Agglomeration and Productivity: New Estimates and Macroeconomic Implications*

Morris A. Davis
University of Wisconsin
mdavis@bus.wisc.edu

Jonas D.M. Fisher
Federal Reserve Bank of Chicago
jfisher@frbchi.org

Toni M. Whited
University of Rochester
toni.whited@simon.rochester.edu

August 24, 2009

Preliminary and Incomplete

Abstract

We construct a dynamic general equilibrium model of cities with aggregate balanced growth and use it to estimate the effects of local agglomeration on per capita consumption growth. Agglomeration affects growth through the density of economic activity: higher production per unit of land raises local productivity. Firms take productivity as given and produce with constant returns in developed land, physical capital and labor. Land and capital may be accumulated. If the relative price of new developed land is rising, as it is empirically, density is rising and contributes to growth. Our model predicts an empirical relationship between wages, housing rents, and labor inputs. We use this to estimate the net effect of agglomeration on local productivity with a panel of US cities. We estimate that doubling output in a location raises productivity of firms in that location by 6.9%. Our preferred estimate indicates that agglomeration increases the growth rate of per capita consumption by 10.7%.

JEL Classification Numbers: E0, E1, E2, O4, R0, R1, R3

Keywords: Balanced Growth, Economic Growth, Productivity, Externalities, Increasing Returns, Agglomeration, Density

*The views expressed herein are those of the authors and do not necessarily represent those of the Federal Reserve Bank of Chicago or the Federal Reserve System.

1 Introduction

There is widespread agreement that cities emerge because of *agglomeration* effects. For example, workers are more productive when they are near other workers and production is more efficient when transportation costs are reduced. Cities emerge because of agglomeration effects, and most aggregate growth actually occurs in cities. So, how important is local agglomeration for aggregate growth? To answer this question, we build and estimate a dynamic general equilibrium model of cities and aggregate growth in which city-based agglomeration affects per capita consumption growth. We estimate the structural parameters of our model using panel data and use these estimates to quantify the impact of local agglomeration on aggregate per capita consumption growth.

Our model embeds into the neoclassical growth model a version of Roback (1982)'s model of cities, augmented with agglomeration effects in the way proposed by Ciccone and Hall (1996). Production is a function of capital, labor and land. Productivity of each firm in a city is increasing in the total quantity of city output per unit of land used in production, but firms do not take this into account when making their decisions. We study the model's stationary competitive equilibrium and show that along the balanced growth path per capita consumption growth depends on exogenous neutral technical change, the rate of increase in the cost of developing new land for production, a parameter governing the net effect of agglomeration on productivity, and labor's share of income. Using our model and panel data for a sample of US cities we estimate the net impact of the density of economic activity on local productivity, thereby yielding the contribution of agglomeration to aggregate per capita consumption growth.

Our model predicts a specific relationship within each city between wages, output prices, land prices, labor inputs, and the level of the neutral technology. We exploit this relationship to estimate the key parameter underlying the impact of agglomeration on local productivity and aggregate growth: Variation in the price of land and wages, holding output prices and the composition of the workforce fixed, provides identification of the size of agglomeration effects, as suggested by Lucas (2001). Our data set merges Metropolitan Statistical Area (MSA) level wage and labor input data from the Current Population Survey, output price data from the Bureau of Economic

Analysis (BEA), and data on housing rental prices constructed using information from the 1990 Decennial Census of Housing (DCH) and the Bureau of Labor Statistics (BLS). This annual panel data-set covers 22 MSAs over the 1985-2006 period. Our twenty-plus year panel samples are particularly amenable to applying modern dynamic panel data methods and we use the Arellano and Bover (1995) estimation strategy.

We find that that local agglomeration effects are small but statistically significant: a doubling of output per unit of productive land in a city increases productivity of each firm in the city, holding all other inputs fixed, by between 6 and 7 percent. In more dense locations, land is more expensive and firms use less land as an input to production, explaining why we find the net impact of agglomeration on wages and productivity of only 2 percent. While seemingly small at the the local level, our model implies agglomeration has a substantial effect on aggregate consumption. Combined with evidence on trends in the cost of new productive land, the agglomeration estimate and our model of balanced growth imply that local agglomeration accounts for about 11% of per capita consumption growth.

Our work is closely related to that of Ciccone and Hall (1996). Using a similar model of agglomeration they estimate a much larger net effect of agglomeration on local productivity. Our analysis goes beyond the analysis of Ciccone and Hall (1996) in several ways. First, we show how to translate estimates of the net impact of agglomeration on local productivity into the impact of agglomeration on aggregate growth. Second, we use 20 years of panel data instead of a single cross-section to identify the size of net agglomeration effects on wages. Third, we allow for alternative uses of land at a location compared to assuming that all land at a location is used in production. In our empirical analysis we show how these differences affect our results.

The paper by Ciccone and Hall (1996) can be viewed as part of a larger literature that measures the impact of agglomeration externalities on firm productivity and wages: See Rosenthal and Strange (2004) for a recent survey. Typically, this literature distinguishes between "urbanization" effects – the effects of the overall size and employment density of a market on wages and productivity in all industries on that market – from "localization" effects, which are the impact of the concentration of a given industry on wages and productivity in that industry.[1] Ciccone and Hall (1996) estimate that a doubling of output density increases wages by 5-1/2 percent. This estimate is within the range of estimates of a set of recent papers, including Carlino and Voith (1992), Ciccone (2002) and Combes et al. (2008). Summarizing, these papers find that a doubling of output or employment density is associated with an increase in worker wages of somewhere between 3 to 8 percent.[2]

Relative to this other work, we find a smaller net impact of agglomeration on wages of 2 percent. We show the impact of agglomeration on per-capita consumption growth is quite sensitive to this estimate. If we use the estimate of Ciccone and Hall (1996), for example, the contribution of agglomeration to per-capita consumption growth increases significantly, from 11 percent to 16 percent.

The rest of the paper is organized as follows. The next section describes the production side of our model economy, the implied balanced growth path, and the equation underlying our estimation of the effect of local agglomeration on per capita consumption growth. Section three describes the rest of our model. After this we outline our empirical strategy and then the data underlying our estimates. Section six discusses our empirical results and section seven quantifies the impact of agglomeration on per capita consumption growth. Section eight concludes.

2 Neo-Classical Growth with Local Production

Our model extends the neo-classical growth model to include location-specific production, consumption and housing; local developed land as a produced, durable input into production and housing; and local agglomeration effects as proposed by Ciccone and Hall (1996). We focus in this section on the production side of the economy. This is sufficient to describe the balanced growth path of the model's competitive equilibrium and derive the equation which underlies our estimation of the impact of local agglomeration on per capita consumption growth. We show how to extend the production side of our model for empirical purposes.

[1] Available evidence suggests localization effects are also important determinants of productivity: See Henderson (2003) and Rosenthal and Strange (2003).

[2] See Ciccone and Peri (2006) for a larger recent estimate.

2.1 The Producers' Environment

Time is discrete and a period is equal to one year. The economy consists of a unit measure of locations, called cities. Cities are indexed on the unit interval or by their current state (s, z), where s denotes the beginning-of-period stock of developed land in the city, and z denotes exogenous city-specific productivity. Developed land is used in housing workers and for production.

Competitive firms in each city produce an intermediate good unique to the city, y, with inputs of developed land, l_b, capital, k_b, and labor, n, using a constant-returns to scale production function,

$$y = a l_b^{1-\phi} k_b^{\alpha\phi} n^{(1-\alpha)\phi}. \tag{1}$$

where a is total factor productivity (TFP), $0 \leq \alpha \leq 1$, and $0 \leq \phi \leq 1$. TFP is taken by firms to be given and is specified as

$$a = z^{(1-\alpha)\phi} \left[\frac{Y}{L_b} \right]^{\frac{\lambda-1}{\lambda}},$$

where Y and L_b are total city output and developed land used in production.[3] Since we normalize the number of intermediate good producers in a city to one, in the competitive equilibrium city-specific aggregates equal those of the representative intermediate good producer in that city, $Y = y$ and $L_b = l_b$.

The ratio of output to land is called the *density of economic activity*. If $\lambda = 1$ density has no impact on productivity and if $\lambda > 1$ firms' productivity is increasing in density. The parameter λ measures the percent increase in productivity achieved by moving a firm holding its inputs of capital, labor, and land fixed) to a city with double the output on the same amount of land. Ciccone and Hall (1996) show how this model of TFP can be derived as the reduced form of a micro-founded model.

We assume the cities share a common trend in productivity and that there are

[3] If the net impact of density on production is positive, $\phi\lambda > 1$, then city output is diminishing in land. This implies that if firm's were to internalize the density externality they would choose as little land as possible to produce, yielding infinite productivity. This is in contrast to Rossi-Hansberg and Wright (2007)'s model where a human capital externality is internalized by the device of a competitive market for developers who use a scheme of subsidies and taxes to induce the first best outcome.

city-specific transitory shocks around the common trend. Specifically,

$$z_t = \gamma^t \tilde{z}_t, \tag{2}$$

where $\tilde{z}_t$ follows a stationary first order Markov process, t denotes calendar time, and $\gamma \geq 1$. Productivity obeys the transition function Q where $Q(z, z')$ is the probability that $z_{t+1} = z'$ conditional on $z_t = z$.

In each city there are also competitive housing service providers. Per capita housing services, h, are produced with developed land, l_h, and physical capital, k_h, according to $h = k_h^{\omega} l_h^{1-\omega}$, where $0 \leq \omega \leq 1$.

Competitive firms produce final goods by combining the output of each city using a constant-elasticity-of-substitution production function in which, for simplicity, the goods produced by each city enter symmetrically. Let $\mu_t(S, Z)$ denote the probability measure of cities with developed land in the set S and productivity shocks in the set Z, at date t. Then, final good production is given by

$$\bar{Y}_t = \left[\int y(s, z)^{\eta} \mu_t(ds, dz) \right]^{\frac{1}{\eta}}. \tag{3}$$

Normalizing the number of final goods producers to unity, $\bar{Y}_t$ corresponds to aggregate output of final goods. Final good output is converted one-for-one into consumption, physical capital investment goods, and into developed land investment goods, x, at the rate v. We assume that land development costs grow at the rate $\tau \geq 1$. Land development costs that increase over time capture the idea that the best land is the first to be developed; the productivity of newly developed land is lower the later it is developed.

Competitive land developers augment the stock of developed land subject to adjustment costs. The land accumulation technology is

$$s_{t+1} = s_t^{1-\zeta} x_t^{\zeta}, \tag{4}$$

where $0 \leq \zeta \leq 1$. This kind of capital augmentation technology is used by Rossi-Hansberg and Wright (2007) since it yields analytical solutions in their environment. We use it because it is a simple way to induce heterogeneity in new developed land prices into our model. This heterogeneity is necessary for there to be cross-sectional variation in land rents, which we use to identify the key agglomeration coefficient.

The developed land, housing and labor services are traded in local markets at land rent $p^l(s,z)$, housing rent $p(s,z)$, wages $w(s,z)$, and a price for developed land $b(s,z)$. Final goods, intermediate goods, and capital services are traded in economy-wide markets at the numeraire price for consumption and physical capital investment goods, the price v for developed land investment goods, the intermediate good prices $q(s,z)$, and the aggregate physical capital rental rate r.

2.2 The Producers' Optimization Problems

We now describe the behavior of local producers of intermediate goods and housing, final good producers and land developers.

2.2.1 Intermediate Goods and Housing Services

Profit maximization by intermediate good producers with TFP $a(s,z)$ yields the following first order conditions, which must be satisfied by the choices $l_b(s,z)$, $k_b(s,z)$, and $n(s,z)$,

$$p^l(s,z) = q(s,z)(1-\phi)a(s,z)l_b(s,z)^{-\phi}k_b(s,z)^{\alpha\phi}n(s,z)^{(1-\alpha)\phi}; \quad (5)$$

$$r = q(s,z)\alpha\phi a(s,z)l_b(s,z)^{1-\phi}k_b(s,z)^{\alpha\phi-1}n(s,z)^{(1-\alpha)\phi}; \quad (6)$$

$$w(s,z) = q(s,z)(1-\alpha)\phi a(s,z)l_b(s,z)^{1-\phi}k_b(s,z)^{\alpha\phi}n(s,z)^{(1-\alpha)\phi-1}. \quad (7)$$

Let $p(s,z)$ denote the price of one unit of rented housing services and $p^l(s,z)$ denote the price of a unit of rented developed land in a city with state vector (s,z). Profit maximization by housing service providers yields

$$p^l(s,z) = \omega^{\frac{\omega}{1-\omega}}(1-\omega)p(s,z)^{\frac{1}{1-\omega}}r^{\frac{-\omega}{1-\omega}}\ . \quad (8)$$

Input use by intermediate good producers and housing service providers is limited by their supply. The local physical capital and developed land market clearing conditions are

$$n(s,z)k_h(s,z) + k_b(s,z) = k(s,z), \quad \forall\,(s,z); \quad (9)$$

$$n(s,z)l_h(s,z) + l_b(s,z) = s, \quad \forall\,(s,z), \quad (10)$$

where $k(s,z)$ denotes the supply of physical capital to a city with state vector (s,z).

Market clearing in the aggregate physical capital and labor markets requires

$$\int k(s,z)\mu_t(ds,dz) \leq K_t; \tag{11}$$

$$\int n(s,z)\mu_t(ds,dz) \leq 1, \tag{12}$$

where K_t is the date t aggregate stock of capital and we have normalized the number of workers in the economy to unity, assuming, for simplicity, that there is no population growth.

2.2.2 Final Goods

Profit maximization of final goods producers yields

$$q(s,z) = \bar{Y}^{1-\eta} y(s,z)^{\eta-1}. \tag{13}$$

The aggregate demand for final goods is limited by their supply. Let C denote aggregate consumption and κ denote the depreciation rate of physical capital. The resource constraint for final goods is that total production of final goods must be at least as large as the sum of its uses:

$$C_t + K_{t+1} - (1-\kappa)K_t + v_t \int x(s,z)\mu_t(ds,dz) \leq \bar{Y}_t. \tag{14}$$

2.2.3 Land Development

Land developers in each city manage the stock of developed land to maximize profits, by renting developed land to intermediate good producers and housing service providers and augmenting the stock of land using (4). Let $W(s,z)$ denote the expected value of a land developer in a city with state vector (s,z) and let '$'$' denote next period's value of a variable. The recursive representation of the land developer problem is then

$$W(s,z) = \max_{\{s',x\}} p^l(s,z)s - vx + \frac{1}{R}\int W(s',z')Q(z,z')dz' \tag{15}$$

subject to the land accumulation technology (4).

The variable R denotes the rate of return on a one period bond. We do not impose that land is irreversibly developed, that is $s_{t+1} \geq s_t$, but we could easily include this constraint.

The first order conditions for the representative land developer can be written

$$b(s,z) = \frac{1}{R} \int \left[p^l(s',z') + (1-\zeta) b(s',z') \left[s'(s,z) \right]^{-\zeta} \left[x(s',z') \right]^{\zeta} \right] Q(z,z') dz'; \quad (16)$$

$$v = \zeta b(s,z) \left[s'(s,z) \right]^{1-\zeta} \left[x(s,z) \right]^{\zeta-1}. \quad (17)$$

Equation (16) says the price of new land in each city, $b(s,z)$, at every date equals the discounted expected value of rent on that land, $p^l(s,z)$, plus the price of the new land in the following period. The price of the new land in the following period is its marginal product in developing new land times the price of new land in the next period. This can be thought of as the price of *existing* developed land, a concept analogous to installed physical capital. Equation (17) equates the marginal benefit of final goods in new land development to its cost. Free entry into land development drives $W(s,z)$ to the value of selling existing developed land to other developers.

2.3 Balanced Growth

We focus on a balanced growth path (with no uncertainty) where employment in all cities grows at the same rate. The variables in the model's stochastic equilibrium fluctuate around this path. This is identical to the concept of balanced growth considered by Rossi-Hansberg and Wright (2007). They show by example how this concept of balanced growth can be made consistent with empirical distributions of cities by population. Their example has i.i.d. idiosyncratic productivity shocks in a model with a locally produced, durable factor of production called human capital. We do not (need to) solve our model and so do not verify that this property is shared by our model.[4]

[4] We can modify our model to guarantee that cities exhibit population trends relative to the aggregate population in finite samples, without changing the models balanced growth properties. For instance, we could assume idiosyncratic productivity follows a random walk with random probability of location destruction and instantaneous replacement, as in Alvarez and Shimer (2008). Building in endogenous trends is more difficult, although there has been some progress in recent years. For example, the multi-sector model of Ngai and Pissarides (2007) might be adaptable to our purposes.

To derive the balanced growth path we use the market clearing conditions, the trend growth rates of productivity and land development costs, and the local and final good production functions. The market clearing condition for final goods, (14), implies per capita consumption, physical capital and the value of land development goods each grow at the same rate in units of the final good. Denote this common rate g_c. An implication of the final good production function, (3), is that output growth in each city corresponds to final good output growth. It follows that output in each city grows at the rate g_c as well. The aggregate physical capital clearing condition, (11) and the local capital market clearing condition conditions, (9) imply that physical capital used to produce intermediate goods in each city also grows at the rate g_c. The trend in land development costs, the land accumulation technology (4), and local developed land market clearing imply that per capita developed land used in intermediate good production grows at the rate g_c/τ in each city.

Combining these growth implications and the local production function we can derive the equilibrium value of g_c as follows. After substituting for a in equation (1), dividing through by n yields an expression involving per capita variables:

$$\frac{y}{n} = z^{(1-\alpha)\phi} \left[\frac{y/n}{l/n}\right]^{\frac{\lambda-1}{\lambda}} \left[\frac{l}{n}\right]^{1-\phi} \left[\frac{k}{n}\right]^{\alpha\phi}.$$

Divide this equation by the equivalent one for the previous period, and make appropriate substitutions for the growth rates of the variables. Solving for g_c in the resulting expression yields that along a balanced growth path, per capita consumption growth is given by

$$g_c = \gamma\tau^{\frac{\delta-1}{(1-\alpha)\delta}}. \tag{18}$$

where $\delta = \phi\lambda$. The parameter δ measures the effect of agglomeration net of congestion due to diminishing returns to land.

Equation (18) shows that per capita consumption growth depends on the rate of productivity growth and, if $\delta \neq 1$, the rate of growth in land development costs as well. According to equation (16) the growth rate of the relative price of new land is identical to that of land development goods along a balanced growth path. If this growth rate is positive, that is $\tau > 1$, then developed land growth will not keep up with output growth. In this case the density of economic activity grows. If agglomeration effects outweigh congestion effects, $\delta > 1$, then this mechanism provides an endogenous source of per capita consumption growth in addition to productivity growth. Without any productivity effect of density, $\lambda = 1$ and $\delta = \phi$, equation (18) implies

$$g_c = \gamma \tau^{\frac{\phi - 1}{(1-\alpha)\phi}}.$$

For $0 < \phi < 1$, per capita consumption growth in this case is lower than the rate of growth of exogenous productivity, because of decreasing returns to developed land in intermediate good production. Finally, if agglomeration and congestion effects cancel out, that is $\delta = 1$, then per capita consumption growth equals the rate of productivity growth, as in the simple neoclassical growth model.

The balanced growth path arises for similar reasons in our model and in the model of Rossi-Hansberg and Wright (2007). In our model land is used for housing and production. In Rossi-Hansberg and Wright (2007) land is not used in production, but each worker consumes a unit amount of land. Rossi-Hansberg and Wright (2007) have diminishing returns to expanding a city due to higher commuting costs. Land is increased when a new city is created and when a given city is made larger. Growth arises through a combination (depending on parameter values) of city size growth and growth in the number of cities. A consequence of this is that growth occurs by adding land. In our model the number of cities is fixed and growth occurs as developed land is added to these cities.

We use equation (18) below to quantify the contribution of agglomeration to per capita consumption growth. For this we need an estimate of δ. We now show how we can use our model to estimate δ with city-level panel data on wages, land rents, and output prices.

2.4 Identifying δ

As in Ciccone and Hall (1996), we use production efficiency to derive an equation which can be used to estimate δ. Consider any two cities, indexed i and j, in some arbitrary period. Solving for output in (1) and dividing output in city i by that in city j yields

$$\frac{y_i}{y_j} = \left[\frac{z_i}{z_j}\right]^{\delta(1-\alpha)} \left[\frac{l_i}{l_j}\right]^{1-\delta} \left[\frac{k_i}{k_j}\right]^{\delta\alpha} \left[\frac{n_i}{n_j}\right]^{\delta(1-\alpha)}.$$

Using (6) to eliminate the capital terms in this last equation we have

$$\frac{y_i/l_i}{y_j/l_j} = \left[\frac{z_i}{z_j}\right]^{\frac{\delta(1-\alpha)}{1-\delta\alpha}} \left[\frac{n_i/l_i}{n_j/l_j}\right]^{\frac{\delta(1-\alpha)}{1-\delta\alpha}} \left[\frac{q_i}{q_j}\right]^{\frac{\delta\alpha}{1-\delta\alpha}}. \tag{19}$$

Equation (19) is closely related to equation (19) in Ciccone and Hall (1996), which is the basis for their estimation equation. It differs in two key respects. First, because Ciccone and Hall (1996) assume intermediate goods are perfect substitutes in producing final goods, $\eta = 1$, their equation does not contain any output prices. So if $\eta < 1$ their equation is subject to an omitted variable bias. Second, they assume all land in a given location is used in production while we have competing uses for land. Since the allocation of land between residential and non-residential uses in U.S. cities is unavailable, we cannot use (19) as a basis for estimation. Further complicating our analysis, annual data on real city output that is produced by the BEA is available for too short a time period (2001-2006) to be useful to us.

We eliminate employment, developed land and output from (19) using (5) and (7). Substituting for developed land rent using (8) yields the equation underlying our estimation of δ :

$$\frac{w_i}{w_j} = \frac{z_i}{z_j}\left[\frac{p_i}{p_j}\right]^{\frac{1}{1-\omega}\frac{\delta-1}{\delta(1-\alpha)}} \left[\frac{q_i}{q_j}\right]^{\frac{1}{\delta(1-\alpha)}}. \tag{20}$$

This equation embeds an idea expressed in Lucas (2001), which is that the size of agglomeration effects can be identified using rental prices. In our case, we use housing rents rather than land rents (since the two are proportional) and we also account for variation across cities in wages due to variation in multi-factor productivity or the price of output. Notice that the exponent on housing rents is proportional to the exponent on the growth rate of land prices in the balanced growth formula, equation (18). The factor of proportionality involves the share of capital in housing, ω. So, with an estimate of ω in hand, we could, in principle, quantify the growth implications of agglomeration without separately estimating δ.[5] If we were to include in the model other local factor inputs with competing uses, such as energy, their prices would also appear in (20) in the same way that land prices do.

[5]That is, all we need to know to estimate the growth implications of agglomeration is $(\delta - 1)/(\delta(1-\alpha))$.

2.5 Extension to Heterogeneous Workers

So we have assumed that workers are homogeneous, which is a questionable assumption from an empirical perspective. Locations with a high density of economic activity may have high wages because of a concentration of high human capital workers, not because of a density effect on productivity. Therefore, not accounting for cross-sectional variation in the distribution of human capital could lead to overstating density effects. To address this issue we follow Ciccone and Peri (2006). They consider high skill and low skill workers as imperfect substitutes in producing total labor services. We now derive a version of the equation underlying our estimation which incorporates heterogeneous workers. Incorporating heterogeneous workers does not affect our balanced growth result.

Suppose that the effective labor input in (1) is a constant elasticity of substitution composite of unskilled, n_u, and skilled labor, n_e:

$$n = \left[\sigma n_u^{\xi} + (1-\sigma) n_e^{\xi}\right]^{1/\xi},$$

where $0 < \sigma < 1$ and $\xi \leq 1$. Composite labor satisfies (7) as before. Let w_u and w_e denote the wages of unskilled and skilled workers. The first order conditions for intermediate good producers' choices of unskilled and skilled labor can be used to express the wages of skilled workers as

$$w_e = (1-\sigma)\sigma^{1/\xi - 1}(1-\alpha)\phi w \chi^{1/\xi-1} m^{\xi-1},$$

where w denotes the implicit wage for the composite labor input, n, $\chi \equiv (w_u n_u + w_e n_e)/(w_u n_u)$, and $m \equiv n_e/n_u$. Substituting for composite wages using (20), we have

$$\frac{w_{ei}}{w_{ej}} = \frac{z_i}{z_j}\left[\frac{p_i}{p_j}\right]^{\frac{1}{1-\omega}\frac{\delta-1}{\delta(1-\alpha)}}\left[\frac{q_i}{q_j}\right]^{\frac{1}{\delta(1-\alpha)}}\left[\frac{\chi_i}{\chi_j}\right]^{1/\xi-1}\left[\frac{m_i}{m_j}\right]^{\xi-1}, \tag{21}$$

for any two cities i and j. The equation (21) reduces to equation (20) if $\xi = 1$, that is if unskilled and skilled labor are perfect substitutes. The variables χ and m are relatively straightforward to measure so this equation can be used to estimate δ taking into account worker heterogeneity.

3 Competitive Equilibrium Growth

This section completes the description of the model. For simplicity, we initially assume zero growth in the aggregate level of productivity, the cost of land development goods, and the number of households and focus on a stationary equilibrium. After describing the competitive equilibrium for this case, we briefly return to the case with growth to verify that the balanced growth path we have described does in fact correspond to a competitive equilibrium. Readers only interested in our empirical findings can skip this section without any loss of continuity.

3.1 The Household's Environment

There is a representative household composed of a unit measure of identical members.[6] The large household structure allows for full risk sharing within each household, a standard device in macroeconomics for studying complete markets allocations. The household owns physical capital as well as producers of intermediate and final goods, and land developers. Each household member supplies one unit of labor inelastically and enjoys utility from consumption and housing. The household derives income from labor and physical capital by allocating its members and its stock of capital across the cities. It also collects dividends from its ownership of firms. The only technology for allocating household members across cities is a random number generator. Each period the household's members are randomly allocated across cities at zero cost, after observing the current level of productivity in each city. A key assumption of the model is that household members must consume goods and housing in the same city that they work.

The household takes as given the same prices as the producers in each city. The household also takes as given the law of motion for the probability distribution of developed land and productivity. The law of motion for this distribution is

$$\mu_{t+1}(S', Z') = \int_{\{(s,z):s_{t+1}(s,z)\in S'\}} Q(z, Z')\mu_t(ds, dz), \tag{22}$$

for all S' and Z'. This equation states that the total number of cities at date $t+1$ with

[6] In the extension to heterogeneous workers we assume that fixed fractions of household members are either skilled or unskilled.

developed land in the set S' and productivity in the set Z' is given by the total of all cities that transit from their current productivity to productivity in Z' and produce developed land so that $s' = s_{t+1}(s, z)$ is in S'.

3.2 The Household's Optimization Problem

The household maximizes the expected present value of average household member utility, where household member utility in any given period is logarithmically separable in consumption and housing. Separability implies the household perfectly insures itself against consumption risk. Therefore every household member receives the same level of consumption, C, which corresponds to aggregate consumption when we normalize the number of household members and households to unity. Let $V(K, \mu)$ denote the expected utility of the household with capital K and distribution of land and productivity μ. The recursive representation of the household's problem is:

$$V(K, \mu) = \tag{23}$$
$$\max_{\left\{\substack{C,K',h(s,z)\\ n(s,z)}\right\}} \left\{ \int n(s,z) \left[\ln C + \psi \ln h(s,z)\right] \mu(ds, dz) + \beta \int V(K', \mu')\mu'(ds', dz') \right\}$$

subject to

$$\begin{aligned} C + K' - (1-\kappa)K \quad &+ \quad \int p(s,z)n(s,z)h(s,z)\mu(ds,dz) \\ &\leq \quad rK + \int w(s,z)n(s,z)\mu(ds,dz) + \pi, \end{aligned} \tag{24}$$

the employment constraint (12) and the transition equation (22), where $\psi \geq 0$.[7] The left-hand-side of the household's budget constraint includes the household's expenditures on consumption, new physical capital and housing. The right-hand-side includes income from physical capital, labor, and dividends from ownership of intermediate good producers, final good producers and land developers, π.

[7]To conserve on notation we have not distinguished between the demand and supply of final goods or factors of production.

The first order conditions for K', $h(s,z)$ and $n(s,z)$ are

$$\begin{aligned} 1 &= \beta \frac{C}{C'}\left[r' + 1 - \kappa\right]; \qquad (25) \\ \frac{\psi}{h(s,z)} &= \frac{p(s,z)}{C}, \ \forall (s,z); \qquad (26) \\ \ln C + \psi \ln h(s,z) &= \theta + \frac{1}{C}\left[p(s,z)h(s,z) - w(s,z)\right], \ \forall (s,z), \qquad (27) \end{aligned}$$

where θ is the Lagrange multiplier for (12). The physical capital accumulation condition (25) is familiar from the simple neoclassical growth model. The simple form of equation (26) is due to our specification of preferences. Its implication of a constant housing share of consumption is supported by evidence in Davis and Ortalo-Magné (2008). Equation (27) is unique to our model. Combined, equations (26) and (27) imply that utility is not identical for each household member but rather household members assigned to relatively high wage cities enjoy a relatively low quantity of housing and thus relatively low utility. This feature of our model contrasts with the static Roback (1982) model where *ex post* utility of workers is equated across cities. In our dynamic model with complete markets *ex ante* expected utility is equalized.

The result that the household assigns lower *ex post* utility to members in more productive locations may be unfamiliar to readers who work with the Roback (1982) model. It is not an implication of our production assumptions, our specification of log-separable preferences, or to any model dynamics, but is a general implication of our market structure. The intuition for the result is easiest to understand if we assume labor is the only factor of production and there is no density externality so that the production function is zn. In addition assume that household members' preferences for consumption and housing are separable and linear in housing. Consider two otherwise identical cities where productivity in city 2 is greater than that in city 1: $z_2 > z_1$. To make utility identical across agents the household would allocate half its members to city 1 and half to city 2. This is not efficient. To see this, notice that the household could move a small fraction of members from city 1 to city 2 and increase total output so that consumption for everyone would increase. Although residents of the high productivity city 2 would have less land for housing, residents of the low productivity city 1 would have *more* land for housing. With utility linear in housing there would be no net impact on average utility from housing from this movement of agents. Therefore, overall average utility rises and the household's objective is increased.

We conclude this subsection by noting that equation (27) reveals the important role played by housing in our model. Consider the case where housing is not valued by household members, $\psi = 0$. Equation (27) is satisfied in this case only if wages are equalized across locations. Since our estimation procedure relies on cross-sectional variation in wages, we require that housing is valued, $\psi > 0$.

3.3 Stationary Competitive Equilibrium

A *stationary competitive equilibrium* consists of value functions $V(K, \mu)$ and $W(s, z)$, decision rules for consumption $g_C(K, \mu)$, physical capital accumulation $g_K(K, \mu)$, land development $s'(s, z)$, housing $h(s, z)$, land in production $l(s, z)$, physical capital in production $k(s, z)$, and workers $n(s, z)$, aggregate quantities $\{C, K'\}$, and a measure $\mu(s, z)$ such that

1. Given prices, households maximize expected utility, so that $V(K, \mu)$ solves the household problem as given by (23) and equations (24)–(27) are satisfied;

2. Given prices, producers maximize profits, so that $W(s, z)$ solves (15) and equations (3) – (7), (13), (16) and (17) are satisfied.;

3. The discount factor $1/R$ is given by β;

4. Aggregates are consistent with individual behavior, so that $\mu(s, z)$ is generated by (22).

5. Markets clear so that equations (9) – (12) and (14) are satisfied.

When $\lambda = 1$ the competitive equilibrium corresponds to the solution of a concave programming problem. Standard arguments can be used to prove in this case that there exists a unique competitive equilibrium and the allocations are Pareto-efficient.

We have not found conditions under which an equilibrium exists with density effects, $\lambda > 1$. However, using the argument in Kehoe et al. (1992) we know that for λ sufficiently close to unity, there exists a unique equilibrium, because the model without density effects has a unique equilibrium. Equilibrium existence and uniqueness with the size of density effects we estimate remain open questions. We do not need to find a solution of the model to estimate the impact of density on local productivity or quantify the effect of agglomeration on per capita consumption growth, and so do not do so in this paper.

So far in this section we have assumed no growth. This is without loss of generality because the preceding discussion applies to a version of the model specified in terms of variables scaled so that they are stationary along a balanced growth path: the balanced growth path corresponds to a competitive equilibrium. This is established by verifying that the conditions stated in the definition of the competitive equilibrium are compatible with the balanced growth path. Since we derived this path using the market clearing conditions, these conditions are obviously compatible. It remains to establish the compatibility of the first order conditions of the firms' and household problems. This is accomplished by verifying that if the first order conditions are satisfied by the scaled variables, they are also satisfied by the unscaled variables.[8]

4 Empirical Strategy

This section describes how we estimate the net agglomeration effect coefficient, δ. We begin by deriving the basic form of the estimation equation we use. After this we describe how our estimation procedure addresses endogeneity and fixed effects.

[8] This is straightforward for all but the first order condition for allocating household members to cities, equation (27). Since C and h appear in both levels and logs in this equation, it must be handled differently. Notice that this equation is the only one in which the Lagrange multiplier on the allocation of household members across cities, θ, appears. In other words, equation (27) determines the equilibrium value of θ, assuming an equilibrium exists. Equation (27) implies that $\theta > 0$ at all dates along a balanced growth path, as long as labor income always exceeds the value of housing services in each city. This condition will be satisfied in any plausible calibration. We conclude that equation (27) is compatible with balanced growth for plausible calibrations.

4.1 The Estimation Equation

Take logs of equation (21) and rearrange terms:

$$\begin{aligned}\ln(w_{ei}) - \ln(w_{ej}) \quad = \quad & \frac{1}{1-\omega}\frac{\delta-1}{\delta(1-\alpha)}\left[\ln(p_i) - \ln(p_j)\right] + \frac{1}{\delta(1-\alpha)}\left[\ln(q_i) - \ln(q_j)\right] \\ & + \frac{1-\xi}{\xi}\left[\ln(\chi_i) - \ln(\chi_j)\right] + (\xi - 1)\left[\ln(m_i) - \ln(m_j)\right] \\ & + \ln(z_i) - \ln(z_j)\end{aligned}$$

This equation holds for any two cities i and j at all dates. Therefore, it must also hold for city i and the average of a sample of size N of cities. Define a "hat" over a variable for city i to mean the difference of the log of that variable and the average of the same logged variable for all N cities in the sample, for example

$$\hat{w}_{ei} \equiv \ln(w_{ei}) - \frac{1}{N}\sum_{j=1}^{N}\ln(w_{ej}).$$

Notice that by construction hatted variables are mean zero in every year. Reintroducing the time subscript, it follows that

$$\hat{w}_{eit} = \frac{1}{1-\omega}\frac{\delta-1}{\delta(1-\alpha)}\hat{p}_{it} + \frac{1}{\delta(1-\alpha)}\hat{q}_{it} + \frac{1-\xi}{\xi}\hat{\chi}_{it} + (\xi-1)\,\hat{m}_{it-1} + \hat{z}_{it}, \tag{28}$$

which also holds at all dates. This is our baseline estimation equation. Because all variables are expressed as deviations from year averages, our identification of the model's parameters is based on changes in the cross-sectional distribution of our panel that occur over time.

We have data on all variables in this equation except for $\hat{z}_{it}$. If we were unconcerned about the correlation of $\hat{z}_{it}$ with the other right-hand-side variables, and were not interested in the magnitudes of the structural parameters, then we could obtain a consistent estimate of the coefficient on housing rents by ordinary least squares (OLS). With an estimate of ω in hand we could then quantify the impact of agglomeration on growth. However, the model predicts that $\hat{z}_{it}$ and the other right-hand-side variables are correlated and our theory imposes restrictions on the coefficients that should be satisfied for our findings to be credible. So, to estimate the structural parameters we need to use a non-linear instrumental variables approach. If we had only one wave of data, we would have to argue for an instrument that is correlated with the observable variables and uncorrelated with the $\hat{z}_{it}$ term and then we could apply non-linear least squares to obtain consistent estimates. This is the empirical approach adopted by Ciccone and Hall (1996). Since we have panel data with a long time dimension, we use a different estimation strategy.

4.2 Addressing Endogeneity

We address endogeneity by making a parametric assumption on the evolution of city-specific productivity. The stationary component in (2) is assumed to be the sum of two random variables, one which is a city-specific first-order auto-regressive process, and one which is common to all cities, but we do not specify explicitly. We include this second term to make it clear that our empirical procedure is consistent with the presence of an aggregate productivity shock. For every city i and all $t \geq 1$, $\ln(z_{it})$ evolves according to:

$$\begin{aligned} \ln(z_{it}) &= \ln(z_{i0}) + \gamma t + \ln(\tilde{z}_{it}) + u_t; \\ \ln(\tilde{z}_{it}) &= \rho \ln(\tilde{z}_{it-1}) + e_{it}, \end{aligned}$$

where γ is the deterministic rate of growth from (2), u_t is the economy-wide stochastic component of productivity, $-1 < \rho < 1$ is the autocorrelation coefficient for city-specific productivity, and e_{it} is a city-specific i.i.d. shock that is orthogonal to all variables dated $t-1$ and earlier. Both γ and ρ are common to all cities. The variable z_{i0} is the initial level of productivity specific to city i.

We now use our assumptions on the evolution of $\tilde{z}_{it}$ to manipulate (28) into an equation with an error term that is uncorrelated with all variables from previous years. This forms the basis of our instrumental variables estimation. Notice that our assumptions on the evolution of z_{it} and $\tilde{z}_{it}$ imply that the growth term and the economy-wide stochastic term drop out of the deviation of $\ln(z_{it})$ from its cross-section average, $\hat{z}_{it}$. Furthermore

$$\hat{z}_{it} - \rho \hat{z}_{it-1} = (1-\rho)\,\hat{z}_{i0} + \epsilon_{it}, \tag{29}$$

where $\epsilon_{it} \equiv [e_{it} - \bar{e}_t]$ with $\bar{e}_t$ equal to the period t average of city-specific productivity shocks for the N cities in our sample. Now, from each variable in equation (28) subtract ρ times its once-lagged value. This is a valid operation since equation (28) holds at all dates. Then, using equation (29), we have

$$\begin{aligned}\hat{w}_{eit} &= (1-\rho)\hat{z}_{i0} + \rho\hat{w}_{eit-1} + \frac{1}{1-\omega}\frac{\delta-1}{\delta(1-\alpha)}\left[\hat{p}_{it} - \rho\hat{p}_{it-1}\right] + \frac{1}{\delta(1-\alpha)}\left[\hat{q}_{it} - \rho\hat{q}_{it-1}\right] \\ &\quad + \frac{1-\xi}{\xi}\left[\hat{\chi}_{it} - \rho\hat{\chi}_{it-1}\right] + (\xi-1)\left[\hat{m}_{it} - \rho\hat{m}_{it-1}\right] + \epsilon_{it}. \qquad (30)\end{aligned}$$

To estimate equation (30) we need to find a set of instruments that are correlated with the wage, price, and skill variables and are uncorrelated with ϵ_{it}. Finding valid instruments is straightforward because ϵ_{it} is an i.i.d. shock; any variable dated $t-1$ or earlier is potentially a valid instrument. In addition, we need to address the unobserved variable $\hat{z}_{i0}$. This variable is a "fixed effect" in the sense that it varies across cities, but is fixed over time in each city.

4.3 Addressing Fixed Effects

Several strategies have been proposed which can in principle address the city fixed effects. The most common one involves eliminating the fixed effect by taking first differences. A second strategy is to assume that the level of the fixed effect, $\hat{z}_{i0}$, is uncorrelated with the first-differences of all model variables. The application of this method to our case would involve using equation (30) with lagged growth rates of model variables as instruments. A third widely used strategy combines the first two strategies. This approach was originally proposed by Blundell and Bond (2000). We are uncomfortable assuming that the fixed effect is uncorrelated with lagged changes of model variables, since we have no theory suggesting this to be the case. Therefore we do not use the Blundell and Bond (2000) estimation procedure.

When more waves of data are available than are the minimum necessary to implement the first-difference strategy, taking first-differences does not use all the available information on the fixed effect and so is inefficient. The procedure proposed by Arellano and Bover (1995) takes advantage of the additional information on the fixed effects which comes with long panel data and so this is what we use. In this approach, each time-t variable in equation (30) is expressed as a deviation from the average of all future observations for city i in the sample. For any variable x_t with observations $t = 1, \ldots, T$, define the Arellano-Bover difference operator as of date t,

Δ_t, as follows:

$$\Delta_t x_t = \varphi_t \left[x_t - \frac{1}{T-t} \sum_{s>t} x_s \right].$$

where

$$\varphi_t = \left(\frac{T-t}{T-t+1} \right)^{1/2}.$$

Applying the Arellano-Bover difference operator to (30) yields

$$\begin{aligned} \Delta_t \hat{w}_{eit} = {} & \rho \Delta_t \hat{w}_{eit-1} + \frac{1}{1-\omega} \frac{\delta - 1}{\delta(1-\alpha)} \left[\Delta_t \hat{p}_{it} - \rho \Delta_t \hat{p}_{it-1} \right] \\ & + \frac{1}{\delta(1-\alpha)} \left[\Delta_t \hat{q}_{it} - \rho \Delta_t \hat{q}_{it-1} \right] + \frac{1-\xi}{\xi} \left[\Delta_t \hat{\chi}_{it} - \rho \Delta_t \hat{\chi}_{it-1} \right] \\ & + (\xi - 1) \left[\Delta_t m_{it} - \rho \Delta_t \hat{m}_{it-1} \right] + \Delta_t \epsilon_{it}. \end{aligned} \tag{31}$$

Including the weight term φ_t in Δ_t guarantees that the error term $\Delta_t \epsilon_{it}$ in (31) has a constant variance.

We use generalized method of moments to estimate equation (31) using the levels of $\hat{w}_{eit-s}$, $\hat{p}_{it-s}$, $\hat{q}_{it-s}$, $\hat{\chi}_{it-s}$, and $\hat{m}_{it-s}$ for $s \in \{3, 4\}$ as instruments. Lags $s \geq 5$ are not used as instruments because typically they do not add much information. We do not use lags $s \in \{1, 2\}$ to accommodate classical measurement error in all the variables. Monte Carlo studies by Blundell and Bond (2000) and Windmeijer (2005) suggest that procedures that use lagged levels of variables as instruments on differenced data may yield biased estimates of parameters and standard errors. To understand if these results apply to our data, we conduct a Monte Carlo study based on our data. When appropriate, we discuss the implications of this study when describing our empirical results. The details and results of the Monte Carlo study are outlined in the Appendix.

5 The Data

In our empirical analysis we exploit the extension of our model to include heterogeneous workers. Therefore, to implement our empirical strategy we need city-level data on housing rents, output prices, and wages and hours of low and high skill workers. Since all our econometric specifications are based on within-year deviations from averages, we do not need to adjust any of our variables for overall price inflation. That is, we can use nominal wage, housing rent, and output price variables directly in our analysis. A city in our data is a metropolitan statistical area (MSA) as defined by the Census Bureau. This section summarizes our MSA-level data and conducts some preliminary analysis of it. A more detailed description of our data is contained in the Appendix.

5.1 Housing, Labor and Output Variables

We construct annual data on housing rents by combining micro-level data from the 1990 Decennial Census of Housing (DCH) with housing rental price indexes from the Bureau of Labor Statistics (BLS). Specifically, we use the data from the 1990 DCH to estimate the level of housing rents by MSA in 1990, and then use MSA-specific rental price indexes from the Bureau of Labor Statistics (BLS) to extrapolate the rental price level backwards to 1985 and forwards to 2006.

We construct our wage and employment variables from the March Current Population Survey (CPS). High skilled workers are identified as those workers with at least four years of college. Low skill workers are identified as workers with less than four years of college. We measure wages for a given type of worker as total wages divided by total hours for that type of worker.

We use data from the Bureau of Economic Analysis (BEA) to construct price indexes for output by MSA. These indexes are created as weighted averages of industry-specific price indexes where the weights are based on the share of annual income earned by each industry in a MSA. Because the mix of industries varies by MSA, and price indexes vary by industry, our price index for output also varies by MSAs. We normalize the price index for MSA-level output to 1.0 in 1969 in every MSA. This arbitrary normalization introduces another MSA-level fixed effect that is differenced-out by the Arellano and Bover (1995) procedure.

We drop MSAs with incomplete or missing wage, employment, output price or house price data. This leaves us with data covering the 1985-2006 period for 22

MSAs. The 22 MSAs are, roughly speaking, the 22 largest MSAs by population in the US. These MSAs include 36 percent of the total US population.

5.2 Summary Statistics

Table 1 shows standard deviations and correlations of our data. By construction, all variables have zero mean in each year and thus are zero mean unconditionally. The reported standard deviations and correlations are estimated from the complete panel; they are not averages of annual statistics.[9] There are two features of this table worth noting.

First, the standard deviation of housing rents, $\hat{p}_{it}$, is larger than the standard deviation of high-skill wages $\hat{w}_{eit}$. This is predicted by our model. To see why, notice that equations (26) and (27) imply that that the percentage deviation of housing rents between any two metro areas i and j approximately satisfies:

$$\frac{p_{it} - p_{jt}}{p_{jt}} = \left(\frac{w_{jt}}{C_t}\right)\left(\frac{1}{\psi}\right)\left(\frac{w_{it} - w_{jt}}{w_{jt}}\right) ,$$

where we approximate $\ln(p_{it}/p_{jt})$ as $(p_{it} - p_{jt})/p_{jt}$. As long as $w_{jt}/c_t > \psi$ then the percentage deviation of rents is greater than the percentage deviation of wages. Davis and Ortalo-Magné (2008) report that the average expenditure share on housing in the United States is 24 percent. According to (26), this suggests $\psi = 0.24$. Since wages everywhere in our sample are greater than 24 percent of the level of consumption, housing rents must be more dispersed than wages paid to labor. In our extension of the model to heterogeneous workers, the same argument can be used to show that housing rents must be more dispersed than high-skill wages as long the ratio of high skill wages to consumption is greater than the household expenditure share on housing.

Second, high-skill wages, $\hat{w}_{eit}$ are positively correlated with housing rents, p_{it}, and the ratio of high-skill to low-skill workers, m_{it}. It is straightforward to show using

[9]The standard deviation of the log deviation of output prices, $\hat{q}_{it}$, cannot be interpreted literally, as output prices in every MSA are arbitrarily normalized to equal 1.0 in 1969. We could have increased the standard deviation of output prices simply by setting the normalized price level to vary across MSAs.

the first order and market clearing conditions that these features of the data are also predicted by the model.

5.3 Preliminary Regression Analysis

Table 2 shows the coefficient estimates and robust standard errors of simple OLS regressions of wages of high-skill workers on housing rents, $\hat{p}_{it}$, output prices, $\hat{q}_{it}$, and the two variables related to the skill composition of the workforce, $\hat{\chi}_{it}$ and $\hat{m}_{it}$. The first column corresponds to a regression in levels including MSA-specific fixed effects; and the second column corresponds to a regression in differences without fixed effects.

From these regressions, we observe three features of the data. First, the model variables account for a significant fraction of the variation in high skill wages. In the regressions in levels the R^2 statistic is 0.79, compared to 0.62 when the regression only includes fixed effects. In the difference regression, where fixed effects have been removed, the R^2 statistic is 0.37.

Second, the coefficient on housing rents is estimated to be positive in both regressions. In the levels regression it is also statistically significant. Since the coefficient on housing rents is proportional to the exponent determining whether there is a role for agglomeration in aggregate growth, these findings suggest we may find a significant role when we examine our structural estimates.

Third, the coefficients on output prices, $\hat{q}_{it}$, are imprecisely estimated and vary substantially across the two regressions. This suggests output prices (as we measure them) do not account for much of the variation in high-skill wages after controlling for skill composition and housing rents. However, we show below that including output prices has a noticeable impact on our estimates of the labor substitutability and serial correlation parameters, ξ and ρ.

6 Estimates of The Model's Parameters

This section describes our estimates of the model's parameters. To implement our GMM procedure, we first need to assign values to α, physical capital's share of production net of land and ω, capital's share of housing when we use housing rents in our estimation. In addition, to translate estimates of δ into estimates of the impact of agglomeration on per capita consumption growth we need a value for ϕ, the total share of capital and labor in production. The first subsection describes how we calibrate these parameters. We then describe our baseline estimates and after this we show how two key model assumptions influence these estimates.

6.1 Calibrated Parameters

The parameters α and ϕ are calibrated as follows. First, we estimate the non-labor share of aggregate income, which corresponds to $1 - \phi + \alpha\phi$ in our model. Next, we estimate the share of non-labor income attributable to land, which in our model is the ratio of $1 - \phi$ to $1 - \phi + \alpha\phi$. We use the estimates of these two shares to solve for the two unknowns, α and ϕ. In these calculations we measure aggregate income twice – first including the service flow from the stock of consumer durables (a common assumption in macroeconomics) as part of aggregate income and then excluding the service flow – to arrive at two sets of estimates for α and ϕ.

The non-labor share of income is measured as follows. Following Cooley and Prescott (1995), we assume that labor's share of "ambiguous income" is the same as that for "unambiguous income." To calculate unambiguous income we subtract from aggregate income those components of income where the factor payments are ambiguous, including taxes less subsidies, proprietors' income and the surplus of government enterprises. We use the ratio of compensation of employees to unambiguous income to measure labor's income share and one minus labor's income share to measure the non-labor income share.[10] Based on data from 1959-2005 these calculations yield non-labor income shares of about .33 (including the service flow from durables as part of aggregate income) and .23.

We estimate the share of non-labor income going to land using data from the Flow of Funds Accounts summarizing the composition of the balance sheets of the nonfinancial corporate sector. Specifically, we identify land's share of non-labor income as the average of the ratio of the value of land used in real estate in the nonfinancial corporate sector to the value of all tangible assets (real estate, equipment and software,

[10]These numbers are obtained from NIPA Table 1.10. The ambiguous income components are lines 9, 10, 15 and 22 of this table and labor income is line 2.

and inventories).[11] Over the 1952-2009 time period for which data are available, we estimate the average of this ratio to be about 15 percent. This estimate of land's share of non-labor income and the two non-labor shares calculated above yield the two sets of estimates: $(\alpha, \phi) = (0.3, 0.95)$ and $(\alpha, \phi) = (0.2, 0.97)$.

To calibrate ω, we use the fact that equation (8) implies for each city i

$$\hat{p}^l_{it} = \frac{1}{1-\omega}\hat{p}_{it}. \tag{32}$$

This equation suggests that if we had data on land rents by MSA then we could estimate ω using the ratio of standard deviations of land rents and housing rents. Data on land rents are unavailable, but data on land *prices* are available. We can use these data to estimate an upper bound on ω. This is useful since lower values of ω yield larger estimates of δ. To arrive at this upper bound notice that equation (16) implies that land prices and land rents are related as

$$\hat{b}_{it} = \hat{p}^l_{it} + \hat{\vartheta}_{it},$$

where $\hat{\vartheta}_{it}$ is stochastic with an unknown distribution. The variable $\hat{\vartheta}_{it}$ will most likely be positively correlated with land rents due to the adjustment costs in developed land. For example, in times when the rental price of land is high, land developers will start developing land, which will raise adjustment costs and drive up the price of land. Using (32), it follows that the upper bound for ω equals one minus the ratio of the standard deviation of housing rents to the standard deviation of land prices. The land price data described in Davis and Palumbo (2008) for the MSAs and years which overlap with our housing rent data (21 MSAs for the years 1985-2004) and our housing rent data combined imply that this upper bound is 0.826. This value compares to the Census Bureau estimate of the share of new home value attributable to the value of land of 0.89 described by Davis and Heathcote (2005). Given the uncertainty involved, we consider $\omega \in \{0.81, 0.85, 0.89\}$.

[11] These numbers are obtained from Table B.102. The value of land used in real estate is line 3 (the market value of real estate) less line 34 (replacement cost value of structures used for nonresidential purposes). The value of all tangible assets held by the nonfinancial corporate sector is line 2.

6.2 GMM Estimates

Table 3 reports parameter estimates, asymptotic standard errors, and specification tests when we estimate our model with GMM using the procedure of Arellano and Bover (1995). These estimates are based on $\alpha = 0.3$. When $\alpha = 0.2$ there are only small differences in the estimates and test statistics (our estimates of the impact of agglomeration on growth are sensitive to the value of α.) We regard $\omega = 0.85$ as the most plausible value for this parameter. The estimates for this baseline case are in the first column. The other two columns report how our estimates change when $\omega = 0.81$ and $\omega = 0.89$.

We employ three different specification tests, summarized in the bottom rows of Table 3. The first is the J-test of the over-identifying restrictions due to Hansen (1982) and Sargan (1958). The p-values corresponding to these test statistics indicate that the model is never rejected by the J-test. Arellano and Bond (1991) argue that the power of the J-test to detect model miss-specification can be quite low. As a more powerful alternative they suggest a test of the serial correlation of the residuals. If the model is correctly specified, the residuals of our estimating equation should only exhibit autocorrelation up to order one. For each case we report the $m2$ test of second order serial correlation from Arellano and Bond (1991), and the analogous test for third-order serial correlation, denoted $m3$. The p-values corresponding to these test statistics indicate that the null hypotheses that the residuals do not display second or third-order serial correlation are never rejected.

All three point estimates of δ exceed unity and our baseline estimate of δ is 1.020. Our Monte Carlo study, described in the Appendix, suggests this estimate is essentially unbiased. Recall that δ is the product of ϕ, capital and labor's share of production, and λ, the density parameter. The parameter δ is thus the impact of the density externality on production, net of the fact that land is also an input in production. Our estimate $\delta > 1$ implies that a reduction in the amount of land used in production, on-net, increases output; the positive impact from the change in output density more than compensates for the reduction of one of the inputs in production. A value of δ greater than unity also implies that agglomeration increases aggregate growth beyond the exogenous growth productivity. The asymptotic standard error of the baseline estimate δ is 0.002. Our Monte Carlo study suggests that the reported standard error on δ may be too low by a factor of about 2.4, such that the "true" standard error of our baseline estimate is closer to 0.005. If this is the case, we easily reject that $\delta = 1$. Similar calculations imply that our other two estimates of δ are also significantly greater than unity.

With an estimate of ϕ, we can determine the direct impact of density on productivity as $\lambda = \delta/\phi$. Given $\phi = 0.95$ or $\phi = 0.97$ we easily reject the hypothesis that $\delta = \phi$ and so that agglomeration increases productivity in a statistically significant way. Our baseline estimate of δ and $\phi = 0.95$ imply $\lambda = 1.074$. A firm located in an area with twice the output density will produce 6.9 percent ($= (\lambda - 1)/\lambda$) more output than an otherwise identical firm, holding all other inputs to production constant.

We estimate ξ and ρ in all cases to be about 0.56 and 0.57. Our Monte Carlo study suggests that if the true value of ρ is 0.6, then both of these estimates are essentially unbiased. This study also suggests that our reported standard errors are too low by a factors of 2 for ξ and 1.4 for ρ.

Our estimate of $\xi = 0.56$ implies an elasticity of substitution of 2.3 between low- and high- skilled labor types. This is slightly higher than the range of recent estimates of the elasticity of substitution between college and non-college educated workers of 1.3 to 1.7 as reported by Autor et al. (1998).[12] However, as Autor et al. (1998) note substantial uncertainty exists concerning the magnitude of the elasticity of substitution. For example, the results of Katz and Murphy (1992) suggest that a two standard error interval around the elasticity of substitution is $[1.01, 2.44]$.

Our estimate of ρ is low when compared to macroeconomic studies which typically assume first-order serial correlation of aggregate productivity between .95 and 1. Recall that $\hat{z}_{it}$ is defined to be the log-deviation of an MSA's exogenous level of productivity from its average. This implies our estimate of ρ is only informative about the persistence of the MSA-specific component of exogenous productivity. Our finding of $\rho = 0.57$ implies that MSA-specific shocks to productivity are less persistent than shocks to aggregate productivity.

[12] See also Katz and Murphy (1992), Heckman et al. (1998) and Krusell et al. (2000).

6.3 Effects of Key Modeling Assumptions

We now investigate how two key model ingredients have affected our estimates of the model's parameters. We consider how our estimates are affected by assuming (a) intermediate goods are perfect substitutes in the production of the final good, and (b) land used in production in an MSA is fixed over time (but can vary across metro areas). These are two assumptions imposed by Ciccone and Hall (1996) in their empirical work that we relax in our study. In all cases, we estimate the parameters by GMM using the procedure of Arellano and Bover (1995) with $\alpha = 0.30$ and $\omega = 0.85$.

Table 4 displays for estimation results for our baseline specification (reproduced from Table 3) and estimates based on two variations on this specification. Recall that our baseline specification is based on equation (28). To derive (28) we assumed that land is a variable input in production and intermediate goods are imperfect substitutes in production, $\eta < 1$. Suppose we had assumed that intermediate goods are perfect substitutes, $\eta = 1$, but that land is still a variable input into production. In this case the estimation equation is

$$\hat{w}_{eit} = \frac{1}{1-\omega}\frac{\delta-1}{\delta(1-\alpha)}\hat{p}_{it} + \frac{1-\xi}{\xi}\hat{\chi}_{it} + (\xi-1)\,\hat{m}_{it-1} + \hat{z}_{it}, \tag{33}$$

Notice this is identical to the baseline specification except that the output price is dropped. The parameter estimates and standard errors for this case are in the second row of Table 4. Our second variation on the estimation maintains the assumption $\eta = 1$ and adds to it the assumption that land is a fixed input in the production of intermediate goods. In this case the estimation equation is

$$\hat{w}_{eit} \;=\; \frac{1}{\xi}\left[1-\xi+\frac{\delta-1}{1-\delta\alpha}\right]\hat{\chi}_{it} + (\xi-1)\,\hat{m}_{it} + \left(\frac{\delta-1}{1-\delta\alpha}\right)\hat{u}_{it} + \check{z}_{it}\ , \tag{34}$$

where $\check{z}_i t$ is proportional to $\hat{z}_{it}$. In this case variation in high-skill wages is a function of labor market variables only and in particular neither housing rents nor output prices appear. Parameter estimates and asymptotic standard errors for this case are shown in the third row of Table 4.

The second row of Table 4 shows that omitting prices has essentially no impact on our estimate of δ, but raises ξ somewhat, and cuts ρ by more than half. The third row shows that assuming land is fixed in addition to omitting prices would lead to a much lower estimates of all three parameters. The estimate of δ in this latter case is particularly striking, since it indicates that the density of economic activity is irrelevant for productivity. Clearly our assumption that land is a variable input in production is important to our finding that density has a significant impact on productivity.

7 Implications for Balanced Growth

In this section, we quantify the impact of agglomeration on per capita consumption growth. For convenience, we reproduce the balanced growth equation (18), which describes the relationship of per-capita consumption growth γ_c, growth in exogenous productivity, γ and the growth in the relative price of new land development τ:

$$\gamma_c = \gamma \tau^{\frac{\delta-1}{(1-\alpha)\delta}}.$$

Our quantification strategy is as follows. Given estimates of γ_c, τ, δ and α we determine the value of γ implied by the balanced growth equation. We then use this value for γ to determine what the balanced growth equation says consumption growth would have been if agglomeration had no affect on productivity, $\lambda = 1$, so that $\delta = \phi$. Let γ_c^* denote this counterfactual value of γ_c and $\tilde{\gamma}_c$ denote the empirical estimate of per capita consumption growth. We measure the impact of agglomeration on per capita consumption growth as the per cent increase in per capita consumption due to agglomeration effects, $\Delta = 100 \times (\tilde{\gamma}_c - \gamma_c^*)/(\gamma_c^* - 1)$.

We already have all the ingredients to do this calculation except for estimates of γ_c and τ. We estimate consumption as the chain weighted aggregate of consumption of non-durables and services (exclusive of housing services) plus government consumption from the NIPA. In the cases corresponding to $\alpha = 0.3$, we also include an estimate of the service flow from consumer durables obtained from the Federal Reserve Board. Since our model focuses on the allocation of workers across cities, we use aggregate employment to put consumption in "per capita" terms. Our measure of aggregate employment is non-farm payroll employment from the Bureau of Labor Statistics (BLS). Over the sample period of our panel data, 1985-2006, we estimate: $\tilde{\gamma}_c = 1.013$ and $\tilde{\gamma}_c = 1.009$, corresponding to $\alpha = 0.3$ and $\alpha = 0.2$.

The parameter τ equals the rate of growth of land rental prices in our model. So one way to measure τ is by estimating average growth of real housing rents and then using equation (32) to rescale growth in housing rents to growth in land rents. To be consistent with how we measure housing rents in our panel dataset, we measure real housing rents as the CPI for shelter divided by the CPI for Urban Consumers on all items excluding shelter. According to this measure real housing rents have increased on average by 0.58 percent per year over the 1985-2006 period. Using equation (5) this translates to τ equal to 1.039, 1.031 and 1.053 corresponding to our three values of ω equal to 0.85, 0.81 and 0.89.

An alternative way to estimate τ is to use the model's implication for the growth rate of per capita developed land along a balanced growth path. The US Department of Agriculture (USDA) uses satellite imagery to estimate the total quantity of developed land in the US, defined as "large urban and built-up areas, small built-up areas, and rural transportation land." Data from the USDA's 2003 National Resources Inventory implies that developed land in the US has increased by 1.88 percent per year between 1982 and 2003. Using non-farm payroll employment this growth translates to per capita developed land growing 0.1 percent per year over this period. Along the model's balanced growth path, the growth rate of developed land per worker equals γ_c/τ. This relationship yields estimates of τ equal to 1.012 and 1.009 corresponding to $\alpha = 0.3$ and $\alpha = 0.2$.

Table 5 contains our estimates of the impact of agglomeration on per capita consumption growth, Δ, for various parameter configurations. For each capital share of housing, ω, we calculate Δ using two values of τ, one derived from USDA data and one from the CPI data, as indicated by the column headings. For each ω we use the corresponding estimate of δ from Table 3 and for each capital share net of land, α, we use the corresponding values of ϕ (see section 6.1) and $\tilde{\gamma}_c$ discussed above. The table also includes estimates of Δ for a value of δ that is in the middle of the range of estimates from the empirical literature on agglomeration and is close to estimates reported by Ciccone and Hall (1996). In these two cases we assume $\omega = 0.85$ when basing our estimates on the CPI data.

Our preferred estimate is the one based on $\omega = 0.85$, $\tau = 1.012$ (from the USDA) and $\alpha = 0.3$. In this case $\Delta = 10.7$ per cent, meaning that per capita consumption growth is that amount higher than in would be in the absence of agglomeration. While this is our preferred estimate, the table reveals that plausible alternative parameter configurations imply that Δ could by as high as 61.8 per cent! The lowest value of Δ based on our estimates of δ is 5.4 per cent, corresponding to $\omega = 0.89$ and $\tau = 1.009$ (USDA). The Ciccone and Hall (1996) estimate of δ implies values of Δ ranging from 10.4 to 78.5 per cent. We conclude that while there is considerable uncertainty regarding the exact magnitude of the contribution of agglomeration to per capita consumption growth, the evidence points to it likely being quite substantial.

8 Conclusions

This paper has three main contributions. First, we extend the neo-classical growth model to include location-specific production, consumption and housing, local developed land as a produced, durable input into production and housing, and local agglomeration effects as proposed by Ciccone and Hall (1996). We derive the balanced growth path of this model and show how per capita consumption growth depends on exogenous productivity growth, growth in land prices, and a parameter measuring the net impact on local productivity of agglomeration. Second, we use our model and panel data from US cities to estimate the key agglomeration parameter and find it to be significantly greater than unity, but substantially lower than estimated by Ciccone and Hall (1996) and the consensus range for this parameter in the literature. Third, we use the balanced growth implications of our model and our estimates of the key agglomeration parameter to quantify the impact of agglomeration on per capita consumption growth. We find that the magnitude of this impact could be quite large. Our our preferred estimate indicates that per capita consumption growth is 10.7 per cent higher than in would be without the effect of agglomeration on productivity.

References

Alvarez, F. and R. Shimer (2008). Search and rest unemployment. Unpublished University of Chicago manuscript.

Arellano, M. and S. Bond (1991). Some tests of specification for panel data: Monte carlo evidence and an application to employment equations. *Review of Economic Studies 58*(2), 277–297.

Arellano, M. and O. Bover (1995). Another look at the instrumental variables estimation of error-components models. *Journal of Econometrics 68*(1), 29–51.

Autor, D. H., L. F. Katz, and A. B. Krueger (1998). Computing inequality: Have computers changed the labor market? *The Quarterly Journal of Economics 113*(4), 1169–1213.

Blundell, R. and S. Bond (2000). Gmm estimation with persistent panel data: An application to production functions. *Econometric Reviews 19*(3), 321–340.

Carlino, G. A. and R. Voith (1992). Accounting for differences in aggregate state productivity. *Regional Science and Urban Economics 22*(4), 597–617.

Ciccone, A. (2002). Agglomeration effects in europe. *European Economic Review 46*(2), 213–227.

Ciccone, A. and R. E. Hall (1996). Productivity and the density of economic activity. *American Economic Review 86*(1), 54–70.

Ciccone, A. and G. Peri (2006). Identifying human-capital externalities: Theory with applications. *Review of Economic Studies 73*(2), 381–412.

Combes, P.-P., G. Duranton, and L. Gobillon (2008). Spatial wage disparities: Sorting matters! *Journal of Urban Economics 63*(2), 723–742.

Cooley, T. F. and E. C. Prescott (1995). Economic growth and business cycles. In T. F. Cooley (Ed.), *Frontiers of Business Cycle Research*, pp. 1–38. Princeton University Press.

Davis, M. and J. Heathcote (2005). Housing and the business cycle. *International Economic Review 46*(3), 751–784.

Davis, M. and F. Ortalo-Magné (2008). Household expenditures, wages, rents. Manuscript.

Davis, M. and M. Palumbo (2008). The price of residential land in large u.s. cities. *Journal of Urban Economics 63*(1), 352–384.

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica 50*(4), 1029–1054.

Heckman, J. J., L. Lochner, and C. Taber (1998). Explaining rising wage inequality: Explorations with a dynamic general equilibrium model of labor earnings with heterogeneous agents. *Review of Economic Dynamics 1*(1), 1–58.

Henderson, J. V. (2003). Marshall's scale economies. *Journal of Urban Economics 53*(1), 1–28.

Holtz-Eakin, D., W. Newey, and H. S. Rosen (1988). Estimating vector autoregressions with panel data. *Econometrica 56*(6), 1371–1395.

Katz, L. F. and K. M. Murphy (1992). Changes in relative wages, 1963-1987: Supply and demand factors. *The Quarterly Journal of Economics 107*(1), 35–78.

Kehoe, T. J., D. K. Levine, and P. M. Romer (1992). On characterizing equilibria of economies with externalities and taxes as solutions to optimization problems. *Economic Theory 2*(1), 43–68.

Krusell, P., L. Ohanian, J. V. Rious-Rull, and G. Violante (2000). Capital-skill complementarity and inequality: A macroeconomic analysis. *Econometrica 68*(5), 1029–1054.

Lucas, R. E. J. (2001). Externalities and cities. *Review of Economic Dynamics 4*(2), 245–274.

Ngai, R. and C. Pissarides (2007). Structural change in a multi-sector model of growth. *American Economic Review 97*(1), 429–443.

Roback, J. (1982). Wages, rents, and the quality of life. *Journal of Political Economy 90*(6), 1257–1278.

Rosenthal, S. S. and W. C. Strange (2003). Geography, industrial organization, and agglomeration. *The Review of Economics and Statistics 85*(2), 377–393.

Rosenthal, S. S. and W. C. Strange (2004). Evidence on the nature and sources of agglomeration economies. In V. Henderson and J.-F. Thisse (Eds.), *Handbook of Regional and Urban Economics*, Volume 4, pp. 2119–2171. North Holland.

Rossi-Hansberg, E. and M. L. J. Wright (2007). Urban structure and growth. *Review of Economic Studies 74*(2), 597–624.

Sargan, J. D. (1958). The estimation of economic relationships using instrumental variables. *Econometrica 26*(3), 393–415.

Windmeijer, F. (2005). A finite sample correction for the variance of linear efficient two-step gmm estimators. *Journal of Econometrics 126*(1), 25–51.

Table 1: Standard Deviations and Correlations in the Data

	Standard	Correlation with:				
Variable	Deviation	$\hat{w}_{eit}$	$\hat{p}_{it}$	$\hat{q}_{it}$	$\hat{\chi}_{it}$	$\hat{m}_{it}$
$\hat{w}_{eit}$	0.082	1.00	0.46	-0.03	0.43	0.21
$\hat{p}_{it}$	0.181		1.00	0.49	0.46	0.34
$\hat{q}_{it}$	0.052			1.00	0.23	0.14
$\hat{\chi}_{it}$	0.172				1.00	0.91
$\hat{m}_{it}$	0.224					1.00

Table 2: Unrestricted OLS Regression Coefficients

Dependent Variable: $\hat{w}_{eit}$		Dependent Variable: $\Delta\hat{w}_{eit}$	
$\hat{p}_{it}$	0.185	$\Delta\hat{p}_{it}$	0.148
	(0.039)		(0.103)
$\hat{q}_{it}$	-0.241	$\Delta\hat{q}_{it}$	0.454
	(0.153)		(0.554)
$\hat{\chi}_{it}$	0.571	$\Delta\hat{\chi}_{it}$	0.512
	(0.043)		(0.044)
$\hat{m}_{it}$	-0.349	$\Delta\hat{m}_{it}$	-0.304
	(0.035)		(0.034)
N	484		462
R^2	0.79		0.37

Note: The equation estimated is (28). Standard errors are in parentheses.

Table 3: GMM Parameter Estimates

Parameter	$\omega = 0.85$	$\omega = 0.81$	$\omega = 0.89$
δ	1.020	1.026	1.015
	(0.002)	(0.003)	(0.002)
ξ	0.563	0.563	0.563
	(0.004)	(0.004)	(0.004)
ρ	0.569	0.566	0.572
	(0.010)	(0.010)	(0.010)
J-test	36.0	36.0	36.0
p-value	1	1	1
m2 test	-0.13	-0.12	-0.14
p-value	0.45	0.45	0.45
m3 test	0.46	0.46	0.45
p-value	0.32	0.32	0.33

Notes: The equation estimated is (28). Standard errors are in parentheses. These estimates are based on $\alpha = 0.30$. The J-tests have 167 degrees of freedom.

Table 4: Effect of Model Assumptions on GMM Parameter Estimates

Assumption for: Land	Goods	δ	ξ	ρ
Variable	$\eta < 1$	1.020	0.563	0.569
		(0.002)	(0.004)	(0.010)
Variable	$\eta = 1$	1.019	0.607	0.256
		(0.002)	(0.004)	(0.013)
Fixed	$\eta = 1$	0.933	0.524	0.182
		(0.008)	(0.007)	(0.018)

Notes: The equations estimated in the first, second and third rows are (28), (33) and (34). Standard errors are in parentheses. These estimates are based on $\alpha = 0.30$ and $\omega = 0.85$.

Table 5: Effect of Agglomeration on Per Capita Consumption Growth

	$\omega = 0.85$		$\omega = 0.81$		$\omega = 0.89$		$\delta = 1.055$	
	USDA	CPI	USDA	CPI	USDA	CPI	USDA	CPI
$\alpha = 0.3$	10.7	29.9	11.6	35.4	9.9	61.8	16.2	78.5
$\alpha = 0.2$	6.1	16.4	6.8	29.7	5.4	45.9	10.4	73.4

Notes: Table entries correspond to the per cent increase in the growth rate of per capita consumption due to agglomeration under the indicated configuration of model parameters. Except for the last two columns, each table entry is based on the estimate of δ in Table 3 corresponding to the indicated value of ω. The table entries in the last two columns correspond to a value for δ that is in the middle of the range of estimates found in empirical studies of agglomeration and is close to estimates reported by Ciccone and Hall (1996). These table entries are based on $\omega = 0.85$. The columns headed by "USDA" correspond to estimates of τ based on USDA data and columns headed by "CPI" correspond to estimates of τ based on CPI data, as discussed in the text. When $\alpha = 0.3$ the corresponding value of ϕ we assume is 0.95, and when $\alpha = 0.2$ the corresponding value of ϕ we assume is 0.97. The values of ϕ are described in section 6.1.

A Data Appendix

In this appendix, we document how we construct key variables from various data sources and then document how we merge our different data sources together to create our data set for estimation.

A.1 CPS Data on Wages and Hours Worked by Skill

The March CPS data are available for download at *http://cps.ipums.org/cps/* as part of the Integrated Public Use Microdata Series (IPUMS-CPS) project at the University of Minnesota Population Center.

We download the March CPS data from 1986 through 2007. We choose 1986 as our starting year because the CPS identifies only 15 metropolitan areas in prior years. The CPS wage and employment questions refer to the "previous calendar year." Therefore, data for any given year's CPS is treated as data appropriate for the previous calendar year. For example, variables generated from the March 2005 CPS are treated as data for the year 2004.

In each year of our data, we use the following criteria to restrict the sample (with IPUMS-CPS variables in italics)

- Respondent lives in a household, not in group quarters or vacant units ($gq = 1$)
- Is aged 20 to 65 ($age \geq 20$ and $age \leq 65$)
- Wage and salary income in the previous calendar year is identified and is nonzero ($incwage > 0$ and $incwage < 999998$)
- Educational attainment is recorded ($educrec \geq 1$ and $educrec \leq 9$)
- Has an identified metro area of residence (*metarea* non missing)[13]

For each MSA, we create the following three variables:

1. Ratio of labor input of high skill to labor input of low skill, n_{ei}/n_{ui}
2. Ratio of total wages paid to total wages paid to low skill workers, χ_i
3. Average weekly wage of high skill workers, $w_{e,i}$.

[13] According to notes from the IPUMS-CPS, the metro area of residence was not collected from respondents, but added by the Census Bureau. The metro areas of residence is based on FIPS codes used in the 1990 census.

We use the *educrec* categorical variable to label respondents as either "low" or "high" skill workers. High skill workers are assumed to have completed 4 years of college. Everyone else in the sample is assumed to be a low skill worker.

The variable n_{ei} is created as the total of weeks worked the previous calendar year ($wkswork1$) multiplied by the number of hours per week the respondent usually worked ($uhrswork$) for high skill workers. The variable n_{ui} is created as the same product, but for low skill workers. For each respondent, we weigh the product of $wkswork1$ and $uhrswork$ using the IPUMS-CPS sampling person weights, $perwt$.

χ_i is computed as

$$\frac{w_{e,i}n_{ei} + w_{u,i}n_{ui}}{w_{u,i}n_{ui}} = \frac{\sum_{j\epsilon MSA_i} perwt_j \cdot wages_j}{\sum_{j\epsilon MSA_i} perwt_j \cdot wages_j \cdot 1\{unskilled_j\}}$$

for respondent j in MSA i, i.e. as the sum of all workers' pre-tax wage and salary income for the previous calendar year (*incwage*) divided by the sum of all low skill workers' pre-tax wage and salary income for the previous calendar year. We weigh pre-tax wage and salary income for all persons using the IPUMS-CPS sampling person weights.

$w_{e,i}$ is created as the sum of all high skill workers' pre-tax wage and salary income for the previous calendar year (created as an input into χ_i) divided by n_{ei}.

A.2 BEA Data on Output Prices

We use two data sources from within the BEA web site: The Annual Industry Accounts, *http://www.bea.gov/industry/index.htm#annual*, and the Regional Economic Accounts data on Local Area Personal Income, *http://www.bea.gov/regional/reis/*.

Chain-type price indexes for industry output are available over the 1947-2007 period in the Annual Industry Accounts. Many industry price indexes are missing in 2007, so we do not use data from that year. To construct a price index for output produced by MSA, we merge this information with MSA-level data on earnings by industry that is available in Tables CA05 and CA05N of the Regional Economic Accounts. Earnings is inclusive of wage and salary disbursements, supplements to wages and salaries, and proprietors' income.

Thus, we assume that the price of output varies across MSAs because industry composition varies across MSAs, and the price index for industry output varies across industries.

We now describe our measurement in detail. For this we use notation that does not correspond with that in the main text. Denote $g_{t,j}$ as the growth rate of the price of industry output j from periods t to $t+1$ and g_t^i as the growth rate of the price of all output produced in MSA i between years t and $t+1$. Assuming output from

$j = 1, \ldots, N$ industries is produced in MSA i in year t, we set the growth rate of the price of output produced in MSA i between years t and $t+1$ as

$$g_t^i = \sum_{j=1}^{N} \omega_{t,j}^i \; g_{t,j} \; . \tag{35}$$

The weight on each industry, $\omega_{t,j}^i$, is the share of total MSA earnings attributable to earnings of industry j in MSA i in year t:

$$\omega_{t,j}^i = \frac{\epsilon_{t,j}^i}{\sum_{k=1}^{N} \epsilon_{t,k}^i} \; , \tag{36}$$

where $\epsilon_{t,j}^i$ stands for total earnings of employees in industry j in MSA i during year t. In these computations, we only consider earnings from non-farm private industries. For each MSA, we construct a price index for output, normalized to 1.0 in the year 1969, that is consistent with the sequence of time-series estimates of g_t^i.

Ideally, we would compute the growth rate of the price of output produced in city i between years t and $t+1$ as

$$\sum_{j=1}^{N} \phi_{t,j}^i \; g_{t,j} \; , \tag{37}$$

with $\phi_{t,j}^i$ equal to the fraction of the nominal value of output in year t in MSA i that is accounted for by industry j. In an environment in which (a) output in each industry is produced by a set of identical firms all using a Cobb-Douglas combination of capital, labor, and land and (b) the labor-share of output is identical in each industry, assumptions that hold in our model, then industry j's share of nominal GDP in MSA i in year t, $\phi_{t,j}^i$, is equal to its earnings share $\omega_{t,j}^i$, and equations (35) and (37) are equivalent. In these calculations, we assume that proprietors' income are payments to labor.[14]

A few details complicate these calculations. First, on a somewhat infrequent basis, Tables CA05 and CA05N do not report estimates of earnings for a given industry in an MSA in a given year. In these cases, we set earnings for this industry-MSA-year cell to zero.[15] Also, some of the industry-MSA-year employment estimates are marked

[14]In the event that proprietors' income includes some payments to capital, equations (35) and (37) are equivalent as long as capital's share of proprietors' income and the fraction of earnings attributable to proprietors' income are both constant across industries.

[15]The three reasons that are listed for omission are (a) avoid disclosure of confidential information (code D), (b) earnings are less than $50,000 (code L), or (c) data not available for this year (code N). These omissions occur in approximately six percent of industry-MSA-year cells from 1969 to the mid-1990s and about thirteen percent of cells from the mid-1990s through 2006.

with code E. According to the BEA web site, these estimates "constitute the major portion of the true estimate." In these cases, we assume that the reported estimate is equal to the actual estimate.

Second, the definition of industries in the Regional Accounts is not consistent across years. Table CA05 reports employment based on SIC-industry classifications over the 1969-2000 period and CA05N reports employment based on NAICS industry classifications spanning the years 2001-2006.

We map SIC and NAICS industry employment from Tables CA05 and CA05N to prices from the Annual Industry Accounts according to the tables below. The tables below list all the categories of nonfarm private employment. The sum of the earnings estimates in each of these categories is considered as total nonfarm private earnings, and is used to compute the denominator of equation (36).

Data for Earnings Weights, $w^i_{t,j}$ Regional Accounts Table CA05, 1969-2000		Data for Growth in Prices, $g^p_{t,j}$ Industry Accounts, 1969-2001	
Line	Label	Line	Label
100	Agricultural services, forestry fishing and other	3	Agriculture, forestry, fishing and hunting
200	Mining	6	Mining
300	Construction	11	Construction
400	Manufacturing	12	Manufacturing
500*	Transportation and public utilities less electric, gas, and sanitary services	36	Transportation and warehousing
570	Electric, gas, and sanitary services	10	Utilities
610	Wholesale trade	34	Wholesale trade
620	Retail trade	35	Retail trade
700	Finance, insurance and real estate	50	Finance, insurance, real estate, rental and leasing
800	Services	59	Professional and business services

Data for Earnings Weights, $w^i_{t,j}$ Regional Accounts Table CA05N, 2001-2005		Data for Growth in Prices, $g^p_{t,j}$ Industry Accounts, 2001-2006	
Line	Label	Line	Label
100	Forestry, fishing, related activities and other	5	Forestry, fishing and related activities
200	Mining	6	Mining
300	Utilities	10	Utilities
400	Construction	11	Construction
500	Manufacturing	12	Manufacturing
600	Wholesale trade	34	Wholesale trade
700	Retail trade	35	Retail trade
800	Transportation and warehousing	36	Transportation and warehousing
900	Information	45	Information
1000	Finance and insurance	51	Finance and insurance
1100	Real estate and rental and leasing	56	Real estate and rental and leasing
1200	Professional, scientific and technical services	60	Professional, scientific and technical services
1300	Management of companies and enterprises	64	Management of companies and enterprises
1400	Administrative and waste services	65	Administrative and waste management services
1500	Educational services	69	Educational services
1600	Health care and social assistance	70	Health care and social assistance
1700	Arts, entertainment and recreation	75	Arts, entertainment and recreation
1800	Accommodation and food services	78	Accommodation and food services
1900	Other services except public administration	81	Other services except government

In all cases except one, there is an exact correspondence of earnings estimates from Tables CA05 and CA05N to prices from the Annual Industry Accounts. For the SIC category of "Transportation and public utilities," line 500 of Table CA05, there is no clean analogous price index in the Annual Industry Accounts. Instead, the Annual Industry Accounts includes separate price indexes for "Transportation and warehousing" and "Utilities." In Table CA05, we therefore separate earnings of the single Transportation and public utilities into earnings in two categories: Earnings from utilities ("electric, gas, and sanitary services", line 570) and earnings from transportation and public utilities less earnings from utilities (i.e. line 500 less line 570).

A.3 BLS and 1990 Decennial Census of Housing Data on Housing Rents

We create annual estimates over the 1985-2006 period of the average rents paid for certain types of rental units, by MSA, using a two-step procedure.

In the first step, we estimate the average rents paid for certain types of rental housing units in 1990 using household-level data from the 1990 Decennial Census of Housing (DCH). These data are available for download at *http://usa.ipums.org/usa/* as part of the Integrated Public Use Microdata Series (IPUMS-USA) project at the University of Minnesota Population Center. We use data from the 1990 DCH because the metropolitan area of residence is identified for many more metropolitan areas than in the 2000 DCH.

With IPUMS-USA variables in italics, we restrict the 1990 DCH sample to renter non-farm households in 2-19 unit residences in a building built between 1940 and 1986 and living in an identifiable MSA ($ownershg = 2$, $farm \neq 1$, $unitsstr \in \{5, 8\}$, $builtyr \in \{3, 7\}$, and $metarea > 0$) who live in households and do not live in group quarters ($gq \in \{3, 4, 6\}$) and where the reported monthly gross rent of the house (rent inclusive of utilities) is nonzero ($rentgrs > 0$). Conditional on these restrictions, we compute the weighted average value of units by MSA using the sampling weight variable $hhwt$. These calculations yield estimates of the average rental price of housing for 272 metro areas as identified in the 1990 DCH. We exclude single-family rented units, rented high-rise units ($>$ 20 units), and units in very old (built before 1940) or very new (built after 1986) apartment buildings to attempt to keep the average characteristics of rental units roughly constant across metropolitan areas without employing hedonic regressions.

In the second step, we extrapolate the annual rental price of housing in each metro area forward from 1990 to 2006 and backwards from 1990 to 1985 using annual MSA-specific constant-quality price indexes for rental units. These price indexes for tenant rents are published by the Bureau of Labor Statistics (BLS) as part of computations for the Consumer Price Index, and are available at *http://www.bls.gov*. The BLS reports rental price indexes for 27 MSAs, but the indexes of three of these MSAs (Phoenix, AZ, Washington, DC, and Tampa Bay, FL) do not extend back to 1985 and we exclude these from our sample.

After merging the 1990 DCH with the BLS price indexes, and eliminating the MSAs for which price indexes are not available back to 1985, we are left with annual estimates of the rental price of housing units over the 1985-2006 period for 24 MSAs. These MSAs are: Anchorage, AK; Atlanta, GA; Boston, MA; Chicago, IL; Cincinnati, OH; Cleveland, OH; Dallas, TX; Denver, CO; Detroit, MI; Honolulu, HI; Houston, TX; Kansas City, MO; Los Angeles, CA; Miami, FL; Milwaukee, WI; Minneapolis, MN; New York, NY; Philadelphia, PA; Pittsburgh, PA; Portland, OR; San Diego, CA; San Francisco, CA; Seattle, WA; St. Louis, MO.

A.4 Merging the Data

To create our data set, we merge the CPS data on wages and employment (section A.1) with the BEA data on output prices (A.2) and the annual data we construct on housing rents by merging the BLS rental price indexes with information on housing rents in the 1990 Decennial Census of Housing (section A.3). After all data are merged, we are left with a balanced panel of 22 MSAs over 22 years, 1985-2006. Anchorage, AK and Portland, OR are excluded from our data due to lack of wage and employment data from the CPS. Over our sample period, 36 percent of the population of the United States live in one of the the 22 MSAs in our study.

In every MSA and date, the minimum number of respondents from the CPS is never less than 200. The median number of respondents is about 540 until about 2000, at which point the median jumps to about 1,000. The maximum number of respondents is always above 3,500 and is typically about 5,000.

All data files are merged by MSA. Note that the MSA definitions are not completely consistent across data sets. The MSA definitions in the CPS data are consistent with the definitions as of the 1990 Census. In the BEA data, MSAs definitions are given by the list in the November, 2007 report of the Office of Management and Budget (OMB).[16] The MSA definitions in the BLS rent indexes are also based on the definitions in OMB, but the definitions can change over time (as the OMB changes its definition of MSAs over time); further, the rent indexes for New York, Los Angeles, and Chicago are based on the concept of the Consolidated Metropolitan Statistical Area (CMSA). CMSAs include more counties than simple MSAs.

B Monte Carlo Study

We perform a Monte Carlo study using our data with the following procedure. First, we generate a normal $N(0,1)$ value for z for the first time period. We then set each of the right-hand side variables (p^h, q, χ, and m) equal to a normal random variable plus $0.5 * z$. We then update z using the specified AR(1) process (using the appropriate value for ρ) and repeat.

Once we generate the entire panel, we restandardize the right-hand side variables so that the have the same first and second moments as our data, and then generate $w_{ei,t}$ using equation (28). In this process, we rescale the variance of z such that the simulated variance of $w_{ei,t}$ is the same as the data. This procedure allows z to be correlated with the right-hand side variables, but also sets lagged values of the right-hand side variables orthogonal to the innovations to z, which is the maintained assumption of the estimator.

We repeat this procedure 10,000 times (thus generating 10,000 data sets) for each of the following four combinations of parameter values: $\delta = \{1.011, 1.040\}$,

[16] For a complete list of the counties comprising each MSA, go to *http://www.census.gov/population/www/metroareas/metrodef.html*.

$\rho = \{0.60, 0.95\}$. In all cases, we set $\xi = 0.60$, $\omega = 0.85$ and $\alpha = 0.30$. We report the results for an unrestricted weight matrix, in which the residuals from different time periods are down-weighted with the Bartlett kernel. This is the method we use when we report parameter estimates and standard errors using the actual data in the main body of the paper.[17]

In table 6, for each combination of parameters we report six statistics. Conceptually, the first three statistics we report – average estimate of the parameter, the mean absolute deviation of the parameter estimate, and the root mean square error of the parameter estimate – are indicative of potential biases to the parameter estimate. The other three statistics we report are indicative of potential bias to estimated standard errors: The probability of rejecting the null (which should equal 0.05 if the test is correctly sized), the 95% coverage probability and a scale factor (to be defined). The 95% coverage probability is the percentage of the simulations in which the distance between the parameter estimate and the true parameter is in the 95% confidence interval. The scale factor represents the adjustment to reported standard errors such that the probability of rejecting the null is 0.05. That is, if the reported standard error is S and the scale factor is F, the probability of rejecting the null when the standard error is of size $S * F$ is 0.05.

Based on the results in this table, we draw two main conclusions. First, our estimates of δ, ρ, and ξ are, for all practical purposes, unbiased. Second, the asymptotic standard errors are too small for all parameter estimates: The probability of rejecting the null is always larger than 0.05 and the 95% coverage probability is always less than 0.95. For the case of $\rho = 0.60$, the scale factors suggest that our standard errors are too small by a factor of 2.35 for δ, 1.44 for ρ, and 2.09 for ξ.

[17] We have also performed the Monte Carlo for using a “restricted” weight matrix in which cross sectional residuals from different time periods are not allowed to be correlated. This is the original formulation of Holtz-Eakin et al. (1988) that results in a block-diagonal weight matrix. The results are very similar and are available on request.

Table 6: Monte Carlo Results

	$\rho = 0.60$			$\rho = 0.95$		
	δ	ρ	ξ	δ	ρ	ξ
			$\delta = 1.011$			
Expected Value	1.0117	0.6030	0.5975	1.0112	0.9406	0.5994
Mean Absolute Deviation	0.0008	0.0119	0.0032	0.0002	0.0117	0.0009
Root Mean Square Error	0.0009	0.0152	0.0039	0.0002	0.0148	0.0011
Prob. Reject Null	0.49	0.17	0.37	0.37	0.28	0.28
95% Coverage Prob.	0.51	0.83	0.63	0.63	0.72	0.72
Scale Factor	2.35	1.44	2.09	2.02	1.75	1.83
			$\delta = 1.040$			
Expected Value	1.0407	0.6030	0.5975	1.0402	0.9406	0.5994
Mean Absolute Deviation	0.0008	0.0119	0.0032	0.0002	0.0117	0.0009
Root Mean Square Error	0.0010	0.0152	0.0039	0.0003	0.0148	0.0011
Prob. Reject Null	0.49	0.17	0.37	0.37	0.28	0.28
95% Coverage Prob.	0.51	0.83	0.63	0.63	0.72	0.72
Scale Factor	2.35	1.44	2.09	2.02	1.75	1.83