.00	0.00	0.00
ΔER_{ct}			0.00	0.00	0.00	0.00	0.00
ΔM_{ict}						0.00	0.00

^aStandard errors, in parentheses, are clustered at the city-year level. The asterisk (*) denotes significance at the 5% level. All models are estimated on a sample of 152 U.S. cities using Census and ACS data for 1970–2007. The dependent variable is the decadal change in regression-adjusted city–industry wages.

ing ΔM_{ict} as an additional regressor in our baseline wage equation. Results for this exercise are reported in Table IV. In columns 1 and 2, the relationship is estimated by OLS. In columns 3–5 we treat ER_{ct} and R_{ct} (and its breakdown) as endogenous, and we use IV1, IV2, and IV3 as instruments. Finally, in columns 6 and 7, we also treat ΔM_{ict} as endogenous and add to our instrument set IV4 = $\sum_j \hat{\eta}_{jct+1} d_{ij} - M_{ict}$, where the $\hat{\eta}_{jct}$ s are predicted city–employment shares as described above. The results indicate that demand effects captured by input–output linkages have some effect on wages in our data. However, in all cases, the inclusion of the linkage variable has very little impact on our estimates of the wage effect of R_{ct} . Accordingly, we infer that input–output demand linkages are unlikely to be driving our results.

6.3. Directly Controlling for the Effects of Housing Costs

Up to this point, we have focused on estimating the general equilibrium (GE) effect of changes in industrial structure on local wages. This aggregate effect reflects two distinct mechanisms. The first, which we call the direct mechanism, relates to the effect of changes in industrial composition on wages, holding constant the local cost of living. This channel captures the effect on bargaining power of having the option to move to other industries within the city. But the wage change induced by this change in bargaining power will also affect housing costs (p_{ct}^h), which will themselves have further effects on wages. We call this second effect the indirect mechanism. The total GE effect is the sum of these two forces. As discussed in Section 2.3, our baseline estimation strategy—where the cost of housing is not included in the wage equation—is the appropriate specification for capturing the full GE effect of industrial composition changes under the assumption that mobility across cities is sufficient to equalize expected utility. In other words, our previous estimates of the effect of ΔR_{ct} on wages should be interpreted as reflecting both of these mechanisms.

In this section, our goal is to decompose the estimate of the GE effects of industrial composition into its two components. To see how to proceed, we return to the Bellman equations that define utility in the presence of mobility and variable housing prices.²³ For simplicity, we focus on the case where $\mu = 0$, which implies that the value function associated with being unemployed in city c at time t (U_{ct}^u) satisfies

$$(17) \quad \rho U_{ct}^u = b + \tau_{ct} - s \cdot p_{ct}^h + \psi_c \cdot \left(\sum_j \eta_{jct} U_{jct}^e - U_{ct}^u \right) + \mu_1 \cdot \left(\max_{c'} U_{c't}^u - U_{ct}^u \right).$$

In expression (17), we see that flow utility is negatively affected by the price of housing (p_{ct}^h) and positively affected by local amenities τ_{ct} (s represents the share of housing in expenditures). Recall that the parameter μ_1 represents the probability that an unemployed worker will have the option to move to another city. If this opportunity arises, the worker chooses the city with the greatest expected utility. In this setup, as there is no on-the-job search, the value function for employed workers in industry i and city c satisfies

$$(18) \quad \rho U_{ict}^e = w_{ict} + \tau_{ct} - s \cdot p_{ct}^h + \delta \cdot (U_{ct}^u - U_{ict}^e),$$

where, again, flow utility depends negatively on the cost of housing.

²³If mobility is sufficient to equalize expected utility across cities, one can show that the direct and indirect mechanisms are not separately identifiable, and this is why our baseline specification does not separate the two mechanisms. To identify the two, it is necessary to assume that expected utility is not perfectly equalized across cities, and that is the assumption we adopt in this subsection.

Using (17) and (18), the surplus received by a worker ($U_{ict}^e - U_{ct}^u$) is given by

$$\begin{aligned} & U_{ict}^e - U_{ct}^u \\ &= \frac{w_{ict} - (1 - \Pi_1)\Pi_3 b - \Pi_1\Pi_3 \left(\sum_j \eta_{jct} w_{jct} \right) - \Pi_2\Pi_3 \max_{c'} U_{c't}^u}{(\rho + \delta)} \\ & \quad - \frac{\Pi_2(s \cdot p_{ct}^h - \tau_{ct})}{(\rho + \Pi_2)(\rho + \delta)}, \end{aligned}$$

where $\Pi_1 = \frac{\psi_c}{\rho + \delta + \psi_c}$, $\Pi_2 = \mu_1(1 - \Pi_1)$, and $\Pi_3 = \frac{\rho}{\rho + \Pi_2}$.

From the above equation, we see that workers receive less surplus in matches, and will, therefore, be able to bargain higher wages at a given job, when housing costs are high. This is the channel through which the indirect mechanism plays out. We can now use the bargaining condition $V_{ict}^f - V_{ict}^v = \kappa(U_{ict}^e - U_{ct}^u)$ to solve for w_{ict} and make this argument explicit. Following the steps used in the [Appendix](#), then taking the log-linear approximation due to the dependence of coefficients on the employment rate, and, finally, taking a first difference, we get the equation for local wages w_{ict} of the form²⁴

$$(19) \quad \Delta \log w_{ict} = \tilde{\alpha}_{it} + \tilde{\alpha}_2 \Delta R_{ct} + \tilde{\alpha}_3 \Delta ER_{ct} + \tilde{\alpha}_4 \Delta \log p_{ct}^h + \Delta \tilde{\xi}_{ict},$$

where $\tilde{\alpha}_2$, $\tilde{\alpha}_3$, and $\tilde{\alpha}_4$ are positive coefficients and the error term, $\Delta \tilde{\xi}_{ict}$, is given by $\Delta \tilde{\xi}_{ict} = \tilde{\alpha}_5 \Delta \varepsilon_{ict} + \tilde{\alpha}_6 \sum_j \frac{1}{j} \varepsilon_{jct} + \tilde{\alpha}_7 \Delta \tau_{ct}$. Note that the term associated with the mobility option, $\max_{c'} U_{c't}^u$, is embedded in the time-varying intercept ($\tilde{\alpha}_{it}$) since it does not vary by city.

As with equation (16), equation (19) maintains the property that wages increase with a higher value of R_{ct} ($\tilde{\alpha}_2 > 0$) or a higher employment rate ($\tilde{\alpha}_3 > 0$). Moreover, the error term has a similar structure as before (with the exception that the city amenity, τ_{ct} , now enters). The main substantive difference is the appearance of the cost of housing services in the wage determination process, with $\tilde{\alpha}_4 > 0$. The latter coefficient captures the indirect housing price effect, while $\tilde{\alpha}_2$ captures the direct effect alone. Therefore, the coefficient on R_{ct} in this specification should be smaller than in the case where housing prices are not held constant.

Equation (19) highlights the roles of the two main frictions operating in our model: a friction across cities and a friction across jobs within a city. Our model predicts that R_{ct} will have a direct, causal effect on wages only if both frictions are present. To see this, consider the case where there are no frictions across cities. In this case, jobs within one's city will not be viewed differently when

²⁴Details of this derivation are given in Appendix S.7.

bargaining from jobs in other cities, and the local composition of employment will not have any distinct effect on workers within a city. This would imply that $\tilde{\alpha}_2 = 0$ and R_{ct} could only have an indirect effect on wages through their effect on the local cost of living. Alternatively, if cross-sector frictions are absent, then all sectors in a city will pay the same wage and a pure shift in industrial composition will have no direct causal effect on wages. Therefore, we only expect to find $\tilde{\alpha}_2 > 0$ in the presence of both frictions.

In Table V we report estimates of equation (19). Following Moretti (2010b), our measure of the price of housing corresponds to the rental cost of a 2 or 3 bedroom apartment. In the OLS estimates in column 1, the coefficients on both ΔR_{ct} and $\Delta \log p_{ct}^h$ are significant and have the predicted sign. However, OLS does not provide consistent estimates of the coefficients in equation (19) since the regressors are likely correlated with the error term. Hence, we pursue the estimation of (19) using instrumental variables. It can be verified that our previous instruments—IV1, IV2, and IV3—remain valid in this setting under the same conditions as before. However, since both ΔR_{ct} and $\Delta \log p_{ct}^h$ can be shifted by changes in industrial composition, instruments working only off variation in this composition (as these do) are not sufficient to identify equation (19) because a rank condition is not met. Therefore, we need to augment our instrument set to include an instrument which is correlated with changes in housing prices, independent of changes in $\tilde{\xi}_{ict}$, and works through a mechanism other than changes in industrial composition.

Over our sample period, natural population growth placed upward pressure on land prices and this was an especially strong force in cities where geographical factors make land scarce. This logic suggests using a measure of the local availability of land as an instrument for housing prices, as this is a variable that should be correlated with increases in housing prices but not correlated with city-specific changes in technological knowledge. Saiz (2010) used satellite-generated data to create a measure of the elasticity of land supply at the city level based on five measures of the amount of land available for residential development. We call Saiz's measure of local land availability instrument IV5. The result from estimating equation (19) by instrumental variables using IV1–IV3–IV5 is reported in column 2, using IV2–IV3–IV5 is in column 3, and using all instruments, IV1–IV2–IV3–IV5, is reported in column 4. The first-stage regression of changes in housing prices indicates that Saiz's measure of local land availability is an important and significant negative predictor of housing prices changes.

The results from IV estimation of equation (19) reported in columns 2–4 of Table V reveal that the direct effect of a 1 percentage-point change in R_{ct} , holding local housing cost constant, is about a 2.1% increase in within-industry wages, and this result is found to be robust across our different sets of instruments. If we compare with our previous results from estimating equation (16), this suggests that around 75% of the total GE effect (i.e., an effect of 2.1 relative to our estimate of the total effect of 2.8) reflects the direct effect of indus-

TABLE V
CONTROLLING FOR LOCAL HOUSING COSTS^a

	OLS	IV			Alternate IV		Local Real Wages	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔR_{ct}	1.15*	2.03*	2.18*	2.17*	2.05*	2.26*	1.14*	0.93*
	(0.18)	(0.57)	(0.52)	(0.52)	(0.79)	(0.76)	(0.16)	(0.31)
$\Delta \log p_{ct}^h$	0.35*	0.15	0.14	0.13	0.17	0.14		
	(0.025)	(0.13)	(0.12)	(0.12)	(0.18)	(0.18)		
ΔER_{ct}	0.19*	0.90*	0.79*	0.83*	0.27	0.24	0.21*	-0.53
	(0.080)	(0.35)	(0.38)	(0.35)	(0.18)	(0.18)	(0.081)	(0.47)
Year \times ind.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	28,109	23,879	23,879	23,879	28,109	28,109	28,109	28,109
R^2	0.60						0.56	
IV set:		IV1-IV3-IV5	IV1-IV3-IV5	IV1-IV2-IV3-IV5	IV1-IV3-IV6	IV1-IV3-IV7		IV1-IV2-IV3
<i>F</i> -statistics								
ΔR_{ct}		32.40	74.82	117.50	44.92	87.40		184.01
ΔER_{ct}		4.89	9.04	7.64	36.00	54.26		12.19
Δp_{ct}^h		16.93	18.32	16.44	20.86	21.35		
AP <i>p</i> -values								
ΔR_{ct}		0.00	0.00	0.00	0.00	0.00		0.00
ΔER_{ct}		0.00	0.00	0.00	0.00	0.00		0.00
Δp_{ct}^h		0.00	0.00	0.01	0.02	0.01		
Over-id. <i>p</i> -value		.	.	0.61	.	.		0.72

^aStandard errors, in parentheses, are clustered at the city-year level. The asterisk (*) denotes significance at the 5% level. All models are estimated on a sample of 152 U.S. cities using Census and ACS data for 1970–2007. The dependent variable is the decadal change in regression-adjusted city–industry wages.

trial composition on wages. In turn, 25% of the total effect is attributable to a feedback process by which industrial composition, through wage effects on land prices and construction costs, changes housing costs and thereby further affects the bargaining power of workers. The results in column 3 are particularly telling, since the identification of the effect of industrial composition on within-industry wages is coming from changes in industrial-wage premia at the national level and, therefore, is very unlikely to be confounded with local changes in aggregate labor demand. The results in this column indicate that in a city where one industry starts paying more, even if there are no changes to housing costs or employment rates, wages increase in all other industries in the city. This pattern is supportive of the idea that wages contain an important bargaining component determined by alternative job opportunities available in one's city.

As a robustness check on the latter claim, we also employed alternative IV strategies where we instrumented for $\Delta \log p_{ct}^h$ using an immigrant enclave instrument to predict population inflows and, separately, climate related instruments such as average temperature and rainfall. These results are presented in columns 5 and 6 of Table V: Details of the construction of the instruments are provided in the data section of the Supplemental Material. Both sets of instruments result in estimates of $\tilde{\alpha}_2$ and $\tilde{\alpha}_4$ that are very similar to those based on the land availability instrument.

These data can also be used to examine whether improvements in industrial composition, as captured by an increase in R_{ct} , are associated with an increase in local purchasing-power wages in a city. To explore this, we follow Moretti (2010b) in constructing a local CPI index that is calculated similarly to the official CPI but with housing costs allowed to vary by metropolitan area. To do so, we assign a fraction of a consumers consumption basket to housing (this varies by year and is given by the Bureau of Labor Statistics (BLS)), and combine this with local housing prices and the CPI-U to obtain a local price index.²⁵ Deflating wages by this index, we can see how locally adjusted wages react to changes in industrial composition. We report results from regressions of this type in columns 7 and 8 of Table V. In these columns, we see that an increase in R_{ct} is associated with a significant increase in local-purchasing-power real wages when estimating by either OLS or IV (for the IV results in column 8 of Table V, we use IV1–IV2–IV3 as our instrument set).

The estimation of equations (16) and (19) provides different perspectives regarding the effect of changes in industrial composition on local wages. Which is more interesting, of course, depends on the question one is trying to address. If the question of interest concerns the role of changes in industrial composition in explaining the cross-city pattern in wages, then the total GE effect as

²⁵See Moretti (2010b) for details on the construction of the local price index as well as a discussion on the measurement of local housing prices using Census data.

reported by the estimation of equation (16) is most relevant. In contrast, if one is more interested in understanding the mechanisms through which industrial composition affects workers, then the estimates of equation (19) are more relevant. Finally, if one is focused on understanding how local workers benefit or lose by changes in industrial composition, then the results from columns 7 and 8 of Table V are more insightful, as they report the net effect of local-purchasing-power parity wages.

6.4. Education Breakdowns

To this point, we have treated workers as representing different bundles of efficiency units of labor. In this subsection, we examine whether this is a reasonable approximation by exploring whether our industrial composition effects are present and similar across different skill groups defined by education and experience. We consider three education groups: workers with at most a high school education, workers with some post secondary education but without a BA, and workers with at least a BA. We further divide each of these groups into young workers (those with less than 10 years of experience) and older workers (those with more than 10 years of experience).²⁶ For each of these groups, we calculate an R_{ct} variable that is specific to the group.²⁷ The estimates of the effects of changes in R_{ct} on wages for each of these groups are presented in Table VI. All the results in this table also control for changes in the employment rate.

The results in Table VI indicate substantial, though not identical, effects of shifts in industrial composition on within-industry wages in all the skill groups. The estimates reveal a slight education gradient, with ΔR_{ct} effects being smaller for those with a BA or more in both experience groups. They also indicate substantial experience differences, with the effects of a change in industrial composition on wages being much greater for younger than older workers. While this observation is not consistent with a strict interpretation of the model, it is potentially consistent with an internal labor market view in which more experienced workers are partially sheltered from direct market comparisons. Overall, the skill group results paint a picture similar to that documented previously, that is, changes in industrial composition have substantial general equilibrium effects on within-industry wages as predicted by search and bargaining theory.

²⁶The measure of experience we use is potential experience, defined by age minus years of school minus 6. Appendix S.1 contains details.

²⁷We create the national-level wage premia, and thus the R_{ct} measures, separately for each skill group. We again estimate in two steps, each step run separately for the six skill groups.

TABLE VI
 BREAKDOWN BY EDUCATION AND POTENTIAL EXPERIENCE^a

		OLS	IV		
		(1)	(2)	(3)	(4)
<10	HS	2.65* (0.20)	2.34* (0.39)	2.65* (0.28)	2.57* (0.22)
	SP	2.69* (0.22)	3.01* (0.54)	3.82* (0.52)	3.59* (0.39)
	BA	1.58* (0.16)	0.68 (0.62)	2.59* (0.33)	2.16* (0.33)
>10	HS	1.43* (0.16)	2.90* (0.61)	1.78* (0.31)	2.06* (0.25)
	SP	1.67* (0.30)	0.92 (1.58)	2.03* (0.72)	2.05* (0.63)
	BA	0.98* (0.21)	0.63 (1.30)	1.65* (0.35)	1.62* (0.34)
Year × ind.		Yes	Yes	Yes	Yes
Observations		59,066	59,066	59,066	59,066
R^2		0.59			
Instrument set			IV1-IV3	IV2-IV3	IV1-IV2-IV3

^aStandard errors, in parentheses, are clustered at the city-year level. The asterisk (*) denotes significance at the 5% level. All models are estimated on a sample of 152 U.S. cities using Census and ACS data for 1970–2007. The dependent variable is the decadal change in regression-adjusted city-industry wages.

7. A QUANTITATIVE ASSESSMENT OF THE IMPORTANCE OF INDUSTRIAL COMPOSITION FOR CITY WAGES

We motivated this paper with the question, “What are the effects of industrial composition on wages?” To help answer this question, we proposed and evaluated a search and bargaining model which formalized how industrial composition and wages can interact. Given that we found considerable empirical support for the model, we are now in a position to use its parameters to shed light on our initial question.

Our goal in this section is to quantify the importance of changes in industrial structure in explaining the cross-city variance in wage growth. Our approach is to investigate how much of the cross-city variance in wage growth in our sample period can be explained by shifts in cities’ industrial structures. The variable we focus on is the cross-section variance of changes in average city wages, that is, the variance across cities of $\sum_i \eta_{ict+1} w_{ict+1} - \sum_i \eta_{ict} w_{ict}$, where w_{ict} is, as before, the human capital adjusted price of labor in industry i in city c . To highlight the implications of our model, we contrast our decomposition of this variance with that obtained from the common approach, which disregards the type of GE effects implied by our bargaining setup. In particular, if changes in industrial structure had no effect on wage determination, then it would be appropriate

TABLE VII
CROSS-CITY VARIANCE IN WAGES EXPLAINED BY INDUSTRIAL COMPOSITION CHANGES^a

Decade	Variance Ratio				
	(1)	(2)	(3)	(4)	(5)
1980s	0.12	0.64	0.02	0.33	0.27
1990s	0.16	0.83	0.04	0.41	0.44
2000s	0.08	0.66	0.01	0.29	0.28

^aColumn 1 denotes % variance explained using partial equilibrium methodology (shift share). Column 2 denotes % variance explained using estimated general equilibrium effects. Column 3 denotes % variance attributed to national-level changes in industrial composition using partial methodology. Column 4 denotes % variance attributed to national-level changes in industrial composition using general equilibrium methodology. Column 5 denotes % variance attributed to national-level changes in industrial composition using general equilibrium methodology netting out effects of international trade.

to calculate the effect of industrial composition changes on cross-sectional city wage growth by examining the ratio of variances

$$(20) \quad \frac{\text{Var}\left(\sum_i \eta_{ict+1} w_{ict} - \sum_i \eta_{ict} w_{ict}\right)}{\text{Var}\left(\sum_i \eta_{ict+1} w_{ict+1} - \sum_i \eta_{ict} w_{ict}\right)}.$$

This is commonly referred to as the “between” component in a shift-share analysis. We report the ratio in column 1 of Table VII for each of our three time periods. The results indicate that this direct effect of industrial structure change on average wage growth is quite small, ranging from 8 to 16%.

Our model, however, indicates that the latter exercise provides only a partial measure since a change in industrial structure is accompanied by within-sector changes in wages due to the induced change in the bargaining position of workers. The model implies that the latter effect is captured by our R_{ct} index. Thus, according to our model, a more meaningful statistic describing the contribution of industrial composition to changes in city wages is

$$(21) \quad \frac{\text{Var}\left(\sum_i \eta_{ict+1} (w_{ict} + \alpha_2 \Delta R_c) - \sum_i \eta_{ict} w_{ict}\right)}{\text{Var}\left(\sum_i \eta_{ict+1} w_{ict+1} - \sum_i \eta_{ict} w_{ict}\right)}.$$

In column 2 of Table VII, we report the statistic associated with (21) when α_2 is set to 2.8, our preferred estimate of the coefficient on R_{ct} from our IV estimates. As can be seen, this measure suggests that changes in industrial composition account for a very substantive fraction of the overall city-level wage

movements, with the fraction ranging from 64 to 83%. This gives a very different picture from the standard decomposition, with industrial composition change now appearing to be a central factor in city-level outcomes.

The effects reported in columns 1 and 2 of Table VII are meant to capture the effects of overall changes in industrial composition, regardless of the source of the change. In particular, in these calculations we are using actual changes in industrial shares from η_{ict} to η_{ict+1} . Within the model, such changes can arise because of location-specific shocks or, alternatively, they can arise as a city adjusts to national-level industrial changes. For example, a national-level shift in demand for textiles versus computers will affect some cities more than others. We can evaluate the latter component by replacing η_{ict+1} in (20) and (21) (including in the calculation of R_{ct+1}) with the $\hat{\eta}_{ict+1}$ used in the construction of IV1. This assigns national-level industrial growth to cities based on their start of decade comparative advantage as reflected in their start of decade industrial shares. The counterfactual city wage changes created in this way leave out the locally driven component of composition shifts. When we do this, the results from the analogue of (20) are quite small (see column 3). However, the results from the analogue of (21) shown in column 4 provide a very different picture. The effects now range from 29 to 44%, indicating that the local changes in industrial structure induced by national-level forces account for a sizable 50% of the overall changes induced by industrial composition shifts, leaving the other 50% as a reflection of city-level developments that influence industrial structure.

Since the results from column 4 in Table VII suggest that national-level forces are quantitatively important for understanding city-level outcomes, we now ask whether international trade, and more especially trade in manufacturing goods, is playing a substantial role. To this end, we use data on manufacturing trade to build a counterfactual to $\hat{\eta}_{ict+1}$. This counterfactual, denoted $\tilde{\eta}_{ict+1}$, is built in two steps. First, we use trade data to build a counterfactual for national-level employment growth by industry, where for manufacturing industries, we subtract from actual rates of growth the observed growth in the net exports to output ratio under the assumption that the displaced employment due to trade is proportional to the change in the net exports to output ratio. We then use the counterfactual national-level growth to predict city-level industrial structure changes in the same manner used to build $\hat{\eta}_{ict+1}$. We finally calculate the analogues to the statistic (21), where we replace η_{ict+1} with $\tilde{\eta}_{ict+1}$ (including in the calculation of R_{ct+1}). The results (reported in column 5) are very similar to those in column 4, implying that manufacturing trade effects played a small role in explaining cross-city outcomes. This occurs because changes in the net exports to output ratio are rather small for most industries and, therefore, $\tilde{\eta}_{ict+1}$ is close to $\hat{\eta}_{ict+1}$. Of course, we recognize that our measure of the impact of trade is quite crude and that trade may still have an impact on national-level wage growth which these calculation neglects.

8. CONCLUSION

In this paper we argue that a standard search and bargaining model, once extended to include multiple sectors, has interesting testable implications that have not been previously recognized. Chief among these is that shifts in the composition of employment in an economy, either toward or away from high-wage industries, should have an impact on wages within all sectors in that economy. This arises because changes in wage options elsewhere in the economy alter the outside options for a worker bargaining with his or her employer. Using data from the U.S. Census and the American Community Survey for the period 1970–2007, we show that these spillover effects are 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 2.5–4 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, controlling for differences in local cost of housing, and focusing on sectors producing highly tradeable goods. Taken together, we view these results as strong evidence in favor of using search and bargaining models to understand aggregate wage determination.

Our results also suggest that policies or events which affect industrial composition should not be evaluated simply using the standard accounting approach, but, instead, evaluations should explicitly take into account the substantial general equilibrium effects implied by the type of social interaction model of wage determination presented here. For example, it is common for technological change, changes in preferences, or the opening up of trade relationships to involve a real location of high- and low-paying jobs across time and space. Our results suggest that a proper evaluation of the effects of such industrial reallocation needs to incorporate the general equilibrium effects on wage bargaining in other sectors, even when the change has no aggregate effect on employment. In particular, recognizing and quantifying these feedback effects should lead to more accurate assessments of the effects of industrial composition changes in explaining wage outcomes.²⁸

APPENDIX: WAGE EQUATION DERIVATION AND CONSISTENCY

A.1. *Wage Equation Derivation*

We begin by rewriting equation (10) so that the sectoral wage is expressed only as a function of nationally determined prices and of the exogenous pro-

²⁸Beaudry, Collard, and Green (2005) and Beaudry and Collard (2006) found that increased openness to international trade over the period 1978–1998 had very uneven effects across countries. In particular, countries that attracted high-capital–high-wage industries gained disproportionately relative to countries that increased employment in low-capital intensive industries. The general equilibrium effects found in this paper offer a potential explanation for the size of the effects found in those two papers.

ductivity terms, ε_{ic} :

$$(22) \quad w_{ic} = d_{ic} + \gamma_{c1} \varepsilon_{ic} + \left(\frac{\gamma_{c2}}{1 - \gamma_{c2}} \right) \gamma_{c1} \sum_j \eta_{jc} (p_j - p_1) + \gamma_{c1} \left(\frac{\gamma_{c2}}{1 - \gamma_{c2}} \right) \sum_j \eta_{jc} \varepsilon_{jc},$$

where $d_{ic} = \gamma_{c0} (1 + \frac{\gamma_{c2}}{1 - \gamma_{c2}}) + \gamma_{c1} (\frac{\gamma_{c2}}{1 - \gamma_{c2}}) p_1 + \gamma_{c1} p_i$. In equation (22), we express prices in relation to the price of an arbitrarily chosen good denoted p_1 to help emphasize how a pure-shift change in industrial composition affects wages (by a pure shift, we mean a change in the η s that does not change the total number of jobs).

We can now use equation (10) or (22) to relate the price differential $p_i - p_1$ to the average national-level wage premium in industry i relative to a baseline industry 1. This is achieved by noting that $w_{ic} - w_{1c}$ is equal to $\gamma_{c1} (p_i - p_1) + \gamma_{c1} (\varepsilon_{ic} - \varepsilon_{1c})$. If we take the average of $w_{ic} - w_{1c}$ across cities, we obtain

$$(23) \quad \nu_i = w_i - w_1 = \gamma_1 (p_i - p_1) + \hat{d}_i,$$

where ν_i is the national-level wage premium in industry i relative to industry 1, γ_1 is the average of γ_{c1} across cities, and \hat{d}_i is an industry-specific constant.²⁹ Substituting (23) into (22), we get

$$(24) \quad w_{ic} = \tilde{d}_{ic} + \left(\frac{\gamma_{c2}}{1 - \gamma_{c2}} \right) \left(\frac{\gamma_{c1}}{\gamma_1} \right) R_c + \gamma_{c1} \varepsilon_{ic} + \gamma_{c1} \left(\frac{\gamma_{c2}}{1 - \gamma_{c2}} \right) \sum_j \eta_{jc} \varepsilon_{jc},$$

where $R_c = \sum_j \eta_{jc} \nu_j$ and $\tilde{d}_{ic} = d_{ic} + (\frac{\gamma_{c2}}{1 - \gamma_{c2}}) (\frac{\gamma_{c1}}{\gamma_1}) \sum_j \eta_{jc} \hat{d}_j$.

In interpreting the effect of R_c on wages from equation (24), we are performing a partial equilibrium exercise, as we treat the employment rate in a city (as reflected in the γ parameters), and the sectoral composition, as given. To capture the dependence of wages on the city's employment rate more explicitly, it is useful to take a linear approximation of (24) around the point where cities have identical employment rates ($ER_c = ER$) and industrial composition ($\eta_{ic} = \eta_i$), which arises when $\varepsilon_{ic} = 0$, $\Omega_{ic} = 0$, $Y_i = Y$, and $a_i = a$. Moreover, at this point, the employment shares, η_i s, are equal across sectors. Differencing the resulting linear approximation yields equation (11) in the text.

²⁹The industry-specific constant \hat{d}_i is equal to $\sum_c (\gamma_{c1} - \gamma_1) (\varepsilon_{ic} - \varepsilon_{1c})$. When we later take a first-order approximation around an equilibrium where the ε s equal zero, this term will be equal to zero.

A.2. Examining Consistency

As described in the text, we are interested in the condition³⁰

$$(25) \quad \lim_{C, I \rightarrow \infty} \frac{1}{I} \frac{1}{C} \sum_{i=1}^I \sum_{c=1}^C \Delta R_c \Delta \xi_{ic},$$

which, using $R = \sum_j \eta_{jc}(w_j - w_1)$, can be written as

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

or

$$(26) \quad \lim_{C, I \rightarrow \infty} \frac{1}{I} \frac{1}{C} \left[\sum_j (w_j - w_1) \sum_c \Delta \eta_{jc} \sum_i \Delta \xi_{ic} \right. \\ \left. + \sum_j \Delta(w_j - w_1) \sum_c \eta_{jc} \sum_i \Delta \xi_{ic} \right].$$

We handle the limiting arguments sequentially, allowing $C \rightarrow \infty$ first. Then we focus our concern to two components in (26), which we handle in turn. The first is

$$(27) \quad \lim_{C \rightarrow \infty} \frac{1}{C} \sum_c \Delta \eta_{jc} \sum_i \Delta \xi_{ic}.$$

Given the decomposition $\varepsilon_{ic} = \bar{\varepsilon}_c + v_{ic}^e$, where $\sum_i v_{ic}^e = 0$, we get

$$\Delta \eta_{jc} = \pi_1(\Delta v_{jc}^e) + \pi_2(\Delta p_j \Omega_{jc} - \Delta p \bar{\Omega}_c),$$

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

Also,

$$(28) \quad \sum_i \Delta \xi_{ic} = \left(\gamma_1 + \frac{\gamma_1 \gamma_2}{1 - \gamma_2} \right) I \Delta \bar{\varepsilon}_c.$$

Then, given that $E(\Delta \bar{\varepsilon}_c) = 0$ (again, recalling that we have removed economywide trends), and if $\Delta \bar{\varepsilon}_c$ is independent of Δv_{ic}^e and $(\Delta p_j \Omega_{jc} - \Delta p \bar{\Omega}_c)$, it is straightforward to show that (27) equals zero.

³⁰Throughout this appendix we omit the t subscript for simplicity.

The second component is

$$(29) \quad \lim_{C \rightarrow \infty} \frac{1}{C} \sum_c \eta_{jc} \sum_i \Delta \xi_{ic},$$

where $\sum_i \Delta \xi_{ic}$ is again given by (28), while η_{jc} is given by (13). For (29) to be zero, we require in addition that $\Delta \bar{\varepsilon}_c$ be independent of past values of v_{ic}^e and of $(p_j \Omega_{jc} - p \bar{\Omega}_c)$. Thus, if $\Delta \bar{\varepsilon}_c$ is independent of the past and is independent of Δv_{ic}^e and $(\Delta p_j \Omega_{jc} - \Delta p \bar{\Omega}_c)$, then (25) equals zero and OLS is consistent.

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

$$(30) \quad \lim_{C, I \rightarrow \infty} \frac{1}{I} \frac{1}{C} \sum_{i=1}^I \sum_{c=1}^C Z_c \Delta \xi_{ic}.$$

For what we call IV1,

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

where $g_j^* = \frac{1+g_j}{\sum_k \eta_{kc} (1+g_k)}$ and g_j is the growth rate in employment in industry j at the national level. Given this, (30) becomes:

$$(31) \quad \lim_{C \rightarrow \infty} \frac{1}{C} \sum_j (w_j - w_1) \sum_c \eta_{jc} (g_j^* - 1) \sum_i \Delta \xi_{ic}.$$

Thus, (31) equals zero under the same conditions under which (29) equaled zero, that is, that $E(\Delta \bar{\varepsilon}_c) = 0$ and $\Delta \bar{\varepsilon}_c$ is independent of past values of v_{ic}^e and $(P_j \Omega_{jc} - P \bar{\Omega}_c)$. Obviously, this condition is satisfied if $\bar{\varepsilon}_c$ behaves as a random walk.

Similarly, the relevant condition when using IV2 is given by

$$(32) \quad \lim_{C \rightarrow \infty} \frac{1}{C} \sum_j \Delta (w_j - w_1) \sum_c \eta_{jc} \sum_i \Delta \xi_{ic},$$

and the same conditions ($\Delta \bar{\varepsilon}_c$ is independent of past values of v_{ic}^e and of $(p_j \Omega_{jc} - p \bar{\Omega}_c)$) ensure that (32) equals zero.

Several points follow from this discussion. First, OLS can provide consistent estimates under the strong assumption that changes in absolute advantage and changes in comparative advantage are independent. Thus, if OLS and the two IV estimates are equal, then this is a test of the stronger assumption about independence in changes. Second, if the key identifying assumption underlying

the IVs is not true (i.e., changes in absolute advantage are not independent of past comparative advantage), then the two IVs weight the problematic correlation (between $\Delta \bar{\varepsilon}_c$ and v_{ic}^e) differently (in particular, IV1 weights using the weights $(w_j - w_1)$, while IV2 uses the weights $\Delta(w_j - w_1)$), and estimates based on the different IVs should differ.

REFERENCES

- ACEMOGLU, D., AND J. ANGRIST (1999): "How Large Are the Social Returns to Education? Evidence From Compulsory Schooling Laws," Working Papers 7444, National Bureau of Economic Research, Inc. [1078,1085]
- BARTIK, T. J. (2002): "Spillover Effects of Welfare Reforms in State Labor Markets," *Journal of Regional Science*, 42 (4), 667–701. [1063]
- BEAUDRY, P., AND F. COLLARD (2006): "Globalization, Returns to Accumulation, and the World Distribution of Income," *Journal of Monetary Economics*, 53 (5), 879–909. [1099]
- BEAUDRY, P., F. COLLARD, AND D. GREEN (2005): "Demographics and Recent Productivity Performance: Insights From Cross-Country Comparisons," *Canadian Journal of Economics*, 38 (2), 309–344. [1099]
- BEAUDRY, P., D. A. GREEN, AND B. SAND (2012): "Supplement to 'Does Industrial Composition Matter for Wages? A Test of Search and Bargaining Theory'," *Econometrica Supplemental Material*, 80, http://www.econometricsociety.org/ecta/Supmat/8659_data_description.pdf; http://www.econometricsociety.org/ecta/Supmat/8659_data_and_programs.zip. [1068,1079]
- BLANCHARD, O. J., AND L. F. KATZ (1992): "Regional Evolutions," *Brookings Papers and Economic Activity*, 1, 1–75. [1066,1076,1085]
- BLUESTONE, B., AND B. HARRISON (1982): *The Deindustrialization of America*. New York: Basic Books. [1063]
- BLUNDELL, R., A. DUNCAN, AND C. MEGHIR (1998): "Estimating Labor Supply Responses Using Tax Reforms," *Econometrica*, 66 (4), 827–862. [1065]
- BOUND, J., AND G. JOHNSON (1992): "Changes in the Structure of Wages in the 1980's: An Evaluation of Alternative Explanations," *American Economic Review*, 82, 371–392. [1064]
- DAHL, G. B. (2002): "Mobility and the Return to Education: Testing a Roy Model With Multiple Markets," *Econometrica*, 70 (6), 2367–2420. [1080,1084]
- FEYRER, J. D., B. SACERDOTE, AND A. D. STERN (2007): "Did the Rust Belt Become Shiny? A Study of Cities and Counties That Lost Steel and Auto Jobs in the 1980s," *Brookings-Wharton Papers on Urban Affairs*, 41–89. [1066]
- GLAESER, E., AND J. GOTTLIEB (2009): "The Wealth of Cities: Agglomeration Economies and Spatial Equilibrium in the United States," *Journal of Economic Literature*, 47 (4), 983–1028. [1064]
- GLAESER, E., H. KALLAL, J. SCHEINKMAN, AND A. SHLEIFER (1992): "Growth in Cities," *Journal of Political Economy*, 100 (6), 1126–1152. [1078,1085]
- GREENSTONE, M., AND E. MORETTI (2003): "Bidding for Industrial Plants: Does Winning a 'Million Dollar Plant' Increase Welfare?" Working Papers 9844, National Bureau of Economic Research, Inc. [1066]
- JENSEN, B., AND L. KLETZER (2005): "Tradable Services: Understanding the Scope and Impact of Services Outsourcing," Working Paper 05-9, Institute for International Economics. [1087]
- LEMIEUX, T. (2002): "Decomposing Changes in Wage Distributions: A Unified Approach," *Canadian Journal of Economics*, 35 (4), 646–688. [1064]
- MANSKI, C. F. (1993): "Identification of Endogenous Social Effects: The Reflection Problem," *Review of Economic Studies*, 60 (3), 531–542. [1065,1071]
- MOFFITT, R. A. (2001): "Policy Interventions, Low-Level Equilibria and Social Interactions," in *Social Dynamics*, 45–82. Cambridge: MIT Press. [1065,1071]

- MORETTI, E. (2004a): “Estimating the Social Return to Higher Education: Evidence From Longitudinal and Repeated Cross-Sectional Data,” *Journal of Econometrics*, 121 (1–2), 175–212. [1078,1085]
- (2004b): “Workers’ Education, Spillovers, and Productivity: Evidence From Plant-Level Production Functions,” *American Economic Review*, 94 (3), 656–690. [1088]
- (2010a): “Local Labor Markets,” Discussion Papers 4905, Institute for the Study of Labor (IZA). [1063,1064]
- (2010b): “Local Multipliers,” *American Economic Review*, 100 (2), 373–377. [1064,1092, 1094]
- MORTENSEN, D. T., AND C. A. PISSARIDES (1999): “New Developments in Models of Search in the Labor Market,” in *Handbook of Labor Economics*, Vol. 3, ed. by O. Ashenfelter and D. Card. Amsterdam: Elsevier, Chap. 39, 2567–2627. [1064]
- OECD (2010): “Competition, State Aids and Subsidies: Contribution From the U.S. Federal Trade Commission,” Report, Global Forum on Competition. [1063]
- SAIZ, A. (2010): “The Geographic Determinants of Housing Supply,” *Quarterly Journal of Economics*, 125 (3), 1253–1296. [1092]
- SAND, B. M. (2006): “The Social Returns to Education,” Working Paper, University of British Columbia. [1085]

Dept. of Economics, University of British Columbia, 997-1873 East Mall, Vancouver, B.C., Canada, V6T 1Z1 and NBER; paulbe@interchange.ubc.ca,
Dept. of Economics, University of British Columbia, 997-1873 East Mall, Vancouver, B.C., Canada, V6T 1Z1 and IFS, London; green@econ.ubc.ca,
and
Dept. of Economics, York University, 4700 Keele Street, Toronto, ON, Canada, M3J 1P3; bmsand@econ.yorku.ca.

Manuscript received June, 2009; final revision received July, 2011.